

Event Structure Semantics of Parallel Extrusion in the Pi-calculus

Silvia Crafa¹, Daniele Varacca², and Nobuko Yoshida³

¹ Dip. di Matematica - Univ. Padova

² PPS - CNRS & Univ. Paris Diderot

³ Dept. of Computing - Imperial College London

Abstract. We give a compositional event structure semantics of the π -calculus. The main issues to deal with are the communication of free names and the extrusion of bound names. These are the source of the expressiveness of the π -calculus, but they also allow subtle forms of causal dependencies. We show that free name communications can be modeled in terms of “incomplete/potential synchronization” events. On the other hand, we argue that it is not possible to satisfactorily model parallel extrusion within the framework of stable event structures. We propose to model a process as a pair (\mathcal{E}, X) where \mathcal{E} is a prime event structure and X is a set of (bound) names. Intuitively, \mathcal{E} encodes the structural causality of the process, while the set X affects the computation on \mathcal{E} so as to capture the causal dependencies introduced by scope extrusion. The correctness of our true concurrent semantics is shown by an operational adequacy theorem with respect to the standard late semantics of the π -calculus.

1 Introduction

In the study of concurrent and distributed systems, the true-concurrent semantics approach takes concurrency as a primitive concept rather than reducing it to nondeterministic interleaving. One of the by-products of this approach is that the causal links between the process actions are more faithfully represented in true-concurrent models.

Prime event structures [14] are a causal model for concurrency which is particularly suited for the traditional process calculi such as CCS and CSP since they directly represent causality and concurrency simply as a partial order and an irreflexive binary relation. Winskel [18] proposed a compositional event structure semantics of CCS, that has been proved to be operationally adequate with respect to the standard labelled transition semantics, hence sound with respect to bisimilarity [20]. Similar results have been proved for variants of the π -calculus, namely for a restricted typed subcalculus [17] and for the internal π I-calculus [5], which are however less expressive than the full calculus. In this paper we extend this result to the full π -calculus.

The main issues when dealing with the full π -calculus are name passing and the extrusion of bound names. These two ingredients are the source of the expressiveness of the calculus, but they are problematic in that they allow complex forms of causal dependencies, as detailed below.

1.1 Free Name Passing

Compared to pure CCS, (either free or bound) name passing adds the ability to dynamically acquire new synchronization capabilities. For instance consider the π -calculus process $P = n(z).(\bar{z}\langle a \rangle \mid m(x))$, that reads from the channel n and uses the received name to output the name a in parallel with a read action on m . Hence a synchronization along the channel m is possible if a previous communication along the channel n substitutes the variable z exactly with the name m . Then, in order to be compositional, the semantics of P must also account for “potential” synchronizations that might be activated by parallel compositions, like the one on channel m .

To account for this phenomenon, we define the parallel composition of event structures so that synchronization events that involve input and output on different channels, at least one of which is a variable, are not deleted straight away. Moreover, the events produced by the parallel composition are relabelled by taking into account their causal history. For instance, the event corresponding to the synchronization pair $(\bar{z}\langle a \rangle, m(x))$ is relabelled into a τ action if, as in the process P above, its causal history contains a synchronization that substitutes the variable z with the name m .

1.2 The Causal Links Created by Name Extrusion

Causal dependencies in π -calculus processes arise in two ways [1, 8]: by nesting prefixes (called *structural* or *prefixing* or *subject* causality) and by using a name that has been bound by a previous action (called *link* or *name* or *object* causality). While subject causality is already present in CCS, object causality is distinctive of the π -calculus. The interactions between the two forms of causal dependencies are quite complex. We illustrate them by means of examples.

Parallel Scope Extrusion. Consider the two processes $P = (\nu n)(\bar{a}\langle n \rangle.n(x))$ and $Q = (\nu n)(\bar{a}\langle n \rangle \mid n(x))$. The causal dependence of the action $n(x)$ on the output $\bar{a}\langle n \rangle$ is clear in the process P (i.e. there is a structural causal link), however, a similar dependence appears also in Q since a process cannot synchronize on the fresh name n before receiving it along the channel a (i.e. there is an objective causal link). Now consider the process $P_1 = (\nu n)(\bar{a}\langle n \rangle \mid \bar{b}\langle n \rangle)$: in the standard interleaving semantics of π -calculus only one output extrudes, either $\bar{a}\langle n \rangle$ or $\bar{b}\langle n \rangle$, and the other one does not. As a consequence, the second (free) output *depends* on the previous extruding output. However, in a true concurrent model we can hardly say that there is a dependence between the two parallel outputs, which in principle could be concurrently executed resulting in the parallel/simultaneous extrusion of the same name n to two different threads reading respectively on channel a and on channel b .

Dynamic Addition of New Extruders. We have seen that a bound name may have multiple extruders. In addition, the coexistence of free and bound outputs allows the set of extruders to dynamically change during the computation. Consider the process $P_2 = (\nu n)(\bar{a}\langle n \rangle \mid n(z)) \mid a(x).(\bar{x}\langle b \rangle \mid \bar{c}\langle x \rangle)$. It can either open the scope of n by extruding it along the channel a , or it can evolve to the process $(\nu n)(n(z) \mid \bar{n}\langle b \rangle \mid \bar{c}\langle n \rangle)$ where the output of the variable x has become a new extruder for both the actions with subject n . Hence after the first synchronization there is still the possibility of opening the scope of n by extruding it along the channel c .

The Lesson we Learned. The examples above show that the causal dependencies introduced by the scope extrusion mechanisms distinctive of the π -calculus can be understood in terms of the two ingredients of extrusion: name restriction and communication.

1. The restriction $(\nu n)P$ adds to the semantics of P a causal dependence between *every* action with subject n and *one of* the outputs with object n .
2. The communication of a restricted name adds *new* causal dependencies since both new extruders and new actions that need an extrusion may be generated by variable substitution.

A causal semantics for the π -calculus should account for such a dynamic additional objective causality introduced by scope extrusion. In particular, the first item above hints at the fact that we have to deal with a form of disjunctive (objective) causality. Prime event structures are stable models that represent disjunctive causality by duplicating events and so that different copies causally depend on different (alternative) events. In our case this amounts to represent different copies of any action with a bound subject, each one causally depending on different (alternative) extrusions. However, the fact that the set of extruders dynamically changes complicates the picture since new copies of any action with a bound subject should be dynamically spawned for each new extruder. In this way the technical details quickly become intractable, as discussed in Section 6.

In this paper we follow a different approach, that leads to an extremely simple technical development. The idea is to represent the disjunctive objective causality in a so-called inclusive way: in order to trace the causality introduced by scope extrusion it is sufficient to ensure that whenever an action with a bound subject is executed, at least one extrusion of that bound name must have been already executed, but it is not necessary to record which output was the real extruder. Clearly, such an inclusive-disjunctive causality is no longer representable with stable structures like prime event structures. However, we show that an operational adequate true concurrent semantics of the π -calculus can be given by encoding a π -process simply as a pair (\mathcal{E}, X) where \mathcal{E} is a prime event structure and X is a set of (bound) names. Intuitively, the causal relation of \mathcal{E} encodes the structural causality of a process. Instead, the set X affects the computation on \mathcal{E} : we define a notion of *permitted configurations*, ensuring that any computation that contains an action whose subject is a bound name in X , also contains a previous extrusion of that name. Hence a further benefit of this semantics is that it clearly accounts for both forms of causality: subjective causality is captured by the causal relation of event structures, while objective causality is implicitly captured by permitted configurations.

2 The π -calculus

In this section we illustrate the synchronous, monadic π -calculus that we consider. We assume a countably-infinite set of names and a countably-infinite set of variables ranged over by m, \dots, q and by x, \dots, z , respectively. Let a, b, c range over both names and variables.

$$\begin{array}{ll}
 \text{Prefixes} & \pi ::= a(x) \mid \bar{a}(b) \\
 \text{Processes} & P, Q ::= \sum_{i \in I} \pi_i.P_i \mid P \mid Q \mid (\nu n)P \mid A(\tilde{x}, \tilde{p} \mid \mathbf{z}, \mathbf{n})
 \end{array}
 \quad
 \begin{array}{l}
 \text{Definitions} \\
 A(\tilde{x}, \tilde{p} \mid \mathbf{z}, \mathbf{n}) = P_A
 \end{array}$$

The syntax consists of the parallel composition, name restriction, finite summation of guarded processes and recursive definition. In $\sum_{i \in I} \pi_i.P_i$, I is a finite indexing set; when

$$\begin{array}{c}
\text{(IN LATE)} \qquad \text{(OUT)} \qquad \text{(COMM)} \\
\frac{}{a(x).P \xrightarrow{a(x)} P} \qquad \frac{}{\bar{a}(b).P \xrightarrow{\bar{a}(b)} P} \qquad \frac{P \xrightarrow{a(x)} P' \quad Q \xrightarrow{\bar{a}(b)} Q'}{P | Q \xrightarrow{\tau} P'\{b/x\} | Q'} \\
\\
\text{(PAR)} \qquad \text{(OPEN)} \qquad \text{(CLOSE)} \\
\frac{P \xrightarrow{\alpha} P'}{P | Q \xrightarrow{\alpha} P' | Q} \qquad \frac{P \xrightarrow{\bar{a}(n)} P' \quad n \neq a}{(vn)P \xrightarrow{\bar{a}(n)} P'} \qquad \frac{P \xrightarrow{a(x)} P' \quad Q \xrightarrow{\bar{a}(n)} Q'}{P | Q \xrightarrow{\tau} (vn)(P'\{n/x\} | Q')} \\
\\
\text{(RES)} \qquad \text{(SUM)} \qquad \text{(REC)} \\
\frac{P \xrightarrow{\alpha} P'}{(vn)P \xrightarrow{\alpha} (vn)P'} \qquad \frac{P_i \xrightarrow{\alpha} P'_i \quad i \in I}{\sum_{i \in I} P_i \xrightarrow{\alpha} P'_i} \qquad \frac{P_A\{\bar{y}, \bar{q}/\bar{x}, \bar{p}\}\{\mathbf{w}, \mathbf{m}/\mathbf{z}, \mathbf{n}\} \xrightarrow{\alpha} P' \quad A(\bar{x}, \bar{p} | \mathbf{z}, \mathbf{n}) = P_A}{A(\bar{y}, \bar{q} | \mathbf{w}, \mathbf{m}) \xrightarrow{\alpha} P'}
\end{array}$$

Fig. 1. Labelled Transition System of the π -calculus

I is empty we write $\mathbf{0}$, or we simply omit it; we denote by $+$ the binary sum. A process $a(x).P$ can perform an input at a and the variable x is the placeholder for the name so received. The output case is symmetric: a process $\bar{a}(b).P$ can perform an output along the channel a . Notice that an output can send a name (either free or restricted) or a variable.

We assume that every constant A has a unique defining equation $A(\bar{x}, \bar{p} | \mathbf{z}, \mathbf{n}) = P_A$. The symbol \bar{p} , resp. \bar{x} , denotes a tuple of distinct names, resp. variables, that correspond to the free names, resp. variables, of P_A . \mathbf{n} , resp. \mathbf{z} , represents an infinite sequence of distinct names $\mathbb{N} \rightarrow \text{Names}$, resp. distinct variables $\mathbb{N} \rightarrow \text{Variables}$, that is intended to enumerate the (possibly infinite) bound names, resp. bound variables, of P_A . The parameters \mathbf{n} and \mathbf{z} do not usually appear in recursive definitions in the literature. The reason we add them is that we want to maintain the following Basic Assumption: *Every bound name/variable is different from any other name/variable, either bound or free.* In the π -calculus, this policy is usually implicit and maintained along the computation by dynamic α -conversion: every time the definition A is unfolded, a copy of the process P_A is created whose bound names and variables must be fresh. This dynamic choice is difficult to interpret in the event structures. Hence, in order to obtain a precise semantic correspondence, our recursive definitions prescribe all the names and variables that will be possibly used in the recursive process (see [5] for some examples).

The sets of free and bound names and free and bound variables of P , denoted by $\text{fn}(P)$, $\text{bn}(P)$, $\text{fv}(P)$, $\text{bv}(P)$, are defined as usual but for constant processes, whose definitions are as follows: $\text{fn}(A(\bar{x}, \bar{p} | \mathbf{z}, \mathbf{n})) = \{\bar{p}\}$, $\text{bn}(A(\bar{x}, \bar{p} | \mathbf{z}, \mathbf{n})) = \mathbf{n}(\mathbb{N})$, $\text{fv}(A(\bar{x}, \bar{p} | \mathbf{z}, \mathbf{n})) = \{\bar{x}\}$ and $\text{bv}(A(\bar{x}, \bar{p} | \mathbf{z}, \mathbf{n})) = \mathbf{z}(\mathbb{N})$. The operational semantics is given in Figure 1 in terms of an LTS (in late style¹) where we let α, β range over the set of labels

¹ We could as well choose the early style semantics. However, in that case the event structure corresponding to a simple input process would be the sum of all possible (infinite) variable instantiations. Instead, the use of late semantics allows a cleaner and more intuitive approach.

$\{\tau, a(x), \bar{a}(b), \bar{a}(n)\}$. The syntax of labels shows that the object of an input is always a variable, whereas the object of a free output is either a variable (e.g. $b(x)$ or $\bar{a}(x)$) or a name. On the other hand, the object of a bound output is always a name, since it must occur under a restriction. Moreover, thanks to the Basic Assumption, the side conditions in rules (PAR) and (RES) are not needed anymore.

3 Event Structures

This section reviews basic definitions of prime event structures [7, 14, 19].

Definition 1 (Labelled Event Structure). Let L be a set of labels. A labelled event structure is a tuple $\mathcal{E} = \langle E, \leq, \smile, \lambda \rangle$ s.t.

- E is a countable set of events;
- $\langle E, \leq \rangle$ is a partial order, called the causal order;
- for every $e \in E$, the set $[e] := \{e' \mid e' < e\}$, called the enabling set of e , is finite;
- \smile is an irreflexive and symmetric relation, called the conflict relation, satisfying the following: for every $e_1, e_2, e_3 \in E$ if $e_1 \leq e_2$ and $e_1 \smile e_3$ then $e_2 \smile e_3$.
- $\lambda : E \rightarrow L$ is a labelling function that associates a label to each event in E .

Intuitively, labels represent *actions*, and events should be thought of as *occurrences of actions*. Labels allow us to identify events which represent different occurrences of the same action. In addition, labels are essential when composing two event structures in a parallel composition, as they identify which events correctly synchronise.

We say that the conflict $e_2 \smile e_3$ is *inherited* from the conflict $e_1 \smile e_3$, when $e_1 < e_2$. If a conflict is not inherited we say that it is *immediate*. If two events are not causally related nor in conflict they are said to be *concurrent*.

The notion of computation is usually captured in event structures in terms of configurations. A *configuration* C of an event structure \mathcal{E} is a conflict free downward closed subset of E , i.e. a subset C of E satisfying: (1) if $e \in C$ then $[e] \subseteq C$ and (2) for every $e, e' \in C$, it is not the case that $e \smile e'$, that is e and e' are either causally dependent or concurrent. In other words, a configuration represents a run of an event structure, where events are partially ordered. The set of configurations of \mathcal{E} , partially ordered by inclusion, is denoted as $\mathcal{L}(\mathcal{E})$. An alternative notion of computation can be defined in terms of labelled transition systems of event structures. Such a definition allows to more directly state (and prove) that the computational steps of a π -calculus process are reflected into its event structure semantics.

Definition 2 (LTS of event structures). Let $\mathcal{E} = \langle E, \leq, \smile, \lambda \rangle$ be a labelled event structure and let e be one of its minimal events with $\lambda(e) = \beta$. Then we write $\mathcal{E} \xrightarrow{\beta} \mathcal{E}|e$, where $\mathcal{E}|e$ is the event structure $\langle E', \leq|_{E'}, \smile|_{E'}, \lambda_{E'} \rangle$ with $E' = \{e' \in E \mid e' \neq e \text{ and } e' \not\smile e\}$.

Roughly speaking, $\mathcal{E}|e$ is \mathcal{E} minus the event e , and minus all events that are in conflict with e . The reachable LTS with initial state \mathcal{E} corresponds to the computations over \mathcal{E} .

Event structures have been shown to be the class of objects of a category [20]. Moreover, it is easily shown that an isomorphism in this category is a label-preserving bijective function that preserves and reflects causality and conflict. We denote by $\mathcal{E}_1 \cong \mathcal{E}_2$ the fact that there is an isomorphism between \mathcal{E}_1 and \mathcal{E}_2 .

We review here an informal description of several operations on labelled event structures, that we are going to use in the next section. See [19] for more details.


- *Prefixing* $a.\mathcal{E}$. This operation adds to the event structure a new minimal element, labelled by a , below every other event in \mathcal{E} .
- *Prefixed sum* $\sum_{i \in I} a_i.\mathcal{E}_i$. This is obtained as the disjoint union of copies of the event structures $a_i.\mathcal{E}_i$. The conflict relation is extended by putting in conflict every pair of events belonging to two different copies of $a_i.\mathcal{E}_i$.
- *Restriction* (or *Hiding*) $\mathcal{E} \setminus X$ where $X \subseteq L$ is a set of labels. This is obtained by removing from \mathcal{E} all events with label in X and all events that are above (i.e., causally depend on) one of those.
- *Relabelling* $\mathcal{E}[f]$ where L and L' are two sets of labels and $f : L \rightarrow L'$. This operation just consists in composing the labelling function λ of \mathcal{E} with the function f . The new event structure is labelled over L' and its labelling function is $f \circ \lambda$.

The parallel composition of two event structures \mathcal{E}_1 and \mathcal{E}_2 gives a new event structure \mathcal{E}' whose events model the parallel occurrence of pairs of events $e_1 \in \mathcal{E}_1$ and $e_2 \in \mathcal{E}_2$. In particular, when the labels of e_1 and e_2 match according to an underlying synchronisation model, \mathcal{E}' records (with an event $e' \in \mathcal{E}'$) that a synchronisation between e_1 and e_2 is possible, and deals with the causal effects of such a synchronisation.

Technically, the parallel composition is defined as the categorical product followed by restriction and relabelling [20]. The categorical product represents all conceivable synchronisations, relabelling implements a synchronisation model by expressing which events are allowed to synchronise, and hiding removes synchronisations that are not permitted. The synchronisation model underlying the relabelling operation is formalised by the notion of *synchronisation algebra*, that is a partial binary operation \bullet_S defined on $L_* := L \uplus \{*\}$ where $*$ is a distinguished label. If α_i are the labels of events $e_i \in \mathcal{E}_i$, then the event $e' \in \mathcal{E}'$ representing the synchronisation of e_1 and e_2 is labelled by $\alpha_1 \bullet_S \alpha_2$. When $\alpha_1 \bullet_S \alpha_2$ is undefined, the synchronisation event e' is given a distinguished label *bad*, indicating that this event is not allowed and should be deleted.

Definition 3 (Parallel Composition of Event Structures). *Let $\mathcal{E}_1, \mathcal{E}_2$ two event structures labelled over L , let \bullet_S be a synchronisation algebra, and let $f_S : L_* \rightarrow L' = L_* \cup \{\text{bad}\}$ be a function defined as $f_S(\alpha_1, \alpha_2) = \alpha_1 \bullet_S \alpha_2$, if S is defined on (α_1, α_2) , and $f_S(\alpha_1, \alpha_2) = \text{bad}$ otherwise. The parallel composition $\mathcal{E}_1 \parallel_S \mathcal{E}_2$ is defined as the categorical product followed by relabelling and restriction: $\mathcal{E}_1 \parallel_S \mathcal{E}_2 = (\mathcal{E}_1 \times \mathcal{E}_2)[f_S] \setminus \{\text{bad}\}$. The subscripts S are omitted when the synchronisation algebra is clear from the context.*

Example 4. We show a simple example of parallel composition. Consider the set of labels $L = \{\alpha, \beta, \bar{\alpha}, \tau\}$ and the synchronisation algebra obtained as the symmetric closure of the following rules: $\alpha \bullet \bar{\alpha} = \tau$, $\alpha \bullet * = \alpha$, $\bar{\alpha} \bullet * = \bar{\alpha}$, $\beta \bullet * = \beta$ and undefined otherwise. Consider the two event structures $\mathcal{E}_1, \mathcal{E}_2$, where $E_1 = \{a, b\}, E_2 = \{a'\}$, with $a \leq_1 b$ and $\lambda_1(a) = \alpha, \lambda_1(b) = \beta, \lambda_2(a') = \bar{\alpha}$. The event structures are represented as follows:



where dotted lines represent immediate conflict, while the causal order proceeds upwards along the straight lines. Then $\mathcal{E}_3 := \mathcal{E}_1 \parallel \mathcal{E}_2$ is the event structure $\langle \mathcal{E}_3, \leq, \sim, \lambda \rangle$ where $\mathcal{E}_3 = \{e := (\emptyset, a, *), e' := (\emptyset, *, a'), e'' := (\emptyset, a, a'), d := (\{e\}, a', *), d'' := (\{e''\}, a', *)\}$, and the ordering, immediate conflict and the labelling are as in the picture above.

We say that an event structure \mathcal{E} is a *prefix* of an event structure \mathcal{E}' , denoted $\mathcal{E} \leq \mathcal{E}'$ if there exists $\mathcal{E}'' \cong \mathcal{E}'$ such that $E \subseteq E''$, no event in $E'' \setminus E$ is below any event of E , and conflict and order in \mathcal{E} are the restriction of those in \mathcal{E}'' . Winskel [18] has shown that the class of event structures with the prefix order is a large CPO, and thus the limits of countable increasing chains exist. Moreover all operators on event structures are continuous. We will use this fact to define the semantics of the recursive definitions.

4 Free Name Passing

We present the event structure semantics of the full π -calculus in two phases, dealing separately with the two main issues of the calculus. We start in this section discussing free name passing, and we postpone to the next section the treatment of scope extrusion.

The core of a compositional semantics of a process calculus is parallel composition. When a process P is put in parallel with another process, new synchronizations can be triggered. Hence the semantics of P must also account for “potential” synchronizations that might be activated by parallel compositions. In Winskel’s event structure semantics of CCS [18], the parallel composition is defined as a product in a suitable category followed by relabelling and hiding, as we have presented in Section 3. For the semantics of the π -calculus, when the parallel composition of two event structures is computed, synchronisation events that involve input and output on different channels cannot be hidden straight away. If at least one of the two channels is a variable, then it is possible that, after prefixing and parallel composition, the two channels will be made equal.

We then resort to a technique similar to the one used in [5]: we consider a generalised notion of relabelling that takes into account the history of a (synchronization) event. Such a relabelling is defined according to the following ideas:

- each pair $(a(x), \bar{a}(b))$ made of two equal names or two equal variables is relabelled $\tau_{x \rightarrow b}$, to indicate that it represents a legal synchronization where b is substituted for x . Moreover, such a substitution must be propagated in all the events that causally depend on this synchronization. However, after all substitutions have taken place, there is no need to remember the extra information carried by the τ action, than the subscripts of the τ events are erased.
- Synchronisations pairs, like $(a(x), \bar{b}(c))$, involving different channels (at least) one of which is a variable, are relabelled $(a(x), \bar{b}(c))_{x \rightarrow c}$, postponing the decision whether they represent a correct synchronization or not.
- Each pair $(n(x), \bar{m}(b))$ made of two different names is relabelled bad to denote a synchronization that is not allowed.

Definition 5 (Generalised Relabelling). *Let L and L' be two sets of labels, and let $\text{Pom}(L')$ be the set of pomsets (i.e., partially ordered multisets) over L' . Given an event structure $\mathcal{E} = \langle E, \leq, \smile, \lambda \rangle$ with labels in L , and a function $f : \text{Pom}(L') \times L \rightarrow L'$, we define the relabelling operation $\mathcal{E}[f]$ as the event structure $\mathcal{E}' = \langle E, \leq, \smile, \lambda' \rangle$ with labels in L' , where $\lambda' : E \rightarrow L'$ is defined by induction on the height of an element of E : if $h(e) = 0$ then $\lambda'(e) = f(\emptyset, \lambda(e))$, if $h(e) = n + 1$ then $\lambda'(e) = f(\lambda'([e]), \lambda(e))$.*

In words, an event e is relabelled with a label $\lambda'(e)$ that depends on the (pomset of) labels of the events belonging to its causal history $[e]$.

In the case of π -calculus with free names, let $L = \{a(x), \bar{a}\langle b \rangle, \tau \mid a, b \in \text{Names} \cup \text{Variables}, x \in \text{Variables}\}$ be the set of labels used in the LTS of π -calculus without restriction. We define the relabelling function needed by the parallel composition operation around the extended set of labels $L' = L \cup \{(\alpha, \beta)_{x \rightarrow b} \mid \alpha, \beta \in L\} \cup \{\tau_{x \rightarrow b}, \text{bad}\}$, where bad is a distinguished label. The relabelling function $f_\pi : \text{Pom}(L') \times (L' \uplus \{*\}) \times L' \uplus \{*\} \longrightarrow L'$ is defined as follows (we omit the symmetric clauses):

$$\begin{aligned}
f_\pi(X, \langle a(y), \bar{a}\langle b \rangle \rangle) &= \tau_{y \rightarrow b} & f_\pi(X, \langle a(x), \bar{y}\langle c \rangle \rangle) &= \begin{cases} \tau_{x \rightarrow c} & \text{if } \alpha_{y \rightarrow a} \in X \\ (a(x), \bar{y}\langle c \rangle)_{x \rightarrow c} & \text{otherwise} \end{cases} \\
f_\pi(X, \langle n(y), \bar{m}\langle b \rangle \rangle) &= \text{bad} & f_\pi(X, \langle y(x), \bar{a}\langle n \rangle \rangle) &= \begin{cases} \tau_{x \rightarrow n} & \text{if } \alpha_{y \rightarrow a} \in X \\ (y(x), \bar{a}\langle n \rangle)_{x \rightarrow n} & \text{otherwise} \end{cases} \\
f_\pi(X, \langle y(x), * \rangle) &= \begin{cases} a(x) & \text{if } \alpha_{y \rightarrow a} \in X \\ y(x) & \text{otherwise} \end{cases} & f_\pi(X, \langle \bar{y}\langle b \rangle, * \rangle) &= \begin{cases} \bar{a}\langle b \rangle & \text{if } \alpha_{y \rightarrow a} \in X \\ \bar{y}\langle b \rangle & \text{otherwise} \end{cases} \\
f_\pi(X, \langle \alpha, * \rangle) &= \alpha & f_\pi(X, \langle \alpha, \beta \rangle) &= \text{bad} \quad \text{otherwise}
\end{aligned}$$

The extra information carried by the τ -actions, differently from that of “incomplete synchronization” events, is only necessary in order to *define* the relabelling, but there is no need to keep it after the synchronization has been completed. Hence we apply a second relabelling er that simply erases the subscript of τ actions.

The semantics of the π -calculus is then defined as follows by induction on processes, where the parallel composition of event structure is defined by

$$\mathcal{E}_1 \parallel_\pi \mathcal{E}_2 = ((\mathcal{E}_1 \times \mathcal{E}_2) [f_\pi][er]) \setminus \{\text{bad}\}$$


To deal with recursive definitions, we use an index k to denote the level of unfolding.

$$\begin{aligned}
\{\mathbf{0}\}_k &= \mathbf{0} & \{\sum_{i \in I} \pi_i.P_i\}_k &= \sum_{i \in I} \pi_i.\{P_i\}_k & \{P \mid Q\}_k &= \{P\}_k \parallel_\pi \{Q\}_k \\
\{A\langle \bar{y}, \bar{q} \mid \mathbf{w}, \mathbf{m} \rangle\}_0 &= \mathbf{0} & \{A\langle \bar{y}, \bar{q} \mid \mathbf{w}, \mathbf{m} \rangle\}_{k+1} &= \{P_A\{\bar{y}, \bar{q} / \bar{x}, \bar{p}\}\{\mathbf{w}, \mathbf{m} / \mathbf{z}, \mathbf{n}\}\}_k
\end{aligned}$$

Recall that all operators on event structures are continuous with respect to the prefix order. It is thus easy to show that, for any k , $\{P\}_k \leq \{P\}_{k+1}$. We define $\{P\}$ to be the limit of the increasing chain $\dots \{P\}_k \leq \{P\}_{k+1} \leq \{P\}_{k+2} \dots$, that is $\{P\} = \sup_{k \in \mathcal{N}} \{P\}_k$. Since all operators are continuous w.r.t. the prefix order we have the following result:

Theorem 6 (Compositionality). *The semantics $\{P\}$ is compositional, i.e. $\{P \mid Q\} = \{P\} \parallel_\pi \{Q\}$, $\{\sum_{i \in I} \pi_i.P_i\} = \sum_{i \in I} \pi_i.\{P_i\}$,*

Example 7. As an example, consider the process $P = n(z).(\bar{z}\langle a \rangle \mid m(x))$. The synchronization along the channel m can be only performed if the previous synchronization along n substitutes the variable z with the name m . Accordingly, the semantics of the process P is the leftmost event structure depicted below, denoted by \mathcal{E}_P . Moreover, the rightmost structure corresponds to the semantics of the process $P \mid \bar{n}\langle m \rangle \mid \bar{n}\langle p \rangle$.



The following theorem shows that the event structure semantics is operationally correct. Indeed, given a process P , the computational steps of P in the LTS of Section 2 are reflected by the semantics $\llbracket P \rrbracket$.

Theorem 8 (Operational Adequacy). *Let $\beta \in \{a(x), \bar{a}(b), \tau\}$. Suppose $P \xrightarrow{\beta} P'$ in the π -calculus. Then there exists \mathcal{E} such that $\llbracket P \rrbracket \xrightarrow{\beta} \mathcal{E}$ and $\mathcal{E} \cong \llbracket P' \rrbracket$. Conversely, suppose $\llbracket P \rrbracket \xrightarrow{\beta} \mathcal{E}'$. Then there exists P' such that $P \xrightarrow{\beta} P'$ and $\llbracket P' \rrbracket \cong \mathcal{E}'$.*

Note that the correspondence holds for the labels in the LTS of the calculus. Labels that identify “incomplete synchronizations” have been introduced in the event structure semantics for the sake of compositionality, but they are not considered in the theorem above since they do not correspond to any operational step. Moreover, the semantics is clearly not fully abstract in any reasonable sense, since interleaving equivalences are less discriminating than the corresponding causal equivalences on event structures.

5 Scope Extrusion

In this section we show how the causal dependencies introduced by scope extrusion can be captured by event structure-based models. As we discussed in Section 1, the communication of bound names implies that any action with a bound subject causally depends on a dynamic set of possible extruders of that bound subject. Hence dealing with scope extrusion requires modelling some form of disjunctive causality. Prime event structures are stable models that represent an action α that can be caused either by the action β_1 or the action β_2 as two different events e, e' that are both labelled α but e causally depends on the event labeled β_1 while e' is caused by the event labeled β_2 . In order to avoid the proliferation of events representing the same action with different extruders, we follow here a different approach, postponing to the next section a more detailed discussion on the use of prime event structures.

5.1 Event Structure with Bound Names

We define the semantics of the full π -calculus in terms of pairs (\mathcal{E}, X) , where \mathcal{E} is a prime event structure, and X is a set of names. We call such a pair an *event structure with bound names*. Intuitively, the causal relation of \mathcal{E} encodes the structural causality of a process, while the set X records bound names. Given a pair (\mathcal{E}, X) we define a notion of *permitted configurations*: a configuration that contains an action whose subject is a bound name, is permitted if it also contains a previous extrusion of that name. Objective causality is then implicitly captured by permitted configurations.

Definition 9 (Semantics). *The semantics of the full π -calculus is inductively defined as follows, where k denote the level of unfolding of recursive definitions, and we write \mathcal{E}_P^k , resp. X_P^k , for the first, resp. the second, projection of the pair $\llbracket P \rrbracket_k$:*


$$\begin{aligned} \llbracket \sum_{i \in I} \pi_i.P_i \rrbracket_k &= (\sum_{i \in I} \pi_i.\mathcal{E}_{P_i}^k, \bigcup_{i \in I} X_{P_i}^k) & \llbracket P \mid Q \rrbracket_k &= (\mathcal{E}_P^k \parallel_{\pi} \mathcal{E}_Q^k, X_P^k \cup X_Q^k) \\ \llbracket A(\bar{y}, \bar{q} \mid \mathbf{w}, \mathbf{m}) \rrbracket_0 &= (\emptyset, \{\mathbf{w}(\mathbb{N})\}) & \llbracket (vn)P \rrbracket_k &= (\mathcal{E}_P^k, X_P^k \cup \{n\}) \\ \llbracket A(\bar{y}, \bar{q} \mid \mathbf{w}, \mathbf{m}) \rrbracket_{k+1} &= (\mathcal{E}_{PA\{\bar{y}, \bar{q}/\bar{x}, \bar{p}\}\{\mathbf{w}, \mathbf{m}/z, \mathbf{n}\}}^k, \{\mathbf{w}(\mathbb{N})\}) & \llbracket \mathbf{0} \rrbracket_k &= (\emptyset, \emptyset) \end{aligned}$$

It is easy to show that, for any k , $\mathcal{E}_P^k \leq \mathcal{E}_P^{k+1}$ and $X_P^k = X_P^{k+1} = X_P$. Then the semantics of a process P is defined as the following limit: $\llbracket P \rrbracket = (\sup_{k \in \mathcal{N}} \mathcal{E}_P^k, X_P)$.

The semantics is surprisingly simple: a restricted process $(\nu n)P$ is represented by a prime event structure that encodes the process P where the scope of n has been opened, and we collect the name n in the set of bound names. As for parallel composition, the semantics $\llbracket P \mid Q \rrbracket$ is a pair (\mathcal{E}, X) where X collects the bound names of both $\llbracket P \rrbracket$ and $\llbracket Q \rrbracket$ (recall that we assumed that bound names are pairwise different), while the event structure \mathcal{E} is obtained exactly as in the previous sections. This is since the event structures that get composed correspond to the processes P and Q where the scope of any bound name has been opened. The following property can be immediately proved.

Proposition 10. *Let P, Q be two processes, then $\llbracket (\nu n)P \mid Q \rrbracket = \llbracket (\nu n)(P \mid Q) \rrbracket$.*

Example 11. Consider the process $P = (\nu n)(\bar{a}\langle n \rangle \mid n(z)) \mid a(x).(\bar{x}\langle b \rangle \mid \bar{c}\langle x \rangle)$, whose first synchronization produces a new extruder $\bar{c}\langle n \rangle$ for the bound name n . The semantics of P is the pair $(\mathcal{E}_P, \{n\})$, where \mathcal{E}_P is the following event structure:



In order to study the operational correspondence between the LTS semantics of the π -calculus and the event structure semantics above, we first need to adapt the notion of computational steps of the pairs (\mathcal{E}, X) . The definition of labelled transitions between prime event structures, i.e., Definition 2, is generalized as follows.

Definition 12 (Permitted Transitions). *Let (\mathcal{E}, X) be a labelled event structure with bound names. Let e be a minimal event of \mathcal{E} with $\lambda(e) = \beta$. We define the following permitted labelled transitions:*

- $(\mathcal{E}, X) \xrightarrow{\beta} (\mathcal{E} \downarrow e, X)$, if $\beta \in \{\tau, a(x), \bar{a}\langle b \rangle\}$ with $a, b \notin X$.
- $(\mathcal{E}, X) \xrightarrow{\bar{a}\langle n \rangle} (\mathcal{E} \downarrow e, X \setminus \{n\})$, if $\beta = \bar{a}\langle n \rangle$ with $a \notin X$ and $n \in X$.

According to this definition, the set of bound names constrains the set of transitions that can be performed. In particular, no transition whose label has a bound subject is allowed. On the other hand, when a minimal event labeled $\bar{a}\langle n \rangle$ is consumed, if the name n is bound, the transition's labels records that this event is indeed a bound output. Moreover, in this case we record that the scope of n is opened by removing n from the set of bound names of the target pair. Finally, observe that the previous definition only allows transitions whose labels are in the set $L = \{\tau, a(x), \bar{a}\langle b \rangle, \bar{a}\langle n \rangle\}$, which is exactly the sets of labels in the LTS of Section 2.

Theorem 13 (Operational Adequacy). *Let be $\beta \in \{a(x), \bar{a}\langle b \rangle, \bar{a}\langle n \rangle, \tau\}$. Suppose $P \xrightarrow{\beta} P'$ in the π -calculus. Then there exists \mathcal{E} s.t. $\llbracket P \rrbracket \xrightarrow{\beta} \mathcal{E}$ and $\mathcal{E} \cong \llbracket P' \rrbracket$. Conversely, suppose $\llbracket P \rrbracket \xrightarrow{\beta} \mathcal{E}'$. Then there exists P' s.t. $P \xrightarrow{\beta} P'$ and $\llbracket P' \rrbracket \cong \mathcal{E}'$.*

5.2 Subjective and Objective Causality

Given an event structure with bound names (\mathcal{E}, X) , Definition 12 shows that some configurations of \mathcal{E} are no longer allowed. For instance, if e is minimal but its label has a subject that is a name in X , e.g. $\lambda(e) = n(x)$ with $n \in X$, then the configuration $\{e\}$ is no longer allowed since the event e requires a previous extrusion of the name n .

Definition 14 (Permitted Configuration). *Let (\mathcal{E}, X) be an event structure with bound names. For an event $e \in E$ define $e \uparrow = \{e' \mid e \leq e'\}$. Given a configuration C of \mathcal{E} , we say that C is permitted in (\mathcal{E}, X) whenever, for any $e \in C$ whose label has subject n with $n \in X$,*

- $C \setminus e \uparrow$ is permitted, and
- $C \setminus e \uparrow$ contains an event whose label is an output action with object n .

The first item of the definition above is used to avoid circular definitions that would allow wrong configurations like $\{\bar{n}\langle m \rangle, \bar{m}\langle n \rangle\}$ with $X = \{n, m\}$. Now, the two forms of causality of the π -calculus can be defined using event structures with bound names and permitted configurations.

Definition 15 (Subjective and Objective Causality). *Let P be a process of the π -calculus, and $\llbracket P \rrbracket = (\mathcal{E}_P, X_P)$ be its semantics. Let be $e_1, e_2 \in E_P$, then*

- e_2 has a subjective dependence on e_1 if $e_1 \leq_{\mathcal{E}_P} e_2$;
- e_2 has a objective dependence on e_1 if (i) the label of e_1 is the output of a name in X which is also the subject of the label of e_2 , and if (ii) there exists a configuration C that is permitted in (\mathcal{E}, X) and that contains both e_1 and e_2 .

Example 16. Let consider again the process P in Example 11. The configurations $C_1 = \{\bar{a}\langle n \rangle, n\langle z \rangle\}$ and $C_2 = \{\tau, \bar{c}\langle n \rangle, n\langle z \rangle\}$ are both permitted by $\llbracket P \rrbracket$, and they witness the fact that the action $n\langle z \rangle$ has an objective dependence on $\bar{a}\langle n \rangle$ and on $\bar{c}\langle n \rangle$. We could also say that $n\langle z \rangle$ objectively depends *either* on $\bar{a}\langle n \rangle$ or on $\bar{c}\langle n \rangle$.

Example 17. Let be $P = (vn)(\bar{a}\langle n \rangle \mid \bar{b}\langle n \rangle \mid n(x))$, then $\llbracket P \rrbracket = (\mathcal{E}_P, \{n\})$ where \mathcal{E}_P has three concurrent events. In this process there is no subjective causality, however the action $n(x)$ has an objective dependence on $\bar{a}\langle n \rangle$ and on $\bar{b}\langle n \rangle$ since both $C_1 = \{\bar{a}\langle n \rangle, n(x)\}$ and $C_2 = \{\bar{b}\langle n \rangle, n(x)\}$ are permitted configurations.

Example 18. Let be $P = (vn)(\bar{a}\langle n \rangle.\bar{b}\langle n \rangle.n(x))$, then $\llbracket P \rrbracket = (\mathcal{E}_P, \{n\})$ where \mathcal{E}_P is a chain of three events. According to the causal relation of \mathcal{E}_P , the action $n(x)$ has a structural dependence on both the outputs. Moreover, the permitted configuration $C = \{\bar{a}\langle n \rangle, \bar{b}\langle n \rangle, n(x)\}$ shows that $n(x)$ has an objective dependence on $\bar{a}\langle n \rangle$ and on $\bar{b}\langle n \rangle$. In this case we do not know which of the two outputs really extruded the bound name, accordingly to the inclusive disjunctive causality approach we are following.

5.3 The Meaning of Labelled Causality

In this paper we focus on compositional semantics, studying a true concurrent semantics that operationally matches the LTS semantics of the π -calculus. Alternatively, one could take as primitive the reduction semantics of the π -calculus, taking the perspective that only τ -events are “real” computational steps of a concurrent system. Therefore one could argue that the concept of causal dependency makes only sense between τ events. In this perspective, we propose to interpret the causal relation between non- τ events as an anticipation of the causal relations involving the synchronizations they will take part in. In other terms, non- τ events (from now on simply called labelled events) represent “incomplete” events, that are waiting for a synchronization or a substitution to be completed. Hence we can prove that in our semantics two labelled events e_1 and e_2 are causally dependent *if and only if* the τ -events they can take part in are causally dependent. This property is expressed by the following theorem in terms of permitted configurations. Recall that the parallel composition of two event structures is obtained by first constructing the cartesian product. Therefore there are projection morphisms π_1, π_2 on the two composing structures. Let call τ -configuration a configuration whose events are all τ -events. Note that every τ -configuration is permitted.

Theorem 19. *Let P be a process. A configuration C is permitted in $\{P\}$ if and only if there exists a process Q and a τ -configuration C' in $\{P \mid Q\}$ such that $\pi_1(C') = C$.*

Let e_1, e_2 be two labelled events of $\{P\} = (\mathcal{E}_P, X_P)$. If e_1, e_2 are structurally dependent, i.e., $e_1 \leq_{\mathcal{E}_P} e_2$, then such a structural dependence is preserved and reflected in the τ -actions they are involved in because of the way the parallel composition of event structures is defined. On the other hand, let be e_1, e_2 objectively dependent. Consider the parallel composition $(\mathcal{E}_P \parallel_{\pi} \mathcal{E}_Q, X_P \cup X_Q)$ for some Q such that there is a τ event e'_2 in $\mathcal{E}_P \parallel_{\pi} \mathcal{E}_Q$ with $\pi_1(e'_2) = e_2$ and $[e'_2]$ is a τ -configuration. Then there must be an event $e'_1 \in [e'_2]$ such that $\pi_1(e'_1) = e_1$.

6 Disjunctive Causality

As we discussed in Section 1, objective causality introduced by scope extrusion requires for the π -calculus a semantic model that is able to express some form of disjunctive causality. In the previous section we followed an approach that just ensures that some extruder (causally) precedes any action with a bound subject. However, we could alternatively take the effort of tracing the identity of the actual extruders. We could do it by duplicating the events corresponding to actions with bound subject and letting different copies depend on different, alternative, extruders. Such a duplication allows us to use prime event structures as semantics models. In this section we discuss this alternative approach showing to what extent it can be pursued.

As a first example, the semantics of the process $P = (\nu n)(\bar{a}\langle n \rangle \mid \bar{b}\langle n \rangle \mid n(x))$, containing two possible extruders for the action $n(x)$, can be represented by left-most prime event structure in Figure 2. When more than a single action use as subject the same bound name, each one of these actions must depend on one of the possible extruders. Then the causality of the process $(\nu n)(\bar{a}\langle n \rangle \mid \bar{b}\langle n \rangle \mid n(x).n(y))$, represented by the right-most event structure in Figure 2, shows that the two read actions might depend either on the same extruder or on two different extrusions.





Fig. 2.

Things get more complicated when dealing with the dynamic addition of new extruders by means of communications. In order to guarantee that there are distinct copies of any event with a bound subject that causally depend on different extruders, we have to consider the objective causalities generated by a communication. More precisely, when the variable x is substituted with a bound name n by effect of a synchronization, (i) any action with subject x appearing in the reading thread becomes an action that requires a previous scope extrusion, and (ii) the outputs with object x become new extruders for any action with subject n or x . To exemplify, consider the process $P' = (\nu n)(\bar{a}\langle n \rangle \mid \bar{b}\langle n \rangle \mid n\langle z \rangle) \mid a(x).(\bar{x}\langle a \rangle \mid \bar{c}\langle x \rangle)$ that initially contains two extruders for n , and with the synchronization along the channel a evolves to $(\nu n)(\bar{b}\langle n \rangle \mid n\langle z \rangle \mid \bar{n}\langle a \rangle \mid \bar{c}\langle n \rangle)$. Its causal semantics can be represented with the following prime event structure:



The read action $n\langle z \rangle$ may depend on one of the two initial extruders $\bar{a}\langle n \rangle$ and $\bar{b}\langle n \rangle$, or on the new extruder $\bar{c}\langle n \rangle$ that is generated by the first communication. Accordingly, three different copies of the event $n\langle z \rangle$ appear over each of the three extruders. On the other hand, the output action on the bound name n is generated by the substitution entailed by the communication along the channel a , hence any copy of that action keeps a (structural) dependence on the corresponding τ event. Moreover, since it is an action with bound subject, there must be a copy of it for each of the remaining extruders of n , that is $\bar{b}\langle n \rangle$ and $\bar{c}\langle n \rangle$. To enhance readability, the event structure resulting from the execution of the communication along the channel a is the leftmost e.s. in Figure 3.




Fig. 3.

So far so good, in particular it seems possible to let the causal relation of prime event structures encode both structural and objective causality of π -processes. However, this is not the case. Consider the process $P = (\nu n)(\bar{a}\langle n \rangle.\bar{b}\langle n \rangle.n\langle z \rangle)$ of Example 18.

If we just duplicate the event $n(z)$ to distinguish the fact that it might depend on an extrusion along a or along b , we obtain the rightmost structure in Figure 3, that we denote \mathcal{E}_p . In particular, even if the two copies intends to represent two different objective causalities, nothing distinguishes them since they both structurally depend on both outputs. This is a problem when we compose the process P in parallel with, e.g., $Q = a(x).\bar{c}\langle x \rangle \mid b(y).\bar{d}\langle y \rangle$. After two synchronizations we would like to obtain two copies of the read actions on n that depend on the two different remaining extruders $\bar{c}\langle n \rangle$ and $\bar{d}\langle n \rangle$. However, in order to obtain such an event structure as the parallel composition of the semantics of P and Q we must be able to record somehow the different objective causality that distinguishes the two copies of $n(z)$ in the semantics of P .

The technical solution would be to enrich the event labels so that the label of an event e also records the identity of the extruder events that e (objectively) causally depends on. A precise account of this approach is technically involved and intractable, so that the intuition on the semantics gets lost. Moreover, we think that this final example sheds light on the fact that structural and objective causality of π -processes cannot be expressed by the sole causal relation of event structures. To conclude, at the price of losing the information about which extruder an event depends on, the approach we developed in the previous section brings a number of benefits: it is technically much simpler, it is operationally adequate, and it gives a clearer account of the two forms of causality distinctive of π -processes.

7 Related Work

There are several causal models for the π -calculus, that use different techniques. There exist noninterleaving semantics in terms of labelled transition systems, where the causal relations between transitions are represented by “proofs” which allow to distinguish different occurrences of the same transition [16, 1, 8]. In [4], a more abstract approach is followed, which involves indexed transition systems. In [11], a semantics of the π -calculus in terms of pomsets is given, following ideas from dataflow theory. The two papers [3, 9] present Petri nets semantics of the π -calculus. However, none of these approaches accounts for parallel extrusion. We finally recall [13] that introduces a graph rewriting-based semantics of the π -calculus that allows parallel extrusions.

Previous work on an event structure semantics of the π -calculus are [17, 5] that study fragments of the calculus, while [2] gives an unlabelled event structure semantics of the full calculus which only corresponds to the *reduction* semantics, hence which is not compositional.

We plan for future work the application of the present semantics to the study of a labelled reversible semantics of the π -calculus that would extend the work of Danos and Krivine [6]. Phillips and Ulidowski [15] noted the strict correspondence between reversible transition systems and event structures. A first step in this direction is [12], which proposes a reversible semantics of the π -calculus that only considers reductions. It would also be interesting to study which kind of configuration structures [10] can naturally include our definition of permitted configuration.

Acknowledgment

We thank Ilaria Castellani, whose curiosity and skepticism pushed us to dig deeper. We also thank Giotto, for his wonderful frescos, and seaside greek iced coffes, for their mind-refreshing power. The third author is partially supported by EPSRC EP/F003757/01 and G015635/01.

References

1. Boreale, M., Sangiorgi, D.: A fully abstract semantics for causality in the π -calculus. *Acta Inf.* 35(5), 353–400 (1998)
2. Bruni, R., Melgratti, H., Montanari, U.: Event structure semantics for nominal calculi. In: *Proc. of CONCUR*. LNCS, vol. 4137, pp. 295–309. Springer (2006)
3. Busi, N., Gorrieri, R.: A petri net semantics for pi-calculus. In: *Proceedings of 6th CONCUR*. LNCS, vol. 962, pp. 145–159. Springer (1995)
4. Cattani, G.L., Sewell, P.: Models for name-passing processes: Interleaving and causal. In: *Proceedings of 15th LICS*. pp. 322–332. IEEE (2000)
5. Crafa, S., Varacca, D., Yoshida, N.: Compositional event structure semantics for the internal pi-calculus. In: *Proc. of CONCUR*. LNCS, vol. 4703, pp. 317–332. Springer (2007)
6. Danos, V., Krivine, J.: Reversible communicating systems. In: *Proceedings of CONCUR 2004*. LNCS, vol. 3170, pp. 292–307. Springer (2004)
7. Degano, P., De Nicola, R., Montanari, U.: On the consistency of truly concurrent operational and denotational semantics. In: *LICS*. pp. 133–141. IEEE (1988)
8. Degano, P., Priami, C.: Non-interleaving semantics for mobile processes. *Theor. Comp. Sci.* 216(1-2), 237–270 (1999)
9. Engelfriet, J.: A multiset semantics for the pi-calculus with replication. *Theor. Comp. Sci.* 153(1&2), 65–94 (1996)
10. van Glabbeek, R.J., Plotkin, G.D.: Configuration structures, event structures and petri nets. *Theor. Comput. Sci.* 410(41), 4111–4159 (2009)
11. Jagadeesan, L.J., Jagadeesan, R.: Causality and true concurrency: A data-flow analysis of the pi-calculus. In: *AMAST*. LNCS, vol. 936, pp. 277–291 (1995)
12. Lanese, I., Mezzina, C.A., Stefani, J.: Reversing higher-order pi. In: *Proceedings of CONCUR 2010*. LNCS, vol. 6269, pp. 478–493. Springer (2010)
13. Montanari, U., Pistore, M.: Concurrent semantics for the pi-calculus. *Electr. Notes Theor. Comput. Sci.* 1, 411–429 (1995), proceedings of MFPS '95
14. Nielsen, M., Plotkin, G.D., Winskel, G.: Petri nets, event structures and domains, part I. *Theor. Comp. Sci.* 13(1), 85–108 (1981)
15. Phillips, I., Ulidowski, I.: Reversibility and models for concurrency. *Electr. Notes Theor. Comput. Sci.* 192(1), 93–108 (2007), proceedings of SOS 2007
16. Sangiorgi, D.: Locality and true-concurrency in calculi for mobile processes. In: *Proceedings of TACS*. LNCS, vol. 789, pp. 405–424. Springer (1994)
17. Varacca, D., Yoshida, N.: Typed event structures and the linear pi-calculus. *Theor. Comput. Sci.* 411(19), 1949–1973 (2010)
18. Winskel, G.: Event structure semantics for CCS and related languages. In: *Proceedings of 9th ICALP*. LNCS, vol. 140, pp. 561–576. Springer (1982)
19. Winskel, G.: Event structures. In: *Advances in Petri Nets 1986, Part II; Proceedings of an Advanced Course*. LNCS, vol. 255, pp. 325–392. Springer (1987)
20. Winskel, G., Nielsen, M.: Models for concurrency. In: *Handbook of logic in Computer Science*, vol. 4. Clarendon Press (1995)