

EVENT STRUCTURES FOR PETRI NETS WITH PERSISTENCE

PAOLO BALDAN^a, ROBERTO BRUNI^b, ANDREA CORRADINI^b, FABIO GADDUCCI^b,
HERNAN MELGRATTI^c, AND UGO MONTANARI^b

^a University of Padova, Italy
e-mail address: baldan@math.unipd.it

^b University of Pisa, Italy
e-mail address: bruni@di.unipi.it
e-mail address: andrea@di.unipi.it
e-mail address: gadducci@di.unipi.it
e-mail address: ugo@di.unipi.it

^c University of Buenos Aires - Conicet, Argentina
e-mail address: hmelgra@dc.uba.ar

ABSTRACT. Event structures are a well-accepted model of concurrency. In a seminal paper by Nielsen, Plotkin and Winskel, they are used to establish a bridge between the theory of domains and the approach to concurrency proposed by Petri. A basic role is played by an unfolding construction that maps (safe) Petri nets into a subclass of event structures, called prime event structures, where each event has a uniquely determined set of causes. Prime event structures, in turn, can be identified with their domain of configurations. At a categorical level, this is nicely formalised by Winskel as a chain of coreflections.

Contrary to prime event structures, general event structures allow for the presence of disjunctive causes, i.e., events can be enabled by distinct minimal sets of events. In this paper, we extend the connection between Petri nets and event structures in order to include disjunctive causes. In particular, we show that, at the level of nets, disjunctive causes are well accounted for by persistent places. These are places where tokens, once generated, can be used several times without being consumed and where multiple tokens are interpreted collectively, i.e., their histories are inessential. Generalising the work on ordinary nets, Petri nets with persistence 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.

Key words and phrases: event structures, disjunctive causes, local connectedness, Petri nets, persistence, concurrency, unfolding, coreflection.

PRELUDE

Among the multitude of his research interests, Furio has been working on the foundations of concurrency with special attention to the mathematical domains required for defining the semantics of concurrent systems. We like to recall his work on hyperuniverses as models for processes and on the final semantics of the pi-calculus achieved by means of an higher order presentation via Logical Framework. It was in the context of the European Project MASK (“Mathematical Structures for Concurrency”), led by Jaco De Bakker, that most of us had the first chance of working closely together with Furio. Surely that was one of the most relevant results of the project, paving the way for a fruitful scientific collaboration that continued in the following years in the framework of several research projects. This paper is heartfully dedicated to him in the occasion of his 60th birthday.

1. INTRODUCTION

Petri nets have been introduced in the Ph.D. Thesis of Carl Adam Petri [27] and soon have become one of the best known models of concurrency [29, 16, 14]. The conceptual simplicity of the model (multiset rewriting) and its intuitive graphical presentation have attracted the interest of both theoreticians and practitioners. Nowadays Petri nets are widely adopted across Computer Science and other disciplines such as Physics, Chemistry, and Biology [1, 18, 17]. They provide a basic model that, on the one hand, offers the ideal playground to study basic concepts of concurrent and distributed computations [26, 13, 31] and, on the other hand, can be readily extended to experiment with advanced features like structured data handling, read and inhibitor arcs, mobility, reflection, time and stochastic behavior [15, 7, 25, 2, 30, 21, 28, 11, 20, 24, 12].

In this paper we are interested in the seminal work of Winskel [33] on net unfolding, which has established a tight connection between Petri nets and (prime algebraic) domains. There it is shown that a chain of coreflections links the category of safe nets to the category of prime event structures, which in turn is equivalent to the category of prime algebraic domains. This is particularly satisfactory since a coreflection essentially establishes that a sub-category of abstract models can be found in a category of concrete models, such that each concrete model can be assigned the best possible abstract model. The first step of the chain is an unfolding construction that maps each net to a special kind of acyclic net (called non-deterministic occurrence net) representing all behaviours of the original net. From this an event structure can be easily defined, by forgetting the places of the net. Later these results have been extended to the more general class of semi-weighted nets [22, 23].

Petri nets semantics is based on consuming and producing data (i.e., tokens) from repositories (i.e., places). Operationally, reading a piece of information can be modelled by a transition that consumes a token from a place and produces it again on the same place. However, from the point of view of concurrency and causality such an encoding is not faithful as it disallows concurrent readings. Moreover, in many situations one is interested in representing persistent information that once created can be read but not consumed and such that its multiplicity (the number of available instances) is not important. This is the case for instance of classical logical conditions that once established to hold can be used repeatedly for proving other conditions. Another example is that of subversioning systems or cloud storage, where data changes are logged and previous versions of stored files remain accessible after an update. Persistent information is also used in [8] to model




FIGURE 1. Running example

security protocols, where the pool of messages exchanged by participants over the network forms a knowledge base that remains available for inspection and processing by attackers: in other words, sent messages cannot be removed. In [6] persistent tokens are used to remove confusion from acyclic nets and equip choices with standard probabilistic distributions. The equivalent framework of CPR nets [5] has been proposed for modelling web services: persistency is needed to capture service availability, and the formalism has been used for describing protocols that are specified with the ontology-based OWL-S language.

Read arcs have been introduced in the literature to handle multiple concurrent readings [7] and their concurrent semantics has been widely investigated [32, 4]. From the point of view of causality, they are not expressive enough to model another interesting phenomenon of persistent information that is the absence of multiplicity: if the same piece of persistent data is created several times, then its causes are merged and the events that read the persistent data can inherit any of them disjunctively. Instead, when read arcs are used, their event structure semantics records the exact causes.

The goal of this paper is to extend Winskel’s construction to Petri nets with persistence, as defined in [8, 9], and to understand what is the right sub-category of general event structures to exploit. Surely, prime event structures are not expressive enough, as they allow for a unique set of minimal causes for each event. Instead, as discussed above, the presence of persistent data leads to events with multiple sets of minimal causes. Consider, e.g., the net with persistence in Figure 1, which will serve as a running example. As usual, places and transitions are represented by circles and boxes, respectively. Persistent places, like o in this case, are represented by double circles. Intuitively, transition c is immediately enabled and can be fired many times (sequentially). As soon as a or b fires, place o becomes marked. The number of tokens in o as well as their causal histories are irrelevant: once o is marked, it will remain so forever. At this point d and e become (concurrently) enabled. The firing of d disables c but not e since the token never leaves the persistent place o .

In a recent work [3] dealing with the semantics of formalisms with fusion and with the corresponding phenomenon of disjunctive causality, the domains of configurations of general event structures are characterised by resorting to a weak form of prime algebraicity. *Connected event structures* are identified as canonical representations of general (possibly non-stable) event structures, in the same way as prime event structures are representations of stable event structures. In particular, the equivalence between the category of prime algebraic domains and the one of prime event structures is generalised to an equivalence between the category of weak prime domains and that of connected event structures. In a connected event structure, an event may have multiple causal histories, but they are required to be “connected”, namely they must not be separable in two classes of pairwise incompatible causal

histories. The idea is that, if in a general event structure an event has conflicting classes of causal histories, then it should split in several copies when generating the corresponding connected event structure. Indeed, as discussed in [3], connected event structures can be alternatively presented as prime event structures where some events (intuitively, those having different but not incompatible causal histories) are deemed equivalent. This establishes a close connection with the work in [10], where, in order to model strategies with disjunctive causes, the authors deal with *prime event structures with equivalence*.

In this paper we rely on the aforementioned work. A major role is played by a weakening of the connectedness property for event structures, referred to as *local connectedness*. The underlying intuition is as follows. A causal history can be seen as a conjunction of its events, thus an event with different causal histories is enabled by a disjunction of conjunctions of events. Connectedness amounts to the fact that the various conjuncts cannot be split in conflicting subclasses. Moving to Petri nets with persistence, a persistent place can be seen as the disjunction of all events that can fill the place. In turn, an event needs all places in its pre-set to be filled, hence it is enabled by a conjunction of disjunctions. The property of local connectedness roughly amounts to the requirements that the different ways of enabling a persistent place cannot be separated into conflicting classes. If this were possible, one should split the place in different copies, one for each class. The notion of local connectedness lifts to event structures and we show that Winskel’s chain of coreflections can be generalised to link the category of Petri nets with persistence to that of locally connected event structures. The latter, in turn, coreflects into the category of connected event structures.

The result can be read from two perspectives. From one viewpoint, where the construction defines the event structure associated with a net, it characterises the “right” concurrent semantics for dealing with formalisms that handle persistent information. The interesting bit is that the notion of connectedness from [3] is relaxed here to local connectedness. From the second viewpoint, the construction builds a standard net that is the best representative for the (locally connected) event structure at hand. The latter is a more interesting viewpoint, because: (i) it shows that Petri nets with persistence are expressive enough to account for disjunctive causes, advancing towards the solution of a long-standing open question about finding the right computational model for general event structures; and (ii) it confirms that Petri nets offer the ideal playground to experiment with concurrency features.

1.1. Structure of the paper. In Section 2 we review the preliminaries on connected event structures and we introduce the original class of locally connected event structures. In Section 3 we introduce Petri nets with persistence (p-nets) and the corresponding category. In Section 4 we introduce occurrence p-nets and define a coreflection between the category of p-nets and the one of occurrence p-nets. The right-adjoint of the coreflection is the unfolding construction that accounts for the description of concurrent computations of a p-net. Technically, it is defined in two steps, going through a category of occurrence p-nets with equivalence. In Section 5 we establish a coreflection between the category of p-nets and the one of locally connected event structures. Some concluding remarks are in Section 6.

2. EVENT STRUCTURES AND (LOCAL) CONNECTEDNESS

In this section we review the basics of event structures [33] and the notion of connected event structure from [3]. Then we single out a wider class of event structures, referred to as *locally connected*, that will play a pivotal role in the paper.

We start by recalling the notion of event structure with binary conflict [33]. In the following, for $m, n \in \mathbb{N}$, we denote by $[m, n]$ the set $\{m, m+1, \dots, n\}$. Also, given a set E , we denote by 2^E the powerset of E and by 2_f^E the set of finite subsets of E .

Definition 2.1 (event structure). An *event structure* (ES for short) is a tuple $\langle E, \vdash, \# \rangle$ such that

- E is a set of events;
- $\vdash \subseteq 2_f^E \times E$ is the *enabling* relation satisfying $X \vdash e$ and $X \subseteq Y$ implies $Y \vdash e$;
- $\# \subseteq E \times E$ is the *conflict* relation.

Two events $e, e' \in E$ are *consistent*, written $e \frown e'$, if $\neg(e\#e')$. A subset $X \subseteq E$ is *consistent* if $e \frown e'$ for all $e, e' \in X$.

An ES $\langle E, \vdash, \# \rangle$ is often denoted simply by E . Computations are captured in the form of configurations.

Definition 2.2 (configuration, live ES, concurrent events). A *configuration* of an ES E is a consistent $C \subseteq E$ which is *secured*, i.e., for all $e \in C$ there are $e_1, \dots, e_n \in C$ with $e_n = e$ such that $\{e_1, \dots, e_{k-1}\} \vdash e_k$ for all $k \in [1, n]$ (in particular, $\emptyset \vdash e_1$). The set of configurations of an ES E is denoted by $\text{Conf}(E)$. An ES is *live* if it has no self-conflicts, i.e., for all $e \in E$ we have $\neg(e\#e)$, and conflict is saturated, i.e., for all $e, e' \in E$, if $\neg(e\#e')$ then there is $C \in \text{Conf}(E)$ such that $\{e, e'\} \subseteq C$. Two events $e, e' \in E$ are *concurrent* if they are consistent ($e \frown e'$) and there is $C \in \text{Conf}(E)$ such that $C \vdash e$ and $C \vdash e'$.

Thus in a live ES conflict is saturated, a property that corresponds to inheritance of conflict in prime ESS, and each event is executable.

Remark 2.3. In the paper we restrict to live ES. Hence the qualification *live* is omitted.

The class of ES can be turned into a category.

Definition 2.4 (category of ES). A morphism of ES $f : \langle E_1, \vdash_1, \#_1 \rangle \rightarrow \langle E_2, \vdash_2, \#_2 \rangle$ is a partial function $f : E_1 \rightarrow E_2$ such that for all $C_1 \in \text{Conf}(E_1)$ and $e_1, e'_1 \in E_1$ with $f(e_1), f(e'_1)$ defined

- if $f(e_1) \#_2 f(e'_1)$ then $e_1 \#_1 e'_1$;
- if $f(e_1) = f(e'_1)$ and $e_1 \neq e'_1$ then $e_1 \#_1 e'_1$;
- if $C_1 \vdash_1 e_1$ then $f(C_1) \vdash_2 f(e_1)$.

We denote by **ES** the category of ES and their morphisms.

Since the enabling predicate is over finite sets of events, we can consider minimal sets of events enabling a given one.

Definition 2.5 (minimal enabling, causality). Let $\langle E, \vdash, \# \rangle$ be an ES. Given $e \in E$ and $C \in \text{Conf}(E)$ such that $C \vdash e$ we say that C is a *minimal enabling* of e , and write $C \vdash_0 e$, when for any configuration $C' \subseteq C$, if $C' \vdash e$ then $C' = C$. We denote by $\mathbf{H}(e) = \{C \mid C \vdash_0 e\}$ the set of minimal enablings of event e . We write $e < e'$ if $e \in C$ for all $C \in \mathbf{H}(e')$.

The configurations of an ES, ordered by subset inclusion, form a partial order that is characterised in [3] as a weak prime algebraic domain, i.e., a coherent finitary partial order where each element is the join of elements satisfying a weak notion of primality. The relation is formalised as a coreflection between the category **ES** and a category **wDom** of weak prime domains. A subclass of ESS can be identified, called *connected* ES, that represents the exact counterpart of weak prime domains, in the same way as prime ESS are the counterpart of prime algebraic domains [33].

Definition 2.6 (connected ES). Let $C, C' \in \mathbf{H}(e)$. We write $C \stackrel{e}{\sim} C'$ if $C \cup C' \cup \{e\}$ is consistent and we denote by $\stackrel{e}{\sim}^*$ the transitive closure of the relation $\stackrel{e}{\sim}$. An ES is *connected* if whenever $C, C' \in \mathbf{H}(e)$ then $C \stackrel{e}{\sim}^* C'$. The full subcategory of ES having connected ES as objects is denoted by **cES**.

The category **cES** is equivalent to the category **wDom** of weak prime domains and thus it coreflects in ES.

Proposition 2.7 (coreflection between ES and **cES** [3]). *The inclusion functor $\mathcal{I} : \mathbf{cES} \rightarrow \mathbf{ES}$ admits a right adjoint $\mathcal{C} : \mathbf{ES} \rightarrow \mathbf{cES}$ establishing a coreflection.*

As mentioned in the introduction, the concurrent semantics of Petri nets with persistence will be given in terms of connected ES through a chain of transformations that first unfolds the net into an acyclic net and then abstracts it to an ES.

The connectedness condition has a natural logical interpretation. Given an event e , we can capture its dependencies by stating that e is caused by the disjunction of its minimal enablings, where each minimal enabling can be seen in turn as a conjunction of events, namely by $\bigvee_{C \in \mathbf{H}(e)} \bigwedge C$. A conflict $e' \# e''$ can be encoded as $\neg e' \vee \neg e''$. In this view, connectedness amounts to the impossibility of partitioning $\mathbf{H}(e)$ in two subsets inducing mutually exclusive minimal causal histories, i.e., we cannot decompose $\mathbf{H}(e) = H_1 \uplus H_2$ in a way that $\neg(\bigvee_{C \in H_1} \bigwedge(C \cup \{e\})) \vee \neg(\bigvee_{C \in H_2} \bigwedge(C \cup \{e\}))$. If this happened, to recover connectedness we should split event e in two events e_1 and e_2 , with $\mathbf{H}(e_i) = H_i$ for $i \in \{1, 2\}$.

At the level of nets, events correspond to transitions. Each transition requires that all the (possibly persistent) places in its pre-set are filled in order to be enabled, hence it is enabled by a *conjunction* of places. In turn, each persistent place can be seen as the *disjunction* of the transitions in its pre-set. Summing up, at the level of nets we can represent conjunctions of disjunctions of events, exploiting persistence. The natural choice, when working with nets, will be to impose the connectedness condition locally to each disjunct. This results in a property weaker than the “global” connectedness from Definition 2.6. For this reason, the extraction of a connected ES from a net will pass through an intermediate class of ES that we call *locally connected*. We next formalise this idea.

Definition 2.8 (ES in disjunctive form). Let E be an ES. Given $e \in E$, a *disjunct* of e is a minimal set $X \subseteq E$ such that $X \cap C \neq \emptyset$ for all $C \in \mathbf{H}(e)$. It is *connected* if for all $e, e' \in X$ there exists $n \geq 1$ and $e_1, e_2, \dots, e_n \in X$ such that $e = e_1 \cap e_2 \cap \dots \cap e_n = e'$.

A *covering* of e is a set of disjuncts $D \subseteq 2^E$ such that for any $C \in \mathbf{Conf}(E)$, if $C \cap X \neq \emptyset$ for all $X \in D$ then $C \vdash e$.

Intuitively, D is a covering of an event e whenever condition $\bigwedge_{X \in D} \bigvee X$ is necessary and sufficient for enabling e , i.e., it is logically equivalent to the disjunction of the minimal enablings $\bigvee_{C \in \mathbf{H}(e)} \bigwedge C$. A disjunct X of e is connected whenever it cannot be partitioned as $X = X_1 \uplus X_2$ with $\neg(\bigvee X_1) \vee \neg(\bigvee X_2)$. Intuitively, we are moving from a \vee - \wedge form of the dependencies to a \wedge - \vee form, and transferring the connectedness condition from the outer to the inner disjunctions. Expressing dependencies as a conjunction of disjunctive causes makes it natural to associate a net with persistent places with the ES: each event e becomes a transition and each disjunct X of e corresponds to a persistent place s in the pre-set of e , filled by the events in X . The guarantee that the disjunctive causes cannot be split into inconsistent subsets will provide a form of canonicity to the construction.

Note that an event enabled by the empty set has no disjuncts (the empty set, which as a disjunct would correspond to “true”, is not admitted).

Definition 2.9 (locally connected ES). An ES E is *locally connected* if for all $e \in E$ there exists a covering D such that any $X \in D$ is connected. We denote by ℓES the full subcategory of ES having locally connected event structures as objects.

It can be easily shown that connectedness implies local connectedness. We first observe that in an ES every event admits a covering, which is the set of all its disjuncts.

Lemma 2.10 (coverings always exist). *Let E be an ES and $e \in E$. Then $D_e = \{X \mid X \subseteq E \wedge X \text{ is a disjunct of } e\}$ is a covering of e .*

Proof. Let $C \in \text{Conf}(E)$ be a configuration such that $C \cap X \neq \emptyset$ for all $X \in D_e$, and suppose by absurd that $C \not\vdash e$. This means that for all $C' \in \mathbf{H}(e)$ there is an event $e_{C'} \in C' \setminus C$. Let X' be a minimal subset of $\{e_{C'} \mid C' \in \mathbf{H}(e)\}$ such that $X' \cap C' \neq \emptyset$ for all $C' \in \mathbf{H}(e)$: then clearly X' is a disjunct in D_e , but $X' \cap C = \emptyset$, yielding a contradiction. \square

Lemma 2.11 (connectedness vs local connectedness). *Let E be an ES. If E is connected then it is locally connected.*

Proof. Let E be a connected ES and let $e \in E$ be an event. Let X be any disjunct of e and consider $e_1, e_2 \in X$. By minimality of a disjunct, we deduce that there must be $C_1, C_2 \in \mathbf{H}(e)$ such that $e_i \in C_i$, for $i \in \{1, 2\}$. Since E is connected, we know that $C_1 \stackrel{e}{\sim} C_2$. Then we can prove that $e_1 \frown^* e_2$ by induction on the length of the chain of consistency $C_1 \stackrel{e}{\sim} C_2$. The base case is trivial. For the inductive step, assume that $C_1 \stackrel{e}{\sim} C'_1 \stackrel{e}{\sim} C_2$. Take any $e'_1 \in C'_1$. By inductive hypothesis we know that $e_1 \frown^* e'_1$. Moreover, since $C'_1 \stackrel{e}{\sim} C_2$ we deduce $e'_1 \frown e_2$. Thus we conclude $e_1 \frown^* e_2$ as desired. \square

The above result shows that cES is a full subcategory of ℓES . Hence the coreflection between ES and cES restricts to a coreflection between ℓES and cES.

Proposition 2.12 (coreflection between ℓES and cES). *The inclusion functor $\mathcal{I} : \text{cES} \rightarrow \ell\text{ES}$ is left adjoint of the restriction $\mathcal{C}_{\ell\text{ES}} : \ell\text{ES} \rightarrow \text{cES}$, establishing a coreflection.*

Proof. Immediate consequence of Proposition 2.7 and Lemma 2.11. \square

Local connectedness is strictly weaker than connectedness. Let $E = \{a, b, c, d, e\}$ be the ES with $\emptyset \vdash_0 a$, $\emptyset \vdash_0 b$, $\emptyset \vdash_0 c$, $\emptyset \vdash_0 d$, $a \# c$, $b \# d$, $\{a, b\} \vdash_0 e$ and $\{c, d\} \vdash_0 e$. This is not so, since $\mathbf{H}(e) = \{\{a, b\}, \{c, d\}\}$ and $\{a, b\} \stackrel{e}{\sim} \{c, d\}$ does not hold. It is locally connected, since $\{\{a, d\}, \{b, c\}\}$ is a covering of e and the disjuncts $\{a, d\}$ and $\{b, c\}$ are connected since $a \frown d$ and $b \frown c$. Logically, the cause of e in $\vee\text{-}\wedge$ form is $(a \wedge b) \vee (c \wedge d)$. It is not connected since $\neg(a \wedge b) \vee \neg(c \wedge d)$ (and thus $\neg(a \wedge b \wedge e) \vee \neg(c \wedge d \wedge e)$). If we put the causes in $\wedge\text{-}\vee$ form we get $(a \vee d) \wedge (b \vee c)$, where neither $\neg a \vee \neg d$ nor $\neg b \vee \neg c$, whence local connectedness.

Dealing with (locally) connected ES will play an essential role for establishing a coreflection between occurrence nets with persistence and ESS (see Theorem 5.9). At an intuitive level, it ensures that or-causality is preserved along morphisms and cannot be transformed in ordinary causality. Instead, consider for instance the ES E_1 and E_2 defined as follows

- $E_1 = \{a_1, a_2, b\}$ with $\emptyset \vdash_0 a_i$, $\{a_i\} \vdash_0 b$ for $i \in \{1, 2\}$ and $a_1 \# a_2$;
- $E_2 = \{a, b\}$ with $\emptyset \vdash_0 a$ and $a \vdash_0 b$

The ES E_1 is not locally-connected, since $\{a_1, a_2\}$ is the only disjunct for b and it is not connected. It is easy to realise that the mapping $f : E_1 \rightarrow E_2$ defined by $f(a_1) = f(a_2) = a$ and $f(b) = b$ is an ES morphism: merging a_1 and a_2 the or-causality of b is transformed into a proper causality $a < b$.

3. NETS WITH PERSISTENCE

In this section we introduce Petri nets with persistence. Since the state of a Petri net will be seen as a multiset, i.e., an element of a suitably defined monoid, we start with some notation on sets and monoids. The irrelevance of the number of tokens in persistent places is modelled by some form of idempotency in the monoid.

Recall that, given a set X , we denote by 2^X the powerset of X and by 2_f^X the set of finite subsets of X . We denote by $X^\oplus = \{u \mid u : X \rightarrow \mathbb{N}\}$ the commutative monoid of multisets on X with the operation \oplus and identity \emptyset defined in an obvious way. Elements of X^\oplus are often represented by formal sums $u = \bigoplus_{x \in X} u(x) \cdot x$. Given $u \in X^\oplus$ we denote by $[u]$ the underlying set $\{x \mid x \in X \wedge u(x) > 0\}$. We write $x \in u$ for $x \in [u]$ and we say that u is *finite* if $[u]$ is finite. We identify a set $u \in 2^X$ and the “corresponding” multiset $\bigoplus_{x \in u} 1 \cdot x$. Given $u, u' \in X^\oplus$ we say that u *covers* u' , written $u' \subseteq u$, if there is $u'' \in X^\oplus$ such that $u = u' \oplus u''$. For $u, u' \in X^\oplus$ we write $u \cap u'$ for the largest u'' such that $u'' \subseteq u$ and $u'' \subseteq u'$.

When X has a chosen subset $X^p \subseteq X$, the commutative monoid with idempotency is $(X, X^p)^\oplus = \{u \mid u : X \rightarrow \mathbb{N} \wedge \forall x \in X^p. u(x) \leq 1\}$. Elements are still seen as formal sums $u = \bigoplus_{x \in X} u(x) \cdot x$, with an idempotency axiom $x \oplus x = x$ for any $x \in X^p$. As before, given $u, u' \in (X, X^p)^\oplus$ we say that u *covers* u' , written $u' \subseteq u$, if there is $u'' \in X^\oplus$ such that $u = u' \oplus u''$. Note however that due to idempotency there can be several u'' such that $u = u' \oplus u''$. More precisely, the set $\{u'' \mid u = u' \oplus u''\}$ forms a lattice with respect to \subseteq and we write $u \ominus u'$ to denote the top element of the lattice. For example, when $a \in X \setminus X^p$ and $b, c \in X^p$, although $(a \oplus b) \oplus (a \oplus c) = 2a \oplus b \oplus c$ we have $(2a \oplus b \oplus c) \ominus (a \oplus b) = a \oplus b \oplus c$.

Given sets X, Y , a monoid homomorphism $f : X^\oplus \rightarrow Y^\oplus$ is called *finitary* if for all $x \in X$, $f(x)$ is finite. A function $f : X \rightarrow Y^\oplus$ that is finitary (i.e., $f(x)$ finite for all $x \in X$) can be extended to a finitary monoid homomorphism denoted $f^\oplus : X^\oplus \rightarrow Y^\oplus$, defined by $f^\oplus(\bigoplus_{x \in X} n_x \cdot x) = \bigoplus_{x \in X} n_x \cdot f(x)$. Analogous notions can be defined for partial functions and when the target or both the source and target are monoids with idempotency.

Definition 3.1 (net with persistence). An (unweighted marked P/T Petri) *net with persistence* (*p-net*, for short) is a tuple $N = (S, S^p, T, \delta_0, \delta_1, u_0)$, where S is a set of *places*, $S^p \subseteq S$ is the set of *persistent* places, T is a set of *transitions*, $\delta_0, \delta_1 : T \rightarrow 2^S$ are functions assigning sets called pre-set and post-set, respectively, to each transition and $u_0 \subseteq S^n$ is the *initial marking*, where S^n is the set of non-persistent places $S^n = S \setminus S^p$.

Given a finite multiset of transitions $v \in T^\oplus$ we write $\bullet v$ and v^\bullet for $\delta_0^\oplus(v)$ and $\delta_1^\oplus(v)$. Given a place $s \in S$ we also write $\bullet s$ for $\{t \mid s \in t^\bullet\}$ and s^\bullet for $\{t \mid s \in \bullet t\}$.

Hereafter, for any p-net N we assume $N = (S, S^p, T, \delta_0, \delta_1, u_0)$, with subscripts and superscripts carrying over the names of the components. Note that we work with nets that are not weighted (pre- and post-sets of transitions are sets, rather than multisets). The results could be trivially extended, as in the ordinary case, to semi-weighted nets (where the pre-set can be a proper multiset). The restriction is adopted to ease the presentation.

The state of a p-net N is represented by some $u \in (S, S^p)^\oplus$, called a *marking* of N . A transition t is *enabled* by a marking u if its pre-set is covered by u , i.e. if $\bullet t \subseteq u$. In this case, t can be fired. The firing of t consumes the tokens in the non-persistent places in the pre-set, leaves untouched the tokens in the persistent places and produces the tokens in the post-set. More generally, this applies to finite multisets of transitions. Formally, given a finite multiset $v \in T^\oplus$ and a marking $u \in (S, S^p)^\oplus$, if $\bullet v \subseteq u$, the firing rule is $u[v](u \ominus \bullet v) \oplus v^\bullet$.

A marking $u \in S^\oplus$ is *reachable* if there exists a firing sequence $u_0[t_1]u_1[t_2]\dots[t_n]u$ from the initial marking to u . The p-net is *safe* if every reachable marking is a set.

Observe now that in the initial marking only non-persistent places can be marked. Indeed, if a persistent place p were marked in the initial marking then its presence would be essentially useless, since any reachable marking would contain one token in p . Therefore removing p and all the incoming and outgoing arcs would lead to an equivalent net that can perform the same firing sequences and such that a marking u is reachable if and only if the marking $u \oplus p$ was reachable in the original net.

Example 3.2. Let us consider the p-net in Figure 1 of our running example, whose initial marking is $u_0 = p \oplus q \oplus r \oplus s$. A sample firing sequence is

$$u_0 [a] q \oplus r \oplus s \oplus o [b] r \oplus s \oplus o [c] r \oplus s \oplus o [d] t \oplus s \oplus o [e] t \oplus u \oplus o$$

Note that once the token in the persistent place o is produced by the firing of a , it is never possible to remove it, not even firing d or e . Also note that after b fires, the multiplicity of o remains 1 due to the idempotency axiom on persistent places. It is immediate to check that the p-net is safe, as all reachable markings are sets.

Remark 3.3. Our presentation of p-nets slightly differs from the original one in [9]. There: (i) arcs carry weight 1 or ∞ ; (ii) an arc has weight ∞ if and only if it goes from a transition to a persistent place; (iii) markings allows having either 0 or infinitely many tokens in each persistent place; (iv) even if a firing removes finitely many tokens from persistent place, there remain infinitely many tokens available. As each marking associates only one bit of information with each persistent place, here we find technically more convenient to represent marked persistent places by assigning them multiplicity 1 and by exploiting idempotency to capture the fact that when infinitely many tokens are added to a marked persistent place there still are infinitely many tokens (indistinguishable one from the other).

Remark 3.4. We use the term “net with persistence” to avoid confusion with the notion of “persistent net” in the literature (see e.g. [19]), a behavioural property defined as follows: A net is persistent if whenever $u[t_1]$ and $u[t_2]$ for a reachable marking u , then $u[t_1 \oplus t_2]$.

When dealing with unfolding and, more generally, with the causal semantics of ordinary Petri nets, it is a standard constraint to assume that transitions have a non-empty pre-set. This avoids the unnatural possibility of firing infinitely many copies of the same transition in parallel. For p-nets, this generalises to the requirement that each transition consumes tokens from at least one non-persistent place. Additionally, since a persistent place will never be emptied once it is filled with a token, whenever a persistent place s is in the post-set of a transition t it is quite natural to forbid the presence of an additional path from t to s . This property is formalised by using the flow relation \rightsquigarrow_N for a net N , defined, for all $x, y \in S \cup T$, by $x \rightsquigarrow_N y$ if $x \in \bullet y$.

Definition 3.5 (well-formed net). A p-net N is *well-formed* if for all $t \in T$, $\delta_0(t) \notin S^{\text{p}\oplus}$ (*t-restrictedness*) and for all $t \in T$, $s \in S^{\text{p}}$, if $t \rightsquigarrow_N s$ then $t \not\rightsquigarrow_N^n s$ for $n \geq 2$ (*irredundancy*).

Observe that, in particular, whenever a p-net is irredundant, it does not include cycles over persistent places, i.e., for any $s \in S^{\text{p}}$, it is not the case that $s \rightsquigarrow_N^+ s$. Hereafter all p-nets will be tacitly assumed to be well-formed.

The notion of p-net morphism naturally arises from an algebraic view, where places and transitions play the role of sorts and operators.

Definition 3.6 (p-net morphism). A *p-net morphism* $f = \langle f_s, f_t \rangle : N \rightarrow N'$ is a pair where

- (1) $f_s : (S, S^p)^\oplus \rightarrow (S', S'^p)^\oplus$ is a finitary monoid homomorphism such that for $s \in S^n$, $f_s(s) \in S'^n$ and the initial marking is preserved, i.e. $f_s(u_0) = u'_0$;
- (2) $f_t : T \rightarrow T'$ is a partial function such that for all finite $v \in T^\oplus$, $\bullet(f_t^\oplus(v)) = f_s(\bullet v)$ and $(f_t^\oplus(v))^\bullet = f_s(v^\bullet)$;
- (3) for all $t \in T$ and $s_1, s_2 \in \bullet t$ or $s_1, s_2 \in t^\bullet$, if $f_s(s_1) \cap f_s(s_2) \neq \emptyset$ then $s_1 = s_2$.

The category of p-nets (as objects) and their morphisms (as arrows) is denoted by PN.

Observe that by the fact that $f_s : (S, S^p)^\oplus \rightarrow (S', S'^p)^\oplus$ is a monoid homomorphism and condition (1) we automatically get that $s \in S^p$ implies $f_s(s) \in S'^p$. Moreover, for all $u_1, u_2 \in (S, S^p)^\oplus$, with $u_2 \subseteq u_1$, it holds that $f_s(u_1 \ominus u_2) = f_s(u_1) \ominus f_s(u_2)$. Condition (2) amounts to require that for any $t \in T$, if $f_t(t)$ is defined then $\bullet f_t(t) = f_s(\bullet t)$ and $f_t(t)^\bullet = f_s(t^\bullet)$, and $f_s(\bullet t) = f_s(t^\bullet) = \emptyset$ otherwise. Finally, condition (3) imposes injectivity of the morphism on the pre-set and post-set of each transition, i.e., the morphism should not mix places in the pre-set (and in the post-set) of the same transition. This is automatically satisfied for non-persistent places, but it could be violated by persistent places (due to idempotency), hence we require it explicitly.

In the sequel, when the meaning is clear from the context, we often omit the subscripts from the morphism components, thus writing f instead of f_s and f_t .

Lemma 3.7 (p-net morphisms are simulations). *Let $f : N \rightarrow N'$ be a p-net morphism. If $u[v]u'$ is a firing in N then $f(u)[f^\oplus(v)]f(u')$ is a firing in N' .*

Proof. Since f is a p-net morphism we have $\bullet(f^\oplus(v)) = f(\bullet v)$ and $(f^\oplus(v))^\bullet = f(v^\bullet)$. Since $u[v]u'$, it must be the case that $u = u'' \oplus \bullet v$ for some u'' , from which $f(u) = f(u'') \oplus f(\bullet v)$, i.e., $f(\bullet v) \subseteq f(u)$ and thus $f^\oplus(v)$ is enabled at $f(u)$. Therefore there is a firing

$$f(u)[f^\oplus(v)](f(u) \ominus \bullet(f^\oplus(v))) \oplus (f^\oplus(v))^\bullet$$

and the target $(f(u) \ominus \bullet(f^\oplus(v))) \oplus (f^\oplus(v))^\bullet = (f(u) \ominus f(\bullet v)) \oplus f(v^\bullet) = f((u \ominus \bullet v) \oplus v^\bullet) = f(u')$, as it was required. \square

4. UNFOLDING NETS WITH PERSISTENCE

In this section we show how a p-net can be unfolded to a suitably defined occurrence p-net that represents all possible occurrences of firing of transitions and their dependencies. We first introduce the class of occurrence p-nets, where each transition can occur at most once in a computation, but possibly with different disjunctive causes. We next observe that occurrence p-nets can be equivalently presented by forcing each transition to have a uniquely determined set of causes and using an equivalence between transitions to account for disjunctive causes. Finally, we present the unfolding construction for a p-net, that works in two steps. First a p-net is unfolded into an occurrence p-net without backward conflicts, where each item has a uniquely determined causal history. The possibility of generating a token in the same persistent place with different histories is captured by means of an equivalence relation on places and transitions. Then the actual unfolding is obtained as the quotient of the pre-unfolding with respect to the equivalence relation.

4.1. Occurrence p-nets. In order to single out the class of occurrence p-nets, we start by defining the possible dependencies in a p-net.

Definition 4.1 (enabling, conflict, causality, dependence, concurrency). Let N be a p-net.

- Enabling, written \vdash_N , is defined by letting, for $X \in 2_f^T$ and $t \in T$, $X \vdash_N t$ if for all $s \in \bullet t$ either $s \in u_0$ or there exists $t' \in X$ such that $s \in t'^\bullet$.
- Conflict $\#_N \subseteq 2^{S \cup T}$ is the least set-relation, closed by superset, defined by
 - (a) if $t \neq t'$ and $\bullet t \cap \bullet t' \not\subseteq S^p$, then $\#_N\{t, t'\}$;
 - (b) if $\bullet s \neq \emptyset$ and $\#_N(X \cup \{t\})$ for all $t \in \bullet s$, then $\#_N(X \cup \{s\})$;
 - (c) if $\#_N(X \cup \{s\})$ and $s \in \bullet t$, then $\#_N(X \cup \{t\})$.

We will often write $x\#_N x'$ instead of $\#_N\{x, x'\}$. We say that X is *consistent* if it is not in conflict. In particular, binary consistency is denoted \frown_N , i.e., $x \frown_N x'$ when $\neg(x\#_N x')$. We say that X is *connected* by \frown_N if for all $x, x' \in X$ there exists $n \geq 1$ and $x_1, x_2, \dots, x_n \in X$ such that $x = x_1 \frown_N x_2 \frown_N \dots \frown_N x_n = x'$.

- Causality \leq_N is the least transitive and reflexive relation \leq_N on $S \cup T$ such that if $s \in \bullet t$ then $s \leq_N t$ and if $\bullet s = \{t\}$ then $t \leq_N s$. We write $x <_N x'$ when $x \leq_N x'$ and $x \neq x'$.

When the context is clear, we will omit the subscript N , writing \vdash , $\#$, \leq and \frown instead of \vdash_N , $\#_N$, \leq_N and \frown_N .

Differently from what happens for ordinary nets, in occurrence p-nets we will allow different ways of enabling the same occurrence of a transition. This is because tokens, once generated in persistent places, cannot be consumed and all tokens in the same place are “merged” into one (by idempotency) in a way that the resulting single token joins all the different possible causal histories. For this reason it is convenient to resort to a general notion of enabling where a transition naturally has several sets of transitions that allow for its execution. Observe that if $X \vdash e$ and $X \subseteq X'$ then $X' \vdash e$. As it happens for ESS, when $X \vdash e$, the set X can be inconsistent. We will later impose a condition forcing each transition to be enabled by at least one minimal consistent set.


Note that direct conflict is only binary, since it is determined by the competition on non-persistent places. However, in order to define properly inheritance of conflict along the enabling relation we need to work with a conflict relation on generic sets. Consider for instance, the p-net in Figure 2A and the set of places $\{p, q, r\}$. The rules for conflict allow us to deduce that $\#\{p, q, r\}$, while $\neg x\#y$ for all $x, y \in \{p, q, r\}$. Intuitively this happens because at most two places in the set $\{p, q, r\}$ can be filled. The formal derivation can be found in Figure 2B. The relation between coverability and absence of conflicts will be characterised later via the notion of concurrency, in Lemma 4.10.

We also introduce causality: when an item x is a cause of an item y the intuition is that the presence of y in a computation implies the presence of x , i.e., x is needed to “enable” y . This will play a role later. However, note that causality alone would not completely characterise the dependencies in the p-net as it does not account for disjunctive causes.

We next define the notion of securing sequence in the context of p-nets.

Definition 4.2 (securing sequence). Let N be a p-net. Given a transition $t \in T$, a *securing sequence* for t in N is a sequence of distinct transitions $t_1, t_2, \dots, t_n = t$ such that for all $i, j \in [1, n]$, $\neg(t_i\#t_j)$ and for all $i \in [1, n]$, $\{t_1 \dots t_{i-1}\} \vdash t_i$.

A securing sequence for t is intended to represent a firing sequence that leads to the execution of t . This fact will be later formalised, for occurrence p-nets, in Lemma 4.6.



$$\begin{array}{c}
 \frac{\bullet b_1 \cap \bullet c_1 \not\subseteq S^p}{\# \{b_1, c_1\}} \quad \frac{\bullet a_1 \cap \bullet c_1 \not\subseteq S^p}{\# \{a_1, c_1\}} \quad \frac{\bullet a_2 \cap \bullet b_2 \not\subseteq S^p}{\# \{a_2, b_2\}} \quad \frac{\bullet a_1 \cap \bullet b_1 \not\subseteq S^p}{\# \{a_1, b_1\}} \quad \frac{\bullet a_2 \cap \bullet c_2 \not\subseteq S^p}{\# \{a_2, c_2\}} \quad \frac{\bullet b_2 \cap \bullet c_2 \not\subseteq S^p}{\# \{b_2, c_2\}} \\
 \hline
 \frac{\# \{p, b_1, c_1\}}{\# \{p, q, c_1\}} \quad \frac{\# \{a_1, b_2, c_1\}}{\# \{p, b_2, c_1\}} \quad \frac{\# \{a_2, b_2, c_1\}}{\# \{p, b_2, c_1\}} \quad \frac{\# \{a_1, b_1, c_2\}}{\# \{p, b_1, c_2\}} \quad \frac{\# \{a_2, b_1, c_2\}}{\# \{p, b_1, c_2\}} \quad \frac{\# \{b_2, c_2\}}{\# \{p, b_2, c_2\}} \\
 \hline
 \# \{p, q, r\}
 \end{array}$$

(B)

FIGURE 2. Non-binary conflict occurrence p-nets

Occurrence p-nets can now be defined as a subclass of p-nets. For historical reasons, places and transitions of occurrence p-nets are called *conditions* and *events*, respectively. We will adopt this convention in the rest of the paper, denoting occurrence p-nets by O , with components $(B, B^p, E, \gamma_0, \gamma_1, v_0)$, where B is the set of conditions and E the set of events.

Definition 4.3 (occurrence p-net). A structure $O = (B, B^p, E, \gamma_0, \gamma_1, v_0)$ is an *occurrence p-net* if it is a p-net such that

- (1) the initial marking v_0 satisfies $v_0 = \{b \in B \mid \bullet b = \emptyset\}$;
- (2) each event $e \in E$ admits a securing sequence in O ;
- (3) for all $e, e' \in E$, if $e \neq e'$ then $e^\bullet \cap e'^\bullet \subseteq B^p$;
- (4) for all $b \in B$, $\bullet b$ is connected by \cap_O .

An occurrence p-net is *without backward conflicts* when for all $e, e' \in E$, if $e \neq e'$ then $e^\bullet \cap e'^\bullet = \emptyset$. We denote by ON the full subcategory of PN with occurrence p-nets as objects.

Observe that by the requiring all events secured (together with the fact that the initial marking is made of the conditions with empty pre-set) an occurrence p-net could only be cyclic for the presence of back-pointers to persistent conditions. This latter possibility is excluded structurally by the irredundancy assumption (Definition 3.5). The connectedness requirement for the pre-set of conditions is trivially satisfied for non-persistent conditions, since they can have at most one event in their pre-set by condition (3). Instead, for persistent conditions it is an actual constraint: it imposes that we cannot split the generators of the condition in two subsets not connected by consistency. Again the idea is that, if this were the case, the condition should be split into two occurrences each having an element of the partition as pre-set. This is reminiscent of the local connectedness requirement for ES and, indeed, it will ensure that the ES extracted from an occurrence p-net is locally connected.

An occurrence p-net inspired by our running example is in Figure 3 (left). Event a has a unique securing sequence consisting of a itself, and similarly for b ; d instead has two minimal securing sequences: ad and bd . There is only one backward conflict, on the persistent condition o , since $a^\bullet \cap b^\bullet = \{o\} \subseteq S^p$. The conflict relation is empty and there is no causality between events. Note in particular, that $\bullet o$ is connected by \cap_O since $a \cap_O b$.




FIGURE 3. An occurrence p-net and its domain of configurations




FIGURE 4. A p-net that is not an occurrence one (left) and its unfolding (right).

Consider now, the p-net in Figure 4 (left). It satisfies all conditions of Definition 4.3 but (4). In fact, now the pre-set of $\bullet o = \{a, b\}$ is not connected by \frown_N since $a \#_N b$. Intuitively, to recover connectedness the place o should be split in two places o_1 and o_2 , with pre-sets $\{a\}$ and $\{b\}$, respectively, thus getting the occurrence p-net in Figure 4 (right). Indeed, this is the occurrence p-net arising from the unfolding construction described below.

Note that, as in the case of ordinary occurrence nets, the initial marking is determined by the structure of the net, i.e., it consists of the set of non-persistent conditions with empty pre-set. Still, for the sake of clarity it will be indicated explicitly. It can be easily seen that, in absence of persistent conditions, Definition 4.3 gives the ordinary notion of occurrence net. We have indeed that

- there is no self-conflict because each event is secured;
- there is no backward conflict because for all $e, e' \in E$, if $e \neq e'$ then $e^\bullet \cap e'^\bullet = \emptyset$.

Some proofs of results on occurrence p-nets will exploit induction on the following notion of depth, that generalises the one of ordinary occurrence nets.

Definition 4.4 (depth). Let $O = (B, B^p, E, \gamma_0, \gamma_1, v_0)$ be an occurrence p-net. We define the *depth* of events and conditions as

$$\begin{aligned} d(e) &= 1 + \max\{d(b) \mid b \in \bullet e\} \\ d(b) &= \begin{cases} 0 & \text{if } b \in v_0 \\ \min\{d(e) \mid e \in \bullet b\} & \text{otherwise} \end{cases} \end{aligned}$$

We introduce a notion of configuration for occurrence p-nets that is intended to capture the concept of a (concurrent) computation.

Definition 4.5 (configuration of occurrence p-nets). Let $O = (B, B^p, E, \gamma_0, \gamma_1, v_0)$ be an occurrence p-net. A configuration is a set of events $C \subseteq E$ such that each $e \in C$ admits a securing sequence in C and for all $e, e' \in C$, $\neg(e \#_N e')$. We denote by $\text{Conf}(O)$ the set of configurations of O .

The configurations of the net in Figure 3 (left) occur on the right, ordered by inclusion.

We next prove that configurations can be interpreted as representations of classes of firing sequences where the order of independent firings is abstracted.

Lemma 4.6 (configurations are executions). *Let $O = (B, B^p, E, \gamma_0, \gamma_1, v_0)$ be an occurrence p-net. Then the following statements are equivalent*

- (1) $C \in \text{Conf}(O)$ is a finite configuration;
- (2) there exists a securing sequence of events e_1, \dots, e_n such that $C = \{e_1, \dots, e_n\}$;
- (3) there exists a firing sequence $v_0[e_1]v_1[e_2] \dots [e_n]v_n$ such that $C = \{e_1, \dots, e_n\}$ and $v_n = (v_0 \cup \bigcup_{i=1}^n e_i^\bullet) \ominus (\bigcup_{i=1}^n \bullet e_i)$.

Proof. (1 \leftrightarrow 2) Let $C \in \text{Conf}(O)$ be a finite configuration. We prove the result by induction on $|C|$. The base case $|C| = 0$ is trivial. If $|C| = n + 1$ for some $n \in \mathbb{N}$, then for each $e \in C$ fix a single securing sequence of minimal length $s_e = e_1, \dots, e_k, e$ in C , which exists by definition of configuration, and call $S = \{s_e \mid e \in C\}$ the set of such securing sequences. Let $\hat{s}_{\hat{e}} \in S$ be one of such sequences of maximal length in S , with \hat{e} as last event. Clearly $C' = C \setminus \{\hat{e}\}$ is again a configuration because its events are not in conflict (as they are in C) and for each $e \in C'$ the securing sequence $s_e \in S$ consists of events in C' only, otherwise the minimality of $\hat{s}_{\hat{e}}$ as a securing sequence for \hat{e} in C would be violated. Then, by inductive hypothesis, there exists a securing sequence $e_1 \dots e_n$ such that $C' = \{e_1, \dots, e_n\}$. It is immediate to see that e_1, \dots, e_n, \hat{e} is the desired securing sequence for C .

The converse implication is immediate, just observing that securing sequences are closed by prefix and conflict free.

(2 \leftrightarrow 3) Assume that there is a securing sequence e_1, \dots, e_n such that $C = \{e_1, \dots, e_n\}$. We proceed by induction on n . The base case $n = 0$ is trivial. For $n > 0$, we know, by inductive hypothesis that there exists a firing sequence $v_0[e_1]v_1[e_2] \dots [e_{n-1}]v_{n-1}$ and $v_{n-1} = (v_0 \cup \bigcup_{i=1}^{n-1} e_i^\bullet) \ominus (\bigcup_{i=1}^{n-1} \bullet e_i)$. By definition of securing sequence, $\{e_1, \dots, e_{n-1}\} \vdash e_n$, we know that for all $b \in \bullet e_n$ either $b \in v_0$ or there exists $i < n$ such that $b \in e_i^\bullet$. Moreover, for all $i < n$, if $b \in \bullet e_n \cap B^n$, certainly $b \notin \bullet e_i$, otherwise we would have $e_i \# e_n$. This allows us to deduce that $\bullet e_n \subseteq v_n$. Hence the firing sequence can be extended by $v_{n-1}[e_n]v_n$, where $v_n = (v_{n-1} \cup e_n^\bullet) \ominus \bullet e_n = (v_0 \cup \bigcup_{i=1}^n e_i^\bullet) \ominus (\bigcup_{i=1}^n \bullet e_i)$.

For the converse implication, let us start showing by induction on n that if $C = \{e_1, \dots, e_n\}$ is the set of events of a given firing sequence, then

$$\text{there cannot be } i, j \in [1, n] \text{ such that } i \neq j \text{ and } \bullet e_i \cap \bullet e_j \cap B^n \neq \emptyset. \quad (4.1)$$

For $n \leq 1$ the statement is trivial. Let $n > 1$ and assume, by contradiction, that there is a non-persistent place $b \in \bullet e_n \cap \bullet e_i \cap B^n$, for $i < n$. Clearly $b \notin v_0$, otherwise $\bullet b = \emptyset$ and after the firing of e_i there would be no way of generating the token in b . Therefore $\bullet b \neq \emptyset$. Let $e \in \bullet b$ be the only event generating a token in b . Then necessarily there are $j < i$ and $k < n$ such that $e_j = e_k = e$ and thus, since O is t-restricted, $\bullet e_j \cap \bullet e_k \cap B^n \neq \emptyset$, which is absurd because we assumed that the statement holds for firing sequences shorter than n .

Now, by Definition 4.2 we have to show that (1) all events of C are distinct, (2) $\neg \#C$, and (3) for all $i \in [1, n]$, $\{e_1 \dots e_{i-1}\} \vdash e_i$. Point (1) follows from (4.1) and t-restrictedness. For (3), suppose by absurd that $\{e_1 \dots e_{i-1}\} \not\vdash e_i$ for some $i \in [1, n]$. Thus there is a place $b \in \bullet e_i$ such that $b \notin v_0$ and $b \notin e_j^\bullet$ for all $j \in [1, i-1]$, but this is impossible because by assumption e_i is enabled in v_{i-1} .

For (2) we proceed by induction on n , the base case $n = 0$ being obvious. Suppose by absurd that $\#C$ holds, and let $k \in [1, n]$ be the minimal index such that $\# \{e_1, \dots, e_k\}$. By this assumption we know that $\neg \# \{e_1, \dots, e_{k-1}\}$, thus by induction hypothesis that $\{e_1, \dots, e_{k-1}\}$ is a securing sequence, and by the implication $(2 \rightarrow 3)$ already proved that $v_0[e_1]v_1[e_2] \dots [e_{k-1}]v_{k-1}$ is a firing sequence and $v_{k-1} = (v_0 \cup \bigcup_{i=1}^{k-1} e_i^\bullet) \ominus (\bigcup_{i=1}^{k-1} \bullet e_i)$.

Now, since $\# \{e_1, \dots, e_k\}$, by the clauses defining conflict in Definition 4.1 either (clause (a)) there is an event e_h with $h \in [1, k-1]$ such that $\bullet e_h \cap \bullet e_k \not\subseteq B^p$, but this is impossible by (4.1) above, or (clause (c)) w.l.o.g. there is a place $b \in \bullet e_k$ such that $\# \{e_1, \dots, e_{k-1}, b\}$. Since e_k is enabled in v_{k-1} , certainly $b \in v_{k-1} = (v_0 \cup \bigcup_{i=1}^{k-1} e_i^\bullet) \ominus (\bigcup_{i=1}^{k-1} \bullet e_i)$, and as it cannot belong to the initial marking, $b \in e_j^\bullet$ for some $j \in [1, k-1]$. But by (clause (b)) from $\# \{e_1, \dots, e_{k-1}, b\}$ we can infer $\# \{e_1, \dots, e_{k-1}, e_j\}$ and thus $\# \{e_1, \dots, e_{k-1}\}$ because $j \in [1, k-1]$, which is absurd by the choice of k . \square

By the above result, it is meaningful to define the marking reached after a configuration for occurrence p-nets.

Definition 4.7 (marking after a configuration). Let O be an occurrence p-net. Given $C \in \text{Conf}(O)$, the marking after C is $m(C) = (v_0 \cup \bigcup_{e \in C} e^\bullet) \ominus (\bigcup_{e \in C} \bullet e)$.

We next observe that all occurrence p-nets are safe.

Proposition 4.8 (occurrence p-nets are safe). *Let O be an occurrence p-net. Then, for each firing sequence $v_0[e_1]v_1[e_2] \dots [e_n]v_n$ all markings v_i are sets.*

Proof. By Lemma 4.6 we know that $\{e_1, \dots, e_i\}$ is a securing sequence, thus all the events in it are pairwise distinct. The fact that all markings are safe immediately follows recalling that for non-persistent places $|\bullet b| \leq 1$, and for persistent places idempotency ensures that at most one token is in the place. \square

We next introduce a notion of concurrency for occurrence p-nets that, as anticipated, is based on non-binary conflict.

Definition 4.9 (concurrency). Let O be an occurrence p-net. A subset of conditions $X \subseteq B$ is called *concurrent*, written $co(X)$, if $\neg \#X$ and for all $b, b' \in X$ if $b < b'$ then $b \in B^p$.

Concurrency, as in the case of ordinary nets, is intended to provide a structural characterisation of coverability. This is formalised below in Lemma 4.10. Quite intuitively, a concurrent set of conditions cannot include conflicts. Note, instead, that causal dependencies from persistent places are admitted, consistently with the fact that using a token in a persistent place does not consume such token.

Lemma 4.10 (coverability vs concurrency). *Let O be an occurrence p-net and $X \subseteq B$. Then X is concurrent iff there is a reachable marking that covers X .*

Proof. (\Rightarrow) Let $co(X)$. In order to prove that X is coverable we proceed by induction on the pairs $h_X = \langle md(X), \#md(X) \rangle$ where $md(X) = \max\{d(b) \mid b \in X\}$, and $\#md(X) = |\{b \in X \mid d(b) = md(X)\}|$, ordered by $\langle n, m \rangle < \langle n', m' \rangle$ if $n < n'$ or $n = n' \wedge m < m'$.

If $h_X = \langle 0, \cdot \rangle$ then $X \subseteq v_0$, hence we conclude immediately. If $md(X) > 0$, take a condition $b \in X$ such that $d(b)$ is maximal. Note that there must be $e \in \bullet b$ such that $Y = (X \setminus e\bullet) \cup \bullet e$ is still concurrent, otherwise either b would not be of maximal depth or, by the rules defining conflict, $\#X$. Clearly $h_Y < h_X$ hence, by inductive hypothesis, Y is coverable and we can conclude that X is coverable.

(\Leftarrow) We show that any reachable marking is concurrent. Since any subset of a concurrent set is concurrent this allows us immediately to conclude. Consider a firing sequence $v_0[e_1]v_1[e_2] \dots [e_n]v_n$. We show by induction on n that $co(v_n)$. For the base case we just need to observe that $co(v_0)$, i.e., the initial marking is clearly concurrent. When $n > 0$, by inductive hypothesis we know that $co(v_{n-1})$. This allows us to deduce that also $v_n = (v_{n-1} \ominus \bullet e_n) \oplus e_n \bullet$ is concurrent. \square

For instance, consider the occurrence p-net in Figure 2a. We already observed that $\# \{p, q, r\}$ and indeed such set is not coverable. Instead, each pair of conditions in $\{p, q, r\}$ is concurrent and thus coverable. Interestingly enough, this shows that differently from what happens for ordinary occurrence nets, pairwise coverability does not imply coverability.

We can now show that morphisms of occurrence p-nets preserve concurrency.

Lemma 4.11 (morphisms preserve concurrency). *Let $f : O \rightarrow O'$ be a morphism of occurrence p-nets and $X \subseteq B$. If $co(X)$ then $co(f_s(X))$.*

Proof. Since morphisms are simulations (see Lemma 3.7) they preserve coverability. By Lemma 4.10, coverability is the same as concurrency. Hence $co(X)$ implies $co(f_s(X))$. \square

4.2. Occurrence p-nets with equivalence. We now introduce the notion of occurrence p-net with equivalence, which will be the target of the pre-unfolding construction. The intuition is that, in an occurrence p-net with equivalence, occurrences of items that depend on different disjunctive causes are kept separate, but related by the equivalence. This is technically useful in the development of the unfolding construction.

Let A be a set and $\sim \subseteq A \times A$ an equivalence relation on A . Given $x \in A$ we write $[x]_\sim = \{y \in A \mid x \sim y\}$ for the equivalence class of x . Moreover, given $X \subseteq A$ we write $[X]_\sim$ for the set $\{[x]_\sim \mid x \in X\}$.

Given two subsets $X, Y \subseteq A$ we write $X \sim Y$ instead of $[X]_\sim = [Y]_\sim$. Should the equivalence be used (as it will always be the case) for sets X and Y each consisting of pairwise non-equivalent elements, from $X \sim Y$ it follows that there is a bijection from X to Y mapping each element $x \in X$ to the only $y \in Y$ such that $x \sim y$. If moreover A is partially ordered by \leq , we write $X \leq Y$ if for all $x \in X, y \in Y$ we have $x \leq y$. The notation will be used only for $|X| \leq 1$. Observe that $\emptyset \leq Y$ trivially holds, while $\{x\} \leq Y$ reduces to $x \leq y$ for all $y \in Y$.

Definition 4.12 (occurrence p-net with equivalence). An occurrence p-net with equivalence is a pair $\langle O, \sim \rangle$ where $O = (B, B^p, E, \gamma_0, \gamma_1, v_0)$ is an occurrence p-net without backward conflicts and $\sim \subseteq (B \times B) \cup (E \times E)$ is an equivalence such that

- (1) for all $b, b' \in B$ with $b \sim b'$
 - (a) either $b, b' \in B^p$ or $b, b' \in B^n$;
 - (b) if $b \neq b'$ then $\neg(\bullet b \leq \bullet b')$ and $\{b, b'\} \not\subseteq \bullet e$ for all $e \in E$;
 - (c) if $b, b' \in B^n$ then $\bullet b \sim \bullet b'$;
 - (d) if $b, b' \in B^p$ then either $\bullet b \sim \bullet b'$ or $b(\sim \setminus \#)^* b'$;




FIGURE 5. An occurrence p-net with equivalence

- (2) for all $e, e' \in E$, if $e \sim e'$ and $e \neq e'$ then $\bullet e \sim \bullet e'$, $\bullet e \neq \bullet e'$, and $e^\bullet \sim e'^\bullet$;
- (3) for all $X, X' \subseteq B$, if $X \sim X'$ then $\{e \mid \bullet e = X\} \sim \{e' \mid \bullet e' = X'\}$.

In words, an occurrence p-net with equivalence is an occurrence p-net where the absence of backward conflicts implies that each item has a uniquely determined causal history. The possibility of joining the different histories of tokens in persistent places is captured by an equivalence that can equate persistent places. Events are equated when they have equivalent pre-sets and, in turn, this implies that also their post-sets are equivalent.

More precisely, by condition (1a) equivalence respects the sort of places: two equivalent places are either both non-persistent or both persistent. By the second part of condition (1b) the pre-set of each event consists of pairwise non-equivalent places. The first part is slightly more complex. First note that whenever $b \in v_0$, i.e. $\bullet b = \emptyset$, the inequality $\bullet b \leq \bullet b'$ is trivially satisfied. Thus, places in the initial marking are not equivalent to any other place and, in particular, the initial marking consists of pairwise non-equivalent places. If instead $\bullet b = \{e\}$, then b cannot be equated to any place b' caused by e . In particular, this implies that the places in the post-set of each event are pairwise non-equivalent. More generally, this condition plays a role in ensuring that the quotiented p-net is irredundant (Definition 3.5). By condition (1c) non-persistent places can be equivalent only when they are generated by equivalent events. Finally, condition (1d) states that persistent places can be equivalent even if they are not generated by equivalent events, but in this case they must be connected by a chain of consistency in the equivalence class. This ensures that, once quotiented, the pre-set of the condition will satisfy the connectedness condition (see Definition 4.3).

Condition (2) says that events are equivalent only when they have equivalent pre-sets and in this case they also have equivalent post-sets. Moreover, equivalent events must differ in their pre-set.

Finally, by condition (3), whenever two sets of conditions X, X' are equivalent then the sets of events having X and X' as pre-sets are equivalent. Note that since $\bullet e = X$, by condition (1b), the set X cannot contain equivalent conditions. The same applies to X' and thus $X \sim X'$ implies that there is a bijection from X to X' mapping each element $b_1 \in X$ to the only $b_2 \in X'$ such that $b_1 \sim b_2$. Moreover, also the set $\{e \mid \bullet e = X\}$ contains pairwise non-equivalent events by condition (2). The same applies to $\{e' \mid \bullet e' = X'\}$, hence, also in this case, the requirement $\{e \mid \bullet e = X\} \sim \{e' \mid \bullet e' = X'\}$ implies a one-to-one correspondence between equivalent events.

An example of occurrence p-net with equivalence is in Figure 5, where equivalent elements are linked by dotted lines.

The fact that occurrence p-nets with equivalence do not have backward conflicts allows us to restrict only to binary conflicts. Formally, the following holds.

Lemma 4.13 (conflicts in occurrence p-nets with equivalence). *Let $\langle O, \sim \rangle$ be an occurrence p-net with equivalence and $X \subseteq B \cup E$. Then $\#X$ holds if and only if there are $x, x' \in X$ such that $x \# x'$.*

Proof. Straightforward using the definition of conflict (see Definition 4.1). \square

For occurrence p-nets with equivalence we will need a notion of concurrency on sets of places stronger than that in Definition 4.9.

Definition 4.14 (strong concurrency). Let $\langle O, \sim \rangle$ be an occurrence p-net with equivalence. A subset of conditions $X \subseteq B$ is *strongly concurrent*, written $sco(X)$, if $co(X)$ and for all $b, b' \in X$, if $b \neq b'$ then $\neg(b \sim b')$.

The idea is that if a set of conditions is strongly concurrent it can be produced by a computation using only a specific instance for each equivalence class of persistent resources.

Lemma 4.15 (pre- and post-sets are strongly concurrent). *Let $\langle O, \sim \rangle$ be an occurrence p-net with equivalence. Then for any $e \in E$ its pre-set $\bullet e$ and post-set e^\bullet are strongly concurrent.*

Proof. Let $e \in E$ be an event of O . Then $\bullet e$ and e^\bullet are concurrent since O is an occurrence p-net, thus e can be fired and hence pre- and post-sets of events are coverable, whence concurrent by Lemma 4.10. They are also strongly concurrent, since, by Definition 4.12, condition (1b), they cannot contain equivalent conditions. \square

Occurrence p-nets with equivalence can be turned into a category by introducing a suitable notion of morphism.

Definition 4.16 (category of occurrence p-nets with equivalence). A morphism of occurrence p-nets with equivalence $f : \langle O, \sim \rangle \rightarrow \langle O', \sim' \rangle$ is a morphism $f : O \rightarrow O'$ of p-nets such that for all $b_1, b_2 \in B$, if $b_1 \sim b_2$ then $f_s(b_1) \sim' f_s(b_2)$. We denote by \mathbf{OE} the category of occurrence p-nets with equivalence.

In words, a morphism of occurrence p-nets is required to preserve the equivalence. This is essential to ensure that it induces a function on the quotient nets.

Definition 4.17 (quotient of an occurrence p-net with equivalence). Let $\langle O, \sim \rangle$ be an occurrence p-net with equivalence. Its *quotient* is the structure

$$\mathcal{Q}(\langle O, \sim \rangle) = ([B]_\sim, [B^p]_\sim, [E]_\sim, \gamma_0^\sim, \gamma_1^\sim, [v_0]_\sim)$$

where $\gamma_j^\sim([e]_\sim) = [\gamma_j(e)]_\sim$ for $j \in \{0, 1\}$.

The quotient of the net in Figure 5 is depicted in Figure 6. Note that the quotient has introduced a backward conflict on the place o .

We want to show that the quotient of an occurrence p-net with equivalence is indeed an occurrence p-net. To this aim we first observe some facts.

Lemma 4.18 (properties of the quotient). *Let $\langle O, \sim \rangle$ be an occurrence p-net with equivalence and $\mathcal{Q}(\langle O, \sim \rangle)$ the corresponding quotient. Then*

- (1) for all $b \in B^n$, $\bullet[b]_\sim = [\bullet b]_\sim$;
- (2) for all $b \in B^p$, $\bullet[b]_\sim = [\bigcup_{b' \sim b} \bullet b']_\sim$;
- (3) for all $e \in E$, $\bullet[e]_\sim = [\bullet e]_\sim$ and $[e]_\sim^\bullet = [e^\bullet]_\sim$;




FIGURE 6. The quotient of the occurrence p-net with equivalence in Figure 5

- (4) for all $X \subseteq E$ and $e \in E$, if $X \vdash e$ then $[X]_{\sim} \vdash [e]_{\sim}$;
- (5) for all $e, e' \in E$, if $e \sim e'$ and $e \neq e'$ then $e \# e'$;
- (6) for all $e, e' \in E$, if $[e]_{\sim} \# [e']_{\sim}$ then $e \# e'$.

Proof. The proofs are mostly routine.

For point (1) observe that, by Definition 4.12(1c), if $b, b' \in B^n$ then $\bullet b \sim \bullet b'$. Point (2) is just a consequence of the definition of quotient. For point (3) observe that, by definition, if $e \sim e'$ then $\bullet e \sim \bullet e'$ and $e \bullet \sim e' \bullet$.

Point (4) follows from the observation that, by (3), $\bullet[e]_{\sim} = [\bullet e]_{\sim}$. Moreover, for all $b \in \bullet e$, if $b \in v_0$ then $[b]_{\sim} \in [v_0]_{\sim}$. Otherwise, there is $e' \in X$ such that $b \in e' \bullet$ and thus, by (3), $[b]_{\sim} \in [e']_{\sim} \bullet$.

Concerning point (5), let $e \sim e'$. Since in a p-net $\bullet e' \not\subseteq B^p$, we can fix some $b \in \bullet e \cap B^n$. By definition of occurrence p-net with equivalence (Definition 4.12(1c)), $\bullet e \sim \bullet e'$, hence there must be $b' \in \bullet e' \cap B^n$ such that $b \sim b'$. We proceed by induction on the depth h of e . When $h = 0$, we have $b \in v_0$. Again, by Definition 4.12(1c), $\bullet b \sim \bullet b'$, and thus also $b' \in v_0$. Since $b \sim b'$ we deduce that $b = b'$ and thus $e \# e'$. If $h > 0$, there are $e_1 \in \bullet b$ and $e'_1 \in \bullet b'$ and, since $b, b' \in B^n$, it holds $e_1 \sim e'_1$. Hence by induction hypothesis $e_1 \# e'_1$ and thus $e \# e'$.

Finally, for point (6), we can proceed by induction on the derivation on rules that define conflict (see Definition 4.1). Let $[e]_{\sim} \# [e']_{\sim}$. If there is $[b]_{\sim} \in \bullet[e]_{\sim} \cap \bullet[e']_{\sim} \cap [B^n]_{\sim}$ then we conclude by point (3) that $e \# e'$. If instead there are $[b]_{\sim} \in \bullet[e]_{\sim} = [\bullet e]_{\sim}$ and $[b']_{\sim} \in \bullet[e']_{\sim} = [\bullet e']_{\sim}$ with $[b]_{\sim} \# [b']_{\sim}$, $b \in \bullet e$ and $b' \in \bullet e'$ it means that for any $[e_1]_{\sim} \in \bullet[b]_{\sim}$ and any $[e_2]_{\sim} \in \bullet[b']_{\sim}$ we have $[e_1]_{\sim} \# [e_2]_{\sim}$. In particular, for $e_1 \in \bullet b$ and $e_2 \in \bullet b'$ (that are uniquely determined as the underlying occurrence p-net as no backward conflict) by inductive hypothesis we have $e_1 \# e_2$ and thus $b \# b'$ and $e \# e'$. \square

We can now reach the desired conclusion.

Lemma 4.19 (quotient is an occurrence p-net). *Let $\langle O, \sim \rangle$ be an occurrence p-net with equivalence. Then $\mathcal{Q}(\langle O, \sim \rangle)$ is an occurrence p-net.*

Proof. Easy consequence of Lemma 4.18. In particular, for any event $e \in E$, by Lemma 4.18(3) we have that $\bullet[e]_{\sim} = [\bullet e]_{\sim}$. Then connectedness follows from Lemma 4.18(6). In order to conclude we have to observe that $\mathcal{Q}(\langle O, \sim \rangle)$ is a well-formed p-net. T-restrictedness follows immediately from t-restrictedness of O . Concerning irredundancy, consider a generic persistent place in $\mathcal{Q}(\langle O, \sim \rangle)$ that will be of the kind $[b]_{\sim}$ for $b \in B^p$ and take any event $[e]_{\sim} \in \bullet[b]_{\sim}$. By Lemma 4.18(2), $\bullet[b]_{\sim} = [\bigcup_{b' \sim b} \bullet b']_{\sim}$, hence we can assume $e \in \bullet b$. In order to violate irredundancy, there should exist $e' \sim e$ and $b' \sim b$ such that $e' \rightsquigarrow_O^n b'$ with $n \geq 2$. It is easy to see that this cannot happen thanks to conditions (1b) and (2) of Definition 4.12. \square

We can thus consider a quotient functor from the category of p-nets with equivalence to the category of occurrence p-nets.

Definition 4.20 (quotient functor). We denote by $\mathcal{Q}() : \text{OE} \rightarrow \text{ON}$ the functor taking an occurrence p-net with equivalence to its quotient.

The fact that the functor is well-defined on objects follows from Lemma 4.19. On arrows it is immediate from the definition of morphism of occurrence net with equivalence (Definition 4.16).

Lemma 4.21 (occurrence p-net with equivalence vs occurrence p-nets). *Let $\langle O, \sim \rangle$ be an occurrence p-net with equivalence. For any finite configuration $C \in \text{Conf}(O)$, $[C]_{\sim} \in \text{Conf}(\mathcal{Q}(O))$ and $m([C]_{\sim}) = [m(C)]_{\sim}$.*

Proof. Immediate from Lemma 4.18, points (4) and (6). \square

4.3. Unfolding. As mentioned above, the first phase of the unfolding construction produces an occurrence p-net with equivalence, which is then quotiented to an occurrence p-net. We will use π to denote the projection on the first component of a pair, i.e., given sets A and B , we let $\pi : A \times B \rightarrow A$ be defined as $\pi(a, b) = a$ for all $(a, b) \in A \times B$.

Definition 4.22 (unfolding). Let $N = (S, S^p, T, \delta_0, \delta_1, u_0)$ be a p-net. Define the *pre-unfolding* $\mathcal{U}^{pre}(N) = (B, B^p, E, \gamma_0, \gamma_1, v_0)$ as the least occurrence p-net with an equivalence \sim_N on $B \cup E$ such that

- $v_0 = \{\langle s, \perp \rangle \mid s \in u_0\} \subseteq B$
- if $t \in T$ and $X \subseteq B$ such that $sco(X)$ and $\pi(X) = \bullet t$ then $e = \langle t, X \rangle \in E$, $Y = \{\langle s, e \rangle \mid s \in t^\bullet\} \subseteq B$ and $\bullet e = X$, $e^\bullet = Y$.

The equivalence \sim_N is the least equivalence relation that satisfies

$$\frac{s \in S^p \quad e \frown_N e'}{\langle s, e \rangle \sim_N \langle s, e' \rangle} \quad \frac{X \sim_N X'}{\langle t, X \rangle \sim_N \langle t, X' \rangle} \quad \frac{e \sim_N e'}{\langle s, e \rangle \sim_N \langle s, e' \rangle}$$

The unfolding $\mathcal{U}(N)$ is obtained as the quotient $\mathcal{U}^{pre}(N)_{\sim_N}$ with the *folding morphism* $\varepsilon_N : \mathcal{U}(N) \rightarrow N$ defined as $\varepsilon_N([x]_{\sim_N}) = \pi(x)$.

Note that if the net N does not contain persistent places, the equivalence relation \sim_N is just the identity relation, strong concurrency coincide with concurrency and the pre-unfolding coincides with the ordinary unfolding. Observe that ε_N is injective on pre- and post-sets of transitions, as implied by the definition of p-net morphism.

As an example, consider the (fragment of the) pre-unfolding of our running example in Figure 7, where we have used convenient names for places and transitions in order to improve readability. Since o_1 and o_2 are instances of the same persistent place o , they are related by the equivalence. Then $\{o_1, r_1\} \sim \{o_2, r_1\}$ and $\{o_1, s\} \sim \{o_2, s\}$, thus $d_1 \sim d_2$ and $e_1 \sim e_2$ and the equivalence is propagated to their post-sets. The same pattern is iterated for any instance of r created by the subsequent firing of (the instances of) c . The corresponding unfolding is obtained as the quotient in Figure 8.

For the p-net in Figure 4 (left), the pre-unfolding is the occurrence p-net in Figure 4 (right). Note that o_1 and o_2 are not equivalent since their generating events are in conflict and are not connected by a chain of consistency. Since \sim_N here is the identity, the unfolding coincides with the pre-unfolding.

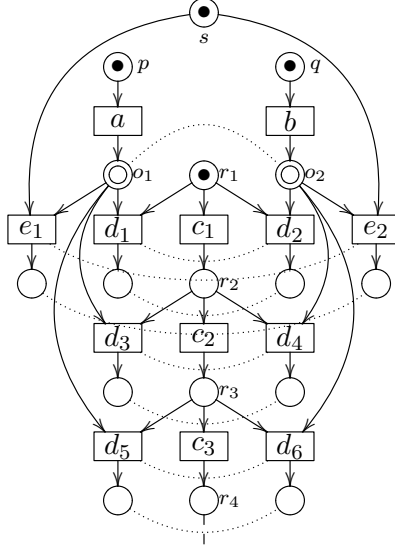


FIGURE 7. Pre-unfolding of our running example




FIGURE 8. Unfolding of our running example

We can consider the inclusion functor $\mathcal{I}() : \mathbf{ON} \rightarrow \mathbf{PN}$ that acts as identity on objects and morphisms. We next observe that the unfolding $\mathcal{U}(N)$ and the folding morphism ε_N are cofree over N . Therefore $\mathcal{U}()$ extends to a functor that is right adjoint of $\mathcal{I}()$ and establishes a coreflection between \mathbf{PN} and \mathbf{ON} .

It is easy to show that the pre-unfolding is an occurrence p-net with equivalence.

Lemma 4.23 (pre-unfolding is an occurrence p-net with equivalence). *Let N be a p-net. Then $\langle \mathcal{U}^{pre}(N), \sim_N \rangle$ is an occurrence p-net with equivalence.*

Proof. The fact that $\langle \mathcal{U}^{pre}(N), \sim_N \rangle$ satisfies the properties in Definition 4.12 follows almost directly by construction. The less immediate property is (1b), specifically the fact that if $b \neq b'$ then $\neg(\bullet b \leq \bullet b')$. This is a consequence of the fact that if $e \sim e'$ and $e \neq e'$ then $e \# e'$. \square

As a preliminary step we show that the pre-unfolding construction $\mathcal{U}^{pre}()$ extends to a functor that, together with the quotient functor $\mathcal{Q}()$, establishes an equivalence between the categories ON and OE.

Lemma 4.24 (mapping from the pre-unfolding). *Let O be an occurrence p-net and $\pi : \mathcal{U}^{pre}(O) \rightarrow O$ the mapping from the pre-unfolding. Then*

- (1) *for all finite $C \in \text{Conf}(O)$ there exists $C' \in \text{Conf}(\mathcal{U}^{pre}(O))$ such that $\pi(C') = C$ and $m(C) = \pi(m(C'))$;*
- (2) *for all $C', C'' \in \text{Conf}(\mathcal{U}^{pre}(O))$, if $\pi(C') = \pi(C'')$ then $C' \sim_N C''$.*

Proof. Point (1) can be shown by induction on $|C|$. The base case $|C| = 0$ is trivial. If $|C| = n > 0$, by Lemma 4.6, we have $C = \{e_1, \dots, e_n\}$ and there is a firing sequence $v_0[e_1]v_1[e_2] \dots [e_n]v_n$. By the same lemma $C_1 = \{e_1, \dots, e_{n-1}\}$ is a configuration of O . Hence by inductive hypothesis there exists a configuration $C'_1 \in \text{Conf}(\mathcal{U}^{pre}(O))$ such that $\pi(C'_1) = C_1$ and $v_n = m(C_1) = \pi(m(C'_1))$ and $co(m(C'_1))$. Since $\bullet e_n \subseteq v_{n-1}$, by definition of $\mathcal{U}^{pre}()$ there is an event $e'_n = \langle e_n, X \rangle$, where $X \subseteq v_{n-1}$ is such that $sco(X)$ and $\pi(X) = \bullet e_n$. Note that we can always choose X such that $sco(X)$ since equivalent conditions have the same π -image. If we define $C' = C_1 \cup \{e'_n\}$, we have that $\pi(C') = C$.

Point (2) can be proved by an easy induction on $|C'|$. \square

Proposition 4.25 (equivalence). *The categories ON and OE are equivalent via the functors $\mathcal{Q}()$ and $\mathcal{U}^{pre}()$.*

Proof. Let us first observe that, on objects, the functors are inverse of each other. First, for an occurrence p-net O let us define $q : \mathcal{Q}(\mathcal{U}^{pre}(O)) \rightarrow O$ by letting $q([x]_\sim) = \pi(x)$.

In order to show that q is an isomorphism we resort to Lemma 4.24. By (1) we can show that q is surjective. In fact, for each event e in O there is e' in $\mathcal{U}^{pre}(O)$ such that $\pi(e') = e$. Hence $q([e']_\sim) = e$. Moreover, it is injective, since by (2) it follows that if $q([x]_\sim) = q([y]_\sim)$ then $x \sim y$.

For the converse, define $\epsilon : \mathcal{U}^{pre}(\mathcal{Q}(\langle O, \sim \rangle)) \rightarrow \langle O, \sim \rangle$ inductively as follows.

The initial marking of $\mathcal{U}^{pre}(\mathcal{Q}(\langle O, \sim \rangle))$ consists of the places $b' = \langle [b]_\sim, \perp \rangle$ with $b \in v_0$. We define $\epsilon(b') = b$.

For the inductive step, consider an event $e' = \langle [e]_\sim, X' \rangle$ where $[e]_\sim$ is an event in $\mathcal{Q}(\langle O, \sim \rangle)$ and $X' \subseteq B'$ such that $sco(X)$ and $\pi(X) = \bullet[e]_\sim$. Since X is strongly concurrent and ϵ preserves strong concurrency (Lemma 4.11), we have that $\epsilon(X)$ is strong concurrent. Moreover, by Lemma 4.18(3) we have $\pi(X) = \bullet[e]_\sim = \bullet[e]_\sim$. Hence $\epsilon(X) \sim \bullet e$ and therefore, by definition of occurrence p-net with equivalence, there exists a unique e_1 in O such that $e_1 \sim e$ and $\bullet e_1 = \epsilon(X)$. It exists by Definition 4.12(3) and it is unique by Definition 4.12(2). We define $\epsilon(e') = e_1$. Finally, ϵ is extended to e'^\bullet by letting $\epsilon(\langle b, e' \rangle) = b_1$, where $b_1 \in e_1^\bullet$ is the unique condition in e_1^\bullet such that $b_1 \sim b$. In this way $\epsilon(e'^\bullet) = e_1^\bullet$.

It is easy to see that ϵ is injective. Surjectivity follows from Lemma 4.21.

We finally show that for any occurrence p-net O , for any occurrence p-net with equivalence $\langle O', \sim \rangle$ and for any morphism $f : \mathcal{Q}(\langle O', \sim \rangle) \rightarrow O$ there exists a unique morphism $h : \langle O', \sim \rangle \rightarrow \mathcal{U}^{pre}(O)$ such that the following diagram commutes

$$\begin{array}{ccc} \mathcal{Q}(\mathcal{U}^{pre}(O)) & \xrightarrow{\varepsilon_O} & O \\ \mathcal{Q}(h) \uparrow & \nearrow f & \\ \mathcal{Q}(\langle O', \sim \rangle) & & \end{array}$$

We define h by induction on the depth k of the items.

($k = 0$) Only conditions in the initial marking have depth 0. Let $b' \in B'$ with $d(b') = 0$. Hence $b' \in v'_0$ and thus $[b']_{\sim}$ is in the initial marking of $\mathcal{Q}(\langle O', \sim \rangle)$. Therefore $f_s([b']_{\sim}) = b \in v_0$ and we can define $h_s(b') = \langle b, \perp \rangle$, which is in the initial marking of $\mathcal{U}^{pre}(O)$.

($k > 0$) Let $e' \in E'$ be an event such that $d(e') = k > 0$. Therefore conditions in $\bullet e'$ have depth less than k . Hence their h -images have been already defined. Moreover, since $\bullet e'$ is concurrent, by Lemma 4.11, also $h_s(\bullet e')$ is. Moreover, on pre-sets morphisms preserve equivalence, hence h is injective on $\bullet e'$ and $h_s(\bullet e')$ does not include equivalent conditions, hence $sco(h_s(\bullet e'))$. Therefore the unfolding contains an event $e = \langle f_t(e'), h_s(\bullet e') \rangle$ and we define $h_t(e') = e$. The mapping is then extended to the post-set of e' by defining, for each $b \in e' \bullet$, $h_s(b) = \langle f_s([b]_{\sim}), e \rangle$.

Uniqueness follows by noticing that at each level we were forced to define h as we did to ensure commutativity. \square

We can finally prove the desired theorem, by showing that $\mathcal{U}^{pre}() : \text{PN} \rightarrow \text{OE}$ is right adjoint to $\mathcal{Q}() : \text{OE} \rightarrow \text{PN}$.

Proposition 4.26. $\mathcal{Q}() \dashv \mathcal{U}^{pre}()$

Proof. Let $N = (S, S^p, T, \delta_0, \delta_1, u_0)$ be a p-net, let $\mathcal{U}^{pre}(N) = \langle O, \sim_N \rangle$ be its pre-unfolding with $O = \langle B, B^p, E, \gamma_0, \gamma_1, v_0 \rangle$, and let $\varepsilon_N : \mathcal{Q}(\mathcal{U}^{pre}(N)) \rightarrow N$ be the folding morphism. We have to show that for any occurrence p-net with equivalence $\langle O', \sim \rangle$ with $O' = \langle B', B'^p, E', \gamma'_0, \gamma'_1, v'_0 \rangle$ and for any morphism $f : \mathcal{Q}(\langle O', \sim \rangle) \rightarrow N$ there exists a unique morphism $h : \langle O', \sim \rangle \rightarrow \mathcal{U}^{pre}(N)$ such that the following diagram commutes:

$$\begin{array}{ccc} \mathcal{Q}(\mathcal{U}^{pre}(N)) & \xrightarrow{\varepsilon_N} & N \\ \mathcal{Q}(h) \uparrow & \nearrow f & \\ \mathcal{Q}(\langle O', \sim \rangle) & & \end{array}$$

We define h by induction on the depth k of the items.

($k = 0$) Only conditions in the initial marking have depth 0. Let $b' \in B'$ with $d(b') = 0$. Hence $b' \in v'_0$ and thus $[b']_{\sim} = \{b'\}$ is in the initial marking of $\mathcal{Q}(\langle O', \sim \rangle)$. Therefore $f_s([b']_{\sim}) = b \in v_0$ and we can define $h_s(b') = \langle b, \perp \rangle$, which is in the initial marking of $\mathcal{U}^{pre}(O)$.

($k > 0$) Let $e' \in E'$ be an event such that $d(e') = k > 0$. Therefore conditions in $\bullet e'$ have depth less than k . Hence their h -images have been already defined. We have to define h_t on e' and h_s on its post-set.

- If $h_t([e']_{\sim})$ is undefined then necessarily also $h_t(e')$ is undefined and $h_s(b') = 0$ for all $b' \in e' \bullet$.

- If instead $h_t([e']_\sim) = t$ observe that since $\bullet e'$ is strongly concurrent by Lemma 4.11 also $h_s(\bullet e')$ is so. And, as argued in the proof of Proposition 4.25, $h_s(\bullet e')$ is strongly concurrent. Now, since f is a p-net morphism $f_s(\bullet[e']_\sim) = \bullet t$. Hence, we have

$$\begin{aligned} f_s(\bullet[e']_\sim) &= \\ &= \varepsilon_{N_s}(\mathcal{Q}(h)_s(\bullet[e']_\sim)) \quad [\text{by commutativity (up to depth } k)] \\ &= \varepsilon_{N_s}([h_s(\bullet e')]_\sim) \quad [\text{by definition of } \mathcal{Q}()] \end{aligned}$$

Recalling the definition of ε_N , the above equality implies that $\pi(\mathcal{Q}(h)_s(\bullet e')) = f_s(\bullet[e']_\sim) = \bullet t$. Hence the unfolding contains an event $e = \langle t, h_s(\bullet e') \rangle$ and we define $h_t(e') = e$. The mapping is then extended to the post-set of e' by defining, for each $b \in e'^\bullet$, $h_s(b) = \langle f_s([b]_\sim), e \rangle$.

Uniqueness follows by noticing that at each level we were forced to define h as we did to ensure commutativity. \square

Corollary 4.27 (coreflection between PN and ON). $\mathcal{I}() \dashv \mathcal{U}()$

Proof. Immediate corollary of Propositions 4.25 and 4.26. \square

5. LOCALLY CONNECTED EVENT STRUCTURES FROM OCCURRENCE P-NETS

In this section we show how an ES can be extracted from an occurrence p-net, thus providing an ES semantics to p-nets via the unfolding semantics. The transformation maps occurrence p-nets to locally connected ESs and it is shown to be functorial. Conversely, we show how a canonical occurrence p-net can be associated to any locally connected ES. The two transformations are shown to establish a coreflection.

An occurrence p-net can be easily mapped to an ES by forgetting the conditions and keeping the events and the enabling and conflict relations on events. The transformation gives rise to a functor from the category of occurrence p-nets to the category of locally connected ES.

Definition 5.1 (ES for an occurrence p-net). The functor $\mathcal{E}() : \text{ON} \rightarrow \ell\text{ES}$ is defined as follows. Let $O = (B, B^p, E, \delta_0, \delta_1, v_0)$ be an occurrence p-net. The corresponding ES is $\mathcal{E}(O) = (E, \vdash_{\mathcal{U}(N)}, \#_{\mathcal{U}(N)})$. For any morphism $f : O_1 \rightarrow O_2$ we let $\mathcal{E}(f) = f_t$.

The ES associated with a p-net is obtained from its unfolding by forgetting the places and keeping the events and their dependencies. We first show that $\mathcal{E}()$ is well-defined on objects.

Lemma 5.2 (occurrence p-nets to locally connected ES). *Let O be an occurrence p-net. Then $\mathcal{E}(O)$ is a locally connected ES.*

Proof. Let $e \in \mathcal{E}(O)$ be an event. Since O is an occurrence p-net, for all $b \in \bullet e$ the pre-set $\bullet b$ is connected. Moreover, if $\bullet b \neq \emptyset$ then $\bullet b$ is a disjunct for e . In fact, clearly, for each configuration $C \in \text{Conf}(\mathcal{E}(O))$ such that $C \vdash e$ we have $C \cap \bullet b \neq \emptyset$. Still $\bullet b$ might not be a disjunct for the failure of minimality, i.e., for the existence of an event $e' \in \bullet b$ such that $\bullet b \setminus \{e'\}$ still intersects any configuration enabling e . However, it is easy to see that this would imply the existence of an event $e'' \in \bullet b$ such that $e'' \rightsquigarrow_O e'$, violating the irredundancy assumption.

It is also immediate to see that $\{\bullet b \mid b \in \bullet e \wedge \bullet b \neq \emptyset\}$ is a covering. Hence we conclude. \square

The fact that $\mathcal{E}()$ is well-defined on morphisms is a consequence of the lemma below. For an occurrence p-net O , let us denote by \prec_O immediate causality on events, i.e., $e \prec_O e'$ if there is $b \in \bullet e'$ such that $\bullet b = \{e\}$. Note that causality \leq_O on events is the transitive closure of \prec_O .

Lemma 5.3 (properties of occurrence p-net morphisms). *Let O, O' be occurrence p-nets and $f : O \rightarrow O'$ a morphism. Then for all $C \in \text{Conf}(O)$ and $e_1, e_2 \in E$ with $f(e_1), f(e_2)$ defined*

- if $f_t(e_1) = f_t(e_2)$ and $e_1 \neq e_2$ then $e_1 \#_O e_2$;
- if $f_t(e_1) \#_{O'} f_t(e_2)$ then $e_1 \#_O e_2$;
- if $C \vdash_O e_1$ then $f_t(C) \vdash_{O'} f_t(e_1)$.

Proof. (1) Let $f_t(e_1) = f_t(e_2)$ and $e_1 \neq e_2$. Consider a causal chain of events in O' , starting from the initial marking and passing through non-persistent conditions only, namely consider $e'_1, e'_2, \dots, e'_{n-1}, e'_n = f_t(e_1) = f_t(e_2)$ and conditions $b'_1, \dots, b'_n \in B^n$ such that $b_1 \in v'_0$ and for all $i \in [1, n-1]$ it holds $b_i \in e'_i \bullet \cap \bullet e'_{i+1}$; note that such a chain exists because O' is t-restricted. It is easy to see that there must be corresponding causal chains $e_j^1, \dots, e_j^n = e_j$ and $b_j^1, \dots, b_j^n \in B^n$, for $j \in \{1, 2\}$, such that $b_j^1 \in v_0$ and for all $i \in [1, n-1]$ it holds $b_j^i \in e_j^i \bullet \cap \bullet e_j^{i+1}$, which are mapped to the causal chain in O' , i.e., $f_t(e_j^i) = e'_i$ and $b'_i \in f_s(b_j^i)$ for all $i \in [1, n]$, $j \in \{1, 2\}$.

Consider the least i such that $e_1^i \neq e_2^i$. If $i = 1$, i.e., $e_1^1 \neq e_2^1$ then, since $b_1^1, b_2^1 \in v_0$, $b_1^1 \in f_s(b_1^1)$ and $b_1^1 \in f_s(b_2^1)$, recalling that $f_s(v_0) = v'_0$ and b_1^1 , by safety of the initial marking, must have multiplicity at most 1, we conclude that necessarily $b_1^1 = b_2^1$ and thus $e_1^1 \#_O e_2^1$, hence, by inheritance, $e_1 = e_1^n \#_O e_2^n = e_2$, as desired.

Otherwise, if $i > 1$, i.e., $e_1^{i-1} = e_2^{i-1}$, then just observe that $b_1^{i-1}, b_2^{i-1} \in e_1^{i-1} \bullet = e_2^{i-1} \bullet$, and, by safety of $e_1^{i-1} \bullet$, as above, we conclude $e_1^i \#_O e_2^i$, and thus $e_1 \#_O e_2$, as desired.

(2) Let $f_t(e_1) \#_{O'} f_t(e_2)$ in O' . We proceed by induction on the length of the derivation of the conflict. If such length is 0, i.e., the conflict is direct, then there is $b' \in \bullet f_t(e_1) \cap \bullet f_t(e_2) \cap B^n$. Thus, for $j \in \{1, 2\}$, there are $b_j \in \bullet e_j \cap B^n$, such that $b' \in f_s(b_j)$. If $b_1 = b_2$ then $e_1 \#_O e_2$, and we conclude. Otherwise, observe that $b_1, b_2 \notin v_0$, otherwise (again by safety of the initial marking) they could not have a common image. Therefore $\bullet b_1 = \{e'_1\}$ and $\bullet b_2 = \{e'_2\}$ for suitable $e_1, e_2 \in E$ (the pre-set is a singleton since b_1 and b_2 are not persistent). Moreover $b' \in f_t(e'_1) \bullet = f_t(e'_2) \bullet$ hence it must be $f_t(e'_1) = f_t(e'_2)$. From point (1), $e'_1 \#_O e'_2$ and thus $e_1 \#_O e_2$, by inheritance.

If instead the conflict $f_t(e_1) \#_{O'} f_t(e_2)$ is inherited, we must have $b'_1 \in \bullet f_t(e_1)$ and $b'_2 \in \bullet f_t(e_2)$ such that $b'_1 \#_{O'} b'_2$. In turn this means that $\bullet b'_1, \bullet b'_2 \neq \emptyset$ and for all $e'_1 \in \bullet b'_1$, $e'_2 \in \bullet b'_2$, $e'_1 \#_{O'} e'_2$.

Consider $b_1 \in \bullet e_1$, $b_2 \in \bullet e_2$ such that $b'_1 \in f_s(b_1)$ and $b'_2 \in f_s(b_2)$. Necessarily, $b_1, b_2 \notin v_0$, hence $\bullet b_1, \bullet b_2 \neq \emptyset$. Moreover, for all $e_3 \in \bullet b_1$, $f_t(e_3) \in \bullet b'_1$ and for all $e_4 \in \bullet b_2$, $f_t(e_4) \in \bullet b'_2$. Hence, by the observation above, $f_t(e_3) \#_{O'} f_t(e_4)$, and the derivation of such conflict is shorter than that of $f_t(e_1) \#_{O'} f_t(e_2)$. Thus, by inductive hypothesis, $e_3 \#_O e_4$. By definition of conflict, this implies that $b_1 \#_O b_2$ and thus $e_1 \#_O e_2$, as desired.

(3) By Definition 4.1 we have to show that for each $b' \in \bullet f_t(e_1)$ either $b' \in u'_0$ or there is an $e' \in f_t(C)$ such that $b' \in e' \bullet$. Given $b' \in \bullet f_t(e_1)$, since f is a morphism there is a $b \in \bullet e_1$ such that $b' \in f_s(b)$. Since $C \vdash_O e_1$, either $b \in u_0$, and in this case $b' \in f_s(u_0) = u'_0$, or there is an $e \in C$ such that $b \in e \bullet$. In the last case we have $b' \in f_s(b) \subseteq f_s(e \bullet) = f_t(e) \bullet$, where $f_t(e) \in f_t(C)$, as desired. \square

Conversely, we show how to freely generate an occurrence p-net from a locally connected ES. Roughly, the idea is to insert suitable conditions that induce exactly the dependencies (enabling and conflict) of the original ES.

Definition 5.4 (occurrence p-net for an ES). Let $(E, \vdash, \#)$ be a locally connected ES. We define the occurrence p-net $\mathcal{O}(E) = (B, B^p, E, \gamma_0, \gamma_1, u_0)$ as follows. The set of places B consists of

- non-persistent places $\langle X, Y \rangle$ with $X, Y \subseteq E$, $|X| \leq 1$ such that $e < e'$ for all $e \in X$, $e' \in Y$ and $e' \# e''$ for all $e', e'' \in Y$, $e' \neq e''$;
- persistent places $\langle X, Y \rangle$ with $X, Y \subseteq E$, X disjunct of all $e \in Y$ and $X \curvearrowright$ -connected.

Furthermore, for all $e \in E$ let $\gamma_0(e) = \{\langle X, Y \rangle \in B \mid e \in Y\}$, and $\gamma_1(e) = \{\langle X, Y \rangle \in B \mid e \in X\}$. Finally, let the initial marking be $u_0 = \{\langle \emptyset, Y \rangle \mid Y \subseteq E \wedge \langle \emptyset, Y \rangle \in B\}$.

The intuition is the following. For any possible set of events Y pairwise in conflict that have a common cause e we insert a non-persistent place $b = \langle \{e\}, Y \rangle$, produced by e and consumed by the events in Y , inducing such dependencies. By the same clause, for any possible set of events Y pairwise in conflict, we insert a non-persistent place $b = \langle \emptyset, Y \rangle$ consumed by the events in Y . Moreover, for any pair of sets of events $\langle X, Y \rangle$ such that X includes an event for each minimal enabling of each $e \in Y$, we introduce a persistent condition $\langle X, Y \rangle$ that is generated by all events in X and used by all events in Y . In this way, whenever a minimal enabling set for some $e \in Y$ has been executed, all the pre-set of e is covered. Conversely, when all the pre-set of e is covered, since we generate conditions that include at least one event for each minimal enabling of each $e \in Y$, at least one minimal enabling has been completely executed. The request that for $b = \langle X, Y \rangle \in B^p$ the set $X = \bullet b$ is a disjunct connected by \curvearrowright will ensure that the pre-set of conditions is connected, as required by the definition of occurrence p-net. The fact that X is a disjunct will guarantee irredundancy. Formally, the fact that the construction above produces a well-defined occurrence p-net will be a consequence of Lemma 5.5.

We next observe that for all locally connected ESS, if we build the corresponding occurrence p-net and then we take the underlying ES we get an ES isomorphic to the original one. First we prove a technical result.

Lemma 5.5 (dependencies in the occurrence p-net for an ES). *Let E be a locally connected ES and $\mathcal{O}(E) = (B, B^p, E, \gamma_0, \gamma_1, v_0)$. Then*

- (1) *for all $e, e' \in E$, $e \# e'$ iff $e \#_{\mathcal{O}(E)} e'$;*
- (2) *for all $C \in \text{Conf}(E)$, $C \vdash e$ implies $C \vdash_{\mathcal{O}(E)} e$;*
- (3) *for all $C \in \text{Conf}(\mathcal{O}(E))$, $C \vdash_{\mathcal{O}(E)} e$ implies $C \vdash e$.*

Proof. (1) *Only if part.* Suppose that $e \# e'$ for $e, e' \in E$. Then, by Definition 5.4, $\mathcal{O}(E)$ contains a non-persistent place $\langle \emptyset, \{e, e'\} \rangle \in \bullet e \cap \bullet e'$, thus $e \#_{\mathcal{O}(E)} e'$.

If part. We prove, more generally, that for $A \subseteq E$, if $\#_{\mathcal{O}(E)} A$ then there is no configuration $C \in \text{Conf}(E)$ such that $A \subseteq C$. Then for binary conflict the thesis follows from the fact that E is saturated.

We proceed by induction on the length of the derivation of the conflict. If the length is 0, i.e., the conflict is direct, there are $e, e' \in A$ and $b \in \bullet e \cap \bullet e' \cap B^n$. By Definition 5.4 this means that $b = \langle \{e''\}, Y \rangle$ with $e, e' \in Y$ and therefore $e \# e'$ by construction and we are done.

If instead the conflict $\#_{\mathcal{O}(E)} A$ is inherited, we must have $e \in A$ such that $\#_{\mathcal{O}(E)} (A \setminus \{e\}) \cup \bullet e$. This means that if we consider any $A_e \subseteq E$ such that $A_e \cap \bullet b \neq \emptyset$ for all $b \in \bullet e$

then $\#_{\mathcal{O}(E)}(A \setminus \{e\}) \cup A_e$. The derivations of the latter conflicts are shorter than that of $\#_{\mathcal{O}(E)} A$, and thus by induction hypothesis we can infer that there is no configuration $C \in \text{Conf}(E)$ such that $(A \setminus \{e\}) \cup A_e \subseteq C$. Now, assume by absurd that $A \subseteq \hat{C}$ for a configuration \hat{C} . Since $e \in \hat{C}$, there is a configuration $\hat{C}_e \subset \hat{C}$ such that $\hat{C}_e \vdash e$. By construction $\{X \mid X \neq \emptyset \wedge b = \langle X, Y \rangle \in \bullet e\} = \{\bullet b \mid b \in \bullet e \wedge \bullet b \neq \emptyset\}$ is a covering of e , therefore $\hat{C}_e \cap \bullet b \neq \emptyset$ for all $b \in \bullet e$. But clearly $(A \setminus \{e\}) \cup \hat{C}_e \subseteq \hat{C}$, contradicting the inductive hypothesis.

(2) Let $C \in \text{Conf}(E)$ such that $C \vdash e$. In order to conclude that $C \vdash_{\mathcal{O}(E)} e$ we have to prove that for all $b \in \bullet e$, either $b \in v_0$ or there exists $e' \in C$ such that $b \in e'^\bullet$. We distinguish two cases.

If b is not persistent, then $b = \langle \{e'\}, Y \rangle$ with $e \in Y$, hence $e' < e$. Then $b \in e'^\bullet$.

If b is persistent, then $b = \langle X, Y \rangle$ with $e \in Y$ and, by definition of $\mathcal{O}(E)$, $X \cap C \neq \emptyset$. Let $e' \in X \cap C$. Then $e' \in C$ and $b \in e'^\bullet$.

(3) We prove that for all finite $C \in \text{Conf}(\mathcal{O}(E))$ it holds that $C \in \text{Conf}(E)$ and, for any $e \in E$, if $C \vdash_{\mathcal{O}(E)} e$ then $C \vdash e$.

The proof proceeds by induction on the cardinality $|C|$. The base case $|C| = 0$ is immediate. If $|C| > 0$, we know by Lemma 4.6 that $C = \{e_1, \dots, e_n\}$ and there exists a firing sequence $v_0[e_1]v_1[e_2] \dots [e_n]v_n$. Since $C' = \{e_1, \dots, e_{n-1}\} \in \text{Conf}(\mathcal{O}(E))$ and $C' \vdash_{\mathcal{O}(E)} e_n$, by inductive hypothesis $C' \in \text{Conf}(E)$ and $C' \vdash e_n$. From this and point (1), we deduce that $C = C' \cup \{e\} \in \text{Conf}(E)$.

For the second part, let $e \in E$ be such that $C \vdash_{\mathcal{O}(E)} e$. By the first part, $C \in \text{Conf}(E)$ and by definition of enabling in occurrence p-nets, for all $b \in \bullet e$, either $b \in v_0$ or there exists $e' \in C$ such that $b \in e'^\bullet$. This means that $C \cap \bullet b \neq \emptyset$ for all $b \in \bullet e$, $\bullet b \neq \emptyset$ and this is a covering of e . Hence by definition of covering, $C \vdash e$. \square

Corollary 5.6 ($\mathcal{O}(E)$ is well-defined). *Let E be a locally connected ES. Then $\mathcal{O}(E)$ is a well-defined occurrence p-net.*

Proof. The fact that $\mathcal{O}(E)$ satisfies the structural properties (a) and (c) of Definition 4.3 is immediate by construction. Property (b), i.e., the fact that each transition $e \in T$ admits a securing sequence follows by the analogous property of ES, recalling that by Lemma 5.5, enabling and conflict coincide in E and $\mathcal{O}(E)$. Similarly, the fact that for each condition $b \in B$, the set $\bullet b$ is connected is true by construction after Lemma 5.5, showing that conflict, and thus the notion of connectedness, coincide in E and $\mathcal{O}(E)$. Finally, we observe that $\mathcal{O}(E)$ is well-formed. T-restrictedness holds by construction. Irredundancy follows from the fact that for each $b = \langle X, Y \rangle \in B^p$, the pre-set $\bullet b = X$ is a disjunct. Hence given $e \in X = \bullet b$ there cannot be an additional path $e \rightsquigarrow_{\mathcal{O}(E)}^n$ with $n \geq 2$ into b . In fact, if this were the case, since by Lemma 5.5, enabling coincides in E and $\mathcal{O}(E)$, event e could be omitted, i.e., $X' = X \setminus \{e\}$ would still intersect any configuration enabling e , contradicting the minimality of X . \square

Corollary 5.7 (unit). *Let E be a locally connected ES. Then $\eta_E : E \rightarrow \mathcal{E}(\mathcal{O}(E))$ defined as the identity on events is an isomorphism.*

Proof. The function η_E is obviously a bijection. The fact that it is an isomorphism of ES follows immediately from Lemma 5.5. \square

In order to conclude we need to show that the construction of the occurrence p-net associated with an ES E and the isomorphism $\eta_E : E \rightarrow \mathcal{E}(\mathcal{O}(E))$ are free over E . The next lemma states some properties of occurrence p-net morphisms that guide the proof.

Lemma 5.8. *Let $h : \mathcal{O} \rightarrow \mathcal{O}'$ be an occurrence p-net morphism. Then for $b \in B$ and $b' \in B'$ such that $b' \in h_s(b)$ it holds*

- (1) *if $b \in B^n$ then (a) $h_t(\bullet b) = \bullet b'$ and $b^\bullet = \{e_1 \in E \mid e_1 \in h_t^{-1}(b^\bullet) \wedge \bullet b \leq e_1\}$, where $\bullet b \leq e_1$ means $e \leq e_1$ when $\bullet b = \{e\}$ and it is a vacuous requirement, otherwise;*
- (2) *if $b \in B^p$ then (a) $\bullet b \subseteq h_t^{-1}(\bullet b')$ and (b) $b^\bullet \subseteq h_t^{-1}(b'^\bullet)$.*

Proof. Point (1) is analogous to the one in the ordinary case. Point (2) easily follows from the definition of morphism. \square

Theorem 5.9 (coreflection between $\mathcal{O}()$ and $\mathcal{E}()$). *The construction $\mathcal{E}()$ extends to a functor that is left adjoint to $\mathcal{O}()$, and they establish a coreflection.*

Proof. Let E be a locally connected ES. We show the freeness of $\eta_E : E \rightarrow \mathcal{E}(\mathcal{O}(E))$ as defined in Corollary 5.7. We have to show that for any occurrence p-net $\mathcal{O}' = \langle B', B'^p, E', \gamma'_0, \gamma'_1, v'_0 \rangle$ and for any ES morphism $f : E \rightarrow \mathcal{E}(\mathcal{O}')$ there exists a unique morphism $h : \mathcal{O}(E) \rightarrow \mathcal{O}'$ such that the following diagram commutes

$$\begin{array}{ccc} E & \xrightarrow{\eta_E} & \mathcal{E}(\mathcal{O}(E)) \\ & \searrow f & \downarrow \mathcal{E}(h) \\ & & \mathcal{E}(\mathcal{O}') \end{array}$$

The transition component of h is determined as $h_t = f$. It can be extended to a p-net morphism by defining the place component as follows:

- (1) on non-persistent places $b = \langle X, Y \rangle$, define

$$h_s(b) = \{b' \in B'^n \mid \bullet b' = h_t(X) \wedge Y = \{e_1 \in E \mid e_1 \in h_t^{-1}(b'^\bullet) \wedge X < e_1\}\}.$$

- (2) on persistent places $b = \langle X, Y \rangle$, say that b is a potential pre-image of $b' \in B'^n$ if $X \subseteq h_t^{-1}(\bullet b')$ and $Y \subseteq h_t^{-1}(b'^\bullet)$. It is maximal if for any potential pre-image $b_1 = \langle X_1, Y_1 \rangle$ of b' , with $X_1 \cap X \neq \emptyset$ it holds $Y_1 \subseteq Y$. Then we define

$$h_s(b) = \{b' \in B'^n \mid b \text{ maximal potential pre-image of } b'\}.$$

It is not difficult to prove that this is a morphism. Uniqueness follows from the fact that, for non-persistent places (point (1)) we defined the morphism in the only possible way, according to Lemma 5.8(1). For persistent places, each $b' \in B'^p$ must be in the image of exactly one potential pre-image. In fact, according to Lemma 5.8(2), it must be in the image of at least one potential pre-image and, by injectivity on pre- and post-sets of transitions, it can be in the image of at most one of the potential pre-images. Using preservation of pre- and post-sets, we then conclude that $h_s(s)$ must be defined as we did in point (2). \square

As a side remark, note that the ES induced by $\mathcal{U}^{pre}(N)$ (before quotienting) is a prime ES. If endowed with the equivalence \sim_N restricted to events, it corresponds to the notion of prime ES with equivalence [10, 3]. More precisely, if we denote by $\#_N^\mu$ the immediate conflict relation, then the prime ES obtained from the pre-unfolding, by forgetting the conditions and the direct conflicts between equivalent events, i.e., $\langle (E, \leq_{\mathcal{U}(N)}, \#), \sim \rangle$, where $\# = \#_{\mathcal{U}(N)} \setminus (\#_{\mathcal{U}(N)}^\mu \cap \sim_N)$ and $\sim = \sim_N \cap (E \times E)$, is a prime ES in the sense of [3]. Its events arise, as explained in [3], as the irreducible elements of the domain associated with N .

6. CONCLUSIONS

Persistence is the continuance of an effect after its causes ceased to exist. In this paper we have studied the effect of adding persistence to Petri nets at the level of ES semantics. Interestingly, we have extended Winskel's chain of coreflection from the category of p-nets to the newly defined category of locally connected ESS, which is a full subcategory of the category of general ESS. Since the category of connected ESS is included in the one of locally connected ESS, the coreflection can serve to explain in basic terms the phenomenon of fusion arising in the context of graph grammars and that induces (connected) disjunctive causes. On the one hand, this confirms our intuition that Petri nets and their natural extensions keep capturing all phenomena of concurrency within easy-to-understand operational models. On the other hand, our results show that while non-prime ESS were actually underestimated in the literature, at least in some cases they are natural, expressive and equipped with an interesting theory even at the operational level.

The result has been proved for the class of well-formed persistent nets, where redundant paths to persistent places are forbidden. Despite the fact that this is a natural restriction, preliminary investigations suggest that the result could be extended, at the price of some technical complications, to the more general setting.

ACKNOWLEDGMENT

The authors wish to thank the anonymous reviewers for their careful reading and their insightful comments on the submitted version on the paper.

REFERENCES

- [1] Samson Abramsky. Petri nets, discrete physics, and distributed quantum computation. In Pierpaolo Degano, Rocco De Nicola, and José Meseguer, editors, *Concurrency, Graphs and Models, Essays Dedicated to Ugo Montanari on the Occasion of His 65th Birthday*, volume 5065 of *Lecture Notes in Computer Science*, pages 527–543. Springer, 2008.
- [2] Andrea Asperti and Nadia Busi. Mobile Petri nets. *Mathematical Structures in Computer Science*, 19(6):1265–1278, 2009.
- [3] Paolo Baldan, Andrea Corradini, and Fabio Gadducci. Domains and event structures for fusions. In *LICS 2017*, pages 1–12. IEEE Computer Society, 2017. Extended version available as arXiv:1701.02394.
- [4] Paolo Baldan, Andrea Corradini, and Ugo Montanari. Contextual Petri nets, asymmetric event structures and processes. *Information and Computation*, 171(1):1–49, 2001.
- [5] Filippo Bonchi, Antonio Brogi, Sara Corfini, and Fabio Gadducci. A net-based approach to web services publication and replaceability. *Fundamenta Informaticae*, 94(3-4):305–330, 2009.
- [6] Roberto Bruni, Hernán C. Melgratti, and Ugo Montanari. Concurrency and probability: Removing confusion, compositionally. In *LICS 2018*, pages 195–204. ACM, 2018.
- [7] Søren Christensen and Niels Damgaard Hansen. Coloured Petri nets extended with place capacities, test arcs and inhibitor arcs. In Marco Ajmone Marsan, editor, *ICATPN 1993*, volume 691 of *Lecture Notes in Computer Science*, pages 186–205. Springer, 1993.
- [8] Federico Crazzolara and Glynn Winskel. Events in security protocols. In Michael K. Reiter and Pierangela Samarati, editors, *CCS 2001*, pages 96–105. ACM, 2001.
- [9] Federico Crazzolara and Glynn Winskel. Petri nets with persistence. In Nadia Busi, Roberto Gorrieri, and Fabio Martinelli, editors, *WISP 2004*, volume 121 of *Electronic Notes in Theoretical Computer Science*, pages 143–155. Elsevier, 2005.
- [10] Marc de Visme and Glynn Winskel. Strategies with parallel causes. In Valentin Goranko and Mads Dam, editors, *CSL 2017*, volume 82 of *LIPICs*, pages 41:1–41:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

- [11] Joanne Bechta Dugan, Kishor S. Trivedi, Robert Geist, and Victor F. Nicola. Extended stochastic Petri nets: Applications and analysis. In Erol Gelenbe, editor, *Performance 1984*, pages 507–519. North-Holland, 1985.
- [12] Yrvann Emzivat, Benoît Delahaye, Didier Lime, and Olivier H. Roux. Probabilistic time Petri nets. In Fabrice Kordon and Daniel Moldt, editors, *ICATPN 2016*, volume 9698 of *Lecture Notes in Computer Science*, pages 261–280. Springer, 2016.
- [13] Claude Girault and Rüdiger Valk. *Petri Nets for Systems Engineering - A Guide to Modeling, Verification, and Applications*. Springer, 2003.
- [14] Roberto Gorrieri. *Process Algebras for Petri nets - The Alphabetization of Distributed Systems*. Springer, 2017.
- [15] Kurt Jensen. Coloured Petri nets. In Wilfried Brauer, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part I*, volume 254 of *Lecture Notes in Computer Science*, pages 248–299. Springer, 1987.
- [16] Kurt Jensen and Lars Michael Kristensen. *Coloured Petri Nets - Modelling and Validation of Concurrent Systems*. Springer, 2009.
- [17] Ina Koch. Petri nets in systems biology. *Software and System Modeling*, 14(2):703–710, 2015.
- [18] Ina Koch, Wolfgang Reisig, and Falk Schreiber. Petri nets in the biosciences. *it - Information Technology*, 56(2):43–45, 2014.
- [19] Lawrence H. Landweber and Edward L. Robertson. Properties of conflict-free and persistent Petri nets. *Journal of the ACM*, 25(3):352–364, 1978.
- [20] Marco Ajmone Marsan, Gianni Conte, and Gianfranco Balbo. A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems. *ACM Transactions on Computer Systems*, 2(2):93–122, 1984.
- [21] Philip M. Merlin. *A study of the recoverability of computing systems*. PhD thesis, University of California, Irvine, CA, 1974.
- [22] José Meseguer, Ugo Montanari, and Vladimiro Sassone. On the semantics of Petri nets. In Rance Cleaveland, editor, *CONCUR 1992*, volume 630 of *Lecture Notes in Computer Science*, pages 286–301. Springer, 1992.
- [23] José Meseguer, Ugo Montanari, and Vladimiro Sassone. On the semantics of place/transition Petri nets. *Mathematical Structures in Computer Science*, 7(4):359–397, 1997.
- [24] Michael K. Molloy. Discrete time stochastic Petri nets. *IEEE Transaction on Software Engineering*, 11(4):417–423, April 1985.
- [25] Ugo Montanari and Francesca Rossi. Contextual nets. *Acta Informatica*, 32(6):545–596, 1995.
- [26] Tadao Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, 1989.
- [27] Carl Adam Petri. *Kommunikation mit Automaten*. PhD thesis, Universität Hamburg, 1962.
- [28] Chander Ramchandani. *Analysis of asynchronous concurrent systems by timed Petri nets*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1973.
- [29] Wolfgang Reisig. *Understanding Petri Nets - Modeling Techniques, Analysis Methods, Case Studies*. Springer, 2013.
- [30] Rüdiger Valk. Petri nets as token objects: An introduction to elementary object nets. In Jörg Desel and Manuel Silva Suárez, editors, *ICATPN 1998*, volume 1420 of *Lecture Notes in Computer Science*, pages 1–25. Springer, 1998.
- [31] Wil M. P. van der Aalst and Christian Stahl. *Modeling Business Processes - A Petri Net-Oriented Approach*. MIT Press, 2011.
- [32] Walter Vogler, Alexei L. Semenov, and Alexandre Yakovlev. Unfolding and finite prefix for nets with read arcs. In Davide Sangiorgi and Robert de Simone, editors, *CONCUR 1998*, volume 1466 of *Lecture Notes in Computer Science*, pages 501–516. Springer, 1998.
- [33] Glynn Winskel. Event Structures. In Wilfried Brauer, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part II*, volume 255 of *Lecture Notes in Computer Science*, pages 325–392. Springer, 1987.