

CONCURRENCY AND PROBABILITY: REMOVING CONFUSION, COMPOSITIONALLY

ROBERTO BRUNI, HERNÁN MELGRATTI, AND UGO MONTANARI

University of Pisa, Italy
e-mail address: bruni@di.unipi.it

ICC - Universidad de Buenos Aires - Conicet, Argentina
e-mail address: hmelgra@dc.uba.ar

University of Pisa, Italy
e-mail address: ugo@di.unipi.it

ABSTRACT. Assigning a satisfactory truly concurrent semantics to Petri nets with confusion and distributed decisions is a long standing problem, especially if one wants to resolve decisions by drawing from some probability distribution. Here we propose a general solution to this problem based on a recursive, static decomposition of (occurrence) nets in loci of decision, called *structural branching cells* (*s-cells*). Each s-cell exposes a set of alternatives, called *transactions*. Our solution transforms a given Petri net, possibly with confusion, into another net whose transitions are the transactions of the s-cells and whose places are those of the original net, with some auxiliary nodes for bookkeeping. The resulting net is confusion-free by construction, and thus conflicting alternatives can be equipped with probabilistic choices, while nonintersecting alternatives are purely concurrent and their probability distributions are independent. The validity of the construction is witnessed by a tight correspondence with the recursively stopped configurations of Abbes and Benveniste. Some advantages of our approach are that: i) s-cells are defined statically and locally in a compositional way; ii) our resulting nets faithfully account for concurrency.

1. INTRODUCTION

Concurrency theory and practice provide a useful abstraction for the design and use of a variety of systems. Concurrent computations (also *processes*), as defined in many models, are equivalence classes of executions, called *traces*, where the order of concurrent (i.e., independent) events is inessential. A key notion in concurrent models is *conflict* (also known as choices or decisions). Basically, two events are in conflict when they cannot occur in the same execution. The interplay between concurrency and conflicts introduces a phenomenon in which the execution of an event can be influenced by the occurrence of another concurrent (and hence independent) event. Such situation, known as *confusion*, naturally arises in concurrent and distributed systems and is intrinsic to problems involving mutual exclusion [30]. When interleaving semantics is considered, the problem is less compelling,

Key words and phrases: Petri nets, confusion, dynamic nets, persistent places, OR causality, concurrency, probabilistic computation.




FIGURE 1. Some nets (top) and their event structures (bottom)

however it has been recognised and studied from the beginning of net research [29], and to address it in a general and acceptable way can be considered as a long-standing open problem for concurrency theory.

To illustrate confusion, we rely on Petri nets [27, 28], which are a basic, well understood model of concurrency. The simplest example of (asymmetric) confusion is the net in Fig. 1a. We assume the reader is familiar with the firing semantics of Petri nets, otherwise see the short summary in Section 2.2. The net has two traces involving the concurrent events a and b , namely $\sigma_1 = a; b$ and $\sigma_2 = b; a$. Both traces define the same concurrent execution. Contrastingly, σ_1 and σ_2 are associated with completely different behaviours of the system as far as the resolution of choices is concerned. In fact, the system makes two choices while executing σ_1 : firstly, it chooses a over d , enabling c as an alternative to b ; secondly, b is selected over c . Differently, the system makes just one choice in σ_2 : since initially c is not enabled, b is executed without any choice; after that, the system chooses a over d . As illustrated by this example, the choices made by two different traces of the same concurrent computation may differ depending on the order in which concurrent events occur.

The fundamental problem behind confusion relates to the description of distributed, global choices. Such problem becomes essential when choices are driven by probabilistic distributions and one wants to assign probabilities to executions, as it is the case with probabilistic, concurrent models. Consider again Fig. 1a and assume that a is chosen over d with probability p_a while b is chosen over c with probability p_b . When driven by independent choices, the trace σ_1 has probability $p_a \cdot p_b$, while σ_2 has probability $1 \cdot p_a = p_a$. Hence, two linear representations of the same concurrent computation, which are deemed equivalent, would be assigned different probabilities.

Different solutions have been proposed in the literature for adding probabilities to Petri nets [12, 23, 24, 13, 19, 16, 6, 17, 7]. As a matter of fact, most of them replace nondeterminism with probability only in part, or take an interleaving semantics approach that disregards concurrency, or introduce time dependent stochastic distributions, thus giving up the abstract flavour of untimed truly concurrent models. Confusion-free probabilistic models have been studied in [32], but this class, which subsumes free-choice nets, is usually considered quite restrictive. More generally, the distributability of decisions has been studied, e.g., in [31, 18], but while the results in [31] apply to some restricted classes of nets, the approach in [18] requires nets to be decorated with agents and produces distributed models




FIGURE 2. AB's dynamic branching cells for the example in Fig. 1a

with both nondeterminism and probability, where concurrency depends on the scheduling of agents.

A substantial advance has been contributed by Abbes and Benveniste (AB) [1, 2, 3]. They consider prime event structures and provide a *branching cell decomposition* that establishes the order in which choices are resolved (see Section 4.2). Intuitively, the event structure in Fig. 1a has the three branching cells outlined in Fig. 2. First a decision between a and d must be taken (Fig. 2a): if a is executed, then a subsequent branching cell $\{b, c\}$ is enabled (Fig. 2b); otherwise (i.e., if d is chosen) the trivial branching cell $\{b\}$ is enabled (Fig. 2c). In this approach, the trace $\sigma_2 = b; a$ is not admissible, because the branching cell $\{b\}$ does not exist in the original decomposition (Fig. 2a): it appears *after* the choice of d over a has been resolved. Branching cells are equipped with independent probability distributions and the probability assigned to a concurrent execution is given by the product of the probabilities assigned by its branching cells. Notably, the sum of the probabilities of maximal configurations is 1. Every decomposition of a configuration yields an execution sequence compatible with that configuration. Unfortunately, certain sequences of events, legal w.r.t. the configuration, are not executable according to AB.

Problem statement. The question addressed in this paper is a foundational one: *can concurrency and general probabilistic distributions coexist in Petri nets? If so, under which circumstances? By coexistence we mean that all the following issues must be addressed:*

- (1) *Time independence:* Truly concurrent semantics usually assumes computation to be independent from the relative speed of processes. In this sense, although truly concurrent models have been extended in the literature with some notion of time such that occurrences of events are studied in terms of stochastic distributions, here we consider the more abstract case of untimed models only.
- (2) *Schedule independence:* Concurrent events must be driven by independent probability distributions. This item is tightly related to the confusion problem, where the set of alternatives, and thus their probability distribution, can be changed by the execution of some concurrent event.
- (3) *Probabilistic computation:* Nondeterministic choices must be replaceable by probabilistic choices. This means that whenever two transitions are enabled, the choice to fire one instead of the other is either inessential (because they are concurrent) or is driven by some probability distribution.
- (4) *Complete concurrency:* It must be possible to establish a bijective correspondence between equivalence classes of firing sequences and a suitable set of concurrent processes. In particular, given a concurrent process it must be possible to recover all its underlying firing sequences.
- (5) *Sanity check #1:* All firing sequences of the same process carry the same probability, i.e., the probability of a concurrent computation is independent from the order of execution.




FIGURE 3. Running example

(6) *Sanity check #2*: The sum of the probabilities assigned to all possible maximal processes must be 1.

In this paper we provide a positive answer for finite occurrence nets: given any such net we show how to define loci of decisions, called *structural branching cells (s-cells)*, and construct another net where independent probability distributions can be assigned to concurrent events. This means that each s-cell can be assigned to a distributed random agent and that any concurrent computation is independent from the scheduling of agents.

Overview of the approach. Following the rationale behind AB’s approach, a net is transformed into another one that postpones the execution of choices that can be affected by pending decisions. According to this intuition, the net in Fig. 1a is transformed into another one that delays the execution of b until all its potential alternatives (i.e., c) are enabled or definitively excluded. In this sense, b should never be executed before the decision between a and d is taken, because c could still be enabled (if a is chosen). As a practical situation, imagine that a and d are the choices of your partner to either come to town (a) or go to the sea (d) and that you can go to the theatre alone (b), which is always an option, or go together with him/her (c), which is possible only when he/she is in town and accepts the invitation. Of course you better postpone the decision until you know if your partner is in town or not. This behaviour is faithfully represented, e.g., by the confusion-free net in Fig. 1b, where two variants of b are made explicit: b_1 (your partner is in town) and b_2 (your partner is not in town). The new place $\neg c$ represents the fact that c will never be enabled.

Now, from the concurrency point of view, there is a single process that comprises both a and b_1 (with a a cause of b_1), whose overall probability is the product of the probability of choosing a over d by the probability of choosing b_1 over c . The other two processes comprise, respectively, d and b_2 (with d a cause of b_2) and a and c (with a a cause of c). As the net is confusion-free all criteria in the desiderata are met.

The general situation is more involved because: i) there can be several ways to disable the same transition; ii) resolving a choice may require to execute several transitions at once. Consider the net in Fig. 3a: i) c is discarded as soon as d or f fires; and ii) when both a and e are fired we can choose to execute c alone or *both* b and g . Likewise the previous example, we may expect to transform the net as in Fig. 3b. Again, the place $\neg c$ represents the permanent disabling of c . This way a probability distribution can drive the choice between c and (the joint execution of) bg , whereas b and g (if enabled) can fire concurrently when $\neg c$ is marked.

A few things are worth remarking: i) a token in $\neg c$ can be needed several times (e.g., to fire b and g), hence tokens should be read but not consumed from $\neg c$ (whence the double headed arcs from $\neg c$ to b and g , called *self-loops*); ii) several tokens can appear in the place $\neg c$ (by firing both d and f). These facts have severe repercussions on the concurrent semantics of the net. Suppose the trace $d; f; b$ is observed. The firings of d and f produce two tokens in the place $\neg c$: Does b causally depend on the token generated from d or from f (or from both)? Moreover, consider the trace $d; e; b; g$, in which b takes and releases a token in $\neg c$. Does g causally depend on b (due to such self-loop)? This last question can be solved by replacing self-loops with *read arcs* [25], so that the firing of b does not alter the content of $\neg c$ and thus no causal dependency arises between b and g . Nevertheless, if process semantics or event semantics is considered, then we should explode all possible combinations of causal dependencies, thus introducing a new, undesired kind of nondeterminism. In reality, we should not expect any causal dependency between b and g , while both have OR dependencies on d and f .

To account for OR dependencies, we exploit the notion of *persistence*: tokens in a persistent place have infinite weight and are collective. Namely, once a token reaches a persistent place, it cannot be removed and if two tokens reach the same persistent place they are indistinguishable. Such networks are a variant of ordinary P/T nets and have been studied in [11]. In the example, we can declare $\neg c$ to be a persistent place and replace self-loops/read arcs on $\neg c$ with ordinary outgoing arcs (see Fig. 3c). Nicely we are able to introduce a process semantics for nets with persistent places that satisfies complete concurrency.

The place $\neg c$ in the examples above is just used to sketch the general idea: our transformation introduces persistent places like $\bar{3}$ to express that a token will never appear in the regular place 3.

Contribution. In this paper we show how to systematically derive confusion-free nets (with persistency) from any (finite, occurrence) Petri net and equip them with probabilistic distributions and concurrent semantics in the vein of AB's construction.

Technically, our approach is based on a structurally recursive decomposition of the original net in s-cells. A simple kind of Asperti-Busi's dynamic nets is used as an intermediate model to structure the coding. While not strictly necessary, the intermediate step emphasises the hierarchical nature of the construction. The second part is a general flattening step independent of our special case. Our definition is purely local (to s-cells), static and

compositional, whereas AB's is dynamic and global (i.e., it requires the entire PES). Using nets with persistency, we compile the execution strategy of nets with confusion in a statically defined, confusion-free, operational model. The advantage is that the concurrency within a process of the obtained p-net is consistent with execution, i.e., all linearizations of a persistent process are executable.

Structure of the paper. After fixing notation in Section 2, our solution to the confusion problem consists of the following steps: (i) we define s-cells in a compositional way (Section 3.1); (ii) from s-cells decomposition and the use of dynamic nets, we derive a confusion-free net with persistency (Section 3.2); (iii) we prove the correspondence with AB's approach (Section 4); (iv) we define a new notion of process that accounts for OR causal dependencies and satisfies complete concurrency (Section 5); and (v) we show how to assign probability distributions to s-cells (Section 6). For the sake of readability, all proofs of main results can be found in Appendix.

2. PRELIMINARIES

2.1. Notation. We let \mathbb{N} be the set of natural numbers, $\mathbb{N}_\infty = \mathbb{N} \cup \{\infty\}$ 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 , a multiset m over S is an element of \mathbb{N}^S , and a bag b over S is an element of \mathbb{N}_∞^S . By overloading the notation, union, difference and inclusion of sets, multisets and bags are all denoted by the same symbols: \cup , \setminus and \subseteq , respectively. In the case of bags, the difference $b \setminus m$ is defined only when the second argument is a multiset, with the convention that $(b \setminus m)(s) = \infty$ if $b(s) = \infty$. Similarly, $(b \cup b')(s) = \infty$ if $b(s) = \infty$ or $b'(s) = \infty$. A set can be seen as a multiset or a bag whose elements have unary multiplicity. Membership is denoted by \in : for a multiset m (or a bag b), we write $s \in m$ for $m(s) \neq 0$ ($b(s) \neq 0$). Given a relation $R \subseteq S \times S$, we let 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, confusion and free-choiceness. A *net structure* N (also *Petri net*) [27, 28] 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 $\bullet t$ and t^\bullet are non empty for every $t \in T$. We write $t : X \rightarrow Y$ for $t \in T$ with $X = \bullet t$ and $Y = t^\bullet$.

A *marking* is a multiset $m \in \mathbb{N}^P$. We say that p is *marked* at m if $p \in m$. We write (N, m) for the net N *marked* by m . We write m_0 for the initial marking of the net, if any.

Graphically, a Petri net is a directed graph whose nodes are the places and transitions and whose set of arcs is F . Places are drawn as circles and transitions as rectangles. The marking m is represented by inserting $m(p)$ tokens in each place $p \in m$ (see Fig. 1).

A transition t is *enabled* at the marking m , written $m \xrightarrow{t}$, if $\bullet t \subseteq m$. The execution of a transition t enabled at m , called *firing*, is written $m \xrightarrow{t} m'$ with $m' = (m \setminus \bullet t) \cup t^\bullet$. A firing sequence from m to m' is a finite sequence of firings $m = m_0 \xrightarrow{t_1} \dots \xrightarrow{t_n} m_n = m'$, abbreviated to $m \xrightarrow{t_1 \dots t_n} m'$ or just $m \rightarrow^* m'$. Moreover, it is *maximal* if no transition is enabled at m' . We write $m \xrightarrow{t_1 \dots t_n}$ if there is m' such that $m \xrightarrow{t_1 \dots t_n} m'$. We say that m'



 (A) Concurrent run (a and b) (B) Concurrent run (b and d) (C) Sequential run (a before c)

FIGURE 4. Maximal processes for the net in Fig. 1a

is *reachable* from m if $m \rightarrow^* m'$. The set of markings reachable from m is written $[m]$. A marked net (N, m) is *safe* if each $m' \in [m]$ is a set.

Two transition t, u are in *direct conflict* if $\bullet t \cap \bullet u \neq \emptyset$. A net is called *free-choice* if for all transitions t, u we have either $\bullet t = \bullet u$ or $\bullet t \cap \bullet u = \emptyset$, i.e., if a transition t is enabled then all its conflicting alternatives are also enabled. Note that free-choiceness is purely structural. Confusion-freeness considers instead the dynamics of the net. A safe marked net (N, m_0) has *confusion* iff there exists a reachable marking m and transitions t, u, v such that:

- (1) (i) t, u, v are enabled at m , (ii) $\bullet t \cap \bullet u \neq \emptyset \neq \bullet u \cap \bullet v$, (iii) $\bullet t \cap \bullet v = \emptyset$ (*symmetric case*);
or
- (2) (i) t and v are enabled at m , (ii) u is not enabled at m but it becomes enabled after the firing of t , and (iii) $\bullet t \cap \bullet v = \emptyset$ and $\bullet v \cap \bullet u \neq \emptyset$ (*asymmetric case*).

In case 1, t and v are concurrently enabled but the firing of one disables an alternative (u) to the other. In case 2, the firing of t enables an alternative to u . An example of symmetric confusion is given by $m = \{2, 3, 8\}$, $t = b$, $u = c$ and $v = g$ in Fig. 3a, while for the asymmetric case take $m = \{1, 2\}$, $t = a$, $v = b$ and $u = c$ in Fig. 1a. A net is *confusion-free* when it has no confusion.

2.3. Deterministic Nonsequential Processes. A *deterministic nonsequential process* (or just *process*) [14] 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 $\pi : \mathcal{D} \rightarrow N$ from a *deterministic occurrence net* \mathcal{D} to N (preserving pre- and post-sets), where a deterministic occurrence net is such that: (1) the flow relation is acyclic, (2) there are no backward conflicts ($\forall p \in P. |\bullet p| \leq 1$), and (3) there are no forward conflicts ($\forall p \in P. |p \bullet| \leq 1$). We let ${}^\circ \mathcal{D} = \{p \mid \bullet p = \emptyset\}$ and $\mathcal{D}^\circ = \{p \mid p \bullet = \emptyset\}$ be the sets of *initial* and *final places* of \mathcal{D} , respectively (with $\pi({}^\circ \mathcal{D})$ be the initial marking of N). When N is an acyclic safe net, the mapping $\pi : \mathcal{D} \rightarrow N$ is just an injective graph homomorphism: without loss of generality, we name the nodes in \mathcal{D} as their images in N and let π be the identity. 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 .

For example, take the net in Fig. 1a. It has three maximal processes that are reported in Fig. 4. The equivalence class of the firing sequences $m_0 \xrightarrow{ab}$ and $m_0 \xrightarrow{ba}$ is the maximal process \mathcal{D} in Fig. 4a with places $\{1, 2, 3, 4\}$ and transitions $\{a : 1 \rightarrow 3, b : 2 \rightarrow 4\}$, where ${}^\circ \mathcal{D} = \{1, 2\}$ and $\mathcal{D}^\circ = \{3, 4\}$. Likewise, the equivalence class of the firing sequences $m_0 \xrightarrow{bd}$

and $m_0 \xrightarrow{db}$ is the maximal process in Fig. 4b. As c can only be executed after a , the corresponding process is in Fig.4c.

Given an acyclic net we let $\leq = F^*$ 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 $t_1 \#_0 t_2$. Then, a *nondeterministic occurrence net* (or just *occurrence net*) is a net $\mathcal{O} = (P, T, F)$ such that: (1) the flow relation is acyclic, (2) there are no backward conflicts ($\forall p \in P. |\bullet p| \leq 1$), and (3) there are no self-conflicts ($\forall t \in T. \neg(t \# t)$). The *unfolding* $\mathcal{U}(N)$ of a safe Petri net N is an occurrence net that accounts for all (finite and infinite) runs of N : its transitions model all the possible instances of transitions in N and its places model all the tokens that can be created in any run. Our construction takes a finite occurrence net as input, which can be, e.g., the (truncated) unfolding of any safe net.

2.4. Nets With Persistency. Nets with persistency (*p-nets*) [11] partition the set of places into regular places P (ranged by p, q, \dots) and persistent places \mathbf{P} (ranged by $\mathbf{p}, \mathbf{q}, \dots$). We use s to range over $\mathbb{S} = P \cup \mathbf{P}$ and write a p-net as a tuple (\mathbb{S}, T, F) . Intuitively, persistent places guarantee some sort of monotonicity about the knowledge of the system. Technically, this is realised by letting states be bags of places $b \in \mathbb{N}_\infty^{\mathbb{S}}$ instead of multisets, with the constraint that $b(p) \in \mathbb{N}$ for any regular place $p \in P$ and $b(\mathbf{p}) \in \{0, \infty\}$ for any persistent place $\mathbf{p} \in \mathbf{P}$. To guarantee that this property is preserved by firing sequences, we assume that the post-set t^\bullet of a transition t is the bag such that: $(t^\bullet)(p) = 1$ if $(t, p) \in F$ (as usual); $(t^\bullet)(\mathbf{p}) = \infty$ if $(t, \mathbf{p}) \in F$; and $(t^\bullet)(s) = 0$ if $(t, s) \notin F$. We say that a transition t is *persistent* if it is attached to persistent places only (i.e. if $\bullet t \cup t^\bullet \subseteq \mathbf{P}$).

The notions of enabling, firing, firing sequence and reachability extend in the obvious way to p-nets (when markings are replaced by bags). For example, a transition t is *enabled* at the bag b , written $b \xrightarrow{t}$, if $\bullet t \subseteq b$, and the firing of an enabled transition is written $b \xrightarrow{t} b'$ with $b' = (b \setminus \bullet t) \cup t^\bullet$.

A firing sequence is *stuttering* if it has multiple occurrences of a persistent transition. Since firing a persistent transition t multiple times is inessential, we consider non-stuttering firing sequences. (Alternatively, we can add a marked regular place p_t to the preset of each persistent transition t , so t fires at most once.)

A marked p-net (N, b_0) is *1- ∞ -safe* if each reachable bag $b \in [b_0]$ is such that $b(p) \in \mathbb{N}$ for all $p \in P$ and $b(\mathbf{p}) \in \{0, \infty\}$ for all $\mathbf{p} \in \mathbf{P}$. Note that in 1- ∞ -safe nets the amount of information conveyed by any reachable bag is finite, as each place is associated with one bit of information (marked or unmarked). Graphically, persistent places are represented by circles with double border (and they are either empty or contain a single token).

The notion of confusion extends to p-nets, by checking direct conflicts w.r.t. regular places only.

As an example, consider the 1- ∞ -safe, confusion-free p-net in Fig. 5. After firing a and c , the firing of b is inessential to enable d , because the persistent place $\mathbf{4}$ is marked by ∞ .

2.5. Dynamic Nets. Dynamic nets [4] are Petri nets whose sets of places and transitions may increase dynamically. We focus on a subclass of persistent dynamic nets that only allows for changes in the set of transitions, which is defined as follows.




FIGURE 5. A marked p-net




FIGURE 6. A dynamic p-net encoded as a p-net

Definition 2.1 (Dynamic p-nets). The set $\text{DN}(\mathbb{S})$ is the *least set* satisfying the recursive equation:

$$\text{DN}(\mathbb{S}) = \{(T, b) \mid T \subseteq 2^{\mathbb{S}} \times \text{DN}(\mathbb{S}) \wedge T \text{ finite} \wedge b \in \mathbb{N}_{\infty}^{\mathbb{S}}\}$$

The definition above is a domain equation for the set of dynamic p-nets over the set of places \mathbb{S} : the set $\text{DN}(\mathbb{S})$ is the least fixed point of the equation. The simplest elements in $\text{DN}(\mathbb{S})$ are pairs (\emptyset, b) with bag $b \in \mathbb{N}_{\infty}^{\mathbb{S}}$ (with $b(p) \in \mathbb{N}$ for any $p \in P$ and $b(\mathbf{p}) \in \{0, \infty\}$ for any $\mathbf{p} \in \mathbf{P}$). Nets (T, b) are defined recursively; indeed any element $t = (S, N) \in T$ stands for a transition with preset S and postset N , which is another element of $\text{DN}(\mathbb{S})$. An ordinary transition from b to b' has thus the form $(b, (\emptyset, b'))$. We write $S \rightarrow N$ for the transition $t = (S, N)$, $\bullet t = S$ for its *preset*, and $t \bullet = N \in \text{DN}(\mathbb{S})$ for its *postset*. For $N = (T, b)$ we say that T is the set of *top transitions* of N . All the other transitions are called *dynamic*.

The firing rule rewrites a dynamic p-net (T, b) to another one. The firing of a transition $t = S \rightarrow (T', b') \in T$ consumes the preset S and releases both the transitions T' and the tokens in b' . Formally, if $t = S \rightarrow (T', b') \in T$ with $S \subseteq b$ then $(T, b) \xrightarrow{t} (T \cup T', (b \setminus S) \cup b')$.

The notion of 1- ∞ -safe dynamic p-net is defined analogously to p-nets by considering the bags b of reachable states (T, b) .

A sample of a dynamic net is shown in Fig. 6a, whose only dynamic transition, which is activated by t_3 , is depicted with dashed border. The arrow between t_3 and t_b denotes the fact that t_b is activated dynamically by the firing of $t_3 : \bar{\mathbf{3}} \rightarrow (\{b : 2 \rightarrow 4\}, \{\bar{\mathbf{5}}\})$.

We show that any dynamic p-net can be encoded as a (flat) p-net. Our encoding resembles the one in [4], but it is simpler because we do not need to handle place creation. Intuitively, we release any transition t immediately but we add a persistent place \mathbf{p}_t to its preset, to enable t dynamically (\mathbf{p}_t is initially empty iff t is not a top transition). Given a set T of transitions, b_T is the bag such that $b_T(\mathbf{p}_t) = \infty$ if $t \in T$ and $b_T(s) = 0$ otherwise.

For $N = (T, b) \in \text{DN}(\mathbb{S})$, we let $\mathbb{T}(N) = T \cup \bigcup_{t \in T} \mathbb{T}(t^\bullet)$ be the set of all (possibly nested) transitions appearing in N . From Definition 2.1 it follows that $\mathbb{T}(N)$ is finite and well-defined.

Definition 2.2 (From dynamic to static). Given $N = (T, b) \in \text{DN}(\mathbb{S})$, the corresponding p-net $\llbracket N \rrbracket$ is defined as $\llbracket N \rrbracket = (\mathbb{S} \cup \mathbf{P}_{\mathbb{T}(N)}, \mathbb{T}(N), F, b \cup b_T)$, where

- $\mathbf{P}_{\mathbb{T}(N)} = \{\mathbf{p}_t \mid t \in \mathbb{T}(N)\}$; and
- F is such that for any $t = S \rightarrow (T', b') \in \mathbb{T}(N)$ then $t : \bullet t \cup \{\mathbf{p}_t\} \rightarrow b' \cup b_T$.

The transitions of $\llbracket N \rrbracket$ are those from N (set $\mathbb{T}(N)$). Any place of N is also a place of $\llbracket N \rrbracket$ (set \mathbb{S}). In addition, there is one persistent place \mathbf{p}_t for each $t \in \mathbb{T}(N)$ (set $\mathbf{P}_{\mathbb{T}(N)}$). The initial marking of $\llbracket N \rrbracket$ is that of N (i.e., b) together with the persistent tokens that enable the top transitions of N (i.e., b_T). Adding b_T is convenient for the statement in Proposition 2.4, but we could safely remove $\mathbf{P}_T \subseteq \mathbf{P}_{\mathbb{T}(N)}$ (and b_T) from the flat p-net without any consequence.

Example 2.3. The dynamic p-net N in Fig. 6a is encoded as the p-net $\llbracket N \rrbracket$ in Fig. 6b, which has as many transitions as N , but the preset of every transition contains an additional persistent place (depicted in grey) to indicate transition's availability. All the new places but p_b are marked because the corresponding transitions are initially available. Contrastingly, p_b is unmarked because the corresponding transition becomes available after the firing of t_3 .

The following result shows that all computations of a dynamic p-net can be mimicked by the corresponding p-net and vice versa. Hence, the encoding preserves also 1-safety over regular places.

Proposition 2.4. *Let $N = (T, b) \in \text{DN}(\mathbb{S})$. Then,*

- (1) $N \xrightarrow{t} N'$ implies $\llbracket N \rrbracket \xrightarrow{t} \llbracket N' \rrbracket$;
- (2) Moreover, $\llbracket N \rrbracket \xrightarrow{t} N'$ implies there exists N'' such that $N \xrightarrow{t} N''$ and $N' = \llbracket N'' \rrbracket$.

Corollary 2.5. $\llbracket N \rrbracket$ is 1- ∞ -safe iff N is 1-safe.

3. FROM PETRI NETS TO DYNAMIC P-NETS

In this section we show that any (finite, acyclic) net N can be associated with a confusion-free, dynamic p-net $\llbracket N \rrbracket$ by suitably encoding loci of decision. The mapping builds on the structural cell decomposition introduced below.

3.1. Structural Branching Cells. A structural branching cell represents a statically determined locus of choice, where the firing of some transitions is considered against all the possible conflicting alternatives. To each transition t we assign an 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 . For convenience, each s-cell $[t]$ also includes the places in the pre-sets of the transitions in $[t]$, i.e., we let the relation Pre^{-1} be also included in \sqsubseteq , with $\text{Pre} = F \cap (P \times T)$. This way, if $(p, t) \in F$ then $p \sqsubseteq t$ because $p \preceq t$ and $t \sqsubseteq p$ because $(t, p) \in \text{Pre}^{-1}$. Formally, we let \sqsubseteq be the transitive closure of the relation $\#_0 \cup \preceq \cup \text{Pre}^{-1}$. Since $\#_0$ is subsumed by the transitive closure of the relation $\preceq \cup \text{Pre}^{-1}$, we equivalently set $\sqsubseteq = (\preceq \cup \text{Pre}^{-1})^*$. Then, we let $\leftrightarrow = \{(x, y) \mid x \sqsubseteq y \wedge y \sqsubseteq x\}$. Intuitively, the choices available in the equivalence class $[t]_{\leftrightarrow}$ of a transition t must be resolved atomically.

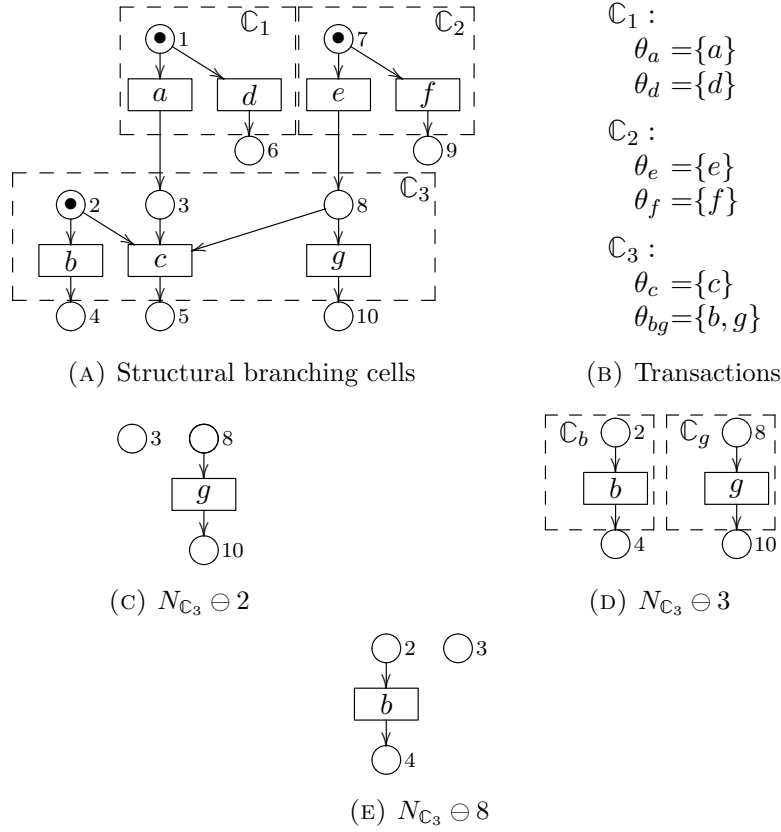


FIGURE 7. Structural branching cells (running example)

Definition 3.1 (S-cells). Let $N = (P, T, F)$ be a finite, nondeterministic occurrence net. The set $\text{BC}(N)$ of *s-cells* is the set of equivalence classes of \leftrightarrow , i.e., $\text{BC}(N) = \{[t]_{\leftrightarrow} \mid t \in T\}$.

Remark 3.2. Exploiting the algebraic structure of monoidal categories, in [9] we have given an alternative characterization of s-cells as those nets that can be decomposed neither in parallel nor in sequence. The alternative definition is maybe more intuitive, but its formalization requires some technical machinery which we prefer to leave out of the scope of the present paper.

We let \mathbb{C} range over s-cells. By definition it follows that for all $\mathbb{C}, \mathbb{C}' \in \text{BC}(N)$, if $\mathbb{C} \cap \mathbb{C}' \neq \emptyset$ then $\mathbb{C} = \mathbb{C}'$. For any s-cell \mathbb{C} , we denote by $N_{\mathbb{C}}$ the subnet of N whose elements are in $\mathbb{C} \cup \bigcup_{t \in \mathbb{C}} t^\bullet$. Abusing the notation, we denote by ${}^\circ\mathbb{C}$ the set of all the initial places in $N_{\mathbb{C}}$ and by \mathbb{C}° the set of all the final places in $N_{\mathbb{C}}$.

Definition 3.3 (Transactions). Let $\mathbb{C} \in \text{BC}(N)$. Then, a *transaction* θ of \mathbb{C} , written $\theta : \mathbb{C}$, is a maximal (deterministic) process of $N_{\mathbb{C}}$.

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.

Example 3.4. The net N in Fig. 3a has the three s-cells shown in Fig. 7a, whose transactions are listed in Fig. 7b. For \mathbb{C}_1 and \mathbb{C}_2 , each transition defines a transaction; \mathbb{C}_3 has one transaction associated with c and one with (the concurrent firing of) b and g .




FIGURE 8. S-cells for the net in Fig. 1a

The following operation \ominus is instrumental for the definition of our encoding and stands for the removal of a minimal place of a net and all the elements that causally depend on it. Formally, $N \ominus p$ is the least set that satisfies the rules (where \circ has higher precedence over set difference):

$$\frac{q \in {}^\circ N \setminus \{p\}}{q \in N \ominus p} \quad \frac{t \in N \quad \bullet t \subseteq N \ominus p}{t \in N \ominus p} \quad \frac{t \in N \ominus p \quad q \in t^\bullet}{q \in N \ominus p}$$

Example 3.5. Consider the net in Fig. 1a. There are two main s-cells: \mathbb{C}_1 associated with $\{a, d\}$, and \mathbb{C}_2 with $\{b, c\}$. There is also a nested s-cell \mathbb{C}_3 that arises from the decomposition of the subnet $N_{\mathbb{C}_2} \ominus 3$. All the above s-cells are shown in Fig. 8.

Example 3.6. Consider the s-cells in Fig. 7a. The net $N_{\mathbb{C}_1} \ominus 1$ is empty because every node in $N_{\mathbb{C}_1}$ causally depends on 1. Similarly, $N_{\mathbb{C}_2} \ominus 7$ is empty. The cases for \mathbb{C}_3 are in Figs. 7c–7e.

3.2. Encoding s-cells as confusion-free dynamic nets. Intuitively, the proposed encoding works by explicitly representing the fact that a place will not be marked in a computation. We denote with \bar{p} the place that models such “negative” information about the regular place p and let $\bar{P} = \{\bar{p} \mid p \in P\}$.¹ The encoding uses negative information to recursively decompose s-cells under the assumption that some of their minimal places will stay empty.

Definition 3.7 (From s-cells to dynamic p-nets). Let $N = (P, T, F, m)$ be a marked occurrence net. Its dynamic p-net $\llbracket N \rrbracket \in \text{DN}(P \cup \bar{P})$ is defined as $\llbracket N \rrbracket = (T_{\text{pos}} \cup T_{\text{neg}}, m)$, where:

$$\begin{aligned} T_{\text{pos}} &= \{ {}^\circ \mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ}) \mid \mathbb{C} \in \text{BC}(N) \text{ and } \theta : \mathbb{C} \} \\ T_{\text{neg}} &= \{ \bar{p} \rightarrow (T', \overline{\mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ}) \mid \mathbb{C} \in \text{BC}(N) \text{ and } p \in {}^\circ \mathbb{C} \\ &\quad \text{and } (T', b) = \llbracket N_{\mathbb{C}} \ominus p \rrbracket \} \end{aligned}$$

For any s-cell \mathbb{C} of N and transaction $\theta : \mathbb{C}$, the encoding generates a transition $t_{\theta, \mathbb{C}} = ({}^\circ \mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ})) \in T_{\text{pos}}$ to mimic the atomic execution of θ . Despite ${}^\circ \theta$ may be strictly included in ${}^\circ \mathbb{C}$, we define ${}^\circ \mathbb{C}$ as the preset of $t_{\theta, \mathbb{C}}$ to ensure that the execution of θ only starts when the whole s-cell \mathbb{C} is enabled. Each transition $t_{\theta, \mathbb{C}} \in T_{\text{pos}}$ is a transition

¹The notation \bar{P} denotes just a set of places whose names are decorated with a bar; it should not be confused with usual set complement.




FIGURE 9. S-cells as dynamic nets and their composed, flattened version (for the net in Fig. 1a)

of an ordinary Petri net because its postset consists of (i) the final places of θ and (ii) the negative versions of the places in $\mathbb{C}^\circ \setminus \theta^\circ$. A token in $\bar{\mathbf{p}} \in \mathbb{C}^\circ \setminus \theta^\circ$ represents the fact that the corresponding ordinary place $p \in \mathbb{C}^\circ$ will not be marked because it depends on discarded transitions (not in θ).

Negative information is propagated by the transitions in T_{neg} . For each cell \mathbb{C} and place $p \in {}^\circ\mathbb{C}$, there exists one dynamic transition $t_{p,\mathbb{C}} = \bar{\mathbf{p}} \rightarrow (T', \mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ)$ whose preset is just $\bar{\mathbf{p}}$ and whose postset is defined in terms of the subnet $N_{\mathbb{C}} \ominus p$. The postset of $t_{p,\mathbb{C}}$ accounts for two effects of propagation: (i) the generation of the negative tokens for all maximal places of \mathbb{C} that causally depend on p , i.e., for the negative places associated with the ones in \mathbb{C}° that are not in $(N_{\mathbb{C}} \ominus p)^\circ$; and (ii) the activation of all transitions T' obtained by encoding $N_{\mathbb{C}} \ominus p$, i.e., the behaviour of the branching cell \mathbb{C} after the token in the minimal place p is excluded. We remark that the bag b in $(T', b) = \llbracket N_{\mathbb{C}} \ominus p \rrbracket$ is always empty, because i) $N_{\mathbb{C}}$ is unmarked and, consequently, $N_{\mathbb{C}} \ominus p$ is unmarked, and ii) the initial marking of $\llbracket N \rrbracket$ corresponds to the initial marking of N .

Example 3.8. We sketch the main ideas over the net in Fig. 1a. We recall that it has two main s-cells (\mathbb{C}_1 associated with $\{a, d\}$, and \mathbb{C}_2 with $\{b, c\}$) and a nested one (\mathbb{C}_3): see Fig. 8. Their dynamic nets are in Figs. 9a–9b, where auxiliary transitions are in grey and unlabeled. Places $\bar{1}$ and $\bar{2}$ (and their transitions) are irrelevant, because the places 1 and 2 are already marked. However, our cells being static, we need to introduce auxiliary places in all cases. Note that in Fig. 9b there is an arc between two transitions. As explained before, this is because the target transition is dynamically created when the other is executed (hence the dashed border). Also note that there are two transitions with the same subscript b : one (t'_b) is associated with the s-cell \mathbb{C}_2 , the other (t_b) with the unique s-cell \mathbb{C}_3 of $N_{\mathbb{C}_2} \ominus 3$ and is released when the place $\bar{3}$ becomes marked.

After the s-cells are assembled and flattened we get the p-net in Fig. 9c (where irrelevant nodes are pruned). Initially, t_a and t_d are enabled. Firing t_a leads to the marking $\{2, 3, \bar{6}\}$ where $t'_b : \{2, 3\} \rightarrow \{4, \bar{5}\}$ and $t_c : \{2, 3\} \rightarrow \{\bar{4}, 5\}$ are enabled (and in conflict). Firing t_d instead leads to the marking $\{2, \bar{3}, 6\}$ where only the auxiliary transition can be fired, enabling $t_b : \{2, \mathbf{p}_b\} \rightarrow 4$. The net is confusion-free, as every conflict involves transitions with the same preset. For example, as the places 3 and $\bar{3}$ (and thus $\bar{\mathbf{p}}_b$) are never marked in a same run, the transitions $t'_b : \{2, 3\} \rightarrow \{4, \bar{5}\}$ and $t_b : \{2, \mathbf{p}_b\} \rightarrow 4$ will never compete for the token in 2.

$$\begin{array}{llll}
t_a : 1 \rightarrow (\emptyset, \{3, \overline{6}\}) & \text{for } \theta_a & t_2 : \overline{2} \rightarrow (\{t_g, t'_8\}, \{\overline{4}, \overline{5}\}) & \\
t_d : 1 \rightarrow (\emptyset, \{6, \overline{3}\}) & \text{for } \theta_d & t_3 : \overline{3} \rightarrow (\{t_b, t'_2, t_g, t'_8\}, \{\overline{5}\}) & \\
t_1 : \overline{1} \rightarrow (\emptyset, \{\overline{3}, \overline{6}\}) & & t_8 : \overline{8} \rightarrow (\{t_b, t'_2\}, \{\overline{5}, \overline{10}\}) & \\
t_e : 7 \rightarrow (\emptyset, \{8, \overline{9}\}) & \text{for } \theta_e & & \text{where} \\
t_f : 7 \rightarrow (\emptyset, \{9, \overline{8}\}) & \text{for } \theta_f & t_b : 2 \rightarrow (\emptyset, \{4\}) & \\
t_7 : \overline{7} \rightarrow (\emptyset, \{\overline{8}, \overline{9}\}) & & t'_2 : \overline{2} \rightarrow (\emptyset, \{\overline{4}\}) & \\
t_{bg} : 2, 3, 8 \rightarrow (\emptyset, \{4, 10, \overline{5}\}) & \text{for } \theta_{bg} & t_g : 8 \rightarrow (\emptyset, \{10\}) & \\
t_c : 2, 3, 8 \rightarrow (\emptyset, \{5, \overline{4}, \overline{10}\}) & \text{for } \theta_c & t'_8 : \overline{8} \rightarrow (\emptyset, \{\overline{10}\}) &
\end{array}$$

FIGURE 10. Encoding of branching cells (running example)

Example 3.9. Consider the net N and its s-cells in Fig. 7a. Then, $\llbracket N \rrbracket = (T, b)$ is defined such that b is the initial marking of N , i.e., $b = \{1, 2, 7\}$, and T has the transitions shown in Fig. 10.

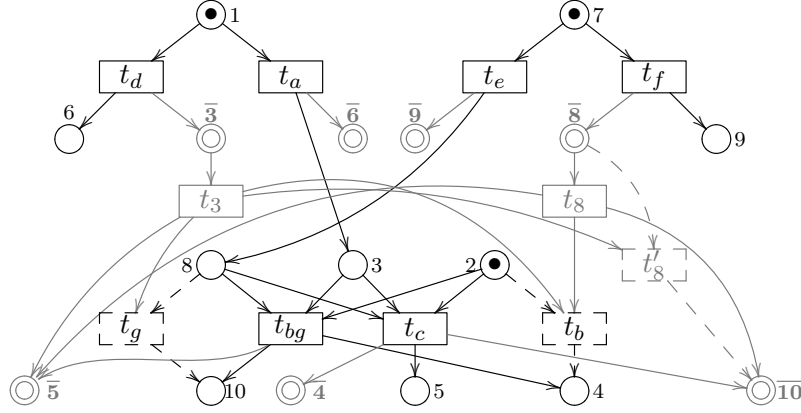
First consider the s-cell \mathbb{C}_1 . T_{pos} contains one transition for each transaction in \mathbb{C}_1 , namely t_a (for $\theta_a : \mathbb{C}_1$) and t_d (for $\theta_d : \mathbb{C}_1$). Both t_a and t_d have ${}^\circ\mathbb{C}_1 = \{1\}$ as preset. By definition of T_{pos} , both transitions have empty sets of transitions in their postsets. Additionally, t_a^\bullet produces tokens in $\theta_a^\circ = \{3\}$ (positive) and $\overline{\mathbb{C}_1}^\circ \setminus \theta_a^\circ = \{\overline{3}, \overline{6}\} \setminus \{3\} = \{\overline{6}\}$ (negative), while t_d^\bullet produces tokens in $\theta_d^\circ = \{6\}$ and $\overline{\mathbb{C}_1}^\circ \setminus \theta_d^\circ = \{\overline{3}\}$. Finally, $t_1 \in T_{\text{neg}}$ propagates negative tokens for the unique place in ${}^\circ\mathbb{C}_1 = \{1\}$. Since $N_{\mathbb{C}_1} \ominus 1$ is the empty net, $\llbracket N_{\mathbb{C}_1} \ominus 1 \rrbracket = (\emptyset, \emptyset)$. Hence, t_1 produces negative tokens for all maximal places of \mathbb{C}_1 , i.e., $\{\overline{3}, \overline{6}\}$. For the s-cell \mathbb{C}_2 we analogously obtain the transitions t_e, t_f and t_7 .

The s-cell \mathbb{C}_3 has two transactions θ_{bg} and θ_c . Hence, $\llbracket N \rrbracket$ has two transitions $t_{bg}, t_c \in T_{\text{pos}}$. Despite θ_{bg} mimics the firing of b and g , which are disconnected from the place 3, it is included in the preset of t_{bg} to postpone the firing of t_{bg} until \mathbb{C}_1 is executed. Transitions $t_2, t_3, t_8 \in T_{\text{neg}}$ propagate the negative information for the places in ${}^\circ\mathbb{C}_3 = \{2, 3, 8\}$. The transition t_3 has ${}^\bullet t_3 = \{\overline{3}\}$ as its preset and its postset is obtained from $N_{\mathbb{C}_3} \ominus 3$, which has two (sub) s-cells \mathbb{C}_b and \mathbb{C}_g (see Fig. 7d). The transitions t_b and t'_2 arise from \mathbb{C}_b , and t_g and t'_8 from \mathbb{C}_g . Hence, ${}^\bullet t_3 = (\{t_b, t'_2, t_g, t'_8\}, \{\overline{5}\})$ because $\llbracket N_{\mathbb{C}_3} \ominus 3 \rrbracket = (\{t_b, t'_2, t_g, t'_8\}, \emptyset)$ and $\overline{\mathbb{C}_3} \setminus (N_{\mathbb{C}_3} \ominus 3)^\circ = \{\overline{5}\}$. Similarly, we derive t_2 from $N_{\mathbb{C}_3} \ominus 2$ and t_8 from $N_{\mathbb{C}_3} \ominus 8$.

We now highlight some features of the encoded net. First, the set of top transitions is free-choice: positive and negative transitions have disjoint presets and the presets of any two positive transitions either coincide (if they arise from the same s-cell) or are disjoint. Recursively, this property holds at any level of nesting. Hence, the only source of potential confusion is due to the combination of top transitions and those activated dynamically, e.g., t_b and either t_{bg} or t_c . However, t_b is activated only when either $\overline{3}$ or $\overline{8}$ are marked, while ${}^\bullet t_{bg} = {}^\bullet t_c = \{2, 3, 8\}$. Then, confusion is avoided if p and \overline{p} can never be marked in the same execution (Lemma 3.10).

The net $\llbracket N \rrbracket$ is shown in Fig. 11, where the places $\{\overline{1}, \overline{2}, \overline{7}\}$ and the transitions $\{t_1, t_7, t_2, t'_2\}$ are omitted because superseded by the initial marking $\{1, 2, 7\}$.

We remark that the same dynamic transition can be released by the firing of different transitions (e.g., t_b by t_3 and t_8) and possibly several times in the same computation. Similarly, the same negative information can be generated multiple times. However this duplication has no effect, since we handle persistent tokens. For instance, the firing sequence


 FIGURE 11. Dynamic net $\llbracket N \rrbracket$ (running example)

$t_d; t_f; t_3; t_8$ releases two copies of t_b and marks $\bar{5}$ twice. This is inessential for reachability, but has interesting consequences w.r.t. causal dependencies (see Section 5).

We now show that the encoding generates confusion-free nets. We start by stating a useful property of the encoding that ensures that an execution cannot generate tokens in both p and \bar{p} .

Lemma 3.10 (Negative and positive tokens are in exclusion). *If $\llbracket N \rrbracket \rightarrow^* (T, b)$ and $\bar{p} \in b$ then $(T, b) \rightarrow^* (T', b')$ implies that $p \notin b'$.*

We now observe from Def. 3.7 that for any transition $t \in \llbracket N \rrbracket \in \text{DN}(P \cup \bar{P})$ it holds that either $\bullet t \subseteq P$ or $\bullet t \subseteq \bar{P}$. The next result says that whenever there exist two transitions t and t' that have different but overlapping presets, at least one of them is disabled by the presence of a negative token in the marking b .

Lemma 3.11 (Nested rules do not collide). *Let $\llbracket N \rrbracket \in \text{DN}(P \cup \bar{P})$. If $\llbracket N \rrbracket \rightarrow^* (T, b)$ then for all $t, t' \in T$ s.t. $\bullet t \neq \bullet t'$ and $\bullet t \cap \bullet t' \cap P \neq \emptyset$ it holds that there is $p \in P \cap (\bullet t \cup \bullet t')$ such that $\bar{p} \in b$.*

The main result states that $\llbracket \cdot \rrbracket$ generates confusion-free nets.

Theorem 3.12. *Let $\llbracket N \rrbracket \in \text{DN}(P \cup \bar{P})$. If $\llbracket N \rrbracket \rightarrow^* (T, b) \xrightarrow{t}$ and $(T, b) \xrightarrow{t'}$ then either $\bullet t = \bullet t'$ or $\bullet t \cap \bullet t' = \emptyset$.*

Corollary 3.13. *Any net $\llbracket N \rrbracket \in \text{DN}(P \cup \bar{P})$ is confusion-free.*

Finally, we can combine the encoding $\llbracket \cdot \rrbracket$ with (\cdot) (from Section 2.5) to obtain a (flat) 1- ∞ -safe, confusion-free, p-net $(\llbracket N \rrbracket)$, that we call the *uniformed net* of N . By Proposition 2.4 we get that the uniformed net $(\llbracket N \rrbracket)$ is also confusion-free by construction.

Corollary 3.14. *Any p-net $(\llbracket N \rrbracket)$ is confusion-free.*

4. STATIC VS DYNAMIC CELL DECOMPOSITION

As mentioned in the Introduction, Abbes and Benveniste proposed a way to remove confusion by dynamically decomposing prime event structures. In Sections 4.1 and 4.2 we recall the basics of the AB's approach as introduced in [1, 2, 3]. Then, we show that there is an operational correspondence between AB decomposition and s-cells introduced in Section 3.1.

4.1. Prime Event Structures. A *prime event structure* (also *PES*) [26, 34] 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.

Consider the nets in Figs. 1a and 3a. The corresponding PESs are shown below each net. Events are in bijective correspondence with the transitions of the nets. Strict causality is depicted by arrows and immediate conflict by curly lines.

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 $e \#_0 e'$ 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.

4.2. 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., $\forall e \in B, e' \in E. e \#_0 e' \Rightarrow e' \in B$. 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 . We assume that any $e \in E$ is contained in a finite stopping prefix.

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 finite 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 . We assume that any finite configuration enables only finitely many events, i.e., the set of minimal elements in E^v w.r.t. \preceq is finite for any $v \in \mathcal{V}_{\mathcal{E}}$.

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




FIGURE 12. AB's branching cell decomposition (running example)

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.

Example 4.1. Consider the PES \mathcal{E}_N in Fig. 3a and its maximal configuration $v = \{a, e, b, g\}$. 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. 12a. There are two possibilities for choosing $v_1 \subseteq v$ and v_1 recursively stopped: either $v_1 = \{a\}$ or $v_1 = \{e\}$. When $v_1 = \{a\}$, the choices for v_2 are determined by the stopping prefixes of $\mathcal{E}_N^{\{a\}}$ (see Fig. 12b) and the only possibility is $v_2 = \{a, e\}$. From $\mathcal{E}_N^{\{a,e\}}$ in Fig. 12c, we take $v_3 = v$. Note that $\{a, e, b\}$ is not recursively stopped because $\{b\}$ is not maximal in the stopping prefix of $\mathcal{E}_N^{\{a,e\}}$ (see Fig. 12c). Finally, note that the branching cells $\mathcal{E}_N^{\{a\}}$ (Fig. 12b) and $\mathcal{E}_N^{\{d\}}$ (Fig. 12d) correspond to different choices in \mathcal{E}_N^\emptyset and thus have different stopping prefixes.

4.3. Relating s-cells and AB's decomposition. The recursively stopped configurations of a net N characterise all the allowed executions of N . Hence, we formally link the recursively stopped configurations of \mathcal{E}_N with the computations of the uniformed net ($\llbracket N \rrbracket$). For technical convenience, we first show that the recursively stopped configurations of \mathcal{E}_N are in one-to-one correspondence with the computations of the dynamic net $\llbracket N \rrbracket$. Then, the desired correspondence is obtained by using Proposition 2.4 to relate the computations of a dynamic net and its associated p-net.

We rely on the auxiliary map $\|-\|$ that links transitions in $\llbracket N \rrbracket$ with events in \mathcal{E}_N . Specifically, $\|-\|$ associates each transition t of $\llbracket N \rrbracket$ with the set $\|t\|$ of transitions of N (also events in \mathcal{E}_N) that are encoded by t . Formally,

$$\|t\| = \begin{cases} ev(\theta) & \text{if } t = t_{\theta, \mathbb{C}} \in T_{\text{pos}} \\ \emptyset & \text{if } t \in T_{\text{neg}} \end{cases}$$

where $ev(\theta)$ is the set of transitions in θ .

Example 4.2. Consider the net N in Fig. 7a which is encoded as the dynamic p-net in Fig. 10. The auxiliary mapping $\|-\|$ is as follows

$$\begin{aligned} \|t_a\| &= \{a\} & \|t_d\| &= \{d\} & \|t_e\| &= \{e\} & \|t_f\| &= \{f\} \\ \|t_{bg}\| &= \{b, g\} & \|t_c\| &= \{c\} & \|t_b\| &= \{b\} & \|t_g\| &= \{g\} \\ \|t\| &= \emptyset & \text{if } t &\in \{t_1, t_7, t_2, t_3, t_8, t'_2, t'_8\} \end{aligned}$$

A transition $t_{\theta, \mathbb{C}}$ of $\llbracket N \rrbracket$ associated with a transaction $\theta : \mathbb{C}$ of N is mapped to the transitions of θ . For instance, t_a is mapped to $\{a\}$, which is the only transition in θ_a . Differently, transitions that propagate negative information, i.e., $t \in \{t_1, t_7, t_2, t_3, t_8, t'_2, t'_8\}$, are mapped to \emptyset because they do not encode any transition of N .

In what follows we write $M \Longrightarrow M'$ for a possibly empty firing sequence $M \xrightarrow{t_1 \cdots t_n} M'$ such that $\|t_i\| = \emptyset$ for all $i \in [1, n]$. If $\|t\| \neq \emptyset$, we write $M \xrightarrow{t} M'$ if $M \Longrightarrow M_0 \xrightarrow{t} M_1 \Longrightarrow M'$ for some M_0, M_1 . Moreover, we write $M \xrightarrow{t_1 \cdots t_n}$ if there exist M_1, \dots, M_n such that $M \xrightarrow{t_1} M_1 \xrightarrow{t_2} \cdots \xrightarrow{t_n} M_n$.

The following result states that the computations of any dynamic p-net produced by $\llbracket _ \rrbracket$ are in one-to-one correspondence with the recursively stopped configurations of Abbes and Benveniste.

Lemma 4.3. *Let N be an occurrence net.*

- (1) *If $\llbracket N \rrbracket \xrightarrow{t_1 \cdots t_n}$, then $v = \bigcup_{1 \leq i \leq n} \|t_i\|$ is recursively stopped in \mathcal{E}_N and $(\|t_i\|)_{1 \leq i \leq n}$ is a valid decomposition of v .*
- (2) *If v is recursively stopped in \mathcal{E}_N , then for any valid decomposition $(v_i)_{1 \leq i \leq n}$ there exists $\llbracket N \rrbracket \xrightarrow{t_1 \cdots t_n}$ such that $\|t_i\| = v_i$.*

Example 4.4. Consider the branching cell decomposition for $v = \{a, e, b, g\} \in \mathcal{E}_v$ discussed in Ex. 4.1. Then, the net $\llbracket N \rrbracket$ in Ex. 3.9 can mimic that decomposition with the following computation

$$(T, \{1, 2, 7\}) \xrightarrow{t_a} (T, \{2, 3, 7, \bar{\mathbf{6}}\}) \xrightarrow{t_e} (T, \{2, 3, 8, \bar{\mathbf{6}}, \bar{\mathbf{9}}\}) \xrightarrow{t_{bg}} (T, \{4, 10, \bar{\mathbf{5}}, \bar{\mathbf{6}}, \bar{\mathbf{9}}\})$$

with $v_1 = \|t_a\| = \{a\}$, $v_2 = \|t_e\| = \{e\}$, and $v_3 = \|t_{bg}\| = \{b, g\}$.

From Lemma 4.3 and Proposition 2.4 we obtain the next result.

Theorem 4.5 (Correspondence). *Let N be an occurrence net.*

- (1) *If $\llbracket \llbracket N \rrbracket \rrbracket \xrightarrow{t_1 \cdots t_n}$, then $v = \bigcup_{1 \leq i \leq n} \|t_i\|$ is recursively stopped in \mathcal{E}_N and $(\|t_i\|)_{1 \leq i \leq n}$ is a valid decomposition of v .*
- (2) *If v is recursively stopped in \mathcal{E}_N , then for any valid decomposition $(v_i)_{1 \leq i \leq n}$ there exists $\llbracket \llbracket N \rrbracket \rrbracket \xrightarrow{t_1 \cdots t_n}$ such that $\|t_i\| = v_i$.*

By (1) above, any computation of $\llbracket \llbracket N \rrbracket \rrbracket$ corresponds to a (recursively stopped) configuration of \mathcal{E}_N , i.e., a process of N . By (2), every execution of N that can be decomposed in terms of AB's branching cells is preserved by $\llbracket \llbracket N \rrbracket \rrbracket$, because any recursively stopped configuration of \mathcal{E}_N is mimicked by $\llbracket \llbracket N \rrbracket \rrbracket$.

5. CONCURRENCY OF THE UNIFORMED NET

In this section we study the amount of concurrency still present in the uniformed net $\llbracket \llbracket N \rrbracket \rrbracket$. Here, we extend the notion of a process to the case of 1- ∞ -safe p-nets and we show that all the *legal* firing sequences of a process of the uniformed net $\llbracket \llbracket N \rrbracket \rrbracket$ are executable.

The notion of deterministic occurrence net is extended to p-nets by slightly changing the definitions of conflict and causal dependency: (i) two transitions are *not* in conflict when all shared places are persistent, (ii) a persistent place can have more than one immediate cause in its preset, which introduces OR-dependencies.

Definition 5.1 (Persistent process). An *occurrence p-net* $O = (P \cup \mathbf{P}, T, F)$ is an acyclic p-net such that $|p^\bullet| \leq 1$ and $|\bullet p| \leq 1$ for any $p \in P$ (but not necessarily for those in \mathbf{P}).

A *persistent process* for N is an occurrence p-net O together with a net morphism $\pi : O \rightarrow N$ that preserves presets and postsets and the distinction between regular and

persistent places. Without loss of generality, when N is acyclic, we assume that O is a subnet of N (with the same initial marking) and π is the identity.

In an ordinary occurrence net, the causes of an item x are all its predecessors. In p-nets, the alternative sets of causes of an item x are given by a formula $\Phi(x)$ of the propositional calculus without negation, where the basic propositions are the transitions of the occurrence net. If we represent such a formula as a sum of products, it corresponds to a set of *collections*, i.e. a set of sets of transitions. Different collections correspond to alternative causal dependencies, while transitions within a collection are all the causes of that alternative and *true* represents the empty collection. Such a formula $\Phi(x)$ represents a monotone boolean function, which expresses, as a function of the occurrences of past transitions, if x has enough causes. It is known that such formulas, based on positive literals only, have a unique DNF (sum of products) form, given by the set of prime implicants. In fact, every prime implicant is also essential [33]. We define $\Phi(x)$ by well-founded recursion:

$$\Phi(x) = \begin{cases} \text{true} & \text{if } x \in P \cup \mathbf{P} \wedge \bullet x = \emptyset \\ \bigvee_{t \in \bullet x} (t \wedge \Phi(t)) & \text{if } x \in P \cup \mathbf{P} \wedge \bullet x \neq \emptyset \\ \bigwedge_{s \in \bullet x} \Phi(s) & \text{if } x \in T \end{cases}$$

The boolean formulas above remind the notion of causal automata [15], where a set of (labelled) events E is accompanied by an enabling function that assigns a formula in the free Boolean algebra generated by E to each event in E . Here note that all formulas are determined by the structure of the p-net and that they are expressed using only true, AND and OR, i.e. they would define a $\{\wedge, \vee\}$ -automata according to the terminology in [15].

Ordinary deterministic processes satisfy *complete concurrency*: each process determines a partial ordering of its transitions, such that the executable sequences of transitions are exactly the linearizations of the partial order. More formally, after executing any firing sequence σ of the process, a transition t is enabled if and only if all its predecessors in the partial order (namely its causes) already appear in σ . In the present setting a similar property holds.

Definition 5.2 (Legal firing sequence). A sequence of transitions $t_1; \dots; t_n$ of a persistent process is *legal* if for all $k \in [1, n]$ we have that $\bigwedge_{i=1}^{k-1} t_i$ implies $\Phi(t_k)$.

It is immediate to notice that if the set of persistent places is empty ($\mathbf{P} = \emptyset$) then the notion of persistent process is the ordinary one, $\Phi(x)$ is just the conjunction of the causes of x and a sequence is legal iff it is a linearization of the process.

Theorem 5.3 (Complete Concurrency). *Let $\sigma = t_1; \dots; t_n$ with $n \geq 0$ be a, possibly empty, firing sequence of a persistent process, and t a transition not in σ . The following conditions are all equivalent: (i) t is enabled after σ ; (ii) there is a collection of causes of t which appears in σ ; (iii) $\bigwedge_{i=1}^n t_i$ implies $\Phi(t)$.*

Corollary 5.4. *Given a persistent process, a sequence is legal iff it is a firing sequence.*

Example 5.5. Figs. 13a–13c show the maximal processes of the net in Fig. 9c. It is evident that all executions are serialized.

Example 5.6. Fig. 14 shows a process for the net $(\llbracket N \rrbracket)$ of our running example (see N in Fig. 3a and $\llbracket N \rrbracket$ in Fig. 11). The process accounts for the firing of the transitions d, f, b in N . Despite they look as concurrent events in N , the persistent place \mathbf{p}_{t_b} introduces some




FIGURE 13. Processes for the net in Fig. 9c




FIGURE 14. A process for $\llbracket N \rrbracket$ (running example)

causal dependencies. In fact, we have: $\Phi(t_d) = \Phi(t_f) = true$, $\Phi(t_3) = t_d$, $\Phi(t_8) = t_f$ and $\Phi(t_b) = (t_3 \wedge t_d) \vee (t_8 \wedge t_f)$, thus t_b can be fired only after either t_d or t_f (or both). The other maximal processes are reported in Appendix E.

Still referring to the net $\llbracket N \rrbracket$ of our running example, a more interesting case to consider is the process in Fig. 15a whose transition t_{bg} stands for the transaction where b and g are executed simultaneously. One may argue that having t_{bg} as an atomic action can reduce the overall concurrency of the system. However, t_{bg} can be expanded with (a fresh copy of) its underlying process as shown in Fig. 15b. We use t_{bg} as a subscript for the new nodes of the process to guarantee they are fresh. The preset of the transition t_{bg} is left unchanged. Its postset takes care of the propagation of negative information and of enabling the initial places of the underlying process. The final places of the underlying process are the (positive) places in the postset of the original transition t_{bg} . This transformation has the side effect to separate the choice of the transaction from its execution, but it increases the amount of concurrency, as $b_{t_{bg}}$ and $g_{t_{bg}}$ can now be executed in any order. While it might be possible in some cases to avoid the additional choice event, the general construction would look cumbersome.

The improvement of the amount of concurrency that can be observed is even more evident if we consider the net N in Fig. 16a. There are two s-cells \mathcal{C}_1 and \mathcal{C}_2 : the former




FIGURE 15. Expanding transactions I




FIGURE 16. Expanding transactions II

has two transactions $\theta_a = \{a\}$ and $\theta_{bc} = \{b, c\}$, the latter is trivial as it has only one transaction $\theta_d = \{d\}$. One process of $(\llbracket N \rrbracket)$ is in Fig. 16b, where t_{bc} and t_d are executed concurrently. However it does not take into account the fact that the execution of b and c can be interleaved with that of d . If we expand t_{bc} as discussed in the previous example, we get the process in Fig. 16c, where d can be executed after b and before c .

Formally, given a process θ and a transition name t , let θ_t be the process where any non final place/transition n is renamed to n_t and any final place is left unchanged. We say a positive transition $t_{\theta, \mathbb{C}}$ of $(\llbracket N \rrbracket)$ is *non-atomic* if the process θ involves more than one transition. Given a net N and its uniformed net $(\llbracket N \rrbracket)$ we let $(\llbracket N \rrbracket)_{conc}$ denote the persistent net obtained from $(\llbracket N \rrbracket)$ by removing each non-atomic transition $t_{\theta, \mathbb{C}}$ and by adding, for each such transition, the places and transitions in $\theta_{t_{\theta, \mathbb{C}}}$ together with a transition $t'_{\theta, \mathbb{C}}$ such




FIGURE 17. A free-choice net

that $\bullet t'_{\theta, \mathbb{C}} = \bullet t_{\theta, \mathbb{C}}$ but whose postset consists of ${}^{\circ}t_{\theta, \mathbb{C}}$ together with the negative places in the postset of $t_{\theta, \mathbb{C}}$.

6. PROBABILISTIC NETS

We can now outline our methodology to assign probabilities to the concurrent runs of a Petri net, also in the presence of confusion. Given a net N , we apply s-cell decomposition from Section 3.1, and then we assign probability distributions to the transactions available in each cell \mathbb{C} (and recursively to the s-cell decomposition of $N_{\mathbb{C}}$). Let $\mathcal{P}_{\mathbb{C}} : \{\theta \mid \theta : \mathbb{C}\} \rightarrow [0, 1]$ denote the probability distribution function of the s-cell \mathbb{C} (such that $\sum_{\theta : \mathbb{C}} \mathcal{P}_{\mathbb{C}}(\theta) = 1$). Such probability distributions are defined locally and transferred automatically to the transitions in T_{pos} of the dynamic p-net $\llbracket N \rrbracket$ defined in Section 3, in such a way that $\mathcal{P}(t_{\theta, \mathbb{C}}) = \mathcal{P}_{\mathbb{C}}(\theta)$. Each negative transition in T_{neg} has probability 1 because no choice is associated with it. Since the uniformed net $\langle \llbracket N \rrbracket \rangle$ has the same transitions of $\llbracket N \rrbracket$, the probability distribution can be carried over $\langle \llbracket N \rrbracket \rangle$ (thanks to Proposition 2.4).

AB's probability distribution. Building on the bijective correspondence in Theorem 4.3, the distribution $\mathcal{P}_{\mathbb{C}}$ can be chosen in such a way that it is consistent with the one attached to the transitions of Abbes and Benveniste's branching cells (if any).

Purely local distribution. Another simple way to define $\mathcal{P}_{\mathbb{C}}$ is by assigning probability distributions to the arcs leaving the same place of the original net, as if each place were able to decide autonomously which transition to fire. Then, given a transaction $\theta : \mathbb{C}$, we can set $\mathcal{Q}_{\mathbb{C}}(\theta)$ be the product of the probability associated with the arcs of N entering the transitions in θ . Of course, in general it can happen that $\sum_{\theta : \mathbb{C}} \mathcal{Q}_{\mathbb{C}}(\theta) < 1$, as not all combinations are feasible. However, it is always possible to normalise the quantities of feasible assignments by setting $\mathcal{P}_{\mathbb{C}}(\theta) = \frac{\mathcal{Q}_{\mathbb{C}}(\theta)}{\sum_{\theta' : \mathbb{C}} \mathcal{Q}_{\mathbb{C}}(\theta')}$ for any transaction $\theta : \mathbb{C}$.

Example 6.1. Take the free-choice net in Fig. 17 and assume that decisions are local to each place. Thus, place 1 lends its token to a with probability $p_1 = \frac{1}{3}$ and to b with $q_1 = \frac{2}{3}$. Similarly, place 2 lends its token to a with probability $p_2 = \frac{1}{3}$ and to b with $q_2 = \frac{2}{3}$. Then one can set $p_a = p_1 \cdot p_2 = \frac{1}{9}$ and $p_b = q_1 \cdot q_2 = \frac{4}{9}$. However their sum is $\frac{5}{9} \neq 1$. This anomaly is due to the existence of deadlocked choices with nonzero probabilities which disappear in the process semantics of nets. To some extent, the probabilities assigned to a and b should be conditional w.r.t. the fact that the local choices performed at places 1 and 2 are compatible, i.e., all non compatible choices are disregarded. This means that we need to normalize the values of p_a and p_b over their sum. Of course, normalisation is possible

only if there is at least one admissible alternative. In this simple example we get $\mathcal{Q}(a) = \frac{1}{9}$, $\mathcal{Q}(b) = \frac{4}{9}$, $\mathcal{P}(a) = \frac{1}{9}/\frac{5}{9} = \frac{1}{5}$ and $\mathcal{P}(b) = \frac{4}{9}/\frac{5}{9} = \frac{4}{5}$.

Example 6.2. Suppose that in our running example we assign uniform distributions to all arcs leaving a place. From simple calculation we have $\mathcal{P}_{\mathcal{C}_1}(\theta_a) = \mathcal{P}_{\mathcal{C}_1}(\theta_d) = \frac{1}{2}$ for the first cell, $\mathcal{P}_{\mathcal{C}_2}(\theta_e) = \mathcal{P}_{\mathcal{C}_2}(\theta_f) = \frac{1}{2}$ for the second cell, $\mathcal{P}_{\mathcal{C}_3}(\theta_c) = \mathcal{P}_{\mathcal{C}_3}(\theta_{bg}) = \frac{1}{2}$ for the third cell. The transactions of nested cells are uniquely defined and thus have all probability 1.

Given a firing sequence $t_1; \dots; t_n$ we can set $\mathcal{P}(t_1; \dots; t_n) = \prod_{i=1}^n \mathcal{P}(t_i)$. Hence firing sequences that differ in the order in which transitions are fired are assigned the same probability. Thanks to Theorem 5.3, we can consider maximal persistent processes instead of firing sequences and set $\mathcal{P}(O) = \prod_{t \in O} \mathcal{P}(t)$. In fact any maximal firing sequence in O includes all transitions of O and its probability is independent from the order of firing. It follows from Theorem 4.3 that any maximal configuration has a corresponding maximal process (and viceversa) and since Abbes and Benveniste proved that the sum of the probabilities assigned to maximal configurations is 1, the same holds for maximal persistent processes.

Example 6.3. Suppose the distributions are assigned as in Example 6.2. Then, the persistent process in Fig. 14 has probability: $\mathcal{P}(O) = \mathcal{P}(t_d) \cdot \mathcal{P}(t_f) \cdot \mathcal{P}(t_3) \cdot \mathcal{P}(t_8) \cdot \mathcal{P}(t_b) \cdot \mathcal{P}(t'_8) = \frac{1}{2} \cdot \frac{1}{2} \cdot 1 \cdot 1 \cdot 1 \cdot 1 = \frac{1}{4}$. There are other four maximal processes shown in Appendix E together with their probabilities. We note that the sum of all probabilities assigned to maximal processes is indeed 1.

7. CONCLUSION AND FUTURE WORK

AB's branching cells are a sort of interpreter (or scheduler) for executing PESs in the presence of confusion. Our main results develop along two orthogonal axis. Firstly, our approach is an innovative construction with the following advantages:

- (1) Compositionality: s-cells are defined statically and locally, while AB's branching cells are defined dynamically and globally (by executing the whole event structure).
- (2) Compilation vs interpretation: AB's construction gives an interpreter that rules out some executions of an event structure. We instead compile a net into another one (with persistency) whose execution is driven by ordinary firing rules.
- (3) Complete concurrency: AB's recursively stopped configurations may include traces that cannot be executed by the interpreter. Differently, our notion of process captures all and only those executable traces of a concurrent computation.
- (4) Simplicity: s-cells definition in terms of a closure relation takes a couple of lines (see Definition 3.1), while AB's branching cell definition is more involved.
- (5) Full matching: we define a behavioural correspondence that relates AB's maximal configurations with our maximal deterministic processes, preserving their probability assignment.

Secondly, we provide the following fully original perspectives:

- (1) Confusion removal: our target model is confusion-free.
- (2) Locally executable model: probabilistic choices are confined to transitions with the same pre-set, and hence can be resolved locally and concurrently. Besides, our target model relies on ordinary firing rules (with persistent places).
- (3) Processes: we define a novel notion of process for nets with persistency that conservatively extends the ordinary notion of process and captures the right amount of concurrency.

(4) Goal satisfaction: our construction meets all requirements in the list of desiderata.

This paper has extended the conference version [8] by including more examples and all detailed proofs of main results. The idea of expanding transactions into the underlying processes is also original to this contribution.

Moreover, the construction presented here has opened the way to other interesting research directions. First, it has led to the implementation of a tool, called `RemConf` [22] after “removal of confusion”, that takes in input an acyclic net N in the standard format `.pnml` and returns the net $\llbracket N \rrbracket$. As persistent places cannot be modeled in `.pnml`, they are implemented using self-loops instead of input arcs when they are part of the preset of a transition. This makes it possible to simulate the execution of $\llbracket N \rrbracket$ using any `.pnml` compatible tool. The tool is available at <http://remconf.di.unipi.it>.

Regarding OR causality, it can be accounted for by general event structures, where events can be enabled by distinct minimal sets of events. Several classes of event structures have been studied to disambiguate causality, in the sense that all alternative causes for an event must be somehow in conflict. This is the case, e.g., of stable event structures [35] and bundle event structures [20]. Five different classes of event structures that allow for causal ambiguity have been proposed in [21]. In [5] we have extended the connection between p-nets and event structures in order to deal with OR causes. Generalising the work on ordinary nets, Petri nets with persistent places are related to a new subclass of general event structures, called locally connected, by means of a chain of coreflections relying on an unfolding construction, as for the original construction by Winskel. The fact that a whole body of theory can be extended from Petri nets to p-nets witnesses that p-nets can be chosen as a general computational model and not just a convenient variant of an existing model. The causal AND/OR-dependencies share some similarities also with the work on connectors and Petri nets with boundaries [10] that we would like to formalize.

We also want to investigate the connection between our s-cell structure and Bayesian networks, so to make forward and backward reasoning techniques available in our setting. Some results in this direction can be found in [9].

Our construction is potentially complex: given a s-cell \mathbb{C} we recursively consider the nested s-cells in $N_{\mathbb{C}} \ominus p$, for any initial place $p \in N_{\mathbb{C}}$. In the worst case, the number of nested s-cells can be exponential in the number of their initial places. However s-cells are typically much smaller than the whole net and it can be the case that the size of all s-cells is bound by some fixed k . In this case, the number of s-cells in our construction can still become exponential on the constant k , but linear w.r.t. the number of places of the net.

A limitation of our approach is that it applies to finite occurrence nets only (or, equivalently, to finite PESs). As a future work, we plan to deal with cycles and unfolding semantics. This requires some efforts and we conjecture it is feasible only if the net is safe and its behaviour has some regularity: the same s-cell can be executed several times in a computation but every instance is restarted without tokens left from previous rounds.

Acknowledgments. The idea of structural cells emerged, with a different goal, in collaboration with Lorenzo Galeotti after his MSc thesis. The tool `RemConf` has been developed as part of the MSc thesis of Gianluca Maraschio. We thank Glynn Winskel, Holger Hermanns, Joost-Pieter Katoen for giving us insightful references. The research has been partially supported by EU H2020 RISE programme under the Marie Marie Skłodowska-Curie grant agreement 778233, by the UBACyT projects 20020170100544BA and 20020170100086BA, by the PIP project 11220130100148CO and by Università di Pisa projects PRA_2016_64 *Through*

the fog and PRA_2018_66 *DECLWARE: Metodologie dichiarative per la progettazione e il deployment di applicazioni*. The third author carried on part of the work while attending a Program on Logical Structures in Computation at Simons Institute, Berkeley, 2016.

REFERENCES

- [1] Samy Abbes and Albert Benveniste. Branching cells as local states for event structures and nets: Probabilistic applications. In *FOSSACS'05*, volume 3441 of *Lect. Notes in Comp. Sci.*, pages 95–109. Springer, 2005.
- [2] Samy Abbes and Albert Benveniste. True-concurrency probabilistic models: Branching cells and distributed probabilities for event structures. *Inf. Comput.*, 204(2):231–274, 2006.
- [3] Samy Abbes and Albert Benveniste. True-concurrency probabilistic models: Markov nets and a law of large numbers. *Theor. Comput. Sci.*, 390(2-3):129–170, 2008.
- [4] Andrea Asperti and Nadia Busi. Mobile Petri nets. *Mathematical Structures in Computer Science*, 19(6):1265–1278, 2009.
- [5] Paolo Baldan, Roberto Bruni, Andrea Corradini, Fabio Gadducci, Hernán C. Melgratti, and Ugo Montanari. Event structures for Petri nets with persistence. *Logical Methods in Computer Science*, 14(3), 2018.
- [6] Anne Bouillard, Stefan Haar, and Sidney Rosario. Critical paths in the partial order unfolding of a stochastic Petri net. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 43–57. Springer, 2009.
- [7] Ed Brinksma, Joost-Pieter Katoen, Rom Langerak, and Diego Latella. A stochastic causality-based process algebra. *Comput. J.*, 38(7):552–565, 1995.
- [8] Roberto Bruni, Hernán C. Melgratti, and Ugo Montanari. Concurrency and probability: Removing confusion, compositionally. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 195–204. ACM, 2018.
- [9] Roberto Bruni, Hernán C. Melgratti, and Ugo Montanari. Unifying inference for Bayesian and Petri nets. *CoRR*, abs/1807.06305, 2018.
- [10] Roberto Bruni, Hernán C. Melgratti, Ugo Montanari, and Pawel Sobocinski. Connector algebras for C/E and P/T nets' interactions. *Logical Methods in Computer Science*, 9(3), 2013.
- [11] Federico Crazzolara and Glynn Winskel. Petri nets with persistence. *Electr. Notes Theor. Comput. Sci.*, 121:143–155, 2005.
- [12] Joanne Bechta Dugan, Kishor S. Trivedi, Robert Geist, and Victor F. Nicola. Extended stochastic Petri nets: Applications and analysis. In *Performance'84*, pages 507–519. North-Holland, 1984.
- [13] Christian Eisentraut, Holger Hermanns, Joost-Pieter Katoen, and Lijun Zhang. A semantics for every GSPN. In *Petri Nets 2013*, volume 7927 of *Lect. Notes in Comp. Sci.*, pages 90–109. Springer, 2013.
- [14] Ursula Goltz and Wolfgang Reisig. The non-sequential behavior of Petri nets. *Information and Control*, 57(2/3):125–147, 1983.
- [15] Jeremy Gunawardena. Causal automata. *Theor. Comput. Sci.*, 101(2):265–288, 1992.
- [16] Stefan Haar. Probabilistic cluster unfoldings. *Fundamenta Informaticae*, 53(3-4):281–314, 2002.
- [17] Joost-Pieter Katoen, Rom Langerak, and Diego Latella. Modeling systems by probabilistic process algebra: An event structures approach. In *FORTE'93*, IFIP Transactions, pages 253–268. North-Holland, 1994.
- [18] Joost-Pieter Katoen and Doron A. Peled. Taming confusion for modeling and implementing probabilistic concurrent systems. In *ESOP'13*, volume 7792 of *Lect. Notes in Comp. Sci.*, pages 411–430. Springer, 2013.
- [19] Manfred Kudlek. Probability in Petri nets. *Fundamenta Informaticae*, 67(1-3):121–130, 2005.
- [20] Rom Langerak. Bundle event structures: a non-interleaving semantics for LOTOS. In Michel Diaz and Roland Groz, editors, *FORTE '92*, volume C-10 of *IFIP Transactions*, pages 331–346. North-Holland, 1993.
- [21] Rom Langerak, Ed Brinksma, and Joost-Pieter Katoen. Causal ambiguity and partial orders in event structures. In Antoni W. Mazurkiewicz and Józef Winkowski, editors, *CONCUR'97*, volume 1243 of *Lecture Notes in Computer Science*, pages 317–331. Springer, 1997.

- [22] Gianluca Maraschio. A tool for confusion removal in probabilistic concurrent models, 2018. MSc thesis, Università degli Studi di Pisa.
- [23] Marco Ajmone Marsan, Gianni Conte, and Gianfranco Balbo. A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems. *ACM Trans. Comput. Syst.*, 2(2):93–122, 1984.
- [24] Michael K. Molloy. Discrete time stochastic Petri nets. *IEEE Trans. Softw. Eng.*, 11(4):417–423, April 1985.
- [25] Ugo Montanari and Francesca Rossi. Contextual nets. *Acta Inf.*, 32(6):545–596, 1995.
- [26] Mogens Nielsen, Gordon D. Plotkin, and Glynn Winskel. Petri nets, event structures and domains, part I. *Theor. Comput. Sci.*, 13:85–108, 1981.
- [27] C.A. Petri. *Kommunikation mit Automaten*. PhD thesis, Institut für Instrumentelle Mathematik, Bonn, 1962.
- [28] Wolfgang Reisig. *Understanding Petri Nets - Modeling Techniques, Analysis Methods, Case Studies*. Springer, 2013.
- [29] Grzegorz Rozenberg and Joost Engelfriet. Elementary net systems. In *Advances in Petri Nets 1996, Part I*, volume 1491 of *Lect. Notes in Comp. Sci.*, pages 12–121. Springer, 1998.
- [30] Einar Smith. On the border of causality: Contact and confusion. *Theor. Comput. Sci.*, 153(1&2):245–270, 1996.
- [31] Rob J. van Glabbeek, Ursula Goltz, and Jens-Wolfhard Schicke-Uffmann. On characterising distributability. *Logical Methods in Computer Science*, 9(3), 2013.
- [32] Daniele Varacca, Hagen Völzer, and Glynn Winskel. Probabilistic event structures and domains. *Theor. Comput. Sci.*, 358(2-3):173–199, 2006.
- [33] Ingo Wegener. *The complexity of Boolean functions*. Wiley-Teubner series in computer science. John Wiley & Sons, Inc., New York, NY, USA, 1987.
- [34] Glynn Winskel. Event structures. In *Advances in Petri Nets 1986, Part II*, volume 255 of *Lect. Notes in Comp. Sci.*, pages 325–392. Springer, 1987.
- [35] Glynn Winskel. An introduction to event structures. In J. W. de Bakker, Willem P. de Roever, and Grzegorz Rozenberg, editors, *Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency, School/Workshop*, volume 354 of *Lecture Notes in Computer Science*, pages 364–397. Springer, 1989.

APPENDIX A. DETAILED PROOFS OF RESULTS IN SECTION 2.4

Proposition A.1 (Proposition 2.4). *Let $N = (T, b) \in \text{DN}(\mathbb{S})$. Then,*

- (1) $N \xrightarrow{t} N'$ implies $\llbracket N \rrbracket \xrightarrow{t} \llbracket N' \rrbracket$;
- (2) Moreover, $\llbracket N \rrbracket \xrightarrow{t} N'$ implies there exists N'' such that $N \xrightarrow{t} N''$ and $N' = \llbracket N'' \rrbracket$.

Proof. We start by showing that $N \xrightarrow{t} N'$ implies $\llbracket N \rrbracket \xrightarrow{t} \llbracket N' \rrbracket$. If $N \xrightarrow{t} N'$ then $t = S \rightarrow (T', b') \in T$, $S \subseteq b$ and $N' = (T \cup T', (b \setminus S) \cup b')$. By definition of $\llbracket _ \rrbracket$, it holds that $\llbracket N \rrbracket = (\mathbb{S} \cup \mathbf{P}_{\mathbb{T}(N)}, \mathbb{T}(N), F, b \cup b_T)$ where $\mathbf{P}_{\mathbb{T}(N)} = \{\mathbf{p}_{t'} \mid t' \in \mathbb{T}(N)\}$. Note that

$$\begin{aligned}
 \mathbb{T}(N) &= T \cup \bigcup_{t' \in T} \mathbb{T}(t'^{\bullet}) && \text{by def. of } \mathbb{T}(_) \\
 &= T \cup \bigcup_{t' \in T} \mathbb{T}(t'^{\bullet}) \cup \mathbb{T}(t^{\bullet}) && t \in T \\
 &= T \cup \bigcup_{t' \in T} \mathbb{T}(t'^{\bullet}) \cup T' \cup \bigcup_{t' \in T'} \mathbb{T}(t'^{\bullet}) && \text{by def. of } \mathbb{T}(_) \\
 &= (T \cup T') \cup \bigcup_{t' \in T \cup T'} \mathbb{T}(t'^{\bullet}) && \text{by assoc. and} \\
 & && \text{comm. of } \cup \\
 &= \mathbb{T}(N') && \text{by def. of } \mathbb{T}(_)
 \end{aligned}$$

Hence, $\llbracket N' \rrbracket = (\mathbb{S} \cup \mathbf{P}_{\mathbb{T}(N)}, \mathbb{T}(N), F, (b \setminus S) \cup b' \cup b_{T \cup T'})$. By the definition of $\llbracket _ \rrbracket$, $t \in \mathbb{T}(N)$ and F is such that $t : S \cup \{\mathbf{p}_t\} \rightarrow b' \cup b_{T'}$. Hence, t is enabled in $b \cup b_T$ because $S \subseteq b$ and $\mathbf{p}_t \in b_T$. Consequently, $b \cup b_T \xrightarrow{t} ((b \cup b_T) \setminus (S \cup \{\mathbf{p}_t\})) \cup b' \cup b_{T'} = (b \setminus S) \cup b_T \cup b' \cup b_{T'} = (b \setminus S) \cup b' \cup b_{T \cup T'}$.

The proof for $\llbracket N \rrbracket \xrightarrow{t} N'$ implies there exists N'' such that $N \xrightarrow{t} N''$ and $N' = \llbracket N'' \rrbracket$ follows by analogous arguments. \square

APPENDIX B. PROOFS OF RESULTS IN SECTION 3

This section presents the proofs of the results in Section 3. Note that we need some auxiliary lemmas that are not present in the main text of the paper. They are marked by the keyword ‘‘Aux’’ to avoid ambiguities. For reviewer’s convenience, the high-level proof sketches are separated from the proofs in full detail, that are included in a separate section.

We start by showing that the encoding of a net into a dynamic net does not add computations. We show that each reachable marking b of the dynamic net can be associated with a reachable marking m of the original net, when disregarding negative information. We remark that in general the relation between such b and m is that $b \cap P \subseteq m$ and not necessarily $b \cap P = m$ (see, e.g., Lemma B.1). This is because the transitions $t_{\theta, \mathbb{C}}$ generated by the encoding (T_{pos}) always consume the tokens in all minimal places of the branching cell \mathbb{C} . This choice is immaterial for the behaviour of the encoded net, as made explicit by the main results in the paper.

Lemma B.1 (Aux.). *Let $N = (P, T, F, m)$. If $\llbracket N \rrbracket \rightarrow^* (T, b)$ then $m \rightarrow^* m'$ and $b \cap P \subseteq m'$.*

Proof. The proof follows by induction on the length of the reduction $\llbracket N \rrbracket \rightarrow^n (T, b)$.

- **Base case ($n=0$).** It follows immediately because $b = m$.
- **Inductive case ($n = k+1$).** Then, $\llbracket N \rrbracket \rightarrow^k (T', b') \xrightarrow{t} (T, b)$. By inductive hypothesis, $m \rightarrow^* m''$ and $b' \cap P \subseteq m''$. We now proceed by case analysis on the shape of t .
 - $t = {}^\circ\mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ})$. Then, ${}^\circ\mathbb{C} \subseteq b'$, $T = T'$ and $b = (b' \setminus {}^\circ\mathbb{C}) \cup \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ}$. Since ${}^\circ\mathbb{C} \subseteq P$, we have ${}^\circ\mathbb{C} \subseteq m''$. Moreover, $\theta : \mathbb{C}$ implies ${}^\circ\theta \subseteq {}^\circ\mathbb{C} \subseteq m''$. Since θ is a deterministic process, $m'' \rightarrow^* (m'' \setminus {}^\circ\theta) \cup \theta^\circ$. Then, take $m' = (m'' \setminus {}^\circ\theta) \cup \theta^\circ$. Note that $b \cap P = ((b' \cap P) \setminus {}^\circ\mathbb{C}) \cup \theta^\circ$. We use $(b' \cap P) \subseteq m''$ and ${}^\circ\theta \subseteq {}^\circ\mathbb{C}$ to conclude that $b \cap P \subseteq m'$.
 - $t = \overline{p} \rightarrow (T', \overline{\mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ})$. It follows immediately because $b' \cap P = b \cap P$. \square

Corollary B.2 (Aux.). *If N is 1-safe then $\llbracket N \rrbracket$ is 1- ∞ -safe.*

Lemma B.3 (Lemma 3.11). *Let $\llbracket N \rrbracket \in \text{DN}(P \cup \overline{P})$. If $\llbracket N \rrbracket \rightarrow^* (T, b)$ then for all $t, t' \in T$ such that $\bullet t \neq \bullet t'$ and $\bullet t \cap \bullet t' \cap P \neq \emptyset$ it holds that there is $p \in P \cap (\bullet t \cup \bullet t')$ such that $\overline{p} \in b$.*

Proof. The proof follows by induction on the length of the firing sequence $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_n} (T, b)$.

- **Base Case $n = 0$.** It holds trivially because any pair of different transitions in T have either the same preset (i.e., if they are taken from T_{pos} and originate from the same s-cell) or disjoint presets (i.e., if they are taken both from T_{pos} but originate from different s-cells, or both from T_{neg} , or one from T_{neg} and the other from T_{pos}).
- **Inductive step $n = k + 1$.** Hence, $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_k} (T', b') \xrightarrow{t_{k+1}} (T, b)$. By inductive hypothesis, for all $t, t' \in T'$ such that $\bullet t \cap \bullet t' \cap P \neq \emptyset$, it holds that there is $p \in P \cap (\bullet t \cup \bullet t')$ such that $\overline{p} \in b'$. Then, we proceed by case analysis on t_{k+1} .
 - $t_{k+1} = {}^\circ\mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ})$. It holds trivially because $T = T'$ and $b' \cap \overline{P} \subseteq b$.
 - $t_{k+1} = \overline{p} \rightarrow (T'', \overline{\mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ})$ for some \mathbb{C} , $p \in {}^\circ\mathbb{C}$, and $(T'', \emptyset) = \llbracket N_{\mathbb{C}} \ominus p \rrbracket$. Then $T = T' \cup T''$. By the definition of $\llbracket _ \rrbracket$, we have that for all $t, t' \in T''$ either (i) $\bullet t = \bullet t'$ or (ii) $\bullet t \cap \bullet t' = \emptyset$ (reasoning analogously to the Base Case). It remains to consider the cases

in which t and t' are taken one from T' and the other from T'' . W.l.o.g., we consider $t \in T''$ and $t' \in T'$ and proceed as follows. By the definition of $\llbracket _ \rrbracket$, $t \in T''$ implies either (i) $\bullet t \subseteq \overline{\mathbf{P}}$ or (ii) $\bullet t = {}^\circ \mathbb{C}_1$ for $\mathbb{C}_1 \in \text{BC}(N_{\mathbb{C}} \ominus p)$. Case (i), follows immediately because there does not exist t' s.t. $\bullet t \cap \bullet t' \cap P \neq \emptyset$. For (ii), we note that $\bullet t' = {}^\circ \mathbb{C}_2$ with $\mathbb{C}_1 \neq \mathbb{C}_2$, ${}^\circ \mathbb{C}_1 \cap {}^\circ \mathbb{C}_2 \neq \emptyset$ and $\bullet t \cup \bullet t' \subseteq P$. We proceed by contradiction and assume $(\bullet t \cup \bullet t') \cap b = \emptyset$. There must exist an s-cell \mathbb{C}_3 such that $\mathbb{C}_1 \cup \mathbb{C}_2 \subseteq \mathbb{C}_3$ (because \mathbb{C}_1 and \mathbb{C}_2 are closed under immediate conflict and their union introduces immediate conflict between the transitions consuming from the shared places in ${}^\circ \mathbb{C}_1 \cap {}^\circ \mathbb{C}_2$). If $\mathbb{C}_2 = \mathbb{C}_3$, then $\mathbb{C}_1 \subset \mathbb{C}_2$. Hence $p \in {}^\circ \mathbb{C}_2$ and $p \in \bullet t'$, which contradicts $(\bullet t \cup \bullet t') \cap b = \emptyset$ because t_{k+1} enabled at b implies $\overline{\mathbf{p}} \in b$. Otherwise, $\mathbb{C}_2 \subset \mathbb{C}_3$. Consequently, there exists (at least) a transition $t'' \in T'$ such that $\bullet t'' = {}^\circ \mathbb{C}_3$ and $\bullet t' \neq \bullet t''$. Since $t' \in T'$ and $t'' \in T'$, we can use inductive hypothesis to conclude that $(\bullet t' \cup \bullet t'') \cap b \neq \emptyset$. The proof is completed by noting that this is in contradiction with the assumption $(\bullet t \cup \bullet t') \cap b = \emptyset$ because $\bullet t'' \supseteq \bullet t \cup \bullet t'$. \square

In what follows we write $p \prec q$ if $p \preceq q$ and $p \neq q$. The following auxiliary result provides some invariants about the configurations that can be reached by an encoded dynamic net.

Lemma B.4 (Aux.). *If $\llbracket N \rrbracket \rightarrow^* (T, b)$ then*

- (1) $p \in b$ implies $\overline{\mathbf{p}} \notin b$;
- (2) if $\overline{\mathbf{p}} \in b$ and $p \preceq q$ then $q \notin b$;
- (3) if $p \preceq q$, $p \in b$ and $\overline{\mathbf{q}} \in b$ then there exists $r \prec q$ and $\overline{\mathbf{r}} \in b$; and
- (4) if $(T, b) \xrightarrow{t}$ and $\bullet t = {}^\circ \mathbb{C}$ for some \mathbb{C} then $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b = \emptyset$.

Proof. The proof follows by induction on the length of the firing sequence $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_n} (T, b)$.

- **Base Case** $n = 0$. Hence $(T, b) = \llbracket N \rrbracket$.

- (1) It follows from $b \subseteq P$.
- (2) Since $b \subseteq P$ there is no $\overline{\mathbf{p}} \in b$.
- (3) Since $b \subseteq P$ there is no $\overline{\mathbf{q}} \in b$.
- (4) It follows from the fact that N is an occurrence net, $b \subseteq {}^\circ N$, and hence there does not exist any $\mathbb{C} \in \text{BC}(N)$ such that $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b \neq \emptyset$.

- **Inductive step** $n = k + 1$. Hence, $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_k} (T', b') \xrightarrow{t_{k+1}} (T, b)$. By inductive hypothesis, (1) $p \in b'$ implies $\overline{\mathbf{p}} \notin b'$; (2) if $\overline{\mathbf{p}} \in b'$ and $p \preceq q$ then $q \notin b'$; (3) if $p \preceq q$, $p \in b'$ and $\overline{\mathbf{q}} \in b'$ then there exists $r \prec q$ and $\overline{\mathbf{r}} \in b'$; and (4) if $(T', b') \xrightarrow{t}$ and $\bullet t = {}^\circ \mathbb{C}$ for some \mathbb{C} then $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b' = \emptyset$. We now proceed by case analysis on t_{k+1} .

- $t_{k+1} = {}^\circ \mathbb{C}_{k+1} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$ for some s-cell \mathbb{C}_{k+1} and transaction $\theta : \mathbb{C}_{k+1}$. Hence, $b = (b' \setminus {}^\circ \mathbb{C}_{k+1}) \cup (\theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$.

- (1) We proceed by contradiction. Assume that there exists p such that $p \in b$ and $\overline{\mathbf{p}} \in b$. Since $p \in b$ we have that either $p \in b' \setminus {}^\circ \mathbb{C}_{k+1}$ or $p \in \theta^\circ$. First, assume $p \in b' \setminus {}^\circ \mathbb{C}_{k+1}$. By inductive hypothesis (1), $\overline{\mathbf{p}} \notin b'$ and, hence, $\overline{\mathbf{p}} \notin b' \setminus {}^\circ \mathbb{C}_{k+1}$. Therefore, it should be the case that $\overline{\mathbf{p}} \in (\theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$. Hence, $\overline{\mathbf{p}} \in \overline{\mathbb{C}_{k+1}^\circ}$ and $p \in \mathbb{C}_{k+1}^\circ$. Since t_{k+1} is enabled at (T', b') , we can use inductive hypothesis (4) on t_{k+1} to conclude $(\mathbb{C}_{k+1}^\circ \cup \overline{\mathbb{C}_{k+1}^\circ}) \cap b' = \emptyset$. Consequently, $p \in \mathbb{C}_{k+1}^\circ$ implies $p \notin b'$. But this is in contradiction with the assumption that $p \in b' \setminus {}^\circ \mathbb{C}_{k+1}$. Assume instead $p \in \theta^\circ$. Then $\overline{\mathbf{p}} \notin \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ$. Hence, it should be the case that $\overline{\mathbf{p}} \in b' \setminus {}^\circ \mathbb{C}_{k+1}$. But this is also in contradiction with the hypothesis (4) $(\mathbb{C}_{k+1}^\circ \cup \overline{\mathbb{C}_{k+1}^\circ}) \cap b' = \emptyset$.

- (2) We proceed by contradiction. Assume there exist p and q such that $\bar{p} \in b$, $p \preceq q$ and $q \in b$.
- * Firstly, consider $\bar{p} \in b' \setminus \circ\mathbb{C}_{k+1}$, which implies $\bar{p} \in b'$. By inductive hypothesis (2), for all q s.t. $p \preceq q$ it holds that $q \notin b'$. Hence, it should be the case that $q \in (\theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$. Hence either (i) $p = q$, (ii) $p \in \circ\mathbb{C}_{k+1}$ or (iii) $p \prec p'$ and $p' \prec q$ for some $p' \in \circ\mathbb{C}_{k+1}$. For (i), note that $q = p \in (\theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$. Hence, $q \in \mathbb{C}_{k+1}^\circ$, which is in contradiction with the assumption $\bar{p} \in b'$ and the inductive hypothesis (4). For (ii), note that it implies $p \in b'$, which is in contradiction with the assumption $\bar{p} \in b'$ and inductive hypothesis (1). For (iii), note that $\bar{p} \in b'$ and $p \prec p'$ imply $p' \notin b'$ by inductive hypothesis (2), which is in contradiction with the fact that $p' \in \circ\mathbb{C}_{k+1}$ and t_{k+1} is enabled.
 - * Assume instead $\bar{p} \in \theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ$. Hence, $p \in \mathbb{C}_{k+1}^\circ \setminus \theta^\circ$ and therefore $p \in \mathbb{C}_{k+1}^\circ$. Suppose there is $q \in b'$ and $p \preceq q$. Note that $p' \preceq p$ for all $p' \in \circ\mathbb{C}_{k+1}$ by definition of branching cells. By transitivity of \preceq , $p' \preceq q$ for all $p' \in \circ\mathbb{C}_{k+1}$. Since t_{k+1} is enabled at b' , $\circ\mathbb{C}_{k+1} \subseteq b'$. By using Lemma B.1, we can conclude that $q \notin b'$ for all q s.t. $p \preceq q$, which contradicts the hypothesis $q \in b'$ and $p \preceq q$. Assume instead $q \in \theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ$. Hence, $q \in \theta^\circ$. Hence, $p \neq q$. Moreover, $p \in \mathbb{C}_{k+1}^\circ$ and $q \in \mathbb{C}_{k+1}^\circ$ contradict the assumption $p \preceq q$.
- (3) We proceed by case analysis. If $p \in b'$ and $\bar{q} \in b'$ then the proof follows by inductive hypothesis. If $p \in b'$ and $\bar{q} \notin b'$, then $\bar{q} \in \theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ$. Therefore, $p \preceq q$ implies $p \in \circ\mathbb{C}$, which contradicts the assumption $p \in b$. If $p \notin b'$ and $\bar{q} \in b'$, then $p \in \theta^\circ$, which contradicts $p \preceq q$.
- (4) Let $t \in T$ such that $\bullet t = \circ\mathbb{C} \subseteq b$ for some \mathbb{C} . Since t is enabled at b and $\llbracket N \rrbracket$ is 1- ∞ -safe by Corollary B.2, then $\mathbb{C}_{k+1} \cap \mathbb{C} = \emptyset$. If t is enabled at $(b' \setminus \circ\mathbb{C}_{k+1})$ then t is enabled at b' . By inductive hypothesis (4), we conclude that $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap (b' \setminus \circ\mathbb{C}_{k+1}) = \emptyset$. If t is not enabled at (T', b') , then it holds that for $x \in \circ\mathbb{C}$ exists $y \in (\mathbb{C}_{k+1}^\circ \cup \overline{\mathbb{C}_{k+1}^\circ})$ such that $y \preceq x$. By inductive hypothesis $(\mathbb{C}_{k+1}^\circ \cup \overline{\mathbb{C}_{k+1}^\circ}) \cap b' = \emptyset$, hence $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b' = \emptyset$. Therefore, $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b = \emptyset$.
- $t_{k+1} = \bar{r} \rightarrow (T'', \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ)$ for some s-cell \mathbb{C}_{k+1} and place $r \in \circ\mathbb{C}_{k+1}$. Then, $T = T' \cup T''$ with $\llbracket N_{\mathbb{C}_{k+1}} \ominus r \rrbracket = (T'', _)$ and $b = b' \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$.
- (1) We proceed by contradiction. Assume that there exists p such that $p \in b$ and $\bar{p} \in b$. Note that $p \in b$ implies $p \in b'$. By inductive hypothesis (1), $\bar{p} \notin b'$. Therefore, it should be the case that $\bar{p} \in \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$. Consequently $p \in \mathbb{C}_{k+1}^\circ$ and $p \notin (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$. Hence, $r \preceq p$. Since t is enabled at b' , $\bar{r} \in b'$. By inductive hypothesis (2), $p \notin b'$ which contradicts the hypothesis $p \in b$.
- (2) We proceed by contradiction. Assume there exist p and q such that $\bar{p} \in b$, $p \preceq q$ and $q \in b$. Note that $q \in b$ implies $q \in b'$. Assume $\bar{p} \in b'$. By inductive hypothesis, for all q s.t. $p \preceq q$ then $q \notin b'$ and, hence it is in contradiction with assumption $q \in b$. Assume instead $\bar{p} \in \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$. As before, we conclude that $r \preceq p$. By transitivity of \preceq , we have $r \preceq q$. By inductive hypothesis (2), $q \notin b'$, which is in contradiction with assumption $q \in b$.
- (3) For $\bar{q} \in b'$, it follows immediately by inductive hypothesis. For $\bar{q} \in \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$, it follows straightforwardly because $r \preceq q$ and $\bar{r} \in b$.
- (4) Assume $\bullet t = \circ\mathbb{C} \subseteq b$ for some \mathbb{C} . Hence, $\bullet t = \circ\mathbb{C} \subseteq b'$. There are two cases:

- * Suppose $t \in T'$. By inductive hypothesis (4), $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b' = \emptyset$. We show that the following holds

$$(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap \overline{\mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ} = \emptyset$$

It is enough to show that

$$\overline{\mathbb{C}^\circ \cap \mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ} = \emptyset$$

We proceed by contradiction and assume there exists q such that $q \in \mathbb{C}^\circ$ and $q \in (\mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ)$. Because $q \in (\mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ)$, $r \preceq q$. Since $q \in \mathbb{C}^\circ$, $r \in {}^\circ\mathbb{C}$ (because \mathbb{C} is closed under causality). Hence $r \in b'$ because t is enabled at b' . By the contrapositive of inductive hypothesis (1), $\bar{r} \notin b'$, but this is in contradiction with the hypothesis that t_{k+1} is enabled at b' .

- * Suppose $t \in T''$. Then, $\bullet t \cap {}^\circ(N_{\mathbb{C}_{k+1}} \ominus r) = \emptyset$. Hence, for all $q \in \mathbb{C}^\circ$ there exists $s \in {}^\circ\mathbb{C}$ s.t. $s \preceq q$. Since t is enabled at b , ${}^\circ\mathbb{C} \subseteq b$ holds. By Lemma B.1, $\mathbb{C}^\circ \cap b = \emptyset$. We show by contradiction that $\overline{\mathbb{C}^\circ} \cap b = \emptyset$ does not hold either. Assume that there exists $\bar{q} \in \overline{\mathbb{C}^\circ}$ and $\bar{q} \in b$. Since there exists $s \in {}^\circ\mathbb{C} \subseteq b$ and $s \preceq q$, we can use inductive hypothesis (3) to conclude that there exist $o \preceq q$ s.t. $\bar{o} \in b$. By the inductive hypothesis (2) $q \notin b$, and this is in contradiction with the assumption of t enabled at b . \square

Lemma B.5 (Lemma 3.10). *If $\llbracket N \rrbracket \rightarrow^* (T, b)$ and $\bar{p} \in b$ then $(T, b) \rightarrow^* (T', b')$ implies that $p \notin b'$.*

Proof. If $\bar{p} \in b$ then $\bar{p} \in b'$ because \bar{p} is persistent. Moreover, $\llbracket N \rrbracket \rightarrow^* (T', b')$. By the contrapositive of Lemma B.4(1), $p \notin b'$. \square

Theorem B.6 (Theorem 3.12). *Let $\llbracket N \rrbracket \in \text{DN}(P \cup \bar{P})$. If $\llbracket N \rrbracket \rightarrow^* (T, b) \xrightarrow{t}$ and $(T, b) \xrightarrow{t'}$ then either $\bullet t = \bullet t'$ or $\bullet t \cap \bullet t' = \emptyset$.*

Proof. By contradiction. Assume t, t' such that $(T, b) \xrightarrow{t}$, $(T, b) \xrightarrow{t'}$, $\bullet t \neq \bullet t'$, and $\bullet t \cap \bullet t' \neq \emptyset$. By construction of the encoding, it must be the case that $\bullet t \subseteq P$ and $\bullet t' \subseteq P$. Hence, $\bullet t \cap \bullet t' \cap P \neq \emptyset$. By Lemma 3.11 there exists $p \in P \cap (\bullet t \cup \bullet t')$ such that $\bar{p} \in b$. By Lemma 3.10, $p \notin b$, which is in contradiction with the assumptions $(T, b) \xrightarrow{t}$ and $(T, b) \xrightarrow{t'}$. \square

B.1. Detailed proofs of results in Section 3.

Lemma B.7 (Lemma B.1). *Let $N = (P, T, F, m)$. If $\llbracket N \rrbracket \rightarrow^* (T, b)$ then $m \rightarrow^* m'$ and $b \cap P \subseteq m'$.*

Proof. The proof follows by induction on the length of the reduction $\llbracket N \rrbracket \rightarrow^n (T, b)$.

- **Base case ($n=0$).** It follows immediately because $b = m$.
- **Inductive case ($n = k+1$).** Then, $\llbracket N \rrbracket \rightarrow^k (T', b') \xrightarrow{t} (T, b)$. By inductive hypothesis, $m \rightarrow^* m''$ and $b' \cap P \subseteq m''$. We now proceed by case analysis on the shape of t .
 - $t = {}^\circ\mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ})$. Then, ${}^\circ\mathbb{C} \subseteq b'$, $T = T'$ and $b = (b' \setminus {}^\circ\mathbb{C}) \cup \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ}$. Since ${}^\circ\mathbb{C} \subseteq P$, we have ${}^\circ\mathbb{C} \subseteq m''$. Moreover, $\theta : \mathbb{C}$ implies ${}^\circ\theta \subseteq {}^\circ\mathbb{C} \subseteq m''$. Since θ is a deterministic process, $m'' \rightarrow^* (m'' \setminus {}^\circ\theta) \cup \theta^\circ$. Then, take $m' = (m'' \setminus {}^\circ\theta) \cup \theta^\circ$. Note that $b \cap P = ((b' \cap P) \setminus {}^\circ\mathbb{C}) \cup \theta^\circ$. We use $(b' \cap P) \subseteq m''$ and ${}^\circ\theta \subseteq {}^\circ\mathbb{C}$ to conclude that $b \cap P \subseteq m'$.
 - $t = \bar{p} \rightarrow (T', \overline{\mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ})$. It follows immediately because $b' \cap P = b \cap P$. \square

Lemma B.8 (Lemma 3.11). *Let $\llbracket N \rrbracket \in \text{DN}(P \cup \bar{P})$. If $\llbracket N \rrbracket \rightarrow^* (T, b)$ then for all $t, t' \in T$ such that $\bullet t \neq \bullet t'$ and $\bullet t \cap \bullet t' \cap P \neq \emptyset$ it holds that there is $p \in P \cap (\bullet t \cup \bullet t')$ such that $\bar{p} \in b$.*

Proof. The proof follows by induction on the length of the firing sequence $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_n} (T, b)$.

- **Base Case** $n = 0$. It holds trivially because any pair of different transitions in T have either the same preset (i.e., if they are taken from T_{pos} and originate from the same s-cell) or disjoint presets (i.e., if they are taken both from T_{pos} but originate from different s-cells, or both from T_{neg} , or one from T_{neg} and the other from T_{pos}).
- **Inductive step** $n = k + 1$. Hence, $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_k} (T', b') \xrightarrow{t_{k+1}} (T, b)$. By inductive hypothesis, for all $t, t' \in T'$ such that $\bullet t \cap \bullet t' \cap P \neq \emptyset$, it holds that there is $p \in P \cap (\bullet t \cup \bullet t')$ such that $\bar{p} \in b'$. Then, we proceed by case analysis on t_{k+1} .
 - $t_{k+1} = \circ \mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \mathbb{C}^\circ \setminus \theta^\circ)$. It holds trivially because $T = T'$ and $b' \cap \bar{P} \subseteq b$.
 - $t_{k+1} = \bar{p} \rightarrow (T'', \overline{\mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ})$ for some $\mathbb{C}, p \in \circ \mathbb{C}$, and $(T'', \emptyset) = \llbracket N_{\mathbb{C}} \ominus p \rrbracket$. Then $T = T' \cup T''$. By the definition of $\llbracket - \rrbracket$, we have that for all $t, t' \in T''$ either (i) $\bullet t = \bullet t'$ or (ii) $\bullet t \cap \bullet t' = \emptyset$ (reasoning analogously to the Base Case). It remains to consider the cases in which t and t' are taken one from T' and the other from T'' . W.l.o.g., we consider $t \in T''$ and $t' \in T'$ and proceed as follows. Note that, by construction of $\llbracket N \rrbracket$, $\bullet t \subseteq \bar{P}$ implies $| \bullet t | = 1$ for any t . Hence, the only possibility is $\bullet t = \circ \mathbb{C}_1$ with $\mathbb{C}_1 \in \text{BC}(N_{\mathbb{C}} \ominus p)$ and $\bullet t' = \circ \mathbb{C}_2$ with $\mathbb{C}_1 \neq \mathbb{C}_2$ and $\circ \mathbb{C}_1 \cap \circ \mathbb{C}_2 \neq \emptyset$. Note that $\bullet t \cup \bullet t' \subseteq P$. We proceed by contradiction and assume $\overline{(\bullet t \cup \bullet t')} \cap b = \emptyset$. There must exist a s-cell \mathbb{C}_3 such that $\mathbb{C}_1 \cup \mathbb{C}_2 \subseteq \mathbb{C}_3$ (because \mathbb{C}_1 and \mathbb{C}_2 are closed under immediate conflict and their union introduces immediate conflict between the transitions consuming from the shared places in $\circ \mathbb{C}_1 \cap \circ \mathbb{C}_2$). If $\mathbb{C}_2 = \mathbb{C}_3$, then $\mathbb{C}_1 \subset \mathbb{C}_2$ and hence $p \in \circ \mathbb{C}_2$ and $p \in \bullet t'$, which contradicts $\overline{(\bullet t \cup \bullet t')} \cap b = \emptyset$ because t_{k+1} enabled at b implies $\bar{p} \in b$. Otherwise, $\mathbb{C}_2 \subset \mathbb{C}_3$. Consequently, there exists (at least) a transition $t'' \in T'$ such that $\bullet t'' = \circ \mathbb{C}_3$ and $\bullet t' \neq \bullet t''$. Since $t' \in T'$ and $t'' \in T'$, we can use inductive hypothesis to conclude that $\overline{(\bullet t' \cup \bullet t'')} \cap b \neq \emptyset$. The proof is completed by noting that this is in contradiction with the assumption $\overline{(\bullet t \cup \bullet t')} \cap b \neq \emptyset$ because $\bullet t'' = \bullet t \cup \bullet t'$. \square

Lemma B.9 (Lemma B.4). *If $\llbracket N \rrbracket \rightarrow^* (T, b)$ then*

- (1) $p \in b$ implies $\bar{p} \notin b$;
- (2) if $\bar{p} \in b$ and $p \preceq q$ then $q \notin b$;
- (3) if $p \preceq q$, $p \in b$ and $\bar{q} \in b$ then there exists $r \prec q$ and $\bar{r} \in b$; and
- (4) if $(T, b) \xrightarrow{t}$ and $\bullet t = \circ \mathbb{C}$ for some \mathbb{C} then $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b = \emptyset$.

Proof. The proof follows by induction on the length of the firing sequence $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_n} (T, b)$.

- **Base Case** $n = 0$. Hence $(T, b) = \llbracket N \rrbracket$.
 - (1) It follows from $b \subseteq P$.
 - (2) Since $b \subseteq P$ there is no $\bar{p} \in b$.
 - (3) Since $b \subseteq P$ there is no $\bar{q} \in b$.
 - (4) It follows from the fact that N is an occurrence net, $b \subseteq \circ N$, and hence there does not exist any $\mathbb{C} \in \text{BC}(N)$ such that $\mathbb{C}^\circ \cap b \neq \emptyset$.
- **Inductive step** $n = k + 1$. Hence, $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_k} (T', b') \xrightarrow{t_{k+1}} (T, b)$. By inductive hypothesis, (1) $p \in b'$ implies $\bar{p} \notin b'$; (2) if $\bar{p} \in b'$ and $p \preceq q$ then $q \notin b'$; (3) if $p \preceq q$, $p \in b'$ and $\bar{q} \in b'$ then there exists $r \prec q$ and $\bar{r} \in b'$; and (4) if $(T', b') \xrightarrow{t}$ and $\bullet t = \circ \mathbb{C}$ for some \mathbb{C} then $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b' = \emptyset$. We now proceed by case analysis on t_{k+1} .

- $t_{k+1} = {}^\circ\mathbb{C}_{k+1} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$ for some s-cell \mathbb{C}_{k+1} and transaction $\theta : \mathbb{C}_{k+1}$. Hence, $b = (b' \setminus {}^\circ\mathbb{C}_{k+1}) \cup (\theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$.
- (1) We proceed by contradiction. Assume that there exists p such that $p \in b$ and $\bar{p} \in b$. Since $p \in b$ we have that either $p \in b' \setminus {}^\circ\mathbb{C}_{k+1}$ or $p \in \theta^\circ$. First, assume $p \in b' \setminus {}^\circ\mathbb{C}_{k+1}$. By inductive hypothesis (1), $\bar{p} \notin b'$ and, hence, $\bar{p} \notin b' \setminus {}^\circ\mathbb{C}_{k+1}$. Therefore, it should be the case that $\bar{p} \in (\theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$. Hence, $\bar{p} \in \overline{\mathbb{C}_{k+1}^\circ}$ and $p \in \mathbb{C}_{k+1}^\circ$. Since t_{k+1} is enabled at (T', b') , we can use inductive hypothesis (4) on t_{k+1} to conclude $(\mathbb{C}_{k+1}^\circ \cup \overline{\mathbb{C}_{k+1}^\circ}) \cap b' = \emptyset$. Consequently, $p \in \mathbb{C}_{k+1}^\circ$ implies $p \notin b'$. But this is in contradiction with the assumption that $p \in b' \setminus {}^\circ\mathbb{C}_{k+1}$. Assume instead $p \in \theta^\circ$. Then $\bar{p} \notin \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ$. Hence, it should be the case that $\bar{p} \in b' \setminus {}^\circ\mathbb{C}_{k+1}$. But this is also in contradiction with the the hypothesis (4) $(\mathbb{C}_{k+1}^\circ \cup \overline{\mathbb{C}_{k+1}^\circ}) \cap b' = \emptyset$.
- (2) We proceed by contradiction. Assume there exist p and q such that $\bar{p} \in b$, $p \preceq q$ and $q \in b$. Assume $\bar{p} \in b'$. By inductive hypothesis (2), for all q s.t. $p \preceq q$ it holds that $q \notin b' \setminus {}^\circ\mathbb{C}_{k+1}$. Moreover, if $q \in (\theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ)$ implies $p' \preceq q$ for all $p' \in {}^\circ\mathbb{C}_{k+1}$ by definition of branching cells. Since t is enabled at b' , ${}^\circ\mathbb{C}_{k+1} \subseteq b'$ and hence $p \in b'$, but this is in contradiction with inductive hypothesis (1), i.e., $\bar{p} \in b'$ implies $p \notin b'$. Assume instead $\bar{p} \in \theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ$. Hence, $p \in \mathbb{C}_{k+1}^\circ \setminus \theta^\circ$ and $p \in \mathbb{C}_{k+1}^\circ$. Suppose there is $q \in b'$ and $p \preceq q$. Note that $p' \preceq p$ for all $p' \in {}^\circ\mathbb{C}_{k+1}$ by definition of branching cells. By transitivity of \preceq , $p' \preceq q$ for all $p' \in {}^\circ\mathbb{C}_{k+1}$. Since t is enabled at b' , ${}^\circ\mathbb{C}_{k+1} \subseteq b'$. By using Lemma B.1, we can conclude that $q \notin b'$ for all q s.t. $p \preceq q$, which contradicts the hypothesis $q \in b'$ and $p \preceq q$. Assume instead $q \in \theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ$. Hence, $q \in \theta^\circ$. Hence, $p \neq q$. Moreover, $p \in \mathbb{C}_{k+1}^\circ$ and $q \in \mathbb{C}_{k+1}^\circ$ contradict the hypothesis $p \preceq q$.
- (3) If $\bar{q} \in b'$ the proof follows by inductive hypothesis and by noting that $p \in \theta^\circ$ and $p \preceq q$ imply there exists $r \in {}^\circ\mathbb{C}$ and $r \preceq q$ (by transitivity of \preceq). If $\bar{q} \in \theta^\circ \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus \theta^\circ$, follows by contradiction because $p \preceq q$ and $p \in b'$ implies $p \in {}^\circ\mathbb{C}$ by Lemma B.1. Therefore, there does not exist p such that $p \preceq q$ and $p \in b$.
- (4) Let $t \in T$ such that $\bullet t = {}^\circ\mathbb{C} \subseteq b$ for some \mathbb{C} . Since t is enabled at b and $\llbracket N \rrbracket$ is 1- ∞ -safe by Corollary B.2, then $\mathbb{C}_{k+1} \cap \mathbb{C} = \emptyset$. If t is enabled at $(b' \setminus {}^\circ\mathbb{C}_{k+1})$ then t is enabled at b' . By inductive hypothesis (2), we conclude that $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap (b' \setminus {}^\circ\mathbb{C}_{k+1}) = \emptyset$. If t is not enabled at (T', b') , then it holds that for $x \in \mathbb{C}$ exists $y \in (\mathbb{C}_{k+1}^\circ \cup \overline{\mathbb{C}_{k+1}^\circ})$ such that $y \preceq x$. By inductive hypothesis $(\mathbb{C}_{k+1}^\circ \cup \overline{\mathbb{C}_{k+1}^\circ}) \cap b' = \emptyset$, hence $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b' = \emptyset$. Therefore, $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b = \emptyset$.
- $t_{k+1} = \bar{r} \rightarrow (T'', \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ)$ for some s-cell \mathbb{C}_{k+1} and place $r \in {}^\circ\mathbb{C}_{k+1}$. Then, $T = T' \cup T''$ with $\llbracket N_{\mathbb{C}_{k+1}} \ominus r \rrbracket = (T'', -)$ and $b = b' \cup \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$.
- (1) We proceed by contradiction. Assume that there exists p such that $p \in b$ and $\bar{p} \in b$. Note that $p \in b$ implies $p \in b'$. By inductive hypothesis (1), $\bar{p} \notin b'$. Therefore, it should be the case that $\bar{p} \in \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$. Consequently $p \in \mathbb{C}_{k+1}^\circ$ and $p \notin (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$. Hence, $r \preceq p$. Since t is enabled at b' , $\bar{r} \in b'$. By inductive hypothesis (2), $p \notin b'$ which contradicts the hypothesis $p \in b'$.
- (2) We proceed by contradiction. Assume there exist p and q such that $\bar{p} \in b$, $p \preceq q$ and $q \in b$. Note that $q \in b$ implies $q \in b'$. Assume $\bar{p} \in b'$. By inductive hypothesis, for all q s.t. $p \preceq q$ then $q \notin b'$ and, hence it is in contradiction with assumption $q \in b$. Assume instead $\bar{p} \in \overline{\mathbb{C}_{k+1}^\circ} \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ$. As before, we conclude that $r \preceq p$.

By transitivity of \preceq , we have $r \preceq q$. By inductive hypothesis (2), $q \notin b'$, which is in contradiction with assumption $q \in b$.

- (3) For $\bar{q} \in b$, it follows immediately by inductive hypothesis. For $\bar{q} \in \overline{\mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ}$, it follows straightforwardly because $r \preceq q$ and $\bar{r} \in b$.
- (4) Assume $\bullet t = {}^\circ \mathbb{C} \subseteq b$ for some \mathbb{C} . Hence, $\bullet t = {}^\circ \mathbb{C} \subseteq b'$. There are two cases:
- * Suppose $t \in T'$. By inductive hypothesis (4), $(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap b' = \emptyset$. We show that the following holds

$$(\mathbb{C}^\circ \cup \overline{\mathbb{C}^\circ}) \cap \overline{\mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ} = \emptyset$$

It is enough to show that

$$\overline{\mathbb{C}^\circ \cap \overline{\mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ}} = \emptyset$$

We proceed by contradiction and assume there exists q such that $q \in \mathbb{C}^\circ$ and $q \in (\mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ)$. Because $q \in (\mathbb{C}_{k+1}^\circ \setminus (N_{\mathbb{C}_{k+1}} \ominus r)^\circ)$, $r \preceq q$. Since $q \in \mathbb{C}^\circ$, $r \in {}^\circ \mathbb{C}$ (because \mathbb{C} is closed under causality). Hence $r \in b'$ because t is enabled at b' . By the contrapositive of inductive hypothesis (1), $\bar{r} \notin b'$, but this is in contradiction with the hypothesis that t_{k+1} is enabled at b' .

- * Suppose $t \in T''$. Then, $\bullet t \cap {}^\circ (N_{\mathbb{C}_{k+1}} \ominus r) = \emptyset$ for some \mathbb{C} . Hence, for all $q \in \mathbb{C}^\circ$ there exists $s \in {}^\circ \mathbb{C}$ s.t. $s \preceq q$. Since t is enabled at b , ${}^\circ \mathbb{C} \subseteq b$ holds. By Lemma B.1, $\mathbb{C}^\circ \cap b = \emptyset$. We show by contradiction that $\overline{\mathbb{C}^\circ} \cap b = \emptyset$ does not hold either. Assume that there exists $\bar{q} \in \overline{\mathbb{C}^\circ}$ and $\bar{q} \in b$. Since there exists $s \in \mathbb{C} \subseteq b$ and $s \preceq q$, we can use inductive hypothesis (3) to conclude that there exist $\bar{s}' \subseteq b$ and $s' \preceq q$ and $s' \in {}^\circ \mathbb{C}$. By the inductive hypothesis (1) $s' \notin b$, and this is in contradiction with the assumption of t enabled at b . qedhere

□

APPENDIX C. PROOFS OF RESULTS IN SECTION 4

This section presents the proof sketches of the results in Section 4. As in Appendix A, we exploit some auxiliary lemmas marked by the keyword ‘‘Aux’’ and full proofs are provided separately. We start by showing that reductions of an encoded net correspond to recursively stopped configurations of the event structure.

Lemma C.1 (Aux.). *Let $N = (P, T, F, m)$ and \mathcal{E} the event structure of N . If $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_n} (T, b)$ and $v = \bigcup_{1 \leq i \leq n} \|t_i\|$, then*

- (1) $b \cap P = {}^\circ \{e \mid e \in \mathcal{E}^v \text{ and } [e] = \{e\}\}$; and
- (2) If $(T, b) \xrightarrow{t}$ then $\|t\| \neq \emptyset$ implies $\|t\|$ is a stopped configuration of \mathcal{E}^v .

Proof. It follows by induction on the length of $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_n} (T, b)$.

- **Base case (n=0).** Then, $v = \emptyset$ and $\mathcal{E}^v = \mathcal{E}$. Moreover, $b = m$.

- (1) It is immediate to notice that m corresponds to the preset of all minimal events of \mathcal{E} .
- (2) Since t is enabled, $\bullet t \subseteq m$. Hence, $\bullet t = {}^\circ \mathbb{C}$ with $\mathbb{C} \in \text{BC}(N)$. Therefore, \mathbb{C} corresponds to a branching cell of \mathcal{E} . By the definition of $\llbracket _ \rrbracket$, t is associated with some $\theta : \mathbb{C}$, which is a maximal, conflict-free set of transitions in \mathbb{C} . Hence, $\|t\|$ is a stopped configuration of \mathcal{E} .

- **Inductive case ($n = k+1$).** Then, $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_k} (T_k, b_k) \xrightarrow{t_{k+1}} (T, b)$. By inductive hypothesis, letting $v_k = \bigcup_{1 \leq i \leq k} \|t_i\|$, we assume (1) $b_k \cap P = \circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } [e] = \{e\}\}$, and (2) If $(T_k, b_k) \xrightarrow{t}$ then $\|t\| \neq \emptyset$ implies $\|t\|$ is a stopped configuration of \mathcal{E}^{v_k} .

We now proceed by case analysis on the shape of the applied rule:

- $t_{k+1} = \circ\mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ})$. Hence, $v = v_k \cup \|\theta\|$ and $b \cap P = (b_k \cap P \setminus \circ\mathbb{C}) \cup \theta^\circ$.

(1) Then:

$$\begin{aligned} & \{e \mid e \in \mathcal{E}^v \text{ and } [e] = \{e\}\} \\ = & \{e \mid e \in \mathcal{E}^{v_k} \text{ and } e \notin \mathbb{C} \text{ and } [e] = \{e\}\} \\ & \cup \{e \mid e \in \mathcal{E}^{v_k} \text{ and } [e] \subseteq \{e\} \cup \|\theta\|\} \end{aligned}$$

The proof is completed by noting that

$$\circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } e \notin \mathbb{C} \text{ and } [e] = \{e\}\} = (b_k \cap P \setminus \circ\mathbb{C})$$

and

$$\circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } [e] \subseteq \{e\} \cup \|\theta\|\} = \circ\theta$$

- (2) Take t such that $\bullet t = \mathbb{C}_t$. Then, $\mathbb{C}_t \subseteq b \cap P$. By Theorem 3.12, there cannot be t' enabled at b and $\bullet t' \neq \circ\mathbb{C}_t$ and $\circ\mathbb{C}_t \cap \bullet t' \neq \emptyset$. By using inductive hypothesis (1), we conclude that all events in direct conflict with \mathbb{C}_t in \mathcal{E}^v are in \mathbb{C} . Hence, $\|\theta\|$ is a stopped configuration of \mathcal{E}^v .

- $t_{k+1} = \overline{\mathbf{p}} \rightarrow (T'', \overline{\mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ})$ for some \mathbb{C} , $p \in \circ\mathbb{C}$, and $(T'', \emptyset) = \llbracket N_{\mathbb{C}} \ominus p \rrbracket$. Then $T = T' \cup T''$.

(1) Immediate because $b_k \cap P = b \cap P$.

(2) It follows analogously to the previous case. □

Lemma C.2 (Aux.). Let $\llbracket N \rrbracket \in \text{DN}(P \cup \overline{\mathbf{P}})$. If $\llbracket N \rrbracket \rightarrow^* (T, b)$ then there exists (T', b') such that $(T, b) \implies (T', b')$ and

- (1) $b' \cap P = b \cap P$;
- (2) for all p, q , if $\overline{\mathbf{p}} \in b$ and $p \preceq q$, then $\overline{\mathbf{q}} \in b'$;
- (3) for all $\mathbb{C} \in \text{BC}(N)$ and $\overline{\mathbf{Q}} \subseteq \overline{\mathbf{P}}$, if $\overline{\mathbf{Q}} \subseteq b'$ then for all $\mathbb{C}' \in \text{BC}(N_{\mathbb{C}} \ominus \mathbf{Q})$ and $\theta : \mathbb{C}'$ there exists $t \in T'$ such that $t = \circ\mathbb{C}' \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}'^\circ \setminus \theta^\circ})$.

Proof. (1) It follows straightforwardly by analysis of the applied rules. They are of the form

$$\overline{\mathbf{p}} \rightarrow (T', \overline{\mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ}),$$

- (2) By induction on the length of the chain $p = p_0 \prec \dots \prec p_n = q$ (this is a finite chain because N is a finite occurrence net). The inductive step follows by straightforward inspection of the shape of the transitions with negative premises.

- (3) By straightforward induction on the number n of elements in $\overline{\mathbf{Q}}$, i.e., $n = |\overline{\mathbf{Q}}|$. □

Lemma C.3 (Aux.). Let $N = (P, T, F, m)$ and \mathcal{E} the event structure of N . If v is recursively stopped configuration and $v = \bigcup_{1 \leq i \leq n} v_i$ is a valid decomposition, then

- (1) $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_n} (T, b)$ and $v = \bigcup_{1 \leq i \leq n} \|t_i\|$;
- (2) $b \cap P = \circ\{e \mid e \in \mathcal{E}^v \text{ and } [e] = \{e\}\}$;
- (3) If v' is a stopped configuration of \mathcal{E}^v , then there exists $t \in T$ s.t. $(T, b) \xrightarrow{t}$ and $\|t\| = v'$;
- (4) For all $e \in \mathcal{E}$, if $e \notin (\mathcal{E}^v \cup v)$ implies $\bullet e \cap b = 0$.

Proof. It follows by induction on the length n of the decomposition $v = \bigcup_{1 \leq i \leq n} v_i$.

- **Base case ($n=0$).** Then, $v = \emptyset$ and $\mathcal{E}^v = \mathcal{E}$. Moreover, $b = m$. Then
 - (1) It is immediate because $(T, b) = \llbracket N \rrbracket$ and $m = b$.
 - (2) Since $b = m$, b corresponds to the preset of all minimal events of $\mathcal{E}^\emptyset = \mathcal{E}$.
 - (3) If v' is a stopped configuration of \mathcal{E} , then there exists $\mathbb{C} \in \text{BC}(N)$ such that $v' \subseteq \mathbb{C}$. Since v' is a maximal configuration, there exists $\theta : \mathbb{C}$ such that $ev(\theta) = v'$. Hence, there exists $t \in T$ such that $\|t\| = v'$. Since, v' is part of an initial prefix, $\bullet t = \circ \mathbb{C} \subseteq m$. Hence, t is enabled.
 - (4) It trivially holds because there does not exist $e \in \mathcal{E}$ and $e \notin (\mathcal{E}^v \cup v)$.
- **Inductive case ($n = k+1$).** Take $v' = \bigcup_{1 \leq i \leq k+1} v_i$ and $v = v_{k+1} \cup v'$. Then, (1) $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_k} (T_k, b_k)$ and $v' = \bigcup_{1 \leq i \leq k} \|t_i\|$; and (2) $b_k \cap P = \circ\{e \mid e \in \mathcal{E}^{v'} \text{ and } [e] = \{e\}\}$; and (3) If v'' is a stopped configuration of \mathcal{E}^{v_k} , then there exists $t \in T$ s.t. $(T_k, b_k) \xrightarrow{t}$ and $\|t\| = v''$; and (4) For all $e \in \mathcal{E}$, if $e \notin (\mathcal{E}^{v'} \cup v')$ implies $\bullet e \cap b = 0$.

By inductive hypothesis (3), there exists t_{k+1} such that $\|t_{k+1}\| = v_{t_{k+1}}$ and $(T_k, b_k) \xrightarrow{t_{k+1}}$. Then, take $(T_k, b_k) \xrightarrow{t} (T', b')$. By using Lemma C.2, we conclude that there exists

(T_{k+1}, b_{k+1}) such that $(T_k, b_k) \xrightarrow{t_{k+1}} (T_{k+1}, b_{k+1})$ where:

- (a): $b_{k+1} \cap P = b_k \cap P$;
- (b): for all p, q , if $\bar{p} \in b_k$ and $p \preceq q$, then $\bar{q} \in b_{k+1}$;
- (c): for all $\mathbb{C} \in \text{BC}(N)$ and $\bar{\mathbb{Q}} \subseteq \bar{\mathbb{P}}$, if $\bar{\mathbb{Q}} \subseteq b_{k+1}$ then for all $\mathbb{C}' \in \text{BC}(N_{\mathbb{C}} \ominus \mathbb{Q})$ and $\theta : \mathbb{C}'$ there exists $t \in T_{k+1}$ such that $t = \circ \mathbb{C}' \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}'^\circ \setminus \theta^\circ})$.

Then,

- (1) It follows immediately because $\|t_{k+1}\| = v_{t_{k+1}}$;
- (2) Then, $t_{k+1} = \circ \mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}'^\circ \setminus \theta^\circ})$. Moreover, $b' \cap P = (b_k \cap P \setminus \circ \mathbb{C}) \cup \theta^\circ$. Hence,

$$\begin{aligned} & \{e \mid e \in \mathcal{E}^v \text{ and } [e] = \{e\}\} \\ &= \{e \mid e \in \mathcal{E}^{v'} \text{ and } e \notin \mathbb{C} \text{ and } [e] = \{e\}\} \\ & \quad \cup \{e \mid e \in \mathcal{E}^{v'} \text{ and } [e] \subseteq \{e\} \cup \|\theta\|\} \end{aligned}$$

The proof is completed by noting that

$$\circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } e \notin \mathbb{C} \text{ and } [e] = \{e\}\} = (b_k \cap P \setminus \circ \mathbb{C})$$

and $\circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } [e] \subseteq \{e\} \cup \|\theta\|\} = \circ \theta$ and by using (a) above.

- (3) It follows from (c).
- (4) It follows from (b). □

Lemma C.4 (Lemma 4.3). *Let N be an occurrence net.*

- (1) If $\llbracket N \rrbracket \xrightarrow{t_0 \dots t_n}$, then $v = \bigcup_{1 \leq i \leq n} \|t_i\|$ is recursively-stopped in \mathcal{E}_N and $(\|t_i\|)_{1 \leq i \leq n}$ is a valid decomposition of v .
- (2) If v is recursively-stopped in \mathcal{E}_N , then for any valid decomposition $(v_i)_{1 \leq i \leq n}$ there exists $\llbracket N \rrbracket \xrightarrow{t_0 \dots t_n}$ such that $\|t_i\| = v_i$.

Proof. 1). It follows from Lemma C.1. 2). It follows from Lemma C.3(1). □

Theorem C.5 (Theorem 4.5). *Let N be an occurrence net.*

- (1) If $(\llbracket N \rrbracket) \xrightarrow{t_1 \dots t_n}$, then $v = \bigcup_{1 \leq i \leq n} \|t_i\|$ is recursively-stopped in \mathcal{E}_N and $(\|t_i\|)_{1 \leq i \leq n}$ is a valid decomposition of v .
- (2) If v is recursively-stopped in \mathcal{E}_N , then for any valid decomposition $(v_i)_{1 \leq i \leq n}$ there exists $(\llbracket N \rrbracket) \xrightarrow{t_1 \dots t_n}$ such that $\|t_i\| = v_i$.

Proof. It follows from Lemma 4.3 and Proposition 2.4. \square

C.1. Detailed proofs of results in Section 4. This section is devoted to prove the main results in Section 4. We start by providing some auxiliary results.

Lemma C.6 (Lemma C.1). *Let $N = (P, T, F, m)$ and \mathcal{E} the event structure of N . If $\llbracket N \rrbracket \xrightarrow{t_1 \cdots t_n} (T, b)$ and $v = \bigcup_{1 \leq i \leq n} \|t_i\|$, then*

- (1) $b \cap P = \circ\{e \mid e \in \mathcal{E}^v \text{ and } \lfloor e \rfloor = \{e\}\}$; and
- (2) If $(T, b) \xrightarrow{t}$ then $\|t\| \neq \emptyset$ implies $\|t\|$ is a stopped configuration of \mathcal{E}^v .

Proof. It follows by induction on the length of the reduction $\llbracket N \rrbracket \xrightarrow{t_1 \cdots t_n} (T, b)$.

- **Base case ($n=0$).** Then, $v = \emptyset$ and $\mathcal{E}^v = \mathcal{E}$. Moreover, $b = m$.
 - (1) It is immediate to notice that m corresponds to the preset of all minimal events of \mathcal{E} .
 - (2) Since t is enabled, $\bullet t \subseteq m$. Hence, $\bullet t = \circ \mathbb{C}$ with $\mathbb{C} \in \text{BC}(N)$. Therefore, \mathbb{C} corresponds to a branching cell of \mathcal{E} . By the definition of $\llbracket \cdot \rrbracket$, t is associated with some $\theta : \mathbb{C}$, which is a maximal, conflict-free set of transitions in \mathbb{C} . Hence, $\|t\|$ is a stopped configuration of \mathcal{E} .
- **Inductive case ($n = k+1$).** Then, $\llbracket N \rrbracket \xrightarrow{t_1 \cdots t_k} (T_k, b_k) \xrightarrow{t_{k+1}} (T, b)$. By inductive hypothesis, letting $v_k = \bigcup_{1 \leq i \leq k} \|t_i\|$, we assume (1) $b_k \cap P = \circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } \lfloor e \rfloor = \{e\}\}$, and (2) If $(T_k, b_k) \xrightarrow{t}$ then $\|t\| \neq \emptyset$ implies $\|t\|$ is a stopped configuration of \mathcal{E}^{v_k} .

We now proceed by case analysis on the shape of the applied rule:

- $t_{k+1} = \circ \mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}^\circ \setminus \theta^\circ})$. Hence, $v = v_k \cup \|\theta\|$ and $b \cap P = (b_k \cap P \setminus \circ \mathbb{C}) \cup \theta^\circ$.
 - (1) Then:
$$\begin{aligned} & \{e \mid e \in \mathcal{E}^v \text{ and } \lfloor e \rfloor = \{e\}\} \\ &= \{e \mid e \in \mathcal{E}^{v_k} \text{ and } e \notin \mathbb{C} \text{ and } \lfloor e \rfloor = \{e\}\} \\ & \quad \cup \{e \mid e \in \mathcal{E}^{v_k} \text{ and } \lfloor e \rfloor \subseteq \{e\} \cup \|\theta\|\} \end{aligned}$$

The proof is completed by noting that

$$\circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } e \notin \mathbb{C} \text{ and } \lfloor e \rfloor = \{e\}\} = (b_k \cap P \setminus \circ \mathbb{C})$$

and

$$\circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } \lfloor e \rfloor \subseteq \{e\} \cup \|\theta\|\} = \circ \theta$$

- (2) Take t such that $\bullet t = \mathbb{C}_t$. Then, $\mathbb{C}_t \subseteq b \cap P$. By Theorem 3.12, there cannot be t' enabled at b and $\bullet t' \neq \circ \mathbb{C}_t$ and $\circ \mathbb{C}_t \cap \bullet t' \neq \emptyset$. By using inductive hypothesis (1), we conclude that all events in direct conflict with \mathbb{C}_t in \mathcal{E}^v are in \mathbb{C} . Hence, $\|\theta\|$ is a stopped configuration of \mathcal{E}^v .
- $t_{k+1} = \overline{p} \rightarrow (T'', \mathbb{C}^\circ \setminus (N_{\mathbb{C}} \ominus p)^\circ)$ for some \mathbb{C} , $p \in \circ \mathbb{C}$, and $(T'', \emptyset) = \llbracket N_{\mathbb{C}} \ominus p \rrbracket$. Then $T = T' \cup T''$.
 - (1) Immediate because $b_k \cap P = b \cap P$.
 - (2) It follows analogously to the previous case.

\square

Lemma C.7 (Lemma C.3). *Let $N = (P, T, F, m)$ and \mathcal{E} the event structure of N . If v is recursively stopped configuration and $v = \bigcup_{1 \leq i \leq n} v_i$ is a valid decomposition, then*

- (1) $\llbracket N \rrbracket \xrightarrow{t_1 \cdots t_n} (T, b)$ and $v = \bigcup_{1 \leq i \leq n} \|t_i\|$;

- (2) $b \cap P = \circ\{e \mid e \in \mathcal{E}^v \text{ and } [e] = \{e\}\}$;
(3) If v' is a stopped configuration of \mathcal{E}^v , then there exists $t \in T$ s.t. $(T, b) \xrightarrow{t}$ and $\|t\| = v'$;
(4) For all $e \in \mathcal{E}$, if $e \notin (\mathcal{E}^v \cup v)$ implies $\bullet e \cap b = 0$.

Proof. It follows by induction on the length n of the decomposition $v = \bigcup_{1 \leq i \leq n} v_i$.

- **Base case ($n=0$).** Then, $v = \emptyset$ and $\mathcal{E}^v = \mathcal{E}$. Moreover, $b = m$. Then
 - (1) It is immediate because $(T, b) = \llbracket N \rrbracket$ and $m = b$.
 - (2) Since $b = m$, b corresponds to the preset of all minimal events of $\mathcal{E}^\emptyset = \mathcal{E}$.
 - (3) If v' is a stopped configuration of \mathcal{E} , then there exists $\mathbb{C} \in \text{BC}(N)$ such that $v' \subseteq \mathbb{C}$. Since v' is a maximal configuration, there exists $\theta : \mathbb{C}$ such that $ev(\theta) = v'$. Hence, there exists $t \in T$ such that $\|t\| = v'$. Since, v' is part of an initial prefix, $\bullet t = \circ \mathbb{C} \subseteq m$. Hence, t is enabled.
 - (4) It trivially holds because there does not exist $e \in \mathcal{E}$ and $e \notin (\mathcal{E}^v \cup v)$.
- **Inductive case ($n = k+1$).** Take $v' = \bigcup_{1 \leq i \leq k+1} v_i$ and $v = v_{k+1} \cup v'$. Then, (1) $\llbracket N \rrbracket \xrightarrow{t_1 \dots t_k} (T_k, b_k)$ and $v' = \bigcup_{1 \leq i \leq k} \|t_i\|$; and (2) $b_k \cap P = \circ\{e \mid e \in \mathcal{E}^{v'} \text{ and } [e] = \{e\}\}$; and (3) If v'' is a stopped configuration of \mathcal{E}^{v_k} , then there exists $t \in T$ s.t. $(T_k, b_k) \xrightarrow{t}$ and $\|t\| = v''$; and (4) For all $e \in \mathcal{E}$, if $e \notin (\mathcal{E}^{v'} \cup v')$ implies $\bullet e \cap b = 0$.

By inductive hypothesis (3), there exists t_{k+1} such that $\|t_{k+1}\| = v_{t_{k+1}}$ and $(T_k, b_k) \xrightarrow{t_{k+1}}$. Then, take $(T_k, b_k) \xrightarrow{t} (T', b')$. By using Lemma C.2, we conclude that there exists (T_{k+1}, b_{k+1}) such that $(T_k, b_k) \xrightarrow{t_{k+1}} (T_{k+1}, b_{k+1})$ where:

- (a): $b_{k+1} \cap P = b_k \cap P$;
- (b): for all p, q , if $\bar{p} \in b_k$ and $p \preceq q$, then $\bar{q} \in b_{k+1}$;
- (c): for all $\mathbb{C} \in \text{BC}(N)$ and $\bar{\mathbb{Q}} \subseteq \bar{\mathbb{P}}$, if $\bar{\mathbb{Q}} \subseteq b_{k+1}$ then for all $\mathbb{C}' \in \text{BC}(N_{\mathbb{C}} \ominus \mathbb{Q})$ and $\theta : \mathbb{C}'$ there exists $t \in T_{k+1}$ such that $t = \circ \mathbb{C}' \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}'^\circ} \setminus \theta^\circ)$.

Then,

- (1) It follows immediately because $\|t_{k+1}\| = v_{t_{k+1}}$;
- (2) Then, $t_{k+1} = \circ \mathbb{C} \rightarrow (\emptyset, \theta^\circ \cup \overline{\mathbb{C}'^\circ} \setminus \theta^\circ)$. Moreover, $b' \cap P = (b_k \cap P \setminus \circ \mathbb{C}) \cup \theta^\circ$. Hence,

$$\begin{aligned} & \{e \mid e \in \mathcal{E}^v \text{ and } [e] = \{e\}\} \\ &= \{e \mid e \in \mathcal{E}^{v'} \text{ and } e \notin \mathbb{C} \text{ and } [e] = \{e\}\} \\ & \quad \cup \{e \mid e \in \mathcal{E}^{v'} \text{ and } [e] \subseteq \{e\} \cup \|\theta\|\} \end{aligned}$$

The proof is completed by noting that

$$\circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } e \notin \mathbb{C} \text{ and } [e] = \{e\}\} = (b_k \cap P \setminus \circ \mathbb{C})$$

and $\circ\{e \mid e \in \mathcal{E}^{v_k} \text{ and } [e] \subseteq \{e\} \cup \|\theta\|\} = \circ \theta$ and by using (a) above.

- (3) It follows from (c).
- (4) It follows from (b). □

APPENDIX D. PROOFS OF RESULTS IN SECTION 5

Theorem D.1 (Theorem 5.3). *Let $\sigma = t_1; \dots; t_n$ with $n \geq 0$ be a, possibly empty, firing sequence of a persistent process, and t a transition not in σ . The following conditions are all equivalent: (i) t is enabled after σ ; (ii) there is a collection of causes of t which appears in σ ; (iii) $\bigwedge_{i=1}^n t_i$ implies $\Phi(t)$.*

Proof. **ii) \Leftrightarrow (iii)::** We have that $\bigwedge_{i=1}^n t_i$ implies $\Phi(t)$ iff there is a prime implicant $\bigwedge_{j=1}^m t_{i_j}$ of $\Phi(t)$ that is implied by $\bigwedge_{i=1}^n t_i$. This is the case iff the collection of causes $\{t_{i_1}, \dots, t_{i_m}\}$ appears in σ .

(i) \Rightarrow (iii):: The proof is by induction on the length n of the sequence.

For the base case, if $n = 0$ it means that t is enabled in the initial marking, i.e., that its pre-set only contains initial places of the process and thus $\Phi(t) = \text{true}$.

For the inductive case, assume the property holds for any shorter sequence $t_1; \dots; t_k$ with $0 \leq k < n + 1$ and let us prove that it holds for $\sigma = t_1; \dots; t_{n+1}$. Let b_0 the initial bag of the process. As t is enabled after σ we have $b_0 \xrightarrow{\sigma} b \xrightarrow{t}$ for some bag b . Since t is enabled in b , we have $\bullet t \subseteq b$, i.e., for any $s \in \bullet t$ we have $b(s) \in \{1, \infty\}$ (by definition of p-net, $\bullet t$ is not empty). We need to prove that $\Phi(t) = \bigwedge_{s \in \bullet t} \Phi(s)$ is implied by $\bigwedge_{i=1}^{n+1} t_i$, i.e., that for any $s \in \bullet t$ the formula $\Phi(s)$ is implied by $\bigwedge_{i=1}^{n+1} t_i$. Take a generic $s \in \bullet t$. Either $\bullet s = \emptyset$, in which case s is initial and $\Phi(s) = \text{true}$, or $\bullet s \neq \emptyset$ and $\Phi(s) = \bigvee_{t' \in \bullet s} (t' \wedge \Phi(t'))$. Since $b(s) \in \{1, \infty\}$, there must exist an index $j \in [1, n + 1]$ such that $t_j \in \bullet s$. Take $t' = t_j$. Since σ is a firing sequence, the transition t_j is enabled after $\sigma' = t_1; \dots; t_{j-1}$. As $k = j - 1 < n + 1$, by inductive hypothesis $\Phi(t_j)$ is implied by $\bigwedge_{i=1}^{j-1} t_i$ and thus also by $\bigwedge_{i=1}^{n+1} t_i$. Since $\bigwedge_{i=1}^{n+1} t_i$ clearly implies t_j we have that $\bigwedge_{i=1}^{n+1} t_i$ implies $\Phi(s) = t_j \wedge \Phi(t_j)$.

(iii) \Rightarrow (i):: Suppose $\bigwedge_{i=1}^n t_i$ implies $\Phi(t) = \bigwedge_{s \in \bullet t} \Phi(s)$. If for all $s \in \bullet t$ we have $\bullet s = \emptyset$, then t is enabled in the initial marking and as the process is deterministic no transition can steal tokens from $\bullet t$ and t remains enabled after the firing of any $\sigma = t_1; \dots; t_n$. Otherwise, $\Phi(t) = \bigwedge_{s \in \bullet t} \bigvee_{t' \in \bullet s \neq \emptyset} (t' \wedge \Phi(t'))$. Thus, for any $s \in \bullet t$ with $\bullet s \neq \emptyset$ there exists some $t' \in \bullet s$ such that $\bigwedge_{i=1}^n t_i$ implies $t' \wedge \Phi(t')$. Since $\bigwedge_{i=1}^n t_i$ implies t' then there exists some index $k \in [1, n]$ such that $t' = t_k$ and s becomes marked during the firing of σ . As the process is deterministic, no transition can steal tokens from s . Since all the places in the pre-set of t becomes marked during the firing of σ , then t is enabled after σ . \square

APPENDIX E. ADDITIONAL PROCESSES OF THE RUNNING EXAMPLE

We show in Fig. 18 the additional processes of the net $(\llbracket N \rrbracket)$ of the running example and their probabilities.




FIGURE 18. Processes of the net $(\llbracket N \rrbracket)$ (running example)