Bayesian Network Semantics for Petri Nets☆

Roberto Bruni[a], Hernán Melgratti[b,*], Ugo Montanari[a]

[a] *Dipartimento di Informatica, Università di Pisa, Italy*
[b] *Instituto de Investigación en Ciencias de la Computación (ICC), Universidad de Buenos Aires & Conicet, Argentina*

Abstract

Recent work by the authors equips Petri occurrence nets (PN) with probability distributions which fully replace nondeterminism. To avoid the so-called confusion problem, the construction imposes additional causal dependencies which restrict choices within certain subnets called structural branching cells (s-cells). Bayesian nets (BN) are usually structured as partial orders where nodes define conditional probability distributions. In the paper, we unify the two structures in terms of Symmetric Monoidal Categories (SMC), so that we can apply to PN ordinary analysis techniques developed for BN. Interestingly, it turns out that PN which cannot be SMC-decomposed are exactly s-cells. This result confirms the importance for Petri nets of both SMC and s-cells.

Keywords: Bayesian nets, Petri nets, conditional probability distributions, confusion, branching cells, Kleisli categories, symmetric monoidal categories, forward and backward inference

☆The first author has been partially supported by the MIUR PRIN 201784YSZ5 *ASPRA: Analysis of program analyses* and by University of Pisa PRA_2018_66 *DECLWARE: Metodologie dichiarative per la progettazione e il deployment di applicazioni.* The second author has been partially supported by UBACyT projects 20020170100544BA and 20020170100086BA, and CONICET project PIP 11220130100148CO. The third author carried on part of the work while attending a Program on Logical Structures in Computation at Simons Institute, Berkeley, 2016.

*Corresponding author

Email addresses: `bruni@di.unipi.it` (Roberto Bruni), `hmelgra@dc.uba.ar` (Hernán Melgratti), `ugo@di.unipi.it` (Ugo Montanari)

1. Introduction

At first sight, Bayesian nets (BN) and Petri Nets (PN) have very different purposes: efficient/intelligent analysis of probabilistic distributions for BN, a concurrent, nondeterministic model of computation for PN. But in fact BN and PN share a similar structure: a partial ordering representing incremental, local evolutions via concurrent firings for PN, the introduction of new variables with independent, conditional probabilities for BN.

A closer comparison can be carried on when equipping also PN with a suitable probability structure. A recent approach [1, 2] aims at fully replacing nondeterministic choices with probability distributions, while keeping concurrency expressiveness as much as possible. The problem here is the so-called *confusion*: in PN with confusion, a concurrent computation may exhibit non stable decision steps: delaying a choice may change the available options, due to the action of a concurrent transition.

The simplest example of confusion is the Petri net in Fig. 1(a). Transitions a and b are enabled but in conflict, because they compete for the token in place 1; transition c is also enabled and concurrent w.r.t. a and b; however the firing of transition a enables the transition d that is in conflict with c. As a consequence, the concurrent run where a and c are executed puts in the same equivalence class two quite different traces, where different decisions are taken: (1) if a is executed first, then two choices are taken (a over b and c over d); (2) if c is executed first, then only one choice is taken (a over b). When choices are taken according to some probability distributions, this makes it impossible to assign a unique probability to the concurrent computation with a and c.

The solution proposed by the authors in [2] is to translate the given PN into an equivalent confusionless net (ClPN). This is done by partitioning the net in *structural branching cells* (*s-cells*) where decisions must be resolved. The s-cells of a PN are the equivalence classes of a preorder $\sqsubseteq$, that introduces some further causal dependencies. The preorder is obtained by closing transitively the relation including prime mutual exclusion and immediate causality. It follows that the

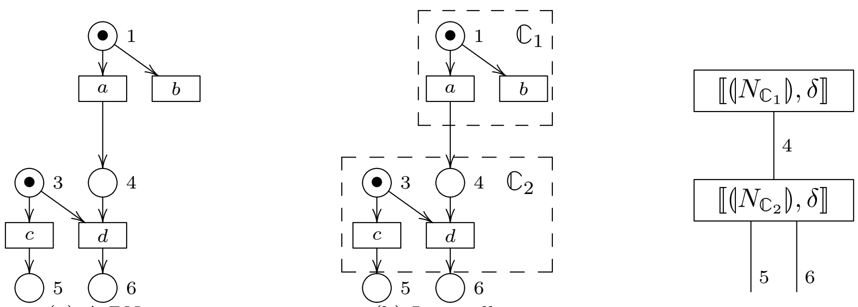

(a) A PN (b) Its s-cells (c) Its BN

Figure 1: A PN with confusion

preorder induces a partial order on s-cells, still denoted $\sqsubseteq$. In the example above there are two s-cells $\mathbb{C}_1 \sqsubseteq \mathbb{C}_2$, meaning that the choice between a and b must be resolved before the one between c and d (see Fig. 1(b)). Each s-cell can then be translated to a confusionless net fragment and all fragments are assembled together, where the dependencies between s-cells are implemented by additional places in a way that corresponds to the execution strategy of [1].

To make confusionless a PN with confusion, it is necessary to delay non stable decisions until any two enabled transitions either do not share any precondition or they share all of them. Then such choice steps are equipped with probability distributions. In practice, our construction introduces a negation place $\overline{p}$ for every place p of the original net, and adds suitable controls to make sure that whenever place $\overline{p}$ becomes inhabited, place p is guaranteed never to become occupied. Thus when the present marking includes $\overline{p}$, all transitions requiring p can be erased and the net simplified. The process is hierarchical, because each s-cell can be further decomposed in smaller s-cells under the assumption that some place $\overline{p}$ becomes inhabited.

The aim of this paper is to show that the partial order of s-cells induces a BN structure. The potential is to develop the countless applications of BN for inference and learning in the context of an expressive model like PN. We propose a strong formal connection between PN and BN via Symmetric Monoidal Categories (SMC).

On the side of BN, convenient categorical presentations have been recently proposed [3, 4, 5] which, in the discrete model, represent BN as string diagrams of a SMC $\mathcal{K}\ell(\mathcal{D})$. Here, objects are natural numbers n which express that 2^n cases are possible, and arrows are rectangular matrices, where rows assign probability distributions on the output cases for every input case. An arrow $f : X \to \mathcal{D}(Y)$ models a conditional probability distribution $P(Y|X)$. Concurrent arrows of string diagrams represent independent probability distributions. Usual inference analysis of BN, like forward and backward inference, Bayesian inversion and disintegration can be made explicit as standard categorical constructions [4].

A ClPN, and thus a PN, can also be mapped to an arrow of $\mathcal{K}\ell(\mathcal{D})$, amenable to the same inference analysis techniques developed for BN. As for our translation PN-ClPN, this mapping is defined by well founded recursion on hierarchical branching cells. Here the effect of positive-negative information $p/\overline{p}$ is played by associating object 1 to a place (that is $2^1 = 2$ cases), which represents explicitly the two options.

Translating a ClPN into a BN is more difficult. In fact, an s-cell may produce several nodes of the BN, since the presence of negative information may break down the cell into a full BN. Thus while in $\mathcal{K}\ell(\mathcal{D})$ associativity of sequential composition takes care of the nested structure, at the level of Bayesian networks it would be necessary to introduce a nested version of BN, which, as far as we know, has not been proposed in the literature.

In Fig. 1(c) we show the BN derived from the PN in Fig. 1(a), represented as a string diagram. There, $N_{\mathbb{C}}$ is the subnet associated with the s-cell $\mathbb{C}$ and δ is the family of probability distributions that rule the choices within $\mathbb{C}_1$ (between a and b) and $\mathbb{C}_2$ (between c and d when place 4 is marked, the trivial choice of c when 4 remains empty, i.e., they are conditional probabilitities depending on the presence/absence of tokens in 4). Roughly, there is one node for each s-cell and wires are associated with places. The first node represents a variable that

$$
\begin{array}{c}
\text{Marked Occurrence net } \mathcal{M} \text{ (Definition 3)} \\
\Big\Downarrow \text{ s-cell decomposition (Proposition 1)} \\
\text{Canonical representation } \mathsf{can}(\mathcal{M}) \text{ (Definition 8)} \\
\Big\Downarrow \text{ compilation (Lemma 4)} \\
\text{Confusion-less representation } (\!|\mathcal{M}|\!) \text{ (Definition 9)} \\
\Big\Downarrow \text{ translation (Proposition 3)} \\
\text{Bayesian Net } [\![(\!|\mathcal{M}|\!), \delta]\!] \text{ (Definition 10)}
\end{array}
$$

Figure 2: Roadmap of technical contribution

may take values $4/\overline{4}$, i.e., it is the arrow

$$
\begin{array}{|c|c|c|}
\hline
 & \emptyset & \{4\} \\
\hline
\emptyset & p_b & p_a \\
\hline
\end{array} : 0 \to 1
$$

where the probabilities p_a and $p_b = 1 - p_a$ are of course determined by δ. The second node represents a variable that may take all combination of values $5/\overline{5}$ and $6/\overline{6}$, conditioned to the value of the first variable, i.e., it is the arrow

$$
\begin{array}{|c|c|c|c|c|}
\hline
 & \emptyset & \{5\} & \{6\} & \{5,6\} \\
\hline
\emptyset & 0 & 1 & 0 & 0 \\
\hline
\{4\} & 0 & p_c & p_d & 0 \\
\hline
\end{array} : 1 \to 2
$$

where, again, the values p_c and $p_d = 1 - p_c$ are drawn by δ. For instance, p_c is the conditional probability that the place 5 is marked given that the place 4 is marked.

To define the arrow in $\mathcal{K}\ell(\mathcal{D})$ that corresponds to a PN we exploit the monoidal category structure of nets and $\mathcal{K}\ell(\mathcal{D})$: first each net is uniquely decomposed in a term of an algebra whose constants are no further hierarchically decomposable s-cells, then a homomorphism returns the corresponding string diagram in $\mathcal{K}\ell(\mathcal{D})$.

More precisely, Fig. 2 summarises the transformation steps that allows us to map a marked occurrence net $\mathcal{M}$ into a Bayesian network. Firstly, we show that any occurrence net $\mathcal{M}$ is an arrow of a strictly symmetric monoidal pre-category, called its canonical representation $\mathsf{can}(\mathcal{M})$, which represents $\mathcal{M}$ as a sequential and parallel composition of s-cells. Then, we show that occurrence nets can be described as terms $(\!|\mathcal{M}|\!)$ of a suitable algebra (see Section 4.1) by decomposing s-cells further: terms make explicit the non-deterministic choices in s-cells and remove confusion. Finally, terms $(\!|\mathcal{M}|\!)$ are mapped into string diagrams $[\![(\!|\mathcal{M}|\!), \delta]\!]$ in the Kleisli category $\mathcal{K}\ell(\mathcal{D})$ of discrete probability distributions.

It is interesting to compare the ClPN and the $\mathcal{K}\ell(\mathcal{D})$ arrow for the same PN. The former model is much more informative in terms of concurrency and causality (see [6] for an event structure theory of persistent nets), while the latter is more straightforward in terms of structure and execution mode. It could be considered a fair algorithmic description of the execution style of [1, 2] original model.

Structure of the paper. In Section 2 we fix the notation, recall the basics of Petri nets and occurrence nets and explain the notion of s-cell from [2]. In Section 3 we provide a novel alternative characterisation of (the pre-oreder induced by) s-cells based on straightforward notion of parallel and sequential (de)composition of nets. This result further justifies the notion of s-cell as basic building block for occurrence nets. In Section 4 we define the mapping from PN to BN. To this aim, an intermediate term algebra is used that builds on the decomposition defined in Section 3 to break s-cells with non-empty initial interface into the hierarchical composition of other terms. Here some sort of case analysis is done: for each marking that can be provided to the s-cell we explore how it can be simplified (the absence of tokens allows for the removal of places and transitions). In Section 5 we show how the Bayesian structure can be exploited to reason about the marking of places of the original PN. Finally, in Section 6 we draw some concluding remarks and give pointers to related and future work.

In Appendix A we show the correspondence between PN decomposition and the approach by Abbes and Benveniste based on event structures, which justifies the assignment of probability distributions to s-cells.

We assume the reader is familiar with some basic concepts from Bayesian networks and category theory.

2. Background

2.1. Notation

We let $\mathbb{N}$ be the set of natural numbers and $2 = \{0, 1\}$. We write U^S for the set of functions from S to U: hence a subset of S is an element of 2^S, and a multiset m over S is an element of $\mathbb{N}^S$. A set can be seen as a multiset whose elements have unary multiplicity. Membership, union, difference and inclusion over sets and multisets are denoted by the (overloaded) symbols: $\in$, $\cup$, $\setminus$ and $\subseteq$, respectively.

Given a relation $R \subseteq S \times S$, we let $R^{-1} = \{(y, x) \mid (x, y) \in R\}$ be its inverse relation, R^+ be its transitive closure and R^* be its reflexive and transitive closure. We say that R is *acyclic* if $\forall s \in S.\ (s, s) \notin R^+$.

2.2. Petri Nets

Definition 1. *A* Petri net *N is a tuple (P, T, F) where: P is the set of places, T is the set of transitions, and $F \subseteq (P \times T) \cup (T \times P)$ is the* flow relation*.*

For $x \in P \cup T$, we denote by ${}^\bullet x = \{y \mid (y, x) \in F\}$ and $x^\bullet = \{z \mid (x, z) \in F\}$ its *pre-set* and *post-set*, respectively. We assume that P and T are disjoint and non-empty and that ${}^\bullet t$ is non empty for every $t \in T$. We write $t : X \to Y$ for $t \in T$ with $X = {}^\bullet t$ and $Y = t^\bullet$. A *marking* is a multiset $m \in \mathbb{N}^P$. A marking denotes a state of a Petri net. We say that the place $p \in P$ is *marked* at m if $p \in m$. We write (N, m) for the net N *marked* by m. In the following we write just N for the marked net $(N, \emptyset)$.

Graphically, a Petri net is a directed bipartite graph whose nodes are the places (circles) and transitions (rectangles) and whose arcs are the elements of F. The marking m is represented by inserting $m(p)$ tokens (bullets) in each place $p \in m$ (see Fig. 3(a)).

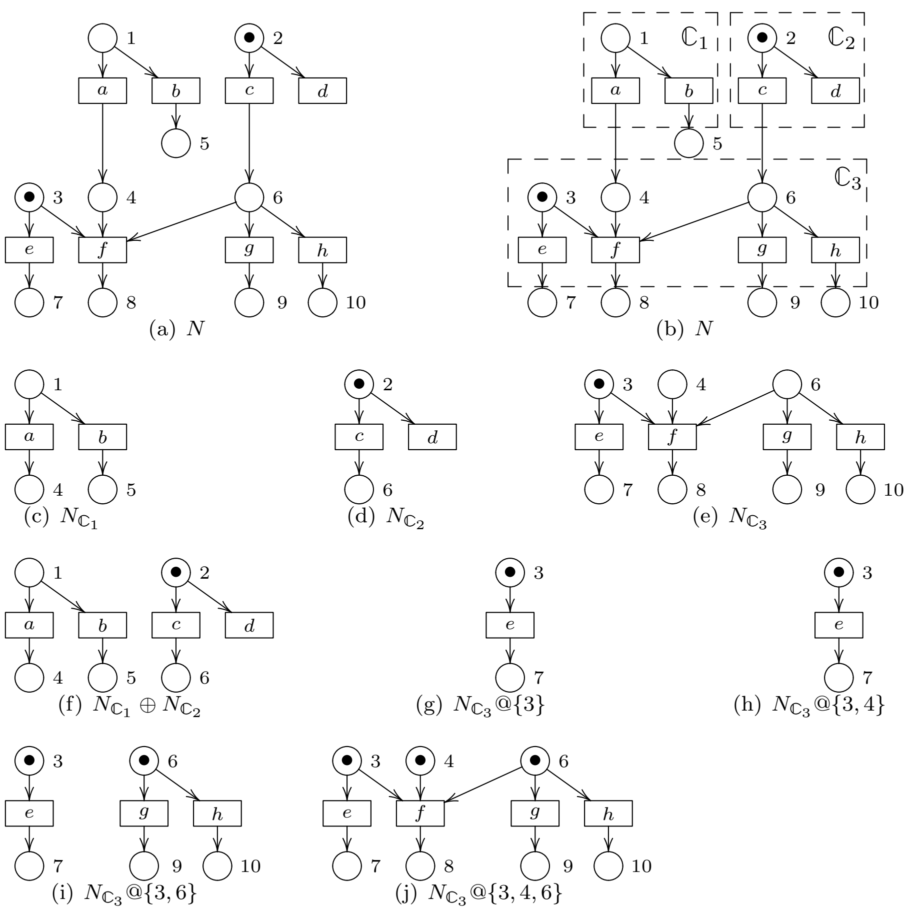

Figure 3: A simple PN

The operational semantics of a Petri net is defined by events called firings. A transition t is *enabled* at the marking m, written $m \xrightarrow{t}$, if ${}^\bullet t \subseteq m$. The *firing* of a transition t enabled at m is written $m \xrightarrow{t} m'$ with $m' = (m \setminus {}^\bullet t) \cup t^\bullet$. A firing sequence $m \xrightarrow{t_1 \cdots t_n} m'$ from m to m' is a finite sequence of firings, sometimes abbreviated $m \rightarrow^* m'$. Moreover, it is *maximal* if no transition is enabled at m'. We say that m' is *reachable* from m if $m \rightarrow^* m'$. The set of markings reachable from m is written $[m\rangle$. A marked net (N, m) is *safe* if each $m' \in [m\rangle$ is a set.

In the rest of the paper we only consider safe nets. More precisely we consider so-called occurrence nets.

2.3. Occurrence nets

We say that a net (P, T, F) is acyclic if its flow relation F is so. Given an acyclic net we let $\preceq = F^*$ be the (reflexive) *causality* relation and say that two transitions t_1 and t_2 are in *immediate conflict*, written $t_1 \#_0 t_2$ if $t_1 \neq t_2 \,\wedge\, {}^\bullet t_1 \cap {}^\bullet t_2 \neq \emptyset$. The *conflict relation* $\#$ is defined by letting $x \# y$ if there are $t_1, t_2 \in T$ such that $(t_1, x), (t_2, y) \in F^+$ and $t_1 \#_0 t_2$.

Definition 2 (Occurrence Net)**.** *A* nondeterministic occurrence net *(or just* occurrence net*) is an acyclic net $\mathcal{O} = (P, T, F)$ such that:*

1. *there are no backward conflicts (i.e., $\forall p \in P.\ |{}^\bullet p| \leq 1$), and*
2. *there are no self-conflicts (i.e., $\forall t \in T.\ \neg(t \# t)$).*

An occurrence net is deterministic if it does not have forward conflicts (i.e., $\forall p \in P.\ |p^\bullet| \leq 1$).

A place p of an occurrence net $\mathcal{O}$ is called *initial* if its pre-set is empty; it is called *final* if its post-set is empty; it is called *isolated* if it is both initial and final. We denote by ${}^\circ\mathcal{O}$ the set of its initial places and by $\mathcal{O}^\circ$ the set of its final places. The net N in Fig. 3(a) is an occurrence net. The sets of its initial and final places respectively are ${}^\circ N = \{1, 2, 3\}$ and $N^\circ = \{5, 7, 8, 9, 10\}$.

Typically it is left implicit that all the initial places of an occurrence net are marked. Here we need to distinguish the cases in which only some initial places are marked.

Definition 3 (Marked Occurrence Net)**.** *A marked occurrence net* $\mathcal{M} = (\mathcal{O}, m)$ *is an occurrence net* $\mathcal{O}$ *together with a subset* m *of initial, non-isolated places.*

The idea is that:

- any initial place in m is already marked (by one token);
- any initial place not in m can receive a token from the context.

Given a marked occurrence net $\mathcal{M} = (\mathcal{O}, m)$, we denote by ${}^\circ\mathcal{M} = {}^\circ\mathcal{O} \setminus m$ the set of its initial (unmarked) places and by $\mathcal{M}^\circ = \mathcal{O}^\circ$ the set of its final places. For the marked occurrence net $(N, \{2, 3\})$ in Fig. 3(a), we have ${}^\circ(N, \{2, 3\}) = \{1\}$ and $(N, \{2, 3\})^\circ = N^\circ = \{5, 7, 8, 9, 10\}$.

A *deterministic nonsequential process* (or just *process*) [7] represents the equivalence class of all firing sequences of a net that only differ in the order in which concurrent firings are executed. It is given as a mapping $\theta : \mathcal{D} \to N$ from a *deterministic occurrence net* $\mathcal{D}$ to N (preserving pre- and post-sets). The firing sequences of a processes $\mathcal{D}$ are its maximal firing sequences starting from the marking ${}^\circ\mathcal{D}$. A process of N is *maximal* if its firing sequences are maximal in N.

When N is an acyclic safe net, the mapping θ is just an injective graph homomorphism: without loss of generality, we name the nodes in $\mathcal{D}$ as their images in N and let θ be the identity.

2.4. Structural Branching Cells

In [2] we have proposed a solution for determining the minimal choice points within an acyclic finite net, called *structural branching cells*: they are subnets where the decision of firing some transition is taken when it is guaranteed that no conflicting transition which is currently not enabled can become enabled in the future.

The construction in [2] takes a (finite) occurrence net as input, which can be, e.g., the (truncated) unfolding of any safe net and returns a partial order of structural branching cells.

To each transition t we assign a unique s-cell $[t]$. This is achieved by taking the equivalence class of t w.r.t. the equivalence relation $\leftrightarrow$ induced by the least preorder $\sqsubseteq$ that includes immediate conflict $\#_0$ and causality $\preceq$. Formally, we let $\sqsubseteq$ be the transitive closure of the relation $\#_0 \cup \preceq \cup \mathsf{Pre}^{-1}$, where $\mathsf{Pre} = F \cap (P \times T)$. This way, each s-cell $[t]$ also includes the places in the pre-sets of the transitions in $[t]$. Since $\#_0$ is subsumed by the transitive closure of the relation $\preceq \cup \mathsf{Pre}^{-1}$, we equivalently set $\sqsubseteq = (\preceq \cup \mathsf{Pre}^{-1})^*$.

Definition 4 (s-cells)**.** *Let $N = (P, T, F)$ be a finite occurrence net and $\sqsubseteq$ defined as above. Let $\leftrightarrow = \{(x, y) \mid x \sqsubseteq y \wedge y \sqsubseteq x\}$. The set $\textsc{bc}(N)$ of* s-cells *is the set of equivalence classes of $\leftrightarrow$, i.e., $\textsc{bc}(N) = \{[t]_{|\leftrightarrow} \mid t \in T\}$.*

We let $\mathbb{C}$ range over s-cells. It is immediate to note that s-cells are ordered by $\sqsubseteq$: we let $\mathbb{C} \sqsubseteq \mathbb{C}'$ if there are $t \in \mathbb{C}, t' \in \mathbb{C}'$ with $t \sqsubseteq t'$.

For any s-cell $\mathbb{C}$, we denote by $N_{\mathbb{C}}$ the subnet of N whose elements are in $\mathbb{C} \cup \bigcup_{t \in \mathbb{C}} t^\bullet$, i.e., we include in $N_{\mathbb{C}}$ also all places in the post-set of some transition in $\mathbb{C}$.

Abusing the notation, we denote by $^\circ\mathbb{C}$ the set of all the initial places in $N_{\mathbb{C}}$ and by $\mathbb{C}^\circ$ the set of all the final places in $N_{\mathbb{C}}$. When the original net (N, m) is marked we sometimes let its cells inherits the marking, i.e., we let the initial marking of $N_{\mathbb{C}}$ be $m \cap {}^\circ\mathbb{C}$.

Example 1. *The net in Fig. 3(a) has three s-cells, which are depicted in Fig. 3(b): $\mathbb{C}_1 = \{1, a, b\}$ concerning the choice between a and b, and $\mathbb{C}_2 = \{2, c, d\}$ concerning the choice between c and d, and $\mathbb{C}_3 = \{3, 4, 6, e, f, g, h\}$. The nets $N_{\mathbb{C}_1}$, $N_{\mathbb{C}_2}$ and $N_{\mathbb{C}_3}$ are respectively shown in Fig. 3(c), 3(d) and 3(e). For $\mathbb{C}_1$, ${}^\circ\mathbb{C}_1 = {}^\circ N_{\mathbb{C}_1} = \{1\}$ and $\mathbb{C}_1^\circ = (N_{\mathbb{C}_1})^\circ = \{4, 5\}$. For $\mathbb{C}_2$, ${}^\circ\mathbb{C}_2 = {}^\circ N_{\mathbb{C}_2} \setminus \{2\} = \{2\} \setminus \{2\} = \emptyset$ and $\mathbb{C}_2^\circ = (N_{\mathbb{C}_2})^\circ = \{6\}$.*

The behaviour of a branching cell is characterised in terms of all its possible executions.

Definition 5 (Transactions)**.** *Let $\mathbb{C} \in \textsc{bc}(N)$ and $m = {}^\circ\mathbb{C}$. Then, a* transaction *θ of $\mathbb{C}$, written $\theta : \mathbb{C}$, is a maximal (deterministic) process of $(N_{\mathbb{C}}, m)$. We* *denote by* $\Theta(\mathbb{C})$ *the set of all the transactions of* $\mathbb{C}$.

Since the set of transitions in a transaction θ uniquely determines the corresponding process in $N_{\mathbb{C}}$, we write a transaction θ simply as the set of its transitions. If $i = {}^{\circ}\theta$ is the set of initial places of θ and $o = \theta^{\circ}$ is the set of its final places, we write $\theta : i \to o$. Note that in general, for $\theta : i \to o \in \Theta(\mathbb{C})$, we have $i \subseteq {}^{\circ}\mathbb{C}$ and $o \subseteq \mathbb{C}^{\circ}$. We write $\mathsf{n}(\theta)$ for the set of transitions and places of θ.

Example 2. *Consider the net* $N_{\mathbb{C}_3}$ *in Fig. 3(e). It has the following three transactions:* $\theta_1 = \{f\}$, $\theta_2 = \{e, g\}$ *and* $\theta_3 = \{e, h\}$, *with* $\theta_1 : \{3, 4, 6\} \to \{8\}$ $\theta_2 : \{3, 6\} \to \{7, 9\}$ $\theta_3 : \{3, 6\} \to \{7, 10\}$.

3. Petri Nets Decomposition

We have already said that s-cells form a partial order. Here we show that it can be seen as a particular monoidal category structure.

We proceed as follows:

1. we define set-theoretical parallel and sequential composition of nets;
2. we show that parallel and sequential composition, together with a suitable notion of identities, induce a strictly symmetric monoidal category structure over occurrence nets;
3. we show that s-cells are neither decomposable in parallel nor in series;
4. we show that each Petri net admits a unique maximal decomposition in terms of parallel and sequence (up to the axioms of strictly symmetric monoidal categories) and that such decomposition coincides with the partial order of s-cells.

This provides a new characterisation of s-cells as the building blocks of occurrence nets that supports our intuition about their relevance.

Intuitively, parallel composition takes two nets and put them side by side.

Definition 6 (Parallel composition)**.** *Let* (P_1, T_1, F_1, m_1) *and* (P_2, T_2, F_2, m_2) *be two Petri nets whose nodes are disjoint (i.e., with* $(P_1 \cup T_1) \cap (P_2 \cup T_2) = \emptyset$*).*

Their parallel composition is given by the element-wise union of their components:

$$(P_1, T_1, F_1, m_1) \oplus (P_2, T_2, F_2, m_2) = (P_1 \cup P_2, T_1 \cup T_2, F_1 \cup F_2, m_1 \cup m_2)$$

Sequential composition is defined over (marked) occurrence nets only.

Definition 7 (Sequential composition)**.** *Let $\mathcal{M}_1 = (\mathcal{O}_1, m_1)$ and $\mathcal{M}_2 = (\mathcal{O}_2, m_2)$ be two marked occurrence nets, with $\mathcal{O}_j = (P_j, T_j, F_j)$ for $j = 1, 2$, whose nodes are disjoint except for the final places of $\mathcal{M}_1$ that are identical to the unmarked initial places of $\mathcal{M}_2$ (i.e., with $\mathcal{M}_1^\circ = (P_1 \cup T_1) \cap (P_2 \cup T_2) = {}^\circ\mathcal{M}_2$). Their sequential composition is given by the element-wise union of their components (but note that the places in $(\mathcal{M}_1^\circ = {}^\circ\mathcal{M}_2$ are shared):*

$$(P_1, T_1, F_1, m_1); (P_2, T_2, F_2, m_2) = (P_1 \cup P_2, T_1 \cup T_2, F_1 \cup F_2, m_1 \cup m_2)$$

Let us write $\mathcal{M} : i \to o$ for a marked occurrence net with $i = {}^\circ\mathcal{M}$ and $o = \mathcal{M}^\circ$ Then we note that for $\mathcal{M}_j : i_j \to o_j$ for $j \in [1, 4]$:

- $\mathcal{M}_1 \oplus \mathcal{M}_2 : i_1 \cup i_2 \to o_1 \cup o_2$, when the parallel composition is defined;
- $\mathcal{M}_1; \mathcal{M}_2 : i_1 \to o_2$, when the sequential composition is defined;
- parallel composition is commutative and associative and has the empty net $\mathbf{0} = (\emptyset, \emptyset, \emptyset, \emptyset) : \emptyset \to \emptyset$ as neutral element, i.e. it forms a commutative monoid;
- sequential composition is associative;
- for each set of places i the identity net $I_i = (i, \emptyset, \emptyset, \emptyset) : i \to i$ consisting just of (unmarked) isolated places i behaves as the identity w.r.t. composition;
- the monoid of parallel composition is functorial: $I_\emptyset = \mathbf{0}$, $I_{i_1 \cup i_2} = I_{i_1} \oplus I_{i_2}$ and $(\mathcal{M}_1; \mathcal{M}_2) \oplus (\mathcal{M}_3; \mathcal{M}_4) = (\mathcal{M}_1 \oplus \mathcal{M}_3); (\mathcal{M}_2 \oplus \mathcal{M}_4)$.

In the following, we assume $\oplus$ has higher precedence over ;, e.g. we write $\mathcal{M}_1 \oplus \mathcal{M}_2; \mathcal{M}_3$ instead of $(\mathcal{M}_1 \oplus \mathcal{M}_2); \mathcal{M}_3$.

From the above we get that marked occurrence nets form the arrows of a strictly symmetric (strict) monoidal pre-category. We get a pre-category and not a category just because parallel and sequential compositions are defined on concrete nets and impose some disjointness requirements on their places and transitions, i.e. they are partial operations instead of total ones. This is not an issue here because we are mainly interested in decomposing concrete nets into parts, not in building new nets and we aim to have a unique decomposition. As the parts are obtained by decomposition, it is guaranteed that they can be reassembled later. Another point that is worth paying attention to is the fact that parallel composition is commutative and thus the monoidal (pre-)category we get is *strictly* symmetric: its symmetries are identities. In Section 4.3, when translating nets to string diagrams in $\mathcal{K}\ell(\mathcal{D})$ that are arrows of a symmetric monoidal category, we need to be parametric w.r.t. some fixed orders of initial and final places. However, we prove that the actual choice of a given order is inessential (see Proposition 3).

One possible alternative to carry out the translation as a proper symmetric monoidal functor between symmetric monoidal categories, would be to consider nets up to isomorphism and equip them with some sort of ordered interface from the very beginning. We prefer not to do so for several reasons. First, standard Petri nets (as well as s-cells) do not come with ordered interfaces. Second, decomposition should take into account also symmetries besides s-cells and identities, so that its uniqueness would only hold up to the axioms of symmetric monoidal categories and there would be special nets (with no transitions) representing symmetries. Third, while in Section 4.1 we propose a syntax for representing concrete nets (in canonical forms), if we choose to work up to net isomorphism, then the description of s-cells becomes more complicated because places and transitions names would have only a local scope and some canonical choice of names would be needed. Additionally, symmetries should be considered in the syntax.

One important issue to bear in mind is that parallel and sequential compositions as defined here give a precise characterisation of s-cells.

In the literature, many other approaches to net composition have been studied. Among the most recent ones we mention [8, 9, 10, 11]. There the objective is typically the generation of all nets starting from a small number of components [9, 10, 11] or the preservation of the semantics [8] via composition. The work in [11] focuses on the algebraic structure of non-deterministic computations but is based on pre-nets [12] rather than Petri nets. The difference is that the pre- and post-sets of transitions are lists and not sets, so that the interfaces for composition are intrinsically ordered. The work in [9, 10] represents nets as some sort of circuits generated by elementary blocks. There interfaces are lists of ports (neither places nor transitions) that are used by sequential composition to perform some non-trivial synchronisation between transitions. If we were to exploit their circuit algebra for decomposition, then the characterisation of s-cells would be difficult to obtain, because decomposition is too fine grained. The work in [8] exploits a colimit-based composition operation that is suited to model interaction between (open) nets and their sorrounding environment. The main result there is that several behavioural equivalences are congruences w.r.t. composition. Again, applying this approach to decompose nets would not allow to characterise s-cells.

Example 3. *Consider the marked occurrence nets* $N_{\mathbb{C}_1} : \{1\} \rightarrow \{4, 5\}$*,* $(N_{\mathbb{C}_2}, \{2\}) : \emptyset \rightarrow \{6\}$*, and* $(N_{\mathbb{C}_3}, \{3\}) : \{4, 6\} \rightarrow \{7, 8, 9, 10\}$ *in Fig. 3(c), 3(d) and 3(e). Note that the parallel composition of* $N_{\mathbb{C}_1}$ *and* $N_{\mathbb{C}_2}$ *is defined because the nets neither share places nor transitions. The resulting net* $N_{\mathbb{C}_1} \oplus (N_{\mathbb{C}_2}, \{2\}) : \{1\} \rightarrow \{4, 5, 6\}$ *is shown in Fig 3(f). We remark that neither* $N_{\mathbb{C}_1} \oplus (N_{\mathbb{C}_3}, \{3\})$ *nor* $(N_{\mathbb{C}_2}, \{2\}) \oplus (N_{\mathbb{C}_3}, \{3\})$ *are defined because* $N_{\mathbb{C}_3}$ *shares the place* 4 *with* $N_{\mathbb{C}_1}$ *and the place* 6 *with* $N_{\mathbb{C}_2}$*. Similarly, note that none of the considered occurrence nets can be composed sequentially, because their interfaces do not match. For instance, the final place* 5 *of* $N_{\mathbb{C}_1} \oplus (N_{\mathbb{C}_2}, \{2\}) : \{1\} \rightarrow \{4, 5, 6\}$ *does not appear as an initial place of* $(N_{\mathbb{C}_3}, \{3\}) : \{4, 6\} \rightarrow \{7, 8, 9, 10\}$*. We can fix this mismatch by considering the net* $I_{\{5\}} : \{5\} \rightarrow \{5\}$ *and noting that* $(N_{\mathbb{C}_3}, \{3\}) \oplus I_{\{5\}} : \{4, 6, 5\} \rightarrow \{7, 8, 9, 10, 5\}$ *is well defined. Then,*

$$N_{\mathbb{C}_1} \oplus (N_{\mathbb{C}_2}, \{2\}); (N_{\mathbb{C}_3}, \{3\}) \oplus I_{\{5\}} : \{1\} \rightarrow \{5, 7, 8, 9, 10\}$$

stands for the net N *in Fig. 3(a).*

A marked occurrence net is called *trivial* if it has no transitions.

We say a marked occurrence net $\mathcal{M}$ is decomposable in parallel if there exists two non-trivial marked occurrence nets $\mathcal{M}_1$ and $\mathcal{M}_2$ such that $\mathcal{M} = \mathcal{M}_1 \oplus \mathcal{M}_2$. Similarly, we say that it is decomposable in series if there exists two non-trivial marked occurrence nets $\mathcal{M}_1$ and $\mathcal{M}_2$ such that $\mathcal{M} = \mathcal{M}_1; \mathcal{M}_2$.

Lemma 1. *Any s-cell $N_{\mathbb{C}}$ cannot be decomposed in series and in parallel.*

Proof. By contraposition, it is immediate to prove that the sequential/parallel composition of two non-trivial nets is not an s-cell. □

Proposition 1. *Any marked occurrence net can be uniquely decomposed as the parallel and sequential composition of its s-cells (and identities), up to the axioms of strictly symmetric monoidal pre-categories.*

Proof. For the existence, the partial order of s-cell (is unique and it) induces a decomposition of the net. For instance this can be done by stratifying the s-cells in layers $L_1, ..., L_n$ where each layer L_j is the (largest) parallel composition of some identity I_{s_j} with all s-cells whose predecessors are in layers $L_1, ..., L_{j-1}$ and then taking their sequential composition $L_1; ...; L_n$.

For uniqueness, suppose two different decompositions can be found, then they must have the same s-cells (because s-cells are not decomposable) ordered in the same way (because the ordering is induced by the places they share), hence they coincide. □

Definition 8 (Canonical form)**.** *Given a marked occurrence net $\mathcal{M}$ we denote by $\mathsf{can}(\mathcal{M})$ its unique decomposition.*

Example 4. *The canonical form of $(N, \{2, 3\})$ in Fig. 3(a) is given by the decomposition below, already discussed in Example 3:*

$$N_{\mathbb{C}_1} \oplus (N_{\mathbb{C}_2}, \{2\}); (N_{\mathbb{C}_3}, \{3\}) \oplus I_{\{5\}} : \{1\} \to \{5, 7, 8, 9, 10\}$$

A s-cell can be itself a fairly complicated fragment. To ease the translation to Bayesian nets we would like to exploit some form of induction over the structure of s-cells themselves. This can be done by some sort of hierarchical decomposition, to be defined next, where each s-cell is studied according to its possible dynamic activations: depending on the conditions under which an s-cell is enabled, some alternatives can be immediately discarded, the structure of the s-cells can be simplified and further decomposed. In this way, the behaviour of an s-cell can be entirely defined by collecting its decompositions for each possible initial marking.

3.1. Place Removal

Given a possibly marked s-cell $N_{\mathbb{C}} : i \to o$ (with $i \neq \emptyset$), we are interested in studying what happens under the hypothesis that some tokens arrive in a subset of places $m \subseteq i$ while the places in $s = i \setminus m$ are guaranteed to stay empty (i.e., they are *dead*). In fact it can happen that the removal of the places in s and of the transitions and places that causally depend on them[1] will allow to further decompose the s-cell.

We let $N_{\mathbb{C}} \ominus s$ be the net obtained by removing all dead nodes as explained above. Additionally, isolated places are also removed. The cancellation of some transitions can break the equivalence class induced by $\sqsubseteq$, which explains why $N_{\mathbb{C}} \ominus s$ is not necessarily an s-cell. Also note that some of the final places of $N_{\mathbb{C}}$ can become dead and canceled. The final dead places can be computed by taking $N_{\mathbb{C}}^{\circ} \setminus (N_{\mathbb{C}} \ominus s)^{\circ}$. Thus in general we have $N_{\mathbb{C}} \ominus s : i' \to o'$ for some $i' \subseteq i \setminus s$ and $o' \subseteq o$. We write $N_{\mathbb{C}}@m$ for the marked net $(N_{\mathbb{C}} \ominus s, {}^{\circ}(N_{\mathbb{C}} \ominus s)) : \emptyset \to o'$, where $N_{\mathbb{C}} : i \to o$ and $s = i \setminus m$, i.e., for the net $N_{\mathbb{C}} \ominus s$ whose initial places are all marked.

To some extent the behaviour of an s-cell is determined by considering its behaviour under all possible initial markings. Consequently we can further explore the behaviour of $N_{\mathbb{C}} : i \to o$ by considering $N_{\mathbb{C}}@m$ for all $m \subseteq i$.

Example 5. *Consider the s-cell* $(N_{\mathbb{C}_3}, \{3\}) : \{4, 6\} \to \{7, 8, 9, 10\}$ *in Fig. 3(e).*

[1]In such cases, all the transitions that depend on some place in s cannot be fired and the places in their post-set are also dead.

The behaviour of $(N_{\mathbb{C}_3}, \{3\})$ *can be explained by considering all the possible ways in which its initial places* 4 *and* 6 *can be marked: none of them is marked (i.e.,* $N_{\mathbb{C}_3}@\{3\}$*), just one of them is marked (i.e., either* $N_{\mathbb{C}_3}@\{3,4\}$ *or* $N_{\mathbb{C}_3}@\{3,6\}$*), or both of them are marked (i.e.,* $N_{\mathbb{C}_3}@\{3,4,6\}$*). Net* $N_{\mathbb{C}_3}@\{3\}$ *depicted in Fig. 3(g) is obtained by removing from* $N_{\mathbb{C}_3}$ *the initial places* 4 *and* 6*, and all the elements that causally depends on them, i.e., the transitions* f*,* g *and* h *and the places* 7*,* 8*,* 9 *and* 10*. The remaining nets are in Fig. 3(h)-3(j). It is worth noticing that in* $N_{\mathbb{C}_3}@\{3,4\}$ *the place* 4 *is also removed from* $N_{\mathbb{C}_3} \ominus \{6\}$ *because, after removing the place* 6 *and thus the transition* f*, the place* 4 *remains isolated.*

4. Compiling nets

In this section we associate each finite occurrence net with a string diagram in the Kleisli category $\mathcal{K}\ell(\mathcal{D})$ of discrete probability distributions. This is achieved in two steps. We first introduce a language for representing occurrence nets and show how the s-cell decomposition can be used to associate each occurrence net with a particular term. Then, we map terms into string diagrams in $\mathcal{K}\ell(\mathcal{D})$.

4.1. Language of nets

The decomposition of a net in branching cells can be described by terms generated by the following grammar, where m, s are sets of places and Θ is a set of transactions:

$$T ::= I_s \mid \bot_s \mid T \oplus T \mid T;T \mid \mathsf{C}(\Theta) \mid \sum_{m \subseteq s} m \triangleright T$$

Here the idea is that $\mathsf{C}(\Theta)$ denotes a basic building block consisting of the set of transactions of an s-cell whose initial places are all marked. The case of an s-cell $\mathbb{C}$ with a set of unmarked initial places s is represented as the formal sum $\sum_{m \subseteq s} m \triangleright T$, where all the possibile ($2^{|s|}$) initial markings m are considered, each paired with the encoding of $N_{\mathbb{C}}@m$. This accounts for the hierarchical decomposition of s-cells. The term I_s denotes the identity net, consisting just of a set of unmarked places with no transitions (i.e., all places are initial and final).

The term $\bot_s$ denote a net with no initial places and no transitions, whose only final places are s (i.e., the places s are dead). The terms $T \oplus T$ and $T; T$ denote respectively the composition in parallel and in series.

The terms of the algebra are taken up to the axioms of strictly symmetric monoidal (pre-)categories, where additionally we have $\bot_\emptyset = I_\emptyset$ and $\bot_{s_1 \cup s_2} = \bot_{s_1} \oplus \bot_{s_2}$.

4.1.1. Typing

Not all terms are valid though. We introduce a type system to discard ill-formed terms. Our types are triples of the form (i, s, o) where i is the set of initial unmarked places, s is the set of all places and transitions appearing in a term and o is the set of final places.

We write $T : i \xrightarrow{s} o$ for $T : (i, s, o)$. The typing rules are in Fig. 4. The rules for I_s and $\bot_s$ are self-explanatory. The rule for $\oplus$ states that a term is well-typed when its subterms are well-typed and do not share place nor transitions (i.e., $s \cap s' = \emptyset$). The case of sequential composition $T; T'$ additionally requires that the set of final places of T coincides with the set of the initial unmarked places of T'. The rule for $\sum_{m \subseteq i} m \triangleright T_m$ requires all subterms T_m to have the same sets of initial and final places (respectively, $\emptyset$ and o), which captures the idea that a sum represents the execution of a s-cell under all possible markings. The rule for $\mathsf{C}(\Theta)$ follows immediately.

Lemma 2. *If $T : i \xrightarrow{s} o$ then $i \cup o \subseteq s$.*

Proof. The proof is by rule induction. □

Typing is unique, as stated by the following result.

Lemma 3. *If $T : i \xrightarrow{s} o$ and $T : i' \xrightarrow{s'} o'$ then $i = i'$, $o = o'$, $s = s'$.*

Proof. The proof is by rule induction. □

Hereafter we assume terms to be well-typed.

$$\frac{}{I_s : s \xrightarrow{s} s} \qquad \frac{}{\bot_s : \emptyset \xrightarrow{s} s} \qquad \frac{T : i \xrightarrow{s} o \quad T' : i' \xrightarrow{s'} o' \quad s \cap s' = \emptyset}{T \oplus T' : i \cup i' \xrightarrow{s \cup s'} o \cup o'}$$

$$\frac{T : i \xrightarrow{s} m \quad T' : m \xrightarrow{s'} o \quad s \cap s' = m}{T;T' : i \xrightarrow{s \cup s'} o} \qquad \frac{\forall m \subseteq i.\ T_m : \emptyset \xrightarrow{s_m} o \quad s = \bigcup_{m \subseteq i} s_m}{\sum_{m \subseteq i} m \triangleright T_m : i \xrightarrow{s} o}$$

$$\frac{o = \bigcup_{\theta \in \Theta} \theta^\circ \qquad s = \bigcup_{\theta \in \Theta} \mathbf{n}(\theta)}{\mathsf{C}(\Theta) : \emptyset \xrightarrow{s} o}$$

Figure 4: Type system

4.2. From Nets to Terms

In this section we introduce a mapping from occurrence nets to terms.

Definition 9. *Let $\mathcal{M}$ be a marked occurrence net. The corresponding term $\llparenthesis \mathcal{M} \rrparenthesis$ is given by the homomorphic extension (w.r.t. identitites, parallel and sequential composition)*[2] *of the encoding defined below over s-cells.*

$$\llparenthesis N_{\mathbb{C}}, i \rrparenthesis \begin{cases} \mathsf{C}(\Theta(N_{\mathbb{C}})) & \textit{if } {}^\circ N_{\mathbb{C}} = i \qquad \text{(1a)} \\ \displaystyle\sum_{m \subseteq {}^\circ(N_{\mathbb{C}}, i)} m \triangleright (\bot_{d_m} \oplus T_m) & \textit{otherwise} \qquad \text{(1b)} \end{cases}$$

$$\textit{where:} \begin{cases} N_m = N_{\mathbb{C}} @ i \cup m \\ T_m = \llparenthesis \mathsf{can}(N_m) \rrparenthesis \\ d_m = N_{\mathbb{C}}^\circ \setminus N_m^\circ \end{cases}$$

The encoding of a marked s-cell $\mathbb{C}$ considers two cases: (i) all initial places of the s-cell are marked (Eq. 1a); and (ii) some initial tokens are unmarked. In the first case, a completely marked s-cell is mapped to the term $\mathsf{C}(\Theta(N_{\mathbb{C}}))$ that

[2]This just means that $\llparenthesis I_s \rrparenthesis = I_s$, $\llparenthesis \mathcal{M}_1 \oplus \mathcal{M}_2 \rrparenthesis = \llparenthesis \mathcal{M}_1 \rrparenthesis \oplus \llparenthesis \mathcal{M}_2 \rrparenthesis$ and $\llparenthesis \mathcal{M}_1 ; \mathcal{M}_2 \rrparenthesis = \llparenthesis \mathcal{M}_1 \rrparenthesis ; \llparenthesis \mathcal{M}_2 \rrparenthesis$.

describes all the possible executions of $N_{\mathbb{C}}$, i.e., its transactions. Differently, when some initial places are unmarked, the corresponding term is obtained by composing the behaviour of the s-cell under each possible marking $m \subseteq {}^\circ(N_{\mathbb{C}}, i)$. The term $m \triangleright (\bot_{d_m} \oplus T_m)$ describes the behaviour of $\mathbb{C}$ when all places in $i \cup m$ are marked and the remaining initial places are dead. For this reason, $\bot_{d_m}$ and T_m are defined in terms of the net $N_m = N_{\mathbb{C}}@i \cup m$. The term $\bot_{d_m}$ stands for the final places that are dead when the initial marking is $i \cup m$. The term T_m encodes the net $N_{\mathbb{C}}@i \cup m$: we just remark here, as already mentioned, that we need to compute the canonical form of N_m, because removing elements from $\mathbb{C}$ may originate a complex net an not an s-cell (as for $N_{\mathbb{C}_3}@\{3, 6\}$ in Fig. 3(i)).

Lemma 4. *For any finite occurrence net N and marking $m \subseteq {}^\circ N$, $(\!|N, m|\!)$ is defined, unique (up-to the structure of strictly symmetric monoidal pre-categories) and well-typed.*

Example 6. *Consider the marked occurrence net $(N, \{2, 3\})$ in Fig. 3(a), whose canonical form is in Example 4*

$$(N, \{2, 3\}) = N_{\mathbb{C}_1} \oplus (N_{\mathbb{C}_2}, \{2\}); (N_{\mathbb{C}_3}, \{3\}) \oplus I_{\{5\}}$$

Then, the corresponding term is obtained by

$$(\!|N, \{2, 3\}|\!) = (\!|N_{\mathbb{C}_1}|\!) \oplus (\!|N_{\mathbb{C}_2}, \{2\}|\!); (\!|N_{\mathbb{C}_3}, \{3\}|\!) \oplus (\!|I_{\{5\}}|\!) \tag{2}$$

The term $(\!|N_{\mathbb{C}_1}|\!)$ is obtained by applying Eq. (1b) *because $i = \emptyset$ and ${}^\circ N_{\mathbb{C}_1} = \{1\} \neq \emptyset$ (see $N_{\mathbb{C}_1}$ in Fig. 3(c)). Then,*

$$(\!|N_{\mathbb{C}_1}|\!) = \emptyset \triangleright (\bot_{d_\emptyset} \oplus T_\emptyset) + \{1\} \triangleright (\bot_{d_{\{1\}}} \oplus T_{\{1\}}) \tag{3}$$

Note that $N_\emptyset = N_{\mathbb{C}_1}@\emptyset$ is obtained from $N_{\mathbb{C}_1}$ by removing all elements that depends on the unique unmarked initial place 1*. Hence, $N_\emptyset = N_{\mathbb{C}_1}@\emptyset = \mathbf{0} = I_\emptyset$. Consequently, $T_\emptyset = (\!|N_m|\!) = I_\emptyset$. Moreover $d_\emptyset = \{4, 5\}$.*

For the marking $\{1\}$, we have $N_{\{1\}} = N_{\mathbb{C}_1}@\{1\} = (N_{\mathbb{C}_1}, \{1\})$. Since $N_{\mathbb{C}_1}$ is an s-cell, $\mathsf{can}(N_{\mathbb{C}_1}@\{1\}) = (N_{\mathbb{C}_1}, \{1\})$*. Therefore, $T_{\{1\}} = (\!|N_{\mathbb{C}_1}, \{1\}|\!)$, which is obtained by using Eq.* (1a)*. The net $N_{\mathbb{C}_1}$ has two transactions, one for each* *transition, i.e.,* $\Theta(N_{\mathbb{C}_1}) = \{\{a\},\{b\}\}$*. Then,* $T_{\{1\}} = \mathsf{C}(\{\{a\},\{b\}\})$*. Moreover,* $d_{\{1\}} = \emptyset$ *because* $(N_{\{1\}})^\circ = (N_{\mathbb{C}_1},\{1\})^\circ = N^\circ_{\mathbb{C}_1}$*. Consequently,*

$$\begin{aligned} \llparenthesis N_{\mathbb{C}_1} \rrparenthesis &= \emptyset \triangleright (\bot_{\{4,5\}} \oplus I_\emptyset) + \{1\} \triangleright (\bot_\emptyset \oplus \mathsf{C}(\{\{a\},\{b\}\})) \\ &= \emptyset \triangleright \bot_{\{4,5\}} + \{1\} \triangleright \mathsf{C}(\{\{a\},\{b\}\}) \end{aligned} \tag{4}$$

Intuitively, the term $\emptyset \triangleright \bot_{\{4,5\}}$ *states that the s-cell* $\mathbb{C}_1$ *does not generate any token in its final places when the initial place* 1 *remains unmarked. Differently,* $\{1\} \triangleright \mathsf{C}(\{\{a\},\{b\}\})$ *describes the behaviour of* $\mathbb{C}_1$ *when its initial place is marked. In this case, the behaviour corresponds to the non-deterministic choice of the transactions* $\{a\}$ *and* $\{b\}$*.*

The encoding of $(N_{\mathbb{C}_2},\{2\})$ *is obtained by using Eq.* (1a),

$$\llparenthesis N_{\mathbb{C}_2},\{2\} \rrparenthesis = \mathsf{C}(\{\{c\},\{d\}\}) \tag{5}$$

For $(N_{\mathbb{C}_3},\{3\})$*, we obtain the following term by analogous calculations*

$$\begin{aligned} \llparenthesis N_{\mathbb{C}_3},\{3\} \rrparenthesis = &\quad \emptyset \triangleright (\bot_{\{8,9,10\}} \oplus \mathsf{C}(\{\{e\}\})) \\ &+ \{4\} \triangleright (\bot_{\{8,9,10\}} \oplus \mathsf{C}(\{\{e\}\})) \\ &+ \{6\} \triangleright (\bot_{\{8\}} \oplus \mathsf{C}(\{\{e\}\}) \oplus \mathsf{C}(\{\{g\},\{h\}\})) \\ &+ \{4,6\} \triangleright \mathsf{C}(\{\{f\},\{e,g\},\{e,h\}\}) \end{aligned} \tag{6}$$

which describes the behaviour of $\mathbb{C}_3$ *for every possible initial marking of its initial places (i.e.,* $\emptyset$*,* $\{4\}$*,* $\{6\}$*, and* $\{4,6\}$*). The most interesting case is the subterm* $\{6\} \triangleright (\bot_{\{8\}} \oplus \mathsf{C}(\{\{e\}\}) \oplus \mathsf{C}(\{\{g\},\{h\}\}))$ *obtained from* $\{6\} \triangleright (\bot_{d_{\{6\}}} \oplus T_{\{6\}})$*. Consider the net* $N_{\{6\}} = (N_{\mathbb{C}_3}@\{3,6\})$ *in Fig. 3(i), which contains two s-cells. Consequently, its canonical form is given by the parallel composition of two s-cells, which are respectively encoded as* $\mathsf{C}(\{\{e\}\})$ *and* $\mathsf{C}(\{\{g\},\{h\}\})$*.*

Finally,

$$\llparenthesis I_{\{5\}} \rrparenthesis = I_{\{5\}} \tag{7}$$

To show that the term $\llparenthesis N, m \rrparenthesis$ is a good representative of the probabilistic semantics of N, we prove that it characterises the configurations allowed by the semantics of Abbes and Benveniste. The interested reader can find all technical details in the Appendix.

4.3. From Terms to $\mathcal{K}\ell(\mathcal{D})$

Given a set X, a discrete probability distribution with finite support over X is a function $\omega : X \to [0,1]$ such that $\sum_{x \in X}^{n} \omega(x) = 1$ and $\mathsf{supp}(\omega) = \{x \in X \mid \omega(x) > 0\}$ is a finite set. The function ω can be sometimes written as the formal convex combination[3]

$$\omega = r_1|x_1\rangle + ... + r_n|x_n\rangle$$

where $\mathsf{supp}(\omega) = \{x_1, ..., x_n\}$ and $r_j = \omega(x_j)$ for $j \in [1, n]$. We let $\mathcal{D}(X)$ be the set of discrete probability distributions ω over X and write $\mathcal{D}$ for the discrete probability monad over the category **Set** of sets (as objects) and functions (as arrows). The category $\mathcal{K}\ell(\mathcal{D})$ is the Kleisli category of the monad $\mathcal{D}$: its objects are sets, its arrows $f : X \to Y$ are functions $f : X \to \mathcal{D}(Y)$. It has been shown in [3] that $\mathcal{K}\ell(\mathcal{D})$ forms a symmetric monoidal category and that Bayesian networks can be seen as special kinds of arrows in $\mathcal{K}\ell(\mathcal{D})$ that can be represented as string diagrams using wire-and-box notation (see also [13]). According to this view, a diagram from n to k represents an arrow from 2^n to 2^k in $\mathcal{K}\ell(\mathcal{D})$.

We next show how to interpret Petri nets as Bayesian networks by exploiting $\mathcal{K}\ell(\mathcal{D})$. To this aim we need to map the arrows of a strictly symmetric monoidal pre-category to those of a symmetric monoidal category: in the first case the objects are sets of places, while in the latter they are natural numbers representing a totally ordered set of ports. Therefore the mapping is defined parametrically on some arbitrarily chosen total orders of initial and final places.

Given a set of places s, we let π_s denote a bijective function $\pi_s : s \to |s|$ that assigns a position to each element of s. We write π when the set s is implicit. Overloading the notation, we let π also denote the string such that the place $p \in s$ appears in position $\pi(p)$. Note that π is without repetitions: each $p \in s$ appears exactly once in π. We let ϵ denote the empty string (over the empty set of places). For $p \in s$ and $m \subseteq s$, we also write $p \in \pi$ and $m \subseteq \pi$ when π is a linearization of s.

[3]The 'ket' notation $r|x\rangle$ has no particular meaning: it is just syntactic sugar.

Given π and π' two such strings over s, we let $\chi^{\pi}_{\pi'} : |s| \to |s|$ denote the unique permutation that swaps π into π', i.e. such that for any $p \in s$ we have $\chi^{\pi}_{\pi'}(\pi(p)) = \pi'(p)$. By coherence of symmetries we have, e.g., $\chi^{\pi}_{\pi'}; \chi^{\pi'}_{\pi''} = \chi^{\pi}_{\pi''}$.

Given two strings π over s and π' over s' with $s \cap s' = \emptyset$ we use juxtaposition to denote the string $\pi\pi'$ over $s \cup s'$ such that $(\pi\pi')(p) = \pi(p)$ if $p \in s$ and $(\pi\pi')(p) = |s| + \pi'(p)$ if $p \in s'$.

As a matter of notation, we assume that a string π over s implicitly defines an ordering over 2^s, e.g., a subset of s can be seen as a binary string of length $|s|$, which are then ordered lexicographically. Correspondingly, the permutation $\chi^{\pi}_{\pi'} : |s| \to |s|$ induces an isomorphism on 2^s, that we denote with the same name $\chi^{\pi}_{\pi'}$.

In the following we assume a function δ is given that associates every constant $\mathsf{C}(\Theta)$ with a finite discrete probability distribution over the elements in Θ. To ease readability, we write $\delta_{\mathsf{C}(\Theta)}$ for the probability distribution $\delta(\mathsf{C}(\Theta))$ over Θ.

Definition 10. *Let $T : i \xrightarrow{s} o$ be a well-typed term, π a string over i, ρ a string over o. Then, $[\![T, \delta]\!]^{\pi}_{\rho}$ stands for an arrow $2^{|i|} \to 2^{|o|}$ in $\mathcal{K}\ell(\mathcal{D})$ (i.e., a diagram from $|i|$ to $|o|$) defined by structural induction as follows:*

$$\begin{aligned}
[\![I_s, \delta]\!]^{\pi}_{\rho} &= \chi^{\pi}_{\rho} && (8)\\
[\![\perp_s, \delta]\!]^{\epsilon}_{\rho} &= \delta^{|s|}_{0} && (9)\\
[\![T_1 \oplus T_2, \delta]\!]^{\pi}_{\rho} &= \chi^{\pi}_{\pi_1\pi_2}; ([\![T_1, \delta]\!]^{\pi_1}_{\rho_1} \otimes [\![T_2, \delta]\!]^{\pi_2}_{\rho_2}); \chi^{\rho_1\rho_2}_{\rho} && (10)\\
[\![T_1; T_2, \delta]\!]^{\pi}_{\rho} &= [\![T_1, \delta]\!]^{\pi}_{\gamma}; [\![T_2, \delta]\!]^{\gamma}_{\rho} && (11)\\
[\![\mathsf{C}(\Theta), \delta]\!]^{\epsilon}_{\rho} &= \lambda m. \sum_{\theta : \emptyset \to m \in \Theta} \delta_{\mathsf{C}(\Theta)}(\theta) && (12)\\
[\![\sum_{m \subseteq i} m \triangleright T_m, \delta]\!]^{\pi}_{\rho} &= [[\![T_{\pi^{-1}(1)}, \delta]\!]^{\epsilon}_{\rho}, ..., [\![T_{\pi^{-1}(2^{|i|})}, \delta]\!]^{\epsilon}_{\rho}] && (13)
\end{aligned}$$

where in Eq. (9) the probability distribution $\delta^{|s|}_{0}$ assigns probabilty 1 to the case $\emptyset$ and 0 to all the remaining $2^{|s|} - 1$ cases and in Eq. (13) the arrows is obtained *as the copairing of each T_m for all $m \subseteq i$.*[4]

The cases in Eqs. (8) and (9) are straightforward. The cases in Eqs. (10) and (11) just exploit the monoidal category structure. It is worth noting that while the operation $\oplus$ is commutative, this is not the case for the monoidal operation of the Kleisli category, hence denoted with a different symbol $\otimes$. The case in Eq. (12) is the most interesting: $[\![\mathsf{C}(\Theta), \delta]\!]^{\epsilon}_{\rho}$ must assign a probability distribution to the elements in the powerset of the places in ρ; given $m \subseteq \rho$ its probability is computed by taking the sum of the probabilities assigned by δ to all processes θ whose final places are exactly m. This is correct as any two such processes are mutually exclusive alternatives. Finally, the case in Eq. (13) is the most complex, as it exploits the hierarchical decomposition of s-cells. Here we take each T_m and compute $2^{|i|}$ arrows $[\![T_m, \delta]\!]^{\epsilon}_{\rho} : 2^0 \to 2^{|\rho|}$. Then, via co-pairing we get an arrow from $2^{|i|}$ to $2^{|\rho|}$. The order of the arrows in the co-pair expression is important to associate them to the right element $m \subseteq i$ (according to the order induced by π).

From the encoding it is maybe not evident that the image of the mapping are string diagrams and not arbitrary arrows in $\mathcal{K}\ell(\mathcal{D})$. However it can be proved inductively that the encoding produces some sort of acyclic graph, in the style of [13]. This is immediately evident for Eq. 8 (symmetries) and Eqs. 9 and 12 (single node diagrams). For Eqs. 10 and 11 we use a simple inductive argument. The most complicated case is that of Eq. 13, which however also leads to the definition of a single node diagram whose probability matrix is obtained by collecting the rows associated with the probability distributions of each hierarchical decomposition.

Proposition 2. $[\![T, \delta]\!]^{\pi}_{\rho} = \chi^{\pi}_{\pi'}; [\![T, \delta]\!]^{\pi'}_{\rho'}; \chi^{\rho'}_{\rho}$.

Proof. The proof is by structural induction on T.

[4] It is important to mention that in Eq. (13) the order of the arrows in the copairing is the one induced by π: remember that π induces an order on 2^i, then $\pi^{-1}(k)$ denotes the k-th subset $m \subseteq i$ according to the order in π.

For the case $T = \bot_s$, we have $\chi^{\epsilon}_{\epsilon}; [\![\bot_s, \delta]\!]^{\epsilon}_{\rho'}; \chi^{\rho'}_{\rho} = [\![\bot_s, \delta]\!]^{\epsilon}_{\rho'}; \chi^{\rho'}_{\rho} = \delta^{|s|}_0; \chi^{\rho'}_{\rho} = \delta^{|s|}_0$.

For the case $T = I_s$, we have $\chi^{\pi}_{\pi'}; [\![I_s, \delta]\!]^{\pi'}_{\rho'}; \chi^{\rho'}_{\rho} = \chi^{\pi}_{\pi'}; \chi^{\pi'}_{\rho'}; \chi^{\rho'}_{\rho} = \chi^{\pi}_{\rho}$ by coherence of symmetries.

For the case $T = T_1 \oplus T_2$, we have

$$\begin{array}{rcl} \chi^{\pi}_{\pi'}; [\![T_1 \oplus T_2, \delta]\!]^{\pi'}_{\rho'}; \chi^{\rho'}_{\rho} & = & \chi^{\pi}_{\pi'}; \chi^{\pi'}_{\pi_1\pi_2}; ([\![T_1, \delta]\!]^{\pi_1}_{\rho_1} \otimes [\![T_2, \delta]\!]^{\pi_2}_{\rho_2}); \chi^{\rho_1\rho_2}_{\rho'}; \chi^{\rho'}_{\rho} \\ & = & \chi^{\pi}_{\pi_1\pi_2}; ([\![T_1, \delta]\!]^{\pi_1}_{\rho_1} \otimes [\![T_2, \delta]\!]^{\pi_2}_{\rho_2}); \chi^{\rho_1\rho_2}_{\rho} \\ & = & [\![T_1 \oplus T_2, \delta]\!]^{\pi}_{\rho} \end{array}$$

by coherence of symmetries.

For the case $T = T_1; T_2$, let us assume that $[\![T_1, \delta]\!]^{\pi_1}_{\rho_1} = \chi^{\pi_1}_{\pi'_1}; [\![T_1, \delta]\!]^{\pi'_1}_{\rho'_1}; \chi^{\rho'_1}_{\rho_1}$ and $[\![T_2, \delta]\!]^{\pi_2}_{\rho_2} = \chi^{\pi_2}_{\pi'_2}; [\![T_2, \delta]\!]^{\pi'_2}_{\rho'_2}; \chi^{\rho'_2}_{\rho_2}$, so that, as a particular case we have $[\![T_1, \delta]\!]^{\pi}_{\gamma} = \chi^{\pi}_{\pi'}; [\![T_1, \delta]\!]^{\pi'}_{\gamma}$ and $[\![T_2, \delta]\!]^{\gamma}_{\rho} = [\![T_2, \delta]\!]^{\gamma}_{\rho'}; \chi^{\rho'}_{\rho}$ (because $\chi^{\gamma}_{\gamma} = I_{|\gamma|}$). Then we have

$$\begin{array}{rcl} \chi^{\pi}_{\pi'}; [\![T_1; T_2, \delta]\!]^{\pi'}_{\rho'}; \chi^{\rho'}_{\rho} & = & \chi^{\pi}_{\pi'}; [\![T_1, \delta]\!]^{\pi'}_{\gamma}; [\![T_2, \delta]\!]^{\gamma}_{\rho'}; \chi^{\rho'}_{\rho} \\ & = & [\![T_1, \delta]\!]^{\pi}_{\gamma}; [\![T_2, \delta]\!]^{\gamma}_{\rho} \\ & = & [\![T_1; T_2, \delta]\!]^{\pi}_{\rho} \end{array}$$

For the case $T = \mathsf{C}(\Theta)$, likewise the case for $\bot_s$, the definition is purely functional.

For the case $T = \sum_{m \subseteq i} m \triangleright T_m$, let us assume that for any $m \subseteq i$ we have $[\![T_m, \delta]\!]^{\epsilon}_{\rho} = \chi^{\epsilon}_{\epsilon}; [\![T_m, \delta]\!]^{\epsilon}_{\rho'}; \chi^{\rho'}_{\rho} = [\![T_m, \delta]\!]^{\epsilon}_{\rho'}; \chi^{\rho'}_{\rho}$. Then, we have

$$\begin{array}{rcl} \chi^{\pi}_{\pi'}; [\![\sum\limits_{m \subseteq i} m \triangleright T_m, \delta]\!]^{\pi'}_{\rho'}; \chi^{\rho'}_{\rho} & = & \chi^{\pi}_{\pi'}; [[\![T_{\pi'^{-1}(1)}, \delta]\!]^{\epsilon}_{\rho'}, ..., [\![T_{\pi'^{-1}(2^{|i|})}, \delta]\!]^{\epsilon}_{\rho'}]; \chi^{\rho'}_{\rho} \\ & = & \chi^{\pi}_{\pi'}; [[\![T_{\pi'^{-1}(1)}, \delta]\!]^{\epsilon}_{\rho'}; \chi^{\rho'}_{\rho}, ..., [\![T_{\pi'^{-1}(2^{|i|})}, \delta]\!]^{\epsilon}_{\rho'}; \chi^{\rho'}_{\rho}] \\ & = & \chi^{\pi}_{\pi'}; [[\![T_{\pi'^{-1}(1)}, \delta]\!]^{\epsilon}_{\rho}, ..., [\![T_{\pi'^{-1}(2^{|i|})}, \delta]\!]^{\epsilon}_{\rho}] \\ & = & [[\![T_{\pi^{-1}(1)}, \delta]\!]^{\epsilon}_{\rho}, ..., [\![T_{\pi^{-1}(2^{|i|})}, \delta]\!]^{\epsilon}_{\rho}] \\ & = & [\![\sum\limits_{m \subseteq i} m \triangleright T_m, \delta]\!]^{\pi}_{\rho} \end{array}$$

□

Proposition 3. *The definition of* $[\![T, \delta]\!]^{\pi}_{\rho}$ *is well given.*

Proof. We must show that: (1) the typing is consistent with the definition, (2) that the choice of $\pi_1, \rho_1, \pi_2, \rho_2$ in Eq. (10) and of γ in Eq. (11) is inessential for the result, and (3) that $\llbracket T_1 \oplus T_2, \delta \rrbracket^{\pi}_{\rho} = \llbracket T_2 \oplus T_1, \delta \rrbracket^{\pi}_{\rho}$.

For (1), we must prove that if $T : i \xrightarrow{s} o$, π is a string over i and ρ is a string over o, then $\llbracket T, \delta \rrbracket^{\pi}_{\rho} : 2^{|i|} \to 2^{|o|}$. The proof is a straightforward rule induction.

For (2), we just exploit Proposition 2. In the case of Eq. (10), we have

$$\begin{array}{rcl}
\llbracket T_1 \oplus T_2, \delta \rrbracket^{\pi}_{\rho} & = & \chi^{\pi}_{\pi_1\pi_2}; (\llbracket T_1, \delta \rrbracket^{\pi_1}_{\rho_1} \otimes \llbracket T_2, \delta \rrbracket^{\pi_2}_{\rho_2}); \chi^{\rho_1\rho_2}_{\rho} \\
& = & \chi^{\pi}_{\pi_1\pi_2}; ((\chi^{\pi_1}_{\pi'_1}; \llbracket T_1, \delta \rrbracket^{\pi'_1}_{\rho'_1}; \chi^{\rho'_1}_{\rho_1}) \otimes (\chi^{\pi_2}_{\pi'_2}; \llbracket T_2, \delta \rrbracket^{\pi'_2}_{\rho'_2}; \chi^{\rho'_2}_{\rho_2})); \chi^{\rho_1\rho_2}_{\rho} \\
& = & \chi^{\pi}_{\pi_1\pi_2}; (\chi^{\pi_1}_{\pi'_1} \otimes \chi^{\pi_2}_{\pi'_2}); (\llbracket T_1, \delta \rrbracket^{\pi'_1}_{\rho'_1} \otimes \llbracket T_2, \delta \rrbracket^{\pi'_2}_{\rho'_2}); (\chi^{\rho'_1}_{\rho_1} \otimes \chi^{\rho'_2}_{\rho_2}); \chi^{\rho_1\rho_2}_{\rho} \\
& = & \chi^{\pi}_{\pi_1\pi_2}; \chi^{\pi_1\pi_2}_{\pi'_1\pi'_2}; (\llbracket T_1, \delta \rrbracket^{\pi'_1}_{\rho'_1} \otimes \llbracket T_2, \delta \rrbracket^{\pi'_2}_{\rho'_2}); \chi^{\rho'_1\rho'_2}_{\rho_1\rho_2}; \chi^{\rho_1\rho_2}_{\rho} \\
& = & \chi^{\pi}_{\pi'_1\pi'_2}; (\llbracket T_1, \delta \rrbracket^{\pi'_1}_{\rho'_1} \otimes \llbracket T_2, \delta \rrbracket^{\pi'_2}_{\rho'_2}); \chi^{\rho'_1\rho'_2}_{\rho}
\end{array}$$

In the case of Eq. (11), we have

$$\begin{array}{rcl}
\llbracket T_1; T_2, \delta \rrbracket^{\pi}_{\rho} & = & \llbracket T_1, \delta \rrbracket^{\pi}_{\gamma}; \llbracket T_2, \delta \rrbracket^{\gamma}_{\rho} \\
& = & \llbracket T_1, \delta \rrbracket^{\pi}_{\gamma'}; \chi^{\gamma'}_{\gamma}; \chi^{\gamma}_{\gamma'}; \llbracket T_2, \delta \rrbracket^{\gamma'}_{\rho} \\
& = & \llbracket T_1, \delta \rrbracket^{\pi}_{\gamma'}; \llbracket T_2, \delta \rrbracket^{\gamma'}_{\rho}
\end{array}$$

Finally, for (3), we have:

$$\begin{array}{rcl}
\llbracket T_1 \oplus T_2, \delta \rrbracket^{\pi}_{\rho} & = & \chi^{\pi}_{\pi_1\pi_2}; (\llbracket T_1, \delta \rrbracket^{\pi_1}_{\rho_1} \otimes \llbracket T_2, \delta \rrbracket^{\pi_2}_{\rho_2}); \chi^{\rho_1\rho_2}_{\rho} \\
& = & \chi^{\pi}_{\pi_1\pi_2}; \chi^{\pi_1\pi_2}_{\pi_2\pi_1}; (\llbracket T_2, \delta \rrbracket^{\pi_2}_{\rho_2} \otimes \llbracket T_1, \delta \rrbracket^{\pi_1}_{\rho_1}); \chi^{\rho_2\rho_1}_{\rho_1\rho_2}; \chi^{\rho_1\rho_2}_{\rho} \\
& = & \chi^{\pi}_{\pi_2\pi_1}; (\llbracket T_2, \delta \rrbracket^{\pi_2}_{\rho_2} \otimes \llbracket T_1, \delta \rrbracket^{\pi_1}_{\rho_1}); \chi^{\rho_2\rho_1}_{\rho} \\
& = & \llbracket T_2 \oplus T_1, \delta \rrbracket^{\pi}_{\rho}
\end{array}$$

□

Example 7. *Consider the net depicted in Fig. 3(a) and the corresponding term calculated in Example 6. We show the encoding of the net by considering a generic distribution δ and use lexicographic order of places. We start from Eq. 2.*

$$\llparenthesis N, \{2,3\} \rrparenthesis = \llparenthesis N_{\mathbb{C}_1} \rrparenthesis \oplus \llparenthesis N_{\mathbb{C}_2}, \{2\} \rrparenthesis; \llparenthesis N_{\mathbb{C}_3}, \{3\} \rrparenthesis \oplus \llparenthesis I_{\{5\}} \rrparenthesis$$

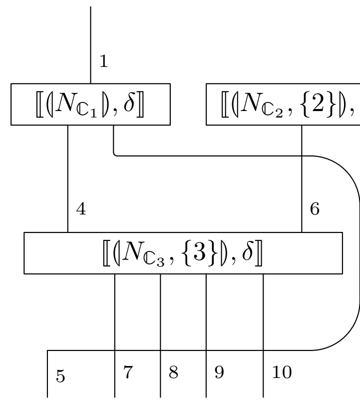

Figure 5: String diagram for $[\![(\!|N, \{2,3\}|\!), \delta]\!]$

Then, the string diagram for $[\![(\!|N, \{2,3\}|\!), \delta]\!]^{1}_{5,7,8,9,10}$ *is shown in Fig. 5 and can be computed as follows.*

$$
\begin{array}{lll}
& [\![(\!|N, \{2,3\}|\!), \delta]\!]^{1}_{5,7,8,9,10} & \\
= & [\![(\!|N_{\mathbb{C}_1}|\!) \oplus (\!|N_{\mathbb{C}_2}, \{2\}|\!); (\!|N_{\mathbb{C}_3}, \{3\}|\!) \oplus (\!|I_{\{5\}}|\!), \delta]\!]^{1}_{5,7,8,9,10} & \textit{by def.} \\
= & [\![(\!|N_{\mathbb{C}_1}|\!) \oplus (\!|N_{\mathbb{C}_2}, \{2\}|\!), \delta]\!]^{1}_{4,5,6}; [\![(\!|N_{\mathbb{C}_3}, \{3\}|\!) \oplus (\!|I_{\{5\}}|\!), \delta]\!]^{4,5,6}_{5,7,8,9,10} & \textit{by } (11) \\
= & \chi^{1}_{1\epsilon}; [\![(\!|N_{\mathbb{C}_1}|\!), \delta]\!]^{1}_{4,5} \otimes [\![(\!|N_{\mathbb{C}_2}, \{2\}|\!), \delta]\!]^{\epsilon}_{6}; \chi^{4,5,6}_{4,5,6}; & \textit{by } (10) \\
& \chi^{4,5,6}_{4,6,5}; [\![(\!|N_{\mathbb{C}_3}, \{3\}|\!), \delta]\!]^{4,6}_{7,8,9,10} \otimes [\![(\!|I_{\{5\}}|\!), \delta]\!]^{5}_{5}; \chi^{7,8,9,10,5}_{5,7,8,9,10} &
\end{array}
$$

We now show the calculation for each of the boxes in Fig. 5. To ease readability, in the following we let

$$
\begin{array}{lllllll}
\mathsf{C}_a & = & \mathsf{C}(\{\{a\},\{b\}\}) & \quad & \mathsf{C}_c & = & \mathsf{C}(\{\{c\},\{d\}\}) \\
\mathsf{C}_e & = & \mathsf{C}(\{\{e\}\}) & & \mathsf{C}_g & = & \mathsf{C}(\{\{g\},\{h\}\})) \\
\mathsf{C}_f & = & \mathsf{C}(\{\{f\},\{e,g\},\{e,h\}\}) & & & &
\end{array}
$$

For $[\![(\!|N_{\mathbb{C}_1}|\!), \delta]\!]^{1}_{4,5}$*, we start from Eq.* (4)*, i.e.,*

$$(\!|N_{\mathbb{C}_1}|\!) = \emptyset \triangleright \bot_{\{4,5\}} + \{1\} \triangleright \mathsf{C}_a$$

By Eq. (13),

$$[\![(\!|N_{\mathbb{C}_1}|\!), \delta]\!]^1_{4,5} = \begin{bmatrix} [\![\perp_{\{4,5\}}, \delta]\!]^\epsilon_{4,5} \\ [\![\mathsf{C}_a, \delta]\!]^\epsilon_{4,5} \end{bmatrix} = \begin{array}{|c|c|c|c|c|} \hline & \emptyset & \{4\} & \{5\} & \{4,5\} \\ \hline \emptyset & 1 & 0 & 0 & 0 \\ \hline \{1\} & 0 & p_a & 1-p_a & 0 \\ \hline \end{array} \tag{14}$$

where the first row in the table corresponds to $\delta_0^{|\{4,5\}|}$*, as prescribed by Eq.* (9)*. The second row is obtained by Eq.* (12)*, by assuming that* $\delta_{\mathsf{C}_a}(\{a\}) = p_a$ *and* $\delta_{\mathsf{C}_a}(\{b\}) = 1 - p_a$*.*

For $[\![(\!|N_{\mathbb{C}_2}, \{2\}|\!), \delta]\!]^\epsilon_6$*, we start from Eq.* (5)*, i.e.,*

$$(\!|N_{\mathbb{C}_2}, \{2\}|\!) = \mathsf{C}_c$$

Then,

$$[\![(\!|N_{\mathbb{C}_2}, \{2\}|\!), \delta]\!]^\epsilon_6 = \quad [\![\mathsf{C}_c, \delta]\!]^\epsilon_6 \quad = \begin{array}{|c|c|c|} \hline & \emptyset & \{6\} \\ \hline \emptyset & 1-p_c & p_c \\ \hline \end{array} \tag{15}$$

where $\delta_{\mathsf{C}_c}(\{c\}) = p_c$ *and* $\delta_{\mathsf{C}_c}(\{d\}) = 1 - p_c$*.*

For $[\![(\!|N_{\mathbb{C}_3}, \{3\}|\!), \delta]\!]^{4,6}_{7,8,9,10}$*, we start from Eq.* (6)*, i.e.,*

$$\begin{aligned} (\!|N_{\mathbb{C}_3}, \{3\}|\!) = \quad & \emptyset \triangleright (\perp_{\{8,9,10\}} \oplus \mathsf{C}_e) \\ & + \{4\} \triangleright (\perp_{\{8,9,10\}} \oplus \mathsf{C}_e) \\ & + \{6\} \triangleright (\perp_{\{8\}} \oplus \mathsf{C}_e \oplus \mathsf{C}_g) \\ & + \{4,6\} \triangleright \mathsf{C}_f \end{aligned}$$

$$[\![(\!|N_{\mathbb{C}_3}, \{3\}|\!), \delta]\!]^\epsilon_{7,8,9,10} = \begin{bmatrix} [\![\perp_{\{8,9,10\}} \oplus \mathsf{C}_e, \delta]\!]^\epsilon_{7,8,9,10} \\ [\![\perp_{\{8,9,10\}} \oplus \mathsf{C}_e, \delta]\!]^\epsilon_{7,8,9,10} \\ [\![\perp_{\{8\}} \oplus \mathsf{C}_e \oplus \mathsf{C}_g, \delta]\!]^\epsilon_{7,8,9,10} \\ [\![\mathsf{C}_f, \delta]\!]^\epsilon_{7,8,9,10} \end{bmatrix}$$

$$= \begin{array}{|c|c|c|c|c|c|c|} \hline & \emptyset & \{7\} & \{7,9\} & \{7,10\} & \{8\} & \dots \\ \hline \emptyset & 0 & 1 & 0 & 0 & 0 & 0 \\ \hline \{4\} & 0 & 1 & 0 & 0 & 0 & 0 \\ \hline \{6\} & 0 & 0 & p_g & 1-p_g & 0 & 0 \\ \hline \{4,6\} & 0 & 0 & p'_g & 1-p_f-p'_g & p_f & 0 \\ \hline \end{array} \tag{16}$$

where the last column (i.e., the one tagged with dots) represents all the remaining nine (inessential) cases. The first two rows are obtained as follows:

$$\llbracket \bot_{\{8,9,10\}} \oplus \mathsf{C}_e, \delta \rrbracket^{\epsilon}_{7,8,9,10} = \llbracket \bot_{\{8,9,10\}}, \delta \rrbracket^{\epsilon}_{8,9,10} \otimes \llbracket \mathsf{C}_e, \delta \rrbracket^{\epsilon}_{7} ; \chi^{8,9,10,7}_{7,8,9,10}$$

$$= \begin{array}{|c|c|c|}\hline & \emptyset & \dots \\ \hline \emptyset & 1 & 0 \\ \hline \end{array} \otimes \begin{array}{|c|c|c|}\hline & \emptyset & \{7\} \\ \hline \emptyset & 0 & 1 \\ \hline \end{array}$$

$$= \begin{array}{|c|c|c|c|}\hline & \emptyset & \{7\} & \dots \\ \hline \emptyset & 0 & 1 & 0 \\ \hline \end{array}$$

The third row is obtained analogously after fixing $\delta_{\mathsf{C}_g}(\{g\}) = p_g$ and $\delta_{\mathsf{C}_g}(\{h\}) = 1 - p_g$. The last row is obtained by Eq. (13) *and taking $\delta_{\mathsf{C}_f}(\{f\}) = p_f$, $\delta_{\mathsf{C}_f}(\{e, g\}) = p'_g$, and $\delta_{\mathsf{C}_f}(\{e, h\}) = 1 - p_f - p'_g$.*

5. Forward and Backward Inference and Disintegration

In this section we illustrate how to perform Bayesian reasoning over Petri nets by following the approach presented in [3, 4].

One of the advantage of the Bayesian net representation is that we can study conditional probability distributions about events happening at the level of the Petri net. For example, we can estimate the probability that a place is eventually marked under different scenarios about the past (e.g. under different hypothesis about other places being marked). Similarly, backward reasoning can be exploited to study the probability that something happened in the past given that some event is observed. For example, if a net represents the behaviour of a possibly faulty system and it is observed a token in a place that represents a malfunctioning of the system, we can estimate the probability that different causes of the malfunctioning happened in the past. In the Baysian nets, the random variables are associated with the initial/final places of s-cells. However, by the structure of the nets we are considering, since each place has at most one incoming arc and each token a unique history, we can transfer the probability distribution to the firing of transitions.

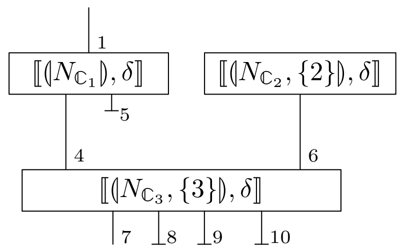

Figure 6: Simplified string diagram for $\llbracket \llparenthesis N, \{2,3\} \rrparenthesis, \delta \rrbracket$

We first recall some notions, which will be used in our reasoning. *Marginalisation* is an operation $\Pi_1 : X \oplus Y \to X$ that projects a joint distribution $P(x,y)$ on $X \oplus Y$ to the marginal distribution on X computed as $P(x) = \sum_y P(x,y)$. Similarly, we have $\Pi_2 : X \oplus Y \to Y$ for the projection of $P(x,y)$ over Y defined as $P(y) = \sum_y P(x,y)$.

Consider the arrow $\llparenthesis N, \{2,3\} \rrparenthesis : 2^1 \to 2^5$ in Fig. 5 and suppose we are interested in reasoning about the probability of producing a token in the place 7. In such case, marginalisation can be used to obtain an arrow $f : 2^1 \to 2^1$ that discards the wires corresponding to the places 5, 8, 9 and 10, as shown in Fig. 6. The wire diagram corresponds to the term:

$$(\llbracket \llparenthesis N_{\mathbb{C}_1} \rrparenthesis, \delta \rrbracket^1_{4,5}; \Pi_1) \otimes \llbracket \llparenthesis N_{\mathbb{C}_2}, \{2\} \rrparenthesis, \delta \rrbracket^{\epsilon}_{6}; (\llbracket \llparenthesis N_{\mathbb{C}_3}, \{3\} \rrparenthesis, \delta \rrbracket^{4,6}_{7,8,9,10}; \Pi_1 \otimes \Pi_1; \Pi_1)$$

From Eq. (14), we obtain

$$\alpha = \llbracket \llparenthesis N_{\mathbb{C}_1} \rrparenthesis, \delta \rrbracket^1_{4,5}; \Pi_1 = \begin{array}{|c|c|c|} \hline & \emptyset & \{4\} \\ \hline \emptyset & 1 & 0 \\ \hline \{1\} & 1 - p_a & p_a \\ \hline \end{array} \tag{17}$$

Analogously, from Eq. (16)

$$\gamma = [\![(\!| N_{\mathbb{C}_3}, \{3\} |\!), \delta]\!]^{\epsilon}_{7,8,9,10}; \Pi_1 \otimes \Pi_1; \Pi_1 = \begin{array}{|c|c|c|} \hline & \emptyset & \{7\} \\ \hline \emptyset & 0 & 1 \\ \hline \{4\} & 0 & 1 \\ \hline \{6\} & 0 & 1 \\ \hline \{4,6\} & p_f & 1 - p_f \\ \hline \end{array} \tag{18}$$

We write β for $[\![(\!| N_{\mathbb{C}_2}, \{2\} |\!), \delta]\!]^{\epsilon}_{6}$ in Eq. 15.

Then, $\alpha \otimes \beta$ is obtained as

$$\alpha \otimes \beta = \begin{array}{|c|c|c|c|c|} \hline & \emptyset & \{4\} & \{6\} & \{4,6\} \\ \hline \emptyset & 1 - pc & 0 & p_c & 0 \\ \hline \{1\} & (1-p_a)(1-p_c) & p_a(1-p_c) & (1-p_a)p_c & p_a p_c \\ \hline \end{array} \tag{19}$$

Finally,

$$\psi = \alpha \otimes \beta; \gamma = \begin{array}{|c|c|c|} \hline & \emptyset & \{7\} \\ \hline \emptyset & 0 & 1 \\ \hline \{1\} & p_a p_c p_f & 1 - p_a p_c p_f \\ \hline \end{array} \tag{20}$$

This means that, given that a token appears in place 1 with probability 1, the place 7 will be marked with probability $1 - p_a p_c p_f$. This is also the probability of firing the transition e, which is the only one producing the token in place 7. Using the notation in [3], this value is computed by precomposing the state $\omega = 1|\{1\}\rangle$ with the arrow ψ, i.e., by letting $\psi_*(\omega) = \omega; \psi = p_a p_c p_f|\emptyset\rangle + (1 - p_a p_c p_f)|\{7\}\rangle$.

As an example of backward reasoning, given the *a priori* probability $\frac{1}{2}$ that a token can appear in place 1, we can compute the probability that place 1 is marked given that a token appears in place 7, which is

$$\frac{1 - p_a p_c p_f}{1 + (1 - p_a p_c p_f)} = \frac{1 - p_a p_c p_f}{2 - p_a p_c p_f}$$

Using the notation in [3], this value is computed by setting (for $\psi : X \to \mathcal{D}(Y)$ and q a predicate on Y)

$$\begin{array}{rcl} \psi^*(q)(x) & = & \sum_{y \in Y} \psi(x)(y) \cdot q(y) \\ & = & \psi(x)(\emptyset) \cdot q(\emptyset) + \psi(x)(\{7\}) \cdot q(\{7\}) \\ & = & \psi(x)(\{7\}) \end{array}$$

where q is the predicate such that $q(\{7\}) = 1$ (and $q(\emptyset) = 0$) and then computing

$$\begin{aligned}
\omega_{|\psi^*(q)} &= \sum_{x \in X} \frac{\omega(x) \cdot \psi^*(q)(x)}{\omega \models \psi^*(q)} |x\rangle \\
&= \frac{\omega(\emptyset) \cdot \psi^*(q)(\emptyset)}{\omega \models \psi^*(q)} |\emptyset\rangle + \frac{\omega(\{1\}) \cdot \psi^*(q)(\{1\})}{\omega \models \psi^*(q)} |\{1\}\rangle \\
&= \frac{\frac{1}{2} \cdot 1}{\omega \models \psi^*(q)} |\emptyset\rangle + \frac{\frac{1}{2} \cdot (1 - p_a p_c p_f)}{\omega \models \psi^*(q)} |\{1\}\rangle \\
&= \frac{\frac{1}{2}}{\omega \models \psi^*(q)} |\emptyset\rangle + \frac{\frac{1 - p_a p_c p_f}{2}}{\omega \models \psi^*(q)} |\{1\}\rangle
\end{aligned}$$

where

$$\begin{aligned}
\omega \models \psi^*(q) &= \sum_{x \in X} \omega(x) \cdot \psi^*(q)(x) \\
&= \omega(\emptyset) \cdot \psi^*(q)(\emptyset) + \omega(\{1\}) \cdot \psi^*(q)(\{1\}) \\
&= \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (1 - p_a p_c p_f) \\
&= \frac{2 - p_a p_c p_f}{2}
\end{aligned}$$

If the presence of a token in place 7 represents a malfunctioning of the system, and the presence of a token in place 1 a possible fault, then the probability we compute is that the malfunctioning is dependent on that fault.

6. Conclusion

In this paper we have shown how to derive a Bayesian network from a probabilistic Petri net in the style of [1, 2]. The construction is computed via an intermediate representation of a PN as a term in a monoidal (pre-)category structure, exploiting the string diagram representation of BN outlined in [3]. As shown in Section 5, the BN representation can then be exploited to reason about conditional probabilities of marking reachability, via forward and backward inference. Notably, when transitions have non-empty post-sets then each marking corresponds to a unique deterministic process (i.e., a unique configuration of the underlying event structure) and thus the inference can be transferred to processes as well.

There are many ways in which PN have been enriched with probabilistic behaviour [14, 15, 16, 17, 18, 19, 20, 21]. To avoid confusion, most of them replace nondeterminism with probability only in part, or focus on interleaved computations, or introduce time dependent stochastic distributions. The approach considered here differs from the others in the literature because: (1) it is purely probabilistic, (2) it deals well with concurrent computations, (3) it addresses confusion.

In the literature, there are very few papers investigating the connections between PN and BN. In [22] the relation is drawn in the opposite direction, i.e., PN are used to encode the reasoning of BN. The connection established in this paper provides two views for the same model: on the one side, the standard token game of the PN view (suitable extended with probabilistic choices) gives a concrete, probabilistic computational model. On the other side, the BN semantics allows us to reason about the properties of the computations of the underlying concrete model.

The relation between Petri nets and Bayesian networks opens the way to several interesting directions for future work. One is about causality, in the sense of Pearl [23]. There the idea is to distinguish between statistical correlation and cause-effect relationship. Causality plays an important role in concurrent semantics and is explicitly represented by the Petri net structure. Using the Bayesian network semantics presented here, we would like to explore the connections between causality in Petri nets and the approach in [13] that shapes Pearl's ideas on a string diagram perspective.

Another direction for future work is in exploring the notion of influence. In Bayesian networks, influence is a relationship between nodes A and B, telling how an observation on node A may have an effect also on node B. In particular, such influence can be quantified by defining suitable metrics as done in [24]. Given the Bayesian network semantics of a Petri net, we can try to recast the notion of influence between markings of the net and study how it is related with causal dependencies arising from the structure of the net.

References

[1] S. Abbes, A. Benveniste, True-concurrency probabilistic models: Branching cells and distributed probabilities for event structures, Inf. Comput. 204 (2) (2006) 231–274. doi:10.1016/j.ic.2005.10.001.

[2] R. Bruni, H. C. Melgratti, U. Montanari, Concurrency and probability: Removing confusion, compositionally, in: A. Dawar, E. Grädel (Eds.), Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018, ACM, 2018, pp. 195–204. doi:10.1145/3209108.3209202.

[3] B. Jacobs, F. Zanasi, A predicate/state transformer semantics for Bayesian learning, Electr. Notes Theor. Comput. Sci. 325 (2016) 185–200. doi:10.1016/j.entcs.2016.09.038.

[4] K. Cho, B. Jacobs, Disintegration and Bayesian inversion, both abstractly and concretely, CoRR abs/1709.00322. arXiv:1709.00322.
URL http://arxiv.org/abs/1709.00322

[5] F. Clerc, V. Danos, F. Dahlqvist, I. Garnier, Pointless learning, in: J. Esparza, A. S. Murawski (Eds.), Foundations of Software Science and Computation Structures - 20th International Conference, FOSSACS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Vol. 10203 of Lecture Notes in Computer Science, 2017, pp. 355–369. doi:10.1007/978-3-662-54458-7_21.

[6] P. Baldan, R. Bruni, A. Corradini, F. Gadducci, H. C. Melgratti, U. Montanari, Event structures for petri nets with persistence, Logical Methods in Computer Science 14 (3). doi:10.23638/LMCS-14(3:25)2018.

[7] U. Goltz, W. Reisig, The non-sequential behavior of Petri nets, Information and Control 57 (2/3) (1983) 125–147. doi:10.1016/S0019-9958(83)80040-0.

[8] P. Baldan, A. Corradini, H. Ehrig, R. Heckel, B. König, Bisimilarity and behaviour-preserving reconfigurations of open petri nets, Logical Methods in Computer Science 4 (4). doi:10.2168/LMCS-4(4:3)2008.

[9] R. Bruni, H. C. Melgratti, U. Montanari, P. Sobocinski, Connector algebras for C/E and P/T nets' interactions, Logical Methods in Computer Science 9 (3). doi:10.2168/LMCS-9(3:16)2013.

[10] F. Bonchi, J. Holland, R. Piedeleu, P. Sobocinski, F. Zanasi, Diagrammatic algebra: from linear to concurrent systems, PACMPL 3 (POPL) (2019) 25:1–25:28.

[11] P. Baldan, F. Gadducci, Petri nets are dioids: a new algebraic foundation for non-deterministic net theory, Acta Inf. 56 (1) (2019) 61–92. doi:10.1007/s00236-018-0314-0.

[12] R. Bruni, J. Meseguer, U. Montanari, V. Sassone, Functorial models for petri nets, Inf. Comput. 170 (2) (2001) 207–236. doi:10.1006/inco.2001.3050.

[13] B. Jacobs, A. Kissinger, F. Zanasi, Causal inference by string diagram surgery, in: M. Boja'nczyk, A. Simpson (Eds.), Foundations of Software Science and Computation Structures - 22nd International Conference, FOSSACS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Vol. 11425 of Lecture Notes in Computer Science, Springer, 2019, pp. 313–329. doi:10.1007/978-3-030-17127-8_18.

[14] J. B. Dugan, K. S. Trivedi, R. Geist, V. F. Nicola, Extended stochastic Petri nets: Applications and analysis, in: Performance'84, North-Holland, 1984, pp. 507–519.

[15] M. A. Marsan, G. Conte, G. Balbo, A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems, ACM Trans. Comput. Syst. 2 (2) (1984) 93–122. doi:10.1145/190.191.

[16] M. K. Molloy, Discrete time stochastic Petri nets, IEEE Trans. Softw. Eng. 11 (4) (1985) 417–423.

[17] C. Eisentraut, H. Hermanns, J. Katoen, L. Zhang, A semantics for every GSPN, in: Petri Nets 2013, Vol. 7927 of Lect. Notes in Comp. Sci., Springer, 2013, pp. 90–109. doi:10.1007/978-3-642-38697-8_6.

[18] M. Kudlek, Probability in petri nets, Fundam. Inform. 67 (1-3) (2005) 121–130.

[19] S. Haar, Probabilistic cluster unfoldings, Fundam. Inform. 53 (3-4) (2002) 281–314.

[20] A. Bouillard, S. Haar, S. Rosario, Critical paths in the partial order unfolding of a stochastic petri net, in: J. Ouaknine, F. W. Vaandrager (Eds.), Formal Modeling and Analysis of Timed Systems, 7th International Conference, FORMATS 2009, Budapest, Hungary, September 14-16, 2009. Proceedings, Vol. 5813 of Lecture Notes in Computer Science, Springer, 2009, pp. 43–57. doi:10.1007/978-3-642-04368-0_6.

[21] J. Katoen, R. Langerak, D. Latella, Modeling systems by probabilistic process algebra: an event structures approach, in: R. L. Tenney, P. D. Amer, M. Ü. Uyar (Eds.), Formal Description Techniques, VI, Proceedings of the IFIP TC6/WG6.1 Sixth International Conference on Formal Description Techniques - FORTE '93, Boston, MA, USA, 26-29 October 1993, Vol. C-22 of IFIP Transactions, North-Holland, 1993, pp. 253–268.

[22] K. Lautenbach, S. Philippi, A. Pinl, Bayesian Networks and Petri Nets, Fachberichte Informatik 2–2006, Universität Koblenz-Landau (2006). URL http://www.uni-koblenz.de/fb4/publikationen/gelbereihe/RR-2-2006.pdf

[23] J. Pearl, Causality: Models, Reasoning and Inference, 2nd Edition, Cambridge University Press, New York, NY, USA, 2009.

[24] B. Jacobs, F. Zanasi, A formal semantics of influence in bayesian reasoning, in: K. G. Larsen, H. L. Bodlaender, J. Raskin (Eds.), 42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017, August 21-25, 2017 - Aalborg, Denmark, Vol. 83 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017, pp. 21:1–21:14. doi:10.4230/LIPIcs.MFCS.2017.21.

[25] M. Nielsen, G. D. Plotkin, G. Winskel, Petri nets, event structures and domains, part I, Theor. Comput. Sci. 13 (1981) 85–108. doi:10.1016/0304-3975(81)90112-2.

[26] G. Winskel, Event structures, in: Advances in Petri Nets 1986, Part II, Vol. 255 of Lect. Notes in Comp. Sci., Springer, 1987, pp. 325–392.

Appendix A. Correctness of mapping to terms

The remaining of this section is devoted to establish a correspondence between the semantics of Abbes and Benveniste for a marked net (N, m) and the corresponding term $(\!| N, m |\!)$.

Appendix A.1. Prime Event Structures

A *prime event structure* (also *PES*) [25, 26] is a triple $\mathcal{E} = (E, \preceq, \#)$ where: E is the set of *events*; the *causality relation* $\preceq$ is a partial order on events; the *conflict relation* $\#$ is a symmetric, irreflexive relation on events such that conflicts are inherited by causality, i.e., $\forall e_1, e_2, e_3 \in E.\ e_1 \# e_2 \preceq e_3 \Rightarrow e_1 \# e_3$.

The PES $\mathcal{E}_N$ associated with a net N can be formalised using category theory as a chain of universal constructions, called coreflections. Hence, for each PES $\mathcal{E}$, there is a standard, unique (up to isomorphism) nondeterministic occurrence net $N_{\mathcal{E}}$ that yields $\mathcal{E}$ and thus we can freely move from one setting to the other.

Given an event e, its *downward closure* $\lfloor e \rfloor = \{e' \in E \mid e' \preceq e\}$ is the set of causes of e. As usual, we assume that $\lfloor e \rfloor$ is finite for any e. Given $B \subseteq E$, we say that B is *downward closed* if $\forall e \in B.\ \lfloor e \rfloor \subseteq B$ and that B is *conflict-free* if

$\forall e, e' \in B.\ \neg(e\#e')$. We let the *immediate conflict* relation $\#_0$ be defined on events by letting $e\#_0e'$ iff $(\lfloor e\rfloor \times \lfloor e'\rfloor) \cap \# = \{(e,e')\}$, i.e., two events are in immediate conflict if they are in conflict but their causes are compatible.

Appendix A.2. Abbes and Benveniste's Branching Cells

In the following we assume that a finite PES $\mathcal{E} = (E, \preceq, \#)$ is given. A *prefix* $B \subseteq E$ is any downward-closed set of events (possibly with conflicts). Any prefix B induces an event structure $\mathcal{E}_B = (B, \preceq_B, \#_B)$ where $\preceq_B$ and $\#_B$ are the restrictions of $\preceq$ and $\#$ to the events in B. A *stopping prefix* is a prefix B that is closed under immediate conflicts, i.e., $\forall e \in B, e' \in E.\ e\#_0e' \Rightarrow e' \in B$. Intuitively, a stopping prefix is a prefix whose (immediate) choices are all available. It is *initial* if the only stopping prefix strictly included in B is $\emptyset$.

A *configuration* $v \subseteq \mathcal{E}$ is any set of events that is downward closed and conflict-free. Intuitively, a configuration represents (the state reached after executing) a concurrent but deterministic computation of $\mathcal{E}$. Configurations are ordered by inclusion and we denote by $\mathcal{V}_\mathcal{E}$ the poset of configurations of $\mathcal{E}$ and by $\Omega_\mathcal{E}$ the poset of maximal configurations of $\mathcal{E}$.

The *future* of a configuration v, written E^v, is the set of events that can be executed after v, i.e., $E^v = \{e \in E \setminus v \mid \forall e' \in v.\neg(e\#e')\}$. We write $\mathcal{E}^v$ for the event structure induced by E^v.

A configuration v is *stopped* if there is a stopping prefix B with $v \in \Omega_B$. and v is *recursively stopped* (or r-stopped) if there is a sequence of configurations $\emptyset = v_0 \subset \ldots \subset v_n = v$ such that for any $i \in [0,n)$ the set $v_{i+1} \setminus v_i$ is a stopped configuration of $\mathcal{E}^{v_i}$ for v_i in $\mathcal{E}$.

A *branching cell* is any initial stopping prefix of the future $\mathcal{E}^v$ of a recursively stopped configuration v. Intuitively, a branching cell is a minimal subset of events closed under immediate conflict. We remark that branching cells are determined by considering the whole (future of the) event structure $\mathcal{E}$ and they are recursively computed as $\mathcal{E}$ is executed. Remarkably, every maximal configuration has a branching cell decomposition.

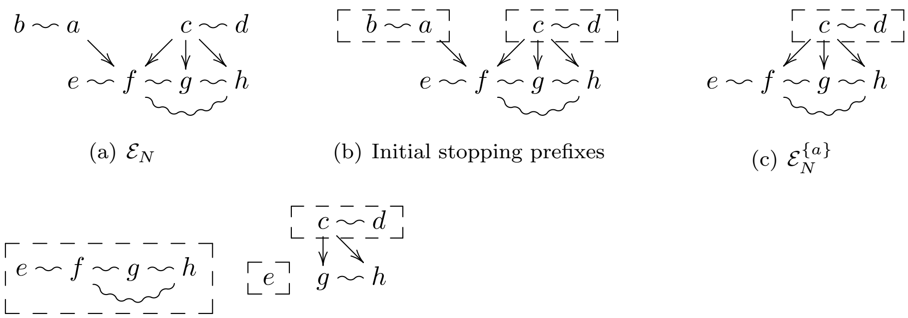

(a) $\mathcal{E}_N$ (b) Initial stopping prefixes (c) $\mathcal{E}_N^{\{a\}}$

(d) $\mathcal{E}_N^{\{a,c\}}$ (e) $\mathcal{E}_N^{\{b\}}$

Figure A.7: AB's branching cell decomposition (running example)

Example 8. *Consider the PES $\mathcal{E}_N$ in Fig. A.7(a) and its maximal configuration $v = \{a, c, e, g\}$. We show that v is recursively stopped by exhibiting a branching cell decomposition. The initial stopping prefixes of $\mathcal{E}_N = \mathcal{E}_N^{\emptyset}$ are shown in Fig. A.7(b). There are two possibilities for choosing $v_1 \subseteq v$ and v_1 recursively stopped: either $v_1 = \{a\}$ or $v_1 = \{c\}$. When $v_1 = \{a\}$, the choices for v_2 are determined by the stopping prefixes of $\mathcal{E}_N^{\{a\}}$ (see Fig. A.7(c)) and the only possibility is $v_2 = \{a, c\}$. From $\mathcal{E}_N^{\{a,c\}}$ in Fig. A.7(d), we take $v_3 = v$. Note that $\{a, c, e\}$ is not recursively stopped because $\{e\}$ is not maximal in the stopping prefix of $\mathcal{E}_N^{\{a,c\}}$ (see Fig. A.7(d)). Finally, note that the branching cells of $\mathcal{E}_N^{\{a\}}$ (Fig. A.7(c)) and $\mathcal{E}_N^{\{b\}}$ (Fig. A.7(e)) correspond to different choices in $\mathcal{E}_N^{\emptyset}$ and thus have different stopping prefixes.*

Appendix A.3. AB's decomposition and terms

The recursively stopped configurations of a marked net (N, m) characterise all the allowed executions of N under the marking m. Hence, we formally link the recursively stopped configurations of $\mathcal{E}_{(N,m)}$ with the deterministic processes associated with $(\!|N, m|\!)$. We start by introducing the notion of configurations associated to a term.

Definition 11. *Given a term $T : i \xrightarrow{s} o$ and a marking $m \subseteq i$, the set of configurations of T under m, written $\mathsf{Conf}(T, m)$, is defined inductively as* *follows.*

$$
\begin{aligned}
&\mathsf{Conf}(I_s, m) &&= \{\emptyset\}\\
&\mathsf{Conf}(\bot_s, \emptyset) &&= \{\emptyset\}\\
&\mathsf{Conf}(T_1 \oplus T_2, m) &&= \{v_1 \cup v_2 \mid \forall j = 1,2.\ T_j : i_j \xrightarrow{s_j} o_j\\
& && \qquad\qquad\qquad\qquad \wedge\ v_j \in \mathsf{Conf}(T_j, m \cap i_j))\}\\
&\mathsf{Conf}(T_1; T_2, m) &&= \{v_1 \cup v_2 \mid v_1 \in \mathsf{Conf}(T_1, m)\ \wedge\ T_2 : i_2 \xrightarrow{s_2} o_2\\
& && \qquad\qquad\qquad\qquad \wedge\ v_2 \in \mathsf{Conf}(T_2, v_1^{\circ} \cap i_2)\}\\
&\mathsf{Conf}(\mathsf{C}(\Theta), \emptyset) &&= \Theta\\
&\mathsf{Conf}(\textstyle\sum_{m \subseteq i}\ m \triangleright T_m, m_j) &&= \mathsf{Conf}(T_j, \emptyset)
\end{aligned}
$$

Proposition 4. *Let $(N, m) : i \to o$ be a finite marked occurrence net and $T = (\!|N, m|\!)$. Then, for $j \subseteq i$, v is a maximal r-stopped configuration of $\mathcal{E}_{(N, m\cup j)}$ iff $v \in \mathsf{Conf}(T, j)$.*

Proof. The proof follows by structural induction on T.

- $T = I_s$. For all $j \subseteq i$, we have $\mathsf{Conf}(I_s, j) = \{\emptyset\}$. Consequently, $v \in \mathsf{Conf}(I_s, j)$ implies $v = \emptyset$. Since $(\!|N, m|\!) = I_s$, $(N, m) = I_s$. Then, $s = i$ and $m = \emptyset$. Therefore, $\mathcal{E}_{(N, m\cup j)} = \emptyset$. Consequently, $v \in \mathcal{E}_{(N, m\cup j)}$ implies $v = \emptyset$.

- $T = \bot_s$. It holds trivially because there is no (N, m) such that $(\!|N, m|\!) = \bot_s$.

- $T = T_1 \oplus T_2$. Then, $(N, m) = (N_1, m_1) \oplus (N_2, m_2)$, $T_1 = (\!|N_1, m_1|\!)$ $T_2 = (\!|N_2, m_2|\!)$. By inductive hypothesis, $v_i \in \mathsf{Conf}(T_i, j_i)$ iff v_i is an r-stopped configuration of $\mathcal{E}_{(N_i, m_i \cup j_i)}$. The proof follows by noting that the union of two disjoint r-stopped configurations is an r-stopped configuration.

- $T = T_1; T_2$. Then, $(N, m) = (N_1, m_1); (N_2, m_2)$, $T_1 = (\!|N_1, m_1|\!)$ $T_2 = (\!|N_2, m_2|\!)$. By inductive hypothesis, $v_i \in \mathsf{Conf}(T_i, j_i)$ iff v_i is an r-stopped configuration of $\mathcal{E}_{(N_i, m_i \cup j_i)}$. The proof follows by noting that v_1 is an r-stopped configuration of $\mathcal{E}_{(N, m\cup j)}$ and v_2 is an r-stopped configuration of $\mathcal{E}^{v_1}_{(N, m\cup j)}$. Consequently, $v = v_1 \cup v_2$ is an r-stopped configuration of $\mathcal{E}_{(N, m\cup j)}$.

- $T = \mathsf{C}(\Theta(N_{\mathbb{C}}))$. Then, $N = N_{\mathbb{C}}$ and $m = {}^\circ\mathbb{C}$. Moreover, $v \in \mathcal{E}_{(\mathbb{C}, {}^\circ\mathbb{C})}$ implies that v is a maximal deterministic process of $(\mathbb{C}, {}^\circ\mathbb{C})$, i.e., a transaction. Hence, $v \in \Theta(N_{\mathbb{C}})$ and $v \in \mathsf{Conf}(T, \emptyset)$.

- $T = \sum_{j \subseteq i} \; j \triangleright \bot_{d_j} \oplus T_j$ with $T_j = (\!|\mathsf{can}(N_{\mathbb{C}} @ m \cup j)|\!)$. Then, $v \in \mathsf{Conf}(T, j)$ iff $v \in \mathsf{Conf}(T_j, \emptyset)$. By inductive hypothesis, v is a maximal r-stopped configuration of $\mathcal{E}_{N_{\mathbb{C}} @ m \cup j}$. The proof is completed by noting that $\mathcal{E}_{N_{\mathbb{C}} @ m \cup j} = \mathcal{E}_{(N_{\mathbb{C}}, m \cup j)}$.

$\square$