

The logo for MIKADO, featuring the word "MIKADO" in a stylized, bold, black font on a yellow rectangular background.The tagline "Mobile Calculi based on Domains" in a white font on a dark blue rectangular background.

MIKADO Global Computing Project
IST-2001-32222

An Instance of the MIKADO Migration Model

MIKADO Deliverable D1.2.2

Editor : G. BOUDOL (INRIA)

Authors : F. MARTINS, L. SALVADOR, L. LOPES and V. VASCONCELOS

Classification : Public

Deliverable no. : D1.2.2

Reference : RR/WP1/6

Date : February 2005

© INRIA, France Telecom R&D, U. of Florence, U. of Sussex, U. of Lisbon



**Project funded by the European Community under the
“Information Society Technologies” Programme (1998–
2002)**

MiKO, An Instance of the MIKADO Migration Model

Deliverable D1.2.2

G rard Boudol

February 16, 2005

This deliverable consists of a contribution by F. Martins, L. Salvador, L. Lopes and V. Vasconcelos (from the Lisbon site of MIKADO) that introduces a specific instance of the parametric model of migration and mobility defined in the deliverable D1.2.1 (see also [3]). To introduce this contribution, let us first recall the main features of the parametric model.

The motivation for the MIKADO migration model is to provide programming constructs for controlling code mobility that are as independent as possible from the particular language used to program the mobile code. The main idea is then to regard a *domain* (or site, or locality), that mobile code may enter and exit, as a *membrane* enclosing running processes, offering the procedures that have to be called for entering or exiting the domain. The membrane may also have a local state, on which decisions may depend regarding the migrating code. The migration model therefore consists in a syntax for domains, namely $a\{S\}[P]$, together with a few constructs for controlling migration from the membrane. When we write $a\{S\}[P]$, we mean that a is the name of the domain, S is the membrane process (with a local state and public procedures for entering and exiting), and P is the pool of processes that are running inside the domain a . In [2], we introduced a minimal syntax for the membrane part S , offering constructs for moving processes from the membrane to the inside of the domain, for sending messages to another domain in the network, and for creating new domains. In [3], this was presented by means of a transition system, rather than a particular syntax.

In the previous deliverable (see also [3]), we sketched a particular instance of the generic MIKADO migration model, based on the π -calculus, and we gave several examples to illustrate what the model is able to express. In the present deliverable, the authors define a more substantial instance of the generic model, that has been implemented on top of the TyCO programming language [5, 6]. (The implementation aspect of this work should be reported upon as part of the WP3 activity.) Moreover, they define a type system for the resulting programming model (this is also relevant to WP2). The TyCO language is based on a refinement of the π -calculus, where the receiver end of a communication channel is actually a named object, which may react to messages according to various methods. This is especially well suited for programming membranes, which have to offer various means (or “services”) for “visiting” a domain, depending on the purpose of the visit. Some examples are given in the text below, and the ones which were given in [2] can easily be formulated in MiKO syntax.

The model presented here is a refinement of the one we introduced in [2], in several respects: first, the membrane process is explicitly split into two parts, a method part defining the interface of the domain, and a process part that implements the behaviour of the membrane (and its local state, following the π -calculus style – or, for that matter, the CCS paradigm –, where stored data are represented as processes). Another refinement is that the messages sent from the contents of a domain to the enclosing membrane are now explicitly directed to the domain’s membrane, by specifying the domain’s name. A domain could therefore be seen as a channel of a specific kind. A benefit of this formulation of the membrane model is that there is a natural notion of type for a domain (similar to a channel type), namely the type associated with its interface. Finally, although this is not presented here, the MiKO model has been implemented, and a prototype implementation should soon be available. Therefore this contribution is a distinct progress over the generic model introduced in [2].

MiKO

Mikado Konkurrent Objects

Francisco Martins¹ Liliana Salvador² Luís Lopes² Vasco T. Vasconcelos³

January 2005

1 A TyCO-calculus based model

In this section we introduce the syntax and operational semantics of MiKO. MiKO is a distributed higher order instance of TyCO-calculus [5] based on the MIKADO's membrane model [1, 2], distributing processes over a flat network of domains. The migration of code from one domain to another passes always through the domains's membranes. The calculus contains special primitives to handle migration between such domains.

Syntax

Figure 1 describes the syntax of the calculus. Consider a set of names \mathcal{N} that do not possess the constants in, out, and mkdom and a set of labels \mathcal{L} , disjoint from \mathcal{N} , ranged over by l . To improve readability we use $r-t$ to range over names of domains, $a-c$ to range over names of channels, and $x-z$ to range over unknowns.

The calculus is organised in two layers: *networks* and *processes*. A network consists of a set of *domains* and *network messages* running independently in parallel. A domain $x\{m\}P_1\}[P_2]$ is a named location x composed by a *membrane* $m\}P_1$ and a *contents* P_2 . Membranes define the interaction between the domain and the outside world, that is, more precisely, between the contents part of the domain and the rest of the network. The set of public methods m defines the interface of the domain, and process P_1 controls the migration of code with the domain. The contents is the computation oracle of the domain. For the following we consistently denote membranes by $R-T$, processes by P and Q , and networks by H, L , and N . A network message $x!M$ denotes a message M heading to domain x . The remaining networks constructs are standard in process calculi.

The syntax of membranes and contents is mostly the standard syntax of TyCO [4, 6] enriched with constructs $\text{in}[P]$ that launches a process P in the contents part of the domain, $\text{out}[x, M]$ that sends a message M to the domain x , and $\text{mkdom}[x, m\}R, P]$ in Q that creates a new domain x with guardian $m\}R$ and contents P . Messages are tuples of values sent on channels and interpreted as method invocations. The values that can be communicated in messages are either names or process abstractions. In the examples, we shall use the standard TyCO style, writing $l(\vec{x}) = P$ instead of $l = (\vec{x})P$. We also adopt TyCO's convention that $c?A$ stands for $c?\{val = A\}$ (and similarly for $c?*A$), in which case a message $c!val[\vec{V}]$ is simply written $c![\vec{V}]$.

Operational Semantics

The operational semantics of the calculus is presented with the help of congruence relations on networks and processes. The bindings of the calculus are scope restriction $\text{new } xP$, abstraction $(\vec{x})P$, both binding free

¹Universidade dos Açores, Departamento de Matemática

²Universidade do Porto, Departamento de Ciência de Computadores

³Universidade de Lisboa, Departamento de Informática

$N ::= \text{inaction} \mid x\{m\}P\}[P] \mid x!M \mid N \mid N \mid \text{new } xN$	(networks)
$P ::= \text{inaction} \mid x!M \mid x?m \mid x?*m \mid P \mid P \mid \text{new } xP \mid$ $A\vec{V} \mid \text{in}[P] \mid \text{out}[x, M] \mid \text{mkdom}[x, m]P, P \text{ in } P$	(processes)
$M ::= l[\vec{V}]$	(messages)
$m ::= \{l_i = A_i\}_{i \in I}$	(methods)
$A ::= (\vec{x})P$	(abstractions)
$V ::= x \mid A$	(values)

Figure 1: Syntax of MiKO

$\text{new } xN \mid L \equiv \text{new } x(N \mid L)$	$x \notin \text{fn}(L)$	(N-SRC)
$s\{m\}S\}[\text{new } cP] \equiv \text{new } cs\{m\}S\}[P]$	$c \notin \text{fn}(m\}S)$	(N-PSR)
$s\{m\}\text{new } cS\}[P] \equiv \text{new } cs\{m\}S\}[P]$	$c \notin (\text{fn}(P) \cup \text{fn}(m))$	(N-MSR)

Figure 2: Structural congruence on networks

occurrences of x and \vec{x} , in process P , and domain creation $\text{mkdom}[x, m]R, P$ in Q where the domain name x is bound in R, P and Q . The set of free names for networks N , processes P , and methods m is denoted by $\text{fn}(N)$, $\text{fn}(P)$ and $\text{fn}(m)$, respectively.

Structural Congruence

The structural congruence relation on networks, \equiv , is the least congruence relation closed under the rules in figure 2 adding α -conversion (that is, the renaming of bound names) and the commutative monoid rules for parallel composition, with inaction as the neutral element. The structural congruence between processes is similar; it is the standard one for the TyCO calculus [5], and is therefore omitted.

The following example illustrates the operational semantics of the calculus. Consider a client-server session manager: a client that establishes a session with a server and finishes the session at the end of the conversation. The protocol we envisage is very simple: (1) the client issues a connect request to the server; (2) the server replies with a session identifier; (3) the client issues a disconnect message to terminate the session. First we comment on the server's membrane. The server's membrane must provide two public methods to handle connection and disconnection: *connect* and *disconnect*. The *connect* method creates a private channel *sessionID*, sends it to the client via an out operation and launches an object at the server's contents to handle the session using the *in* process. The object that handles the session just implements the *quit* operation. (A more elaborate session handler is presented in the *mathServer* example later in section 2.) The *disconnect* method triggers the end of the session by selecting the *quit* operation from the handler. A possible implementation of the server's membrane is given in figure 3. The client's membrane needs to offer a *connect* and *disconnect* methods, to interact with the server's counterpart. Moreover, it must provide a method (that we name *enter*) to receive the responses from the server. The methods *connect* and *disconnect* just invoke the server operations. Both client and server do not need a state, so the body of the membrane is the *inaction* process. In figure 4 we provide a possible implementation of the client's membrane. An interaction between the client and the server may be written as displayed in the figure 5.

```

servermemb = {
  connect (client, replyTo) =
    new sessionID
    out [client, enter [(replyTo ! sessionID)] |
    in [
      sessionID ? {
        quit() = inaction
      }
    ]
  ]
  disconnect (client, sessionID) =
    in [sessionID ! quit[]]
  }
}
}
{
  inaction
}
}

```

Figure 3: Server's membrane

```


clientmemb = {
  connect (server, replyTo) =
    out [server, connect [myDomain, replyTo]]
  enter (x) =
    in [x []]
  disconnect (server, sessionID) =
    out [server, disconnect [sessionID] ]
  }
}
}
{
  inaction
}
}

```

Figure 4: Client's membrane

Reduction Rules

The reduction rules for networks and processes are given in figures 6 and 7, where we use the notation $m.M$ to denote $A_j.\vec{V}$ when $m = \{l_i = A_i\}_{i \in I}$ and $M = l_j[\vec{V}]$ with $j \in I$. The migration of messages between domains proceeds as depicted below



The N-SEND rule delivers a message from the contents part of a domain to its membrane. The message M initiates the execution of a method from the domain's interface in its membrane. Rules N-OUT and N-COM

```

clientcont =
  new connection
  client ! connect [ server, connection ] |
  connection ? (sessionID) = myDomain ! disconnect [sessionID]

```

Figure 5: Client-server interaction

$$\begin{array}{l}
s\{m\}P\}[s!M | Q] \rightarrow s\{m\}m.M | P\}[Q] \quad (\text{N-SEND}) \\
s\{m\}\text{out}[r, l[\vec{V}]] | S\}[P] \rightarrow r!l[s, \vec{V}] | s\{m\}S\}[P] \quad (\text{N-OUT}) \\
s!M | s\{m\}S\}[P] \rightarrow s\{m\}S | m.M\}[P] \quad (\text{N-COM}) \\
s\{m\}\text{in}[P] | S\}[Q] \rightarrow s\{m\}S\}[Q | P] \quad (\text{N-IN}) \\
s\{m\}\text{mkdom}[r, m'\}S, P] \text{ in } R | T\}[Q] \rightarrow \\
\quad \text{new } r(r\{m'\}S\}[P] | s\{m\}R | T\}[Q]) \quad \text{if } r \notin (\text{fn}(T) \cup \text{fn}(Q)) \quad (\text{N-MKD}) \\
\frac{S \rightarrow T}{s\{m\}S\}[P] \rightarrow s\{m\}T\}[P] \quad (\text{N-MEMB}) \\
\frac{P \rightarrow Q}{s\{m\}S\}[P] \rightarrow s\{m\}S\}[Q] \quad \frac{N \rightarrow L}{N | H \rightarrow L | H} \quad (\text{N-PROC, N-PAR}) \\
\frac{N \rightarrow L}{\text{new } xN \rightarrow \text{new } xL} \quad \frac{N \equiv L \quad N \rightarrow H}{L \rightarrow H} \quad (\text{N-SRES, N-STR})
\end{array}$$

Figure 6: Reduction rules: networks

route the message through the network. Upon reception, the message interacts with the target membrane by selecting a public method. Rule N-IN drives processes from membranes to the contents of the domain.

A domain can only be created from the membrane of an existing domain. To create a domain, rule N-MKD, we provide a domain name (r), a guardian ($m'\}S$), a contents (P) that is going to run in the domain, and a part of the membrane (R) of the creator domain where the new domain is visible. Rules N-MEMB, N-PROC, N-PAR, and N-RESC allow reductions inside the domain, the parallel composition of networks, and channel restriction, respectively. Finally, N-STRC brings structural congruence on networks to the reduction relation.

As for processes, figure 7, the reception of messages by objects (followed by the selection of the appropriate method and the instantiation of the method's body) constitutes the basic communication mechanism of the calculus, as captured by axiom P-COM. The remaining rules are straightforward. As an example of a reduction in MiKO, consider

$$\text{server}\{\text{server}_{\text{memb}}\}[\text{inaction}] | \text{client}\{\text{client}_{\text{memb}}\}[\text{client}_{\text{cont}}]$$

It reduces as follows:

1. The communication between $\text{client}_{\text{cont}}$ and $\text{client}_{\text{memb}}$ is made by the rule N-SEND. The message $\text{client} ! \text{connect} [\text{server}, \text{connection}]$ is sent to the membrane and interacts with the *connect* method.
2. Membrane $\text{server}_{\text{memb}}$ sends messages to the network by the rule N-OUT. The *connect* method sends the message $\text{server} ! \text{connect} [\text{myDomain}, \text{replyTo}]$ to the network towards the server domain.

$$\begin{array}{c}
c?m \mid c!M \rightarrow m.M \quad (\text{P-COM}) \\
c?*m \mid c!M \rightarrow c?*m \mid m.M \quad (\text{P-COMR}) \\
((\vec{x})P)\vec{V} \rightarrow P/\vec{V}\vec{x} \quad \frac{P \rightarrow Q}{P \mid R \rightarrow Q \mid R} \quad (\text{P-SUBS, P-PAR}) \\
\frac{P \rightarrow Q}{\text{new } cP \rightarrow \text{new } cQ} \quad \frac{P \equiv Q \quad P \rightarrow R}{Q \rightarrow R} \quad (\text{P-SRESC, P-STRC})
\end{array}$$

Figure 7: Reduction rules: processes

$$\alpha ::= \{l_i : \vec{\alpha}_i\}_{i \in I} \mid t \mid \mu t. \alpha$$

Figure 8: Syntax of types

3. Network messages interact with $\text{client}_{\text{memb}}$ by the rule N-COM. The message sent to the server enters in the membrane and interacts with the *enter* method.
4. $\text{server}_{\text{memb}}$ interacts with its contents $\text{server}_{\text{cont}}$ by the rule N-IN. The *connect* method when triggered launches the process `replyTo ! sessionID` to the server's contents.
5. By rule P-COM, the process launched in the contents reduces with the object connection `? (sessionID) = myDomain ! disconnect [sessionID]` and the $\text{server}_{\text{cont}}$ sends, by rules N-SEND and N-OUT, a *disconnect* message `server!disconnect[sessionID]` to the client.
6. The message enters in $\text{client}_{\text{memb}}$, reduces with the *disconnect* method and the process `sessionID ! quit[]` is launched to the contents, by rules N-COM, P-COM and P-IN.

2 The Type System

In this section, we define the syntax of types and the type inference system for MiKO.

Syntax

We fix a countable set of type variables ranged over by t , and let α, β range over types. A sequence of types $\alpha_1 \dots \alpha_n$ is abbreviated as $\vec{\alpha}$. The syntax of types is given in figure 8. A type of the form $\{l_i : \vec{\alpha}_i\}_{i \in I}$ describes locations of objects or domains containing n methods labelled $l_1 \dots l_n$ ($n > 0$) with types $\vec{\alpha}_1 \dots \vec{\alpha}_n$. Types are interpreted as rational (regular infinite) trees. A type of the form $\mu t. \alpha$ (with $\alpha \neq t$) denotes the rational tree solution of the equation $t = \alpha$. If α is a type, denote by α^* its associated rational tree. An interpretation of recursive types as infinite trees naturally induces an equivalence relation \uparrow on types, by putting $\alpha \uparrow \beta$ if $\alpha^* = \beta^*$.

Typing rules

Type assignments to names are formulae $x : \alpha$, for x a name and α a type. Typings, denoted by Γ , are finite partial functions from names to types (that is, finite sets of formulae $x : \alpha$ with distinct subjects). We write $\text{dom}(\Gamma)$ for the domain of Γ . When $x \notin \text{dom}(\Gamma)$, we write $\Gamma, x : \alpha$ for the type environment Γ' such that $\text{dom}(\Gamma') = \text{dom}(\Gamma) \cup \{x\}$, $\Gamma'(x) = \alpha$, and $\Gamma'(y) = \Gamma(y)$, for $y \neq x$. So, the expression $\Gamma(x)$ denotes the type α

if $x: \alpha$ and $\Gamma(\vec{x})$ represents the sequence of types $\vec{\alpha}$. We denote by x^{dom} , x^c , and x^\bullet , respectively, identifiers that are domains, channels, and that can be both.

The type system, described in figure 9 includes four kinds of judgements: (a) judgement $\Gamma \vdash N$ asserts that network N is well typed under typing assumptions Γ ; (b) judgement $\Gamma \vdash_s P$ means that process P is running at domain s and is well typed under typing assumptions Γ ; (c) judgement $\Gamma \vdash_s \vec{V}: \vec{\alpha}$ assigns types to sequences of values; and finally (d) judgement $\Gamma \vdash_s m: \alpha$ states that the object (set of methods) m has type α . In the rules $\{l_i: \vec{\alpha}_i\}_{i \in I}.l_j$ denotes $\vec{\alpha}_j$ if $j \in I$.

Networks. To type a domain $s\{m\}S\{P\}$, rule TN-NET, one has to type the membrane $m\}S$ and the contents P . Domain s possesses the type of its interface m . Rule TN-MSG expresses that s is a domain that has a method with the same name as the message heading to it and with a compatible type.

Processes. When typing processes, we record (under the turnstile) the domain that hosts the process. The information is relevant when typing the out operation, since we stamp messages with the name of the sending domain. Rule TP-OBJ expresses the fact that c must be a channel and has type compatible with method m . Rule TP-MSG is similar to TN-MSG, but here x can be a channel or a domain, since we use the same constructor to send messages to membranes or interact with objects. To type the in constructor in domain s , we type P in s using rule TN-IN. Rule TP-OUT types the out constructor in the origin domain s and the sequence of values $s\vec{V}$ in the destination domain r , since abstractions run in the destination domain. (Names are network-wide, so may be typed in any domain.) Notice that one adds the origin domain s to the list of parameters. To type creation of domains, rule TM-MKD, the created domain $r\{m\}S\{P\}$ must be itself well-typed as well as membrane R .

Methods. When typing a method, by rule TM-METH, we type each abstraction with the type of the arguments of the method l_i .

Some properties of the type system

We now present some properties of the type system for MiKO. The following lemmas are proved by straightforward induction on the depth of the derivation of the judgements in the hypothesis.

Lemma 2.1 (Strengthening Lemma)

- i. If $\Gamma, x: \alpha \vdash N$ and $x \notin \text{fn}(N)$, then $\Gamma \vdash N$
- ii. If $\Gamma, x: \alpha \vdash_s P$ and $x \notin \text{fn}(P)$, then $\Gamma \vdash_s P$

Lemma 2.2 (Weakening Lemma)

- i. If $\Gamma \vdash N$, then $\Gamma, x: \alpha \vdash N$
- ii. If $\Gamma \vdash_s P$, then $\Gamma, x: \alpha \vdash_s P$

Lemma 2.3 (Substitution Lemma)

If $\Gamma, \vec{x}: \vec{\alpha} \vdash_s P$ and $\Gamma \vdash_s \vec{V}: \vec{\beta}$ with $\vec{\alpha} \uparrow \vec{\beta}$, then $\Gamma \vdash_s P/\vec{V}\vec{x}$.

Subject-reduction (theorem 2.1) expresses a consistency property between the operational semantics and the typing rules. We prove the theorem by induction on the depth of the derivations $N \rightarrow L$ and $P \rightarrow Q$, respectively.

Theorem 2.1 (Subject-Reduction)

- i. If $\Gamma \vdash N$ and $N \rightarrow L$, then $\Gamma \vdash L$.
- ii. If $\Gamma \vdash_s P$ and $P \rightarrow Q$, then $\Gamma \vdash_s Q$.

Networks: $\Gamma \vdash N$

$$\begin{array}{c}
\text{(TN-NET)} \\
\frac{\Gamma \vdash_s m : \alpha \quad \Gamma \vdash_s P \quad \Gamma \vdash_s Q \quad \Gamma(s^{\text{dom}}) \uparrow \alpha}{\Gamma \vdash_s \{m\}P\{Q\}} \\
\\
\text{(TN-INAC)} \qquad \text{(TN-PAR)} \qquad \text{(TN-RESC)} \\
\frac{}{\Gamma \vdash \text{inaction}} \qquad \frac{\Gamma \vdash N \quad \Gamma \vdash L}{\Gamma \vdash N \mid L} \qquad \frac{\Gamma, x^\bullet : \alpha \vdash N}{\Gamma \vdash \text{new } x N}
\end{array}$$

Processes: $\Gamma \vdash_s P$

$$\begin{array}{c}
\text{(TP-INAC)} \qquad \text{(TP-PAR)} \qquad \text{(TP-OBJ)} \\
\frac{}{\Gamma \vdash_s \text{inaction}} \qquad \frac{\Gamma \vdash_s P \quad \Gamma \vdash_s Q}{\Gamma \vdash_s P \mid Q} \qquad \frac{\Gamma \vdash_s m : \alpha \quad \Gamma(c^c) \uparrow \alpha}{\Gamma \vdash_s c?m} \\
\\
\text{(TP-OBJR)} \qquad \text{(TP-RESC)} \qquad \text{(TP-MSG)} \\
\frac{\Gamma \vdash_s c?m}{\Gamma \vdash_s c?*m} \qquad \frac{\Gamma, c^c : \alpha \vdash_s P}{\Gamma \vdash_s \text{new } c P} \qquad \frac{\Gamma \vdash_s \vec{V} : \vec{\alpha} \quad \Gamma(x^\bullet) \uparrow \vec{\alpha}}{\Gamma \vdash_s x ! l[\vec{V}]} \\
\\
\text{(TP-APP)} \qquad \text{(TP-IN)} \\
\frac{\Gamma \vdash_s A : \vec{\alpha} \quad \Gamma \vdash_s \vec{V} : \vec{\beta} \quad (\vec{\alpha} \uparrow \vec{\beta})}{\Gamma \vdash_s A \vec{V}} \qquad \frac{\Gamma \vdash_s P}{\Gamma \vdash_s \text{in}[P]} \\
\\
\text{(TP-OUT)} \qquad \text{(TP-MKD)} \\
\frac{\Gamma \vdash_r s \vec{V} : \vec{\alpha} \quad \Gamma(r^{\text{dom}}) \uparrow \vec{\alpha}}{\Gamma \vdash_s \text{out}[r, l[\vec{V}]]} \qquad \frac{\Gamma, r^{\text{dom}} : \alpha \vdash r\{m\}S\{P\} \quad \Gamma, r^{\text{dom}} : \alpha \vdash_s R}{\Gamma \vdash_s \text{mkdom}[r, m]S, P \text{ in } R}
\end{array}$$

Values: $\Gamma \vdash_s \vec{V} : \vec{\alpha}$

$$\begin{array}{c}
\text{(TV-VAR)} \qquad \text{(TV-ABS)} \qquad \text{(TV-SEQ)} \\
\frac{}{\Gamma, x^\bullet : \alpha \vdash_s x^\bullet : \alpha} \qquad \frac{\Gamma, \vec{x}^\bullet : \vec{\alpha} \vdash_s P}{\Gamma \vdash_s (\vec{x})P : \vec{\alpha}} \qquad \frac{\Gamma \vdash_s V_1 : \alpha_1 \dots \Gamma \vdash_s V_n : \alpha_n}{\Gamma \vdash_s \vec{V} : \vec{\alpha}}
\end{array}$$

Methods: $\Gamma \vdash_s m : \alpha$

$$\text{(TM-METH)} \\
\frac{\forall i \in I, \quad \Gamma \vdash_s A_i : \vec{\alpha}_i}{\Gamma \vdash_s \{l_i = A_i\}_{i \in I} : \{l_i : \vec{\alpha}_i\}_{i \in I}}$$

Figure 9: Typing rules

```

mathServer_memb = {
  connect (client, replyTo) =
    controller ! connect [client, replyTo]
  eval (client, x) =
    in [x [] ]
  replyResult (client, x) =
    out [client, enter [x]]
  disconnect (client, sessionID) =
    controller ! disconnect [sessionID]
}
|
new localStatus
localStatus ?* (availSessions) =
  controller ? {
    connect (client, replyTo) =
      if availSessions > 0 then
        sessionIDManager |
        localStatus ! [availSessions - 1]
      else
        localStatus ! [availSessions]
    disconnect (sessionID) =
      in [sessionID ! disconnect [] ] |
      localStatus ! [availSessions + 1]
  } |
  localStatus ! [5]
}

```

Figure 10: Math server's membrane

Example. Our second example illustrates a mathematical server, math server for short. A math server accepts computation requests from clients, compute the data, and reply back the answers. To build the math server, we elaborate upon the first server example we showed. Therefore, the client establishes a session with the server and issues repeatedly computation requests. Upon completion it closes the session with the server. The server's membrane now offers two additional methods: *eval* and *replyResult*. The *eval* method accepts computation requests and routes them to the respective session handler that sits in the computational area of the math server. The *replyResult* method delivers the computation results to the client. Our math server accounts for the number of concurrent sessions it handles, illustrating a membrane with a state associated. Therefore, the *connect* operation consumes one session, restored when a *disconnect* operation occurs. The state is recorded via a private replicated object *localStatus*, updated via an object that offers two methods called *connect* and *disconnect* to easily identify them with the domain's public methods. The math server membrane may be programmed as shown in figure 10.

Notice that the *connect* and *disconnect* methods redirect their requests to update the number of sessions handled concurrently by the server. The number of simultaneous sessions starts with five. Then, whenever a connection request arrives, the server verifies if there are still resources available. If there are, a new session is created and the number of available sessions decreased, otherwise the request is simply discarded. When disconnecting from the math server, the number of available resources is restored.

The *sessionIDManager* is responsible for generating a new session identifier and for setting a session handler at the contents part of the math server. Since the client may issue repeated requests, the session handler needs to be a persistent object. The object instantiates itself until a disconnect message is received.

Every time a new session is created, the `sessionHandler` is launched in the math server's contents where it is responsible for computing the mathematical operations selected by the client and for sending the results back to it. The client sends a message to the server that, after being filtered by the membrane, selects the operation method in the corresponding object. The `sessionIDManager` is shown in figure 11

```

sessionIDManager =
  new sessionID
  out [client, enter [() replyTo ! [sessionID] ] ] |
  in [
    new sessionHandler
    sessionHandler ?* (self, client) =
      self ? {
        add (n, m, replyTo) =
          mathServer ! replyResult [client, () replyTo ! [n + m]] |
          sessionHandler ! [self, client]
        neg (n, replyTo) =
          mathServer ! replyResult [client, () replyTo ! [0 - n]] |
          sessionHandler ! [self, client]
        disconnect () =
          inaction
      } |
    sessionHandler ! [sessionID, client]
  ]

```

Figure 11: Session manager

The client's membrane provides the same `connect`, `enter`, and `disconnect` methods as in the first example. Additionally, there is a new method `eval` that sends a computation to be executed by the math server. As before, there is no status associated with the client's membrane. The client's membrane is given in figure 12.

To interact with the math server, we provide a process that establishes a session and then requests the addition of two values and the negation of the result of the previous operation. Finally, after receiving from the server the result of the negation, it closes the session with the math server. The process is shown in figure 13.

```

mathClientmemb = {
  connect (server, replyTo) =
    out [ server, connect [replyTo] ]
  enter (server, x) =
    in [ x [] ]
  eval (server, x) =
    out [server, eval [x]]
  disconnect (server, sessionID) =
    out [server, disconnect [sessionID] ]
}
|
inaction
}

```

Figure 12: Math client's membrane

```

mathClientcont =
  new replyTo
  client ! connect [ mathServer, replyTo ] |
  replyTo ? (sessionID) =
    new result
    client ! eval [ mathServer, () sessionID ! add [3, 4, result] ] |
    result ? (x) =
      client ! eval [ mathServer, () sessionID ! neg [x, result]] |
      result ? (x) =
        io ! printi [x] |
        client ! disconnect [mathServer, sessionID]

```

Figure 13: Math client

References

- [1] G. Boudol. Core programming model, release 0. Mikado Deliverable D1.2.0, 2002.
- [2] G. Boudol. A parametric model of migration and mobility, release 1. Mikado Deliverable D1.2.1, 2003.
- [3] G. Boudol. A generic membrane model. In *Second Global Computing Workshop*, 2004. To appear in *Lecture Notes in Computer Science*, Vol. 3267.
- [4] Vasco T. Vasconcelos. Core-tyco appendix to the language definition, version 0.2. Technical report, Faculty of Sciences of the University of Lisbon, 2001.
- [5] Vasco T. Vasconcelos. Tyco gently. Technical report, Faculty of Sciences of the University of Lisbon, 2001.
- [6] Vasco T. Vasconcelos and Rui Bastos. Core-tyco the language definition, version 0.1. Technical report, Faculty of Sciences of the University of Lisbon, 1998.