

A Type System for Components

Ornela Dardha, Elena Giachino, Michael Lienhardt

► **To cite this version:**

Ornela Dardha, Elena Giachino, Michael Lienhardt. A Type System for Components. Robert M. Hierons and Mercedes G. Merayo and Mario Bravetti. SEFM - International Conference on Software Engineering and Formal Methods - 2013, 2013, Madrid, Spain. Springer, 8137, pp.167-181, 2013, Lecture Notes in Computer Science. <10.1007/978-3-642-40561-7_12>. <hal-00909310>

HAL Id: hal-00909310

<https://hal.inria.fr/hal-00909310>

Submitted on 26 Nov 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Type System for Components^{*}

Ornela Dardha¹ Elena Giachino¹ Michaël Lienhardt²

¹ INRIA Focus Team / University of Bologna, Italy
{dardha, giachino}@cs.unibo.it

² University of Paris Diderot, France
lienhar@inria.fr

Abstract. In modern distributed systems, dynamic reconfiguration, i.e., changing at runtime the communication pattern of a program, is challenging. Generally, it is difficult to guarantee that such modifications will not disrupt ongoing computations. In a previous paper, a solution to this problem was proposed by extending the object-oriented language ABS with a component model allowing the programmer to: *i*) perform updates on objects by means of *communication ports* and their *rebinding*; and *ii*) precisely specify when such updates can safely occur in an object by means of *critical sections*. However, improper rebinding operations could still occur and lead to runtime errors. The present paper introduces a type system for this component model that extends the ABS type system with the notion of ports and a precise analysis that statically enforces that no object will attempt illegal rebinding.

1 Introduction

In modern complex distributed scenarios, unplanned dynamic reconfiguration, i.e., changing at runtime the communication pattern of a program, is challenging as it is difficult to ensure that such modifications will not disrupt ongoing computations. In [14] the authors propose to solve the problem by integrating notions coming from component models [2–4, 8] within the actor-based *Abstract Behavioral Specification* programming language (ABS) [13]. ABS is designed for distributed object-oriented systems and integrates concurrency and synchronization mechanisms to solve data races. Actors, called *cogs* or simply *groups*, are dynamic collections of collaborating objects. Cogs offer consistency by guaranteeing that at most one method per cog is executing at any time. Within a cog, objects collaborate using (*synchronous*) method calls and *collaborative concurrency* with the **suspend** and **await** operations which can suspend the execution of the current method, and thus allow another one to execute. Between cogs, collaboration is achieved by means of *asynchronous* method calls that return *future*, i.e., a placeholder where the result of the call is put when its computation finishes.

^{*} This research is partly funded by the EU project FP7-231620 HATS and by the French National Research Agency (ANR), projects REVER ANR 11 INSE 007

On top of the ABS language, [14] adds the notions of *ports*, *bindings* and *safe state* to deal with dynamic reconfiguration. Ports define variability points in an object and can be *rebound* (i.e., modified) from outside the object (on the contrary, fields, which represent the inner state of the object, can only be modified by the object itself). To ensure consistency of the **rebind** operation, [14] enforces two constraints on its application: *i*) it is only possible to rebind an object’s port when the object is in a *safe state*; and *ii*) it is only possible to rebind an object’s port from *any* object within the *same* cog. Safe states are modeled by annotating methods as **critical**, specifying that while one or more **critical** methods are executing, the object is *not* in a safe state. The resulting language offers a consistent setting for dynamic reconfigurations, which means performing modifications on a program at runtime while still ensuring consistency of its execution. Consistency is based on two constraints: both synchronous method calls and rebinding operations must involve two objects in the same cog. These constraints are enforced at runtime; therefore, programs may encounter unexpected runtime errors during their execution.

In this paper, we define a type system for the component model that statically ensures the legality of both synchronous method calls and port rebindings, guaranteeing that well-typed programs will always be consistent. Our approach is based on a static tracking of group membership of the objects. The difficulty in retrieving this kind of information is that cogs as well as objects are dynamic entities. Since we want to trace group information statically, we need a way to identify and track every group in the program. To this aim, we define a technique that associates to each group creation a fresh *group name*. Then, we keep track of which cog an object is allocated to, by associating to each object a *group record*. The type system checks that objects indeed have the specified group record, and uses this information to ensure that synchronous calls and rebindings are always performed locally to a cog. The type system is proven to be sound with respect to the operational semantics. We use this result to show that well-typed programs do not violate consistency during execution.

Motivating example. In the following we present a running example that gives a better understanding of the ABS language and the component extension, and most importantly, motivates our type system. Consider the following typical distributed scenario: suppose that we have several clients working together in a specific workflow and using a central server for their communications. Updating the server is a difficult task, as it requires to update its reference in all clients at the same time in order to avoid communication failures.

First, in Fig. 1 we consider how this task is achieved in ABS. The programmer declares two interfaces **Server** and **Client** and a class **Controller**. Basically, the class **Controller** updates the server in all the clients c_i by synchronously calling their setter method. All the clients are updated at the same time: since they are in the same cog as the controller they cannot execute until the execution of method `updateServer` has terminated.

```

interface Server { ... }
interface Client { Unit setServer(Server s); ... }

class Controller {
  Client c1, c2, ... cn;

  Unit updateServer(Server s2) {
    c1.setServer(s2);
    c2.setServer(s2);
    ...
    cn.setServer(s2);
  }}

```

Fig. 1. Workflow in ABS.

However, this code does not ensure that the update is performed when the clients are in a safe state. This can lead to inconsistencies because clients that are using the server are not aware of the modification taking place. This problem can be solved by using the notions of **port** and **rebind** [14] as shown in Fig. 2. Here, the method `updateServer` first waits for all clients to be in a safe state (**await** statement performed on the conjunction of all clients) and then updates their reference one by one (**rebind** server `s` which is declared to be a **port**).

```

interface Server { ... }
interface Client { port Server s; ... }

class Controller {
  Client c1, c2, ... cn;
  ...
  Unit updateServer(Server s2) {
    await ||c1|| ^ ||c2|| ^ ... ^ ||cn||;
    rebind c1.s = s2;
    rebind c2.s = s2;
    ...
    rebind cn.s = s2;
  }}

```

Fig. 2. Workflow using the Component Model.

However, runtime errors can still occur. For instance, if the clients and the controller are not in the same cog, the update will fail. Consider the code in Fig. 3. Method `main` instantiates classes `Client` and `Controller` –and possibly other classes, like `Server`, present in the program– by creating objects `c1, c2, ..., cn, c`. These objects are created in the same cog by the **new** command,

except for client c_1 , which is created and placed in a new cog by the `new cog` command. Now, suppose that the code in Fig. 2 is executed. At runtime, the program will check if the controller and the client belong to the same cog to respect the consistency constraints on rebinding. In case of c_1 this check will fail by leading to a runtime error.

The present paper addresses this problem in order to avoid these runtime errors and the overhead in dealing with them. We present a type system that tracks cog membership of objects thus permitting to typecheck only programs where rebinding is consistent. So, the code presented above would not typecheck, as shown in § 3, thus discarding the program at compile time instead of leading to a runtime error.

```

Unit main () {
    ...
    Client c1 = new cog Client (s);
    Client c2 = new Client (s);
    ...
    Client cn = new Client (s);
    Controller c = new Controller (c1, c2, ... cn);
}

```

Fig. 3. Client and Controller objects creation.

Roadmap. The rest of the paper is structured as follows: § 2 introduces the calculus, types and terms; § 3 presents our type system and its properties; and § 4 concludes the paper and discusses future and related works.

2 The calculus

In this section we present the calculus underlying our approach, which is a component extension of the ABS language³. We present formally only the syntax of the calculus which is necessary for specifying the type system. We already gave some intuitions about the operational semantics of the calculus in the introduction and through the example, whereas for the formal definition we refer to the original paper [14] and the extended version of this paper [9].

The syntax of the calculus is given in Fig. 4 and corresponds to the original one, except for types, which are here extended in order to store also group information. This syntax is based on several categories of names: **I** and **C** range over interface and class names; **V** ranges over type variables for polymorphism; **G**

³ For the sake of readability, the calculus we consider is a subset of [14]. The notion of *location* has been dropped, since it is orthogonal to ports and rebinding. The validity of our approach and of our type system still holds for the full calculus.

$P ::= \overline{Dl} \{ s \}$	Program
$Dl ::= D \mid F \mid I \mid C$	Declarations
$T ::= V \mid D[\overline{T}] \mid (I, r)$	Type
$r ::= \perp \mid G[\overline{f : T}] \mid \alpha \mid \mu\alpha.r$	Record
$D ::= \mathbf{data} D[\overline{T}] = \mathbf{Co}[\overline{T}] \mid \mathbf{Co}[\overline{T}];$	Data Type
$F ::= \mathbf{def} T \mathbf{fun}[\overline{T}](\overline{T} x) = e;$	Function
$I ::= \mathbf{interface} I [\mathbf{extends} \overline{I}] \{ \mathbf{port} T x; \overline{S} \}$	Interface
$C ::= \mathbf{class} C [(\overline{T} x)] [\mathbf{implements} \overline{I}] \{ \overline{Fl} \overline{M} \}$	Class
$Fl ::= [\mathbf{port}] T x$	Field Declaration
$S ::= [\mathbf{critical}] (G, r) T m(\overline{T} x)$	Method Header
$M ::= S \{ s \}$	Method Definition
$s ::= \mathbf{skip} \mid s; s \mid T x \mid x = z \mid \mathbf{await} g$ $\mid \mathbf{if} e \{ s \} \mathbf{else} \{ s \} \mid \mathbf{while} e \{ s \} \mid \mathbf{return} e$ $\mid \mathbf{rebind} e.p = z$	Statement
$z ::= e \mid \mathbf{new} [\mathbf{cog}] C(\overline{e}) \mid e.m(\overline{e}) \mid e.lm(\overline{e}) \mid \mathbf{get}(e)$	Expression with Side Effects
$e ::= v \mid x \mid \mathbf{fun}(\overline{e}) \mid \mathbf{case} e \{ \overline{p} \Rightarrow e_p \}$	Expression
$v ::= \mathbf{null} \mid \mathbf{Co}[\overline{v}]$	Value
$p ::= - \mid x \mid \mathbf{null} \mid \mathbf{Co}[\overline{p}]$	Pattern
$g ::= x \mid x? \mid \ x\ \mid g \wedge g$	Guard

Fig. 4. Core ABS Language and Component Extension.

ranges over cog names, which will be explained thoroughly in § 3; D , \mathbf{Co} and \mathbf{fun} range respectively over data type, constructor and function names; m , f and p range respectively over method, field and port names (in order to have a uniform presentation, we will often use f for both fields and ports); and x ranges over variables, with the addition of the special variable **this** indicating the current object. For the sake of readability, we use the following notations: an overlined element corresponds to any finite, possibly empty, sequence of such element; and an element between square brackets is optional.

A program P consists of a sequence of declarations ended by a main block, namely a statement s to be executed. Declarations include data type declarations D , function declarations F , interface declarations I and class declarations C . A type T can be: a type variable V ; a data type D like **Bool** or futures $\mathbf{Fut}\langle T \rangle$, used to type data structures; or a pair consisting of an interface name I and a *record* r to type objects. Note that the ABS type system only uses interface names to type objects, but here we add records to track in which cog an object is located. Records can be: \perp , meaning that the structure of the object is unknown; $G[\overline{f : T}]$, meaning that the object is in the cog G and its fields \overline{f} are typed with \overline{T} ; or regular terms, using the standard combination of variables α and the μ -binder. Data types D have at least one constructor, with name \mathbf{Co} , and possibly a list of type parameters \overline{T} . Functions F are declared with a return type T , a name \mathbf{fun} , a list of parameters $\overline{T} x$ and a code e . Interfaces I declare methods and ports that can be modified at runtime. Classes C implement interfaces; they have a list of fields and ports \overline{Fl} and implement all declared

methods. Method headers S are used to declare methods with their classic type annotation, and *i*) the possible annotation **critical** that ensures that no rebinding will be performed on that object during the execution of that method; and *ii*) a *method signature* $(\mathcal{G}, \mathbf{r})$ which will be described and used in our type system section. Method declarations M consist of a header and a body, the latter being a sequential composition of local variables and commands. Statements s are standard except for **await** g , which suspends the execution of the method until the guard g is **true**, and **rebind** $e.p = z$, which rebinds the port p of the object e to the value stored in z . Expressions z include: expressions without side effects e ; **new** $\mathbf{C}(\bar{e})$ and **new cog** $\mathbf{C}(\bar{e})$ that instantiate a class \mathbf{C} and place the object in the current cog and in a new cog, respectively; synchronous $e.\mathbf{m}(\bar{e})$ and asynchronous $e!\mathbf{m}(\bar{e})$ method calls, the latter returning a future that will hold the result of the method call when it will be computed; and **get** (e) which gives the value stored in the future e , or actively waits for it if it is not computed yet. Pure expressions e include values v , variables x , function call $\mathbf{fun}(\bar{e})$ and pattern matching **case** $e \{ \bar{p} \Rightarrow e_p \}$ that tests e and execute e_p if it matches p . Patterns p are standard: $_$ matches everything, x matches everything and binds it to x , **null** matches a null object and $\mathbf{Co}(\bar{p})$ matches a value $\mathbf{Co}(\bar{e}_p)$ where p matches e_p . Finally, a guard g can be: a variable x ; $x?$ which is **true** when the future x is completed, **false** otherwise; $\|x\|$ which is **true** when the object x is in a safe state, i.e., it is not executing any **critical** method, **false** otherwise; and the conjunction of two guards $g \wedge g$ has the usual meaning.

3 Type System

The goal of our type system is to statically check whether synchronous method calls and rebindings are performed locally to a cog. Since cogs and objects are entities created at runtime, we cannot know statically their identity. We address this issue by using a *linear* type system approach on names of cogs $\mathbf{G}, \mathbf{G}', \mathbf{G}'' \dots$ that abstracts the runtime identity of cogs. This type system associates to every cog creation a unique cog name, which makes it possible to check if two objects are in the same cog or not. Precisely, we associate objects to their cogs using records \mathbf{r} , having the form $\mathbf{G}[\bar{f} : \bar{T}]$, where \mathbf{G} denotes the cog in which the object is located and $[\bar{f} : \bar{T}]$ maps any object's fields \bar{f} to its type \bar{T} . In order to correctly track cog membership of each expression, we also need to keep information about the cog of the object's fields in a record. This is needed, for instance, when an object stored in a field is accessed within the method body and then returned by the method; in this case one needs a way to bind the cog of the accessed field to the cog of the returned value. Let us now explain the method signature $(\mathcal{G}, \mathbf{r})$ annotating a method header. The record \mathbf{r} is used as the record of **this** during the typing of the method, i.e., \mathbf{r} is the binder for the cog of the object **this** in the scope of the method body, as we will see in the typing rules in the following. The set of cog names \mathcal{G} is used to keep track of the fresh cogs that the method creates. In particular, when we deal with recursive method calls, the set \mathcal{G} gathers the fresh cogs of every call, which is then returned to the main

$$\begin{array}{c}
\text{S:DATA} \\
\frac{\forall i \ T_i \leq T'_i}{D\langle \bar{T} \rangle \leq D\langle \bar{T}' \rangle} \\
\\
\text{S:PORTS} \\
\frac{\forall i \ T_i \leq T'_i \quad f \in \text{ports}(L)}{(L, \mathbf{G}[f : \bar{T}]) \leq (L, \mathbf{G}[f : \bar{T}'])} \\
\\
\text{S:TYPE} \\
\frac{L \leq L' \in CT}{(L, \mathbf{r}) \leq (L', \mathbf{r})}
\end{array}
\quad
\begin{array}{c}
\text{S:BOT} \\
(L, \mathbf{r}) \leq (L, \perp) \\
\\
\text{S:FIELDS} \\
\frac{\forall i \ T_i \leq T'_i \quad f \notin \text{ports}(L)}{(L, \mathbf{G}[f : \bar{T}; f : \bar{T}']) \leq (L, \mathbf{G}[f : \bar{T}'])}
\end{array}$$

Fig. 5. Subtyping Relation⁴

execution. Moreover, when it is not necessary to keep track of cog information about an object, because the object is not going to take part in any synchronous method call or any rebind operation, it is possible to associate to this object the *unknown* record \perp . This special record does not keep any information about the cog where the object or its fields are located, and it is to be considered different from any other cog, thus to ensure the soundness of our type system. Finally, note that data types also can contain records: for instance, a list of objects is typed with $\text{List}\langle T \rangle$ where T is the type of the objects in the list and it includes also the record of the objects.

A *typing environment* Γ is a partial function from names to typings, which assigns types T to variables, a pair (\mathbf{C}, \mathbf{r}) to **this**, and arrow types $\bar{T} \rightarrow T'$ to function symbols like **Co** or **fun**.

3.1 Subtyping Relation

The subtyping relation \leq on types is a preorder and is presented in Fig. 5. Rule S:DATA states that data types are covariant in their type parameters. Rule S:BOT states that every record \mathbf{r} is a subtype of the unknown record \perp . Rules S:FIELDS and S:PORTS use *structural* subtyping on records. Fields, like methods, are what the object provides, hence it is sound to forget about the existence of a field in an object. This is why the rule S:FIELDS allows to remove fields from records. Ports on the other hand, model the *dependencies* the objects have on their environment, hence it is sound to consider that an object may have more dependencies than it actually has during execution. This is why the rule S:PORTS allows to add ports to records. Notice that in the standard object-oriented setting this rule would not be sound, since trying to access a non-existing attribute would lead to a null pointer exception. Therefore, to support our vision of port behavior, we add a REBIND-NONE reduction rule to the component calculus semantics which simply permits the rebind to succeed without modifying anything if the port is not available. Finally, rule S:TYPE adopts *nominal* subtyping between classes and interfaces.

⁴ For readability, we let L be either a class name \mathbf{C} or an interface name \mathbf{I} .

$$\begin{array}{c}
\text{tmatch}(T, T) = id \qquad \text{tmatch}(r, r) = id \qquad \text{tmatch}(V, T) \triangleq [V \mapsto T] \\
\hline
\forall i. \text{tmatch}(T_i, T'_i) = \sigma_i \quad \forall i, j, \sigma_i |_{\text{dom}(\sigma_j)} = \sigma_j |_{\text{dom}(\sigma_i)} \qquad \text{tmatch}(r, r') = \sigma \\
\text{tmatch}(D(\overline{T}), D(\overline{T}')) \triangleq \bigcup_i \sigma_i \qquad \text{tmatch}((\mathbf{I}, r), (\mathbf{I}, r')) \triangleq \sigma \\
\hline
\forall i. \text{tmatch}(T_i, T'_i) = \sigma_i \quad \forall i, j, \sigma_i |_{\text{dom}(\sigma_j)} = \sigma_j |_{\text{dom}(\sigma_i)} \quad \forall i, \sigma(\mathbf{G}) \in \{\mathbf{G}, \mathbf{G}'\} \\
\text{tmatch}(\mathbf{G}[\overline{f} : \overline{T}], \mathbf{G}'[\overline{f} : \overline{T}']) \triangleq [\mathbf{G} \mapsto \mathbf{G}'] \bigcup_i \sigma_i \\
\hline
\text{pmatch}(_, T) \triangleq \emptyset \qquad \text{pmatch}(x, T) \triangleq \emptyset; x : T \qquad \text{pmatch}(\mathbf{null}, (\mathbf{I}, r)) \triangleq \emptyset \\
\Gamma(\mathbf{Co}) = \overline{T} \rightarrow T' \quad \text{tmatch}(T', T'') = \sigma \quad \forall i. \text{pmatch}(p_i, \sigma(T_i)) = \Gamma_i \\
\text{pmatch}(\mathbf{Co}(\overline{p}), T'') \triangleq \biguplus_i \Gamma_i \\
\hline
\mathbf{C} \subseteq \mathbf{I} \in \mathbf{CT} \quad \text{dom}(\sigma') \cap \text{dom}(\sigma) = \emptyset \quad \text{fields}(\mathbf{C}) = (\overline{f} : (\mathbf{I}, r); f' : \mathbf{D}(\dots)) \\
(\mathbf{I}, \mathbf{G}[\sigma \uplus \sigma'(\overline{f} : (\mathbf{I}, r))]) \in \text{crec}(\mathbf{G}, \mathbf{C}, \sigma) \\
\hline
\text{equals}(\mathbf{G}, \mathbf{G}') \\
\text{coloc}(\mathbf{G}[\dots], (\mathbf{C}, \mathbf{G}'[\dots])) \\
\hline
\text{ports}(\mathbf{C}) \subseteq \text{ports}(\mathbf{I}) \text{ and } \forall p \in \text{ports}(\mathbf{C}). \text{ptype}(p, \mathbf{C}) \leq \text{ptype}(p, \mathbf{I}) \\
\text{heads}(\mathbf{I}) \subseteq \text{heads}(\mathbf{C}) \text{ and } \forall \mathbf{m} \in \mathbf{I}. \text{mtype}(\mathbf{m}, \mathbf{I}) = \text{mtype}(\mathbf{m}, \mathbf{C}) \\
\hline
\text{implements}(\mathbf{C}, \mathbf{I}) \\
\hline
\text{ports}(\mathbf{I}) \subseteq \text{ports}(\mathbf{I}') \text{ and } \forall p \in \text{ports}(\mathbf{I}). \text{ptype}(p, \mathbf{I}) \leq \text{ptype}(p, \mathbf{I}') \\
\text{heads}(\mathbf{I}') \subseteq \text{heads}(\mathbf{I}) \text{ and } \forall \mathbf{m} \in \mathbf{I}'. \text{mtype}(\mathbf{m}, \mathbf{I}) = \text{mtype}(\mathbf{m}, \mathbf{I}') \\
\hline
\text{extends}(\mathbf{I}, \mathbf{I}')
\end{array}$$

Fig. 6. Auxiliary functions and predicates.

3.2 Functions and Predicates

The type system makes use of several auxiliary functions and predicates presented in Fig. 6⁵. Function *tmatch* returns a substitution σ of the formal parameters to the actual ones. It is defined both on types and on records. The matching of a type T to itself, or of a record r to itself, returns the identity substitution *id*; the matching of a type variable V to a type T returns a substitution of V to T ; the matching of data type D parameterized on formal types \overline{T} and on actual types \overline{T}' returns the union of substitutions that correspond to the matching of each type T_i with T'_i in such a way that substitutions coincide when applied to the same formal types; the matching of records follows the same idea as that of data types. Finally, *tmatch* applied on types (\mathbf{I}, r) , (\mathbf{I}, r') returns the same substitution obtained by matching r with r' . Function *pmatch*, performs matchings on patterns and types by returning a typing environment Γ . In particular, *pmatch* returns an empty set when the pattern is $_$ or **null**, or $x : T$ when applied on a variable x and a type T . Otherwise, if applied to a constructor expression $\mathbf{Co}(\overline{p})$ and a type T'' it returns the union of typing environments corresponding to patterns in \overline{p} . Function *crec* asserts that $(\mathbf{I}, \mathbf{G}[\sigma \uplus \sigma'(\overline{f} : (\mathbf{I}, r))])$ is a member of $\text{crec}(\mathbf{G}, \mathbf{C}, \sigma)$ if class \mathbf{C} implements interface \mathbf{I} and σ' and σ are substitutions defined on disjoint sets of names. Function *fields*(\mathbf{C}) returns the

⁵ For readability reasons, the lookup functions like *ports*, *fields*, *ptype*, *mtype*, *heads* are written in italics, whether the auxiliary functions and predicates are not.

$$\begin{array}{c}
\text{T:VAR/FIELD} \\
\frac{\Gamma(x) = T}{\Gamma \vdash x : T} \\
\\
\text{T:CONSTRUCTOR} \\
\frac{\Gamma(\text{Co}) = \bar{T} \rightarrow T' \quad \text{tmatch}(\bar{T}, \bar{T}') = \sigma \quad \Gamma \vdash \bar{v} : \bar{T}'}{\Gamma \vdash \text{Co}(\bar{v}) : \sigma(T')} \\
\\
\text{T:FUN} \\
\frac{\Gamma(\text{fun}) = \bar{T} \rightarrow T' \quad \text{tmatch}(\bar{T}, \bar{T}') = \sigma \quad \Gamma \vdash \bar{v} : \bar{T}'}{\Gamma \vdash \text{fun}(\bar{v}) : \sigma(T')} \\
\\
\text{T:CASE} \\
\frac{\Gamma \vdash e : (T, r) \quad \Gamma \vdash \bar{p} \Rightarrow \bar{e}_p : (T', r')}{\Gamma \vdash \text{case } e \{ \bar{p} \Rightarrow \bar{e}_p \} : (T', r')} \\
\\
\text{T:BRANCH} \quad \text{T:SUB} \\
\frac{\Gamma \vdash p : (T, r) \quad \Gamma; \text{pmatch}(p, (T, r)) \vdash e_p : (T', r')}{\Gamma \vdash p \Rightarrow e_p : (T, r) \rightarrow (T', r')} \quad \frac{\Gamma \vdash e : T \quad T \leq T'}{\Gamma \vdash e : T'} \\
\\
\text{T:FGUARD} \quad \text{T:CGUARD} \quad \text{T:LGUARD} \\
\frac{\Gamma \vdash x : \text{Fut}\langle T \rangle}{\Gamma \vdash x? : \text{Bool}} \quad \frac{\Gamma \vdash x : (I, r)}{\Gamma \vdash \|x\| : \text{Bool}} \quad \frac{\Gamma \vdash g_1 : \text{Bool} \quad \Gamma \vdash g_2 : \text{Bool}}{\Gamma \vdash g_1 \wedge g_2 : \text{Bool}}
\end{array}$$

Fig. 7. Typing Pure Expressions and Guards.

typed fields and ports of a class C . Function *port* instead, returns only the typed ports. Predicate *coloc* states the equality of two cog names. Predicates *implements* and *extends* check when a class implements an interface and an interface extends another one properly. A class C implements an interface I if the ports of C are at *most* the ones of I . This follows the intuition: since ports indicate services then an object has at most the services declared in its interface. Then, any port in C has a subtype of the respective port in I . Instead, for methods, C may define at *least* the methods declared in I having the same signature. The *extends* predicate states when an interface I properly extends another interface I' and it is defined similarly to the *implements* predicate.

3.3 Typing Rules

In this section we present the typing rules. Typing judgments use a typing environment Γ and possibly a set \mathcal{G} which indicates the set of new cogs created by the term being typed. They have the following forms: $\Gamma \vdash g : \text{Bool}$ for guards, $\Gamma \vdash e : T$ for pure expressions, $\Gamma, \mathcal{G} \vdash z : T$ for expressions with side effects and $\Gamma, \mathcal{G} \vdash s$ for statements. Finally, typing judgments for method, class and interface declarations are $\Gamma \vdash M$, $\Gamma \vdash C$ and $\emptyset \vdash I$, respectively.

Pure Expressions. Typing rules for pure expressions are given in Fig. 7. Rule **T:VAR/FIELD** states that a variable is of type the one assumed in the typing environment. Rule **T:NULL** states that **null** is of type any interface I declared in the CT and any record r . Rule **T:CONSTRUCTOR** states that constructor Co applied to a list of values \bar{v} is of type $\sigma(T')$ where the constructor is of a functional type $\bar{T} \rightarrow T'$ and the values are of type \bar{T}' obtained by the auxiliary function *tmatch*. Rule **T:FUN** for function expressions is the same as the previous one for constructor expressions. Rule **T:CASE** states that if all branches are well-typed and have the same type, then the case expression is also well-typed. Rule

$$\begin{array}{c}
\text{T:EXP} \\
\frac{\Gamma \vdash e : T}{\Gamma, \emptyset \vdash e : T} \\
\\
\text{T:NEW} \\
\frac{\text{params}(\mathbf{C}) = \overline{T} \overline{f} \quad \Gamma \vdash e : \overline{T}' \quad \frac{\Gamma(\mathbf{this}) = (\mathbf{C}', \mathbf{G}[\dots])}{\text{tmatch}(\overline{T}, \overline{T}') = \sigma} \quad T \in \text{crec}(\mathbf{G}, \mathbf{C}, \sigma)}{\Gamma \vdash \mathbf{new} \ \mathbf{C}(\overline{e}) : T} \\
\\
\text{T:COG} \\
\frac{\text{params}(\mathbf{C}) = \overline{T} \overline{f} \quad \Gamma \vdash \overline{e} : \overline{T}' \quad \text{tmatch}(\overline{T}, \overline{T}') = \sigma \quad T \in \text{crec}(\mathbf{G}, \mathbf{C}, \sigma)}{\Gamma, \{\mathbf{G}\} \vdash \mathbf{new} \ \mathbf{cog} \ \mathbf{C}(\overline{e}) : T} \\
\\
\text{T:SCALL} \\
\frac{\text{mtype}(\mathbf{m}, \mathbf{I}) = (\mathcal{G}, \mathbf{r})(\overline{T} \overline{x}) \rightarrow T \quad \Gamma \vdash e : (\mathbf{I}, \sigma(\mathbf{r})) \quad \Gamma \vdash \overline{e} : \overline{\sigma(T)} \quad \text{coloc}(\sigma(\mathbf{r}), \Gamma(\mathbf{this}))}{\Gamma \vdash e.m(\overline{e}) : \sigma(T)} \\
\\
\text{T:ACALL} \\
\frac{\text{mtype}(\mathbf{m}, \mathbf{I}) = (\mathcal{G}, \mathbf{r})(\overline{T} \overline{x}) \rightarrow T \quad \Gamma \vdash e : (\mathbf{I}, \sigma(\mathbf{r})) \quad \Gamma \vdash \overline{e} : \overline{\sigma(T)}}{\Gamma \vdash e!m(\overline{e}) : \text{Fut}(\sigma(T))} \\
\\
\text{T:GET} \\
\frac{\Gamma \vdash e : \text{Fut}\langle T \rangle}{\Gamma \vdash \mathbf{get}(e) : T}
\end{array}$$

Fig. 8. Typing Expressions.

T:BRANCH states that a branch $p \Rightarrow e_p$ is well-typed if the pattern p is well-typed and the expression e_p is well-typed in the extension of Γ with typing assertions for the pattern. Rule **T:SUB** is the standard subsumption rule.

Guards. Typing rules for guards are given in Fig. 7. Rule **T:FGUARD** states that if a variable x has type $\text{Fut}\langle T \rangle$, the guard $x?$ has type **Bool**. Rule **T:CGUARD** states that $\|x\|$ has type **Bool** if x is an object. Rule **T:LGUARD** states that if each g_i has type **Bool** for $i = 1, 2$ then the conjunction $g_1 \wedge g_2$ has type **Bool**.

Expressions. The typing rules for expressions with side effects are given in Fig. 8. These are different w.r.t. the previous ones as they keep track of the new cogs created. Rule **T:EXP** is a weakening rule which asserts that a pure expression e is well-typed in a typing context Γ and an empty set of cogs, if it is well-typed in Γ . Rule **T:NEW** assigns type T to the object $\mathbf{new} \ \mathbf{C}(\overline{e})$ if the actual parameters have types compatible with the formal ones, by applying function tmatch , the cogs of the object and \mathbf{this} coincide and the type T is in the crec predicate. Rule **T:COG** is similar to the previous one, except for the creation of a new cog \mathbf{G} where the new object is placed. Rules **T:SCALL** and **T:ACALL** type synchronous and asynchronous method invocations, respectively. Both rules use mtype to obtain the method signature as well as the method's typed parameters and the return type, i.e., $(\mathcal{G}, \mathbf{r})(\overline{T} \overline{x}) \rightarrow T$. The group record \mathbf{r} , the parameters types and the return type of the method are the “formal” ones. In order to obtain the “actual” ones, we use σ that maps formal cog names to actual cog names. Consequently, the callee e has type $(\mathbf{I}, \sigma(\mathbf{r}))$ and the actual parameters \overline{e} have types $\overline{\sigma(T)}$. Finally, the invocations are typed in the substitution $\sigma(T)$. The rules differ in that the former also checks whether the group of \mathbf{this} and the group of the callee coincide, by using the auxiliary function coloc , and also the types of the returned value are $\sigma(T)$ and $\text{Fut}\langle \sigma(T) \rangle$, respectively. Rule **T:GET** states that $\mathbf{get}(e)$ is of type T , if expression e is of type $\text{Fut}\langle T \rangle$.

Statements. The typing rules for statements are presented in Fig. 9. Rule **T:SKIP** states that \mathbf{skip} is always well-typed. Rule **T:DECL** states that $T \ x$ is

$$\begin{array}{c}
\text{T:SKIP} \\
\Gamma, \emptyset \vdash \mathbf{skip} \\
\hline
\text{T:DECL} \\
\frac{\Gamma(x) = T}{\Gamma, \emptyset \vdash T x} \\
\text{T:SEMI} \\
\frac{\Gamma, \mathcal{G}_1 \vdash s_1 \quad \Gamma, \mathcal{G}_2 \vdash s_2}{\Gamma, \mathcal{G}_1 \uplus \mathcal{G}_2 \vdash s_1; s_2} \\
\text{T:ASSIGN} \\
\frac{\Gamma(x) = T \quad \Gamma, \mathcal{G} \vdash z : T}{\Gamma, \mathcal{G} \vdash x = z} \\
\text{S:AWAIT} \\
\frac{\Gamma \vdash g : \mathbf{Bool}}{\Gamma, \emptyset \vdash \mathbf{await } g} \\
\text{S:COND} \\
\frac{\Gamma \vdash e : \mathbf{Bool} \quad \Gamma, \mathcal{G}_1 \vdash s_1 \quad \Gamma, \mathcal{G}_2 \vdash s_2}{\Gamma, \mathcal{G}_1 \uplus \mathcal{G}_2 \vdash \mathbf{if } e \{ s_1 \} \mathbf{else } \{ s_2 \}} \\
\text{S:WHILE} \\
\frac{\Gamma \vdash e : \mathbf{Bool} \quad \Gamma, \mathcal{G} \vdash s}{\Gamma, \mathcal{G} \vdash \mathbf{while } e \{ s \}} \\
\text{S:RETURN} \\
\frac{\Gamma \vdash e : T \quad \Gamma(\mathbf{destiny}) = T}{\Gamma, \emptyset \vdash \mathbf{return } e} \\
\text{REBIND} \\
\frac{\Gamma(\mathbf{this}) = (\mathbf{c}, \mathbf{g}[\dots]) \quad T p \in \mathit{ports}(\mathbf{I}) \quad \Gamma \vdash e : (\mathbf{I}, \mathbf{r}) \quad \Gamma, \mathcal{G} \vdash z : T \quad \mathit{coloc}(\mathbf{r}, \Gamma(\mathbf{this}))}{\Gamma, \mathcal{G} \vdash \mathbf{rebind } e.p = z}
\end{array}$$

Fig. 9. Typing Statements.

$$\begin{array}{c}
\text{T:METHOD} \\
\frac{\Gamma, \bar{x} : \bar{T}, \mathbf{destiny} : \mathbf{Fut}(T), \mathbf{this} : (\mathbf{C}, \mathbf{r}), \mathcal{G} \vdash s}{\Gamma \vdash [\mathbf{critical}] (\mathcal{G}, \mathbf{r}) T \mathbf{m}(\bar{T} x) \{ s \} \mathit{in } \mathbf{C}} \\
\text{T:CLASS} \\
\frac{\forall \mathbf{I} \in \bar{\mathbf{I}}. \mathit{implements}(\mathbf{C}, \mathbf{I}) \quad \Gamma, \bar{x} : \bar{T} \vdash \bar{M} \mathit{in } \mathbf{C}}{\Gamma \vdash \mathbf{class } \mathbf{C} (\bar{T} x) \mathbf{implements } \bar{\mathbf{I}} \{ \bar{F} \bar{L} \bar{M} \}} \\
\text{T:INTERFACE} \\
\frac{\forall \mathbf{I}' \in \bar{\mathbf{I}}. \mathit{extends}(\mathbf{I}, \mathbf{I}')}{\vdash \mathbf{interface } \mathbf{I} \mathbf{extends } \bar{\mathbf{I}} \{ \mathbf{port } T x; \bar{S} \}}
\end{array}$$

Fig. 10. Typing Declarations.

well-typed if variable x is of type T in Γ . Rule T:SEMI types the composition of statements, if s_1 and s_2 are well-typed in the same typing environment and, like in linear type systems, they use distinct cog names. Hence, their composition uses the disjoint union \uplus of the corresponding sets. Rule T:ASSIGN asserts the well-typedness of the assignment $x = z$ if both x and z have the same type T . Rule T:AWAIT asserts that **await** g is well-typed whenever the guard g has type **Bool**. Rules T:COND and T:WHILE are quite standard, except for the presence of the linear set of cog names. Rule T:RETURN asserts that **return** e is well-typed if expression e has the same type as the variable **destiny**. Finally, rule T:REBIND types statement **rebind** $e.p = z$ by checking that: *i*) p is a port of the right type, and *ii*) z is in the same group as **this**.

Method, class and interface declarations. The typing rules are presented in Fig. 10. Rule T:METHOD states that method \mathbf{m} is well-typed in class \mathbf{C} if the method's body s is well-typed in a typing environment augmented with the method's typed parameters; type information about **destiny** and the current object **this**; and cog names as specified by the method signature. Rule T:CLASS states that a class \mathbf{C} is well-typed when it implements all the interfaces $\bar{\mathbf{I}}$ and all its methods are well-typed. Rule T:INTERFACE states that an interface \mathbf{I} is well-typed if it extends all interfaces in $\bar{\mathbf{I}}$.

Remark. The typing rule for assignment requires the group of the variable and the group of the expression being assigned to be the same. This restriction applies to rule for rebinding, as well. To see why this is needed let us consider

$$\frac{\Gamma(\mathbf{this}) = (\mathbf{Controller}, \mathbf{G}[\dots]) \quad (\mathbf{Server}, \mathbb{r}) \quad s \in \mathit{ports}(\mathbf{Client}) \quad \forall i = 2, \dots, n \quad \Gamma \vdash c_i : (\mathbf{Client}, \mathbf{G}[\dots, s : (\mathbf{Server}, \mathbb{r})]) \quad \Gamma, \emptyset \vdash s2 : (\mathbf{Server}, \mathbb{r}) \quad \mathit{coloc}(\mathbf{G}[\dots, s : (\mathbf{Server}, \mathbb{r})], \Gamma(\mathbf{this}))}{\forall i \Gamma, \emptyset \vdash \mathbf{rebind} \ c_i.s = s2}$$

Fig. 11. REBIND derivation.

a sequence of two asynchronous method invocations $\mathbf{x}!\mathbf{m}(); \mathbf{x}!\mathbf{n}()$, both called on the same object and both modifying the same field. Say \mathbf{m} does $\mathbf{this.f} = \mathbf{z}_1$ and \mathbf{n} does $\mathbf{this.f} = \mathbf{z}_2$. Because of asynchronicity, there is no way to know the order in which the updates will take place at runtime. A similar example may be produced for the case of rebinding. Working statically, we can either force the two expressions \mathbf{z}_1 and \mathbf{z}_2 to have the same group as \mathbf{f} , or keep track of all the different possibilities, thus the type system must assume for an expression a set of possible objects it can reduce to. In this paper we adopt the former solution, we let the exploration of the latter as a future work. We plan to relax this restriction following a similar idea to the one proposed in [11].

Example Revisited. We now recall the example of the workflow given in Fig. 2 and Fig. 3. We show how the type system works on this example: by applying the typing rule for **rebind** we have the derivation in Fig. 11 for any clients from c_2 to c_n . For client c_1 , if we try to typecheck the rebinding, we would have the following typing judgments in the premise of REBIND:

$$\Gamma(\mathbf{this}) = (\mathbf{Controller}, \mathbf{G}[\dots]) \quad \Gamma, \emptyset \vdash c_1 : (\mathbf{Client}, \mathbf{G}'[\dots, s : (\mathbf{Server}, \mathbb{r})])$$

But then, the predicate $\mathit{coloc}(\mathbf{G}'[\dots, s : (\mathbf{Server}, \mathbb{r})], \Gamma(\mathbf{this}))$ is false, since $\mathit{equals}(\mathbf{G}, \mathbf{G}')$ is false. Then one cannot apply the typing rule REBIND, by thus not typechecking **rebind** $c_1.s = s2$.

3.4 Properties of the type system

In this section we briefly overview the properties of the type system and we outline the runtime system devised in order to provide the proofs of those properties. The full technical treatment with proofs can be found in [9]. Before stating the properties that our type system enjoys, we first introduce the following notions:

Runtime typing environments Δ are obtained by augmenting typing environments Γ with runtime information about objects and futures, namely $\mathbf{o} : (\mathbf{C}, \mathbb{r})$ and $\mathbf{f} : \mathbf{Fut}\langle T \rangle$ where \mathbf{o} and \mathbf{f} are object and future variables, respectively.

Runtime configurations N extend the language with constructs used during execution, mainly with objects. An object $ob(\mathbf{o}, \sigma, K_{\mathbf{idle}}, Q)$ has a name \mathbf{o} ; a substitution σ mapping the object's fields, ports and special variables like **this**, **destiny**, to values; a running process $K_{\mathbf{idle}}$, that is **idle** if the object is idle; and

a queue of *suspended processes* Q . A process K is $\{ \sigma \mid s \}$ where σ maps the local variables to their values and s is a list of statements.

Reduction relation $N \rightarrow N'$ is defined over runtime configurations and follows the definition of such relation in [13, 14].

Runtime judgments are of the form $\Delta, \mathcal{G} \vdash_R N$ meaning that the configuration N is well-typed in the typing context Δ by using a set \mathcal{G} of new cogs.

Our type system enjoys the classical properties of subject reduction and type correction stated in the following.

Theorem 1 (Subject Reduction). *If $\Delta, \mathcal{G} \vdash_R N$ and $N \rightarrow N'$ then $\exists \Delta', \mathcal{G}'$ such that $\Delta' \supseteq \Delta$, $\mathcal{G}' \subseteq \mathcal{G}$ and $\Delta', \mathcal{G}' \vdash_R N'$.*

Proof. The proof is done by induction over the operational semantics rules. \square

Theorem 2 (Correction). *If $\Delta, \mathcal{G} \vdash_R N$, then for all objects $ob(\circ, \sigma, \{ \sigma_k \mid s \}, Q) \in N$ with either $s = \mathbf{rebind} \ x.f_i = x'; s'$ or $s = x.m(\bar{x}); s'$, there exists an object $ob(\circ', \sigma', K_{\mathbf{idle}}, Q') \in N$ such that $\sigma \circ \sigma_k(x) = \circ'$ and $\sigma(\mathbf{cog}) = \sigma'(\mathbf{cog})$. Where \circ defines the composition of substitutions.*

Proof. The proof is done by induction over the structure of N . \square

As a consequence of the previous results, rebinding and synchronous method calls are always performed between objects of the same cog:

Corollary 1. *Well-typed programs do not perform i) an illegal rebinding or ii) a synchronous method call outside the cog.*

4 Conclusions, Future and Related Works

This paper presents a type system for a component-based calculus [14], an extension of ABS [13] with **ports** and **rebind** operations. Ports denote the access point to the functionalities provided by the environment and can be modified by performing a rebind operation. There are two consistency issues involving ports: *i*) ports cannot be modified while in use; this problem is solved in [14] by combining the notions of ports and critical section; *ii*) it is forbidden to modify a port of an object outside the cog; this problem is solved in the present paper by designing a type system that guarantees the above requirement. The type system tracks an object's membership to a certain cog by adopting group records. Rebind statement is well-typed if there is compatibility of groups between objects involved in the operation.

Regarding future work, we want to investigate several directions. First, as discussed in Section 3 our current approach imposes a restriction on assignments, namely, it is possible to assign to a variable/field/port only an object belonging to the same cog. We plan to relax this restriction following an idea similar to the one proposed in [11], where instead of having just one group associated to a variable, it is possible to have a set of groups. Second, we want to deal with runtime misbehavior. For instance, deadlocks are intrinsically related to the semantic model, which requires a component to be in a safe state when rebinded,

thus introducing synchronization points between the rebinder and the rebindee. For this reason deadlocks may arise. How to detect and avoid this kind of misbehavior is left as future work, possibly following [10]. Moreover, in this paper we showed how to use our technique for a very specific safety problem in the context of a component-based language, but we believe the tracking of object/-group identities/memberships is useful for other problems (deadlock detection, race detection, resource consumption) and other settings (business processes and web-services languages). We plan to investigate this direction further.

Related Works. Most component models [2–4, 8] have a notion of component distinct from that of object. The resulting language is structured in two separate layers, one using objects for the main execution of the program and the other using components for the dynamic reconfiguration. This separation makes it harder for the reconfiguration requests to be integrated in the program’s workflow. The component model used in the present paper has a unified description of objects and components by exploiting the similarities between them. This brings several benefits w.r.t. previous component models: *i*) the integration of components and objects strongly simplifies the reconfiguration requests handling, *ii*) the separation of concepts (component and object, port and field) makes it easier to reason about them, for example, in static analysis, and *iii*) ports are useful in the deployment phase of a system by facilitating, for example, the connection to local communication. Various type systems have been designed for components. The type system in [15] categorizes references to be either Near (i.e., in the same cog), Far (i.e., in another cog) or Somewhere (i.e., we don’t know). The goal is to automatically detect the distribution pattern of a system by using the inference algorithm, and also control the usage of synchronous method calls. It is more flexible than our type system since the assignment of values of different cogs is allowed, but it is less precise than our analysis: consider two objects o_1 and o_2 in a cog c_1 , and another one o_3 in c_2 ; if o_1 calls a method of o_3 which returns o_2 , the type system will not be able to detect that the reference is Near. In [1] the authors present a tool to statically analyze concurrency in ABS. Typically, it analyses the concurrency structure, namely the cogs, but also the synchronization between method calls. The goal is to get tools that analyze concurrency for actor-like concurrency model, instead of the traditional thread-based concurrency model. On the other hand, our type system has some similarities with the type system in [5] which is designed for a process calculus with *ambients* [6], the latter roughly corresponding to the notion of components in a distributed scenario. The type system is based on the notion of group which tracks communication between ambients as well as their movement. However, groups in [5] are a “flat” structure whereas in our framework we use group records defined recursively; in addition, the underlying language is a process calculus, whereas ours is a concurrent object-oriented one. As object-oriented languages are concerned, another similar work to ours is the one on *ownership types* [7], where basically, a type consists of a class name and a context representing object ownership: each object owns a context and is owned by the context it resides in. The goal of the

type system is to provide alias control and invariance of aliasing properties, like role separation, restricted visibility etc. [12].