

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

O'Klaim: a coordination language with mobile mixins

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/28796> since 2015-10-12T09:30:10Z

Publisher:

Springer

Published version:

DOI:10.1007/978-3-540-24634-3_5

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

O'Klaim: a coordination language with mobile mixins

Lorenzo Bettini¹ Viviana Bono² Betti Venneri¹

¹Dipartimento di Sistemi e Informatica, Università di Firenze,
`{bettini,venneri}@dsi.unifi.it`

²Dipartimento di Informatica, Università di Torino, `bono@di.unito.it`

Abstract. This paper presents O'KLAIM (Object-Oriented KLAIM), a linguistic extension of the higher-order calculus for mobile processes KLAIM with object-oriented features. Processes interact by an asynchronous communication model: they can distribute and retrieve resources, sometimes structured as incomplete classes, i.e., mixins, to and from distributed tuple spaces. This mechanism is coordinated by providing a subtyping relation on classes and mixins, which become polymorphic items during communication. We propose a static typing system for: (i) checking locally each process in its own locality; (ii) decorating object-oriented code that is sent to remote sites with its type. This way, tuples can be dynamically retrieved only if they match by subtyping with the expected type. If this pattern matching succeeds, the so retrieved code can be composed with local code, dynamically and automatically, in a type-safe way. Thus a global safety condition is guaranteed without requiring any additional information on the local reconfiguration of local and foreign code, and, in particular, without any further type checking. Finally, we present main issues concerning the implementation of O'KLAIM.

1 Introduction

Mixins [14, 22, 2] are (sub)class definitions parameterized over a superclass and were introduced as an alternative to standard class inheritance. A mixin could be seen as a function that, given one class as an argument, produces another class, by adding or overriding specific sets of methods. The same mixin can be used to produce a variety of classes with the same functionality and behavior, since they all have the same sets of methods added and/or redefined. The superclass definition is not needed at the time of writing the mixin definition, thus improving modularity. The uniform extension and modification of classes is instead absent from the classical class-based languages. Moreover, in class-based languages, parentage is determined statically at compile time instead of at run-time.

Due to their dynamic nature, mixin inheritance can be fruitfully used in a *mobile code* setting [28, 17]. In [7], we introduced MoMi (Mobile Mixins), a coordination model for mobile processes that exchange object-oriented code. The

underlying idea motivating MOMI is that standard class-based inheritance mechanisms, which are often used to implement distributed systems, do not appear to scale well to a distributed context with mobility. MOMI's approach consists in structuring mobile object-oriented code by using mixin-based inheritance, and this is shown to fit into the dynamic and open nature of a mobile code scenario. For example, a downloaded mixin, describing a mobile agent that has to access some files, can be completed with a base class in order to provide access methods that are specific of the local file system. Conversely, critical operations of a mobile agent, enclosed in a downloaded class, can be redefined by applying a local mixin to it (e.g., in order to restrict the access to sensible resources, as in a *sand-box*). Therefore, MOMI is a combination of a core coordination calculus and an object-oriented mixin-based calculus.

MOMI highly relies on typing. The most important feature of MOMI's typing is the *subtyping* relation that guarantees safe, yet flexible, code communication. We assume that the code that is sent around has been successfully compiled in its own site (independently from the other sites), and it travels together with its static type. When the code is received on a site (whose code has been successfully compiled, too), it is accepted only if its type is compliant with respect to the one expected, where compliance is based on subtyping. Thus, dynamic type checking is performed only at communication time. This is a crucial matter for mobility, since mobile code and in particular mobile agents are expected to be autonomous: once the communication successfully occurred, transmitted code behaves remotely in a (type) safe way (no failure communication to the sender will ever be required). This makes the code exchange an *atomic* action.

This paper presents the experimental language O'KLAIM that is obtained by applying MOMI's approach [7] to the language KLAIM [18,4], which is specifically designed to program distributed systems where mobile components interact through multiple distributed tuple spaces and mobile code. A preliminary design that led to O'KLAIM was introduced in [6]. KLAIM offers a much more sophisticated, complete, and effective coordination mechanism of mobile processes than the toy language of MOMI, where the focus was mainly on the subtyping relation on classes and mixins. O'KLAIM integrates the mixin-based object-oriented features into a mobile code calculus with an asynchronous coordination mechanism. To this aim, the notion of "tuple" is extended to include object-oriented code, therefore KLAIM processes can retrieve from and insert into tuple spaces object-oriented components (in particular, classes, mixins and objects) as well as standard KLAIM processes. A type system is designed for checking statically the extended notion of processes, so that compiled processes contain some static type information which is used dynamically. Indeed, the tuples that are added to a tuple space are decorated with their type. Therefore, a process willing to retrieve a tuple from a tuple space will employ an extended pattern matching mechanism that uses also this tuple type information. This matching essentially consists in checking subtyping on object-oriented components. If the code is successfully accepted, it can interact with the local code in a safe way (i.e., no run-time errors) without requiring any further type checking of the whole code.

Type safety of the communication results from the static type soundness of local and foreign code and a (global) subject reduction property. In particular, we show that the subject reduction property is based on a crucial property of substitutivity by subtyping. The underlying substitution operation requires specific methods renaming, in order to avoid problems that arise when classes and mixins are used as first-class data in a mobile code setting where matching relies on subtyping. These new metatheoretical results about the precise concept of substitution to be used extend and improve the results presented for MOMI in [7].

Summarizing, O’KLAIM aims at two complementary goals. Firstly, subtyping on classes and mixins (as designed for MOMI) is successfully experimented as a coordination mechanism within a mobile code calculus with a more sophisticated communication mechanism. Secondly, the language KLAIM is enriched with object-oriented code. This casts some light on how the same approach can be fruitfully used for extending many other mobile code languages with safe object-oriented code exchange. Finally, the implementation of O’KLAIM is presented. This consists in a Java package, `momi`, providing the run-time systems for classes and mixins that can be dynamically manipulated and composed. The programming language X-KLAIM (that implements the basic concepts of KLAIM) has been extended in order to be compiled into Java code exploiting the `momi` package.

2 O’Klaim: an object-oriented Klaim

O’KLAIM is a linguistic integration of KLAIM with object-oriented features, following the design of MOMI [7]. The coordination part and the object-oriented part are orthogonal, so that, in principle, such an integration would work for any extension/restriction of KLAIM (as discussed in [4]) and also for other calculi for mobility and distribution, such as *DJoin* [23]. We first recall the main features of KLAIM and MOMI and then we show how they are integrated in order to build O’KLAIM.

2.1 The basics of Klaim

KLAIM (*Kernel Language for Agent Interaction and Mobility*) [18, 4] is a coordination language inspired by the Linda model [24], hence it relies on the concept of *tuple space*. A tuple space is a multiset of *tuples*; these are sequences of information items (called *fields*). There are two kinds of fields: *actual fields* (i.e., expressions, processes, localities, constants, identifiers) and *formal fields* (i.e., variables). Syntactically, a formal field is denoted with *!ide*, where *ide* is an identifier. Tuples are anonymous and content-addressable; *pattern-matching* is used to select tuples in a tuple space:

- two tuples match if they have the same number of fields and corresponding fields have matching values or formals;
- formal fields match any value of the same type, but two formals never match, and two actual fields match only if they are identical.

For instance, tuple ("foo", "bar", 300) matches with ("foo", "bar", !val). After matching, the variable of a formal field gets the value of the matched field: in the previous example, after matching, *val* (an integer variable) will contain the value 300.

Tuple spaces are placed on *nodes* (or *sites*), which are part of a *net*. Each node contains a single tuple space and processes in execution, and can be accessed through its *locality*. The distinction between logical and physical locality (and thus the concept of “allocation environment”), the creation of new nodes and process definitions are not relevant in the O’KLAIM context, thus, for the sake of simplicity, we omit them in the present formal presentation. Notice, however, that their integration, being orthogonal, is completely smooth.

KLAIM processes may run concurrently, both at the same node or at different nodes, and can perform five basic operations over nodes. **in**(*t*)@*ℓ* looks for tuple *t'* that matches with *t* in the tuple space located at *ℓ*. Whenever the matching tuple *t'* is found, it is removed from the tuple space. The corresponding values of *t'* are then assigned to the formal fields of *t* and the operation terminates. If no matching tuple is found, the operation is suspended until one is available. **read**(*t*)@*ℓ* differs from **in**(*t*)@*ℓ* only because the tuple *t'*, selected by pattern-matching, is not removed from the tuple space located at *ℓ*. **out**(*t*)@*ℓ* adds the tuple *t* to the tuple space located at *ℓ*. **eval**(*P*)@*ℓ* spawns process *P* for execution at node *ℓ*.

KLAIM is higher-order in that processes can be exchanged as primary class data. While **eval**(*P*)@*ℓ* spawns a process for (remote) evaluation at *ℓ*, processes sent with an **out** must be retrieved explicitly at the destination site. The receiver can then execute the received process locally, as in the following process: **in**(!X)@self.**eval**(X)@self.

2.2 MoMi and O’Klaim

MoMi was introduced in [7], where mixin inheritance is shown to be more flexible than standard class inheritance to fit into the dynamic nature of a mobile code scenario. The key rôle in MoMi’s typing is played by a *subtyping* relation that guarantees safe, yet flexible and scalable, code communication. MoMi’s subtyping involves not only object subtyping, but also a form of class subtyping and mixin subtyping: therefore, subtyping hierarchies are provided along with the inheritance hierarchies. It is important to notice that we are not violating the design rule of keeping inheritance and subtyping separated, since mixin and class subtyping plays a pivotal role only during the communication, when classes and mixins become genuine run-time polymorphic values.

In synthesis, MoMi consists of:

1. the definition of an object-oriented “surface calculus” containing essential features that an object-oriented language must have to write mixin-based code;
2. the definition of a new subtyping relation on class and mixin types to be exploited dynamically at communication time;

$exp ::= v$	(value)
$\text{new } exp$	(object creation)
$exp \Leftarrow m$	(method call)
$exp_1 \diamond exp_2$	(mixin appl.)
$v ::= \{m_i = b_i \mid i \in I\}$	(record)
x	(variable)
$\text{class } [m_i = b_i \mid i \in I] \text{ end}$	(class def)
mixin	
$\text{expect}[m_i : \tau_{m_i} \mid i \in I]$	
$\text{redef}[m_k : \tau_{m_k} \text{ with } b_k \mid k \in K]$	(mixin def)
$\text{def}[m_j = b_j \mid j \in J]$	
end	

Table 1. Syntax of object-oriented terms.

3. a very primitive coordination language consisting in a synchronous send/receive mechanism, to study the communication of the mixin-based code among different site.

O'KLAIM integrates the object-oriented component of MOMI, in particular the subtyping relation on classes and mixins (both described in the next section), within KLAIM, which offers a much more sophisticated, complete, and effective coordination mechanism than the toy one of MOMI.

2.3 O'Klaim: object-oriented expressions

In this section we present the object-oriented part of O'KLAIM, which is defined as a class-based object-oriented language supporting mixin-based class hierarchies via *mixin definition* and *mixin application*. It is important to notice that specific incarnations of most object-oriented notions (such as, e.g., functional or imperative nature of method bodies, object references, cloning, etc.) are irrelevant in this context, where the emphasis is on the structure of the object-oriented mobile code. Hence, we work here with a basic syntax of the kernel object-oriented calculus.

Object-oriented expressions offer object instantiation, method call and mixin application; \diamond denotes the mixin application operator. An object-oriented value, to which an expression reduces, is either an object, which is a (recursive) record $\{m_i = b_i \mid i \in I\}$, or a class definition, or a mixin definition, where $[m_i = b_i \mid i \in I]$ denotes a sequence of method definitions, $[m_k : \tau_{m_k} \text{ with } b_k \mid k \in K]$ denotes a sequence of method re-definitions, and I, J and K are sets of indexes. Method bodies, denoted here with b (possibly with subscripts), are closed terms/programs and we ignore their actual structure. A mixin can be seen as an abstract class that is parameterized over a (super)class. Let us describe informally the mixin use through a tutorial example:

M = mixin	C = class	
expect $[n : \tau]$	$[n = \dots$	(new (M \diamond C)) $\Leftarrow m_1()$
redef $[m_2 : \tau_2 \text{ with } \dots \text{ next } \dots]$	$m_2 = \dots]$	
def $[m_1 = \dots n() \dots]$	end	
end		

$\tau ::= \Sigma \mid \iota \mid \tau_1 \rightarrow \tau_2$
$\mid \text{class}\langle \Sigma \rangle$
$\mid \text{mixin}\langle \Sigma_{new}, \Sigma_{red}, \Sigma_{exp} \rangle$
$\Sigma ::= \{m_i : \tau_{m_i} \mid i \in I\}$

Table 2. Syntax of types.

$\frac{}{\Gamma, x : \tau \vdash x : \tau} (proj)$	$\frac{\Gamma \Vdash \{m_i = b_i \mid i \in I\} : \{m_i : \tau_{m_i} \mid i \in I\}}{\Gamma \vdash \{m_i = b_i \mid i \in I\} : \{m_i : \tau_{m_i} \mid i \in I\}} (rec)$
$\frac{\Gamma \vdash \{m_i = b_i \mid i \in I\} : \{m_i : \tau_{m_i} \mid i \in I\}}{\Gamma \vdash \text{class } [m_i = b_i \mid i \in I] \text{ end} : \text{class}\langle \{m_i : \tau_{m_i} \mid i \in I\} \rangle} (class)$	
$\frac{\begin{array}{l} \Gamma, \bigcup_{i \in I} m_i : \tau_{m_i}, \bigcup_{k \in K} m_k : \tau_{m_k} \vdash \{m_j = b_j \mid j \in J\} : \{m_j : \tau_{m_j} \mid j \in J\} \\ \Gamma, \bigcup_{i \in I} m_i : \tau_{m_i}, \bigcup_{k \in K} m_k : \tau_{m_k}, \bigcup_{j \in J} m_j : \tau_{m_j}, \text{next} : \tau_{m_r} \Vdash b_r : \tau'_{m_r} \quad \tau_{m_r} <: \tau_{m_r} \quad \forall r \in K \\ \text{Subj}(\Sigma_{new}) \cap \text{Subj}(\Sigma_{exp}) = \emptyset \quad \text{Subj}(\Sigma_{new}) \cap \text{Subj}(\Sigma_{red}) = \emptyset \\ \text{Subj}(\Sigma_{red}) \cap \text{Subj}(\Sigma_{exp}) = \emptyset \end{array}}{\begin{array}{l} \text{mixin} \\ \text{expect}[m_i : \tau_{m_i} \mid i \in I] \\ \Gamma \vdash \text{redef}[m_k : \tau_{m_k} \text{ with } b_k \mid k \in K] : \text{mixin}\langle \Sigma_{new}, \Sigma_{red}, \Sigma_{exp} \rangle \\ \text{def}[m_j = b_j \mid j \in J] \\ \text{end} \end{array}} (mixin)$	
$\text{where } \Sigma_{new} = \{m_j : \tau_{m_j} \mid j \in J\}, \Sigma_{red} = \{m_k : \tau_{m_k} \mid k \in K\}, \Sigma_{exp} = \{m_i : \tau_{m_i} \mid i \in I\}$	

Table 3. Typing rules for object-oriented values

Each mixin consists of three parts:

1. methods *defined* in the mixins, like m_1 ;
2. *expected methods*, like n , that must be provided by the superclass;
3. *redefined methods*, like m_2 , where *next* can be used to access the implementation of m_2 in the superclass.

The application $M \diamond C$ constructs a class, which is a subclass of C .

The typing for the object-oriented code essentially refines the typing rules sketched in [7]. The set \mathcal{T} of types is defined in Table 2. ι is a basic type and \rightarrow is the functional type operator. Σ (possibly with a subscript) denotes a record type of the form $\{m_i : \tau_{m_i} \mid i \in I\}$. if $m_i : \tau_{m_i} \in \Sigma$ we say that the *subject* m_i *occurs* in Σ (with type τ_{m_i}). $\text{Subj}(\Sigma)$ is the set of the *subjects* of Σ and $\text{Meth}(\Sigma)$ is the set of all the method names occurring in Σ (e.g., if $\Sigma = \{m : \{n : \tau\}\}$ then $\text{Subj}(\Sigma) = \{m\}$ while $\text{Meth}(\Sigma) = \{m, n\}$). As we left method bodies unspecified (see above), we must assume that there is an underlying system to type method bodies and records. We will denote this typing with \Vdash . Rules for \Vdash are obviously not specified, but \Vdash -statements are used as assumptions in other typing rules. The typing rules for values are in Table 3.

Mixin types, in particular, encode the following information:

1. record types Σ_{new} and Σ_{red} contain the types of the mixin methods (new and redefined, respectively);

$\frac{\Gamma \vdash \text{exp} : \{m_i : \tau_{m_i}^{i \in I}\} \quad j \in I}{\Gamma \vdash \text{exp} \Leftarrow m_j : \tau_{m_j}} \text{ (lookup)}$	$\frac{\Gamma \vdash \text{exp} : \text{class}\langle \{m_i : \tau_{m_i}^{i \in I}\} \rangle}{\Gamma \vdash \text{new exp} : \{m_i : \tau_{m_i}^{i \in I}\}} \text{ (new)}$
$\frac{\begin{array}{l} \Gamma \vdash \text{exp}_1 : \text{mixin}\langle \Sigma_{\text{new}}, \Sigma_{\text{red}}, \Sigma_{\text{exp}} \rangle \\ \Gamma \vdash \text{exp}_2 : \text{class}\langle \Sigma_b \rangle \\ \Sigma_b <: (\Sigma_{\text{exp}} \cup \Sigma_{\text{red}}) \\ \text{Meth}(\Sigma_b) \cap \text{Meth}(\Sigma_{\text{new}}) = \emptyset \end{array}}{\Gamma \vdash \text{exp}_1 \diamond \text{exp}_2 : \text{class}\langle \Sigma_b \cup \Sigma_{\text{new}} \rangle} \text{ (mixin app)}$	

Table 4. Typing rules for object-oriented expressions.

2. record type Σ_{exp} contains the *expected* types, i.e., the types of the methods expected to be supported by the superclass;
3. well typed mixins are well formed, in the sense that name clashes among the different families of methods are absent (the last three clauses of the (*mixin*) rule).

The typing rules for expressions are in Table 4.

Rule (*mixin app*) relies strongly on a subtyping relation $<:$. The subtyping relation rules depend obviously on the nature of the object-oriented language we choose, but an essential constraint is that it must contain the *subtyping-in-width* rule for record types: $\Sigma_2 \subseteq \Sigma_1 \Rightarrow \Sigma_1 <: \Sigma_2$.

We consider $m : \tau_1$ and $m : \tau_2$ as distinct elements, and $\Sigma_1 \cup \Sigma_2$ is the standard record union. Σ_1 and Σ_2 are considered *equivalent*, denoted by $\Sigma_1 = \Sigma_2$, if they differ only for the order of their pairs $m_i : \tau_{m_i}$.

In the rule (*mixin app*), Σ_b contains the type signatures of all methods supported by the superclass to which the mixin is applied. The premises of the rule (*mixin app*) are as follows:

- i) $\Sigma_b <: (\Sigma_{\text{exp}} \cup \Sigma_{\text{red}})$ requires that the superclass provides all the methods that the mixin expects and redefines.
- ii) $\text{Meth}(\Sigma_b) \cap \text{Meth}(\Sigma_{\text{new}}) = \emptyset$ guarantees that name clashes cannot occur during the mixin application.

Notice that the superclass may have more methods than those required by the mixin constraints. Thus, the type of the mixin application expression is a class type containing both the signatures of all the methods supplied by the superclass (Σ_b) and those of the new methods defined by the mixin (Σ_{new}).

Now the key point is the introduction of a novel subtyping relation, denoted by \sqsubseteq , defined on class and mixin types. This subtyping relation is used to match dynamically the actual parameter's types against the formal parameter's types during communication. The part of the operational semantics of O'KLAIM, which describes communication formally, is presented in Section 2.6. The subtyping relation \sqsubseteq is defined in Table 5; rule (\sqsubseteq *class*) is naturally induced by the (width) subtyping on record types, while rule (\sqsubseteq *mixin*): permits the subtype to define more 'new' methods; prohibits to override more methods; and enables a subtype to require less expected methods.

$\frac{\Sigma' <: \Sigma}{\text{class}\langle \Sigma' \rangle \sqsubseteq \text{class}\langle \Sigma \rangle} (\sqsubseteq \text{ class})$	
$\frac{\Sigma'_{\text{new}} <: \Sigma_{\text{new}} \quad \Sigma_{\text{exp}} <: \Sigma'_{\text{exp}} \quad \Sigma'_{\text{red}} = \Sigma_{\text{red}}}{\text{mixin}\langle \Sigma'_{\text{new}}, \Sigma'_{\text{red}}, \Sigma'_{\text{exp}} \rangle \sqsubseteq \text{mixin}\langle \Sigma_{\text{new}}, \Sigma_{\text{red}}, \Sigma_{\text{exp}} \rangle} (\sqsubseteq \text{ mixin})$	

Table 5. Subtype on class and mixin types.

$P ::= \mathbf{nil}$	(null process)
$\text{act}.P$	(action prefixing)
$P_1 \mid P_2$	(parallel composition)
X	(process variable)
$\text{let } x = \text{exp} \text{ in } P$	(object-oriented expression)
$\text{act} ::= \mathbf{out}(t)@l \mid \mathbf{in}(t)@l \mid \mathbf{read}(t)@l \mid \mathbf{eval}(P)@l$	
$l ::= l \mid \chi$	
$t ::= f \mid f, t$	
$f ::= \text{arg} \mid !id : \sigma$	
$id ::= x \mid X \mid \chi$	
$\text{arg} ::= id \mid e \mid P \mid l \mid v$	
$N ::= l :: P$	(single node)
$l :: \langle t \rangle$	(located tuple)
$N_1 \parallel N_2$	(net composition)

Table 6. O’KLAIM syntax (see Table 1 for the syntax of *exp* and *v* and Section 2.5 for types σ).

2.4 O’Klaim: processes and nets

O’KLAIM syntax is defined in Table 6. In order to obtain O’KLAIM, we extend the KLAIM syntax of tuples t to include any object-oriented value v (defined in Table 1). In particular, differently from KLAIM, formal fields are now explicitly typed. Actions $\mathbf{in}(t)@l$ (and $\mathbf{read}(t)@l$) and $\mathbf{out}(t)@l$ can be used to move object-oriented code (together with the other KLAIM items) from/to a locality l , respectively. Moreover, we add to KLAIM processes the construct $\text{let } x = \text{exp} \text{ in } P$ in order to pass to the sub-process P the result of computing an object-oriented expression *exp* (for *exp* syntax see Table 1). We use the following syntactic convention: x , X and χ are variables representing object-oriented values, processes and localities, respectively. A constant locality (e.g., IP:port) is denoted by l . Moreover, e is a basic expression (i.e., not object-oriented).

A *Net* is a finite collection of *nodes*. A node is a pair where the first component is a (constant) locality and the second component is either a process P or a tuple $\langle t \rangle$. Thus, a tuple space is represented by the parallel composition of located tuples.

$\frac{}{\Gamma, id : \sigma \vdash id : \sigma} (proj)$	$\frac{}{\Gamma \vdash l : loc} (loc)$	$\frac{}{\Gamma \vdash \mathbf{nil} : proc} (nil)$
$a \equiv \mathbf{in, read, out}$	$\frac{\begin{array}{l} \Gamma \vdash \ell : loc \\ \Gamma \vdash f_i : \sigma_i \quad i = 1, \dots, n \wedge f_i \equiv arg \\ \Gamma \cup ftypes(f_1, \dots, f_n) \vdash P : proc \end{array}}{\Gamma \vdash a(f_1, \dots, f_n)@l.P : proc} (action)$	
$ftypes(f, t) = \begin{cases} \{id : \sigma\} \cup ftypes(t) & \text{if } f \equiv !id : \sigma \\ ftypes(t) & \text{otherwise} \end{cases}$		
$\frac{\Gamma \vdash Q : proc \quad \Gamma \vdash \ell : loc}{\Gamma \vdash \mathbf{eval}(Q)@l.P : proc} (eval)$		
$\frac{\Gamma \vdash P_1 : proc \quad \Gamma \vdash P_2 : proc}{\Gamma \vdash (P_1 \mid P_2) : proc} (comp)$	$\frac{\Gamma \vdash exp : \tau \quad \Gamma, x : \tau \vdash P : proc}{\Gamma \vdash \mathbf{let } x = exp \mathbf{ in } P : proc} (let)$	

Table 7. Typing rules for processes

2.5 Typing for O’Klaim

In order to type processes and nets, we extend the set of types \mathcal{T} of Table 2 to $\mathcal{T}^* = \mathcal{T} \cup \{\mathbf{proc}, \mathbf{loc}\}$. σ will range over \mathcal{T}^* ; in particular, \mathbf{loc} is used to type localities and \mathbf{proc} for well-typed processes. Typing rules for processes are defined in Table 7. O’Klaim type system is not concerned with access rights and capabilities, as it is instead the type system for Klaim presented in [19]. In the O’Klaim setting, types serve the purpose of avoiding the “message-not-understood” error when merging local and foreign object-oriented code in a site. Thus, we are not interested in typing actions inside processes: from our perspective, an O’Klaim process is well typed when it has type \mathbf{proc} , which only means that the object-oriented code that the process may contain is well typed.

O’Klaim requires that every process is statically type-checked separately on its site and annotated with its type. The annotation process, not formally presented here, can be performed by the compiler during type checking: namely, every tuple item t_i that takes part in the information exchange (which may be an object-oriented value) must be decorated with its type information, denoted by $t_i^{\sigma_i}$. The types of the tuples are built statically by the compiler, while the types of tuple formal fields have to be written explicitly by the programmer. In a process of the form $\mathbf{in}(!id : \sigma)@l.P$, the type σ is used to statically type check the continuation P , where id is possibly used. More generally, concerning (*action*) rule, in a process performing an operation with a tuple (i.e., **out**, **read** and **in**), the actual fields of the tuple are type checked, and the types of formal fields (collected by the function $ftypes$) are used to type check the continuation.

We observe that the typing rules for object-oriented expressions are syntax-driven and do not contain an explicit subsumption rule. Thus, they define an algorithm to assign a principal type to each expression, in a given environment Γ .

$exp_1 \twoheadrightarrow$	$\left(\begin{array}{l} \text{mixin} \\ \text{expect}[m_i : \tau_{m_i} \mid i \in I] \\ \text{redef}[m_k : \tau_{m_k} \text{ with } b_k \mid k \in K] \\ \text{def}[m_j = b_j \mid j \in J] \\ \text{end} \end{array} \right)$
$exp_2 \twoheadrightarrow$	$\text{class } [m_l = b_l \mid l \in L] \text{ end}$
$exp_1 \diamond exp_2 \twoheadrightarrow$	$\left(\begin{array}{l} \text{class} \\ [m_j = b_j \mid j \in J] \cup \\ \text{override}([m_k = b_k \mid k \in K], [m_l = b_l \mid l \in L]) \\ \text{end} \end{array} \right)$

Table 8. The (*mixinapp*) operational rule

Analogously, both subtyping and typing rules for processes are in an algorithmic shape.

2.6 Operational semantics for O'Klaim

The operational semantics of O'KLAIM involves two sets of rules. The first set of rules describes how object-oriented expressions reduce to values and is denoted by \twoheadrightarrow . We omit here most of the rules because they are quite standard; they can be found in [8]. However, we want to discuss the rule concerning mixin application, that produces a new class containing all the methods which are added and redefined by the mixin and those defined by the superclass. The rule (*mixinapp*) is presented in Table 8. The function *override*, defined below and used by rule (*mixinapp*), takes care of introducing in the new class the overridden methods, and of binding the special variable *next* to the implementations provided by the super class in the mixin's redefined method bodies: these “old” method implementations are given a fresh name, denoted by $m_{i'}$. Dynamic binding is then implemented for redefined methods, and old implementations from the super class are basically hidden in the derived class, since they are given a fresh name.

Definition 1. *Given two method sets, ϱ_1 and ϱ_2 , the result of $\text{override}(\varrho_1, \varrho_2)$ is the method set ϱ_3 defined as follows:*

- for all $m_i = b_i \in \varrho_2$ such that $m_i \neq m_j$ for all $m_j = b_j \in \varrho_1$, then $m_i = b_i \in \varrho_3$;
- for all $m_i = b_i \in \varrho_1$ such that $m_i = b'_i \in \varrho_2$, let $m_{i'}$ be a fresh method name: then $m_{i'} = b'_i \in \varrho_3$ and $m_i = b_i[m_{i'}/\text{next}] \in \varrho_3$.

Notice that name clashes among methods during the application will never take place, since they have already been solved during the typing of mixin application.

The second set of rules for O'KLAIM, shown in Table 9, concerns processes and it is an extension of the operational semantics of KLAIM. Notice that the O'KLAIM's operational semantics must be defined on typed *compiled* processes, i.e., processes where each object-oriented value and tuples are decorated with their types, as explained in Section 2.5, because the crucial point is the dynamic matching of types. Indeed, an **out** operation adds a tuple decorated with a (static) type to a tuple space. Conversely, a process can perform an **in** action

$\frac{}{l :: \mathbf{out}(t)@l'.P \parallel l' :: P' \longrightarrow l :: P \parallel l' :: P' \mid \langle t \rangle} \text{ (OUT)}$
$\frac{match(t, t')}{l :: \mathbf{in}(t)@l'.P \parallel l' :: \langle t' \rangle \longrightarrow l :: P[t'/t] \parallel l' :: \mathbf{nil}} \text{ (IN)}$
$\frac{match(t, t')}{l :: \mathbf{read}(t)@l'.P \parallel l' :: \langle t' \rangle \longrightarrow l :: P[t'/t] \parallel l' :: \langle t' \rangle} \text{ (READ)}$
$\frac{}{l :: \mathbf{eval}(P)@l'.P' \parallel l' :: P'' \longrightarrow l :: P' \parallel l' :: P' \mid P} \text{ (EVAL)}$
$\frac{exp \twoheadrightarrow v}{l :: \mathbf{let } x = exp \text{ in } P \longrightarrow l :: P[v/x]} \text{ (LET)}$
$\frac{N \equiv N_1 \quad N_1 \longrightarrow N_2 \quad N_2 \equiv N'}{N \longrightarrow N'} \text{ (NET)}$

Table 9. O'KLAIM operational rules

$match(e, e)$	$match(l, l)$
$\frac{match(t_2, t_1)}{match(t_1, t_2)}$	$\frac{match(t_1, t_2) \quad match(t_3, t_4)}{match((t_1, t_3), (t_2, t_4))}$
	$\frac{match(\sigma, \sigma_i)}{match(!id : \sigma, t_i^{\sigma_i})}$
$match(\sigma_1, \sigma_2) = \begin{cases} \sigma_1 \sqsubseteq \sigma_2 & \text{if } \sigma_1 \text{ and } \sigma_2 \text{ are mixin or class types} \\ \sigma_1 <: \sigma_2 & \text{otherwise} \end{cases}$	

Table 10. Matching rules (with **proc** <: **proc** and **loc** <: **loc**)

by synchronizing with a process which represents a matching typed tuple. The rule for **let** $x = exp$ in P relies on the reduction relation for object-oriented expressions \twoheadrightarrow .

The predicate for tuples, $match$, is presented in Table 10. The matching rule exploits the static type information, delivered together with the tuple items, in order to dynamically check that the received item is correct with respect to the type of the formal field, say τ . Therefore, an item is accepted if and only if it is subtyping-compliant with the expected type of the formal field. Informally speaking, one can accept any class containing more resources than expected, and any mixin with weaker requests about methods expected from the superclass can be accepted. This subtyping checking is analogous to the one we would perform in a sequential language where mixins and classes could be passed as parameters to

$$\begin{array}{c}
N_1 \parallel N_2 = N_2 \parallel N_1 \\
(N_1 \parallel N_2) \parallel N_3 = N_1 \parallel (N_2 \parallel N_3) \\
l :: P = l :: P \mid \mathbf{nil} \\
l :: (P_1 \mid P_2) = l :: P_1 \parallel l :: P_2
\end{array}$$

Table 11. Congruence laws

methods. In a sequential setting, this dynamic checking might look as a burden¹, but in a distributed mobile setting the burden seems well-compensated by the added flexibility in communications.

Finally, the semantics for the distributed part is based on structural congruence and reduction relations. Reduction represents individual computation steps, and is defined in terms of structural congruence. The structural congruence \equiv (defined as the least congruence relation closed under the rules in Table 11) allows the rearrangement of the syntactic structure of a term so that reduction rules may be applied.

A final remark. Let us observe that we do not define a matching predicate for actual fields containing object-oriented values and processes since this would require to decide equalities on classes, mixins and objects (e.g., equality on their interfaces) and on processes (e.g., a bisimulation). This issue is out of the scope of the present work, since matching between two actual fields does not involve any substitution and then does not cause problems w.r.t. typing.

3 Typing issues and subject-reduction property

The important point in O’KLAIM semantics is that if a process P (statically well-typed) retrieves a tuple by the subtyping matching mechanism and the retrieved value is merged in the continuation of P , then the evaluation proceeds without any additional type-checking. Thus, in order to obtain the subject-reduction theorem, we need to prove that substitution preserves well-typedness, in particular when classes and mixins are replaced to variables inside object-oriented expressions. In the following, we address this issue and we outline the main technical steps, skipping proofs and details for lack of space.

The crucial case concerns mixin application expressions; namely if class/mixin variables are replaced by classes/mixins having a subtype, accidental overrides can occur because of names of the new methods added by the replacing value (see the definition of \sqsubseteq). This matter is related to the “width subtyping versus method addition” problem (well known in the object-based setting, see for instance [21]), that in our case boils down to a careful management of these *dynamic name clashes*. Thus, we have to define a suitable capture-avoid-substitution, $[]$, requiring possible renaming of methods with fresh names.

Definition 2 (Substitution by refresh). *If x is a class variable of type $\text{class}\langle\Sigma\rangle$ and C is a class value of type $\text{class}\langle\Sigma'\rangle$ such that $\text{class}\langle\Sigma'\rangle \sqsubseteq \text{class}\langle\Sigma\rangle$, then $[C/x]$ denotes the replacement of C' to x , where C' is obtained from C*

¹ For example, in [13], mixins and classes are first-order entities, i.e., they can be passed as parameters in methods, but the matching among formal and actual parameters is by syntactic equality on types and not by subtyping.

by renaming all methods belonging to $\text{Meth}(\Sigma') - \text{Meth}(\Sigma)$ with fresh names. For mixins variables and values, the renaming acts on all and only the methods belonging to $\text{Meth}(\Sigma'_{\text{new}}) - \text{Meth}(\Sigma_{\text{new}})$.

With our solution, new methods added by a class or a mixin value during substitution are hidden by renaming, for each occurrence of the variable to be replaced (this is very similar to the “privacy via subsumption” of [27]). On the other hand, we only rename methods that do not appear in the type of the variable x . This second constraint ensures a basic property: the refreshed version C' of C has a type τ' such that $\tau' \sqsubseteq \text{class}(\Sigma)$. The same holds for refreshed mixins.

Now, using this property in an essential way, we can prove that substitution is type safe.

Lemma 1 (Substitution by narrowing). *Let v , exp and P be an object-oriented value, an object-oriented expression and a process, respectively,*

1. *if $x : \tau_1 \vdash \text{exp} : \tau$ and $\Gamma \vdash v : \tau_2$ where $\tau_2 \sqsubseteq \tau_1$, then $\Gamma \vdash \text{exp}[v/x] : \tau'$ with $\tau' \sqsubseteq \tau$;*
2. *if $x : \tau_1 \vdash P : \text{proc}$ and $\Gamma \vdash v : \tau_2$ where $\tau_2 \sqsubseteq \tau_1$, then $\Gamma \vdash P[v/x] : \text{proc}$.*

Sketch of proof:

1. By induction on typing rules for expressions. The only crucial case is when exp is a mixin application and v is a class value or a mixin. Notice that, exp is well-typed and no method occurring in the type x is renamed; then the last condition in (*mixin app*) rule is preserved and ensures that no name clash can occur after substitution.
2. By induction on typing rules for processes using the previous point.

Summarizing, the type safety of the communication results from two main properties: (i) static type soundness of local and foreign code; (ii) the preservation of well-typedness under substitution by subtyping. It is standard to verify that all the other rules concerning \rightarrow preserve well typedness and so we obtain the *subject-reduction theorem*. Thus, the local evaluation of a process cannot produce errors like “message-not-understood” even if it retrieves data from foreign sites and merges it in the local configuration. In other words, a *well-typed net* (i.e., a net where each process in each site is well-typed) remains well typed during its evolution (*global safety condition*).

We remark that, from the point of view of the implementation, the above treatment of “global” fresh names can be solved with static binding for the mentioned methods. The technique of using the static types of variables and the actual types of substituted mixin and class definitions may recall the approach of [22] of allowing overriding, i.e., dynamic binding, only for methods declared in the mixin’s *inheritance interface*.

4 The implementation

We recall that the implementation we present is based on X-KLAIM [10, 11] (extended with the proper object-oriented mixin-based primitives), used both as the “surface” object-oriented calculus and as the coordination language, with the added bonus of being able to write methods that can perform KLAIM actions, all the same guaranteeing absence of message-not-understood run-time errors, as shown in the previous section.

The implementation of the O’KLAIM object-oriented component in Java consists in a package `momi` presented in details in [3]. This package provides the run-time system, or the virtual machine, for classes, mixins and objects that can be downloaded from the network and dynamically composed (via the mixin application operation). It thus provides functionalities for checking subtyping among classes and among mixins and for building at run-time new subclasses. Since we abstract from the specific communication and mobility features, this package does not provide means for code mobility and network communication, so that `momi` can be smoothly integrated into existing Java mobility frameworks. We would like to stress that this package should be thought of as an “assembly” language that is the target of a compiler for a high-level language (in our case the language is X-KLAIM). If `momi`, as it is, was used for directly writing object-oriented expressions, the programmer would be left with the burden of writing methods containing Java statements dealing with `momi` objects, classes and mixins, and to check manually that they are well typed. Basically these are the same difficulties a programmer has to face when using an assembly language directly, instead of a high-level language. We could say that `momi` enhances the functionalities of the Java virtual machine: while the latter already provides useful mechanisms for dynamically loading new classes into a running application, the former supplies means for dynamically building class hierarchies (based on mixins) and for inserting new subclasses into existing hierarchies (which is not possible in Java).

In order to implement O’KLAIM we extended the KLAIM programming framework that consists in the programming language X-KLAIM [10, 11], which extends KLAIM with high-level programming constructs, and KLAVA [12] a Java package that implements the run-time system for X-KLAIM operations (X-KLAIM programs are compiled into Java programs that use KLAVA). The package KLAVA already provided all the primitives for network communication, through distributed tuple spaces, and, in particular, for code mobility, not supplied by `momi`. Thus the package has been modified in order to be able to exchange code that is based on `momi`, and for performing subtyping on `momi` elements during pattern matching by relying on the `MoMiType` classes and the associated subtyping. On the other hand, the X-KLAIM compiler generates code that uses both the KLAVA package and `momi`. Obviously, before generating code, it also performs type checking according to the type system defined by MoMi. All this software is freely available at <http://music.dsi.unifi.it>.

The programming example shown in this section involves mixin code mobility, and implements “dynamic inheritance” since the received mixin is applied to a

<pre> mixin MyPrinterAgent expect print_doc(doc : str) : str; def start_agent() : str begin return this.print_doc (this.preprocess("my document")) end; def preprocess(doc : str) : str begin return "preprocessed(" + doc + ")" end end rec SendPrinterAgent[server : loc] declare var response : str begin out(MyPrinterAgent)@server; in(!response)@server; print "response is " + response end </pre>	<pre> mixin PrinterAgent expect print_doc(doc : str) : str; def start_agent() : str; end class LocalPrinter print_doc(doc : str) : str begin # real printing code omitted :-) return "printed " + doc end; init() begin nil # foo init end end rec ReceivePrinterAgent[] declare var rec_mixin : mixin PrinterAgent; var result : str begin in(!rec_mixin)@self; result := (new rec_mixin <> LocalPrinter).start_agent(); out(result)@self end </pre>
--	---

Listing 1: The printer agent example.

local parent class at run-time. We assume that a site provides printing facilities for local and mobile agents. Access to the printer requires a driver that the site itself has to provide to those that want to print, since it highly depends on the system and on the printer. Thus, the agent that wants to print is designed as a mixin, that expects a method for actually printing, `print_doc`, and defines a method `start_agent` through which the site can start its execution. The actual instance of the printing agent is instantiated from a class dynamically generated by applying such mixin to a local superclass that provides the method `print_doc` acting as a wrapper for the printer driver. However the system is willing to accept any agent that has a compatible interface, i.e., any mixin that is a subtype of the one used for describing the printing agent. Thus any client wishing to print on this site can send a mixin that is subtyping compliant to the one expected. In particular such a mixin can implement finer printing formatting capabilities.

Listing 1, where `rec` is the X-KLAIM keyword for defining a process, presents a possible implementation of the printing client node (on the left) and of the printer server node (on the right). The printer client sends to the server a mixin `MyPrinterAgent` that complies with (it is a subtype of) the mixin that the server expects to receive, `PrinterAgent`. In particular `MyPrintedAgent` mixin will print a document on the printer of the server after preprocessing it (method `preprocess`). On the server, once the mixin is received, it is applied to the local (super)class `LocalPrinter`, and an object (the agent) is instantiated from the resulting class, and started so that it can actually print its document. The result of the printing task is then retrieved and sent back to the client.

We observe that the sender does not actually know the mixin name `PrinterAgent`: it only has to be aware of the mixin type expected by the server. Furthermore, the sent mixin can also define more methods than those specified

in the receiving site, thanks to the mixin subtyping relation. This adds a great flexibility to such a system, while hiding these additional methods to the receiving site (since they are not specified in the receiving interface they are actually unknown statically to the compiler).

5 Conclusions and related work

We have presented the kernel language O’KLAIM, which extends the higher-order calculus KLAIM for mobile processes with mixin-based object-oriented code.

The novel contributions of this paper, with respect to [7] where we firstly presented design motivations for a mixin-based approach in a mobile context, can be summarized as follows:

1. we integrate the basic ideas of [7] into a mobile process calculus with an asynchronous and more sophisticated communication mechanism;
2. we refine the typing for the object-oriented component, we define a new typing system for KLAIM processes, and we study main typing concerns (in particular, a notion of substitution with renaming) in order to demonstrate the soundness of the proposed solution;
3. we present an implementation of O’KLAIM.

Keeping the O’KLAIM object-oriented calculus and the O’KLAIM processes separated may appear a limitation, but in fact this is not true. Our system consists of three components: the “surface” object-oriented component, a mixin/class subtyping relation, and a coordination calculus. If the object-oriented component is chosen to be an object-oriented concurrent/mobile language, the two components (object-oriented and concurrent/mobile) may interleave in a deeper way. A good example is the O’KLAIM implementation presented in Section 4: in there, X-KLAIM (extended with the proper object-oriented mixin-based primitives) is both the “surface” object-oriented calculus and the coordination language, so that method bodies can perform KLAIM actions. The matching mechanism that allows safe interactions during code exchange is based on the subtyping relation that acts as a general glue to glue together the two language components, of whichever nature they are (as long as the object-oriented one implements classes and mixins).

In the literature, there are several proposals of combining objects with processes and/or mobile agents. *Obliq* [16] is a lexically-scoped language providing distributed object-oriented computation. Mobile code maintains network references and provides transparent access to remote resources. In [15], a general model for integrating object-oriented features in calculi of mobile agents is presented where agents are extended with constructs for remote method invocations. Other works, such as, e.g., [20, 26, 25] do not deal explicitly with mobile distributed code. Our approach is more related to papers, as [29], where properties of distributed systems are enforced by a typing system equipped with subtyping. In our case the property we address is a flexible and type-safe coordination for exchanging code among processes, up- and down-loading classes and mixins from different sites.

A topic for future research is an extension of O’KLAIM with *subtyping-in-depth* (a preliminary investigation of this subtyping be found in [9]). Subtyping-in-depth offers a more flexible communication pattern, but it complicates the object-oriented code exchange for problems similar to the “subtyping-in-depth versus override” matter of the object-based languages (see [1, 9] for examples). Another direction is to enrich the object-oriented component, for example with *incomplete objects*, i.e., objects instantiated from mixins [5]; this would imply a further refinement of the subtyping relation.

References

1. M. Abadi and L. Cardelli. *A Theory of Objects*. Springer, 1996.
2. D. Ancona, G. Lagorio, and E. Zucca. Jam - designing a java extension with mixins. *ACM Transaction on Programming Languages and Systems*, 2003. To appear.
3. L. Bettini. *Linguistic Constructs for Object-Oriented Mobile Code Programming & their Implementations*. PhD thesis, Dip. di Matematica, Università di Siena, 2003. Available at <http://music.dsi.unifi.it>.
4. L. Bettini, V. Bono, R. De Nicola, G. Ferrari, D. Gorla, M. Loreti, E. Moggi, R. Pugliese, E. Tuosto, and B. Venneri. The KLAIM Project: Theory and Practice. In *Global Computing – Trento*, LNCS. Springer, 2003. To appear.
5. L. Bettini, V. Bono, and S. Likavec. A core calculus of mixin-based incomplete objects. Submitted, 2003.
6. L. Bettini, V. Bono, and B. Venneri. Towards Object-Oriented KLAIM. In M. Lenisa and M. Miculan, editors, *TOSCA 2001*, volume 62 of *ENTCS*. Elsevier, 2001.
7. L. Bettini, V. Bono, and B. Venneri. Coordinating Mobile Object-Oriented Code. In F. Arbarb and C. Talcott, editors, *Proc. of Coordination Models and Languages*, volume 2315 of *LNCS*, pages 56–71. Springer, 2002.
8. L. Bettini, V. Bono, and B. Venneri. MoMi - A Calculus for Mobile Mixins. Manuscript, 2003.
9. L. Bettini, V. Bono, and B. Venneri. Subtyping Mobile Classes and Mixins. In *Proc. of Int. Workshops on Foundations of Object-Oriented Languages, FOOL 10*, 2003.
10. L. Bettini, R. De Nicola, G. Ferrari, and R. Pugliese. Interactive Mobile Agents in X-KLAIM. In P. Ciancarini and R. Tolksdorf, editors, *Proc. of the 7th Int. IEEE Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE)*, pages 110–115. IEEE Computer Society Press, 1998.
11. L. Bettini, R. De Nicola, and R. Pugliese. X-KLAIM and KLAVA: Programming Mobile Code. In M. Lenisa and M. Miculan, editors, *TOSCA 2001*, volume 62 of *ENTCS*. Elsevier, 2001.
12. L. Bettini, R. De Nicola, and R. Pugliese. KLAVA: a Java package for distributed and mobile applications. *Software – Practice and Experience*, 32(14):1365–1394, 2002.
13. V. Bono, A. Patel, and V. Shmatikov. A Core Calculus of Classes and Mixins. In R. Guerraoui, editor, *Proceedings ECOOP’99*, number 1628 in *LCNS*, pages 43–66. Springer-Verlag, 1999.
14. G. Bracha and W. Cook. Mixin-based inheritance. In *Proc. OOPSLA ’90*, pages 303–311, 1990.
15. M. Bugliesi and G. Castagna. Mobile Objects. In *Proc. of FOOL*, 2000.
16. L. Cardelli. A Language with Distributed Scope. *Computing Systems*, 8(1):27–59, 1995.

17. A. Carzaniga, G. Picco, and G. Vigna. Designing Distributed Applications with Mobile Code Paradigms. In R. Taylor, editor, *Proc. of the 19th Int. Conf. on Software Engineering (ICSE '97)*, pages 22–33. ACM Press, 1997.
18. R. De Nicola, G. Ferrari, and R. Pugliese. KLAIM: a Kernel Language for Agents Interaction and Mobility. *IEEE Transactions on Software Engineering*, 24(5):315–330, 1998.
19. R. De Nicola, G. Ferrari, R. Pugliese, and B. Venneri. Types for Access Control. *Theoretical Computer Science*, 240(1):215–254, 2000.
20. P. Di Blasio and K. Fisher. A Calculus for Concurrent Objects. In U. Montanari and V. Sassone, editors, *CONCUR '96: Concurrency Theory, 7th Int. Conf.*, volume 1119 of *LNCS*, pages 655–670. Springer, 1996.
21. K. Fisher and J. C. Mitchell. A Delegation-based Object Calculus with Subtyping. In *Proc. of FCT*, volume 965 of *LNCS*, pages 42–61. Springer, 1995.
22. M. Flatt, S. Krishnamurthi, and M. Felleisen. Classes and mixins. In *Proc. POPL '98*, pages 171–183, 1998.
23. C. Fournet, G. Gonthier, J. J. Levy, L. Maranget, and D. Remy. A Calculus of Mobile Agents. In U. Montanari and V. Sassone, editors, *Proc. of 7th Int. Conf. on Concurrency Theory (CONCUR'96)*, volume 1119 of *LNCS*, pages 406–421. Springer-Verlag, 1996.
24. D. Gelernter. Generative Communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, 1985.
25. A. Gordon and P. Hankin. A Concurrent Object Calculus: Reduction and Typing. In U. Nestmann and B. C. Pierce, editors, *Proc. of HLCL '98: High-Level Concurrent Languages*, volume 16.3 of *ENTCS*. Elsevier, 1998.
26. B. C. Pierce and D. N. Turner. Concurrent Objects in a Process Calculus. In T. Ito and A. Yonezawa, editors, *Proc. Theory and Practice of Parallel Programming (TPPP 94)*, volume 907 of *LNCS*, pages 187–215. Springer, 1995.
27. J. Riecke and C. Stone. Privacy via Subsumption. *Information and Computation*, 172:2–28, 2002. 3rd special issue of Theory and Practice of Object-Oriented Systems (TAPOS).
28. T. Thorn. Programming Languages for Mobile Code. *ACM Computing Surveys*, 29(3):213–239, 1997. Also Technical Report 1083, University of Rennes IRISA.
29. N. Yoshida and M. Hennessy. Subtyping and Locality in Distributed Higher Order Mobile Processes (extended abstract). In *Proc. of CONCUR'99*, volume 1664 of *LNCS*, pages 557–572. Springer-Verlag, 1999.