# SAFE DYNAMIC BINDING IN THE JOIN CALCULUS

Alan Schmitt
*INRIA Rocquencourt*
alan.schmitt@inria.fr

**Abstract** This paper presents an extension of the distributed Join Calculus with messages dynamically bound to definitions according to their location. New dynamic channels and new definitions for dynamic channels may be created at runtime. A dynamic message is rebound when the location containing the message migrates. A sound type system is introduced to guarantee that every dynamic message is bound to a definition. A longer version of this paper with additional examples is available [17].

## 1. Introduction

In a distributed world, channel-based communication usually involves senders and receivers that may reside in different locations. A receiver is a *definition*, an association between a channel name and a process that is spawned upon reception of a message on this channel. A message sent on a channel is routed to a location containing a definition for this channel. Thus the binding of senders to receivers is dependent upon the routing of messages. On the one hand, the binding may be *static*, in the sense that the routing does not depend on where the message originates, greatly simplifying remote communication. The distributed Join Calculus [10, 9] makes such a choice, by restricting the definition of a given channel to a single location. On the other hand, the routing may depend on where the message is. This *dynamic* binding models location-dependent behaviors. This paper enriches the distributed Join Calculus with channels whose messages are routed dynamically. Because of space constraints, we do not present the distributed Join Calculus. Readers who do not feel familiar with this calculus can find a tutorial in [11].

The JoCaml system [15], an implementation of the distributed Join Calculus, provides some form of dynamic binding to functions defined in the current runtime. Modeling this behavior is also one of our goals.

One requirement for our system is the ability to create new definitions for existing dynamic channels. However, when considering such a system of first

R. Baeza-Yates et al. (eds.), *Foundations of Information Technology in the Era of Network and Mobile Computing* 

class dynamic channels where definitions may be added and new channel names may be created at runtime, it is difficult to check manually that every message is bound to a definition. This *receptiveness* property should prevent mistakes like trying to send a message on a channel when it is not bound. This paper aims at such a safe calculus of local resources using a static type system.

An important design choice is the relation between the process being rebound to new definitions and the process triggering this rebinding. As in [6], we distinguish between *objective* and *subjective* rebinding. The rebinding is objective if it is triggered by a guard, and the (guarded) process that is being rebound is not active. On the opposite, rebinding is subjective when it involves running processes that are in the environment of the process triggering the rebinding. As subjective dynamic calculi have less control on the scope of the rebinding, they let the programmer modify the execution environment of running processes. However, this makes the receptiveness property significantly harder to prove statically, since the process being rebound is not explicitly available when typing the process triggering the rebinding.

Our extension adds dynamic channels and corresponding definitions. The definition bound to a dynamic channel at a given location is the closest definition of this channel in the enclosing locations. We say that a dynamic channel is *available* when there is such a definition. Since migration in the distributed Join Calculus is subjective, our calculus is a subjective dynamic calculus. As a location may redefine a channel already defined in an enclosing location, this model is an extension of the JoCaml model, where only *runtimes*—top level locations that do not migrate—provide definitions for dynamic channels.

The dynamic Join Calculus is designed to be implemented in a distributed setting. In the distributed Join Calculus, each static channel is defined in a single location. This property does not hold for dynamic channels, as a dynamic channel may be defined in several locations. However, the routing of a message on a dynamic channel is deterministic as there is only one closest enclosing location defining this channel. Moreover, we require the determination of the destination of a message to be local. This insures that there is no need for distributed synchronization.

In section 2 we present the syntax and semantics of the dynamic Join Calculus; in section 3 we present a type system that guarantees the presence of definitions, and we state its soundness in section 4; we conclude in section 5.

## 2. Syntax and semantics

The chemical syntax of the dynamic Join Calculus is the following:

$$
\begin{array}{rcl}
\mathcal{P} & ::= & \mathbf{0} \mid \mathcal{P} \mid \mathcal{P}' \mid n\langle \widetilde{n_i} \rangle; P \mid \mathbf{go}\, n; P \mid \nu \mathbf{n}.P \mid a.n\langle \widetilde{n_i} \rangle \\
\mathcal{D} & ::= & \top \mid \mathcal{D}, \mathcal{D}' \mid J \triangleright P \mid a[\mathcal{D} : \mathcal{P}]^{\Delta, I} \\
J & ::= & n\langle \widetilde{y} \rangle \mid J \mid J' \\
\mathcal{S} & ::= & \Omega \mid \mathcal{S} \parallel \mathcal{S}' \mid \mathcal{D} \vdash_{\varphi a}^{\Delta, I, F} \mathcal{P} \\
\Delta, I & ::= & \emptyset \mid \{\mathbf{n}\} \mid \{y\} \mid \Delta \cup \Delta
\end{array}
$$

The structural and reduction semantics are given in figure 1. We presuppose the existence of an infinite set of names ranged over by $m, n, x$. Location names are ranged over by $a, b, c$, dynamic channel names are ranged over by $\mathbf{m}, \mathbf{n}, \mathbf{x}$, static names are ranged over by $\mathtt{m}, \mathtt{n}, \mathtt{x}$, and variables are ranged over by $u, y$. We write $\widetilde{n}$ for possibly empty tuples of names. We write $P$ (resp. $D$) for processes (resp. definitions) that do not contain any occurrence of resolved messages (of the form $a.n\langle\widetilde{n_i}\rangle$). We write $\mathcal{P}$ (resp. $\mathcal{D}$) for processes (resp. definitions) which may contain occurrences of resolved messages. Free names, received names, and defined names are defined as usual. A local definition $\mathtt{def}\ D\ \mathtt{in}\ P$ binds within $D$ and $P$ the static channel names and location names defined in $D$. However, it does not bind the dynamic channel names defined in $D$. A restriction $\nu\mathbf{n}.P$ binds the dynamic name $\mathbf{n}$ in $P$. A reaction rule $J \triangleright P$ binds the received names of $J$ in $P$. The formal definitions can be found in [17]. We also introduce the notion of *defined local names* (*dln*) as names that are defined in a given location, *defined static names* (resp. *defined dynamic names*) (*dsn*) (resp. *ddn*) as defined names that are static (resp. defined names that are dynamic).

Intuitively, a configuration consists of several concurrently running locations. Each location contains a multiset of definitions $\mathcal{D}$ and a multiset of running processes $\mathcal{P}$. As in the Join Calculus, locations are structured as a tree, and each location has a unique static name. Since the chemical semantics acts on a flat structure of running locations, the tree structure is reflected in the names of the running locations: a location has name $\varphi a$ if its name is $a$ and if the path from the root of the location tree to this location is $\varphi$.

In order to account for the different routings of static and dynamic messages, we split the routing in two steps (much as in [12]). The first step, called the *name lookup* step, resolves the location where to route the message, and prepends this location to the message. The destination is the location containing the definition the message is bound to. For static channels, it is the unique location defining the channel; for dynamic channels, it is the closest enclosing location containing a definition of the channel. The second step is the *communication* step, it corresponds to the migration of the resolved message to its destination (rule COMM). We remark that, unlike [19], our semantics only focuses on name lookup and does not deal with the actual routing.

To show that the name lookup of dynamic channels is local, each running location bears a lookup function $F$ from dynamic channels to the the name of the closest enclosing location defining the channel. This function is used in the dynamic name lookup rule (NL-DYN), where $\perp$ is the undefined location. The static name lookup rule (NL-STAT) resolves the unique location defining the channel. Since the name lookup step is local, we let the programmer define a continuation to unresolved messages that is spawned when the lookup occurs (rules NL-STAT and NL-DYN). However, the delivery and consumption of the message are asynchronous. In the following we may write $n\langle\widetilde{m}\rangle$ for $n\langle\widetilde{m}\rangle; \mathbf{0}$.

A definition has the form $n_1\langle\widetilde{y_1}\rangle \mid \ldots \mid n_k\langle\widetilde{y_k}\rangle \triangleright P$ where the $n_i$ are the channel names, the $\widetilde{y_i}$ are the received names, and $P$ is the guarded process. A channel name in a join pattern may either be of the form $\mathbf{n}_i$, if the channel

$$\frac{\mathcal{S} =_\alpha \mathcal{S}'}{\mathcal{S} \equiv \mathcal{S}'}\ [\textsc{Str-}\alpha] \qquad \frac{\forall \psi \in loc(\mathcal{S}), a \notin \psi \qquad G = Lookup(F, I, \Delta, a)}{a[\mathcal{D} : \mathcal{P}]^{\Delta,I} \vdash_{\varphi}^{\Delta',I',F} \equiv \vdash_{\varphi}^{\Delta',I',F} \parallel \mathcal{D} \vdash_{\varphi a}^{\Delta,I,G} \mathcal{P}}\ [\textsc{Str-Loc}]$$

$$\frac{dsn(D) \cap (bn(\mathcal{S}) \cup bn(\mathcal{D}, D) \cup bn(\mathcal{P} \mid P)) = \emptyset \qquad dln(D) \cap ddn(D) = \emptyset}{\mathcal{S} \parallel \mathcal{D} \vdash_{\varphi a}^{\Delta,I,F} \mathcal{P} \mid \mathtt{def}\ D\ \mathtt{in}\ P \longrightarrow \mathcal{S} \parallel \mathcal{D}, D \vdash_{\varphi a}^{\Delta,I,F} \mathcal{P} \mid P}\ [\textsc{Def}]$$

$$\frac{\{\mathbf{n}\} \cap (bn(\mathcal{S}) \cup bn(\mathcal{D}) \cup bn(\mathcal{P} \mid P)) = \emptyset}{\mathcal{S} \parallel \mathcal{D} \vdash_{\varphi a}^{\Delta,I,F} \mathcal{P} \mid \nu\mathbf{n}.P \longrightarrow \mathcal{S} \parallel \mathcal{D} \vdash_{\varphi a}^{\Delta,I,F} \mathcal{P} \mid P}\ [\textsc{Nu}]$$

$$\frac{\mathbf{n} \in dln(b)}{\vdash_{\varphi a}^{\Delta,I,F} \mathbf{n}\langle\widetilde{v}\rangle; P \longrightarrow \vdash_{\varphi a}^{\Delta,I,F} b.\mathbf{n}\langle\widetilde{v}\rangle \mid P}\ [\textsc{NL-Stat}]$$

$$\frac{F(\mathbf{n}) = b \wedge b \neq \bot}{\vdash_{\varphi a}^{\Delta,I,F} \mathbf{n}\langle\widetilde{v}\rangle; P \longrightarrow \vdash_{\varphi a}^{\Delta,I,F} b.\mathbf{n}\langle\widetilde{v}\rangle \mid P}\ [\textsc{NL-Dyn}]$$

$$\frac{dom(\sigma_{rn}) = rn(J)}{J \triangleright P \vdash_{\varphi a}^{\Delta,I,F} a.J\sigma_{rn} \longrightarrow J \triangleright P \vdash_{\varphi a}^{\Delta,I,F} P\sigma_{rn}}\ [\textsc{Join}]$$

$$\vdash_{\varphi a}^{\Delta_a,I_a,F_a} b.n\langle\widetilde{v}\rangle \parallel \vdash_{\psi b}^{\Delta_b,I_b,F_b} \longrightarrow \vdash_{\varphi a}^{\Delta_a,I_a,F_a} \parallel \vdash_{\psi b}^{\Delta_b,I_b,F_b} b.n\langle\widetilde{v}\rangle\ [\textsc{Comm}]$$

$$a[\mathcal{D} : \mathcal{P} \mid \mathtt{go}\ b; Q]^{\Delta_a,I_a} \vdash_{\varphi}^{\Delta,I,F} \parallel \vdash_{\psi b}^{\Delta_b,I_b,F_b} \longrightarrow \vdash_{\varphi}^{\Delta,I,F} \parallel a[\mathcal{D} : \mathcal{P} \mid Q]^{\Delta_a,I_a} \vdash_{\psi b}^{\Delta_b,I_b,F_b}\ [\textsc{Go}]$$

*Figure 1.* Semantics of the dynamic Join Calculus

is static, or $\mathbf{n}_i$ if it is dynamic, or $y_i$ if it is a variable. A definition is triggered when among the running processes of the location there are messages on each of the $n_i$. These messages are consumed, and the guarded process is spawned, replacing the formal names (the received names) by the arguments of the messages using the substitution $\sigma_{rn}$ (as described in the JOIN rule). Note that we use a slightly different JOIN rule: since only resolved messages may be consumed, we write $a.J$ for the join pattern where every message pattern has the prefix $a$ (*i.e.* $a.(J \mid J') = a.J \mid a.J'$). New definitions are introduced using the $\mathtt{def}\ D\ \mathtt{in}\ P$ construct, where the defined static names of $D$ have scope $D$ and $P$. New dynamic channels are introduced using the $\nu\mathbf{n}.P$ construct.

Locations, either folded or running, gather in the set $\Delta$ the dynamic channels they *define*, and in the set $I$ the dynamic channels they *import* (*i.e.* the dynamic names they require to be defined in enclosing locations).

When a process $\mathtt{go}(b); P$ is evaluated, the current location as well as all its sublocations migrate to location $b$ (rule GO).

Some running location $\varphi a$ may be folded in its parent location $\varphi$ (for subsequent migration, for instance) using rule STR-LOC. The first condition of this rule insures that there is no running sublocation of $\varphi a$, in order to preserve the tree structure. Conversely, when unfolding a location, its lookup function needs to be computed using the operator *Lookup*, that takes the lookup function of the enclosing location and patches it to correspond to the current location.

**Definition 2.1** (*Lookup* **operator**) *The operator Lookup takes the lookup function $F$ of the enclosing location, the imported dynamic names $I$, the locally defined dynamic names $\Delta$, and the name of the current location $a$, to create a lookup function $G = Lookup(F, I, \Delta, a)$ that associates the current location to locally defined dynamic names and the result of the enclosing lookup function for imported names (we write $\perp$ for the undefined location).*

*We have:*

$$G(n) = \begin{cases} a & n \in \Delta \\ F(n) & n \notin \Delta \wedge n \in I \\ \perp & n \notin \Delta \wedge n \notin I \end{cases}$$

We define $\alpha$-conversion as in the Join Calculus (renaming of defined static names bound by a `def` and renaming of received names bound by join patterns), with the additional renaming of dynamic names bound by a $\nu$ operator.

In the following, we only consider a restricted class of processes: every location must have a unique name; every defined static name is defined in a single location; join patterns are linear, *i.e.* no defined name nor received name may occur more than once in a given join pattern; free and bound names are distinct ($fn(\mathcal{S}) \cap bn(\mathcal{S}) = \emptyset$). We call this last condition the *hygienic* condition.

The condition of rule DEF enforces the preservation of the hygienic condition. Since it must be true before the reduction, the defined names of $D$, that are free afterward, were bound. Thus they could not occur free in the initial configuration and the reduction cannot capture free names. The rule simply checks that these names are not bound in the final configuration, and that no defined local name of $D$ is a dynamic name (well typed definitions satisfy this property). The hygienic condition is also enforced through rule NU.

In figure 1, only rules DEF and NU explicitly mention the context. In the other rules, the other running locations, definitions, and running processes in the locations involved in the reduction are left implicit.

The structural equivalence $\equiv$ is the smallest reflexive, symmetric and transitive relation generated by rules STR-$\alpha$ and STR-LOC, with the parallel operator "|" (resp. the definition composition operator ",") being associative, commutative and having $\mathbf{0}$ (resp. $\top$) as neutral element. The reduction relation $\rightarrow$ is the smallest relation generated by rules of figure 1 such that $\equiv \rightarrow \equiv \subseteq \rightarrow$. We recall that most of these rules include implicit contexts.

Discussion and examples of the dynamic Join Calculus are available in [17].

# 3. Safe dynamic binding

We describe a type system that allows only configurations where dynamic messages are bound to a dynamic definition in some enclosing location. This type system is similar to the one for the distributed Join Calculus. It uses the same generalization criterion as the one implemented in JoCaml and formalized in [12]. We use the following types:

$$\begin{array}{rcl} \tau & ::= & \widetilde{\tau} \mid \langle\tau\rangle_w^+ \mid \langle\tau\rangle_{\mathbf{\Delta}} \mid loc(\mathbf{\Delta}) \mid \alpha \\ \mathbf{\Delta}, \mathbf{I} & ::= & \emptyset \mid \{w\} \mid \mathbf{\Delta} \cup \mathbf{\Delta} \end{array} \qquad \begin{array}{rcl} w & ::= & \mathbf{n} \mid \delta \\ \sigma & ::= & \forall \widetilde{\alpha}\widetilde{\delta}.\tau \end{array}$$

$$\frac{}{\langle\tau\rangle_w^+ \leq \langle\tau\rangle_{\{w\}}}\ [\textsc{Plus}] \qquad \frac{\tau' \leq \tau \qquad \Delta \subseteq \Delta'}{\langle\tau\rangle_\Delta \leq \langle\tau'\rangle_{\Delta'}}\ [\textsc{N-Sub}] \qquad \frac{\Delta' \subseteq \Delta}{loc(\Delta) \leq loc(\Delta')}\ [\textsc{L-Sub}]$$

*Figure 2.* Subtyping rules

Name type variables $\delta$ occur in the types of processes guarded by a join pattern, and represent a dynamic channel name (as a name variable $y$ represents a channel or location name).

Intuitively, locations provide dynamic definitions, either by directly defining them, or by requesting enclosing locations to define them. The type of a location is $loc(\Delta)$, where $\Delta$ is the set of dynamic names that are available in the location. Thus, inside such a location a message sent on a dynamic name of $\Delta$ is correct, whereas a message sent on any other dynamic name should be considered as incorrectly typed.

The type of dynamic channel names reflects the required dynamic definitions. A message on a channel of type $\langle\tau\rangle_\Delta$ carries an argument of type $\tau$, and requires the availability of definitions for the names of $\Delta$. For instance, a dynamic channel $\mathbf{n}$ that does not carry any argument has type $\langle\rangle_{\{\mathbf{n}\}}$, since a message on such a channel requires a definition for $\mathbf{n}$ to be available. Since static channels do not require the presence of any dynamic definition in enclosing locations, their type is of the form $\langle\tau\rangle_\emptyset$ which is simply written $\langle\tau\rangle$.

Since channel names are first class values, it is possible to write a definition such as $\mathbf{send}\langle x\rangle \triangleright x\langle\rangle$, that receives a name and sends a message on it. Any use of **send** with a static name is correct, and using **send** with a dynamic name is correct only if the dynamic name is defined in an enclosing location. In order to represent this behavior, we say that the type of the argument of **send** is $\langle\rangle_\Delta$ if $\Delta$ is the set of available dynamic channels. Thus **send** has the type $\langle\langle\rangle_\Delta\rangle$. As $\Delta$ represents an upper bound of the definitions that may be used, we have an immediate notion of subtyping (written $\leq$) on dynamic channels: $\langle\tau\rangle_{\Delta'} \leq \langle\tau\rangle_\Delta$ if $\Delta' \subseteq \Delta$; a channel that may access fewer dynamic definitions is a subtype of a channel that may access more dynamic definitions. Thus a static channel has a type that is a subtype of any dynamic channel carrying the same type of arguments. The subtyping rule N-Sub of figure 2 also introduces contravariant subtyping on the argument type.

Since a location that provides more dynamic definitions may be used instead of one that provides fewer dynamic definitions, we introduce a notion of subtyping on location types, in rule L-Sub.

One important condition for insuring soundness of the type system is to forbid subtyping on dynamic names that are redefined. We write $\langle\tau\rangle_w^+$ for the type of these channels. Since a redefinable channel may be used instead of a plain dynamic channel for message sending, we introduce the Plus subtyping rule. In the following we consider $\leq$ to be the smallest reflexive transitive closure generated by the rules of figure 2.

A type scheme $\forall\widetilde{\alpha}\widetilde{\delta}.\tau$ is composed of generalized type variables $\widetilde{\alpha}$, generalized name type variables $\widetilde{\delta}$, and type $\tau$. An instantiation of this type scheme, written $Inst(\forall\widetilde{\alpha}\widetilde{\delta}.\tau)$, is a type $\tau\theta$, where $\theta$ is a substitution from the type variables $\widetilde{\alpha}$ to types and from the name type variables $\widetilde{\delta}$ to dynamic name types $w$.

For soundness reasons, we do not allow polymorphic redefinable channel types. Similarly, we do not allow polymorphic location types. All other types are said to be *well formed.*

A *type environment* $B$ (resp. a *type scheme environment* $A$ or $\Gamma$) is an association map between names and types (resp. names and type schemes), where each name occurs at most once.

In the following, we call $ftv(\tau)$ the free type variables in $\tau$, $fnv(\tau)$ the free name type variables in $\tau$. We write $fv(\tau)$ for $ftv(\tau) \cup fnv(\tau)$. We also extend $ftv$ and $fnv$ to type environments and to type scheme environments.

In the following typing rules, we use a *generalization* operator $Gen(B, \Lambda, \Theta)$ defined as: $Gen(B, \Lambda, \Theta) = \bigcup_{m:\tau \in B}\{m : \forall\widetilde{\alpha}\widetilde{\delta}.\tau\}$, where $\widetilde{\alpha} = ftv(\tau) \setminus (\Lambda \cup \Theta)$ and $\widetilde{\delta} = fnv(\tau) \setminus (\Lambda \cup \Theta)$. The set $\Lambda$ contains the names and type variables that occur in the typing environment (in typing rule DEF, $\Lambda$ is $fv(\Gamma)$); the set $\Theta$ contains the names and type variables that may not be generalized because they are shared in a join pattern (as in [12]).

A typing judgment has one of the following forms: $\Gamma \Vdash \widetilde{n} : \widetilde{\tau}$, $\mathbf{\Delta}; \mathbf{I}; \Gamma \Vdash P$, $\mathbf{\Delta}; \mathbf{I}; \Gamma \Vdash \mathcal{D} :: B; \Delta_1; \Theta$, $\Gamma \Vdash \mathcal{S}$, or $\mathbf{\Delta}; \mathbf{I}; \Gamma \Vdash \mathcal{D} : \Delta_1$ where $\Gamma$ is the type scheme environment; $\mathbf{\Delta}$ is a set of dynamic name types, the dynamic channels defined in the current location; $\mathbf{I}$ is a set of dynamic names, the dynamic channels imported by the current location; $B$ gathers the types of the defined static names of $\mathcal{D}$; $\Delta_1$ collects the dynamic channels locally defined by $\mathcal{D}$; $\Theta$ is a set of type variables and name type variables that cannot be generalized.

The four main typing rules to guarantee the presence of a dynamic definition for every dynamic message are MSG, LOC, SOUP-LOC, and GO. Rule MSG checks that the channel used does not require more dynamic definitions than the ones available locally (*i.e.* the ones specified in $\mathbf{\Delta}$ and $\mathbf{I}$). Rules LOC and SOUP-LOC are very similar, the former being more complex as it can occur in the process guarded by a join pattern. These rules check that the specified dynamic channels are defined and that the imported dynamic channels are available locally (in rule CONF for the SOUP-LOC case). The contents of the location are typed in an environment where the defined and imported dynamic channels are available (these rules modify $\mathbf{\Delta}$ and $\mathbf{I}$). Rule GO checks that the target of the migration provides at least the dynamic channels imported by the current location.

The two typing rules for messages MSG and R-MSG are very similar, and follow the different states of a message as it is first resolved, then sent to its destination. In these rules, the name on which the message is sent needs to satisfy the typing judgment $\Gamma \Vdash n : \langle\widetilde{\tau}\rangle_{\mathbf{\Delta} \cup \mathbf{I}}$. Because of subtyping on channel names, this judgment gives an upper bound on the dynamic names that the message may use, thus on the definitions accessed; this upper bound consists of the available dynamic definitions at this point, thus insuring that every

$$\frac{m : \forall \tilde{\alpha}\tilde{\delta}.\tau' \in \Gamma \qquad \tau = Inst(\forall \tilde{\alpha}\tilde{\delta}.\tau')}{\Gamma \Vdash m : \tau} \text{ [NAME]} \qquad \frac{\Gamma \Vdash m : \tau \qquad \tau \leq \tau'}{\Gamma \Vdash m : \tau'} \text{ [SUB]}$$

$$\frac{\Gamma \Vdash m_i : \tau_i \text{ for } i \in [1..n]}{\Gamma \Vdash m_1, \ldots, m_n : \tau_1, \ldots, \tau_n} \text{ [TUPLE]}$$

$$\frac{\Gamma \Vdash n : \langle \tilde{\tau} \rangle_{\Delta \cup \mathbf{I}} \qquad \Gamma \Vdash \tilde{m} : \tilde{\tau} \qquad \Delta; \mathbf{I}; \Gamma \Vdash P}{\Delta; \mathbf{I}; \Gamma \Vdash n\langle \tilde{m} \rangle; P} \text{ [MSG]}$$

$$\frac{\Gamma \Vdash n : \langle \tilde{\tau} \rangle_{\Delta \cup \mathbf{I}} \qquad \Gamma \Vdash \tilde{m} : \tilde{\tau} \qquad \Gamma \Vdash a : loc(\emptyset) \qquad n \in dln(a)}{\Delta; \mathbf{I}; \Gamma \Vdash a.n\langle \tilde{m} \rangle} \text{ [R-MSG]}$$

$$\frac{\begin{array}{c}\Delta; \mathbf{I}; \Gamma + A \Vdash D :: B; \emptyset; \Theta \\ A = Gen(B, fv(\Gamma), \Theta) \qquad \Delta; \mathbf{I}; \Gamma + A \Vdash P \qquad dom(A) \cap dom(\Gamma) = \emptyset\end{array}}{\Delta; \mathbf{I}; \Gamma \Vdash \mathbf{def}\ D\ \mathbf{in}\ P} \text{ [DEF]}$$

$$\frac{\Delta; \mathbf{I}; \Gamma \Vdash \mathcal{P}_1 \qquad \Delta; \mathbf{I}; \Gamma \Vdash \mathcal{P}_2}{\Delta; \mathbf{I}; \Gamma \Vdash \mathcal{P}_1 \mid \mathcal{P}_2} \text{ [PAR]} \qquad \frac{\Gamma \Vdash n : loc(\mathbf{I}) \qquad \Delta; \mathbf{I}; \Gamma \Vdash P}{\Delta; \mathbf{I}; \Gamma \Vdash \mathbf{go}\ n; P} \text{ [GO]}$$

$$\frac{\mathbf{d} \notin fn(\Gamma) \qquad \Delta; \mathbf{I}; \Gamma + \mathbf{d} : \langle \tau \rangle_{\mathbf{d}}^{+} \Vdash P}{\Delta; \mathbf{I}; \Gamma \Vdash \nu \mathbf{d}.P} \text{ [NU]} \qquad \frac{}{\Delta; \mathbf{I}; \Gamma \Vdash \mathbf{0}} \text{ [NIL]}$$

$$\frac{\begin{array}{c}\Delta; \mathbf{I}; \Gamma + (\tilde{u_i} : \tilde{\tau_i})^i + (\tilde{y_j} : \tilde{\tau_j})^j \Vdash P \\ \Gamma \Vdash \mathbf{x_i} : \langle \tilde{\tau_i} \rangle \qquad \Gamma \Vdash m_j : \langle \tilde{\tau_j} \rangle_{w_j}^{+} \qquad (\bigcup_i \tilde{u_i} \cup \bigcup_j \tilde{y_j}) \cap dom(\Gamma) = \emptyset \\ \forall (n:\tau), (n':\tau') \in \left( \{\mathbf{x_i} : \langle \tilde{\tau_i} \rangle\} \cup \{m_j : \langle \tilde{\tau_j} \rangle_{w_j}^{+}\} \right) . n \neq n' \implies fv(\tau) \cap fv(\tau') \subseteq \Theta\end{array}}{\Delta; \mathbf{I}; \Gamma \Vdash (\mathbf{x_i}\langle \tilde{u_i} \rangle)^i \mid (m_j\langle \tilde{y_j} \rangle)^j \triangleright P :: (\mathbf{x_i} : \langle \tilde{\tau_i} \rangle)^i; \bigcup_j \{m_j\}; \Theta} \text{ [JOIN]}$$

$$\frac{}{\Delta; \mathbf{I}; \Gamma \Vdash \top :: \emptyset; \emptyset; \Theta} \text{ [TOP]} \qquad \frac{\Delta; \mathbf{I}; \Gamma \Vdash \mathcal{D}_1 :: B_1; \Delta_1; \Theta \qquad \Delta; \mathbf{I}; \Gamma \Vdash \mathcal{D}_2 :: B_2; \Delta_2; \Theta}{\Delta; \mathbf{I}; \Gamma \Vdash \mathcal{D}_1, \mathcal{D}_2 :: B_1 \oplus B_2; \Delta_1 \cup \Delta_2; \Theta} \text{ [AND]}$$

$$\frac{\begin{array}{c}\forall m_i \in \Delta\ .\ m_i : \langle \tau_i \rangle_{w_i}^{+} \in \Gamma \qquad \forall n_j \in I\ .\ n_j : \langle \tau_j \rangle_{w'_j}^{+} \in \Gamma \qquad \Delta' = \bigcup_i w_i \qquad \mathbf{I}' = \bigcup_j w'_j \\ \Delta'; \mathbf{I}'; \Gamma \Vdash \mathcal{D} :: B; \Delta; \Theta \qquad \Delta'; \mathbf{I}'; \Gamma \Vdash \mathcal{P} \qquad \Gamma \Vdash a : loc(\Delta' \cup \mathbf{I}') \qquad \mathbf{I}' \subseteq (\Delta'' \cup \mathbf{I}'')\end{array}}{\Delta''; \mathbf{I}''; \Gamma \Vdash a[\mathcal{D} : \mathcal{P}]^{\Delta, I} :: B + a : loc(\Delta' \cup \mathbf{I}'); \emptyset; \Theta} \text{ [LOC]}$$

$$\frac{\begin{array}{c}\{\varphi_i\} \text{ form a tree with root } b \qquad \forall \psi a \in \{\varphi_i\}.I_{\psi a} \subseteq \Delta_\psi \cup I_\psi \qquad I_b = \emptyset \\ \forall \psi a \in \{\varphi_i\}.F_{\psi a} = Lookup(F_\psi, I_{\psi a}, \Delta_{\psi a}, a) \qquad (\Gamma \Vdash \mathcal{D}_i \vdash_{\varphi_i}^{\Delta_{\varphi_i}, I_{\varphi_i}, F_{\varphi_i}} \mathcal{P}_i)^i\end{array}}{\Gamma \Vdash \prod_i (\mathcal{D}_i \vdash_{\varphi_i}^{\Delta_{\varphi_i}, I_{\varphi_i}, F_{\varphi_i}} \mathcal{P}_i)} \text{ [CONF]}$$

$$\frac{\begin{array}{c}n \in \Delta \cup I \implies n = \mathbf{n} \wedge \mathbf{n} : \langle \tau \rangle_{\mathbf{n}}^{+} \in \Gamma \\ a : loc(\Delta \cup I) \in \Gamma \qquad \Delta; I; \Gamma \Vdash \mathcal{P} \qquad \Delta; I; \Gamma \Vdash \mathcal{D} : \Delta\end{array}}{\Gamma \Vdash \mathcal{D} \vdash_{\varphi a}^{\Delta, I, F} \mathcal{P}} \text{ [SOUP-LOC]}$$

$$\frac{\Delta; \mathbf{I}; \Gamma \Vdash \mathcal{D} :: B; \Delta_{\mathcal{D}}; \Theta \qquad A = Gen(B, fv(\Gamma), \Theta) \qquad A \subseteq \Gamma}{\Delta; \mathbf{I}; \Gamma \Vdash \mathcal{D} : \Delta_{\mathcal{D}}} \text{ [CHEM-DEF]}$$

*Figure 3.* Typing rules

dynamic message is bound to a definition. The R-MSG rule checks that the location specified in the prefix of the message is present with a location type in $\Gamma$ (every location type is a subtype of $loc(\emptyset)$), and that the channel name is a defined local name of this location.

The typing rule DEF checks that no local dynamic name is defined in $D$ ($\Delta_1 = \emptyset$), and also checks that the defined static names of $D$—which are the domain of $B$ and $A$—do not clash with names of the typing environment. We remark that we use polymorphic recursion here, which drastically simplifies the subject reduction proof for the distributed Join Calculus.

Rule GO requires that the destination location has type $loc(\mathbf{I})$, where $\mathbf{I}$ is the set of imported dynamic channels. By definition of subtyping on location, this set is a lower bound, and any location providing more dynamic definitions may be the target of migration.

Rule NU introduces a new dynamic channel, which is monomorphic, and which has the type of dynamic channels sending messages on their own name. Since every dynamic channel created has a monomorphic redefinable type, every subsequent definition must exactly follow this type.

Rule JOIN is used to type one join pattern. The defined names of the join pattern are partitioned into two sets: the static names and the dynamic names. Static names $\mathbf{x_i}$ are given the type $\langle \widetilde{\tau_i} \rangle$, and are collected in the typing environment $B$. Dynamic names need to be present in $\Gamma$, with a redefinable type $\langle \widetilde{\tau_j} \rangle^{+}_{w_j}$, as they are redefined. Dynamic names are not written $\mathbf{m}_j$ since they may be variables bound by an enclosing join pattern. In evaluation context, we always have $m_j = w_j = \mathbf{m}_j$. They may however be different if the join pattern occurs in the guarded process of another join pattern that receives a dynamic name and then redefines it. All dynamic names are collected in the set of local dynamic names. The set of non generalized variables $\Theta$ is checked to be big enough: any type variable or name that is shared between two types cannot be generalized (as in [12]).

The typing rule AND uses the $\oplus$ operator in $B_1 \oplus B_2$, that requires names that are both in the domain of $B_1$ and $B_2$ to have the same type.

The rule LOC extracts from the typing environment the types associated to the names of $\Delta$ and $I$, in order to collect the dynamic name types associated to these channels. As in rule JOIN, in evaluation context the names are the same. They may be different if the typing rule is used to type the guarded process of a join pattern that receives a dynamic name that is redefined. The typing of the definition $\mathcal{D}$ must yield a set of local dynamic names $\Delta$ identical to the one declared in the location. The typing of $\mathcal{D}$ and $\mathcal{P}$ may use the available dynamic channels associated to the ones declared by the location being typed. This rule also checks that the imported names are available in the enclosing location.

The CONF rule checks that the running locations form a tree, all running locations import names that are available in the enclosing location (the root location does not import any name), every name lookup function is correctly computed, and all running locations are well typed. To do this, the SOUP-LOC rule is used. It first checks that all the names declared in $\Delta$ and $I$ are

dynamic channels (not variables), and that the type of these channels in $\Gamma$ is the type of redefinable channels sending messages on their own name. It then types the definition (using rule CHEM-DEF) and process of the location, using the available dynamic channels declared by the location, and checks that the location is present with the correct type in the typing environment. This rule is similar to rule LOC, although simpler as running locations occur only in evaluation contexts. The CHEM-DEF rule types the definition and checks that the resulting generalization is present in the environment.

Typing examples are available in [17].

# 4. Type Soundness

To prove the soundness of our system, we first prove a subject reduction theorem, then we prove a progress property that insures that well-typed configurations do not go wrong. We first specify *well-formed* typing environments.

**Definition 4.1** *A typing environment $\Gamma$ is* well formed *if and only if its types are well formed and every type binding is of the form $\mathbf{n} : \langle\tau\rangle^{+}_{\mathbf{n}}$, $\mathrm{n} : \forall\widetilde{\alpha}\widetilde{\delta}.\langle\tau\rangle$, and $a : loc(\Delta)$ where $\Delta$ contains no name type variable.*

We now prove that structural equivalence and reductions preserve typing.

**Lemma 4.2** *Let $\mathcal{S}$ be a configuration and $\Gamma \Vdash \mathcal{S}$ a typing of this configuration where $\Gamma$ is well-formed. If $\mathcal{S} \equiv \mathcal{S}'$, there is a well-formed $\Gamma'$ such that $\Gamma' \Vdash \mathcal{S}'$.*

**Theorem 1 (Subject reduction)** *Let $\mathcal{S}$ be a configuration and $\Gamma$ be a well formed environment. If $\Gamma \Vdash \mathcal{S}$ and $\mathcal{S} \to \mathcal{S}'$, then there exists a well-formed environment $\Gamma'$ such that $\Gamma' \Vdash \mathcal{S}'$.*

We remark that $\Gamma$ and $\Gamma'$ need not be related, since reduction steps involve configurations including implicit contexts. We now prove that the NL-DYN step resolves messages to the closest enclosing location defining them (where $dyn(b)$ is the set of dynamic names locally defined in $b$). This proof insures that our way of computing the name lookup function corresponds to our specification.

**Lemma 4.3** *Let $\Gamma \Vdash \mathcal{S}$ be a typing derivation. For any dissolved location $\varphi a$ of $\mathcal{S}$, if we have $F_{\varphi a}(n) = b \neq \bot$, then we have $n \in dyn(b)$, $\varphi a = \psi b \psi'$ with $\forall c \in \psi'.n \notin dln(c)$.*

We define the notion of a *stuck configuration.*

**Definition 4.4** *We say that a configuration $\mathcal{S}$ is* stuck *when one of the following is true: there is a message $n\langle\widetilde{m}\rangle$ or $a.n\langle\widetilde{m}\rangle$ in evaluation context where $n$ is not a channel name; there is a $n \in dn(\mathcal{S})$ that is not a channel name or a location name; there is a process* go $n; P$ *where $n$ is not a location name; there is a message $n\langle\widetilde{m}\rangle$ or $a.n\langle\widetilde{m}\rangle$ and a definition of $n$ with different arities; there is a message $\mathbf{n}\langle\widetilde{m}\rangle$ in evaluation context that cannot be reduced by rule* NL-DYN*; there is a message $a.n\langle\widetilde{m}\rangle$ with $n \notin dln(a)$.*

We can now state that no well-typed configuration is stuck.

**Theorem 2 (Progress)** *Let $\mathcal{S}$ be a configuration and $\Gamma$ a well-formed environment. If $\Gamma \Vdash \mathcal{S}$, then $\mathcal{S}$ is not stuck.*

Combining lemma 4.2, 4.3, and theorems 1 and 2, we prove that well-typed configurations cannot become stuck, thus dynamic messages are always sent to the closest enclosing location providing a definition and there is always such a location. The proofs are fairly complex, especially the substitution lemma, as channel names occur in the types. They are available in [16].

## 5. Future work, related work, and conclusion

One of our first priorities is the integration of dynamic channels into JoCaml. A form of dynamic binding at the module level is already present in JoCaml, but is difficult to use. The implementation would provide the programmers with easier management of local resources. As our system was designed with implementation in mind, this should not present any difficulty. In particular, it seems clear that maintaining the lookup function $F$ (that resolves dynamic names to the closest enclosing location defining them) would be cheap because it does not require more locking than the one already present in JoCaml.

On the theoretical side, it is unfortunate that the construct $a.n\langle\widetilde{m}\rangle$, very similar to the message construct of [19], cannot be made available to programmers. To use such a construct soundly, it is necessary to check that $n$ is indeed defined in location $a$, where $a$ may be a variable bound by an enclosing join pattern. We have been designing such a type system, which is very similar to the one introduced in this paper, and we are finishing the soundness proofs.

Another extension is the design of a type reconstruction algorithm. Since our type system uses polymorphic recursion, which greatly simplifies the proof of subject reduction, a type reconstruction algorithm requires a different type system (see [17]). This issue is already present in the distributed Join Calculus. We are also modifying our system to use a constraint based type system to recover principal types, adapting the B($T$) framework of [8] to this end.

In this paper, our strategy for handling a message on a dynamic channel name is to resolve in one step the closest enclosing location defining the channel, and then to send the message to this location in an other step. An alternate approach would simply consist in sending a message on a dynamic name that is not defined locally to the parent location, which would then deal with the message (this second incremental approach is similar to [6] and [18]). The two semantics yield two different behaviors but, surprisingly enough, the type system presented in this paper is also sound for the second system.

Part of this work could have been achieved in a distributed $\pi$ Calculus. However, this would have been more subtle since the association between senders and receivers is more dynamic than in the Join Calculus, since receivers may disappear. It is therefore more complex to distinguish between deadlock freedom [1, 14] and availability of receivers. This implies that our type system is simpler as we do not guarantee deadlock freedom because of join patterns.

However, the resulting type system of [1] requires that names passed on channels cannot be used as receivers, unlike our dynamic channels as shown in [17].

Several works deal with resource control (as opposed to resource availability) in a $\pi$-calculus with localities and objective rebinding. In [13], localities have a flat structure, and processes may migrate objectively between localities. A type system insures that no agent may access a resource if it was not given the capability to do so. In the *local area calculus* [7], localities form a fixed hierarchy of levels that do not migrate. Channels have a level of operation, meaning that no communication on such a channel may cross the boundary of an higher level area. In the box-$\pi$ calculus [18], localities also form a fixed hierarchy, and communication may cross only one locality boundary at a time. This calculus aims at controlling the flow of information between localities.

The higher order $\pi$-calculus of [21, 20] also deals with access control, by explicitly specifying for each input which resources may be accessed by the input process, distinguishing between read and write accesses. However, this calculus is objective, guarantees locality of resources instead of availability, and uses dependent types instead of polymorphism with name type variables.

The work on secrecy and groups [5] also deals with controlling access to some names through the creation of fresh groups, and the assignment of channels to these groups. However, the different intent (secrecy vs receptiveness) leads to different type systems. We use polymorphism to let dynamic names escape the scope in which they are created, to type more processes, whereas in [5] secrecy is achieved by checking that groups cannot escape their initial scope.

A very simple calculus of localized resources is the Ambient Calculus [6]. Recent works on the Ambient Calculus add type systems to analyze the behavior of ambients or enforce security policies. In [4], groups are introduced to refine mobility types of ambients in order to control the escape of capabilities. In [2], safe ambients are typed according to a security policy. The type system takes into account the capabilities that the ambient may acquire. Boxed ambients [3] drops the *open* capabilities of the Ambient Calculus, but allows communication between parent and children. A type system allows the analysis of ambient behavior, distinguishing local communications from communications with the context. In these ambient calculi, the emphasis is more on restricting the behavior of processes than on checking the availability of resources. Moreover, there is no notion of static names.

We have presented an extension of the distributed Join Calculus that features the creation of definitions accessible through dynamic channels that are rebound at migration. New definitions for a given channel may be created at runtime, and new dynamic channels may be generated as well. A type system that has the subject reduction property was presented, which guarantees the presence of a definitions for any dynamic channel that may be used.

## Acknowledgments

We would like to thank Sylvain Conchon, Cédric Fournet, James Leifer, Jean-Jacques Lévy, François Pottier, and Didier Rémy for comments.

# References


[1] R. M. Amadio, G. Boudol, and C. Lhoussaine. The receptive distributed pi-calculus (extended abstract). In *FST-TCS'99*, LNCS, 1999.

[2] M. Bugliesi and G. Castagna. Secure safe ambients. In *Proceedings of POPL '01*, pages 222–235. ACM Press, 2001.

[3] M. Bugliesi, G. Castagna, and S. Crafa. Boxed ambients. In *TACS'01*, LNCS, 2001.

[4] L. Cardelli, G. Ghelli, and A. D. Gordon. Ambient groups and mobility types. In *IFIP TCS 2000 (Sendai, Japan)*, LNCS. IFIP, Springer, Aug. 2000.

[5] L. Cardelli, G. Ghelli, and A. D. Gordon. Secrecy and group creation. In *CONCUR 2000 (University Park, PA, USA)*, LNCS. Springer, Aug. 2000.

[6] L. Cardelli and A. D. Gordon. Mobile ambients. In *Foundations of Software Science and Computational Structures*. Springer (LNCS), 1998.

[7] T. Chothia and I. Stark. A distributed calculus with local areas of communication. In *Proceedings of HLCL '00*, 2001.

[8] S. Conchon and F. Pottier. JOIN(X): Constraint-Based Type Inference for the Join-Calculus. In *Proceedings of ESOP'01*, LNCS, Apr. 2001.

[9] C. Fournet. *The Join-Calculus: a Calculus for Distributed Mobile Programming*. PhD thesis, Ecole Polytechnique, Palaiseau, Nov. 1998. INRIA, TU-0556.

[10] C. Fournet, G. Gonthier, J.-J. Lévy, L. Maranget, and D. Rémy. A calculus of mobile agents. In *Proceedings of CONCUR'96*, Aug. 1996. LNCS 1119.

[11] C. Fournet, J.-J. Lévy, and A. Schmitt. A distributed implementation of Ambients. Draft of long version, `http://join.inria.fr/ambients.html`, 1999.

[12] C. Fournet, L. Maranget, C. Laneve, and D. Rémy. Inheritance in the join calculus. In *FST-TCS'00*, LNCS, Dec. 2000.

[13] M. Hennessy and J. Riely. Resource access control in systems of mobile agents. *Information and Computation*, To appear.

[14] N. Kobayashi, S. Saito, and E. Sumii. An implicitly-typed deadlock-free process calculus. In *Proceedings of CONCUR 2000*, LNCS, Aug. 2000.

[15] F. Le Fessant. The JoCAML system prototype. Software and documentation available from `http://pauillac.inria.fr/jocaml`, 1998.

[16] A. Schmitt. Safe Dynamic Binding in the Join Calculus, Draft. Available from `http://pauillac.inria.fr/~aschmitt/publications.html`, 2001.

[17] A. Schmitt. Safe dynamic binding in the join calculus. Version of this paper with discussions and examples, available from `http://pauillac.inria.fr/~aschmitt/publications.html`, 2002.

[18] P. Sewell and J. Vitek. Secure composition of insecure components. In *Proceedings of CSFW 99 (Mordano, Italy)*, June 1999.

[19] A. Unyapoth and P. Sewell. Nomadic Pict: Correct communication infrastructure for mobile computation. In *Proceedings of POPL 2001 (London)*, Jan. 2001.

[20] N. Yoshida and M. Hennessy. Subtyping and locality in distributed higher-order processes. In *Proceedings CONCUR 99, LNCS no 1664*, 1999.

[21] N. Yoshida and M. Hennessy. Assigning types to processes (extended abstract). In *Fifteenth Annual IEEE Symposium on Logic in Computer Science*, 2000.