Probabilistic Event Structures and Domains

Daniele Varacca[1]*, Hagen Völzer[2], and Glynn Winskel[3]

[1] *LIENS - École Normale Supérieure, France*
[2] *Institut für Theoretische Informatik - Universität zu Lübeck, Germany*
[3] *Computer Laboratory - University of Cambridge, UK*

Abstract. This paper studies how to adjoin probability to event structures, leading to the model of probabilistic event structures. In their simplest form probabilistic choice is localised to cells, where conflict arises; in which case probabilistic independence coincides with causal independence. An application to the semantics of a probabilistic CCS is sketched. An event structure is associated with a domain—that of its configurations ordered by inclusion. In domain theory probabilistic processes are denoted by continuous valuations on a domain. A key result of this paper is a representation theorem showing how continuous valuations on the domain of a confusion-free event structure correspond to the probabilistic event structures it supports. We explore how to extend probability to event structures which are not confusion-free via two notions of probabilistic runs of a general event structure. Finally, we show how probabilistic correlation and probabilistic event structures with confusion can arise from event structures which are originally confusion-free by using morphisms to rename and hide events.

1 Introduction

There is a central divide in models for concurrent processes according to whether they represent parallelism by nondeterministic interleaving of actions or directly as causal independence. Where a model stands with respect to this divide affects how probability is adjoined. Most work has been concerned with probabilistic interleaving models [LS91,Seg95,DEP02]. In contrast, we propose a probabilistic causal model, a form of probabilistic event structure.

An event structure consists of a set of events with relations of causal dependency and conflict. A configuration (a state, or partial run of the event structure) consists of a subset of events which respects causal dependency and is conflict free. Ordered by inclusion, configurations form a special kind of Scott domain [NPW81].

The first model we investigate is based on the idea that all conflict is resolved probabilistically and locally. This intuition leads us to a simple model based on *confusion-free* event structures, a form of concrete data structures [KP93], but where computation proceeds by making a probabilistic choice as to which event occurs at each currently accessible cell. (The probabilistic event structures which arise are a special case of those studied by Katoen [Kat96]—though our concentration on the purely probabilistic case

* Work partially done as PhD student at BRICS - Aarhus, Denmark
Basic Research in Computer Science (`www.brics.dk`),
funded by the Danish National Research Foundation.

and the use of cells makes the definition simpler.) Such a probabilistic event structure immediately gives a "probability" weighting to each configuration got as the product of the probabilities of its constituent events. We characterise those weightings (called *configuration valuations*) which result in this way. Understanding the weighting as a true probability will lead us later to the important notion of probabilistic test.

Traditionally, in domain theory a probabilistic process is represented as a continuous valuation on the open sets of a domain, i.e., as an element of the probabilistic powerdomain of Jones and Plotkin [JP89]. We reconcile probabilistic event structures with domain theory, lifting the work of [NPW81] to the probabilistic case, by showing how they determine continuous valuations on the domain of configurations. In doing so however we do not obtain all continuous valuations. We show that this is essentially for two reasons: in valuations probability can "leak" in the sense that the total probability can be strictly less than 1; more significantly, in a valuation the probabilistic choices at different cells need not be probabilistically independent. In the process we are led to a more general definition of probabilistic event structure from which we obtain a key representation theorem: continuous valuations on the domain of configurations correspond to the more general probabilistic event structures.

How do we adjoin probabilities to event structures which are not necessarily confusion-free? We argue that in general a probabilistic event structure can be identified with a probabilistic run of the underlying event structure and that this corresponds to a probability measure over the maximal configurations. This sweeping definition is backed up by a precise correspondence in the case of confusion-free event structures. Exploring the operational content of this general definition leads us to consider probabilistic tests comprising a set of finite configurations which are both mutually exclusive and exhaustive. Tests do indeed carry a probability distribution, and as such can be regarded as finite probabilistic partial runs of the event structure.

Finally we explore how phenomena such as probabilistic correlation between choices and confusion can arise through the hiding and relabeling of events. To this end we present some preliminary results on "tight" morphisms of event structures, showing how, while preserving continuous valuations, they can produce such phenomena.

2 Probabilistic Event Structures

2.1 Event Structures

An *event structure* is a triple $\mathcal{E} = \langle E, \leq, \# \rangle$ such that

- E is a countable set of *events*;
- $\langle E, \leq \rangle$ is a partial order, called the *causal order*, such that for every $e \in E$, the set of events $\downarrow e$ is finite;
- $\#$ is an irreflexive and symmetric relation, called the *conflict relation*, satisfying the following: for every $e_1, e_2, e_3 \in E$ if $e_1 \leq e_2$ and $e_1 \# e_3$ then $e_2 \# e_3$.

Causal dependence and conflict are mutually exclusive. If two events are not causally dependent nor in conflict they are said to be *concurrent*.

A *configuration* x of an event structure $\mathcal{E}$ is a conflict-free downward closed subset of E, i.e. a subset x of E satisfying: (1) whenever $e \in x$ and $e' \leq e$ then $e' \in x$ and (2) for every $e, e' \in x$, it is not the case that $e \# e'$. Therefore, two events of a configuration are either causally dependent or concurrent, i.e., a configuration represents a run of an event structure where events are partially ordered. The set of configurations of $\mathcal{E}$, partially ordered by inclusion, is denoted as $\mathcal{L}(\mathcal{E})$. The set of finite configurations is written by $\mathcal{L}_{\text{fin}}(\mathcal{E})$. We denote the empty configuration by $\perp$.

If x is a configuration and e is an event such that $e \notin x$ and $x \cup \{e\}$ is a configuration, then we say that e is *enabled* at x. Two configurations x, x' are said to be *compatible* if $x \cup x'$ is a configuration. For every event e of an event structure $\mathcal{E}$, we define $[e] := \downarrow e$, and $[e) := [e] \setminus \{e\}$. It is easy to see that any event e is enabled at $[e)$.

We say that events e_1 and e_2 are in *immediate* conflict, and write $e_1 \#_\mu e_2$ when $e_1 \# e_2$ and both $[e_1) \cup [e_2]$ and $[e_1] \cup [e_2)$ are configurations. Note that the immediate conflict relation is symmetric. It is also easy to see that a conflict $e_1 \# e_2$ is immediate if and only if there is a configuration where both e_1 and e_2 are enabled.

2.2 Confusion-free Event Structures

The most intuitive way to add probability to an event structure is to identify "probabilistic events", such as coin flips, where probability is associated locally. A probabilistic event can be thought of as probability distribution over a *cell*, that is, a set of events (the outcomes) that are pairwise in immediate conflict and that have the same set of causal predecessors. The latter implies that all outcomes are enabled at the same configurations, which allows us to say that the probabilistic event is either enabled or not enabled at a configuration.

Definition 2.1. *A* partial cell *is a set c of events such that $e, e' \in c$ implies $e \#_\mu e'$ and $[e) = [e')$. A maximal partial cell is called a* cell.

We will now restrict our attention to event structures where each immediate conflict is resolved through some probabilistic event. That is, we assume that cells are closed under immediate conflict. This implies that cells are pairwise disjoint.

Definition 2.2. *An event structure is* confusion-free *if its cells are closed under immediate conflict.*

Proposition 2.3. *An event structure is confusion-free if and only if the reflexive closure of immediate conflict is transitive and inside cells, the latter meaning that $e \#_\mu e' \implies [e) = [e')$.*

It follows that, in a confusion-free event structure, the reflexive closure of immediate conflict is an equivalence with cells being its equivalence classes. If an event $e \in c$ is enabled at a configuration x, all the events of c are enabled as well. In which case we say that the cell c is *accessible* at x. Confusion-free event structures correspond to deterministic concrete data structures [NPW81,KP93] and to confusion-free occurrence nets [NPW81].

We find it useful to define cells without directly referring to events. To this end we introduce the notion of *covering*.

Definition 2.4. *Given two configurations $x, x' \in \mathcal{L}(\mathcal{E})$ we say that x'* covers *x if there exists $e \notin x$ such that $x' = x \cup \{e\}$. For every finite configuration x of a confusion-free event structure, a* partial covering *at x is a set of pairwise incompatible configurations that cover x. A* covering *at x is a maximal partial covering at x.*

Proposition 2.5. *In a confusion-free event structure if C is a covering at x, then $c = \{e \mid x \cup \{e\} \in C\}$ is a cell accessible at x. Conversely, if c is accessible at x, then $C := \{x \cup \{e\} \mid e \in c\}$ is a covering at x.*

2.3 Probabilistic Event Structures with Independence

Once an event structure is confusion-free, we can associate a probability distribution with each cell. Intuitively it is as if we have a die local to each cell, determining the probability with which the events at that cell occur. In this way we obtain our first definition of a probabilistic event structure, a definition in which dice at different cells are assumed probabilistically independent.

Definition 2.6. *When $f : X \to [0, +\infty]$ is a function, for every $Y \subseteq X$, we define $f[Y] := \sum_{x \in Y} f(x)$. A* cell valuation *on a confusion-free event structure $\langle E, \leq, \# \rangle$ is a function $p : E \to [0, 1]$ such that for every cell c, we have $p[c] = 1$.*

Assuming probabilistic independence of all probabilistic events, every finite configuration can be given a "probability" which is obtained as the product of probabilities of its constituent events. This gives us a function $\mathcal{L}_{\text{fin}}(\mathcal{E}) \to [0, 1]$ which we can characterise in terms of the order-theoretic structure of $\mathcal{L}_{\text{fin}}(\mathcal{E})$ by using coverings.

Proposition 2.7. *Let p be a cell valuation and let $v : \mathcal{L}_{\text{fin}}(\mathcal{E}) \to [0, 1]$ be defined by $v(x) = \Pi_{e \in x} p(e)$. Then we have*

(a) (Normality) $v(\perp) = 1$;
(b) (Conservation) if C is a covering at x, then $v[C] = v(x)$;
(c) (Independence) if x, y are compatible, then $v(x) \cdot v(y) = v(x \cup y) \cdot v(x \cap y)$.

Definition 2.8. *A* configuration valuation with independence *on a confusion-free event structure $\mathcal{E}$ is a function $v : \mathcal{L}_{\text{fin}}(\mathcal{E}) \to [0, 1]$ that satisfies normality, conservation and independence. The configuration valuation associated with a cell valuation p as in Prop. 2.7 is denoted by v_p.*

Proposition 2.9. *If v is a configuration valuation with independence and $p : E \to [0, 1]$ is a mapping such that $v([e]) = p(e) \cdot v([e))$ for all $e \in E$, then p is a cell valuation such that $v_p = v$.*

Condition *(c)* from Proposition 2.7 is essential to prove Proposition 2.9. We will show later (Theorem 5.3) the sense in which this condition amounts to probabilistic independence.

We give an example. Take the following confusion-free event structure $\mathcal{E}_1$: $E_1 = \{a, b, c, d\}$ with the discrete causal ordering and with $a \#_\mu b$ and $c \#_\mu d$.

We define a cell valuation on $\mathcal{E}_1$ by $p(a) = 1/3, p(b) = 2/3, p(c) = 1/4, p(d) = 3/4$. The corresponding configuration valuation is defined as

- $v_p(\perp) = 1$;
- $v_p(\{a\}) = 1/3, v_p(\{b\}) = 2/3, v_p(\{c\}) = 1/4, v_p(\{d\}) = 3/4$;
- $v_p(\{a, c\}) = 1/12, v_p(\{b, c\}) = 1/6, v_p(\{a, d\}) = 1/4, v_p(\{b, d\}) = 1/2$.

In the event structure above, a covering at $\perp$ consists of $\{a\}, \{b\}$, while a covering at $\{a\}$ consists of $\{a, c\}, \{a, d\}$.

We conclude this section with a definition of a probabilistic event structure. Though, as the definition indicates, we will consider a more general definition later, one in which there can be probabilistic correlations between the choices at different cells.

Definition 2.10. *A* probabilistic event structure with independence *consists of a confusion-free event structure together with a configuration valuation with independence.*

3 A Process Language

Confusion-freeness is a strong requirement. But it is still possible to give a semantics to a fairly rich language for probabilistic processes in terms of probabilistic event structures with independence. The language we sketch is a probabilistic version of value passing CCS. Following an idea of Milner, used in the context of confluent processes [Mil89], we restrict parallel composition so that there is no ambiguity as to which two processes can communicate at a channel; parallel composition will then preserve confusion-freeness.

Assume a set of channels L. For simplicity we assume that a common set of values V may be communicated over any channel $a \in L$. The syntax of processes is given by:

$$P ::= 0 \mid \sum_{v \in V} a!(p_v, v).P_v \mid a?(x).P \mid P_1 \| P_2 \mid P \setminus A \mid$$

$$P[f] \mid \textbf{if } b \textbf{ then } P_1 \textbf{ else } P_2 \mid X \mid \textbf{rec}\, X.P$$

Here x ranges over value variables, X over process variables, A over subsets of channels and f over injective renaming functions on channels, b over boolean expressions (which make use of values and value variables). The coefficients p_v are real numbers such that $\sum_{v \in V} p_v = 1$.

A closed process will denote a probabilistic event structure with independence, but with an additional labelling function from events to output labels $a!v$, input labels $a?v$ where a is a channel and v a value, or τ. At the cost of some informality we explain the probabilistic semantics in terms of CCS constructions on the underlying labelled event structures, in which we treat pairs of labels consisting of an output label $a!v$ and input label $a?v$ as complementary. (See e.g. the handbook chapter [WN95] or [Win82,Win87] for an explanation of the event structure semantics of CCS.) For simplicity we restrict attention to the semantics of closed process terms.

The nil process 0 denotes the empty probabilistic event structure. A closed output process $\sum_{v \in V} a!(p_v, v).P_v$ can perform a synchronisation at channel a, outputting a value v with probability p_v, whereupon it resumes as the process P_v. Each P_v, for $v \in V$, will denote a labelled probabilistic event structure with underlying labelled event structure $\mathcal{E}[\![P_v]\!]$. The underlying event structure of such a closed output process is got by the juxtaposition of the family of prefixed event structures

$$a!v.\mathcal{E}[\![P_v]\!] \ ,$$

with $v \in V$, in which the additional prefixing events labelled $a!v$ are put in (immediate) conflict; the new prefixing events labelled $a!v$ are then assigned probabilities p_v to obtain the labelled probabilistic event structure.

A closed input process $a?(x).P$ synchronises at channel a, inputting a value v and resuming as the closed process $P[v/x]$. Such a process $P[v/x]$ denotes a labelled probabilistic event structure with underlying labelled event structure $\mathcal{E}[\![P[v/x]]\!]$. The underlying labelled event structure of the input process is got as the parallel juxtaposition of the family of prefixed event structures

$$a?v.\mathcal{E}[\![P[v/x]]\!] \ ,$$

with $v \in V$; the new prefixing events labelled $a?v$ are then assigned probabilities 1.

The probabilistic parallel composition corresponds to the usual CCS parallel composition followed by restricting away on all channels used for communication. In order for the parallel composition $P_1 \| P_2$ to be well formed the set of input channels of P_1 and P_2 must be disjoint, as must be their output channels. So, for instance, it is not possible to form the parallel composition

$$\sum_{v \in V} a!(p_v, v).0 \| a?(x).P_1 \| a?(x).P_2 \ .$$

In this way we ensure that no confusion is introduced through synchronisation.

We first describe the effect of the parallel composition on the underlying event structures of the two components, assumed to be E_1 and E_2. This is got by CCS parallel composition followed by restricting away events in a set S:

$$(E_1 \mid E_2) \setminus S$$

where S consists of all labels $a!v$, $a?v$ for which $a!v$ appears in E_1 and $a?v$ in E_2, or vice versa. In this way any communication between E_1 and E_2 is forced when possible. The newly introduced τ-events, corresponding to a synchronisation between an $a!v$-event with probability p_v and an $a?v$-event with probability 1, are assigned probability p_v.

A restriction $P \setminus A$ has the effect of the CCS restriction

$$\mathcal{E}[\![P]\!] \setminus \{a!v, a?v \mid v \in V \ \& \ a \in A\}$$

on the underlying event structure; the probabilities of the events which remain stay the same. A renaming $P[f]$ has the usual effect on the underlying event structure, probabilities of events being maintained. A closed conditional (**if** b **then** P_1 **else** P_2) has the denotation of P_1 when b is true and of P_2 when b is false.

The recursive definition of probabilistic event structures follows that of event structures [Win87] carrying the extra probabilities along. Though care must be taken to ensure that a confusion-free event structure results: one way to ensure this is to insist that for **rec** $X.P$ to be well-formed the process variable X may not occur under a parallel composition.

4 Probabilistic Event Structures and Domains

The configurations $\langle \mathcal{L}(\mathcal{E}), \subseteq \rangle$ of a confusion-free event structure $\mathcal{E}$, ordered by inclusion, form a domain, specifically a *distributive concrete domain* (cf. [NPW81,KP93]). In traditional domain theory, a probabilistic process is denoted by a *continuous valuation*. Here we show that, as one would hope, every probabilistic event structure with independence corresponds to a unique continuous valuation. However not all continuous valuations arise in this way. Exploring why leads us to a more liberal notion of a configuration valuation, in which there may be probabilistic correlation between cells. This provides a representation of the normalised continuous valuations on distributive concrete domains in terms of probabilistic event structures. (The Appendix includes a brief survey of the domain theory we require. The rather involved proofs of this section can be found in [Var03].)

4.1 Domains

The probabilistic powerdomain of Jones and Plotkin [JP89] consists of continuous valuations, to be thought of as denotations of probabilistic processes. A *continuous valuation* on a DCPO D is a function ν defined on the Scott open subsets of D, taking values on $[0, +\infty]$, and satisfying:

- (Strictness) $\nu(\emptyset) = 0$;
- (Monotonicity) $U \subseteq V \Longrightarrow \nu(U) \leq \nu(V)$;
- (Modularity) $\nu(U) + \nu(V) = \nu(U \cup V) + \nu(U \cap V)$;
- (Continuity) if $\mathcal{J}$ is a directed family of open sets, $\nu\left(\bigcup \mathcal{J}\right) = \sup_{U \in \mathcal{J}} \nu(U)$.

A continuous valuation ν is *normalised* if $\nu(D) = 1$. Let $\mathcal{V}^1(D)$ denote the set of normalised continuous valuations on D equipped with the pointwise order: $\nu \leq \xi$ if for all open sets U, $\nu(U) \leq \xi(U)$. $\mathcal{V}^1(D)$ is a DCPO [JP89,Eda95].

The open sets in the Scott topology represent observations. If D is an algebraic domain and $x \in D$ is compact, the *principal* set $\uparrow x$ is open. Principal open sets can be thought of as basic observations. Indeed they form a basis of the Scott topology.

Intuitively a normalised continuous valuation ν assigns probabilities to observations. In particular we could think of the probability of a principal open set $\uparrow x$ as representing the probability of x.

4.2 Continuous and Configuration Valuations

As can be hoped, a configuration valuation with independence on a confusion-free event structure $\mathcal{E}$ corresponds to a normalised continuous valuation on the domain $\langle \mathcal{L}(\mathcal{E}), \subseteq \rangle$, in the following sense.

Proposition 4.1. *For every configuration valuation with independence v on $\mathcal{E}$ there is a unique normalised continuous valuation ν on $\mathcal{L}(\mathcal{E})$ such that for every finite configuration x, $\nu(\uparrow x) = v(x)$.*

Proof. The claim is a special case of the subsequent Theorem 4.4. □

While a configuration valuation with independence gives rise to a continuous valuation, not every continuous valuation arises in this way. As an example, consider the event structure $\mathcal{E}_1$ as defined in Section 2.3. Define

- $\nu(\uparrow\{a\}) = \nu(\uparrow\{b\}) = \nu(\uparrow\{c\}) = \nu(\uparrow\{d\}) = 1/2$;
- $\nu(\uparrow\{a,d\}) = \nu(\uparrow\{b,c\}) = 1/2$;
- $\nu(\uparrow\{a,c\}) = \nu(\uparrow\{b,d\}) = 0$;

and extend it to all open sets by modularity. It is easy to verify that it is indeed a continuous valuation on $\mathcal{L}(\mathcal{E}_1)$. Define a function $v : \mathcal{L}_{\mathrm{fin}}(\mathcal{E}_1) \to [0,1]$ by $v(x) := \nu(\uparrow x)$. This is not a configuration valuation with independence; it does not satisfy condition *(c)* of Proposition 2.7. If we consider the compatible configurations $x := \{a\}, y := \{c\}$ then $v(x \cup y) \cdot v(x \cap y) = 0 < 1/4 = v(x) \cdot v(y)$.

Also continuous valuations "leaking" probability do not arise from probabilistic event structures with independence.

Definition 4.2. *Denote the set of maximal elements of a DCPO D by $\Omega(D)$. A normalised continuous valuation ν on D is* non-leaking *if for every open set $O \supseteq \Omega(D)$, we have $\nu(O) = 1$.*

This definition is new, although inspired by a similar concept in [Eda95]. For the simplest example of a leaking continuous valuation, consider the event structure $\mathcal{E}_2$ consisting of one event e only, and the valuation defined as $\nu(\emptyset) = 0$, $\nu(\uparrow \bot) = 1$, $\nu(\uparrow\{e\}) = 1/2$. The corresponding function $v : \mathcal{L}_{\mathrm{fin}}(\mathcal{E}_2) \to [0,1]$ violates condition *(b)* of Proposition 2.7. The probabilities in the cell of e do not sum up to 1.

We analyse how valuations without independence and leaking valuations can arise in the next two sections.

4.3 Valuations Without Independence

Definition 2.10 of probabilistic event structures assumes the probabilistic independence of choice at different cells. This is reflected by condition *(c)* in Proposition 2.7 on which it depends. In the first example above, the probabilistic choices in the two cells are not independent: once we know the outcome of one of them, we also know the outcome of the other. This observation leads us to a more general definition of a configuration valuation and probabilistic event structure.

Definition 4.3. *A* configuration valuation *on a confusion-free event structure $\mathcal{E}$ is a function $v : \mathcal{L}_{\mathrm{fin}}(\mathcal{E}) \to [0,1]$ such that:*

(a) $v(\bot) = 1$*;*
(b) if C is a covering at x, then $v[C] = v(x)$.

A probabilistic event structure *consists of a confusion-free event structure together with a configuration valuation.*

Now we can generalise Proposition 4.1, and provide a converse:

Theorem 4.4. *For every configuration valuation v on $\mathcal{E}$ there is a unique normalised continuous valuation ν on $\mathcal{L}(\mathcal{E})$ such that for every finite configuration x, $\nu(\uparrow x) = v(x)$. Moreover ν is non-leaking.*

Theorem 4.5. *Let ν be a non-leaking continuous valuation on $\mathcal{L}(\mathcal{E})$. The function $v : \mathcal{L}_{\text{fin}}(\mathcal{E}) \to [0, 1]$ defined by $v(x) = \nu(\uparrow x)$ is a configuration valuation.*

The two theorems above provide a representation of non-leaking continuous valuations on distributive concrete domains—see [Var03], Thm. 6.4.1 and Thm. 7.6.2 for their proof. Using this representation result, we are also able to characterise the maximal elements in $\mathcal{V}^1(\mathcal{L}(\mathcal{E}))$ as precisely the non-leaking valuations—a fact which is not known for general domains.

Theorem 4.6. *Let $\mathcal{E}$ be a confusion-free event structure and let $\nu \in \mathcal{V}^1(\mathcal{L}(\mathcal{E}))$. Then ν is non-leaking if and only if it is maximal.*

Proof. See [Var03], Prop. 7.6.3 and Thm. 7.6.4. □

4.4 Leaking Valuations

There remain leaking continuous valuations, as yet unrepresented by any probabilistic event structures. At first sight it might seem that to account for leaking valuations it would be enough to relax condition *(b)* of Definition 4.3 to the following

(b') if C is a covering at x, then $v[C] \leq v(x)$.

However, it turns out that this is not the right generalisation, as the following example shows. Consider the event structure $\mathcal{E}_3$ where $E_3 = \{a, b\}$ with the flat causal ordering and no conflict. Define a "leaking configuration valuation" on $\mathcal{E}_3$ by $v(\perp) = v(\{a\}) = v(\{b\}) = 1$, $v(\{a, b\}) = 0$.

The function v satisfies conditions *(a)* and *(b')*, but it cannot be extended to a continuous valuation on the domain of configurations. However, we can show that the leaking of probability is attributable to an "invisible" event.

Definition 4.7. *Consider a confusion-free event structure $\mathcal{E} = \langle E, \leq, \# \rangle$. For every cell c we consider a new "invisible" event ∂_c such that $\partial_c \notin E$ and if $c \neq c'$ then $\partial_c \neq \partial_{c'}$. Let $\partial = \{\partial_c \mid c \text{ is a cell}\}$. We define $\mathcal{E}_\partial$ to be $\langle E_\partial, \leq_\partial, \#_\partial \rangle$, where*

- *$E_\partial = E \cup \partial$;*
- *$\leq_\partial$ is $\leq$ extended by $e \leq_\partial \partial_c$ if for all $e' \in c$, $e \leq e'$;*
- *$\#_\partial$ is $\#$ extended by $e \#_\partial \partial_c$ if there exists $e' \in c$, $e' \leq e$.*

So $\mathcal{E}_\partial$ is $\mathcal{E}$ extended by an extra invisible event at every cell. Invisible events can absorb all leaking probability, as shown by Theorem 4.9 below.

Definition 4.8. *Let $\mathcal{E}$ be a confusion-free event structure. A* generalised configuration valuation *on $\mathcal{E}$ is a function $v : \mathcal{L}_{\text{fin}}(\mathcal{E}) \to [0, 1]$ that can be extended to a configuration valuation on $\mathcal{E}_\partial$.*

It is not difficult to prove that, when such an extension exists, it is unique.

Theorem 4.9. *Let $\mathcal{E}$ be a confusion-free event structure. Let $v : \mathcal{L}_{\mathrm{fin}}(\mathcal{E}) \to [0, 1]$. There exists a unique normalised continuous valuation ν on $\mathcal{L}(\mathcal{E})$ with $v(x) = \nu(\uparrow x)$, if and only if v is a generalised configuration valuation.*

Proof. See [Var03], Thm. 6.5.3. □

The above theorem completely characterises the normalised continuous valuations on distributive concrete domains in terms of probabilistic event structures.

5 Probabilistic Event Structures as Probabilistic Runs

In the rest of the paper we investigate how to adjoin probabilities to event structures which are not confusion-free. In order to do so, we find it useful to introduce two notions of probabilistic run.

Configurations represent runs (or computation paths) of an event structure. What is a probabilistic run (or probabilistic computation path) of an event structure? One would expect a probabilistic run to be a form of probabilistic configuration, so a probability distribution over a suitably chosen subset of configurations. As a guideline we consider the traditional model of probabilistic automata [Seg95], where probabilistic runs are represented in essentially two ways: as a probability measure over the set of maximal runs [Seg95], and as a probability distribution over finite runs of the same length [dAHJ01].

The first approach is readily available to us, and where we begin. As we will see, according to this view probabilistic event structures over an underlying event structure $\mathcal{E}$ correspond precisely to the probabilistic runs of $\mathcal{E}$.

The proofs of the results in this section are omitted, but they can be found in the technical report [VVW04].

5.1 Probabilistic Runs of an Event Structure

The first approach suggests that a probabilistic run of an event structure $\mathcal{E}$ be taken to be a probability measure on the maximal configurations of $\mathcal{L}(\mathcal{E})$.

To do so requires some notions from measure theory. A measurable space is a pair $\langle \Omega, \mathcal{S} \rangle$, where Ω is a set and $\mathcal{S}$ is a σ-algebra over Ω. A measure over a measurable space $\langle \Omega, \mathcal{S} \rangle$ is a countably additive function $\mu : \mathcal{S} \to [0, +\infty]$. If $\mu(\Omega) = 1$, we talk of a probability measure. Let D be an algebraic domain. Recall that $\Omega(D)$ denotes the set of maximal elements of D and that for every compact element $x \in D$ the *principal* set $\uparrow x$ is Scott open. The set $K(x) := \uparrow x \cap \Omega(D)$ is called the *shadow* of x. We shall consider the σ-algebra $\mathcal{S}$ on $\Omega(D)$ generated by the shadows of the compact elements. The configurations of an event structure form a coherent ω-algebraic domain, whose compact elements are the finite configurations [NPW81].

Definition 5.1. *A* probabilistic run *of an event structure $\mathcal{E}$ is a probability measure on $\langle \Omega(\mathcal{L}(\mathcal{E})), \mathcal{S} \rangle$, where $\mathcal{S}$ is the σ-algebra generated by the shadows of the compact elements.*

There is a tight correspondence between non-leaking valuations and probabilistic runs.

Theorem 5.2. *Let ν be a non-leaking normalised continuous valuation on a coherent ω-algebraic domain D. Then there is a unique probability measure μ on $\mathcal{S}$ such that for every compact element x, $\mu(K(x)) = \nu(\uparrow x)$.*
Let μ be a probability measure on $\mathcal{S}$. Then the function ν defined on open sets by $\nu(O) = \mu(O \cap \Omega(D))$ is a non-leaking normalised continuous valuation.

According to the result above, probabilistic event structures over a common event structure $\mathcal{E}$ correspond precisely to the probabilistic runs of $\mathcal{E}$. Among these we can characterise probabilistic event structures *with independence* in terms of the standard measure-theoretic notion of independence. In fact, for such a probabilistic event structure, every two compatible configurations are probabilistically independent, given the common past:

Proposition 5.3. *Let v be a configuration valuation on a confusion-free event structure $\mathcal{E}$. Let μ_v be the corresponding measure as of Propositions 4.1 and Theorem 5.2. Then, v is a configuration valuation* with independence *iff for every two finite compatible configurations x, y*

$$\mu_v\Big(K(x) \cap K(y) \mid K(x \cap y)\Big) = \mu_v\Big(K(x) \mid K(x \cap y)\Big) \cdot \mu_v\Big(K(y) \mid K(x \cap y)\Big).$$

Note that the definition of probabilistic run of an event structure does not require that the event structure is confusion-free. It thus suggests a general definition of a probabilistic event structure as an event structure with a probability measure μ on its maximal configurations, even when the event structure is not confusion-free. This definition, in itself, is however not very informative and we look to an explanation in terms of finite probabilistic runs.

5.2 Finite Runs

What is a finite probabilistic run? Following the analogy heading this section, we want it to be a probability distribution over finite configurations. But which sets are suitable to be the support of such distribution? In interleaving models, the sets of runs of the same length do the job. For event structures this won't do.

To see why consider the event structure with only two concurrent events a, b. The only maximal run assigns probability 1 to the maximal configuration $\{a, b\}$. This corresponds to a configuration valuation which assigns 1 to both $\{a\}$ and $\{b\}$. Now these are two configurations of the same size, but their common "probability" is equal to 2! The reason is that the two configurations are compatible: they do not represent *alternative* choices. We therefore need to represent alternative choices, and we need to represent them all. This leads us to the following definition.

Definition 5.4. *Let $\mathcal{E}$ be an event structure. A* partial test *of $\mathcal{E}$ is a set C of pairwise incompatible configurations of $\mathcal{E}$. A* test *is a maximal partial test. A test is* finitary *if all its elements are finite.*

Maximality of a partial test C can be characterised equivalently as *completeness*: for every maximal configuration z, there exists $x \in C$ such that $x \subseteq z$. The set of tests, endowed with the Egli-Milner order has an interesting structure: the set of all tests is a complete lattice, while finitary tests form a lattice.

Tests were designed to support probability distributions. So given a sensible valuation on finite configurations we expect it to restrict to probability distributions on tests.

Definition 5.5. *Let v be a function $\mathcal{L}_{\mathrm{fin}}(\mathcal{E}) \to [0, 1]$. Then v is called a* test valuation *if for all finitary tests C we have $v[C] = 1$.*

Theorem 5.6. *Let μ be a probabilistic run of $\mathcal{E}$. Define $v : \mathcal{L}_{\mathrm{fin}}(\mathcal{E}) \to [0, 1]$ by $v(x) = \mu(K(x))$. Then v is a test valuation.*

Note that Theorem 5.6 is for general event structures. We unfortunately do not have a converse in general. However, there is a converse when the event structure is confusion-free:

Theorem 5.7. *Let $\mathcal{E}$ be a confusion-free event structure. Let v be a function $\mathcal{L}_{\mathrm{fin}}(\mathcal{E}) \to [0, 1]$. Then v is a configuration valuation if and only if it is a test valuation.*

The proof of this theorem hinges on a property of tests. The property is that of whether partial tests can be completed. Clearly every partial test can be completed to a test (by Zorn's lemma), but there exist finitary partial tests that cannot be completed to *finitary* tests.

Definition 5.8. *A finitary partial test is* honest *if it is part of a finitary test. A finite configuration is honest if it is honest as partial test.*

Proposition 5.9. *If $\mathcal{E}$ is a confusion-free event structure and if x is a finite configuration of $\mathcal{E}$, then x is honest in $\mathcal{L}(\mathcal{E})$.*

So confusion-free event structures behave well with respect to honesty. For general event structures, the following is the best we can do at present:

Theorem 5.10. *Let v be a test valuation on $\mathcal{E}$. Let $\mathcal{H}$ be the σ-algebra on $\Omega(\mathcal{L}(\mathcal{E}))$ generated by the shadows of honest finite configurations. Then there exists a unique measure μ on $\mathcal{H}$ such that $\mu(K(x)) = v(x)$ for every honest finite configuration x.*

Theorem 5.11. *If all finite configurations are honest, then for every test valuation v there exists a unique continuous valuation ν, such that $\nu(\uparrow x) = v(x)$.*

But, we do not know whether in all event structures, every finite configuration is honest. We conjecture this to be the case. If so this would entail the general converse to Theorem 5.6 and so characterise probabilistic event structures, allowing confusion, in terms of finitary tests.

6 Morphisms

It is relatively straightforward to understand event structures with independence. But how can general test valuations on a confusion-free event structures arise? More generally how do we get runs of arbitrary event structures? We explore one answer in this section. We show how to obtain test valuations as "projections" along a morphism from a configuration valuation with independence on a confusion-free event structure. The use of morphisms shows us how general valuations are obtained through the hiding and renaming of events.

Definition 6.1 ([Win82,WN95]). *Given two event structures $\mathcal{E}, \mathcal{E}'$, a* morphism $f : \mathcal{E} \to \mathcal{E}'$ *is a partial function $f : E \to E'$ such that*

- *whenever $x \in \mathcal{L}(\mathcal{E})$ then $f(x) \in \mathcal{L}(\mathcal{E}')$;*
- *for every $x \in \mathcal{L}(\mathcal{E})$, for all $e_1, e_2 \in x$ if $f(e_1), f(e_2)$ are both defined and $f(e_1) = f(e_2)$, then $e_1 = e_2$.*

A morphism $f : \mathcal{E} \to \mathcal{E}'$ expresses how the occurrence of an event in $\mathcal{E}$ induces a synchronised occurrence of an event in $\mathcal{E}'$. Some events in $\mathcal{E}$ are hidden (if f is not defined on them) and conflicting events in $\mathcal{E}$ may synchronise with the same event in $\mathcal{E}'$ (if they are identified by f).

The second condition in the definition guarantees that morphisms of event structures "reflect" reflexive conflict ($\# \cup Id_E$). We now introduce morphisms that reflect tests; such morphisms enable us to define a test valuation on $\mathcal{E}'$ from a test valuation on $\mathcal{E}$. To do so we need some preliminary definitions. Given a morphism $f : \mathcal{E} \to \mathcal{E}'$, we say that an event of $\mathcal{E}$ is f-invisible, if it is not in the domain of f. Given a configuration x of $\mathcal{E}$ we define x_f to be x minus all its maximal f-invisible events. Clearly x_f is still a configuration and $f(x) = f(x_f)$. If $x = x_f$, we say that x is f-minimal.

Definition 6.2. *A morphism of event structures $f : \mathcal{E} \to \mathcal{E}'$ is* tight *when*

- *if $y = f(x)$ and if $y' \supseteq y$, there exists $x' \supseteq x_f$ such that $y' = f(x')$;*
- *if $y = f(x)$ and if $y' \subseteq y$, there exists $x' \subseteq x_f$ such that $y' = f(x')$;*
- *all maximal configurations are f-minimal (no maximal event is f-invisible).*

Proposition 6.3. *A tight morphism of event structures is surjective on configurations. Given $f : \mathcal{E} \to \mathcal{E}'$ tight, if C' is a finitary test of $\mathcal{E}'$ then the set of f-minimal inverse images of C' along f is a finitary test in $\mathcal{E}$.*

We now study the relation between valuations and morphisms. Given a function $v : \mathcal{L}_{\text{fin}}(\mathcal{E}) \to [0, +\infty]$ and a morphism $f : \mathcal{E} \to \mathcal{E}'$ we define a function $f(v) : \mathcal{L}_{\text{fin}}(\mathcal{E}') \to [0, +\infty]$ by $f(v)(y) = \sum\{v(x) \mid f(x) = y$ and x is f-minimal$\}$.

Proposition 6.4. *Let $\mathcal{E}, \mathcal{E}'$ be event structures, v be a test valuation on $\mathcal{E}$, and $f : \mathcal{E} \to \mathcal{E}'$ a tight morphism. Then the function $f(v)$ is a test valuation on $\mathcal{E}'$.*

Therefore we can obtain a run of a general event structure by projecting a run of a probabilistic event structure with independence. Presently we don't know whether every run can be generated in this way.

7 Related and Future Work

In his PhD thesis, Katoen [Kat96] defines a notion of probabilistic event structure which includes our probabilistic event structures with independence. But his concerns are more directly tuned to a specific process algebra. So in one sense his work is more general—his event structures also possess nondeterminism—while in another it is much more specific in that it does not look beyond local probability distributions at cells. Völzer [Voe01] introduces similar concepts based on Petri nets and a special case of Theorem 5.10. Benveniste et al. have an alternative definition of probabilistic Petri nets in [BFH03], and there is clearly an overlap of concerns though some significant differences which require study.

We have explored how to add probability to the independence model of event structures. In the confusion-free case, this can be done in several equivalent ways: as valuations on configurations; as continuous valuations on the domain of configurations; as probabilistic runs (probability measures over maximal configurations); and in the simplest case, with independence, as probability distributions existing locally and independently at cells. Work remains to be done on a more operational understanding, in particular on how to understand probability adjoined to event structures which are not confusion-free. This involves relating probabilistic event structures to interleaving models like Probabilistic Automata [Seg95] and Labelled Markov Processes [DEP02].

Acknowledgments

The first author wants to thank Mogens Nielsen, Philippe Darondeau, Samy Abbes and an anonymous referee.

References

[AJ94] S. Abramsky and A. Jung. Domain theory. In *Handbook of Logic in Computer Science*, volume 3. Clarendon Press, 1994.

[Alv00] M. Alvarez-Manilla. *Measure Theoretic Results for Continuous Valuations on Partially Ordered Spaces*. PhD thesis, University of London - Imperial College of Science, Technology and Medicine, 2000.

[AES00] M. Alvarez-Manilla, A. Edalat, and N. Saheb-Djaromi. An extension result for continuous valuations. *Journal of the London Mathematical Society*, 61(2):629–640, 2000.

[BFH03] A. Benveniste, E. Fabre, and S. Haar. Markov nets: Probabilistic models for distributed and concurrent systems. *IEEE Transactions on Automatic Control*, 48(11):1936–1950, 2003.

[dAHJ01] L. de Alfaro, T. A. Henzinger, and R. Jhala. Compositional methods for probabilistic systems. In *Proc. 12th CONCUR*, volume 2154 of *LNCS*, pages 351–365, 2001.

[DEP02] J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labelled markov processes. *Information and Computation*, 179(2):163–193, 2002.

[Eda95] A. Edalat. Domain theory and integration. *Theoretical Computer Science*, 151(1):163–193, 13 November 1995.

[JP89] C. Jones and G. D. Plotkin. A probabilistic powerdomain of evaluations. In *Proceedings of 4th LICS*, pages 186–195, 1989.

[Kat96] J.-P. Katoen. *Quantitative and Qualitative Extensions of Event Structures*. PhD thesis, University of Twente, 1996.
[KP93] G. Kahn and G. D. Plotkin. Concrete domains. *Theoretical Computer Science*, 121(1-2):187–277, 1993.
[LS91] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
[Mil89] R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
[NPW81] M. Nielsen, G. D. Plotkin, and G. Winskel. Petri nets, event structures and domains, part I. *Theoretical Computer Science*, 13(1):85–108, 1981.
[Seg95] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.
[Voe01] H. Völzer. Randomized non-sequential processes. In *Proceedings of 12th CONCUR*, volume 2154 of *LNCS*, pages 184–201, 2001. Extended version as Technical Report 02-28 - SVRC - University of Queensland.
[Var03] D. Varacca. *Probability, nondeterminism and Concurrency. Two denotational models for probabilistic computation*. PhD thesis, BRICS - Aarhus University, 2003. Available at http://www.brics.dk/∼varacca.
[VVW04] D. Varacca, H. Völzer and G. Winskel. *Probabilistic Event Structures and Domains*. BRICS Technical Report RS-04-10 - Aarhus University, 2004.
[Win82] G. Winskel. Event structure semantics for CCS and related languages. In *Proceedings of 9th ICALP*, volume 140 of *LNCS*, pages 561–576. Springer, 1982.
[Win87] G. Winskel. Event structures. In *Advances in Petri Nets 1986, Part II; Proceedings of an Advanced Course*, Bad Honnef, September 1986, volume 255 of *LNCS*, pages 325–392. Springer, 1987.
[WN95] G. Winskel and M. Nielsen. Models for concurrency. In *Handbook of logic in Computer Science*, volume 4. Clarendon Press, 1995.

Appendix: Domain Theory—Basic Notions

We briefly recall some basic notions of domain theory (see e.g. [AJ94]). A *directed complete partial order (DCPO)* is a partial order where every directed set Y has a least upper bound $\bigsqcup Y$. An element x of a DCPO D is *compact* (or *finite*) if for every directed Y and every $x \leq \bigsqcup Y$ there exists $y \in Y$ such that $x \leq y$. The set of compact elements is denoted by $Cp(D)$. A DCPO is an *algebraic domain* if or every $x \in D$, x is the directed least upper bound of $\downarrow x \cap Cp(D)$. It is ω-*algebraic* if $Cp(D)$ is countable.

In a partial order, two elements are said to be *compatible* if they have a common upper bound. A subset of a partial order is *consistent* if every two of its elements are compatible. A partial order is *coherent* if every consistent set has a least upper bound.

The *Egli-Milner* order on subsets of a partial order is defined by $X \leq Y$ if for all $x \in X$ there exists $y \in Y$, $x \leq y$ and for all $y \in Y$ there exists $x \in X$, $x \leq y$. A subset X of a DCPO is *Scott open* if it is upward closed and if for every directed set Y whose least upper bound is in X, then $Y \cap X \neq \emptyset$. Scott open sets form the *Scott topology*.