

High-Level Petri Nets as Type Theories in the Join Calculus

Marzia Buscemi and Vladimiro Sassone

DMI, Università di Catania, Italy
buscemi@dmi.unict.it, vs@dmi.unict.it

Abstract. We study the expressiveness of the join calculus by comparison with (generalised, coloured) Petri nets and using tools from type theory. More precisely, we consider four classes of nets of increasing expressiveness, Π_i , introduce a hierarchy of type systems of decreasing strictness, Δ_i , $i = 0, \dots, 3$, and we prove that a join process is typeable according to Δ_i if and only if it is (strictly equivalent to) a net of class Π_i . In the details, Π_0 and Π_1 contain, resp., usual place/transition and coloured Petri nets, while Π_2 and Π_3 propose two natural notions of high-level net accounting for dynamic reconfiguration and process creation and called reconfigurable and dynamic Petri nets, respectively.

Introduction

The join calculus [5, 7] is an algebra of mobile processes with asynchronous name-passing communication that simplifies the π -calculus by enforcing the hypothesis of *unique receptors*. This means that there is at most one process receiving messages on a name, and is a distinctive feature of the join calculus that makes it suitable for distributed implementations, as channels may be allocated at their receptor process. The present work focuses on a version of the join calculus, studying its expressiveness by establishing a tight link to Petri nets.

Petri nets [16, 17] are a fundamental model in concurrency, representing basic distributed machines that, although rudimentary, may exhibit complex interaction behaviours when processes (transitions) compete for shared resources (tokens). Operational in nature, Petri nets have been studied extensively from the semantic viewpoint (see [18] for some references).

The analogy between join terms and nets is relatively simple, and was first noticed in [1] and, recently, in [14]. Names, messages, and elementary definitions of the join calculus (cf. §1) correspond respectively to places, tokens, and transitions of Petri nets (cf. §2). The correspondence, however, runs short soon because nets are *not* a value-passing formalism and can express *no* mobility. Regrettably, they have a static, immutable network topology. This suggests to look for suitable extensions of Petri nets – in particular nets with mobility – that might be profitably applied to the study of mobile networks.

Supported by MURST project *TOSCA*. The authors wish to thank **BRICS**, Basic Research in Computer Science and project MIMOSA, INRIA Sophia Antipolis.

In the present paper we consider three extensions of place/transition Petri nets (Π_0) obtained by adding one by one, in a hierarchical fashion, the features needed to achieve the full expressiveness of join calculus. Namely, *value-passing*, accounted for by (a version of) *coloured nets* (Π_1), mobility as *network reconfigurability*, achieved by introducing *reconfigurable nets* (Π_2), and dynamically growing, *open networks*, modelled by the notion of *dynamic nets* (Π_3).

High level nets exist in several shapes and versions. The distinctive features of the mobile nets presented here is that input places of transitions are private and presets are immutable. A reconfigurable net can deliver tokens in different places at each firing, according to the input it receives, while a dynamic net can spawn a new net to life. But they *cannot* change the input arcs of any transitions. As in [13], where an algebra with such properties has been considered, we believe that this is the key to a tractable compositional semantic framework, the key to control generality. Most noticeably, it corresponds to unique receptiveness in the join calculus. A further naturality criteria in our view is that, as it turns out, dynamic nets can be proved *equivalent* to the join calculus.

The technical bulk of this work is characterising classes of join terms corresponding to the classes of the high-level nets discussed above. More precisely, our approach consists of designing four type systems Δ_i , $i = 0, \dots, 3$, that single out the terms corresponding to the nets in Π_i , respectively for $i = 0, \dots, 3$. In particular, for a fixed, semantic-preserving, well-behaved translation $\llbracket _ \rrbracket$, a join term P is typeable in Δ_i if and only if $\llbracket P \rrbracket$ belongs to Π_i . System Δ_0 is aimed at constraining terms to place/transition nets and, therefore, forbids dynamic process creation and nontrivial messages, Δ_1 relaxes the latter limitation, but enforces a strict distinction between channels and values, supporting value-passing but forbidding any mobility. Such distinction is then relaxed in Δ_2 . System Δ_3 relaxes the dynamicity constraint and, therefore, turns out to be trivial, i.e., each term is typeable in Δ_3 . That is, dynamic nets *coincide* with join terms. Finally, we provide a system Δ_4 that sums the features of all the others.

It is worth remarking that these systems are very rudimentary from the type theoretic viewpoint: there are no complex types or rules, nor sophisticated issues such as polymorphism [8, 15] or similar. However, we believe that our formalisation is quite interesting, suggestive, and worth pursuing, because the nature of join-terms makes it natural to express our conditions in systems of rules. Nevertheless, this paper remains a paper about comparing models. As a matter of fact, relating join calculus and Petri nets may be beneficial for both. On one of the edges of the connection, in fact, the join calculus may provide Petri nets with a compositional framework, together with behavioural semantics such as testing [10] and bisimulation [3]. Also, it may suggest interesting, semantically well-founded extensions of Petri nets, such as reconfigurable and dynamic nets here, mobile nets in [1], and functional nets in [14]. On the other edge, it opens the join calculus to an entire body of results on the semantics of noninterleaving concurrency, such as those supported by monoidal categories (see, e.g., [4]).

Here we actually consider a generalisation of the join calculus in that join patterns may not be linear, a choice that we discuss at length in §1.

Related Work. Although focus and approach here differ radically from [1], our models are clearly related to those of Asperti and Busi. While we aim at representing the join calculus precisely, the *dynamic nets* of [1] are more loosely inspired by it and are more general than ours. In particular, they do not enforce privateness of input places – that is unique receptors – as our nets do. The analogy between Petri nets and join calculus has also provided the inspiration for Odersky’s *functional nets* [14]. Although *loc. cit.* is quite different in spirit from the present paper, the relationships between Odersky’s framework and our ideas here are worth of further investigation. This, together with the fine comparison with [1], we leave to a later paper.

Structure of the Paper. The paper is organised as follows. In §1 we recall the basic definitions of the join calculus, while §2 introduces Petri nets, Δ_0 and relates typeability to PT nets. The following sections repeat the same pattern for coloured, reconfigurable, dynamic nets and the respective type systems. Finally, §6 introduces Δ_4 . Due to space limitations all proofs are omitted.

1 The join calculus

We shall focus on a monadic version of the join calculus [5, 7], writing its operational semantics in terms of a reduction system as in [15]. For a thorough introduction the reader is referred to the literature. Let \mathcal{Nm} be an infinite set of names and let $x, y \dots, u, v, \dots$ range over \mathcal{Nm} . Join terms, definitions, and patterns are given by the following grammar.

TERMS	$P, Q ::= 0 \mid x\langle y \rangle \mid P \mid Q \mid \text{def } D \text{ in } P$
DEFINITIONS	$D, E ::= J \triangleright P \mid D \wedge E$
JOIN PATTERNS	$J, K ::= x\langle y \rangle \mid J \mid K$

Thus, a join term P is either the ‘null’ term, an emission $x\langle y \rangle$ of message y on channel x , a parallel composition of terms, or a local definition. A definition D is a set of elementary definitions $J \triangleright P$ matching join patterns J to terms P .

The sets of defined names dn , received names rn , and free names fn are defined below, and then extended to sets of terms in the obvious way.

$rn(x\langle y \rangle) = \{y\}$	$rn(J \mid K) = rn(J) \cup rn(K)$
$dn(x\langle y \rangle) = \{x\}$	$dn(J \mid K) = dn(J) \cup dn(K)$
$dn(D \wedge E) = dn(D) \cup dn(E)$	$dn(J \triangleright P) = dn(J)$
$fn(x\langle y \rangle) = \{x\} \cup \{y\}$	$fn(P \mid Q) = fn(P) \cup fn(Q)$
$fn(D \wedge E) = fn(D) \cup fn(E)$	$fn(J \triangleright P) = dn(J) \cup (fn(P) \setminus rn(J))$
$fn(\text{def } D \text{ in } P) = (fn(P) \cup fn(D)) \setminus dn(D)$	

A *renaming* σ is a map from names to names that is the identity except on a finite number of names. We indicate by $dom(\sigma)$ the set of names on which σ

is not the identity, and by $cod(\sigma)$ its image, using function application for the renaming of free names in terms.

Definition 1. The *structural congruence* \equiv is the smallest substitutive (i.e., closed for contexts) equivalence relation on terms that satisfies the following rules, where θ and σ are *one-to-one* renamings such that $dom(\theta) \subseteq rn(D)$, $cod(\theta) \cap fn(D) = \emptyset$ and $dom(\sigma) \subseteq dn(D)$, $cod(\sigma) \cap fn(\theta D, P) = \emptyset$.

$$\text{def } D \text{ in } P \equiv \text{def } \sigma\theta D \text{ in } \sigma P \quad (1)$$

$$P \mid \text{def } D \text{ in } Q \equiv \text{def } D \text{ in } P \mid Q \quad fn(P) \cap dn(D) = \emptyset \quad (2)$$

$$\text{def } D \text{ in def } E \text{ in } P \equiv \text{def } D \wedge E \text{ in } P \quad dn(E) \cap (dn(D) \cup fn(D)) = \emptyset \quad (3)$$

plus equations stating that 0 is the unit for \mid and that the operators \mid and \wedge are associative and commutative.

Taking $dom(\sigma) = \emptyset$, rule (1) expresses the α -equivalence of definitions up to a renaming of their received names, while considering $dom(\theta) = \emptyset$ we obtain α -equivalence of terms up to renaming of defined names. Rule (2) formalises the scope extrusion of names, and rule (3) states that, under conditions that avoid name clashes, definitions can equivalently be gathered together by the \wedge connective. Interestingly, rule (3) is problematic in the design of polymorphic type systems for the join calculus [8, 15], essentially because it may introduce polymorphic (mutual) recursion. This, however, is not an issue in the present setting. On the contrary, the law is instrumental in establishing our results.

The operational semantics adopted here corresponds to the original one based on the chemical abstract machine [2]. Terms within a reduction context play the role of ‘molecules’, definitions determine the ‘reactions’. The reduction context determines which join pattern is matched in a definition and binds its received names. The reduction rule replaces in the context the matched pattern by the right-hand side of its definition.

Definition 2 (cf. [15]). The *reduction* \longrightarrow is the smallest substitutive relation on \equiv -classes that satisfies the following rule:

$$\text{def } D \wedge J \triangleright P \text{ in } \mathcal{C}[\sigma J] \longrightarrow \text{def } D \wedge J \triangleright P \text{ in } \mathcal{C}[\sigma P] \quad \text{if } dom(\sigma) \subseteq rn(J),$$

where D may be absent (if $J \triangleright P$ is the only definition), \mathcal{C} is a *reduction context*, i.e.,

$$\mathcal{C} ::= [] \mid \text{def } D \text{ in } \mathcal{C} \mid P \mid \mathcal{C},$$

and $\mathcal{C}[P]$ denotes the reduction context \mathcal{C} with its hole filled by P .

On Linearity. Here we no longer require join patterns to be linear, i.e., we allow names to occur several times in patterns. There are two aspects to this. The first one consists of renouncing the linearity of defined names in join patterns. It presents no further difficulties for implementations [11], and we see no reason not to consider it. In terms of Petri nets it corresponds to adding multiplicities to

arcs. More problematic in distributed implementations is relaxing the linearity of received names, because this amounts to equating values on different channels. Nevertheless we decided to adopt it, as it corresponds to a form of pattern matching essential in coloured Petri nets. It is worth remarking that our work is largely independent of this choice: we could as well keep linearity of received names at the price of disallowing matching in our coloured nets, and our results would stay, *mutatis mutandis*.

2 Place Transition Nets

A *multiset* on a set P is a function $\mu: P \rightarrow \mathbb{N}$. We shall use $\mu(P)$ to denote the set of finite multisets, i.e., markings, on P . The *sum* of $\mu_0, \mu_1 \in \mu(P)$ is the multiset $\mu = \mu_0 \oplus \mu_1$ such that $\forall p \in P. \mu(p) = \mu_0(p) + \mu_1(p)$.

Petri nets consists of places, transitions, and tokens. In view of relating nets and join terms, we identify places with names. In addition to \mathcal{Nm} – that will play the role of *public* places (*free* names) – we shall consider an infinite set ω of distinguished *private* places (*bound* names). In the following, \mathcal{Nm}_ω stands for the union $\mathcal{Nm} + \omega$.

Definition 3 (PT Nets). A *place/transition net* is a tuple $N = (T, \partial_0, \partial_1, \mu_0)$, where T is a finite set of transitions, $\partial_0, \partial_1: T \rightarrow \mu(\mathcal{Nm}_\omega)$ are the pre- and post-set functions, and $\mu_0 \in \mu(\mathcal{Nm}_\omega)$ is the initial marking. We shall use Π_0 to refer to PT Petri nets.

Thus, to simplify notations, we include in every net all the names of \mathcal{Nm}_ω with the understanding that *only* marked places, i.e., those carrying tokens, and places connected to some transitions are effectively to be considered in the net. In particular, the *empty net*, that we denote by \emptyset , formally consists of all the places in \mathcal{Nm}_ω , but no tokens and no transitions. Analogously, the net consisting of a distribution of tokens on a *multiset* of places μ , has all places but only $\mu(x)$ tokens in each $x \in \mathcal{Nm}_\omega$. With abuse of notations, we shall use μ to denote both the multisets and the corresponding net. Also, for $\mu_0, \mu_1 \in \mu(\mathcal{Nm}_\omega)$, **tran** $\mu_0 \mu_1$ stands for the net with a unique transition t with $\partial_0(t) = \mu_0$ and $\partial_1(t) = \mu_1$.

Following our intuition, isomorphisms preserve free names. This is detailed in the definition below, where we use *id* for identities, $+$ for function coproducts, and take the liberty of identifying functions with their extensions to multisets.

Definition 4. Nets $(T, \partial_0, \partial_1, \mu_0)$ and $(T', \partial'_0, \partial'_1, \mu'_0)$ are *isomorphic* if there exist isomorphisms $f_t: T \rightarrow T'$ and $f_p: \omega \rightarrow \omega$ such that $\partial'_i \circ f_t = (id_{\mathcal{Nm}} + f_p) \circ \partial_i$, for $i = 0, 1$, and $(id_{\mathcal{Nm}} + f_p)(\mu_0) = \mu'_0$.

Observe that nets that only differ for the use of names in ω are isomorphic, yielding a form of α -conversion. In the following, isomorphic nets will be regarded as equal. In particular, as we shall see, net isomorphism corresponds to structural congruence of join terms.

Definition 5. Let $N = (T, \partial_0, \partial_1, \mu_0)$ and $N' = (T', \partial'_0, \partial'_1, \mu'_0)$ be nets. Denoted by \otimes , *parallel composition* juxtaposes nets, merging public places without confusing the private ones. Formally,

$$N \otimes N' = (T + T', \partial_0 + \partial'_0, \partial_1 + \partial'_1, \mu_0 \oplus \mu'_0), \quad \text{if } \omega(N) \cap \omega(N') = \emptyset,$$

where $\omega(N)$ denotes the names in ω that are used in N , i.e., that are either marked or connected to some transition. Observe that the side condition on private names can always be achieved up to isomorphism. Clearly, \otimes is commutative and associative.

Denoted by ν_x , *restriction* ‘hides’ the place $x \in \mathcal{Nm}$ replacing it by a *fresh* name $i \in \omega$, i.e., a name not occurring in $\omega(N)$. Formally, for $N = (T, \partial_0, \partial_1, \mu_0)$

$$\nu_x N = N[x \leftrightarrow i], \quad i \text{ is fresh in } N$$

where $N[x \leftrightarrow i]$ is the net $(T, \partial'_0, \partial'_1, \mu'_0)$ whose μ'_0 and ∂'_j coincide respectively with μ_0 and ∂_j ($j = 0, 1$), but for the values they yield on x and i , that are exchanged, i.e., for $\delta \in \{\partial_0, \partial_1, \mu_0\}$, we have that $\delta'(k)$ is equal to $\delta(i)$ if $k = x$, to $\delta(x)$ if $k = i$, and to $\delta(k)$ otherwise.

For $X = \{x_1, \dots, x_n\} \subseteq \mathcal{Nm}$, we use $\nu_X N$ to mean $\nu_{x_1}(\dots \nu_{x_n} N)$. This is a correct definition because $\nu_x \nu_y N = \nu_y \nu_x N$ for each x, y , up to isomorphism.

The evolution of nets is described in terms of the ‘firing’ of its transitions. As usual, the firing of t consumes and produces resources, as prescribed by the pre- and post-set functions ∂_i . In the present setting, using \otimes and ν_X to form net contexts, this can be expressed as the least *substitutive* relation $[\cdot]$ such that

$$\mathbf{tran} \mu_0 \mu_1 \otimes \mu_0 [\cdot] \mathbf{tran} \mu_0 \mu_1 \otimes \mu_1.$$

This should be read as saying that, in any context, a marking matching the preset of a transition can be replaced by the associated post-set.

Observe that we stop at the single transition semantics of nets, rather than considering also the usual step semantics. We do so in order to match the standard reduction semantics of the join calculus, and it is not a hard limitation. We could easily extend the results to the classical step semantics, provided we equip join terms with multiple concurrent reductions.

The Type System Δ_0

The purpose of Δ_0 is to single out those join terms that are PT Petri nets. In order to achieve this we need to prevent any form of name passing and mobility, imposing a static network structure to terms. In particular, since definitions represent transitions, it is fundamental that their right-hand sides consist of messages only. Also, corresponding to the fact that there is only one atom (*token*) of information delivered in PT nets, we enforce that only empty messages may be exchanged. For notational convenience, we assume in the following the existence of a distinguished name $\bullet \in \mathcal{Nm}$ that represents the empty tuple.

The types of Δ_0 are \diamond and $\diamond\diamond$, with τ ranging over them. There are three kinds of type judgement:

$$\begin{aligned} \vdash P : \diamond\diamond & \quad (P \text{ is ok and no } \text{def_in_} \text{ occurs in it}) \\ \vdash P : \diamond & \quad (P \text{ is ok}) \quad \vdash D : \diamond \quad (D \text{ is ok}). \end{aligned}$$

The typing rules are the following.

$\begin{array}{c} \text{(P-ZERO)} \\ \vdash 0 : \diamond\diamond \end{array}$	$\begin{array}{c} \text{(P-MESS)} \\ \vdash x\langle\bullet\rangle : \diamond\diamond \end{array}$	$\begin{array}{c} \text{(P-PAR)} \\ \frac{\vdash P : \tau \quad \vdash Q : \tau}{\vdash P \mid Q : \tau} \end{array}$
$\begin{array}{c} \text{(P-DEF)} \\ \frac{\vdash D : \diamond \quad \vdash P : \diamond}{\vdash \text{def } D \text{ in } P : \diamond} \end{array}$	$\begin{array}{c} \text{(P-SUB)} \\ \frac{\vdash P : \diamond\diamond}{\vdash P : \diamond} \end{array}$	
$\begin{array}{c} \text{(D-PAT)} \\ \frac{\vdash P : \diamond\diamond}{\vdash J \triangleright P : \diamond} \quad \text{rn}(J)=\{\bullet\}$	$\begin{array}{c} \text{(D-AND)} \\ \frac{\vdash D : \diamond \quad \vdash E : \diamond}{\vdash D \wedge E : \diamond} \end{array}$	

The rules are elementary. The key is (D-PAT) that allows *only* simple processes as right-hand sides of definitions, while (P-MESS) ensures that only emissions of \bullet are well typed and, together with (P-PAR), that only parallel composition of messages can be used inside definitions. We use $\Delta_0 \vdash P$ as a shorthand for typeability, i.e., $\vdash P : \diamond$, in the system Δ_0 above.

Proposition 1 (Subject Reduction). *If $\Delta_0 \vdash P$ and $P \longrightarrow Q$, then $\Delta_0 \vdash Q$.*

Translating Δ_0 terms to PT Nets

We intend to define a translation from Δ_0 -typed join terms to PT nets that can be extended to other type systems and other classes of nets. Therefore, we first give a general map on all join terms, and then prove that it restricts to a well defined encoding of terms typeable in Δ_0 into PT nets. By induction on the structure of terms, $\llbracket _ \rrbracket$ is defined as follows.

$\llbracket 0 \rrbracket$	$= \emptyset$	$\llbracket J \triangleright P \rrbracket$	$= \mathbf{tran}[\llbracket J \rrbracket] [\llbracket P \rrbracket]$
$\llbracket x\langle v \rangle \rrbracket$	$= (x, v)$	$\llbracket D \wedge E \rrbracket$	$= \llbracket D \rrbracket \otimes \llbracket E \rrbracket$
$\llbracket P \mid Q \rrbracket$	$= \llbracket P \rrbracket \otimes \llbracket Q \rrbracket$	$\llbracket \text{def } D \text{ in } P \rrbracket = \nu_{dn(D)}(\llbracket D \rrbracket \otimes \llbracket P \rrbracket)$	

where $\llbracket J \rrbracket$ is J seen as a multiset, i.e., $\llbracket x\langle v \rangle \rrbracket = (x, v)$ and $\llbracket J \mid K \rrbracket = \llbracket J \rrbracket \oplus \llbracket K \rrbracket$.

Roughly speaking, we map names to places, messages to markings, and definitions to (groups of) transitions. Since we shall soon consider nets whose tokens carry information, we use (x, v) to denote a marking of place x with value v . Of course, in the current case v always equals \bullet , and (x, \bullet) should be read as

an exotic way to mean x . In more details, $J \triangleright P$ is mapped into a transition whose pre-set is essentially J and whose post-set is the encoding of P , while $|$ and \wedge are mapped homomorphically to \otimes . Similarly for $\text{def } D \text{ in } P$, where the use of restriction ensures the correct treatment of local definitions. Here it is important to understand the different roles of $x\langle v \rangle$ because, depending on the context, (x, \bullet) is treated as the net consisting of a single token in place x , or as an arc connecting place x to a transition. We shall get back to this in §3.

Example 1. Let P be the Δ_0 -typeable term $\text{def } x(\bullet) \triangleright y(\bullet) \text{ in } x(\bullet)$. Then, using the standard representation of Petri nets enriched with labels that decorate *public* places with their name, the translation of P is

$$\llbracket P \rrbracket = \nu_x (\text{tran}\{x\}\{y\} \otimes (x, \bullet)) = \textcircled{x} \rightarrow \text{█} \rightarrow \textcircled{y}.$$

Observe that, in force of the restriction ν_x , the name x does not appear.

As announced, the encoding is well defined on well-typed terms.

Proposition 2 (Correctness). *If $\Delta_0 \vdash P$, then $\llbracket P \rrbracket$ is a PT net.*

In addition to the previous proposition, we can prove a completeness theorem for Δ_0 , expressed as follows.

Theorem 1 (Completeness). *The translation $\llbracket _ \rrbracket$ is an isomorphism between the set of terms typeable in Δ_0 and PT nets in Π_0 .*

The rest of this section is devoted to illustrating that the $\llbracket _ \rrbracket$ is very tight a connection, preserving both structural congruence and dynamic behaviour.

Proposition 3. *Let P and Q be join terms typeable in Δ_0 . Then,*

$$P \equiv Q \quad \text{if and only if} \quad \llbracket P \rrbracket = \llbracket Q \rrbracket.$$

Observe that there is a one-to-one correspondence between transitions in $\llbracket P \rrbracket$ and simple definitions $J \triangleright Q$ in P . Here and in the following we use $N [t] N'$ to mean that N becomes N' by firing transitions t , and use $P -t \rightarrow Q$ to signify that P reduces to Q by means of the simple definition from which the transition t of $\llbracket P \rrbracket$ was generated.

Proposition 4. *Let P be a join term typeable in Δ_0 . Then,*

$$P -t \rightarrow Q \quad \text{if and only if} \quad \llbracket P \rrbracket [t] \llbracket Q \rrbracket$$


This implies that $P \longrightarrow^* Q$ if and only if there exist transitions t_1, \dots, t_n such that $\llbracket P \rrbracket [t_1] \dots [t_n] \llbracket Q \rrbracket$. Also, P has a barb $x \in \text{fn}(P)$, i.e., $P \equiv \mathcal{C}[x\langle v \rangle]$ for \mathcal{C} a context that does not bind x , if and only if x is marked in $\llbracket P \rrbracket$. Together with surjectivity and compositionality of $\llbracket _ \rrbracket$, this describes a very strong relationship, allowing to identify (the system represented by) P and (that represented by) $\llbracket P \rrbracket$. In other words, it allows us to identify typeability in Δ_0 with Π_0 .

3 Coloured Nets

In coloured nets, tokens carry information: the ‘colours’. Arcs between places and transitions are labelled by *arc expressions*, which evaluate to multisets of coloured tokens after binding free variables to colours. Expressions on input and output arcs describe respectively the resources needed for the firing and those generated by it. Although our treatment below attempts at providing an essential version of coloured Petri nets, we maintain one of their fundamental features. Namely, we permit a name to label several of the input arcs of a transition, with the intended meaning that firing is allowed only if the same colour can be fetched along all those arcs. This form of pattern matching is the reason why we relaxed the linearity requirement of join patterns. Technically, our model is obtained simply by allowing PT Nets to circulate tokens other than ‘•’. The obvious choice in the present context is to identify colours with names.

Definition 6 (CNets). A *coloured net* is a tuple $(T, \partial_0, \partial_1, \mu_0)$, where T is a finite set of transitions, $\partial_0, \partial_1: T \rightarrow \mu(\mathcal{Nm}_\omega \times \mathcal{Nm})$, are the pre and post-set functions, and $\mu_0 \in \mu(\mathcal{Nm}_\omega \times \mathcal{Nm})$ is the initial marking. We use Π_1 to denote the class of coloured Petri nets.

Example 2. In the above definition we interpret the pairs (p, c) of $\mathcal{Nm}_\omega \times \mathcal{Nm}$ as representing a place p that carries a colour (name) c . The coloured net $\mathbf{tran}\{(x, v)\}\{(y, v)\} \otimes (x, c)$ is represented graphically as below on the left.



As usual, the role of names is twofold: in markings they represent colours, e.g., the constant c in x ; in transitions they represent bound variables. The net above fetches any token from x , binds it to v , and then delivers the actual value of v to y , as illustrated by the picture above and formalised by the definition below.

Definition 7. The set of *received colours* of a marking $\mu \in \mu(\mathcal{Nm}_\omega \times \mathcal{Nm})$ is

$$\mathbf{rc}(\mu) = \{x \in \mathcal{Nm} \mid x \neq \bullet, \exists p. (p, x) \in \mu\}.$$

A *binding* for μ is a function $\mathbf{b}: \mathcal{Nm} \rightarrow \mathcal{Nm}$ which is the identity on $\mathcal{Nm} \setminus \mathbf{rc}(\mu)$. We shall let $\mu(\mathbf{b})$ denote the multiset $\{(p, \mathbf{b}(c)) \mid (p, c) \in \mu\}$.

The firing rule for coloured nets is the *substitutive* relation generated by the reduction rule below, where \mathbf{b} is a binding for μ_0

$$\mathbf{tran} \mu_0 \mu_1 \otimes \mu_0(\mathbf{b}) [\cdot] \mathbf{tran} \mu_0 \mu_1 \otimes \mu_1(\mathbf{b}).$$

Isomorphisms of coloured nets are the obvious extension of those of PT nets given in Definition 4. They allow renaming of private places, but map public ones identically. It is important to remark that coloured nets are considered up to α -conversion, that is renaming of received colours. In particular, coloured nets are isomorphic if there is an isomorphism between them after possibly renaming received names.

The Type System Δ_1

The purpose of type system Δ_1 is to characterise coloured Petri nets among join terms. Having introduced names, we now face the issue of distinguishing among two kind of names: channels and parameters. In fact, coloured nets are a strict value-passing formalism. They are not allowed to use values received along channels as channels themselves, nor to send private names as messages. In order to enforce this, we consider typing environments holding assumptions on free names, *viz.*, if they are channels or messages. Type environments are therefore pairs of *disjoint* sets Γ (the channels) and ∇ (the messages). Type judgements are exactly as before, but for the presence of type environments.

$\text{(P-ZERO)} \quad \Gamma; \nabla \vdash 0 : \diamond$	$\text{(P-MESS)} \quad \Gamma; \nabla \vdash x(y) : \diamond$	$\text{(P-PAR)} \quad \frac{\Gamma; \nabla \vdash P : \tau \quad \Gamma; \nabla \vdash Q : \tau}{\Gamma; \nabla \vdash P \mid Q : \tau}$
$\text{(P-DEF)} \quad \frac{\Gamma, dn(D); \nabla \vdash D : \diamond \quad \Gamma, dn(D); \nabla \vdash P : \diamond}{\Gamma; \nabla \vdash \text{def } D \text{ in } P : \diamond}$		$\text{(P-SUB)} \quad \frac{\Gamma; \nabla \vdash P : \diamond}{\Gamma; \nabla \vdash P : \diamond}$
$\text{(D-PAT)} \quad \frac{\Gamma; \nabla, rn(J) \vdash P : \diamond}{\Gamma; \nabla \vdash J \triangleright P : \diamond}$		$\text{(D-AND)} \quad \frac{\Gamma; \nabla \vdash D : \diamond \quad \Gamma; \nabla \vdash E : \diamond}{\Gamma; \nabla \vdash D \wedge E : \diamond}$

The structure of the rules matches exactly those of Δ_0 , the only difference being the use of Γ in (P-DEF) and of ∇ in (D-PAT), so to be able to control the use of names in (P-MESS).

Proposition 5. $\Delta_0 \vdash P$ implies $\Delta_1 \vdash P$.

Proposition 6 (Subject Reduction). If $\Delta_1 \vdash P$ and $P \longrightarrow Q$, then $\Delta_1 \vdash Q$.

A simple inspection of the rules shows that, as Δ_0 , system Δ_1 does not allow processes of the form $\text{def } _ \text{ in } _$ to appear inside definitions. The other fundamental properties of Δ_1 are expressed in the proposition below.


Proposition 7. If $\Delta_1 \vdash P$, then P never emits messages on received names nor emits bound names out of the scope of their definitions.

Translating Δ_1 terms to C Nets

It follows as a consequence of Proposition 7 that the translation $\llbracket _ \rrbracket$ extends to a well defined map from Δ_1 typeable terms to coloured Petri nets.

Proposition 8 (Correctness). If $\Delta_1 \vdash P$, then $\llbracket P \rrbracket$ is a coloured net.

Here more than before it is important to understand the three different roles played by $x\langle v \rangle$ in the translation. First, when $x\langle v \rangle$ appears outside all definitions, at top level, it is translated as the net consisting of a token v in the place x . Secondly, when it appears in a join pattern J , it is translated in $\llbracket J \rrbracket$ as (x, v) , representing the input from place x on a bound variable v . Finally, when $x\langle v \rangle$ appears in the body of a definition, it is translated as (x, v) by the same clause above but, considered as a marking, it represents the output of a free or bound variables in the place x . This is summarised by the pictures below.



Example 3. Let P be the Δ_1 typeable term $\text{def } x\langle v \rangle \triangleright y\langle v \rangle \text{ in } x\langle c \rangle$. Then $\llbracket P \rrbracket$ is the coloured net of Example 2.

Our results about Δ_1 match exactly those for Δ_0 . They are listed below.

Theorem 2 (Δ_1 vs Π_1). *The translation $\llbracket _ \rrbracket$ is an isomorphism between the set of terms typeable in Δ_1 and coloured nets in Π_1 . Moreover, if P and Q are typeable in Δ_1 , then*

$$\begin{aligned}
 P \equiv Q & \quad \text{if and only if} \quad \llbracket P \rrbracket = \llbracket Q \rrbracket \\
 P -t \rightarrow Q & \quad \text{if and only if} \quad \llbracket P \rrbracket [t] \llbracket Q \rrbracket
 \end{aligned}$$

4 Reconfigurable Nets

Coloured Petri nets have a static structure, i.e., their firings only affects the marking. While this accounts faithfully for value-passing theories, it is completely inadequate to represent dynamically changing structures. Reconfigurable nets generalise coloured ones by adding precisely one ingredient: each firing of a transition may deliver to a different set of output places. This equips nets with a mechanism to model networks with reconfigurable topologies, that is networks in which the set of components is fixed, but the connectivity among them may change in time. Formally, this is achieved by a very smooth alteration of the definitions of coloured nets – one that allows markings and messages to contain private names – and of their firing rule.

Definition 8 (RNets). A *reconfigurable net* is a tuple $(T, \partial_0, \partial_1, \mu_0)$, where T is a finite set of transitions, $\partial_0: T \rightarrow \mu(\mathcal{N}m_\omega \times \mathcal{N}m)$ and $\partial_1: T \rightarrow \mu(\mathcal{N}m_\omega \times \mathcal{N}m_\omega)$, are, respectively, the pre and post-set functions, and $\mu_0 \in \mu(\mathcal{N}m_\omega \times \mathcal{N}m_\omega)$ is the initial marking. We use Π_2 to denote the set of reconfigurable Petri nets.

Restriction ν_x and parallel composition \otimes extends straightforwardly to reconfigurable nets. Also, isomorphisms of reconfigurable Petri nets are obtained extending those of coloured nets in the obvious way.


The firing rule is generalised by extending the scope of the binding of received names to include the output places in addition to the names on the output arcs.

Definition 9. Let \mathbf{b} be a binding. For $\mu \in \mu(\mathcal{Nm}_\omega \times \mathcal{Nm}_\omega)$, let $X\langle\langle\mathbf{b}\rangle\rangle$ denote the multiset $\{(\mathbf{b}_\omega(p), \mathbf{b}_\omega(c)) \mid (p, c) \in \mu\}$, where $\mathbf{b}_\omega = \mathbf{b} + id_\omega$.

For \mathbf{b} a binding for μ_0 , the firing is generated closing by net contexts the rule

$$\mathbf{tran} \mu_0 \mu_1 \otimes \mu_0 \langle \mathbf{b} \rangle [\cdot] \mathbf{tran} \mu_0 \mu_1 \otimes \mu_1 \langle \langle \mathbf{b} \rangle \rangle.$$

Example 4. The simple reconfigurable net represented below binds v to any name found in x – in this case c – and delivers the name x to the place to which v is bound.



The Type System Δ_2

In order to design a type system Δ_2 corresponding to Π_2 we only need to remove the difference between channels and messages upon which Δ_1 is based, which can be done easily by removing Γ and ∇ . This yields a system identical to Δ_0 , apart from (P-MESS) that allows any message to be transmitted on any name.

(P-ZERO) $\vdash 0 : \diamond$	(P-MESS) $\vdash x(y) : \diamond$	(P-PAR) $\frac{\vdash P : \tau \quad \vdash Q : \tau}{\vdash P \mid Q : \tau}$
(P-DEF) $\frac{\vdash D : \diamond \quad \vdash P : \diamond}{\vdash \text{def } D \text{ in } P : \diamond}$	(P-SUB) $\frac{\vdash P : \diamond}{\vdash P : \diamond}$	
(D-PAT) $\frac{\vdash P : \diamond}{\vdash J \triangleright P : \diamond}$	(D-AND) $\frac{\vdash D : \diamond \quad \vdash E : \diamond}{\vdash D \wedge E : \diamond}$	

Proposition 9 (Subject Reduction). *If $\Delta_2 \vdash P$ and $P \longrightarrow Q$, then $\Delta_2 \vdash Q$.*

Proposition 10. *$\Delta_1 \vdash P$ implies $\Delta_2 \vdash P$.*

Translating Δ_2 terms to R Nets

The property guaranteed by Δ_2 is that there is no process of the kind `def _ in _` inside definitions, i.e., the topology of the net may change by redirecting output edges, but the components of the net are fixed once and for all. In force of this, we can extend to Δ_2 the correspondence between well-typed terms and nets.

Theorem 3 (Δ_2 vs Π_2). *The translation $\llbracket _ \rrbracket$ is an isomorphism between the set of terms typeable in Δ_2 and reconfigurable nets in Π_2 . Moreover, if P and Q are typeable in Δ_2 , then*

$$\begin{aligned} P \equiv Q & \text{ if and only if } \llbracket P \rrbracket = \llbracket Q \rrbracket \\ P -t \rightarrow Q & \text{ if and only if } \llbracket P \rrbracket [t] \llbracket Q \rrbracket \end{aligned}$$

5 Dynamic Nets

The obvious generalisation of reconfigurable nets is to allow the dynamic creation of components, that we achieve by means of the notions of *dynamic Petri nets*. The idea behind such structures is that the firing of a transition allocates a new net parametric in the actual values of the received names. As the net may consist simply of a marking and no transitions, this includes the standard definition of PT nets. Also coloured and reconfigurable nets are, of course, special kinds of dynamic nets. The characteristic feature of dynamic nets, that to our best knowledge distinguishes them by other approaches in the literature and draws a connection to [13], is that, as for reconfigurable nets, input arcs are never modified. While it is possible to modify dynamically the post-set of a transition, and also to spawn new subnets, it is not possible to add places to the presets of transitions. This allows us to formalise our intuition by simply generalising the post-set functions of nets, allowing $\partial_1(t)$ to be a dynamic net. Of course, this means that the definition of nets becomes recursive.

Definition 10 (DNets). Let $DNets$ be the least set satisfying the equation

$$DNets \cong \{(T, \partial_0: T \rightarrow \mu(Nm_\omega \times Nm), \partial_1: T \rightarrow DNets, \mu_0)\},$$


where T is a finite set and μ_0 , the initial marking, is in $\mu(Nm_\omega \times Nm_\omega)$. A *dynamic net* DN is an element of $DNets$. We use Π_3 to refer to the set of dynamic nets.

Once again, the operations of parallel composition and restriction, and the notions of α -conversion and isomorphism lift smoothly to this setting.

Definition 11. Dynamic nets $(T, \partial_0, \partial_1, \mu_0)$ and $(T', \partial'_0, \partial'_1, \mu'_0)$ are isomorphic if, up to α -conversion, there exists a pair $(f_t: T \rightarrow T', f_p: \omega \rightarrow \omega)$ of isomorphisms such that $(id_{Nm} + f_p \times id_{Nm} + f_p)(\mu_0) = \mu'_0$, $\partial'_0 \circ f_t = (id_{Nm} + f_p \times id_{Nm}) \circ \partial_0$, and, recursively, $\partial_1(t)$ and $\partial'_1(f_t(t))$ are isomorphic for each $t \in T$.

Example 5. For $P = \text{def } x\langle v \rangle \triangleright (\text{def } y\langle \bullet \rangle \triangleright x\langle v \rangle \text{ in } y\langle \bullet \rangle) \text{ in } x\langle a \rangle \mid x\langle b \rangle \mid x\langle c \rangle$, the structure $\llbracket P \rrbracket$ is the dynamic net illustrated below together with one of its possible firing sequences. The transition may fire three times, binding v in turn to a , b , and c , and spawning three new nets instantiating v appropriately in N . The free name x of N maintains its meaning across instantiations. This rules

are formalised below, and amounts to a recursive extension of binding.



Definition 12. Let \mathbf{b} be a binding. For $X = (T, \partial_0, \partial_1, \mu_0) \in DNets$, let $X \lll \mathbf{b} \rrr$ denote the dynamic net obtained from X by substituting free names according to \mathbf{b} in μ_0 , $\partial_0(t)$ and, recursively, applying $\lll \mathbf{b} \rrr$ to $\partial_1(t)$, for all $t \in T$.

The firing for dynamic nets is generated as follows, for \mathbf{b} a binding for μ_0 .

$$\mathbf{tran} \mu_0 X \otimes \mu_0 \langle \mathbf{b} \rangle [\cdot] \mathbf{tran} \mu_0 X \otimes X \lll \mathbf{b} \rrr$$

The Type System Δ_3

In order to capture dynamic nets in terms of the join calculus, we simply have to relax the restriction that disallows definitions inside definitions. At the level of type system, this can be achieved by removing the distinction between types \diamond and $\diamond\diamond$ from Δ_2 . As an obvious consequence, any join term results typeable.

Proposition 11. For each join term P , $\Delta_3 \vdash P$.

Thus, extending our results for $\llbracket _ \rrbracket$ to Δ_3 amounts to proving that dynamic nets coincide with join terms, which is the content of the following.

Theorem 4 (Δ_3 vs Π_3). The translation $\llbracket _ \rrbracket$ is an isomorphism between join terms (terms typeable in Δ_3) and dynamic nets in Π_3 . Moreover, if P and Q are typeable in Δ_3 , then

$$\begin{aligned} P \equiv Q & \quad \text{if and only if} \quad \llbracket P \rrbracket = \llbracket Q \rrbracket \\ P \rightarrow t \rightarrow Q & \quad \text{if and only if} \quad \llbracket P \rrbracket [t] \llbracket Q \rrbracket \end{aligned}$$

6 The Type System Δ_4

The information contained in Δ_i , $i = 0, \dots, 3$ can be summarised in a single type system Δ_4 that tags terms with indices that characterise them as nets of Π_i . Types range over $i^{\diamond\diamond} \in \{0^{\diamond\diamond}, 1^{\diamond\diamond}, 2^{\diamond\diamond}, 3^{\diamond\diamond}\}$ and $i^\diamond \in \{0^\diamond, 1^\diamond, 2^\diamond, 3^\diamond\}$. For $\tau = i^?$ a type, we write $\tau \downarrow$ for i and $\tau \uparrow$ for $?$. By $\tau \leq \tau'$ we mean that $\tau \downarrow \leq \tau' \downarrow$ and $\tau \uparrow \leq \tau' \uparrow$, with the convention that $\diamond \leq \diamond$. Type environments are pairs $\Gamma; \nabla$ as for Δ_1 , and type judgements are as follows:

$$\begin{aligned} \Gamma; \nabla \vdash P : i^{\diamond\diamond} & \quad (P \text{ well-typed, in } \Pi_i, \text{ and containing no def } _ \text{ in } _) \\ \Gamma; \nabla \vdash P : i^\diamond & \quad (P \text{ well-typed and in } \Pi_i) \\ \Gamma; \nabla \vdash D : i^\diamond & \quad (D \text{ well-typed and containing terms in } \Pi_i) \end{aligned}$$

The typing rules are the following.

$\text{(P-MESS}_0\text{)} \quad \Gamma; \nabla \vdash x(\bullet) : 0^\diamond \quad (x \notin \nabla)$	$\text{(P-MESS}_1\text{)} \quad \Gamma; \nabla \vdash x\langle y \rangle : 1^\diamond \quad \left(\begin{array}{l} x \notin \nabla \\ y \notin \Gamma \end{array} \right)$	$\text{(P-MESS}_2\text{)} \quad \Gamma; \nabla \vdash x\langle y \rangle : 2^\diamond$
$\text{(P-DEF)} \quad \frac{\Gamma, dn(D); \nabla \vdash D : i^\diamond \quad \Gamma, dn(D); \nabla \vdash P : i^\diamond}{\Gamma; \nabla \vdash \text{def } D \text{ in } P : i^\diamond}$		$\text{(P-ZERO)} \quad \Gamma; \nabla \vdash 0 : 0^\diamond$
$\text{(P-PAR)} \quad \frac{\Gamma; \nabla \vdash P : \tau \quad \Gamma; \nabla \vdash Q : \tau}{\Gamma; \nabla \vdash P \mid Q : \tau}$		$\text{(P-SUB)} \quad \frac{\Gamma; \nabla \vdash P : \tau \quad \tau \leq \tau'}{\Gamma; \nabla \vdash P : \tau'}$
$\text{(D-PATT}_0\text{)} \quad \frac{\Gamma; \nabla \vdash P : 0^\diamond}{\Gamma; \nabla \vdash J \triangleright P : 0^\diamond}$	$\text{(D-PATT}_1\text{)} \quad \frac{\Gamma; \nabla, rn(J) \vdash P : i^\diamond}{\Gamma; \nabla \vdash J \triangleright P : i^\diamond}$	$\text{(D-PATT}_2\text{)} \quad \frac{\Gamma; \nabla, rn(J) \vdash P : i^\diamond}{\Gamma; \nabla \vdash J \triangleright P : 3^\diamond}$
$\text{(D-AND)} \quad \frac{\Gamma; \nabla \vdash D : i^\diamond \quad \Gamma; \nabla \vdash E : i^\diamond}{\Gamma; \nabla \vdash D \wedge E : i^\diamond}$		$\text{(D-SUB)} \quad \frac{\Gamma; \nabla \vdash D : i^\diamond \quad i < j}{\Gamma; \nabla \vdash D : j^\diamond}$

Here (P-MESS) is split into three rules, each behaving as in the corresponding Δ_i ; (P-MESS₂) ignores type environments, so achieving the effect of Δ_2 . Definitions can be typed by i if the terms they contain are typeable by i ; (D-PATT₂) constrains to 3 the type of definitions containing a `def_in` term, so forcing to 3 the type of enclosing processes by means of rule (P-DEF). For Δ_4 we have the following results.

Proposition 12 (Subject Reduction). *If $\Delta_4 \vdash P$ and $P \longrightarrow Q$, then $\Delta_4 \vdash Q$.*

Theorem 5. $\Delta_4 \vdash P : i^\diamond$ if and only if $P \in \Pi_i$, for $i = 0, \dots, 3$.

Proof. It follows by proving that $\Delta_4 \vdash P : i^\diamond$ if and only if $\Delta_i \vdash P$.

Conclusions and Future Work

We have provided a full correspondence between join calculus and a hierarchy of Petri net classes. It would be interesting to study relevance and adaptability to coloured nets of the existing polymorphic type systems [8, 15] and extensional semantic equivalence [10, 3, 6] for join terms. On the other hand, the body of work on the semantics of nets suggests a noninterleaving semantics based on monoidal categories for the join-calculus.

Our version of coloured nets simplifies considerably those in the literature. It is worth investigating whether it can be interesting for the coloured Petri net community by putting it at work on suitable applications.

Also, we plan to compare our nets to Milner's named nets [12] – clearly closely related to the nets of §2 – to Asperti and Busi's dynamic nets [1], and to Odersky's functional nets [14].

Acknowledgements. We heartily thank Cedric Fournet, Alan Schmitt, and Emilio Tuosto for their useful comments. Special thanks to Roberto Bruni, who made substantial suggestions, and to an anonymous referee, who spotted some technical flaws.

References

1. A. ASPERTI AND N. BUSI (1996), *Mobile Petri Nets*, Technical Report UBLCS 96-10, Università di Bologna.
2. G. BERRY AND G. BOUDOL (1992), The Chemical Abstract Machine, *Theoretical Computer Science*, 96:217–248.
3. M. BOREALE, C. FOURNET, AND C. LANEVE (1998), Bisimulations for the Join-Calculus, In *Proc. PROCOMET'98*, D. Gries and W.P. de Roever (Eds.), 68–86, IFIP, Chapman & Halls.
4. R. BRUNI, J. MESEGUER, U. MONTANARI, AND V. SASSONE (2000), Functorial Models for Petri Nets, *Information and Computation*. To appear.
5. C. FOURNET AND G. GONTHIER (1996), The Reflexive Chemical Abstract Machine and the Join-Calculus, In *Proc. POPL'96*, 372–385, ACM.
6. C. FOURNET AND G. GONTHIER (1996), A Hierarchy of Equivalences for Asynchronous Calculi, In *Proc. ICALP'98*, Lecture Notes in Computer Science 1443, 844–855, Springer.
7. C. FOURNET, G. GONTHIER, J. LÉVY, L. MARANGET, AND D. RÉMY (1996), A Calculus of Mobile Agents, In *Proc. CONCUR'96*, Lecture Notes in Computer Science 1119, 406–421, Springer.
8. C. FOURNET, C. LANEVE, L. MARANGET, AND D. RÉMY (1997), Implicit Typing à la ML for the Join-Calculus, In *Proc. CONCUR'97*, Lecture Notes in Computer Science 1243, 196–212, Springer.
9. K. JENSEN (1992), *Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use*. Monographs on Theoretical Computer Science, Springer.
10. C. LANEVE (1996), May and Must Testing in the Join-Calculus, Technical Report UBLCS 96-04, Università di Bologna.
11. F. LE FESSANT AND L. MARANGET (1998), Compiling Join-Patterns, In *Proc. HLCL'98*, Electronic Notes in Computer Science 16(3), Elsevier.
12. R. MILNER (1993), Action Calculi, or Syntactic Action Structures, In *Proc. MFCS'93*, Lecture Notes in Computer Science 711, 105–121, Springer.
13. M. NIELSEN, L. PRIESE, AND V. SASSONE (1995), Characterizing Behavioural Congruences for Petri Nets, in *Proc. CONCUR 95*, Lecture Notes in Computer Science 962, 175–189, Springer.
14. M. ODERSKY (2000), Functional Nets, In *Proc. ESOP'2000*, Lecture Notes in Computer Science 1782, 1–25, Springer.
15. M. ODERSKY, C. ZENGER, M. ZENGER, AND G. CHEN (1999), A Functional View of Join, Technical Report ACRC-99-016, University of South Australia.
16. C.A. PETRI (1962), *Kommunikation mit Automaten*. Ph.D. thesis, Institut für Instrumentelle Mathematik, Bonn.
17. W. REISIG (1985), *Petri Nets: An Introduction*. EATCS Monographs on Theoretical Computer Science, Springer.
18. V. SASSONE (2000), On the Algebraic Structure of Petri Nets, *Bulletin of EATCS* 72, 133–148.