



True-concurrency probabilistic models Branching cells and distributed probabilities for event structures

Samy Abbas^{a,*}, Albert Benveniste^b

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

^b*IRISA, Campus de Beaulieu, 35042 Rennes Cedex, France*

Received 15 February 2005; revised 1 August 2005

Abstract

This paper is devoted to probabilistic models for concurrent systems under their true-concurrency semantics. Here we address probabilistic event structures. We consider a new class of event structures, called *locally finite*, that extend confusion-free event structure. In locally finite event structures, maximal configurations can be tiled with *branching cells*: branching cells are minimal and finite sub-structures capturing the choices performed while scanning a maximal configuration. The probabilistic event structures that we introduce have the property that “concurrent processes are independent in the probabilistic sense.”

© 2005 Elsevier Inc. All rights reserved.

1. Introduction

For concurrent systems, there is a fundamental difference according to the underlying *interleaving* or *true-concurrency* semantics considered. True-concurrency probabilistic models belong to a

* Corresponding author.

E-mail address: sabbes@isr.umd.edu (S. Abbas), albert.benveniste@irisa.fr, *URL:* www.irisa.fr/distribcom/benveniste (A. Benveniste).

recent research area [1–4] where partially ordered processes are randomized, not their interleavings. This paper addresses probabilistic event structures [5,6] in this framework.

Probability and stochastic process theories provide the mathematical foundations for randomizing systems. The general model is the following: first consider the space Ω of all possible histories of the system. Randomizing the system consists in defining a probability measure on Ω , meaning that a particular execution of the system will occur at random, according to the probability distribution chosen. For true-concurrency systems, we still conform to this general concept. In particular, for systems represented by event structures, a history consists of a maximal configuration of the event structure.

How can such a probability be constructed? And, among all probabilities that can be constructed, are not there some particular ones to be preferred, with some desirable properties?

For sequential systems such as, say, Markov chains, classical constructions follow the following natural steps. Regard a history as resulting from a sequence of successive choices (choosing the next state). Assign a given probability to each choice. Any partial execution, seen as a finite stack of choices, is given a probability by the usual chain rule. Measure theoretic arguments show that this construction indeed extends to a “limiting” probability measure on the space of all histories of the system. Thus, randomizing a system amounts to:

- (1) isolating the choices performed by the system—this typically relates to combinatorics;
- (2) giving a method for assigning a probability to each choice, and then to each finite stack of choices—this is the central job of probability theory;
- (3) obtaining a “limit probability”—by using arguments from measure theory.

For true-concurrency systems, we wish to follow the same steps. However, the main difference with the classical setting is that choices are no more totally ordered in time. Indeed, some choices occur *concurrently*. The interleaving semantics ignores this issue by assigning a particular order to the choices made, which results in assigning a probability when choosing among different interleavings. For a true-concurrency randomization, we may not use this trick. Instead, we wish to consider stacks of possibly concurrent choices. An advantage is that concurrent choices can then be made probabilistically independent. Such a requirement is natural for distributed systems: local components act asynchronously and without communication for some limited amount of “time”, and during this period, the actions inside local components shall be independent in the probabilistic sense since the local components do not communicate. To summarize, our first task consists in decomposing partial executions of an event structure, i.e., finite configurations, as stacks of possibly concurrent choices. Then, we will associate a probabilistic interpretation of this decomposition by making concurrent choices independent in the probabilistic sense.

We analyze a simple example to make these ideas precise.

Consider the event structure depicted at left in Fig. 1. Consider first the triple $a\#b\#c$ (where $\#$ denotes the conflict relation). Note that a and c are concurrent. If event a is selected in an execution of the system, the only possibility is that event c eventually fires in the same execution. We have thus a mutual implication $a \Leftrightarrow c$, although events a and c are concurrent. Consequently, events a and c must be considered jointly for the randomization of maximal configurations. Furthermore, unlike for sequential systems like, e.g., Markov chains, we cannot assume in general that the probability decomposes multiplicatively over events.




Fig. 1. Two event structures. Curved lines denote conflict and directed arrows denote causality.

Add now three events $d\#e\#f$ to the event structure, not related to events a, b, c . Then the choices involving d, e, f are unrelated to the choices involving a, b, c . Thus, we can multiplicatively decompose our desired probability as a product of the two probabilities governing the choices in a, b, c and d, e, f , respectively. By this way, we make use of our policy according to which, as much as we can, parallel processes shall be made independent in probability. For example, the probability of selecting the pair $\{a, c\}$ jointly with $\{e\}$ is:

$$\mathbb{P}(\{a, c\}, \{e\}) = q_{a\#b\#c}(\{a, c\}) \times q_{d\#e\#f}(\{e\}),$$

where probabilities $q_{a\#b\#c}$ and $q_{d\#e\#f}$ are governing the choices in a, b, c and d, e, f , respectively. Sub-event structures $a\#b\#c$ and $d\#e\#f$ are called *branching cells*. We have implicitly used the fact that, in this example, every maximal configuration of the event structure decomposes as a union of maximal configurations of each branching cell. This remark will be instrumental in our theory.

Consider now Example 2 of the same figure. It coincides with Example 1, except that causality $a \leq d$ has been added. Remark that any maximal execution of the whole system still induces, by restriction, a maximal configuration of the branching cell $a\#b\#c$. Consider the alternatives “firing e ” versus “firing f .” Both are allowed, whatever the decision taken in branching cell $a\#b\#c$ is. However, if branching cell $a\#b\#c$ produces $\{b\}$, then d is disabled, so that e and f compete alone. Whereas, if the result is $\{a, c\}$ instead of $\{b\}$, then d is enabled, and now the competition involves d, e and f . Therefore we have to consider the two possible branching cells $e\#f$ and $d\#e\#f$. Typically, probabilities are now computed by:

$$\begin{aligned} \mathbb{P}(\{b\}, \{e\}) &= q_{a\#b\#c}(\{b\}) \times q_{e\#f}(\{e\}), \\ \mathbb{P}(\{a, c\}, \{e\}) &= q_{a\#b\#c}(\{a, c\}) \times q_{d\#e\#f}(\{e\}). \end{aligned}$$

As we see, any maximal configuration may be decomposed through stacks of finite configurations, each finite configuration being maximal in some sub-event structure. The different sub-event structures encountered are called branching cells. Branching cells isolate the choices performed to obtain a maximal configuration of the event structure. Although branching cells involved in the decomposition of some given maximal configuration do not overlap, there may be branching cells with a nonempty intersection— $e\#f$ and $d\#e\#f$ in our example. We interpret this fact by saying that the decomposition through branching cells is *dynamic*. This means that an event, when occurring in different executions of the system, may be considered in different branching cells, depending on the prior context. If one performs the decomposition on trees instead of event structures, or even on confusion-free event structures, one finds that branching cells are globally disjoint, i.e., the decompositions are not dynamic. We summarize the properties of branching cells as follows:

- (i) Branching cells isolate in a recursive and dynamic way the independent choices performed while constructing a maximal configuration of an event structure;
- (ii) Branching cells support a randomization where concurrent branching cells are made independent in the probabilistic sense.

In this paper, we propose an analysis of event structures that generalizes the above example. We define *branching cells* as finite sub-event structures possessing properties (i,ii) above. Hence branching cells represent the atomic parts in the stacks of choices that we are seeking for event structures. Branching cells must be dynamically defined. Therefore, for event structures arising from the unfolding of a Petri net, branching cells differ from clusters [7], which are statically defined on the net. Branching cells and their properties constitute the first contribution of this paper.

Our study does not encompass all prime event structures. To ensure the finiteness of branching cells, we add an assumption called *local finiteness*. Under some very mild conditions, trees and confusion-free event structures are locally finite. Indeed, locally finite event structures can be seen as event structures with a kind of “bounded confusion.” Extending the decompositions that we propose to general event structures—i.e., considering non locally finite event structures—requires some more work involving in particular transfinite arguments.

The probabilistic construction performed in the above example generalizes to any locally finite event structure. We attach a *local transition probability* to each branching cell. Then we show that local transition probabilities can be combined using a chain rule where concurrent choices are made independent. This amounts to define a probability measure on the space of maximal configurations of the event structure, seen as a kind of product of local transition probabilities. This special construction of probabilities is called a *distributed product*—this term is reminiscent from the fact that concurrent choices are made probabilistically independent. Probabilities reached by this way are called *distributed probabilities*. For distributed probabilities, parallel processes are made independent in probability, *at the grain of branching cells*. Moreover, we will show that it is not possible, in general, to get the same property at a finer granularity than branching cells.

Conversely, we show that the “concurrency/independence” matching property is characteristic of distributed probabilities. In other words, if a probability satisfies this property, then it is a distributed product. Distributed probabilities and their recursive construction through distributed products are the second main contribution of the paper.

The paper is organized as follows. In Section 2, we introduce locally finite event structures and stopping prefixes, which are basic objects for studying locally finite event structures. Section 3 is devoted to branching cells and the associated decomposition of configurations. Probabilistic event structures are defined in Section 4 and we also show how to reduce the construction of locally finite probabilistic event structures to that of finite probabilistic event structures. The special class of distributed probabilities is investigated in Section 5. Two appendices collect the longest proofs.

Related work. In [4], the randomization of event structures is studied from the domain theory point of view, by using continuous valuations defined on the domain of configurations of an event structure. This viewpoint is closely related to the probabilistic powerdomains from Jones and Plotkin [8,9]. The authors use the one-to-one correspondence between continuous valuations and Borel measures on the space of configurations [10]. However, it is not easy to determine when a measure on the domain of configurations has its support in the space of maximal configurations. This is

the role of the non-leaking valuations in [4]. Non-leaking valuations are constructed explicitly for confusion-free event structures, where the example of valuations with independence is given. For confusion-free event structures, non-leaking valuations with independence coincide with our own distributed probabilities.

The present approach has its roots in the work [1] where stopping prefixes were first proposed for true-concurrency systems—stopping prefixes were called “stopping times” in the above reference—and the principle of concurrency matching probabilistic independence was first stated. However, it is only with the new notion of branching cell first proposed in [3] that the preliminary ideas of [1] could really be developed. The reader is referred to [1] for motivations of the present work related to applications.

2. Stopping prefixes of event structures

2.1. Prerequisites on event structures

Throughout this paper we consider only prime event structures. We will say *event structures* for short, always meaning *prime event structures*. We list in Table 1 the notations used throughout Sections 2–3.

Event structures. Let (\mathcal{E}, \preceq) be a partially ordered set. Elements of \mathcal{E} are called *events*, and we assume that \mathcal{E} is at most countable. The order relation \preceq is called the *causality* relation. $<$ and \succeq are obvious notations for relations derived from \preceq . The *downward closure* of a subset $A \subseteq \mathcal{E}$ is defined

Table 1
Notations for event structures and branching cells

Symbol	Meaning
Configurations of event structures	
$\mathbf{co}, \#_{\mu}$	Concurrency and immediate conflict relations
\mathcal{V} or $\mathcal{V}_{\mathcal{E}}$	Poset of finite configurations of \mathcal{E}
$\overline{\mathcal{V}}$ or $\overline{\mathcal{V}}_{\mathcal{E}}$	Poset of configurations of \mathcal{E}
Ω or $\Omega_{\mathcal{E}}$	Set of maximal configurations of \mathcal{E}
<i>R</i> -stopped configurations, branching cells	
\mathcal{W} or $\mathcal{W}_{\mathcal{E}}$	Poset of finite <i>R</i> -stopped configurations of \mathcal{E}
$\overline{\mathcal{W}}$ or $\overline{\mathcal{W}}_{\mathcal{E}}$	Poset of <i>R</i> -stopped configurations of \mathcal{E}
\mathcal{C} or $\mathcal{C}_{\mathcal{E}}$	Set of branching cells of \mathcal{E}
$\delta(v)$ or $\delta_{\mathcal{E}}(v)$	Set of branching cells of \mathcal{E} enabled by v
$\Delta(v)$ or $\Delta_{\mathcal{E}}(v)$	Covering of v in \mathcal{E} , for $v \in \overline{\mathcal{W}}$
Future of a configuration	
\mathcal{V}^u	Poset of finite configurations of \mathcal{E}^u
$\overline{\mathcal{V}}^u$	Poset of configurations of \mathcal{E}^u
\mathcal{W}^u	Poset of finite <i>R</i> -stopped configurations of \mathcal{E}^u
$\overline{\mathcal{W}}^u$	Poset of <i>R</i> -stopped configurations of \mathcal{E}^u
\mathcal{C}^u	Set of branching cells of \mathcal{E}^u
$\Delta^u(v)$	Covering of v in \mathcal{E}^u , for $v \in \overline{\mathcal{W}}^u$

by $\downarrow A = \{e \in \mathcal{E} : \exists e' \in A, e \preceq e'\}$. For a singleton, we note $\downarrow e = \downarrow\{e\}$. We assume that $\downarrow e$ is a finite subset of \mathcal{E} for every event e . An *event structure* is a triple $(\mathcal{E}, \preceq, \#)$, where (\mathcal{E}, \preceq) is a partially ordered set as above, and $\#$ is a binary symmetric and irreflexive relation on \mathcal{E} , called the *conflict relation*. It is assumed that the conflict relation satisfies the so-called *inheritance axiom*, i.e.:

$$\forall e_1, e_2, e_3 \in \mathcal{E}, \quad (e_1 \# e_2 \text{ and } e_2 \preceq e_3) \implies e_1 \# e_3.$$

With a slight abuse of notations, we shall identify an event structure $(\mathcal{E}, \preceq, \#)$ and its set \mathcal{E} of events. Remark that our definition includes the empty set \emptyset as an event structure. We say that an event $e \in \mathcal{E}$ is *minimal* in \mathcal{E} if e is a minimal element of the partial order (\mathcal{E}, \preceq) . We denote by $\text{Min}_{\preceq}(\mathcal{E})$ the set of minimal events of \mathcal{E} .

The *concurrency relation* is the binary relation on \mathcal{E} denoted by the symbol \mathbf{co} , and defined by: $\mathbf{co} = (\mathcal{E} \times \mathcal{E}) \setminus (\preceq \cup \succeq \cup \#)$. Hence two events e, e' are concurrent if they are neither in conflict nor causally related.

Prefixes and configurations. A subset $P \subseteq \mathcal{E}$ is called a *prefix* of \mathcal{E} if it is downward closed, i.e., if $P = \downarrow P$.

A subset $v \subseteq \mathcal{E}$ is said to be *conflict-free* if it does not contain any two elements in conflict, i.e., if $\# \cap (v \times v) = \emptyset$. A subset $v \subseteq \mathcal{E}$ is said to be a *configuration* of \mathcal{E} if v is a conflict-free prefix of \mathcal{E} . Remark that \emptyset is always a configuration of \mathcal{E} . We say that u is a *sub-configuration* of v if u and v are two configurations such that $u \subseteq v$.

We denote by $\overline{\mathcal{V}}_{\mathcal{E}}$, or by $\overline{\mathcal{V}}$ for short, the set of configurations of \mathcal{E} . $(\overline{\mathcal{V}}, \subseteq)$ is a partial order. We denote by $\mathcal{V}_{\mathcal{E}}$, or by \mathcal{V} for short, the sub-poset of *finite* configurations of \mathcal{E} . Remark that, for every event $e \in \mathcal{E}$, the finite subset $\downarrow e$ is the smallest configuration that contains e .

We say that two configurations v, v' are *compatible* if $v \cup v'$ is a configuration. Otherwise we say that v and v' are *incompatible*. We say that two events e and e' are compatible if $\downarrow e$ and $\downarrow e'$ are compatible, and that an event e is compatible with a configuration v if $\downarrow e$ and v are compatible.

Maximal configurations. Any union of pairwise compatible configurations is a configuration. In particular, any chain of configurations admits an upper bound. Furthermore, the set $\overline{\mathcal{V}}$ of configurations is nonempty since $\emptyset \in \overline{\mathcal{V}}$. As a consequence, by virtue of Zorn's Lemma, $\overline{\mathcal{V}}$ has maximal elements, i.e., configurations ω such that, for every configuration v , $v \supseteq \omega \implies v = \omega$. We denote the nonempty set of maximal configurations by $\Omega_{\mathcal{E}}$, or by Ω for short. The notation is indeed reminiscent to the Ω from probability theory, the reason will be given in Section 4. Any configuration of \mathcal{E} is a sub-configuration of some maximal configuration.

Sub-event structures. Let F be a subset of \mathcal{E} . Let $\preceq|_F$ and $\#|_F$ denote, respectively, the restrictions of causality and conflict to F , defined by:

$$\preceq|_F = \preceq \cap (F \times F), \quad \#|_F = \# \cap (F \times F).$$

Then the triple $(F, \preceq|_F, \#|_F)$ is an event structure, we denote it by $(F, \preceq, \#)$ for short. Implicitly, every subset $F \subseteq \mathcal{E}$ will be considered as an event structure with this convention.

Sequential event structures: trees of events. Event structures are a model for concurrency, where \mathbf{co} captures the concurrency properties of \mathcal{E} . Accordingly, sequential systems, seen as particular cases of concurrent systems, shall be characterized by a trivial concurrency relation.

Therefore, we say that an event structure \mathcal{E} is a *tree of events* [11] if $\mathbf{co} = \emptyset$. This is equivalent to \mathcal{E} be an at most countable union of disjoint oriented trees in the usual sense, with the conflict relation as follows: all roots are pairwise in conflict, and for every event e , the immediate successors of e are pairwise in conflict.

2.2. Future of a configuration. Concatenation of configurations

Let v be any configuration of \mathcal{E} . We introduce the notion of future to analyze the set of events that can occur “after” v , in a sense to be made precise, since event structures involve concurrency.

Definition 2.1. For v a configuration of \mathcal{E} , we define the following subset of \mathcal{E} :

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

\mathcal{E}^v is called the *future* of v .

Note the extremal cases: $\mathcal{E}^\emptyset = \mathcal{E}$, and $\mathcal{E}^v = \emptyset$ if and only if v is a maximal configuration of \mathcal{E} .

Notations: We note \mathcal{V}^v for short instead of $\mathcal{V}_{\mathcal{E}^v}$ to denote the poset of finite configurations of \mathcal{E}^v . Similarly, $\bar{\mathcal{V}}^v$ denotes the poset of configurations of \mathcal{E}^v .

Example. Let \mathcal{E} be a tree of events, and let v be a configuration of \mathcal{E} , i.e., v is a path in \mathcal{E} . If v is infinite, then $\mathcal{E}^v = \emptyset$, otherwise v can be written as $v = \{e_1, \dots, e_n\}$ with $e_1 < \dots < e_n$. The future \mathcal{E}^v is given by: $\mathcal{E}^v = \{e \in \mathcal{E} : e_n < e\}$. See an illustration in Fig. 2.

Example. In general, and because of concurrency, events in the future \mathcal{E}^v need not be causally related to events of v . This is illustrated in Fig. 3. Indeed, for $v = \{e_2, e_5\}$, events e_3 and e_6 belong to \mathcal{E}^v without being causally related to any event of v .

Concatenation. It follows from Definition 2.1 that, for any two configurations, u of \mathcal{E} , and v of \mathcal{E}^u , the union of subsets $u \cup v$ is a configuration of \mathcal{E} . To distinguish this kind of union from union of compatible configurations of \mathcal{E} , we call $u \cup v$ the *concatenation* of u and v , and we use the following special (non commutative) notation:

$$u \oplus v = u \cup v, \quad \text{only defined for } u \in \bar{\mathcal{V}} \text{ and } v \in \bar{\mathcal{V}}^u.$$




Fig. 2. Left, a tree of events. Immediate successors of a same events are pairwise in conflict. The future of $v = \{e_1, e_3\}$ is depicted on the right.




Fig. 3. An event structure \mathcal{E} and a configuration v of \mathcal{E} . The future of $v = \{e_2, e_5\}$ is shown on the right.

When it is well defined, it satisfies $u \oplus v \supseteq u$. Conversely, for any configuration w containing u , $w \setminus u$ is a configuration of \mathcal{E}^u , which is the “tail of w after u .” We use the following notation:

$$w \ominus u = w \setminus u, \quad \text{defined for all } u, w \in \bar{\mathcal{V}} \text{ such that } u \subseteq w.$$

To summarize, if the following objects are well-defined, we have:

$$u \oplus v \in \bar{\mathcal{V}}, \quad w \ominus u \in \bar{\mathcal{V}}^u, \quad (u \oplus v) \ominus u = v.$$

Clearly, the following formula holds, for the composition of futures:

$$\forall u \in \bar{\mathcal{V}}, \quad \forall v \in \bar{\mathcal{V}}^u, \quad (\mathcal{E}^u)^v = \mathcal{E}^{u \oplus v}. \quad (1)$$

Pre-regular event structures. As event structures of particular interest, we find the event structures arising from unfoldings of 1-safe Petri nets [5]. The systematic analysis of the notions introduced here, applied to the case of unfoldings is out of the scope of this paper—this is the topic of *Markov nets* [12]. However, we define pre-regular event structures in order to capture an important property of unfoldings. According to the following terminology, unfoldings of 1-safe Petri nets are uniformly pre-regular. We choose this terminology since pre-regularity (actually, uniform pre-regularity) is a condition for an event structure to be *regular* in the sense of Thiagarajan [13].

Definition 2.2. We say that \mathcal{E} is *pre-regular* if, for every finite configuration u of \mathcal{E} , the set $\text{Min}_{\preceq}(\mathcal{E}^u)$ is finite. We say that \mathcal{E} is *K-uniformly pre-regular* if for any finite configuration u of \mathcal{E} , $\text{Min}_{\preceq}(\mathcal{E}^u)$ has at most K elements (in which case this is true for any configuration u). We say that \mathcal{E} is *uniformly pre-regular* if there is a constant K such that \mathcal{E} is K -uniformly pre-regular.

If \mathcal{E} is pre-regular, it follows from the composition formula (1) that the future \mathcal{E}^u of any finite configuration u is then pre-regular. If \mathcal{E} is K -uniformly pre-regular, this is also the case for any future \mathcal{E}^u .

2.3. Stopping prefixes and stopped configurations

In probability, the notion of *choice* is central, as the very purpose of probabilities is to randomize choices. Choice is therefore a key concept in this paper. Choice in event structures relies on the




Fig. 4. An event structure with minimal conflicts depicted by curved arcs. Some stopping prefixes are depicted by dashed frames.

notion of immediate conflict we recall next. The *immediate conflict* relation is the following binary relation on \mathcal{E} , denoted by $\#_\mu$, and defined by:

$$\forall e_1, e_2 \in \mathcal{E}, \quad e_1 \#_\mu e_2 \iff \# \cap (\downarrow e_1 \times \downarrow e_2) = \{(e_1, e_2)\}.$$

Informally, “stopping” is the action of “cutting” a prefix of an event structure in such a way that choices remain internal to the considered prefix.

Definition 2.3. A prefix B is called a **stopping prefix** if B is $\#_\mu$ -closed, i.e., if for all $(e_1, e_2) \in B \times \mathcal{E}$, $e_1 \#_\mu e_2 \Rightarrow e_2 \in B$.

If B is a stopping prefix, a configuration v of \mathcal{E} is called **B -stopped** if v is a maximal configuration of B . A configuration v of \mathcal{E} is called **stopped** if there is a stopping prefix B such that v is B -stopped.

We give some examples to illustrate the notions of stopping prefix and of stopped configuration.

Example. Let \mathcal{E} be the event structure consisting of $\mathcal{E} = \{e_1, e_2, e_3\}$ with empty causality relation, and with conflict relation defined by $e_1 \# e_2$ and $e_2 \# e_3$. Then these conflict are also minimal conflicts, and therefore, the only two stopping prefixes of \mathcal{E} are \emptyset and \mathcal{E} itself. It follows that stopped configurations of \mathcal{E} are either \emptyset or maximal configurations of \mathcal{E} . The latter are $(e_1 e_3)$ and (e_2) . Here, configuration (e_1) is an example of configuration which is not stopped.

Example. Fig. 4 depicts stopping prefixes of an event structure.

The following example analyzes the case of trees of events. Confusion-free event structures are treated in Section 2.6.

Example. If \mathcal{E} is a tree of events, a prefix B is a stopping prefix if and only if B satisfies: for every event $e \in B$, if v_e denotes the configuration $v_e = \downarrow e \setminus \{e\}$, B contains all the events minimal in \mathcal{E}^{v_e} . Fig. 5, left, depicts a stopping prefix in a tree of events, while the prefix depicted on the right is not a stopping prefix. We see that stopping prefixes are given by unions of groups of events that can be simultaneously enabled.

Stopping prefixes satisfy the following crucial property:

Lemma 2.4. Recall that Ω_B denotes the set of maximal configurations of B .

(1) For every stopping prefix B and every maximal configuration ω , the intersection $\omega \cap B$ is a maximal configuration of B . Hence, every stopping prefix B induces a mapping:

$$\pi_B : \Omega \rightarrow \Omega_B, \quad \omega \mapsto \pi_B(\omega) = \omega \cap B.$$




Fig. 5. Left: a stopping prefix of a tree of events. Right: a prefix which is *not* a stopping prefix.

(2) For every pair B, B' of stopping prefixes with $B \subseteq B'$, we have a mapping:

$$\pi_{B,B'} : \Omega_{B'} \rightarrow \Omega_B, \quad v \mapsto \pi_{B,B'}(v) = v \cap B,$$

making the following diagram commutative:

$$\begin{array}{ccc}
 \Omega & \xrightarrow{\pi_{B'}} & \Omega_{B'} \\
 \searrow \pi_B & & \downarrow \pi_{B,B'} \\
 & & \Omega_B
 \end{array} \tag{2}$$

Proof. 1. We only have to show that, for every $\omega \in \Omega$, $\omega \cap B$ is maximal in B . Put $\omega_B = \omega \cap B$, and assume that $\omega_B \notin \Omega_B$. Then there is an event $e \in B$ such that $e \notin \omega_B$ and $\omega_B \cup \{e\}$ is a configuration of B . Since $e \in B \setminus \omega_B$, e is not an event of ω . Since ω is maximal, this implies that e is incompatible with ω . Two incompatible configurations contain events in immediate conflict (see Lemma 2.5). Therefore, there are events $x \in \downarrow e$ and $y \in \omega$ such that $x \#_\mu y$. B is $\#_\mu$ -closed, and since $x \in B$, this implies that $y \in B$, and thus $y \in \omega_B$. But then, $\omega_B \cup \{e\}$ contains the events x and y which are in conflict, contradicting that $\omega_B \cup \{e\}$ is a configuration. This shows that $\omega_B \in \Omega_B$.

2. Point 2 follows from Point 1, applied with $\mathcal{E} = B'$. The diagram (2) is obviously commutative (by associativity of “ \cap ”). \square

Remark. The key point in Lemma 2.4 is that, for every maximal configuration ω , $\omega \cap B$ is maximal in B . This is not the case in general if B is any prefix. Take for instance the event structure $\mathcal{E} = \{a, b\}$ with $a \# b$, the prefix $P = \{a\}$, and $\omega = \{b\}$. Then $\omega \cap P = \emptyset$ is not maximal in P .

2.4. Concurrent stopping prefixes

Stopping prefixes are defined in such a way that choices performed in a stopping prefix B remains internal to B . To formalize this, we show first that disjoint stopping prefixes are concurrent. As a consequence, configurations of disjoint stopping prefixes do not interact with each other. That is to say, every configuration of some stopping prefix B , seen as a choice made in B , is compatible with any choice made “beside” B .

We begin by recalling a well-known result.

Lemma 2.5. *If v, v' are two incompatible configurations, then there are events $e \in v$ and $e' \in v'$ such that $e \#_\mu e'$.*

Then the following lemma provides the key step for studying disjoint stopping prefixes.

Lemma 2.6. *Let P be a prefix of \mathcal{E} , and let B be a stopping prefix of \mathcal{E} . Assume that $P \cap B = \emptyset$. Then (e, f) holds for every pair of events $(e, f) \in P \times B$.*

Proof. Let $(e, f) \in P \times B$. Since $P \cap B = \emptyset$, and since both P and B are downward closed, e and f are not causally related. Assume that $e \# f$. Then, according to Lemma 2.5, there are events $x \in \downarrow e$ and $y \in \downarrow f$ such that $x \#_{\mu} y$. Then y belong to B . Since B is $\#_{\mu}$ -closed, x also belongs to B , and thus $x \in P \cap B$, a contradiction. Thus, e and f are not causally related and neither in conflict, hence *ecof*. \square

We can now state the result showing that choices performed in disjoint stopping prefixes do not interact with each other. Hence, for stopping prefixes, concurrency fits independence.

Proposition 2.7. *Let B be a stopping prefix of \mathcal{E} , given as a union of distinct stopping prefixes $B = \bigcup_{i \in I} B_i$, where I is some set of indices. Then the sets of configurations $\bar{\mathcal{V}}_B$ and of maximal configurations Ω_B of B , respectively, decompose as:*

$$\bar{\mathcal{V}}_B = \prod_{i \in I} \bar{\mathcal{V}}_{B_i}, \quad \Omega_B = \prod_{i \in I} \Omega_{B_i}. \quad (3)$$

Proof. Let $B = \bigcup_{i \in I} B_i$, with B_i distinct stopping prefixes. Consider the following mapping:

$$\Phi : \bar{\mathcal{V}}_B \rightarrow \prod_{i \in I} \bar{\mathcal{V}}_{B_i}, \quad v \mapsto \Phi(v) = (v \cap B_i)_{i \in I}. \quad (4)$$

Φ is indeed well defined, since it is clear that $v \cap B_i \in \bar{\mathcal{V}}_{B_i}$ for each $v \in \bar{\mathcal{V}}_B$. Then Φ is injective, since we have the reconstruction formula $v = \bigcup_{i \in I} v \cap B_i$ for all $v \in \bar{\mathcal{V}}_B$. Φ is also surjective. Indeed, for every element $(v_i)_{i \in I} \in \prod_{i \in I} \bar{\mathcal{V}}_{B_i}$, the subset:

$$v = \bigcup_{i \in I} v_i,$$

is a prefix of B , and Lemma 2.6 implies that v is conflict-free. Hence v is a configuration of B . Since the B_i are pairwise disjoint, we get that $\Phi(v) = (v_i)_{i \in I}$, which shows that Φ is surjective, and thus bijective. We equip the product $\prod_{i \in I} \bar{\mathcal{V}}_{B_i}$ with the product order (i.e., $(u_i)_{i \in I} \leq (v_i)_{i \in I}$ if $u_i \subseteq v_i$ for all $i \in I$). This makes Φ an isomorphism of partial orders, and thus:

$$\bar{\mathcal{V}}_B = \prod_{i \in I} \bar{\mathcal{V}}_{B_i}.$$

In particular, Φ respects maximal elements. We obtain thus by restriction of Φ to Ω_B the identification:

$$\Omega_B = \prod_{i \in I} \Omega_{B_i}.$$

This completes the proof of the proposition. \square

2.5. Locally finite event structures

The set of stopping prefixes is obviously a complete lattice, and the event structure \mathcal{E} is itself a stopping prefix. Therefore, for every event $e \in \mathcal{E}$, there exists a unique minimal stopping prefix that contains e , namely the intersection of all stopping prefixes containing e . We denote this stopping prefix by $B(e)$. A typical difficulty with concurrency models is that, in general, stopping prefixes $B(e)$ can be infinite. The following restriction is considered:

Definition 2.8. An event structure \mathcal{E} is called *locally finite* if for every event e , there exists a *finite* stopping prefix of \mathcal{E} containing e . The lattice of finite stopping prefixes of \mathcal{E} is denoted by \mathcal{B} .

Equivalently, $B(e)$ is finite for every event $e \in \mathcal{E}$. Equivalently also, for every finite set $A \subseteq \mathcal{E}$, there is a finite stopping prefix B containing A . The lattice of finite stopping prefixes plays a fundamental role for locally finite event structures.

Remark (*finite and finitely stopped configurations are the same*). For a general event structure, if v is a finite stopped configuration, it is generally not true that v is B -stopped for B a *finite* stopping prefix. We should thus distinguish between *finite stopped configurations* and *finitely stopped configurations* (those $v \in \Omega_B$ for some finite stopping prefix B). However the two notions coincide if \mathcal{E} is locally finite. Indeed, let v be finite and stopped, with $v \in \Omega_C$ (here, stopping prefix C may not be finite). Since v is finite, and since we assume that \mathcal{E} is locally finite, the smallest finite stopping prefix B that contains v is finite and satisfies $B \subseteq C$. According to Lemma 2.4, Point 2, it implies that $v \cap B$ is maximal in B , and since $v \subseteq B$ by construction, we finally get: $v \in \Omega_B$, what was to be shown. This justifies that we refer, for locally finite event structures, to *finite stopped configurations*, without any further precaution.

To show that local finiteness is stable when taking the future, first observe the following:

Lemma 2.9. *If B is a stopping prefix of \mathcal{E} , then $B \cap \mathcal{E}^v$ is a stopping prefix of \mathcal{E}^v for every configuration v of \mathcal{E} .*

Proof. Denote by $\#_\mu^v$ the immediate conflict relation in \mathcal{E}^v . Then the lemma follows from the following identity:

$$\#_\mu^v = \#_\mu \cap (\mathcal{E}^v \times \mathcal{E}^v). \quad (5)$$

□

From the above result, we immediately deduce:

Proposition 2.10. *If \mathcal{E} is locally finite, then \mathcal{E}^v is locally finite for every configuration v .*

Remark. Although the unfolding of a safe finite Petri net is always pre-regular, and even uniformly pre-regular, it is not true in general that the unfolding of a safe finite Petri net is locally finite. Fig. 6 depicts a uniformly pre-regular event structure that is not locally finite.

2.6. The particular case of confusion-free event structures

We open a parenthesis to illustrate the notions introduced above in the case of confusion-free event structures. Confusion-free event structures are defined as those event structures whose




Fig. 6. An event structure uniformly pre-regular, but non locally finite: the event at right-bottom is in immediate conflict with infinitely many events.

domain of configurations satisfies the so-called Q -axiom [5,14]. Equivalently, we shall follow [4, Prop. 2.4] and define an event structure \mathcal{E} to be confusion-free as follows: Let F be the binary relation $F = \#_{\mu} \cup D$, where D is the diagonal $D = \{(e, e), e \in \mathcal{E}\}$. \mathcal{E} is said to be confusion-free if F is transitive, and if the following holds:

$$\forall e, e' \in \mathcal{E}, \quad e \#_{\mu} e' \Rightarrow \downarrow e \setminus \{e\} = \downarrow e' \setminus \{e'\}. \tag{6}$$

It is known from [5] that so-called confusion-free Petri nets unfold to confusion-free event structures. Let \mathcal{E} be a confusion-free event structure. For $e \in \mathcal{E}$, define:

$$G(e) = \{e' \in \mathcal{E} : e F e'\}.$$

Then it is easy to verify that $B(e)$, the smallest stopping prefix of \mathcal{E} that contains an event $e \in \mathcal{E}$, is given by:

$$B(e) = \bigcup_{e' \in \downarrow e} G(e'). \tag{7}$$

Assume moreover that \mathcal{E} is pre-regular. Then it follows from (6) that $G(e)$ is finite for every $e \in \mathcal{E}$. In turn, $B(e)$ is also finite for every $e \in \mathcal{E}$, and thus \mathcal{E} is locally finite. We have obtained the following: *Let \mathcal{E} be a confusion-free event structure. If \mathcal{E} is pre-regular, then \mathcal{E} is locally finite. As a corollary: The unfolding of a confusion-free Petri net is locally finite.*

We leave as an exercise to the reader to prove the following, by making use of the form (7) for stopping prefixes $B(e)$: *In a confusion-free event structure, every configuration is stopped.*

2.7. Recursive stopping

Next, we analyze the effect of concatenation on stopped configurations. The following example shows that the class of stopped configurations is not closed under concatenation in general. This motivates extending this class.

Example. Let \mathcal{E} be the event structure depicted in Fig. 7, left. \mathcal{E} has two nonempty stopping prefixes, $B_1 = \{e_1, e_2\}$ and $B_2 = \mathcal{E}$. Let $v_1 = (e_1)$; v_1 is B_1 -stopped. The future \mathcal{E}^{v_1} is depicted in Fig. 7, right.




Fig. 7. Left, an event structure \mathcal{E} . Right, the future of configuration (e_1) . All nonempty stopping prefixes are depicted by dashed frames. The concatenation of two stopped configurations is not stopped: (e_1e_3) is the concatenation of two stopped configurations, but (e_1e_3) is not stopped in \mathcal{E} .

Configuration $z = (e_3)$ is stopped in \mathcal{E}^{v_1} since $\{e_3\}$ is a stopping prefix of \mathcal{E}^{v_1} . However the concatenation $v = v_1 \oplus z = (e_1e_3)$ is not stopped in \mathcal{E} . Indeed, if v was stopped, then v would be maximal in $B_2 = \mathcal{E}$, which is not the case. Hence, the concatenation of two stopped configurations is not a stopped configuration in general.

Definition 2.11. A configuration v of \mathcal{E} is said to be *R-stopped* in \mathcal{E} (R for Recursively stopped) if for some integer $N > 0$ or for $N = \infty$, there is a non-decreasing sequence $(v_n)_{0 \leq n < N}$ of configurations with $v_0 = \emptyset$ and $v = \bigcup_{0 \leq n < N} v_n$, and such that:

$$\forall n \geq 0, \quad n < N \implies v_{n+1} \ominus v_n \text{ is finite stopped in } \mathcal{E}^{v_n}.$$

The sequence $(v_n)_{0 \leq n < N}$ is called a *valid decomposition* of v . If v has a valid decomposition with $N < \infty$, we say that v is *finite R-stopped*.

We denote by $\overline{\mathcal{W}}_{\mathcal{E}}$, or by $\overline{\mathcal{W}}$ for short if no confusion can occur, the set of R -stopped configurations of \mathcal{E} . $\mathcal{W}_{\mathcal{E}}$ and shortly \mathcal{W} denote the set of finite R -stopped configurations.

We use the same conventions as before in denoting by \mathcal{W}^v and $\overline{\mathcal{W}}^v$, respectively, the sets of finite R -stopped configurations and R -stopped configurations of the future \mathcal{E}^v .

Proposition 2.13 below relates R -stopped configurations of \mathcal{E} with R -stopped configurations in stopping prefixes of \mathcal{E} , and in futures of R -stopped configurations. For this we need the following lemma:

Lemma 2.12. *Let B be a stopping prefix of \mathcal{E} , and let v be a configuration of B . Then we have:*

- (1) D is a stopping prefix of $B^v \implies D$ is a stopping prefix of \mathcal{E}^v .
- (2) D is a stopping prefix of $\mathcal{E}^v \implies D \cap B$ is a stopping prefix of B^v .

Proof. 1. Since B is in particular a prefix of \mathcal{E} , it is immediate from Definition 2.1 that we have: $B^v = B \cap \mathcal{E}^v$. Therefore, by Lemma 2.9, B^v is a stopping prefix of \mathcal{E}^v . Point 1 follows then from the fact that a stopping prefix of a stopping prefix is a stopping prefix.

2. Let D be a stopping prefix of \mathcal{E}^v . Then $D \cap B$ is obviously a prefix of B^v . The immediate conflict relation in B^v is the restriction of $\#_{\mu}^v$ (immediate conflict in \mathcal{E}^v) to $B^v \times B^v$. Therefore $D \cap B$ is a stopping prefix of B^v . \square

Proposition 2.13.

(1) If B is a stopping prefix of \mathcal{E} , R -stopped configurations of B are those R -stopped configurations of \mathcal{E} contained in B . Moreover, we have:

$$v \in \mathcal{W}_{\mathcal{E}} \Rightarrow v \cap B \in \mathcal{W}_B.$$

(2) For every pair u, v of configurations, we have:

$$u \in \mathcal{W}_{\mathcal{E}}, v \in \mathcal{W}^u \Rightarrow u \oplus v \in \mathcal{W}_{\mathcal{E}}.$$

Proof. Let B be a stopping prefix of \mathcal{E} . Then a configuration v of B is R -stopped in B if and only if v is R -stopped in \mathcal{E} . Indeed, it follows from Lemma 2.12 above that valid decompositions are obtained from one another by:

$$(v_n)_n \rightarrow (v_n \cap B)_n, \quad (v_n)_n \rightarrow (v_n)_n.$$

This shows point 1. For point 2: the concatenation of valid decompositions is obviously a valid decomposition. \square

A more precise result than point 2 will be stated below, showing that any R -stopped configuration $w \in \mathcal{W}_{\mathcal{E}}$ containing $u \in \mathcal{W}_{\mathcal{E}}$ has the form $w = u \oplus v$ with $v \in \mathcal{W}^u$.

Finally, since finite stopped configurations are clearly R -stopped, R -stopped configurations form the smallest class of configurations containing all finite stopped configurations, and closed under concatenation. It is not clear at this point whether stopped configurations in general are R -stopped, and in particular if maximal configurations of \mathcal{E} are R -stopped. This will be examined in Section 3.5.

3. Branching cells

So far, we have considered stopped configurations, and R -stopped configurations that are obtained by concatenations of stopped configurations. However, a given R -stopped configuration shall certainly have several valid decompositions. Branching cells, we introduce in this section, will allow decomposing R -stopped configurations in a canonical way.

3.1. Initial stopping prefixes

Definition 3.1. We say that a stopping prefix B is an *initial stopping prefix* of \mathcal{E} if B is nonempty, and if \emptyset is the only stopping prefix strictly included in B .

Hence B is initial if B is minimal among nonempty prefixes of \mathcal{E} .

Although we will latter on focus on event structures that are both locally finite and pre-regular, we state the following result in a more general case. This shows that the foundations of our approach do not collapse when we relax the local finiteness assumption.

Theorem 3.2. *If they exist, initial stopping prefixes of \mathcal{E} are disjoint. If \mathcal{E} satisfies one of the following two conditions:*

- (1) \mathcal{E} is locally finite, or
- (2) \mathcal{E} is pre-regular,

then every nonempty stopping prefix of \mathcal{E} contains an initial stopping prefix.

Proof. Since stopping prefixes are stable under intersection, distinct minimal nonempty stopping prefixes, if they exist, are disjoint.

If \mathcal{E} is locally finite, any nonempty stopping prefix B contains a finite nonempty stopping prefix C . Then C contains an initial stopping prefix, and thus the same holds for B .

We now show that the same holds if \mathcal{E} is pre-regular. Let B be a nonempty stopping prefix of a pre-regular event structure \mathcal{E} . Let \mathcal{B}^\star be the set of nonempty stopping prefixes included in B . \mathcal{B}^\star is nonempty. We use Zorn's Lemma to show that \mathcal{B}^\star has a minimal element. Such a minimal element will be an initial stopping prefix included in B .

Hence, let $(I, <)$ be some totally ordered set, and let $(B_i)_{i \in I}$ be a decreasing family in \mathcal{B}^\star , indexed by I . That is, $i > j \Rightarrow B_i \subseteq B_j$. We show that $C = \bigcap_{i \in I} B_i$ is a lower bound in \mathcal{B}^\star for the family $(B_i)_{i \in I}$. Since C is a stopping prefix, we only have to show that C is nonempty. Assume that $C = \emptyset$. Fix ω a maximal configuration of \mathcal{E} . Then $\omega \cap B$ is maximal in B , thanks to Lemma 2.4, and in particular, $\omega \cap B \neq \emptyset$. Pick e_0 an event minimal in $\omega \cap B$. By induction, we construct a sequence of events $(e_n)_{n \geq 0}$, and an increasing sequence of indices $(i_n)_{n \geq 0}$, such that e_n is a minimal event of $\omega \cap B_{i_n}$, and $e_0, \dots, e_{n-1} \notin B_{i_n}$. In particular the events e_n are pairwise distinct. Since they are all minimal in \mathcal{E} , this contradicts that \mathcal{E} is pre-regular. This contradiction shows that $C \neq \emptyset$, what was to be shown. \square

We leave as an exercise to the reader to construct an event structure that does not have any initial stopping prefix.

The above result specializes as follows.

Proposition 3.3. *If \mathcal{E} is locally finite, every initial stopping prefix of \mathcal{E} is finite. If \mathcal{E} is pre-regular, initial stopping prefixes of \mathcal{E} are finitely many. If \mathcal{E} is K -uniformly pre-regular, the number of initial stopping prefixes of \mathcal{E} is lesser than or equal to K .*

Proof. It is obvious that initial stopping prefixes of \mathcal{E} are finite if \mathcal{E} is locally finite. Observe that each initial stopping prefix of \mathcal{E} contains events minimal in \mathcal{E} . Pick one of those minimal events e_c for each initial stopping prefix c . Then the e_c are pairwise distinct since the initial stopping prefixes are pairwise disjoint. The remaining of the proposition follows. \square

Example. In Fig. 4, the initial stopping prefixes are depicted as follows: the left dashed frame, and the smallest of the right dashed frames. Remark that:

- (1) some minimal event may not belong to any initial stopping prefix;
- (2) some events of an initial stopping prefix may not be minimal.

Since initial stopping prefixes are disjoint, we get, as a particular case of Proposition 2.7:

Proposition 3.4. *Let B be a stopping prefix of \mathcal{E} given by the union of finite family of initial stopping prefixes $(c_i)_{i \in I}$. Then Ω_B decomposes as the following product:*

$$\Omega_B = \prod_{i \in I} \Omega_{c_i}.$$

3.2. Branching cells

We will now exclusively focus on the case of locally finite event structures. Throughout the remaining of the paper, the following assumption is in force:

Assumption 3.5. Event structure \mathcal{E} is locally finite.

According to Proposition 2.10, all futures \mathcal{E}^v are then locally finite. In turn, Theorem 3.2 shows that this is a sufficient condition to guarantee that any \mathcal{E}^v has initial stopping prefixes whenever $\mathcal{E}^v \neq \emptyset$.

Definition 3.6. A **branching cell** of \mathcal{E} is any initial stopping prefix of \mathcal{E}^v , where v ranges over $\mathcal{W}_{\mathcal{E}}$ (i.e., over finite R -stopped configurations of \mathcal{E}). We denote by $\mathcal{C}_{\mathcal{E}}$, or by \mathcal{C} for short, the set of all branching cells of \mathcal{E} .

The set of branching cells that are initial stopping prefixes of \mathcal{E}^v , with $v \in \mathcal{W}_{\mathcal{E}}$, is denoted by $\delta_{\mathcal{E}}(v)$, or shortly $\delta(v)$. Branching cells in $\delta(v)$ are called the branching cells **enabled by** v .

Hence, $\delta(\emptyset)$ for example represents the set of initial stopping prefixes of the event structure.

As usual, \mathcal{C}^v shall denote the set of branching cells of the future \mathcal{E}^v , for any configuration v . Propositions 2.10 and 3.3 together have the following consequence (recall that we assume \mathcal{E} to be locally finite).

Proposition 3.7. *Every branching cell of \mathcal{E} is finite.*

It is not easy at this point to describe all branching cells of an event structure. This requires to examine all R -stopped configurations, but the definition that we have given suffers from a large combinatorial complexity. We will thus present examples only after having provided some more efficient ways to describe R -stopped configurations.

As a first property of branching cells, we examine how branching cells of \mathcal{E} are related to branching cells in stopping prefixes of \mathcal{E} , and in futures of R -stopped configurations.

Proposition 3.8. *If B is a stopping prefix of \mathcal{E} , then $\mathcal{C}_B \subseteq \mathcal{C}_{\mathcal{E}}$. If v is a finite R -stopped configuration of \mathcal{E} , then $\mathcal{C}^v \subseteq \mathcal{C}_{\mathcal{E}}$.*

Proof. This is an immediate consequence of Proposition 2.13. \square

3.3. Covering through branching cells

Lemma 3.9 below is the key of our study of branching cells. It shows that branching cells decompose R -stopped configurations in an intrinsic manner. The proof is postponed in Appendix 6.

Lemma 3.9. *Let v be some R -stopped configuration of \mathcal{E} . Then there exists a valid decomposition $(v_n)_{0 \leq n < N}$ of v , $N \leq \infty$, and a sequence of branching cells $(c_n)_{0 < n < N}$ such that, for every integer n with $0 \leq n < N - 1$:*

- (1) c_{n+1} is a branching cell enabled by v_n ;
- (2) $v_{n+1} \ominus v_n$ is maximal in c_{n+1} .

For any such pair of sequences $(v_n)_{0 \leq n < N}$ and $(c_n)_{0 < n < N}$, the c_n are pairwise disjoint. If $(v'_n)_{0 \leq n < N'}$ and $(c'_n)_{0 < n < N'}$ is another pair of such sequences then we have the equality of sets:

$$\{c_n, 0 < n < N\} = \{c'_n, 0 < n < N'\}.$$

In particular, $N = N'$.

We may thus define the covering of R -stopped configurations as follows:

Definition 3.10. The covering $\Delta_{\mathcal{E}}(v)$ (or $\Delta(v)$ for short) of a R -stopped configuration v is defined as the set of branching cells:

$$\Delta_{\mathcal{E}}(v) = \{c_n, 0 < n < N\},$$

where $(v_n)_{0 \leq n < N}$ and $(c_n)_{0 < n < N}$ are two sequences associated with v as in Lemma 3.9.

Example. Let \mathcal{E} be the event structure depicted in Fig. 8, top-left, and let ω be the maximal configuration given by $\omega = \{e_1, \dots, e_5\}$. Since ω is finite stopped, ω is R -stopped in \mathcal{E} . To find $\Delta(\omega)$, it is enough to follow any decomposition of ω as described in Lemma 3.9.

The initial stopping prefixes of \mathcal{E} are depicted by frames in Fig. 8, top-left. We start the decomposition for example with $c_1, v_1 = \omega \cap c_1 = (e_1)$. Then \mathcal{E}^{v_1} is depicted top-right. There is a unique initial stopping prefix $c_2 \in \delta(v_1)$, so that the next step is necessarily $v_2 = (e_1 e_2 e_3)$. The two following steps are depicted in bottom-left and bottom-right, respectively. Each step has a unique initial stopping prefix (c_3 and c_4 , respectively).




Fig. 8. Decomposition of $\omega = (e_1 e_2 e_3 e_4 e_5)$ through branching cells to determine $\Delta(\omega)$. The possible choices of branching cells at each step are depicted by rectangles.




Fig. 9. Showing the collection of branching cells $\Delta(\omega)$, for ω as in Fig. 8.

We obtain thus $\Delta(\omega) = \{c_1, \dots, c_4\}$. The collection of branching cells $\Delta(\omega)$ can be represented as in Fig. 9. Any enumeration of c_1, \dots, c_4 that stacks all of them in a “tetris-compliant” way corresponds to a valid decomposition of ω , namely: (c_1, c_2, c_3, c_4) , (c_2, c_1, c_3, c_4) and (c_2, c_3, c_1, c_4) . As an exercise, the reader can verify that each of these enumerations corresponds indeed to a valid decomposition of ω . Observe the invariance of the branching cells encountered.

Example (branching cells in confusion-free event structures). Let \mathcal{E} be a pre-regular and confusion-free event structure. Then every configuration is stopped, and therefore every finite configuration is R -stopped. Recall that we have defined in Section 2.6 the relation F on \mathcal{E} as the reflexive closure of $\#_\mu$. Then F is an equivalence relation. The following is readily checked: *Branching cells of \mathcal{E} are the equivalence classes of F .* We recognize thus in the branching cells of \mathcal{E} the cells defined in [4] for confusion-free event structures. In particular, branching cells of a confusion-free event structures globally do not overlap. The following example shows that this is not the case in general.

Example (branching cells may overlap). Lemma 3.9 states that branching cells involved in the decomposition of a given configuration are disjoint. However, in general, the whole collection of branching cells of an event structure may contain branching cells $c \neq c'$ such that $c \cap c' \neq \emptyset$. This is shown by the following example:

Consider the event structure \mathcal{E} depicted in Fig. 7, left. Fig. 7-right depicts two branching cells $c = \{e_3\}$ and $c' = \{e_5\}$ of \mathcal{E} , obtained by $\delta(e_1) = \{c, c'\}$. Consider the stopped configuration (e_2) . The future of (e_2) is given by $\mathcal{E}^{(e_2)} = \{e_3, e_4, e_5\}$, with an empty causality relation and with $e_3 \# e_4 \# e_5$. $\mathcal{E}^{(e_2)}$ has thus a unique branching cell $c'' = \{e_3, e_4, e_5\}$, which intercepts c and c' without being equal to c nor to c' .

This example shows that *branching cells of an event structure may globally overlap*. We interpret this fact by saying that *the decomposition through branching cells is dynamic*. Indeed, a same event may belong to different branching cells. Which branching cell is actually selected in an execution including this event depends on the execution (until a certain extend), not only on the event itself.

3.4. Properties of the covering

We shall now study the properties of the covering map. Following our usual method, we study the relationship between the covering map $\Delta_{\mathcal{E}}$ and the analogous Δ_B and Δ^u defined, respectively, in a stopping prefix B and in the future \mathcal{E}^u of some $u \in \mathcal{W}_{\mathcal{E}}$. The proof of the following theorem is found in Appendix 6.

Theorem 3.11. *Let B be a stopping prefix of \mathcal{E} , and let $u \in \mathcal{W}_{\mathcal{E}}$.*

(1) *The covering Δ_B defined on \mathcal{W}_B coincides with the restriction of $\Delta_{\mathcal{E}}$ to \mathcal{W}_B . In symbols:*

$$\forall v \in \mathcal{W}_B, \quad \Delta_B(v) = \Delta_{\mathcal{E}}(v).$$

(2) *Let Δ^u be the covering associated with \mathcal{E}^u , and defined on \mathcal{W}^u . Then we have, for any $v \in \mathcal{W}^u$:*

$$\Delta(u \oplus v) = \Delta(u) \cup \Delta^u(v), \quad \Delta(u) \cap \Delta^u(v) = \emptyset. \quad (8)$$

(3) *The covering map covers R -stopped configurations, i.e.:*

$$\forall u \in \mathcal{W}, \quad u = \bigcup_{c \in \Delta(u)} u \cap c. \quad (9)$$

Moreover, $u \cap c \in \Omega_c$ for each $c \in \Delta(u)$, and:

$$\forall c, c' \in \Delta(u), \quad c \neq c' \Rightarrow c \cap c' = \emptyset. \quad (10)$$

(4) *The covering has the following expression, for u a finite R -stopped configuration:*

$$\forall u \in \mathcal{W}_{\mathcal{E}}, \quad \Delta(u) = \{c \in \delta(w) : w \in \mathcal{W}, w \subseteq u\} \setminus \delta(u), \quad (11)$$

and for u any R -stopped configuration:

$$\forall u \in \overline{\mathcal{W}}_{\mathcal{E}}, \quad \Delta(u) = \bigcup_{w \in \mathcal{W}, w \subseteq u} \Delta(w). \quad (12)$$

(5) *Let w be a configuration of \mathcal{E} . The following set of sub-configurations of w :*

$$F_w = \{v \in \overline{\mathcal{W}} : v \subseteq w\}$$

is a lattice. Moreover, if u, v are two R -stopped configurations of \mathcal{E} such that $u \subseteq v$, then $v \ominus u$ is R -stopped in \mathcal{E}^u .

As an application, we derive the following result which is quite intuitive, but not obvious when inspecting directly Definition 3.10.

Corollary 3.12. *Let $B = \bigcup_{i \in I} c_i$ be a stopping prefix given by a finite union of pairwise distinct initial stopping prefixes $(c_i)_{i \in I}$. For each $i \in I$, let $z_i \in \Omega_{c_i}$, and let $v = \bigcup_{i \in I} z_i$. Then v is R -stopped, and the covering of v is given by $\Delta(v) = \{c_i, i \in I\}$.*

Proof. Since B is a finite union of finite prefixes (each c_i is finite according to Proposition 3.7), B is finite. It is easy to check that v is maximal in B . Therefore v is a finite stopped configuration, and in particular v is R -stopped.

We determine $\Delta(v)$ as follows. For each $i \in I$, z_i is R -stopped, and $z_i \subseteq v$. Therefore, it follows from point 5 of Theorem 3.11 that $v \ominus z_i$ is R -stopped in \mathcal{E}^{z_i} . Applying (8), we obtain:

$$\Delta(v) = \Delta(z_i) \cup \Delta^{z_i}(v \ominus z_i).$$

Since $\Delta(z_i) = \{c_i\}$, we have in particular $c_i \in \Delta(v)$. Observe that $v \cap c_i = z_i$ for each $i \in I$, since z_i is maximal in c_i . Therefore $v = \bigcup_{i \in I} (v \cap c_i)$. It follows from point 3 in Theorem 3.11 that any branching cell $c \in \Delta(v)$ must be one of the c_i , $i \in I$. We conclude that $\Delta(v) = \{c_i, i \in I\}$, what was to be shown. \square

3.5. Max-initial decomposition

So far we have studied R -stopped configurations without knowing “how far” they may go. In other words, we still do not know whether maximal configurations, for instance, are R -stopped. It turns out that the answer is “yes.” In getting this answer, we make a critical use of the local finiteness assumption. A key step is to introduce some particular decomposition of maximal configurations that we call the *max-initial decomposition*. This construction has also useful applications for the study of Markov nets—this is beyond the scope of the present paper.

For simplicity, we restrict ourselves to the case of a pre-regular event structures \mathcal{E} . We still assume also that \mathcal{E} is locally finite.

Definition 3.13. Let \mathcal{E} be a pre-regular event structure. The *max-initial stopping prefix* of \mathcal{E} is the union of all the initial stopping prefixes of \mathcal{E} . We denote it by $B_0(\mathcal{E})$. We also take the convention that $B_0(\emptyset) = \emptyset$.

According to Proposition 3.3, since \mathcal{E} is both locally finite and pre-regular, initial stopping prefixes are finite and finitely many. Therefore, $B_0(\mathcal{E})$ itself is finite. More generally, $B_0(\mathcal{E}^v)$ is finite for every configuration v of \mathcal{E} .

Theorem 3.14. Let \mathcal{E} be a pre-regular event structure. Every maximal configuration ω is R -stopped. A valid decomposition of $\omega \in \Omega$ is given by the sequence $(v_n)_{n \geq 0}$ defined by:

$$v_0 = \emptyset, \quad \forall n \geq 0, \quad v_{n+1} = v_n \oplus z_{n+1}, \quad z_{n+1} = \omega \cap B_0(\mathcal{E}^{v_n}),$$

where $B_0(\mathcal{E}^{v_n})$ denotes the max-initial stopping prefix of \mathcal{E}^{v_n} . The sequence $(v_n)_{n \geq 0}$ is called the *max-initial decomposition* of ω .

Proof. We have $v_n \subseteq \omega$ for each $n \geq 0$. Moreover, for each $n \geq 0$, we have $z_{n+1} = (\omega \setminus v_n) \cap B_0(\mathcal{E}^{v_n})$. Since ω is maximal in \mathcal{E} , $\omega \setminus v_n$ is maximal in \mathcal{E}^{v_n} , and therefore, by Lemma 2.4, z_{n+1} is maximal in $B_0(\mathcal{E}^{v_n})$, which is a finite stopping prefix of \mathcal{E}^{v_n} . To get that $(v_n)_{n \geq 0}$ is a valid decomposition of ω , it remains thus only to show the following:

$$\omega \subseteq \bigcup_{n \geq 0} v_n. \tag{13}$$

The proof of (13) decomposes in three steps. Let v denote the configuration $v = \bigcup_{n \geq 0} v_n$.

Step 1. We claim that (13) holds if \mathcal{E} is finite. Indeed, assume that \mathcal{E} is finite. Since $(v_n)_{n \geq 0}$ is non-decreasing, there is an integer $N \geq 0$ such that $v = v_N = v_{N+1}$. Then $z_{N+1} = v_{N+1} \ominus v_N = \emptyset$. Since

z_{N+1} is maximal in $B_0(\mathcal{E}^{v_N})$, this implies that $B_0(\mathcal{E}^{v_N}) = \emptyset$, which in turns implies that $\mathcal{E}^{v_N} = \emptyset$, i.e., v_N is maximal in \mathcal{E} . Since $\omega \supseteq v_N$, we get that $v = v_N = \omega$.

Step 2. Let B be a finite stopping prefix of \mathcal{E} , and let $\omega_B = \omega \cap B$. Then we claim that the max-initial decomposition $(v'_n)_{n \geq 0}$ of ω_B is given by $v'_n = v_n \cap B$. Indeed, this is a consequence of Proposition 2.13.

Step 3. Let B be a finite stopping prefix of \mathcal{E} , and let $(v'_n)_{n \geq 0}$ be the max-initial decomposition of $\omega_B = \omega \cap B$. Then we have, according to Step 2:

$$v = \bigcup_{n \geq 0} v_n \supseteq \bigcup_{n \geq 0} (v_n \cap B) = \bigcup_{n \geq 0} v'_n = \omega_B,$$

the latter equality by Step 1. Since this holds for any finite stopping prefix B of \mathcal{E} , we have:

$$v \supseteq \bigcup_{B \in \mathcal{B}} \omega \cap B, \tag{14}$$

where B ranges over the lattice of finite stopping prefixes of \mathcal{E} . Now let $e \in \omega$. Since \mathcal{E} is locally finite, there is a finite stopping prefix D such that $e \in D$. We have $\omega \cap D \ni e$, and thus, from (14), $e \in v$. Since this holds for any $e \in \omega$, we conclude that $\omega \subseteq v$, which is (13). This completes the proof. \square

Corollary 3.15. *Every stopped configuration of a pre-regular (and locally finite) event structure is R-stopped.*

We illustrate the theorem on two examples.

Example. Consider the event structure depicted in Fig. 8. We have already examined the decompositions of $\omega = (e_1 \dots e_5)$ through branching cells. The max-initial decomposition of ω is given by: $v_0 = \emptyset$, $z_1 = (e_1 e_2 e_3)$, $v_1 = z_1$, involving branching cells c_1 and c_2 , then $z_2 = (e_4)$, $v_2 = z_1 \oplus z_2$, involving branching cell c_3 , and finally $z_3 = e_5$, $v_3 = z_1 \oplus z_2 \oplus z_3 = \omega$, involving branching cell c_4 .

Example (decomposition in trees of events). If \mathcal{E} is a tree of events, and if $\omega = (e_1, e_2, \dots)$, the max-initial decomposition $(v_n)_{n \geq 0}$ of ω is given by $v_n = (e_1, \dots, e_n)$, for $n \geq 0$.

4. Probabilistic event structures

In this section, we define probabilistic event structures. Then, we develop the key tool that allows us reducing the construction of *locally finite* probabilistic event structures to that of *finite* probabilistic event structures. We list in Table 2 the notations used throughout Sections 4–5.

4.1. Definition of probabilistic event structures

We first recall some basic definitions from probability theory (see for example [15]). Then, we apply these definitions to the case of event structures.

Probability spaces. A σ -algebra \mathfrak{F} on a set Ω is a collection of subsets of Ω , such that $\emptyset \in \mathfrak{F}$, \mathfrak{F} is stable under complement, and \mathfrak{F} is stable under countable union. A *measurable space* is a pair (Ω, \mathfrak{F}) , where \mathfrak{F} is a σ -algebra on Ω . The elements of \mathfrak{F} are called the \mathfrak{F} -measurable subsets of Ω , or

Table 2
Notations for probabilistic event structures and distributed probabilities

Symbol	Meaning
Probabilistic event structures	
$(\mathcal{E}, \mathbb{P})$	Probabilistic event structures
p	Likelihood of \mathbb{P} , $p(v) = \mathbb{P}(\mathcal{S}(v))$
$\mathcal{S}(v)$	Shadow of configuration v
Probabilistic future	
$(\mathcal{E}^u, \mathbb{P}^u)$	Probabilistic future of u , if $p(u) > 0$
p^u	Likelihood of \mathbb{P}^u , $p^u(v) = \frac{1}{p(u)} p(u \oplus v)$
Distributed products	
$(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$	Locally randomized event structure
p	Likelihood associated with $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$
\mathbb{P}	Distributed product associated with $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$
$Z_c : \Omega \rightarrow \Omega_c$	Random variable defined for $c \in \delta(\emptyset)$
$Z = (Z_c)_{c \in \delta(\emptyset)}$	Product random variable
$Z_c^v : \mathcal{S}(v) \rightarrow \Omega_c$	Conditional random variable, for $c \in \delta(v)$
$Z^v = (Z_c^v)_{c \in \delta(v)}$	Product random variable

shortly the *measurable* subsets if no confusion can occur on the σ -algebra. A *measure* on a measurable space (Ω, \mathfrak{F}) is a real valued function $\mathbb{P} : \mathfrak{F} \rightarrow \mathbb{R}$ such that $\mathbb{P}(\emptyset) = 0$, $\mathbb{P}(A)$ is nonnegative for every $A \in \mathfrak{F}$, and such that for every sequence of pairwise disjoint sets $A_n \in \mathfrak{F}$:

$$\mathbb{P}\left(\bigcup_{n \geq 0} A_n\right) = \sum_{n \geq 0} \mathbb{P}(A_n).$$

Finally, \mathbb{P} is said to be a *probability measure* if $\mathbb{P}(\Omega) = 1$, and in this case $(\Omega, \mathfrak{F}, \mathbb{P})$ is called a *probability space*. If the singletons $\{x\}$ are measurable, we simply note $\mathbb{P}(x)$ for $\mathbb{P}(\{x\})$. If Ω is finite, we usually consider the *discrete σ -algebra* on Ω , which is just the powerset $\wp(\Omega)$ of Ω . In this case, the singletons are measurable, and \mathbb{P} is entirely determined by the values $\mathbb{P}(x)$, for x ranging over Ω . We shortly say that \mathbb{P} is a *finite probability*.

Let (Ω, \mathfrak{F}) and (Ω', \mathfrak{F}') be two measurable spaces. Following the traditional terminology from Probability theory, we say that a mapping $f : \Omega \rightarrow \Omega'$ is a *random variable* if it is $\mathfrak{F}/\mathfrak{F}'$ -measurable, i.e., if $f^{-1}(A) \in \mathfrak{F}$ for every $A \in \mathfrak{F}'$. If (Ω, \mathfrak{F}) is equipped with a probability \mathbb{P} , the set function $\mathbb{Q} : \mathfrak{F}' \rightarrow \mathbb{R}$ defined by $\mathbb{Q}(A) = \mathbb{P}(f^{-1}(A))$ is a probability on (Ω', \mathfrak{F}') , which is called the *image probability* of \mathbb{P} under f . \mathbb{Q} is also called the *law* of f under \mathbb{P} , and is denoted by $\mathbb{Q} = f\mathbb{P}$. This is indeed a left action on measures, i.e., $(f \circ g)\mathbb{P} = f(g\mathbb{P})$.

Probabilistic event structures. Let \mathcal{E} be an event structure, and denote as in Section 2 by Ω the set of maximal configurations of \mathcal{E} . Let τ be the restriction to Ω of the Scott topology on $\overline{\mathcal{V}}$, with $(\overline{\mathcal{V}}, \subseteq)$ seen as a DCPO [5,16]. We denote by \mathfrak{F} the **Borel σ -algebra on Ω** associated with τ . That is, \mathfrak{F} is the smallest σ -algebra on Ω that contains the (countable) collection of subsets of the form:

$$\mathcal{S}(v) \stackrel{[\text{def}]}{=} \{\omega \in \Omega : \omega \supseteq v\}, \tag{15}$$

where v ranges over the set of finite configurations of \mathcal{E} . Hence an event structure \mathcal{E} naturally defines a measurable space (Ω, \mathfrak{F}) . For every configuration v —not necessarily finite—, the subset $\mathcal{S}(v)$ defined by (15) is then measurable (write $\mathcal{S}(v)$ as the countable intersection of $\mathcal{S}(u)$, with u finite and contained in v). $\mathcal{S}(v)$ is called the *shadow* of v .

Definition 4.1. A *probabilistic event structure* is a pair $(\mathcal{E}, \mathbb{P})$, where \mathbb{P} is a probability measure on the measurable space (Ω, \mathfrak{F}) .

With this definition, the space (Ω, \mathfrak{F}) is interpreted as the *sample space* associated with event structure \mathcal{E} . Ω represents indeed the set of histories of the system modeled by \mathcal{E} .

The intuitive interpretation of a probabilistic event structure $(\mathcal{E}, \mathbb{P})$ is as follows: if v is any configuration of \mathcal{E} , the probability that v occurs in an execution of \mathcal{E} is given by the number $\mathbb{P}(\mathcal{S}(v))$. Whence the following definition:

Definition 4.2. If $(\mathcal{E}, \mathbb{P})$ is a probabilistic event structure, we define the *likelihood* associated with \mathbb{P} as the nonnegative function $p : \bar{\mathcal{V}} \rightarrow \mathbb{R}$ given by:

$$\forall v \in \bar{\mathcal{V}}, \quad p(v) = \mathbb{P}(\mathcal{S}(v)).$$

Let B be a stopping prefix of an event structure, and recall the mapping $\pi_B : \Omega \rightarrow \Omega_B, \omega \rightarrow \omega \cap B$, given by Lemma 2.4, point 1. Then π_B is measurable w.r.t. the Borel σ -algebras \mathfrak{F} and \mathfrak{F}_B on Ω and Ω_B , respectively. Assume that $(\mathcal{E}, \mathbb{P})$ is a probabilistic event structure. Then the image probability $\pi_B \mathbb{P}$ defines a probabilistic event structure $(B, \pi_B \mathbb{P})$. It follows from the very definition of the image probability $\pi_B \mathbb{P}$ that, if p_B denotes the likelihood on $\bar{\mathcal{V}}_B$ associated with $\pi_B \mathbb{P}$, we have:

$$\forall v \in \bar{\mathcal{V}}_B, \quad p_B(v) = p(v). \tag{16}$$

Finally, note that if \mathcal{E} is a finite event structure, the Borel σ -algebra on Ω is simply the powerset $\wp(\Omega)$.

4.2. Prerequisites on projective systems of probabilities

We introduce some background material on projective systems of probabilities. Next subsection will show how to apply this material to the case of probabilistic event structures.

Our goal is to state a simplified version of Prokhorov's extension theorem, adapted to our needs. We first recall some definitions. Let (I, \leq) be a directed poset, at most countable. For each $i \in I$, let A_i be a finite set, and for each $i, j \in I$ with $i \leq j$, let $\alpha_{i,j} : A_j \rightarrow A_i$ be a mapping such that $\alpha_{i,k} = \alpha_{i,j} \circ \alpha_{j,k}$ for all $i, j, k \in I$ with $i \leq j \leq k$, and $\alpha_{i,i} = \text{Id}_{A_i}$ for all $i \in I$. The data $(A_i)_{i \in I}$ together with the collection of mappings $\alpha_{i,j}$ is called a *projective system*. Let Y denote the product space $Y = \prod_{i \in I} A_i$. The *projective limit* X of the projective system is defined as the following subset of Y :

$$X = \{(a_i)_{i \in I} : i \leq j \Rightarrow a_i = \alpha_{i,j}(a_j)\}.$$

X is denoted by $X = \lim_{\leftarrow i \in I} A_i$. X is equipped with the topology τ , called *projective topology*, restriction to X of the product topology on Y . The *Borel σ -algebra* \mathfrak{F} on X is defined as the σ -algebra generated by the projective topology. Finally, we denote by $\alpha_i : X \rightarrow A_i$ the natural projection.

Assume moreover that each set A_i is equipped with a finite probability \mathbb{P}_i . The collection $(\mathbb{P}_i)_{i \in I}$ is said to be a *projective system of probabilities* if $\mathbb{P}_i = \alpha_{i,j} \mathbb{P}_j$ for all $i, j \in I$ with $i \leq j$.

Theorem 4.3 (Prokhorov, [17, Th. 2 p. 53]). *Within the above framework, there is a unique probability measure \mathbb{P} on (X, \mathfrak{F}) such that $\mathbb{P}_i = \alpha_i \mathbb{P}$ for all $i \in I$.*

4.3. Extension of probabilistic event structures

In this section, we show how to reduce the construction of locally finite probabilistic event structures to that of finite probabilistic event structures. This is achieved by using Prokhorov’s extension theorem recalled above.

Any event structure \mathcal{E} gives rise to a projective system in the above sense, as follows: take $I = \mathcal{B}$, the lattice of finite stopping prefixes of \mathcal{E} . \mathcal{B} is indeed directed and at most countable. Consider then the sets $(\Omega_B)_{B \in \mathcal{B}}$, together with the collection of mappings $\pi_{B,B'} : \Omega_{B'} \rightarrow \Omega_B$, defined for $B, B' \in \mathcal{B}$ with $B \subseteq B'$ as in Lemma 2.4, point 2. It is obvious that $\pi_{B,B''} = \pi_{B,B'} \circ \pi_{B',B''}$ for any $B \subseteq B' \subseteq B''$, and that $\pi_{B,B} = \text{Id}_{\Omega_B}$ for all $B \in \mathcal{B}$. Hence, $(\Omega_B)_{B \in \mathcal{B}}$ is a projective system. According to the following result, its projective limit is closely related to the space Ω . A (sketch of) proof is found in Appendix B.1. The reader is referred to [3, Ch. 2] for more details.

Lemma 4.4. *Let \mathcal{E} be a locally finite event structure, and let X be the projective limit $X = \lim_{\leftarrow B \in \mathcal{B}} \Omega_B$. The mapping $\Phi : \Omega \rightarrow X$, defined by $\Phi(\omega) = (\pi_B(\omega))_{B \in \mathcal{B}}$, is a homeomorphism. Moreover, for each $B \in \mathcal{B}$, the projection $\alpha_B : X \rightarrow \Omega_B$ and the mapping $\pi_B : \Omega \rightarrow \Omega_B$ are conjugated through Φ , i.e., $\pi_B = \alpha_B \circ \Phi$.*

In particular, the Borel σ -algebra on Ω corresponds through Φ to the Borel σ -algebra of X (i.e., Φ and Φ^{-1} send measurable sets to measurable sets).

Since $(\Omega_B)_{B \in \mathcal{B}}$ is a projective system, we say that a collection $(\mathbb{P}_B)_{B \in \mathcal{B}}$ of probabilities, with \mathbb{P}_B a finite probability on Ω_B for each $B \in \mathcal{B}$, is a projective system of probabilities if we have:

$$\forall B, B' \in \mathcal{B}, \quad B \subseteq B' \Rightarrow \mathbb{P}_B = \pi_{B,B'} \mathbb{P}_{B'}.$$

Combining Lemma 4.4 and Prokhorov’s theorem (Theorem 4.3), we obtain the following:

Theorem 4.5. *Let \mathcal{E} be a locally finite event structure. If $(\mathbb{P}_B)_{B \in \mathcal{B}}$ is a projective system of probabilities, there is a unique probabilistic event structure $(\mathcal{E}, \mathbb{P})$ such that $\mathbb{P}_B = \pi_B \mathbb{P}$ for every $B \in \mathcal{B}$. \mathbb{P} is called the **extension** of $(\mathbb{P}_B)_{B \in \mathcal{B}}$.*

The theorem can be seen as a probabilistic interpretation of the commutative diagram (2). We have indeed for $B \subseteq B'$ the following new commutative diagram of probability spaces, where the σ -algebras are understood:

$$\begin{array}{ccc} (\Omega, \mathbb{P}) & \xrightarrow{\pi_{B'}} & (\Omega_{B'}, \mathbb{P}_{B'}) \\ & \searrow \pi_B & \downarrow \pi_{B,B'} \\ & & (\Omega_B, \mathbb{P}_B) \end{array}$$

5. Distributed probabilities

It follows from Theorem 4.5 of previous section, that, if \mathcal{E} is locally finite, the construction of a probabilistic event structure $(\mathcal{E}, \mathbb{P})$ reduces to the construction of a projective system of *finite* probabilistic event structures $(B, \mathbb{P}_B)_{B \in \mathcal{B}}$. This is our next objective. We shall in fact construct special classes of probabilistic event structure that are adequate models of probabilistic distributed concurrent systems: corresponding probabilities are called *distributed*, their construction is tightly bound to branching cells.

In Section 3, we have introduced branching cells as supports for exercising choice in event structures: choice is internal to branching cells and branching cells are minimal subsets of events having this property. It is therefore natural to use branching cells in constructing probabilities on event structures, based on the following policy:

- (1) attach to each branching cell c an agent α_c , responsible for the choices made within branching cell c . Agent α_c has a dice to take random decisions according to probability distribution q_c on Ω_c ;
- (2) different agents throw their dice independently.

Recall that branching cells are dynamic, hence so are the agents. The following procedure is therefore recursively applied:

- (1) assume that finite configuration v has been given some likelihood $p(v)$;
- (2) for each branching cell c enabled by v (i.e., $c \in \delta(v)$), and for each $\omega_c \in \Omega_c$, the likelihood of $v \oplus \omega_c$ is equal to $p(v \oplus \omega_c) = p(v)q_c(\omega_c)$.

The properties of branching cells play a fundamental role in the construction of the probabilities. Indeed, consider two different branching cells c and c' continuing the same v . As a property of branching cells, we have that $c' \in \delta(v \oplus \omega_c)$. Therefore:

$$p(v \oplus \omega_c \oplus \omega_{c'}) = p(v \oplus \omega_c)q_{c'}(\omega_{c'}) = p(v)q_c(\omega_c)q_{c'}(\omega_{c'}). \quad (17)$$

Formula (17) has been established by selecting c' to act first; but selecting c to act first would have brought the same result. Hence, the consistency of decompositions using branching cells makes the above construction meaningful.

Since c and c' are concurrent, formula (17) expresses that “concurrency matches probabilistic independence”, at the granularity of branching cells, reflecting point 2 of the above policy. We shall see in Section 5.3 that it is not possible in general to have the same property at a finer granularity than branching cells. Probability distributions over Ω that are constructed in this way are called *distributed*, since they result from chaining distributed agents throwing their dice independently.

5.1. Local transitions probabilities and distributed products

Let \mathcal{E} be a locally finite event structure. Recall that \mathcal{C} denotes the set of branching cells of \mathcal{E} . We shall define a probabilistic event structure from the new notion of *locally randomized* event structure.

Definition 5.1. Let \mathcal{E} be a locally finite event structure. For every branching cell c of \mathcal{E} , we say that a finite probability q_c on Ω_c is a **local transition probability** on c . We say that \mathcal{E} is **locally randomized** if each $c \in \mathcal{C}$ is equipped with a local transition probability q_c .

We fix a locally randomized event structure $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$, and we proceed with the construction of a projective system of probabilities $(\mathbb{P}_B)_{B \in \mathcal{B}}$. We define a real-valued function $p : \mathcal{W} \rightarrow \mathbb{R}$ as follows:

$$\forall v \in \mathcal{W}, \quad p(v) = \prod_{c \in \Delta(v)} q_c(v \cap c), \quad (18)$$

where $\Delta(v)$ denotes the covering of v in \mathcal{E} . The function p is well defined since, on the one hand, the product in (18) is finite, and on the other hand $v \cap c \in \Omega_c$ for every $c \in \Delta(v)$. For each $B \in \mathcal{B}$, we define the function $\mathbb{P}_B : \Omega_B \rightarrow \mathbb{R}$ by:

$$\forall v \in \Omega_B, \quad \mathbb{P}_B(v) = p(v).$$

The construction of the so-called *distributed product* breaks down into two steps, summarized in the following results. The proofs are found in Appendix B.2.

Lemma 5.2 (and definition). *The collection $(\mathbb{P}_B)_{B \in \mathcal{B}}$ is a projective system of probabilities. The extension \mathbb{P} of the projective system $(\mathbb{P}_B)_{B \in \mathcal{B}}$ (see Theorem 4.5) is called the **distributed product** of the collection $(q_c)_{c \in \mathcal{C}}$.*

For \mathbb{P} the distributed product thus constructed, the likelihood of some finite stopped configuration is given by formula (18). According to the following result, this formula also holds for finite R -stopped configurations.

Theorem 5.3. *Let $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$ be a locally randomized event structure, and let \mathbb{P} be the associated distributed product. Then $(\mathcal{E}, \mathbb{P})$ is the unique probabilistic event structure such that the likelihood function $p : \bar{\mathcal{V}} \rightarrow \mathbb{R}$ associated with \mathbb{P} is given by (18) on finite R -stopped configurations.*

5.2. Compositional properties of distributed products

In this section, we study how distributed products behave when we restrict them to stopping prefixes and to futures of configurations. We reuse the techniques we developed to manipulate branching cells and extend them to dealing with probabilities.

Universal property of distributed product w.r.t. the past. Let $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$ be a locally randomized event structure. According to Proposition 3.8, $\mathcal{C}_B \subseteq \mathcal{C}$ holds for every stopping prefix B of \mathcal{E} . Hence the pair $(B, (q_c)_{c \in \mathcal{C}_B})$ defines a locally randomized event structure. By construction, we have the following relationship between the distributed product on B and on \mathcal{E} :

Proposition 5.4. *Let \mathbb{P} denote the distributed product of $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$, and let \mathbb{P}_B denote the distributed product of $(B, (q_c)_{c \in \mathcal{C}_B})$. \mathbb{P} and \mathbb{P}_B are related by:*

$$\mathbb{P}_B = \pi_B \mathbb{P},$$

where $\pi_B : \Omega \rightarrow \Omega_B$ is the mapping defined in Lemma 2.4.

This result is obvious. Yet, it has the following interesting consequence:

Corollary 5.5. *For each initial branching cell $c \in \delta(\emptyset)$, let $Z_c : \Omega \rightarrow \Omega_c$ be the random variable defined by $Z_c(\omega) = \omega \cap c, \omega \in \Omega$. Then the family $(Z_c)_{c \in \delta(\emptyset)}$ is a family of independent random variables, and Z_c has law q_c in Ω_c , for each $c \in \delta(\emptyset)$. Equivalently:*

$$\forall (z_c)_{c \in \delta(\emptyset)} \in \prod_{c \in \delta(\emptyset)} \Omega_c, \quad \mathbb{P}\{\omega \in \Omega : \forall c \in \delta(\emptyset), \omega \cap c = z_c\} = \prod_{c \in \delta(\emptyset)} q_c(z_c). \tag{19}$$

Proof. Let B_0 be the max-initial stopping prefix of \mathcal{E} , defined by $B_0 = \bigcup_{c \in \delta(\emptyset)} c$. Let $Z : \Omega \rightarrow \prod_{c \in \delta(\emptyset)} \Omega_c$ be the product random variable $Z = (Z_c)_{c \in \delta(\emptyset)}$. According to Proposition 3.4, Z identifies with the random variable $\omega_{B_0} = \omega \cap B_0$. Applying Proposition 5.4 to B_0 , the law of Z is given by the distributed product constructed in B_0 . It follows from Corollary 3.12 that $\Delta(\omega_{B_0}) = \delta(\emptyset)$ holds for every $\omega \in \Omega$. Therefore, formula (18) yields (19) and proves the corollary. \square

Conditional probability and probabilistic future. Recall the notion of *conditional probability*: Let $(\Omega, \mathfrak{F}, \mathbb{P})$ be a probability space, and let A be a measurable subset of Ω such that $\mathbb{P}(A) > 0$. The σ -algebra induced by \mathfrak{F} on A is the σ -algebra \mathfrak{F}^A on A which elements are those $B \subseteq A$ such that $B \in \mathfrak{F}$. We define a probability \mathbb{P}^A on (A, \mathfrak{F}^A) by putting:

$$\forall B \in \mathfrak{F}^A, \quad \mathbb{P}^A(B) = \frac{\mathbb{P}(B)}{\mathbb{P}(A)}.$$

\mathbb{P}^A is called the probability \mathbb{P} *conditionally on* A .

Assume that $(\mathcal{E}, \mathbb{P})$ is a probabilistic event structure. Let p be the likelihood associated with \mathbb{P} , and assume that u is a configuration of \mathcal{E} satisfying: $p(u) > 0$. In other words, the shadow $\mathcal{S}(u)$ has positive probability, and we define thus the conditional probability $\mathbb{P}^{\mathcal{S}(u)}$ on $\mathcal{S}(u)$. Remark that $\mathcal{S}(u)$ is isomorphic, as a measurable space, with the space Ω^u of maximal configurations of \mathcal{E}^u , equipped with its Borel σ -algebra. Therefore $\mathbb{P}^{\mathcal{S}(u)}$ is equivalently defined on Ω^u . Denote for short $\mathbb{P}^{\mathcal{S}(u)}$ by \mathbb{P}^u . Denote also by $\mathcal{S}^u(v)$ the shadow in \mathcal{E}^u of a configuration $v \in \overline{\mathcal{V}}^u$. Then we have:

$$u \oplus \mathcal{S}^u(v) = \mathcal{S}(u \oplus v). \tag{20}$$

Denote by p^u the likelihood of probability \mathbb{P}^u . It follows from (20) that p^u , defined on $\overline{\mathcal{V}}^u$, is given by:

$$\forall v \in \overline{\mathcal{V}}^u, \quad p^u(v) = \frac{p(u \oplus v)}{p(u)}. \tag{21}$$

We have obtained:

Lemma 5.6 (and definition). *If $(\mathcal{E}, \mathbb{P})$ is a probabilistic event structure, then for every configuration u such that $p(u) > 0$, the future \mathcal{E}^u inherits the structure of a probabilistic event structure $(\mathcal{E}^u, \mathbb{P}^u)$, that we call the **probabilistic future** of u , and which likelihood p^u is given by formula (21).*

Universal property of distributed products w.r.t. the future. Let $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$ be a locally randomized event structure. Fix u a finite R -stopped configuration of \mathcal{E} . According to Proposition 3.8, $\mathcal{C}^u \subseteq \mathcal{C}$, so that $(\mathcal{E}^u, (q_c)_{c \in \mathcal{C}^u})$ is a well-defined locally randomized event structure. Consider the probabilistic

event structure $(\mathcal{E}, \mathbb{P})$ constructed from the distributed product of $(q_c)_{c \in \mathcal{C}}$, and assume that $p(u) > 0$. We have two ways to construct a probabilistic event structure on \mathcal{E}^u : first, we have the probabilistic future defined in Lemma 5.6, and second, the distributed product of the family $(q_c)_{c \in \mathcal{C}^u}$. They actually coincide:

Proposition 5.7. *Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, such that \mathbb{P} is the distributed product arising from a locally randomized event structure $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$. Let $u \in \mathcal{W}$, and assume that u has positive likelihood. Then the probabilistic future $(\mathcal{E}^u, \mathbb{P}^u)$ coincides with the distributed product of $(\mathcal{E}^u, (q_c)_{c \in \mathcal{C}^u})$.*

The proof of the proposition is found in Appendix B.2. We can then sharpen Corollary 5.5 as follows:

Corollary 5.8. *Let u be a finite R -stopped configuration of a locally randomized event structure $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$, with distributed product \mathbb{P} . For each $c \in \delta(u)$, let $Z_c^u : \mathcal{S}(u) \rightarrow \Omega_c$ be the random variable defined by $Z_c^u(\omega) = \omega \cap c$ for $\omega \in \mathcal{S}(u)$. Then, conditionally on $\mathcal{S}(u)$, the collection $(Z_c^u)_{c \in \delta(u)}$ is a family of independent random variables and the law of Z_c^u is q_c . Equivalently:*

$$\forall (z_c)_{c \in \delta(u)} \in \prod_{c \in \delta(u)} \Omega_c, \quad \mathbb{P}^u \{ \omega \in \mathcal{S}(u) : \forall c \in \delta(u), \omega \cap c = z_c \} = \prod_{c \in \delta(u)} q_c(z_c). \quad (22)$$

Proof. According to Proposition 5.7, $(\mathcal{E}^u, \mathbb{P}^u)$ is the distributed product associated with $(\mathcal{E}^u, (q_c)_{c \in \mathcal{C}^u})$. Applying Corollary 5.5 brings the result. \square

5.3. Concurrency and probabilistic independence

This section analyzes two questions: First, in which extend the construction of distributed products achieves the goal that “concurrent processes are independent in probability”? The answer is that the matching concurrency/independence holds for processes bound to R -stopped configurations, and thus, implicitly, bound to branching cells. Second, would it be possible to have the same property at a finer grain than branching cells? The answer is *no* in general.

The probabilistic independence of concurrent processes can be expressed in the following form: If p denotes the likelihood of a distributed product, and if u and v are two disjoint and compatible R -stopped configurations, then we have:

$$p(u \cup v) = p(u)p(v). \quad (23)$$

Indeed, this follows from the likelihood formula (18), combined with the facts that $\Delta(u \cup v) = \Delta(u) \cup \Delta(v)$ and $\Delta(u) \cap \Delta(v) = \emptyset$. Note that (23) also holds if we only assume that u and v are not necessarily R -stopped, but are sub-configurations respectively of u' and v' , where u' and v' are finite, R -stopped and compatible.

Can we further relax the assumption about R -stopped configurations? What about disjoint and compatible configurations *inside a same branching cell*? In other words, can we have a matching between concurrency and probabilistic independence at a finer grain than branching cells? In general, the answer is “no,” except for trivial probabilities.

Here is a simple example to illustrate this claim. Let $\mathcal{E} = \{e_1, e_2, e_3\}$, with an empty causality relation, and with conflict defined by $e_1 \# e_2, e_2 \# e_3$. Consider the two compatible configurations $u = (e_1)$

and $v = (e_3)$. Assume that the formula $p(u \cup v) = p(u)p(v)$ holds. Remark that we have, for every $\omega \in \Omega$: $\omega \supseteq u \iff \omega \supseteq v$. In other words: $\mathcal{S}(u) = \mathcal{S}(v) = \mathcal{S}(u \cup v)$, and therefore:

$$p(u) = p(u \cup v) = p(u)p(v) = p(u)^2.$$

Hence $p(u) = p(u)^2$, and thus $p(u) = 0$ or $p(u) = 1$. That is, for this example, any probability \mathbb{P} on Ω such that $p(u \cup v) = p(u)p(v)$ is trivial.

The conclusion is thus the following:

- (1) distributed products allow concurrent processes to be independent in the probabilistic sense, at the grain of branching cells;
- (2) it is not possible, in general and for any probability, to have the same property at a finer grain than branching cells.

Conditional matching of concurrency and probabilistic independence. We give below the conditional formulation of Eq. (23), still for distributed products. Let u be some finite R -stopped configuration, with $p(u) > 0$. Since the probabilistic future $(\mathcal{E}^v, \mathbb{P}^v)$ is given by a distributed product (Proposition 5.7), the likelihood p^u satisfies the above property (23). Namely, if v, w are any two finite configurations of \mathcal{E}^u , that we assume disjoint and R -stopped, we have $p^u(v \cup w) = p^u(v)p^u(w)$. Equivalently, since $p^u(\cdot) = p(u \oplus \cdot)/p(u)$ according to (21), we get $p(u)p(u \oplus (v \cup w)) = p(u \oplus v)p(u \oplus w)$. Remark that the latter also holds even if $p(u) = 0$.

This can be re-expressed as follows. Let v, w be two finite and R -stopped compatible configurations of \mathcal{E} . According to point 5 of Theorem 3.11, the intersection $u = v \cap w$ is R -stopped in \mathcal{E} , and moreover $v \ominus u$ and $w \ominus u$ are R -stopped and disjoint in \mathcal{E}^u . Therefore, $p(u)p(v \cup w) = p(v)p(w)$. Hence we get, for any two R -stopped compatible configurations v, w of \mathcal{E} :

$$p(v \cap w)p(v \cup w) = p(v)p(w). \quad (24)$$

If $v \cap w = \emptyset$, (24) reduces to (23) since $p(\emptyset) = 1$.

The case of confusion-free event structures. In a confusion-free event structures every configuration is stopped, and thus R -stopped. Hence, for a distributed product defined on a confusion-free event structure, formula (24) above holds for any finite configurations u and v . This particular result for confusion-free event structures is stated in [4], in the framework of so-called “non-leaking valuations with independence”. These valuations for confusion-free event structures correspond 1–1 with distributed products.

5.4. Distributed probabilities and distributed products

In this subsection, we give a characterization of those probabilities that can be obtained as a distributed product. For this, we sharpen the condition (23) discussed above and define by this way *distributed probabilities*. We obtain then an equivalence between distributed probabilities and distributed products. As a corollary, we get that the local transition probabilities that give rise to a given distributed product are unique.

Induced local transition probabilities. Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure. To check whether \mathbb{P} is a distributed product, we first need candidates for the branching probabilities $(q_c)_{c \in \mathcal{C}}$. For this, we proceed as follows (proofs are found in Appendix B.3).

Let $c \in \mathcal{C}$, and consider the following subset of Ω :

$$\mathcal{H}^c = \{\omega \in \Omega : c \in \Delta(\omega)\}.$$

In other words, $\omega \in \mathcal{H}^c$ if there exists a finite R -stopped configuration $u \subseteq \omega$ such that $c \in \Delta(u)$. We call \mathcal{H}^c the *thick shadow* of c .

Exercise : Show that, if \mathcal{E} is confusion-free, \mathcal{H}^c coincides with the shadow $\mathcal{S}(v)$ of some finite configuration. Why is \mathcal{H}^c called a thick shadow in general?

Lemma 5.9. *For each $c \in \mathcal{C}$, \mathcal{H}^c is a measurable subset of Ω . We equip \mathcal{H}^c with the σ -algebra $\mathfrak{F}^{\mathcal{H}^c}$ induced from the Borel σ -algebra on Ω . The function Y^c defined by $Y^c(\omega) = \omega \cap c$ is a random variable $Y^c : \mathcal{H}^c \rightarrow \Omega_c$.*

Assume that c is a branching cell of \mathcal{E} such that $\mathbb{P}(\mathcal{H}^c) > 0$. Then, since \mathcal{H}^c is measurable, we equip \mathcal{H}^c with the conditional probability $\mathbb{P}^{\mathcal{H}^c}$. This makes $(\mathcal{H}^c, \mathfrak{F}^{\mathcal{H}^c}, \mathbb{P}^{\mathcal{H}^c})$ a probability space. Since Y^c is a random variable with values in Ω_c , the law of Y^c is a probability on Ω_c , i.e., a local transition probability on c . We define thus:

Definition 5.10. Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, and let c be a branching cell such that $\mathbb{P}(\mathcal{H}^c) > 0$. We define the *local transition probability on c induced by \mathbb{P}* as the probability r_c on Ω_c , image of $\mathbb{P}^{\mathcal{H}^c}$ under Y^c :

$$r_c = Y^c \mathbb{P}^{\mathcal{H}^c}, \quad \text{i.e.,} \quad \forall \omega_c \in \Omega_c, \quad r_c(\omega_c) = \mathbb{P}^{\mathcal{H}^c} \{\omega \in \mathcal{H}^c : \omega \cap c = \omega_c\}.$$

The induced local transition probability r_c is indeed a good candidate, as shown by the following result:

Lemma 5.11. *Let \mathbb{P} be the distributed product of a locally randomized event structure $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$, and let c be a branching cell of \mathcal{E} such that $\mathbb{P}(\mathcal{H}^c) > 0$. Then the induced local transition probability r_c is given by $r_c = q_c$.*

Distributed probabilities. For each branching cell c , the random variable Y^c and the induced local transition probability r_c are defined in a way intrinsic to c . There is also an alternative way of defining a random variable with values in Ω_c . Recall that we have defined in Corollary 5.8 for $v \in \mathcal{W}$ and $c \in \delta(v)$, the random variable:

$$Z_c^v : \mathcal{S}(v) \rightarrow \Omega_c, \quad Z_c^v(\omega) = \omega \cap c.$$

Fixing $v \in \mathcal{W}$, and letting c range over $\delta(v)$, we define the joint random variable Z^v as follows:

$$Z^v = (Z_c^v)_{c \in \delta(v)}, \quad Z^v : \mathcal{S}(v) \rightarrow \prod_{c \in \delta(v)} \Omega_c. \quad (25)$$

Since, according to Lemma 5.11, $q_c = r_c$ for distributed products, we equivalently reformulate Corollary 5.8 by saying that the law of Z^v is given by the product probability $\bigotimes_{c \in \delta(v)} r_c$. This suggests the following definition. (The requirement below that “ $p(v) > 0$ for every finite configuration v ” is stated for the sake of simplicity—it can be removed with some more technical effort, see [3] for details.)

Definition 5.12. Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, with likelihood p , such that $p(v) > 0$ for every finite configuration v . We say that $(\mathcal{E}, \mathbb{P})$, or shortly that the probability \mathbb{P} , is *distributed*, if for any finite and R -stopped configuration v of \mathcal{E} , the law of Z^v in $\prod_{c \in \delta(v)} \Omega_c$ is given by the product $\bigotimes_{c \in \delta(v)} r_c$ of local transition probabilities r_c induced by \mathbb{P} . Equivalently:

$$\forall (z_c)_{c \in \delta(v)} \in \prod_{c \in \delta(v)} \Omega_c, \quad \mathbb{P}^v \{ \omega \in \mathcal{S}(v) : \forall c \in \delta(v), \omega \cap c = z_c \} = \prod_{c \in \delta(v)} r_c(z_c). \tag{26}$$

In this definition, not only we require the variables Z_c^v to be independent, when c ranges over $\delta(v)$, we also require that the law of Z_c^v , for c fixed, is independent of v . Hence, we require a little bit more than the independence form (23) that we obtained in Section 5.3 when discussing the matching of concurrency and independence.

Remark that, if \mathbb{P} originates from a distributed product $(\mathcal{E}, (q_c)_{c \in C})$, and according to formula (18), the requirement that “ $p(v) > 0$ for every finite configuration v ” in Definition 5.12 is fulfilled if and only if the branching probabilities q_c satisfy:

$$\forall c \in C, \quad \forall z \in \Omega_c, \quad q_c(z) > 0.$$

A distributed product gives rise to a distributed probability (compare (26) with (22)). The following theorem addresses the converse problem:

Theorem 5.13. Let $(\mathcal{E}, \mathbb{P})$ be a probabilistic event structure, with likelihood p , such that $p(v) > 0$ for every finite configuration v . Then \mathbb{P} is a distributed product if and only if \mathbb{P} is distributed. In this case \mathbb{P} is the distributed product of the family $(r_c)_{c \in C}$ of local transition probabilities induced by \mathbb{P} . The decomposition of \mathbb{P} as a distributed product is unique.

Without the positivity assumption, the result of Theorem 5.13 remains valid, except that uniqueness is not guaranteed anymore.

Remark that, in general, not every probabilistic event structure is distributed. Consider for example two discrete random variables $X \in \{a, b\}$, $Y \in \{c, d\}$, non independent, and the event structure $\{a, b, c, d\}$ without causality relations, and with $a \# b$ and $c \# d$. The probability law of the pair (X, Y) is not given by a distributed product, since the independence condition between X and Y is not fulfilled.

6. Conclusion

In this paper, we have proposed *locally finite* event structures, a new class of event structures that support an explicit construction of *probabilistic event structures*, i.e., of models where partial orders are randomized, not interleavings.

Our construction relies on a dynamic decomposition of configurations by means of *branching cells*. Branching cells decompose maximal configurations in a recursive and dynamic way, such that a maximal configuration can be seen as a stack of choices performed inside branching cells. The *distributed probabilities* we construct are such that, at the granularity of branching cells, parallel local processes are made independent in the probabilistic sense, conditionally on their common past: informally, “concurrency matches probabilistic independence.” In general, no finer grain is possible for concurrency matching probabilistic independence.

Branching cells and distributed probabilities are the two main contributions of this work. Their use is illustrated in [18,12], where the tools developed here are applied to the particular case of event structures arising from unfoldings of safe Petri nets. This yields the model of *Markov nets*, and a first sample of probabilistic and statistical results are stated in these references. Besides their application to probabilistic event structures, we believe that branching cells are of interest per se, as they adequately capture some notion of choice.

A challenging direction for future work deals with further relaxing local finiteness in the construction of probabilistic event structures. This can be tackled by handling infinite branching cells directly—the difficulty is that the max-initial decomposition of a maximal configuration ω is not anymore guaranteed to converge to ω . Alternatively, considering products in the category of event structures is a second approach, allowing to reach event structures that are not locally finite.

Appendix A. Proof of Lemma 3.9 and of Theorem 3.11

We state first some intermediate results.

Lemma A.1. *Let v be a R -stopped configuration, and let c be an initial stopping prefix of \mathcal{E} . Then either $v \cap c = \emptyset$ or $v \cap c \in \Omega_c$.*

Proof. It follows from Proposition 2.13 that $u = v \cap B$ is R -stopped in B for every stopping prefix B , and in particular for $B = c$. Now, since c is an initial stopping prefix, it is clear that $v \cap c$ is either empty or maximal in c , as claimed. \square

Lemma A.2. *Let u be a configuration of \mathcal{E} , and let c be an initial stopping prefix of \mathcal{E} . If $u \cap c = \emptyset$, then c is an initial stopping prefix of \mathcal{E}^u .*

Proof. We first prove that $c \subseteq \mathcal{E}^u$. Let $e \in c$, and assume that $e \notin \mathcal{E}^u$. e does not belong to u since $u \cap c = \emptyset$, hence e is incompatible with u . It follows from Lemma 2.5 that there are events $e' \preceq e$ and $e'' \in u$ such that $e' \#_\mu e''$. Then $e' \in c$, and since c is $\#_\mu$ -closed, this implies that $e'' \in c$. This contradicts that $u \cap c = \emptyset$. Hence we have shown that $c \subseteq \mathcal{E}^u$.

According to Lemma 2.9, this implies that $c \cap \mathcal{E}^u = c$ is a stopping prefix of \mathcal{E}^u . Since $c \neq \emptyset$, to show that c is an initial stopping prefix of \mathcal{E}^u , it remains only to show that c is minimal among non-empty stopping prefixes of \mathcal{E}^u . For this, let γ be a nonempty stopping prefix of \mathcal{E}^u , and assume that $\gamma \subseteq c$. Then we claim that γ is a stopping prefix of \mathcal{E} . Denote by $\#_\mu^u$ the minimal conflict relation in \mathcal{E}^u . First, it is clear that γ is a prefix of \mathcal{E} , since $\gamma \subseteq c$ and since c is a prefix of \mathcal{E} . Second, we show that γ is $\#_\mu$ -closed in \mathcal{E} . Let $e \in \gamma$, and let $e' \in \mathcal{E}$ with $e \#_\mu e'$. Then e' belongs to c since c is $\#_\mu$ -closed

in \mathcal{E} , and therefore $e' \in \mathcal{E}^u$. According to Eq. (5), this implies that e and e' are in minimal conflict in \mathcal{E}^u . Since γ is chosen to be $\#_\mu^u$ -closed, this implies that $e' \in \gamma$. This shows that γ is $\#_\mu$ -closed in \mathcal{E} . Finally, γ is a stopping prefix of \mathcal{E} as claimed. Since c is initial in \mathcal{E} , we get that $c = \gamma$. We have thus shown that c is minimal in \mathcal{E}^v , which completes the proof. \square

It will be convenient to use the following terminology:

Definition A.3. We say that a configuration z of \mathcal{E} is a **germ** of \mathcal{E} if there is an initial stopping prefix c such that $z \in \Omega_c$. A valid decomposition of some R -stopped configuration v satisfying conditions 1 and 2 of Lemma 3.9 is said to be a **germ-decomposition** of v .

Lemma A.4. Every R -stopped configuration v has a germ decomposition.

Proof. We first show the result when v is a finite stopped configuration, i.e., $v \in \Omega_B$ with B a finite stopping prefix of \mathcal{E} . According to point 1 in Proposition 2.13, there is no loss of generality if we assume that $B = \mathcal{E}$, and \mathcal{E} is a finite event structure.

Consider the following inductive construction: set $v_0 = \emptyset$. Assume that the sequence $(v_j)_{0 \leq j \leq n}$ has been constructed, such that $(v_j)_{0 \leq j \leq n}$ is a germ decomposition of v_n , for $n \geq 0$, with $v_n \subseteq v$. Then:

Case (a): If $v_n = v$, stop the construction.

Case (b): Otherwise, consider $w = v \ominus v_n$. Then w is maximal in \mathcal{E}^{v_n} since v is maximal in \mathcal{E} . Pick c any initial stopping prefix of \mathcal{E}^{v_n} , and put: $z = c \cap v = c \cap w$. Then, since w is maximal in \mathcal{E}^{v_n} , Lemma 2.4 implies that $z \in \Omega_c$. Define $v_{n+1} = v_n \oplus z$. Then $(v_j)_{0 \leq j \leq n+1}$ is a germ decomposition of v_{n+1} . Repeat the procedure.

We claim that this construction eventually enters in case (a). Indeed, each time we are in case (b), the branching cell c is nonempty; therefore $z \in \Omega_c$ is nonempty, and therefore the cardinal $|v_{n+1}|$ satisfies $|v_{n+1}| \geq |v_n| + 1$. Since v is finite, and since $v_n \subseteq v$ for all $n \geq 0$, case (b) can only be reached finitely many times. When case (a) is reached, say at step n , $(v_j)_{0 \leq j \leq n}$ is a germ decomposition of v .

For the general case, let v be some R -stopped configuration with $(v_n)_{0 \leq n < N}$ a valid decomposition of v , $N \leq \infty$. We apply the above construction to each finite configuration $v_{n+1} \ominus v_n$, stopped in \mathcal{E}^{v_n} . We get a finite germ decomposition $(v_{n,j})_{0 \leq j \leq N_n}$ for each $n < N$. The concatenation of these germ decomposition yields a germ decomposition of v . \square

The following lemma is illustrated in Fig. A.1.

Lemma A.5 (First exchange lemma). Let v_0 be a finite R -stopped configuration of \mathcal{E} , let ζ be a germ of \mathcal{E}^{v_0} , and let ξ be a germ of \mathcal{E} . Assume that ξ and $v_0 \oplus \zeta$ are compatible and set:

$$v \stackrel{[\text{def}]}{=} v_0 \cup \xi, \quad v' \stackrel{[\text{def}]}{=} (v_0 \oplus \zeta) \cup \xi, \quad \zeta' \stackrel{[\text{def}]}{=} v' \setminus v. \quad (\text{A.1})$$

Then ζ' is stopped in \mathcal{E}^v .

Proof. Let c be the (unique) initial stopping prefix of \mathcal{E} such that $\xi \in \Omega_c$. We distinguish two cases.




Fig. A.1. Illustrating Lemma A.5. The first diagram shows $v_0 \oplus \zeta$, with two possible positions for ζ . The other diagrams show the possible situations for ξ , corresponding to “first case,” “second case (a),” and “second case (b)” of the proof, which yield $\zeta' = \zeta$, $\zeta' = \emptyset$, and $\zeta' = \zeta$, respectively.

First case : $v_0 \cap \xi \neq \emptyset$. Then $v_0 \cap c \neq \emptyset$. According to Lemma A.1, this implies that $v_0 \cap c \in \Omega_c$. Since $v_0 \cap c$ and ξ are two maximal compatible configurations of c , they coincide. Hence, by (A.1), $\xi \subseteq v_0$, $v = v_0$ and $\zeta' = \zeta$. So ζ' is a germ of $\mathcal{E}^v = \mathcal{E}^{v_0}$, and thus ζ' is stopped in \mathcal{E}^v , as requested.

Second case : $v_0 \cap \xi = \emptyset$. We claim that we have: $v_0 \cap c = \emptyset$. Indeed, $v_0 \cap c$ is either empty or maximal in c according to Lemma A.1. In the latter case, since ξ is also maximal in c , and compatible with v_0 , both coincide, which contradicts $v_0 \cap \xi = \emptyset$. Hence, $v_0 \cap c = \emptyset$, as claimed.

Applying Lemma A.2 with $u = v_0$, we get that c is an initial stopping prefix of \mathcal{E}^{v_0} . Hence ζ and ξ are two compatible germs of \mathcal{E}^{v_0} . Let c' be the initial stopping prefix of \mathcal{E}^{v_0} such that $\zeta \in \Omega_{c'}$. Since distinct initial stopping prefixes are disjoint by Theorem 3.2, we either have $c = c'$ or $c \cap c' = \emptyset$.

- (a) $c = c'$. Then ξ and ζ are compatible and maximal in c , so $\xi = \zeta$. Then, from (A.1), $\zeta' = \emptyset$ is trivially stopped in \mathcal{E}^v .
- (b) $c \cap c' = \emptyset$. This implies that $\xi \cap c' = \emptyset$. Hence, by Lemma A.2, c' is an initial stopping prefix of $(\mathcal{E}^{v_0})^\xi = \mathcal{E}^{v_0 \oplus \xi} = \mathcal{E}^v$. We also have $\zeta \cap \xi = \emptyset$, whence $\zeta' = \zeta$ by (A.1). Therefore, $\zeta' \in \Omega_{c'}$, and thus ζ' is a germ of \mathcal{E}^v . In particular, ζ' is stopped in \mathcal{E}^v , what was to be shown. \square

Lemma A.6 (Second exchange lemma). *Let u, u' be two finite R -stopped configurations of \mathcal{E} . Assume that u and u' are compatible. Then $(u \cup u') \ominus u'$ is R -stopped in $\mathcal{E}^{u'}$.*

Proof. Assume first that u' is a germ of \mathcal{E} . According to Lemma A.4, we can choose a germ decomposition $(u_n)_{0 \leq n \leq N}$ of u . Set $u'_0 = \emptyset$, and for each integer $1 \leq n \leq N$:

$$u'_n = u' \cup u_n, \quad z_n = u_n \ominus u_{n-1}, \quad z'_n = u'_n \ominus u'_{n-1}.$$

Then we have, for all integers $1 \leq n \leq N$:

$$z'_n = (u_{n-1} \cup u' \cup z_n) \setminus (u_{n-1} \cup u').$$

We apply Lemma A.5 with $v_0 = u_{n-1}$, $\zeta = z_n$ and $\xi = u'$ to get that z'_n is stopped in $\mathcal{E}^{u_{n-1} \cup u'} = \mathcal{E}^{u'_{n-1}}$. This defines $(u'_n)_{0 \leq n \leq N}$ as a valid decomposition of $u'_N = u \cup u'$ in \mathcal{E} , such that $u'_n \supseteq u'$ for all $n \geq 1$. Therefore, $(u'_{n+1} \ominus u')_{0 \leq n \leq N-1}$ is a valid decomposition of $(u \cup u') \ominus u'$ in $\mathcal{E}^{u'}$. This completes the proof for the case where u' is a germ of \mathcal{E} .

For the general case, let $(v_n)_{0 \leq n \leq K}$ be a germ decomposition of u , such a decomposition exists according to Lemma A.4. Then, applying the first part of the proof shows that $(u \cup v'_1) \setminus v'_1$ is R -stopped in $\mathcal{E}^{v'_1}$. Since v'_2 is a germ of $\mathcal{E}^{v'_1}$, we apply again the first part of the proof to get that $(u \cup v'_2) \setminus v'_2$ is R -stopped in $\mathcal{E}^{v'_2}$, and so on. After K steps, we obtain that $(u \cup u') \setminus u'$ is R -stopped in $\mathcal{E}^{u'}$. \square

Corollary A.7. *If $(u_n)_{n \geq 0}$ is a nondecreasing sequence of finite R -stopped configurations, then $u = \bigcup_{n \geq 0} u_n$ is R -stopped.*

Proof. It follows from Lemma A.6 that, for each $n \geq 1$, $u_n \ominus u_{n-1}$ is finite R -stopped in \mathcal{E}^{u_n} . The sequence $(u_n \ominus u_{n-1})_{n \geq 1}$ brings thus, after decomposition of each term, a valid decomposition of $\bigcup_{n \geq 0} u_n$. \square

We still need two more lemmas before we can complete the proof of Lemma 3.9.

Lemma A.8. *Let v, v' be two compatible and finite R -stopped configurations. Let $c \in \delta(v)$ and $c' \in \delta(v')$. If $c \cap c' \neq \emptyset$, then $c = c'$.*

Proof. Thanks to Lemma A.6, it is enough to show the result for $v' = \emptyset$. Assume that $c \cap c' \neq \emptyset$. Since c' is an initial stopping prefix of \mathcal{E} , and since v is R -stopped in \mathcal{E} , by Lemma A.1, $v \cap c'$ is either empty or maximal in c' . The latter case cannot occur: otherwise, since $c \in \delta(v)$, this would imply that $c \subseteq \mathcal{E}^{v \cap c'}$, and then $c \cap c' = \emptyset$. It follows therefore that $v \cap c' = \emptyset$. According to Lemma A.2, this implies that $c' \in \delta(v)$. Hence c and c' are two initial stopping prefixes of \mathcal{E}^v satisfying $c \cap c' \neq \emptyset$. Since distinct initial stopping prefixes are disjoint (Theorem 3.2), this implies that $c = c'$. \square

Lemma A.9. *Let v be a finite R -stopped configuration of \mathcal{E} . Define:*

$$\overline{\Delta}(v) = \{c \in \delta(w), w \in \mathcal{W}, w \subseteq v\}. \quad (\text{A.2})$$

Then we have, for any germ-decomposition $(v_n)_{0 \leq n \leq N}$ of v :

$$\overline{\Delta}(v) = \bigcup_{n=0}^N \delta(v_n). \quad (\text{A.3})$$

Proof. It follows from the definition (A.2) of $\overline{\Delta}$ that $\bigcup_n \delta(v_n) \subseteq \overline{\Delta}(v)$. Conversely, let $c \in \overline{\Delta}(v)$, and let $u \subseteq v$ be a finite R -stopped configuration such that $c \in \delta(u)$. On the one hand, c is an initial stopping prefix of \mathcal{E}^u . On the other hand, it follows from Lemma A.6 that $v \ominus u$ is R -stopped in \mathcal{E}^u . Hence, applying Lemma A.1 in the event structure \mathcal{E}^u implies that $(v \ominus u) \cap c$ is either empty or maximal in Ω_c . We analyze the two cases:

a) $(v \ominus u) \cap c = \emptyset$.

Applying Lemma A.2 in event structure \mathcal{E}^u shows that c is an initial stopping prefix of $(\mathcal{E}^u)^{v \ominus u} = \mathcal{E}^v$. Therefore $c \in \delta(v) = \delta(v_N)$, so that $c \in \bigcup_n \delta(v_n)$.

b) $(v \ominus u) \cap c \in \Omega_c$.

Let k be the greatest integer such that $v_k \cap c = \emptyset$; k is well defined since $v_0 \cap c = \emptyset$. And $k < N$ since $v \cap c \neq \emptyset$. Thus, v_{k+1} is defined. Let c' be the initial stopping prefix of \mathcal{E}^{v_k} such that $v_{k+1} \ominus v_k \in \Omega_{c'}$. Then $c' \cap c \neq \emptyset$ by construction. Since u and v_k are compatible, it follows from Lemma A.8 that $c = c'$. Hence $c \in \bigcup_n \delta(v_n)$. \square

Proof of Lemma 3.9. Let v be a R -stopped configuration of \mathcal{E} . The existence of a germ decomposition of v is stated by Lemma A.4. Let $(v_n)_n$ be such a germ decomposition, and let $(c_n)_n$ be the associated sequence of branching cells, so that $c_n \in \delta(v_n)$ for all n . It follows from Lemma A.8 that the branching cells $(c_n)_n$ are pairwise disjoint.

We now show the invariance of the set of branching cells $C = \{c_1, c_2, \dots\}$. We first assume that v is finite, so that in the decomposition $(v_n)_n$, n ranges over the finite set $n = 0, \dots, N$ for some integer N . Consider the set of branching cells $\overline{\Delta}(v)$ defined by (A.2) in Lemma A.9. Then $C \subseteq \overline{\Delta}(v)$. Conversely, it follows from Eq. (A.3) that a branching cell $c \in \overline{\Delta}(v)$ satisfies $c \in C$ if and only if $c \cap v \neq \emptyset$. Therefore:

$$C = \overline{\Delta}(v) \setminus \delta(v).$$

The right member of the latter expression does not depend on the germ decomposition $(v_n)_n$. This completes the proof of the invariance of C if v is finite, and we also get in that case:

$$\Delta(v) = \overline{\Delta}(v) \setminus \delta(v). \tag{A.4}$$

It remains only to show that C is invariant for any R -stopped configuration v . For this, we show the following expression for C :

$$C = \bigcup_{w \in \mathcal{W}, w \subseteq v} \Delta(w). \tag{A.5}$$

Let $(c_n)_{n \in I}$ be the sequence of branching cells associated with the germ decomposition $(v_n)_n$ of v . We have:

$$\begin{aligned} C &= \{c_n, n \in I\} \\ &= \bigcup_{n \in I} \{c_j, 1 \leq j \leq n\} \\ &= \bigcup_{n \in I} \Delta(v_n), \quad \text{by the above result.} \end{aligned}$$

This implies the “ \subseteq ” inclusion in Eq. (A.5). Conversely, let $c \in \Delta(w)$ for some $w \in \mathcal{W}$ such that $w \subseteq v$. Then there is an integer n such that $w \subseteq v_n$. According to the above result for finite configurations, applied to v_n , there is an integer $k \leq n$ such that $c = c_k$. Thus $c \in C$, which shows the “ \supseteq ” inclusion in Eq. (A.5), and completes the equality. Hence C is independent of the germ decomposition chosen. We have also shown:

$$\forall v \in \overline{\mathcal{W}}_{\mathcal{E}}, \quad \Delta(v) = \bigcup_{w \in \mathcal{W}, w \subseteq v} \Delta(w). \tag{A.6}$$

□

Proof of Theorem 3.11. 1. Given the definition of the covering, this is a simple consequence of Proposition 2.13.

2. Choose $(u_n)_{0 \leq n \leq N}$ a germ decomposition of u , $N < \infty$, and $(v_k)_{0 \leq k \leq K}$ a germ decomposition of v in \mathcal{E}^u , $K \leq \infty$. Then the concatenation:

$$u_0, u_1, \dots, u_N = u, u \oplus v_1, u \oplus v_2, \dots,$$

is a germ decomposition of $u \oplus v$. This shows that $\Delta(u \oplus v) = \Delta(u) \cup \Delta^u(v)$. Since any $c \in \Delta^u(v)$ satisfies $c \subseteq \mathcal{E}^u$, we also have $\Delta(u) \cap \Delta^u(v) = \emptyset$.

3. This is a re-writing of Lemma 3.9 using the notion of covering.

4. Eqs. (11) and (12) have been shown above in the proof of Lemma 3.9, in Eqs. (A.4) and (A.6), respectively.

5. The fact that $v \ominus u$ is R -stopped in \mathcal{E}^u , whenever u and v are two R -stopped configurations with $u \subseteq v$, is a direct consequence of Lemma A.6 and its Corollary A.7.

We now show that F_w is a lattice. Let $u, v \in F_w$, we prove that $u \cap v \in F_w$. We only have to show that $u \cap v$ is R -stopped. Thanks to Corollary A.7, we assume without loss of generality that u and v are finite. Let $(u_n)_{0 \leq n \leq N}$, $N < \infty$, be a germ decomposition of u , with $(c_n)_{0 < n \leq N}$ the associated branching cells. Put $\xi_n = v \cap u_n$ for $n = 0, \dots, N$. According to Lemma A.1, $v \cap c_1$ is either empty or maximal in c_1 . Therefore $v \cap c_1 = v \cap u_1 = \xi_1$, and thus ξ_1 is stopped, and in particular R -stopped. Moreover $v \ominus \xi_1$ is R -stopped in \mathcal{E}^{ξ_1} . For the same reasons, ξ_2 is R -stopped in \mathcal{E} , and $v \ominus \xi_2$ is R -stopped in \mathcal{E}^{ξ_2} . Continuing N times, we find that $\xi_N = u \cap v$ is R -stopped in \mathcal{E} , what was to be shown.

Finally, we show that $u \cup v$ is R -stopped in \mathcal{E} . Again, thanks to Corollary A.7, we assume without loss of generality that u and v are finite. We have just seen that $u \cap v$ is R -stopped. We know from Lemma A.6 that $(u \cup v) \ominus (u \cap v)$ is R -stopped in $\mathcal{E}^{u \cap v}$. By concatenation, we find that $(u \cap v) \oplus ((u \cup v) \ominus (u \cap v)) = u \cup v$ is R -stopped in \mathcal{E} , which completes the proof. \square

Appendix B. Proofs of main theorems for probabilistic event structures

B.1. Extension of probabilities

Proof of Lemma 4.4. First, the mapping Φ is well-defined. Indeed, let $\omega \in \Omega$, and let $\omega_B = \pi_B(\omega)$ for each $B \in \mathcal{B}$. Then, for $B \subseteq B'$, we have $\pi_B = \pi_{B,B'} \circ \pi_{B'}$, and therefore $\omega_B = \pi_{B,B'}(\omega_{B'})$. Hence $(\omega_B)_{B \in \mathcal{B}} \in X$, and $\Phi : \Omega \rightarrow X$ is well defined.

Φ is 1-1: Indeed, since \mathcal{E} is locally finite, we have the reconstruction formula:

$$\forall \omega \in \Omega, \quad \omega = \bigcup_{B \in \mathcal{B}} \omega_B.$$

Let us show that Φ is onto. For this, let $\xi = (\xi_B)_{B \in \mathcal{B}}$ be an element of X . Let $v = \bigcup_{B \in \mathcal{B}} \xi_B$. Then v is a prefix of \mathcal{E} . Assume, if possible, that v contains two events e and e' in conflict. Then there are two finite stopping prefixes B and B' such that $e \in \xi_B$ and $e' \in \xi_{B'}$. Let $B'' = B \cup B'$. Then $\xi_{B''}$ contains ξ_B and $\xi_{B'}$, and therefore $\xi_{B''}$ is a configuration containing both events in conflict e and e' , a contradiction. This shows that v is conflict-free, and thus a configuration of \mathcal{E} . Now pick any maximal configuration ω that contains v . Then $\omega \cap B \supseteq \xi_B$ for each $B \in \mathcal{B}$, and since ξ_B is maximal in B , this

implies that $\omega \cap B = \xi_B$. Hence $\Phi(\omega) = \xi$, and this shows that Φ is a bijection. By construction, the formula $\pi_B = \alpha_B \circ \Phi$ holds.

Furthermore, routine verifications show that Φ is both continuous and open when X is equipped with the projective topology and Ω is equipped with the restricted Scott topology. Hence Φ is a homeomorphism. \square

B.2. Construction and compositional properties of the distributed product

We consider a locally randomized event structure $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$. Before we proceed with the proofs of Lemma 5.2 and Theorem 5.3, we need to extend the range of definition of the function p , originally defined by (18) on $\mathcal{W}_{\mathcal{E}}$. Say that a sub-event structure $F \subseteq \mathcal{E}$ is **well-formed** if the collection \mathcal{C}_F of branching cells of F satisfies: $\mathcal{C}_F \subseteq \mathcal{C}_{\mathcal{E}}$. For any well-formed event structure F , denoting by Δ_F the covering map in F , we define the real-valued function $p_F : \mathcal{W}_F \rightarrow \mathbb{R}$ by:

$$\forall v \in \mathcal{W}_F, \quad p_F(v) = \prod_{c \in \Delta_F(v)} q_c(v \cap c). \tag{B.1}$$

The function p_F is well defined for the same reasons making p well defined. If F is finite and well-formed, we define \mathbb{P}_F on Ω_F by $\mathbb{P}_F(\omega_F) = p_F(\omega_F)$ for $\omega_F \in \Omega_F$. From Proposition 3.8, we have that every $B \in \mathcal{B}$ is well-formed. Moreover, since $\Delta_B = \Delta$ on \mathcal{W}_B , we have $p_B = p$ on \mathcal{W}_B . In particular, if B is finite, the new definition of \mathbb{P}_B coincides with the original definition.

Lemma B.1. Let B be any stopping prefix of \mathcal{E} given as a union of disjoint initial stopping prefixes: $B = c_1 \cup \dots \cup c_n, c_i \in \delta(\emptyset)$. Then \mathbb{P}_B is a probability on Ω_B .

Proof. We have the identification given in Proposition 3.4: $\Omega_B = \prod_{i=1}^n \Omega_{c_i}$. Hence we recognize in formula (B.1) for \mathbb{P}_B the product probability: $\mathbb{P}_B = q_{c_1} \otimes \dots \otimes q_{c_n}$. In particular, \mathbb{P}_B is a probability. \square

Lemma B.2. For every finite and well-formed sub-event structure $F \subseteq \mathcal{E}$, \mathbb{P}_F is a probability on Ω_F .

Proof. Since $\mathbb{P}_F : \Omega_F \rightarrow \mathbb{R}$ is a nonnegative function, we only have to show:

$$\sum_{u \in \Omega_F} \mathbb{P}_F(u) = 1. \tag{B.2}$$

For each finite and well-formed event structure F , we set

$$N_F = \max_{v \in \Omega_F} (\text{Card}(\Delta_F(v)) < \infty,$$

and we proceed by induction on N_F . Eq. (B.2) is trivial for $N_F = 0$ (i.e., if $F = \emptyset$), assume that it holds for all F finite and well-formed with $N_F \leq n$, and let F be a finite and well-formed sub-event structure of \mathcal{E} with $N_F \leq n + 1$. Set B the max-initial stopping prefix of F (Definition 3.13), and let $\pi_B : \Omega_F \rightarrow \Omega_B$ be the mapping defined in Lemma 2.4. Since π_B is onto Ω_B , Ω_F decomposes as the following (disjoint) union of sets:

$$\Omega_F = \bigcup_{v \in \Omega_B} \{v \oplus w, w \in \Omega_{F^v}\}, \tag{B.3}$$

where Ω_{F^v} denotes the set of maximal configurations of the future F^v . For $v \in \Omega_B$ and $w \in \Omega_{F^v}$, we have the decomposition from point 2 in Theorem 3.11:

$$\Delta_F(v \oplus w) = \Delta_F(v) \cup \Delta_{F^v}(w), \quad \Delta_F(v) \cap \Delta_{F^v}(w) = \emptyset. \tag{B.4}$$

We have $\Delta_F(v) = \Delta_B(v)$ from point 1 of Theorem 3.11. Moreover, the future F^v is well-formed thanks to Proposition 3.8. Hence the decomposition (B.4) brings, with formula (B.1):

$$p_F(v \oplus w) = p_B(v)p_{F^v}(w) = \mathbb{P}_B(v)\mathbb{P}_{F^v}(w). \tag{B.5}$$

The decomposition of Ω_F in (B.3) is a union of disjoint sets. This, together with (B.5), brings:

$$\sum_{u \in \Omega_F} \mathbb{P}_F(u) = \sum_{v \in \Omega_B} \sum_{w \in \Omega_{F^v}} p_F(v \oplus w) = \sum_{v \in \Omega_B} \mathbb{P}_B(v) \left(\sum_{w \in \Omega_{F^v}} \mathbb{P}_{F^v}(w) \right). \tag{B.6}$$

We claim that $N_{F^v} \leq n$ for each $v \in \Omega_B$. Indeed, without loss of generality we can assume that $N_F \geq 1$, and therefore $\text{Card}(\Delta_B(v)) \geq 1$ since $B \neq \emptyset$. We get thus from (B.4), for each $w \in \Omega_{F^v}$:

$$\text{Card}(\Delta_{F^v}(w)) \leq \text{Card}(\Delta_F(v \oplus w)) - 1 \leq N_F - 1 \leq n,$$

as we claimed. Therefore, the induction hypothesis implies:

$$\sum_{w \in \Omega_{F^v}} \mathbb{P}_F(w) = 1. \tag{B.7}$$

From Lemma B.1., we also have:

$$\sum_{v \in \Omega_B} \mathbb{P}_B(v) = 1. \tag{B.8}$$

From (B.6), (B.7) and (B.8) together, we get (B.2), what was to be shown. \square

Proof of Lemma 5.2. Let $B, B' \in \mathcal{B}$ with $B \subseteq B'$. We have to show that $\mathbb{P}_B = \pi_{B,B'}\mathbb{P}_{B'}$, or equivalently:

$$\forall v \in \Omega_B, \quad \mathbb{P}_B(v) = \sum_{u \in \Omega_{B'}, u \supseteq v} \mathbb{P}_{B'}(u). \tag{B.9}$$

Fix $v \in \Omega_B$. For each $u \in \Omega_{B'}$ with $u \supseteq v$, we set $w = u \ominus v$, and w ranges over Ω_{B^v} . And we still have the multiplicative formula (B.5):

$$\mathbb{P}_{B'}(u) = \mathbb{P}_B(v)\mathbb{P}_{B^v}(w). \tag{B.10}$$

Summing (B.10) over u brings:

$$\sum_{u \in \Omega_{B'}, u \supseteq v} \mathbb{P}_{B'}(u) = \mathbb{P}_B(v) \sum_{w \in \Omega_{B^v}} \mathbb{P}_{B^v}(w). \tag{B.11}$$

It follows from Lemma B.2. that the sum in the right member of (B.11) equals 1. This gives (B.9). \square

Proof of Theorem 5.3. If \mathbb{P} exists, then the likelihood of \mathbb{P} is determined on finite R -stopped configurations, and thus on finite stopped configurations. In other words, $\pi_B \mathbb{P}$ is determined for each $B \in \mathcal{B}$. According to Theorem 4.5, this implies the uniqueness of \mathbb{P} .

Now we show that \mathbb{P} has the required property. Let q be the likelihood of \mathbb{P} . By construction, p and q coincide on finite stopped configurations. It remains to show that q and p also coincide on finite R -stopped configurations. For this, let v be finite and R -stopped in \mathcal{E} . Since \mathcal{E} is locally finite, there is a finite stopping prefix B such that $v \subseteq B$. We have then, as in (16):

$$q(v) = \pi_B \mathbb{P}(\mathcal{S}_B(v)),$$

where $\mathcal{S}_B(v)$ denotes the shadow of v in Ω_B . Hence:

$$q(v) = \sum_{u \in \Omega_B, u \supseteq v} \pi_B \mathbb{P}(u) = \sum_{u \in \Omega_B, u \supseteq v} p(u) = \sum_{u \in \Omega_B, u \supseteq v} p_B(u).$$

Reasoning as in the proof of Lemma 5.2, we get by factorization:

$$\sum_{u \in \Omega_B, u \supseteq v} p_B(u) = p_B(v) = p(v).$$

Therefore $q(v) = p(v)$, and this completes the proof. \square

Proof of Proposition 5.7. Let \mathbb{P} be the distributed product of $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$, with likelihood p . We fix $u \in \mathcal{W}$, and we assume that $p(u) > 0$. Let \mathbb{P}^u be the probabilistic future defined according to Lemma 5.6, with p^u the associated likelihood. For any $v \in \mathcal{W}^u$, the following multiplicative formula holds true, and is shown in the same way than (B.5):

$$p(u \oplus v) = \prod_{c \in \Delta(u)} q_c(u \cap c) \cdot \prod_{c \in \Delta^u(v)} q_c(v \cap c) = p(u) \prod_{c \in \Delta^u(v)} q_c(v \cap c).$$

Therefore, from the formula (21) for the likelihood p^u , we get:

$$\forall v \in \mathcal{W}^u, \quad p^u(v) = \prod_{c \in \Delta^u(v)} q_c(v \cap c). \tag{B.12}$$

We recognize in the right member of (B.12) the formula analogous to (18), that defines the likelihood of the distributed product of $(\mathcal{E}^u, (q_c)_{c \in \mathcal{C}^u})$. It follows from the uniqueness stated in Theorem 5.3 that \mathbb{P}^u is the distributed product of $(\mathcal{E}^u, (q_c)_{c \in \mathcal{C}^u})$, which completes the proof. \square

B.3. Characterization of distributed probabilities

Before we proceed with the proof of Lemma 5.11, we need to introduce some material. Fix c a branching cell of \mathcal{E} . For each $\omega \in \mathcal{H}^c$, consider the following set of compatible R -stopped configurations:

$$F^c(\omega) = \{v \in \mathcal{W} : v \subseteq \omega, c \in \Delta(v)\}.$$

We claim that $F^c(\omega)$ is stable under finite intersections. Indeed, let $v, v' \in F^c(\omega)$. It follows from point 5 of Theorem 3.11 that $v \cap v' \in \mathcal{W}$. Clearly, $v \cap v' \subseteq \omega$, and finally $c \in \Delta(v \cap v')$ is a consequence of point 3 in Theorem 3.11. Hence $v \cap v' \in F^c(\omega)$, and this shows that $F^c(\omega)$ is stable under finite intersections, as we claimed. Since $F^c(\omega)$ consists of finite configurations, it follows that $F^c(\omega)$ has a unique minimal element. We denote it by:

$$R^c(\omega) = \min(F^c(\omega)).$$

It is a consequence of Theorem 3.11 that R^c satisfies the two following properties:

(1) For all pairs $\omega, \omega' \in \Omega$, we have:

$$\omega \in \mathcal{H}^c, \omega' \supseteq R^c(\omega) \implies \omega' \in \mathcal{H}^c, R^c(\omega') = R^c(\omega).$$

(2) For any stopping prefix B of \mathcal{E} such that $c \subseteq B$, denote by \mathcal{H}_B^c, F_B^c and R_B^c the objects \mathcal{H}^c, F^c and R^c defined in event structure B . Then $R^c = R_B^c \circ \pi_B$.

Fix B a finite stopping prefix such that $c \subseteq B$ —such a B exists since c is finite by Proposition 3.7, and since \mathcal{E} is locally finite. Then R_B^c has obviously finitely many values, say $\{v_1, \dots, v_n\}$. It follows from Point 2 above that R^c takes the same values than R_B^c . We have thus the following decomposition of \mathcal{H}^c into a disjoint union $\mathcal{H}^c = \bigcup_{i=1}^n \mathcal{H}^c \cap \{\omega \in \Omega : R^c(\omega) = v_i\}$. From Point 1 above, we get that each of these subsets is actually a shadow, as follows:

$$\forall i = 1, \dots, n, \quad \mathcal{H}^c \cap \{\omega \in \Omega : R^c(\omega) = v_i\} = \mathcal{S}(v_i).$$

Therefore, \mathcal{H}^c decomposes through a disjoint union of shadows:

$$\mathcal{H}^c = \bigcup_{i=1}^n \mathcal{S}(v_i). \tag{B.13}$$

We are now ready for the proofs of Lemma 5.9 and 5.11.

Proof of Lemma 5.9. Since \mathcal{H}^c has the form (B.13) of a finite union of shadows, it is clear that \mathcal{H}^c is measurable. It follows from point 3 of Theorem 3.11 that $Y^c(\omega) = \omega \cap c$ is maximal in c for every $\omega \in \mathcal{H}^c$. Therefore Y^c is defined as a mapping $Y^c : \mathcal{H}^c \rightarrow \Omega_c$. To show that Y^c is a random variable, fix $z \in \Omega_c$. Clearly, the set $\{\omega \in \Omega : \omega \supseteq z\}$ is measurable. Therefore, the set:

$$\{\omega \in \mathcal{H}^c : Y^c(\omega) = z\} = \mathcal{H}^c \cap \{\omega \in \Omega : \omega \supseteq z\}$$

is measurable. This shows that Y^c is a random variable. \square

Proof of Lemma 5.11. Let c be a branching cell of \mathcal{E} , and let r_c be the branching probability induced by the distributed product \mathbb{P} of $(\mathcal{E}, (q_c)_{c \in \mathcal{C}})$. We have seen that R^c has finitely many values $\{v_1, \dots, v_n\}$, leading to the decomposition (B.13) of \mathcal{H}^c . As a consequence, we get this decomposition through a disjoint union:

$$\{\omega \in \mathcal{H}^c : Y^c(\omega) = z\} = \bigcup_{i=1}^n \mathcal{S}(v_i \oplus z).$$

Therefore, with p the likelihood of \mathbb{P} :

$$\begin{aligned} \mathbb{P}(\omega \in \mathcal{H}^c, Y^c(\omega) = z) &= \sum_{i=1}^n \mathbb{P}(\mathcal{S}(v_i \oplus z)) = \sum_{i=1}^n p(v_i \oplus z) \\ &= q_c(z) \sum_{i=1}^n p(v_i) = q_c(z) \mathbb{P}(\mathcal{H}^c). \end{aligned}$$

Since $\mathbb{P}(\mathcal{H}^c) > 0$, we get:

$$r_c(z) = \frac{1}{\mathbb{P}(\mathcal{H}^c)} \mathbb{P}(\omega \in \mathcal{H}^c, Y^c = z) = q_c(z).$$

This completes the proof. \square

Proof of Theorem 5.13. We have seen that every distributed product is distributed in the sense of Definition 5.12. Conversely, let \mathbb{P} be a distributed probability. For each finite stopping prefix B of \mathcal{E} , let \mathbb{P}_B denote the image probability $\mathbb{P}_B = \pi_B \mathbb{P}$ on Ω_B , where $\pi_B : \Omega \rightarrow \Omega_B$ is the mapping defined in Lemma 2.4. We have already seen that \mathbb{P} and \mathbb{P}_B have the same likelihood on \mathcal{V}_B , we denote it by p . A simple computation shows that for each $c \in \mathcal{C}_B$, \mathbb{P} and \mathbb{P}_B induce the same local transition probability r_c , from which follows that \mathbb{P}_B is distributed. Consider $\omega_B \in \Omega_B$, and let $(v_n)_{n \geq 0}$ be the max-initial decomposition of ω_B given by Theorem 3.14. Let $n \geq 0$. Using the random variable Z_{B, v_n} defined by (25) in B , we apply the chain rule to get:

$$p(v_{n+1}) = p(v_n) \mathbb{P}_B(Z_{B, v_n}(\omega_B) = v_{n+1} \ominus v_n \mid \mathcal{S}(v_n)).$$

Since \mathbb{P}_B is distributed, the law of Z_{B, v_n} is the product $\bigotimes_{c \in \delta(v_n)} r_c$. Hence:

$$p(v_{n+1}) = p(v_n) \prod_{c \in \delta(v_n)} r_c(\omega_B \cap c) \stackrel{\text{by induction}}{=} \prod_{c \in \Delta(v_{n+1})} r_c(\omega_B \cap c). \tag{B.14}$$

Since $(v_n)_n$ is eventually constant equals to ω_B according to Theorem 3.14, letting n grow to ∞ in (B.14) brings:

$$p(\omega_B) = \prod_{c \in \Delta(\omega_B)} r_c(\omega_B \cap c). \tag{B.15}$$

Now let \mathbb{Q} be the distributed product of the locally randomized event structure $(\mathcal{E}, (r_c)_{c \in \mathcal{C}})$. Eq. (B.15) shows that \mathbb{P} and \mathbb{Q} have the same likelihoods on finite stopped configurations. The uniqueness in Theorem 4.5 implies that $\mathbb{P} = \mathbb{Q}$. This shows that \mathbb{P} is indeed a distributed product.

We finally show the uniqueness of the decomposition of \mathbb{P} as a distributed product. Indeed, if \mathbb{P} is the distributed product of a family of local transition probabilities $(s_c)_{c \in \mathcal{C}}$, then it follows from Lemma 5.11 that s_c is the local transition probability induced by \mathbb{P} in c (note the use of the positivity assumption). This completes the proof. \square

References

- [1] A. Benveniste, S. Haar, E. Fabre, Markov nets: probabilistic models for distributed and concurrent systems, *IEEE Trans. Automatic Control* 48 (11) (2003) 1936–1950.
- [2] H. Völzer, Randomized non-sequential processes, in: K. Larsen, M. Nielsen (Eds.), *Proceedings of CONCUR 01*, Lecture Notes in Computer Science, vol. 2154, Springer, Berlin, 2001, pp. 184–201.
- [3] S. Abbes, Probabilistic model for concurrent and distributed systems. limit theorems and application to statistical parametric estimation, Ph.D. thesis, IRISA/Université de Rennes 1, available from <ftp://ftp.irisa.fr/techreports/theses/2004/abbes.pdf> (2004).
- [4] D. Varacca, H. Völzer, G. Winskel, Probabilistic event structures and domains, in: P. Gardner, N. Yoshida (Eds.), *Proceeding of CONCUR 2004*, Lecture Notes in Computer Science, vol. 3170, Springer, 2004, pp. 481–496.
- [5] M. Nielsen, G. Plotkin, G. Winskel, Petri nets, event structures and domains, part 1, *Theor. Comput. Sci.* 13 (1980) 86–108.
- [6] M. Nielsen, G. Winskel, *The Book of Traces*, World Scientific, Singapore, 1995, Chapter. Trace structures and other models for concurrency.
- [7] J. Desel, R. Esparza, *Free-choice Petri nets*, Cambridge University Press, Cambridge, 1995.
- [8] C. Jones, G. Plotkin, A probabilistic powerdomain of evaluations, in: *Proceedings of LICS 1989*, IEEE Computer Society Press, Silver Spring, MD, 1989.
- [9] R. Tix, K. Keimel, G. Plotkin, Semantic domains for combining probability and non-determinism, *Electronic Notes Theor. Comput. Sci.* 129 (2004) 1–104.
- [10] M. Alvarez-Manilla, A. Edalat, N. Saheb-Djahromi, An extension result for continuous valuations, *J. Lond. Math. Soc.* 61 (2) (2000) 629–640.
- [11] G. Winskel, Event structure semantics for CCS and related languages., Tech. rep., DAIMI, pB-159 (1983).
- [12] S. Abbes, A. Benveniste, Probabilistic true-concurrency models: Markov nets and a Law of large numbers, IRISA, 2005, available from http://www.irisa.fr/distribcom/benveniste/pub/Abbes_PIfossacs05-tcs.pdf (2004).
- [13] P. Thiagarajan, Regular trace event structures, Tech. Rep. RS-96-32, BRICS, Aarhus University, Denmark, 1996.
- [14] G. Kahn, G. Plotkin, Concrete domains, *Theor. Comput. Sci.* 121 (1–2) (1993) 187–277.
- [15] L. Breiman, *Probability*, SIAM, 1992.
- [16] G. Gierz, K. Hofmann, K. Keimel, J. Lawson, M. Mislove, D. Scott, *Continuous Lattices and Domains*, Cambridge University Press, Cambridge, 2003.
- [17] N. Bourbaki, *Intégration*, Chapitre ix, *Éléments de mathématiques*, fascicule xxxv, Hermann, troisième édition, 1969.
- [18] S. Abbes, A. Benveniste, Branching cells as local states for event structures and nets: probabilistic applications, in: V. Sassone (Ed.), *Proceedings of FOSSACS 05*, Lecture Notes in Computer Science, vol. 3441, Springer, Berlin, 2005, pp. 95–109.