

# Resource Access Control in Systems of Mobile Agents

(Extended Abstract)

Matthew Hennessy<sup>1</sup>

*School of Cognitive and Computing Sciences*  
*Univ. of Sussex*  
*Brighton, UK*  
matthewh@cogs.susx.ac.uk

James Riely<sup>2</sup>

*Department of Computer Science*  
*North Carolina State Univ.*  
*Raleigh, NC, USA*  
riely@csc.ncsu.edu

---

## Abstract

We describe a typing system for a distributed  $\pi$ -calculus which guarantees that distributed agents cannot access the *resources* of a system without first being granted the *capability* to do so. The language studied allows agents to move between distributed *locations* and to augment their set of capabilities via communication with other agents. The type system is based on the novel notion of a *location type*, which describes the set of resources available to an agent at a location. Resources are themselves equipped with capabilities, and thus an agent may be given permission to send data along a channel at a particular location without being granted permission to read data along the same channel. We also describe a *tagged* version of the language, where the capabilities of agents are made explicit in the syntax. Using this tagged language we define *access violations* as runtime errors and prove that well-typed programs are incapable of such errors.

---

<sup>1</sup> Research funded by CONFER II and EPSRC project GR/K60701.

<sup>2</sup> Research funded by NSF grant EIA-9805604 and EPSRC project GR/K60701.

*This is a preliminary version. The final version can be accessed at*  
URL: <http://www.elsevier.nl/locate/entcs/volume16.html>

## 1 Introduction

*Mobile computation*, where independent agents roam widely distributed networks in search of resources and information, is fast becoming a reality. A number of programming languages, APIs and protocols have recently emerged which seek to provide high-level support for mobile agents. These include Java [24], Odyssey [11], Aglets [15], Voyager [20] and the latest revisions of the Internet protocol [21,1]. In addition to these commercial efforts, many prototype languages have been developed and implemented within the programming language research community — examples include Linda [5,6], Facile [12], Obliq [4], Infospheres [8], and the join calculus [9]. In this paper we address the issue of resource access control for such languages.

Central to the paradigm of mobile computation are the notions of agent, resource and location. *Agents* are effective entities that perform computation and interact with other agents. Interaction is achieved using shared *resources* such as memory cells, M-structures, objects (with shared methods and state) or communication channels. The use of the term “mobile” implies that agents are bound to particular *locations* and that this binding may vary over time, *i.e.* agents can *move*. Resources, on the other hand, are often fixed to a single location, although proxies and mirrors may be set up in order to distribute their contents.

In *open* distributed systems, such as the internet, it is unwise to assume that all agents are benign, and thus a certain amount of effort must be spent to ensure that vital resources are protected from unauthorized access. This can be accomplished by using a system of *capabilities* and by predicating resource access on possession of the appropriate capability. It is unreasonable, however, to expect that *every* use of *every* resource in a network be thus verified dynamically; such a requirement surely would degrade performance unacceptably. Thus it is attractive to develop static analyses, or *typing systems* that guarantee controlled access to network resources.

We present a typed language for mobile agents which allows fine control over the use of resources in a network. We also define a *tagged* version of the language in which agents explicitly carry the sets of capabilities which they have acquired. Using this tagged language, we capture resource access violations as *runtime errors* and show that well-typed terms are incapable of such errors.

The language studied in this paper, called  $D\pi$ , is a distributed variant of the  $\pi$ -calculus [19], and thus the *resources* of interest are *channels* which support binary communication between agents. We take agents to be located *threads*, which are simply terms of the ordinary polyadic  $\pi$ -calculus [18], extended with primitives for *movement* between locations and for the creation of *new* locations.

The type system is based on the notion of *location types* of the form:

$$\text{loc}\{\kappa_1, \dots, \kappa_n\}$$

where each  $\kappa_i$  is a *location capability*. These may take the form of primitive capabilities, typical examples being *go*, the ability to move to the location, and *newc*, the ability to create a new local channel; or a capability associated with a particular

channel,  $a:A$ . Here  $A$  represents a set of *channel capabilities* which are of the form

- $r\langle T \rangle$ , the capability to receive values  $V$  from a channel and then to use each  $V$  with *at most* the capabilities specified by the type  $T$ , or
- $w\langle T \rangle$ , the capability to send values  $V$  into a channel, as long as that agent has, on each  $V$  sent, *at least* the capabilities specified by  $T$ .

Agents may restrict access to a resource by controlling the type of the channel over which the name of the resource is sent. Thus if an agent sends the name of a location,  $\ell$ , over a channel of type  $\text{res}\{r\langle \text{loc}\{a:A, b:B\} \rangle\}$ , then the recipient gains capabilities on the channels  $a$  and  $b$  at  $\ell$ , as specified by the capability sets  $A$  and  $B$  respectively. Instead, when the same name is communicated over a channel of type  $\text{res}\{r\langle \text{loc}\{a:A'\} \rangle\}$ , the recipient gains access only to channel  $a$  at  $\ell$ , with permissions determined by  $A'$ .

The remainder of this extended abstract is organized as follows. In the next section we define the language  $D\pi$  and its reduction semantics; we also describe several examples that highlight some of the main features of the language. The following section gives a description of the typing system and states a Subject Reduction Theorem; we show the application of the typing system to an example program, a *cell server*. In Section 4 we give an informal outline the Type Safety Theorem, which formalizes the idea that well-typed networks can not misuse resources. We end with a brief comparison with related work.

In this extended abstract all proofs are omitted, as are many other details, including some definitions. The reader is referred to the published technical report [14] for a full account.

## 2 The Language

A typical  $D\pi$  network is the following:

$$\ell[[P]] \mid (\nu_k a:A) (\ell[[Q]] \mid k[[R]])$$

There are three agents running in parallel:  $\ell[[P]]$  and  $\ell[[Q]]$  running at location  $\ell$  and  $k[[R]]$  running at location  $k$ . Moreover  $Q$  and  $R$  share a private channel  $a$ , declared at location  $k$ . The syntax for the agents is a mild extension of that of the  $\pi$ -calculus; *structured values* may be exchanged along channels, and there is a new command for *code movement* ‘ $\text{go } \ell.P$ ’, which causes the agent to move to location  $\ell$  and then execute  $P$ .

The syntax of the language is given in Table 2, where letters  $a$ – $m$  range over the set *Name* of names,  $x$ – $z$  range over the disjoint set *Var* of variables, and  $u$ – $w$  range over *identifiers* in  $\text{Name} \cup \text{Var}$ . The syntax for types  $T$  will be explained in the following section. Types appear in the binders for variables and names: the input construct ‘ $u?(X:T)Q$ ’ binds each of the variables in the pattern  $X$ ; the restriction constructs ‘ $(\nu e:T)P$ ’ and ‘ $(\nu_k e:T)N$ ’ bind the name  $e$ .

The typing system will distinguish the location of resources, leading us to define

**Table 1** Syntax of patterns  $X$ , values  $U$ , threads  $P$  and networks  $M$ 

$X, Y ::= x$	Variable	$U, V ::= u$	Id
$(X_1, \dots, X_n)$	Tuple	$(U_1, \dots, U_n)$	Tuple
$P, Q, R ::= \text{stop}$	Termination	$M, N ::= \mathbf{0}$	Empty
$P \mid Q$	Composition	$M \mid N$	Composition
$(\nu e:T)P$	Restriction	$(\nu_k e:T)N$	Restriction
$\text{go } u.P$	Movement	$k[[P]]$	Agent
$u!\langle V \rangle P$	Output		
$u?(X:T)P$	Input		
$*P$	Replication		
$\text{if } U = V \text{ then } P \text{ else } Q$	Matching		

dependent tuple types. In examples, we use the notation

$$u[v_1, \dots, v_n] \stackrel{\text{def}}{=} (u, v_1, \dots, v_n)$$

to indicate that the identifiers  $v_i$  refer to resources at location  $u$ .

The reduction semantics is defined as a reduction relation between networks; thus judgments are of the form  $M \longrightarrow M'$  where  $M$  and  $M'$  are (closed) network terms, *i.e.* terms which contain no free occurrences of variables. The semantics is given in Table 2 using two relations; a structural equivalence  $\equiv$  and a primitive reduction relation  $\longmapsto$ . The main relation of interest is:

$$(\longrightarrow) \stackrel{\text{def}}{=} (\equiv \cdot \longmapsto \cdot \equiv)$$

The primitive reduction relation is defined to be the least precongruence on networks that satisfies the axioms given in Table 2. Most of the rules are familiar from the  $\pi$ -calculus, with a few changes to accommodate the fact that agents are explicitly located. The main new rule is that for code movement, (r-move), which allows an agent to move from one location to another, say from  $\ell$  to  $k$ :  $\ell[[\text{go } k.P]] \longmapsto k[[P]]$ . The rule (r-comm) for communication allows two agents running at the same location  $\ell$  to exchange a value  $V$  along a common channel  $a$ :

$$\ell[[a!\langle V \rangle P]] \mid \ell[[a?(X)Q]] \longmapsto \ell[[P]] \mid \ell[[Q\{V/X\}]]$$

It is worth emphasizing that the agents must be co-located for communication to occur; agents that wish to communicate on a remote channel must first move to the remote location using the asynchronous “move” operation. Nevertheless we can easily implement a form of remote asynchronous output by using  $\ell.a!\langle V \rangle$  to denote  $\text{go } \ell.a!\langle V \rangle \text{ stop}$ . In our reduction semantics we then have:

$$k[[\ell.a!\langle V \rangle]] \mid \ell[[a?(X)Q]] \longmapsto \longmapsto \ell[[Q\{V/X\}]]$$

The purpose of the structural equivalence is to abstract from the static structure of terms, *i.e.* from the irrelevant details of the syntactic relation between composition ( $P \mid Q$ ), restriction ( $(\nu e)P$ ) and location ( $\ell[[P]]$ ). The structural equivalence

**Table 2** Reduction

(s-extr)	$N \mid (ve)M \equiv (ve)(N \mid M)$ if $e \notin \text{fn}(N)$
(s-garbage <sub>1</sub> )	$(\nu_{\ell}e:T)\mathbf{0} \equiv \mathbf{0}$
(s-garbage <sub>2</sub> )	$\ell[\text{stop}] \equiv \mathbf{0}$
(s-split)	$\ell[P \mid Q] \equiv \ell[P] \mid \ell[Q]$
(s-itr)	$\ell[*P] \equiv \ell[P] \mid \ell[*P]$
(s-new)	$\ell[(\nu_{\ell}e:T)P] \equiv (\nu_{\ell}e:T)\ell[P]$ if $e \neq \ell$
(r-move)	$\ell[\text{gok}.P] \mapsto k[P]$
(r-comm)	$\ell[a!\langle V \rangle P] \mid \ell[a?(X)Q] \mapsto \ell[P] \mid \ell[Q\{V/x\}]$
(r-eq <sub>1</sub> )	$\ell[\text{if } U = U \text{ then } P \text{ else } Q] \mapsto \ell[P]$
(r-eq <sub>2</sub> )	$\ell[\text{if } U = V \text{ then } P \text{ else } Q] \mapsto \ell[Q]$ if $U \neq V$

is defined to be the least congruence over networks that satisfies the commutative monoid laws for composition and the axioms given in Table 2. In addition to the standard axiom for name extrusion (s-extr), the structural equivalence includes axioms that allow restriction and composition to be lifted from threads to networks. The most important of these is the rule (s-split) which allows an agent to split into two independent agents ( $\ell[P \mid Q] \equiv \ell[P] \mid \ell[Q]$ ). The rule (s-garbage<sub>2</sub>) allows for garbage collection of terminated agents, whereas (s-itr) provides a standard interpretation of iteration. Note that when a channel name is extracted from a thread using (s-new) ( $\ell[(\nu_{\ell}e:T)P] \equiv (\nu_{\ell}e:T)\ell[P]$ ) it is necessary to note the location where the name is defined. This in fact determines the syntactic form for channel restriction at the network level. In  $(\nu_{\ell}a:A)M$  the channel  $a$  defined at location  $\ell$  and its scope is restricted to the network  $M$ .

**Example 2.1 (A Cell)** A simple system consisting of a user and a cell may be described as follows:

$$\begin{aligned}
\text{Net}_1 &\Leftarrow \ell[\text{Cell}(v)] \mid h[\text{User}] \\
\text{Cell}(n) &\Leftarrow (\nu s)s!\langle n \rangle \\
&\quad \mid *g?(z) \quad s?(v) \quad (s!\langle v \rangle \mid z.\text{ret}!\langle v \rangle) \\
&\quad \mid *p?(z,x) \quad s?(v) \quad (s!\langle x \rangle \mid z.\text{ack}!\langle \rangle) \\
\text{User} &\Leftarrow \ell.p!\langle h, 2 \rangle \mid \text{ack}?(()) \quad (g!\langle h \rangle \mid \text{ret}?(x) \text{ print}!\langle x \rangle)
\end{aligned}$$

The cell has an internal channel  $s$  in which the contents is stored and two public channels (or methods) for accessing the contents,  $p$  for putting values into the cell and  $g$  for getting the current contents; to make the example more accessible we assume the existence of some primitive values such as integers. The get method receives a return address from the user, which is assumed to be a location, reads the current contents and sends it along the channel  $\text{ret}$  at the callers site. The get method acts in a similar manner; it receives a value and a return address, updates the contents and sends an acknowledgement along the channel  $\text{ack}$  at the return address.

According to our reduction semantics the user and the cell may interact twice,

after which the *print* channel at the users site  $\ell$  will have the value 2 available on it.

**Example 2.2 (A refined Cell)** The cell in the previous example has the disadvantage that it may only be used by users which have the two (global) methods *ret*, *ack* available at their sites. Here we improve on this by using structured values:

$$\begin{aligned} \text{Net}_2 &\Leftarrow \ell \llbracket \text{Cell}(n) \rrbracket \mid h \llbracket \text{User} \rrbracket \\ \text{Cell}(n) &\Leftarrow (vs) s! \langle n \rangle \\ &\quad \mid *g?(z[y]) \quad s?(v) \quad (s! \langle v \rangle \mid z.y! \langle v \rangle) \\ &\quad \mid *p?(z[y], x) \quad s?(v) \quad (s! \langle x \rangle \mid z.y! \langle \rangle) \\ \text{User} &\Leftarrow (vr_1) \ell.p! \langle h[r_1], 2 \rangle \mid r_1?() (vr_2) (g! \langle h[r_2] \rangle \mid r_2?(x) \text{print!} \langle x \rangle) \end{aligned}$$

On the get method, for example, the cell receives a structured value consisting of a location, bound to  $z$ , and a channel  $y$  *at that location*, reads the current contents and sends it along the newly acquired channel. When the cell is defined in this manner the user may generate new channels  $r_1, r_2$ , local to its site  $h$ , for the purpose of communicating with the cell. This interaction strategy on makes the cell less dependent on global assumptions.

**Example 2.3 (A Cell Server)** A server for generating new cells may be defined as  $\text{serv} \llbracket S \rrbracket$  where  $S$  is given by:

$$S \Leftarrow *req?(z[y]) \quad (vcell) \quad z.y! \langle cell \rangle \mid go \text{cell}. \text{Cell}(2)$$

Upon receiving a new request, the server creates a new cell location  $cell$ , spawns the cell code at that location, initialized to 2, and then sends the name of the cell location to the user. A typical user would take the form  $h \llbracket cU \rrbracket$ , where:

$$cU \Leftarrow (vr) \text{serv}.req! \langle h[r] \rangle \mid r?(z) \quad U(z)$$

Many variations of cell servers can be described in our language. For example the following code describes a server which spawns a new cell at a location specified by the user; moreover the put and get methods are no longer global, but private to the new cell and the calling user:

$$\begin{aligned} cS' &\Leftarrow *req?(z[x]) \quad goz. (v)p, g \quad (\text{Cell}(2) \mid x! \langle p, g \rangle) \\ cU' &\Leftarrow (vr) \text{serv}.req! \langle h[r] \rangle \mid r?(p, g) \quad U'(p, g) \end{aligned}$$

**Example 2.4 (Routed Forwarding)** Here we write a program  $\text{Forwarder}(h[in], d[s])$  which establishes a connection between the local channel  $in$  and the (possibly remote) channel  $s$ . By “connection” we mean that messages sent into  $in$  should eventually find their way to the service channel  $s$  at destination  $d$ . Such a program is trivial to write in  $\mathcal{D}\pi$ :

$$*in?(x) \quad go d.s! \langle x \rangle$$

The unpleasant part of the problem specification is that we are not allowed to assume that there is a direct connection from the current location to  $d$ . Instead, the program must consult the local method  $\text{route}(d)$  which returns the name of the neighboring location that is closest to  $d$ , *i.e.* somewhere between the current location and  $d$ . To make the program readable, we assume some additional syntactic

conventions, including recursive definitions and let-expressions.

$$\begin{aligned}
 \text{Forwarder}(h[in], d[s]) \Leftarrow & \text{ if } h = d \text{ then} \\
 & \quad *in?(x) \ s!\langle x \rangle \\
 & \text{ else} \\
 & \quad \text{let } n \leftarrow \text{route}(d) \\
 & \quad \text{in } \text{gon.}(\text{vc}) \ \text{Forwarder}(n[c], d[s]) \\
 & \quad \quad | \ \text{go } h. *in?(x) \ \text{gon.} \ c!\langle x \rangle \\
 & \text{ endif}
 \end{aligned}$$

When the *Forwarder* is started, it checks to see if the destination  $d$  is the same as the current location  $h$ . If  $h$  and  $d$  are the same, then there is no need for routing, and the program can simply set up a forwarding process from  $in$  to  $s$ :  $*in?(x) \ s!\langle x \rangle$ . If  $h$  and  $d$  are different, then the name of a neighbor  $n$  is retrieved, where  $n$  is between  $h$  and  $d$  on the network. Then a new copy of the code is started at  $n$ , and a forward process is set up between  $in$  and  $n$ .

### 3 Typing

An informal description of the types for  $D\pi$  was given in the introduction. Formally they are a subset of the *pre-types* defined in Table 3 which satisfy some consistency constraints. These pre-types belong to three distinct syntactic categories:

- location types,  $K, L$ , of the form  $\text{loc}\{\tilde{\kappa}\}$ , where  $\kappa_i$  are location capabilities.
- channel types,  $A, B, C$ , of the form  $\text{res}\{\tilde{\alpha}\}$ , where  $\alpha_i$  are channel capabilities.
- transmission types,  $S, T$ , which can be of the form  $K$  for locations,  $A$  for local resources,  $\tilde{T}$  for tuples, or  $K[\tilde{A}]$  for dependent tuples with non-local resources.

Location and channel types are identified up to reordering of capabilities; in fact, they may be viewed simply as sets of capabilities. We also routinely drop brackets when they are empty.

The types come equipped with a subtyping relation, also defined in Table 3. For location pre-types we have  $K <: L$  if for every capability  $\lambda \in L$  there exists a capability  $\kappa \in K$  which is “at least as good”, *i.e.*  $\kappa <: \lambda$ . Here the location capabilities  $\kappa$  and  $\lambda$  are compared inductively using the associated types, *e.g.*  $a:A <: a:B$  if  $A <: B$ . Subtyping for channels is just as for locations:  $A <: B$  if for every capability  $\beta \in B$  there exists a capability  $\alpha \in A$  such that  $\alpha <: \beta$ . But the subtyping relation on channel capabilities is more interesting:

$$\begin{aligned}
 r\langle S \rangle <: r\langle T \rangle & \text{ if } S <: T \\
 w\langle S \rangle <: w\langle T \rangle & \text{ if } T <: S
 \end{aligned}$$

As one should expect from [22], the read capability is covariant, whereas the write capability is contravariant. Thus a receiver can always take *fewer* capabilities than specified, whereas a sender can always send *more*.

**Table 3** Pre-Types

Capabilities:	Subtyping:
$\kappa ::= \text{go} \mid \text{newc}$	$\kappa <: \kappa$
$\mid a:A$	$a:A <: a:B$ if $A <: B$
$\alpha ::= r\langle T \rangle$	$r\langle S \rangle <: r\langle T \rangle$ if $S <: T$
$\mid w\langle T \rangle$	$w\langle S \rangle <: w\langle T \rangle$ if $T <: S$
Pre-Types:	
$K ::= \text{loc}\{\tilde{\kappa}\}$	$K <: L$ if $\forall \lambda \in L: \exists \kappa \in K: \kappa <: \lambda$
$A ::= \text{res}\{\tilde{\alpha}\}$	$A <: B$ if $\forall \beta \in B: \exists \alpha \in A: \alpha <: \beta$
$T ::= K \mid A \mid (T_1, \dots, T_n)$	$\tilde{S} <: \tilde{T}$ if $\forall i: S_i <: T_i$
$\mid K[A_1, \dots, A_n]$	$K[\tilde{A}] <: L[\tilde{B}]$ if $K <: L$ and $\tilde{A} <: \tilde{B}$

**Definition 3.1 (Types)**

- (i) A location pre-type  $K$  is a type if  $a:A \in K$  and  $a:A' \in K$  imply  $A = A'$ .
- (ii) A channel pre-type  $A$  is a type if:

$$\begin{aligned} r\langle T \rangle \in A \text{ and } r\langle T' \rangle \in A &\text{ imply } T = T' \\ w\langle S \rangle \in A \text{ and } w\langle S' \rangle \in A &\text{ imply } S = S' \\ r\langle T \rangle \in A \text{ and } w\langle S \rangle \in A &\text{ imply } S <: T \end{aligned}$$

- (iii) Pre-types of the form  $\tilde{T}$  and  $K[\tilde{A}]$  are types if their constituent components are types.  $\square$

Thus location types are allowed at most one capability for each channel. Channel types are also constrained to have at most one read and one write capability. The final constraint on channel types is a consistency requirement. It prevents agents from “fabricating” capabilities. For example, it prevents an agent from sending a value at type  $\text{loc}\{a:A\}$  and then receiving the same value at type  $\text{loc}\{a:A, b:B\}$ .

Readers familiar with [22] will notice that Pierce and Sangiorgi’s channel types — “PS” types — are also representable in our type system (ignoring recursion). The PS read type  $[T]^r$  is identified with  $\text{res}\{r\langle T \rangle\}$ , the PS write type  $[T]^w$  is identified with  $\text{res}\{w\langle T \rangle\}$ , and the PS read/write type  $[T]^{rw}$  is identified with  $\text{res}\{w\langle T \rangle, r\langle T \rangle\}$ , which we abbreviate by  $\text{rw}\langle T \rangle$ . For these PS types, our definition of subtyping coincides with that of Pierce and Sangiorgi.

Our channel types include many types that are not definable using the system of Pierce and Sangiorgi, however. For example, the type

$$C = \text{res}\{r\langle \text{loc}\{a:A\} \rangle, w\langle \text{loc}\{a:A, b:B\} \rangle\}$$

is not expressible as a PS type. Nonetheless, it is easy to see how such types arise when agents are granted different permissions on the names in a network.

In addition, our subtyping relation induces a partial *meet* operator ‘ $\sqcap$ ’. No such operator exists for PS types — consider the types  $[\square]^r$  and  $[\square]^{rw}$ .



**Table 4** A Type System

Values:

$$\frac{\Gamma(u) <: \mathbf{K}}{\Gamma \vdash_w u : \mathbf{K}} \quad \frac{\Gamma(w) <: \text{loc}\{u : \mathbf{T}\}}{\Gamma \vdash_w u : \mathbf{T}} \quad \frac{\Gamma \vdash_w U_i : \mathbf{T}_i \ (\forall i)}{\Gamma \vdash_w \tilde{U} : \tilde{\mathbf{T}}} \quad \frac{\Gamma \vdash_w u : \mathbf{K} \quad \Gamma \vdash_u \tilde{v} : \tilde{\mathbf{A}}}{\Gamma \vdash_w (u, \tilde{v}) : \mathbf{K}[\tilde{\mathbf{A}}]}$$

Threads:

$$\frac{\Gamma \vdash_w u : \text{res}\{w \langle \mathbf{T} \rangle\} \quad \Gamma \vdash_w V : \mathbf{T} \quad \Gamma \vdash_w P}{\Gamma \vdash_w u ! \langle V \rangle P} \quad \frac{\Gamma \vdash_w u : \text{res}\{r \langle \mathbf{T} \rangle\} \quad \text{fv}(X) \text{ disjoint } \text{fv}(\Gamma) \quad \Gamma \sqcap \{w X : \mathbf{T}\} \vdash_w Q}{\Gamma \vdash_w u ? (X : \mathbf{T}) Q} \quad \frac{\Gamma \vdash_w U : \mathbf{S} \quad \Gamma \vdash_w V : \mathbf{T} \quad \Gamma \sqcap \{w U : \mathbf{T}\} \sqcap \{w V : \mathbf{S}\} \vdash_w P \quad \Gamma \vdash_w Q}{\Gamma \vdash_w \text{if } U = V \text{ then } P \text{ else } Q}$$

$$\frac{\Gamma \vdash_w u : \text{loc}\{\text{go}\} \quad \Gamma \vdash_u P}{\Gamma \vdash_w \text{go } u . P} \quad \frac{k \notin \text{fn}(\Gamma) \quad \Gamma \sqcap \{k : \mathbf{K}\} \vdash_w P}{\Gamma \vdash_w (\mathbf{v}k : \mathbf{K}) P} \quad \frac{\Gamma \vdash_w w : \text{loc}\{\text{newc}\} \quad a \notin \text{fn}(\Gamma) \quad \Gamma \sqcap \{w a : \mathbf{A}\} \vdash_w P}{\Gamma \vdash_w (\mathbf{v}a : \mathbf{A}) P} \quad \frac{\Gamma \vdash_w P \quad \Gamma \vdash_w Q}{\Gamma \vdash_w \text{stop}, P \mid Q, *P}$$

Networks:

$$\frac{\Gamma \vdash_k k : \text{loc} \quad \Gamma \vdash_k P}{\Gamma \vdash k \llbracket P \rrbracket} \quad \frac{\Gamma \vdash_k k : \text{loc} \quad \ell \notin \text{fn}(\Gamma) \quad \Gamma \sqcap \{\ell : \mathbf{L}\} \vdash M}{\Gamma \vdash (\mathbf{v}_k \ell : \mathbf{L}) M} \quad \frac{\Gamma \vdash_k k : \text{loc}\{\text{newc}\} \quad a \notin \text{fn}(\Gamma) \quad \Gamma \sqcap \{k a : \mathbf{A}\} \vdash M}{\Gamma \vdash (\mathbf{v}_k a : \mathbf{A}) M} \quad \frac{\Gamma \vdash M \quad \Gamma \vdash N}{\Gamma \vdash \mathbf{0}, M \mid N}$$

The primary judgments of the type system are of the form  $\Gamma \vdash M$  where  $\Gamma$  is a *type environment* and  $M$  is a network term. Type environments are taken to be maps from identifiers to *open location types*, which have the form  $\text{loc}\{\tilde{u} : \tilde{\mathbf{T}}\}$ . The typing system is given in Table 3 and uses auxiliary judgments for threads, identifiers and values. For threads, judgments have the form  $\Gamma \vdash_w P$ , indicating that the thread  $P$  is well-typed to run at location  $w$ , where  $w \in \text{dom}(\Gamma)$ . This in turn uses judgments of the form  $\Gamma \vdash_w V : \mathbf{T}$ , which indicates that the value  $V$  is well formed at  $w$  and has at least the capabilities specified by  $\mathbf{T}$ .

In this extended abstract we do not explain the various rules in detail. Instead, we briefly look at some examples.

At the thread level to deduce that  $\text{go } u . P$  is well-typed to run at  $w$ , that is  $\Gamma \vdash_w \text{go } u . P$ , we need to establish  $\Gamma(u)$  is a location with go capability and that  $P$  is well-typed to run at  $u$ , *i.e.*  $\Gamma \vdash_u P$ . At the network level to deduce that  $u \llbracket P \rrbracket$  is well-typed,  $\Gamma \vdash u \llbracket P \rrbracket$ , we need to show that  $u$  is a location and  $P$  is well-typed to run at  $u$ , *i.e.*  $\Gamma \vdash_u P$ .

At the thread level to deduce  $\Gamma \vdash_w u ? (X : \mathbf{T}) Q$  we must establish that  $u$  can be assigned type  $\text{res}\{r \langle \mathbf{T} \rangle\}$  at location  $w$ ,  $\Gamma \vdash_w u : \text{res}\{r \langle \mathbf{T} \rangle\}$ , and that  $Q$  is well-typed to run at  $w$ . But in showing the later, we may augment the environment  $\Gamma$  with the information that  $X$  is of type  $\mathbf{T}$ , that is we must show  $\Gamma \sqcap \{w X : \mathbf{T}\} \vdash_w Q$ . The formal definition of this environment extension uses the *partial meet operator*  $\sqcap$ ,

mentioned above. Since the pattern  $X$  may include structured values, the definition of environment extension is somewhat non-standard. For example, if  $X:T$  is  $(x, z[y]):(\mathbf{B}, \text{loc}\{a:A'\}[\mathbf{C}])$  then:

$$\{wX:T\} = \{w:\text{loc}\{x:\mathbf{B}\}, z:\text{loc}\{a:A', y:\mathbf{C}\}\}$$

If further  $\Gamma$  is  $\{w:\text{loc}\{a:A\}\}$ , then  $\Gamma \sqcap \{wX:T\}$  denotes  $\{w:\text{loc}\{a:A, x:\mathbf{B}\}, z:\text{loc}\{a:A', y:\mathbf{C}\}\}$ . The same notation is used in the rules for restriction.

To deduce  $\Gamma \vdash_w$  if  $u = v$  then  $P$  else  $Q$ , where in  $\Gamma$  both  $u$  and  $v$  have location types — say  $\Gamma(u) <: \mathbf{K}$  and  $\Gamma(v) <: \mathbf{L}$  — then it is necessary to establish that  $Q$  is well-typed to run at  $w$ ,  $\Gamma \vdash_w Q$ , and that  $P$  is well-typed to run at  $w$ , relative to an augmented version of  $\Gamma$  in which both  $u$  and  $v$  have inherited each others type information:  $\Gamma \sqcap \{u:\mathbf{L}, v:\mathbf{K}\} \vdash_w P$ . It is worth noting that the Routed Forwarding example of the previous section cannot be typed using the standard typing rule for matching, which requires  $\Gamma \vdash_w P$ ; other examples are discussed in the full version of the paper.

The main result of this section is the following:

**Theorem 3.2 (Subject reduction)** *If  $\Gamma \vdash M$  and  $M \longrightarrow M'$  then  $\Gamma \vdash M'$ .  $\square$*

**Example 3.3 (A Typed Cell Server)** As an example of the use of these types to control access to capabilities, consider again the cell server from Example 2.3, this time annotated with types.

$$S \Leftarrow *req?(z[y]) (\mathbf{v}cell:L_{\text{cell}}) z.y!\langle cell \rangle \mid go\ cell. Cell(0)$$

where “Cell(0)” represents the code for the cell initialized to 0.

Let us use the abbreviations for PS types introduced on page 8. The *allocation type*  $L_{\text{cell}}$  of the cell location  $cell$  can then be written:

$$\begin{aligned} L_{\text{cell}} &= \text{loc}\{go, newc, g:rw\langle T_g \rangle, p:rw\langle T_p \rangle\} \\ T_g &= \text{loc}\{go\}[w\langle \text{int} \rangle] \\ T_p &= (\text{loc}\{go\}[w\langle \rangle], \text{int}) \end{aligned}$$

Location  $cell$  must be given at least the type  $L_{\text{cell}}$  in order for the cell code to typecheck. Note that the channels  $g$  and  $p$  must be declared with both read and write capabilities as the server reads from them and a user must be able to write to them. The cell requires only the write capability on the response channels it receives on  $p$  and  $g$ .

The user’s capabilities on the cell are determined by the *transmission type*  $T_{\text{req}}$  of channel  $req$  (which must have type  $rw\langle T_{\text{req}} \rangle$ ). If one takes

$$\begin{aligned} T_{\text{req}} &= \text{loc}\{go\}[w\langle L'_{\text{cell}} \rangle] \\ L'_{\text{cell}} &= \text{loc}\{go, g:w\langle T_g \rangle, p:w\langle T_p \rangle\} \end{aligned}$$

then this type ensures that a cell user cannot “redefine” the methods  $p$  or  $g$  (by intercepting messages sent on these channels), nor can it create new channels at the cell location. We should point out that this typing also affords some level of protection

to the user. The response channel  $r$  is sent to the server with write capability only; thus the server may not intercept other messages that the user may wish to receive on  $r$ . Perhaps more important, the user's location is sent without the privilege to create new channels there, keeping the server from performing any computation at the users location.

To emphasize the restrictions imposed by these capabilities consider the following user:

$$U \Leftarrow (\nu r) \text{serv.req!}\langle h[r] \rangle \mid r?(z) U'(z)$$

$U$  requests a cell using the response channel  $r$ . Then the network  $\text{serv}[\mathbb{S}] \mid h[\mathbb{U}]$  can reduce to

$$\text{serv}[\mathbb{S}] \mid (\nu \text{cell}:\mathbb{L}_{\text{cell}}) h[\mathbb{U}'(\text{cell})] \mid \text{cell}[\mathbb{C}_{\text{cell}}]$$

If  $T_{\text{req}}$  is as above, then one can be sure that the agent  $U'(\text{cell})$  has restricted access to  $\text{cell}$  in this network. For example, if  $U'$  has the form

$$U'(\text{cell}) \Leftarrow \text{go cell.p?}(X) \dots$$

then  $U$  will be untypable.

The user may pass on to its clients the capabilities it has received for the cell, or restrictions of them. For example if  $U'$  has the form

$$U'(\text{cell}) \Leftarrow \text{req}_{\text{low}}?(z[y]) z.y!\langle \text{cell} \rangle \mid \text{req}_{\text{high}}?(z[y]) z.y!\langle \text{cell} \rangle \mid \dots$$

then the capabilities sent to *low* and *high* priority clients can be controlled by the types of the channels  $\text{req}_{\text{low}}$  and  $\text{req}_{\text{high}}$ . For example if  $\text{req}_{\text{low}}$  has the type  $\text{loc}\{\text{go}\}[\text{w}\langle \text{loc}\{\text{go}, g:\text{w}\langle T_g \rangle\} \rangle]$  then low priority clients will not have any access to the put method at the cell.

## 4 Type Safety

Due to lack of space in this extended abstract we can only briefly outline the Type Safety theorem for  $\mathcal{D}\pi$ .

We first define a tagged version of the language, where threads are explicitly annotated with the permissions/capabilities they have accumulated for locations. The syntax of threads and values is unchanged from that of Table 2; only the network level is affected, and here only the clause for agents. Each agent of the original language  $\ell[\mathbb{P}]$  is tagged with a *closed* type environment  $\Gamma$  which represents the capabilities (or permissions) of the agent. For example, the agent

$$\ell[\mathbb{P}]\{\ell:\text{loc}\{a:A, b:B\}, k:\text{loc}\{a:A'\}\}$$

has knowledge of resources  $a$  and  $b$  at  $\ell$  and of resource  $a$  at  $k$ . In addition to recording the *names* of available resources, the tag also records the *permissions* that the agent has acquired for the use of that resource (the types  $A$ ,  $B$  and  $A'$ ). This additional information allows fine control in the definition of runtime error.

**Table 5** Runtime Errors

(e-move)	$\ell\llbracket\text{go } k.P\rrbracket_\Gamma \xrightarrow{err}$	if $\Gamma(k) \not\prec \text{loc}\{\text{go}\}$
(e-newc)	$\ell\llbracket(\text{va}) P\rrbracket_\Gamma \xrightarrow{err}$	if $\Gamma(k) \not\prec \text{loc}\{\text{newc}\}$
(e-snd)	$\ell\llbracket a!\langle V \rangle Q\rrbracket_\Gamma \xrightarrow{err}$	if $\Gamma_\ell(V) \not\prec \text{wobj}(\Gamma(\ell, a))$
(e-rcv)	$\ell\llbracket a?(X:T) P\rrbracket_\Gamma \xrightarrow{err}$	if $\text{robj}(\Gamma(\ell, a)) \not\prec T$
(e-comm)	$\ell\llbracket a!\langle V \rangle P\rrbracket_\Delta \mid \ell\llbracket a?(X:T) Q\rrbracket_\Gamma \xrightarrow{err}$	if $\text{wobj}(\Delta(\ell, a)) \not\prec \text{robj}(\Gamma(\ell, a))$
(e-new)	$\frac{M \xrightarrow{err}}{(\text{ve}) M \xrightarrow{err}}$	(e-str) $\frac{M \xrightarrow{err}}{M \mid N \xrightarrow{err}} \quad \frac{M \equiv N \quad N \xrightarrow{err}}{M \xrightarrow{err}}$

Next the reduction semantics of Table 2 is adapted to show how tags evolve over time. To avoid confusion, we write  $M \mapsto M'$  for tagged reduction. The only non-trivial change is to the rule (r-comm) which is revised to

$$\ell\llbracket a!\langle V \rangle P\rrbracket_\Gamma \mid \ell\llbracket a?(X:T) Q\rrbracket_\Delta \mapsto \ell\llbracket P\rrbracket_\Gamma \mid \ell\llbracket Q\{V/X\}\rrbracket_{\Delta \sqcap \{\ell V:T\}}$$

Note that here the receiver can accumulate new capabilities from the sender as  $\Delta \sqcap \{\ell V:T\}$  denotes the result of augmenting  $\Delta$  with the information that the value  $V$  at  $\ell$  has acquired the capabilities described by  $T$ . As an example, let  $T = \text{loc}\{[C]$  in the following tagged network:

$$\ell\llbracket a!\langle k[c] \rangle P\rrbracket_{\{\dots, k:\text{loc}\{b:B, c:C\}\}} \mid \ell\llbracket a?(z[x]:T) Q\rrbracket_{\{\dots, k:\text{loc}\{d:D\}\}}$$

After the communication the network is:

$$\ell\llbracket P\rrbracket_{\{\dots, k:\text{loc}\{b:B, c:C\}\}} \mid \ell\llbracket Q\{k[c]/z[x]\}\rrbracket_{\{\dots, k:\text{loc}\{d:D, c:C\}\}}$$

The receptor has gained extra capabilities through this communication, as mediated through the reception type  $T$ .

Next because of the presence of these tags we can easily define a notion of run-time error; informally  $M \xrightarrow{err}$  means that somewhere in the (tagged network)  $M$  a thread can use a resource in some manner which contradicts the explicit permissions it has accumulated over that resource. The formal definition is given in Table 4, where  $\text{robj}(\text{res}\{r\langle T \rangle, \dots\}) = T$  and  $\text{wobj}(\text{res}\{w\langle T \rangle, \dots\}) = T$ .

The final step in the formalization is to extend the typing system of Table 3 to tagged networks,  $\Gamma \Vdash M$ . This is achieved by adding the following rule, where  $\Gamma \prec: \Delta$  if  $\Gamma(w) \prec: \Delta(w)$  for every  $w$  in  $\text{dom}(\Delta)$ .

$$\frac{\Delta \vdash_k k:\text{loc} \quad \Delta \vdash_k P}{\Gamma \Vdash k\llbracket P\rrbracket_\Delta} \Gamma \prec: \Delta$$

Within this framework we can prove the following results:

- **SUBJECT REDUCTION FOR TAGGED NETWORKS:** For all tagged networks  $N$ ,  $\Gamma \Vdash N$  and  $N \mapsto N'$  then  $\Gamma \Vdash N'$

- **TYPE SAFETY FOR TAGGED NETWORKS:** For all tagged networks  $N$ ,  $\Gamma \Vdash N$  implies  $N \xrightarrow{err}$
- **STRONG EQUIVALENCE OF TAGGED AND UNTAGGED REDUCTION:** For every well-typed (untagged) network,  $\Gamma \vdash M$ , we can define a canonical well-typed tagged network  $\Gamma \Vdash \text{tag}_\Gamma(M)$  that is strongly bisimilar to  $M$ .

These results, together with Theorem 3.2, imply that a well-typed network  $\Gamma \vdash M$  is strongly bisimilar to the tagged network  $\text{tag}_\Gamma(M)$  and that this explicitly tagged network will never raise a runtime error, *i.e.* no agent will ever misuse a resource during its execution.

## 5 Related Work

There are numerous languages now in the literature for describing distributed systems;  $D\pi$  is perhaps closest in spirit to [9,2,23,3] which also take as their point of departure the  $\pi$ -calculus, although with each there are significant differences. For example in the join calculus [9] message routing is *automatic* as the restricted syntax ensures that all channels have a unique location at which they are serviced. In  $D\pi$ , to send a message to a remote location, an agent must first spawn a sub-agent which moves to that location; locations are more *visible* in  $D\pi$ . In addition, several of these languages [9,23,3] adopt *location movement* as the mechanism for agent mobility. Location movement allows groups of running threads to be moved about the network asynchronously (*i.e.* without each thread performing an explicit go); for further discussion, see the full version.

Many channel-based typing systems for  $\pi$ -calculi and related languages have been proposed. For example in [22], Pierce and Sangiorgi define a type system for the  $\pi$ -calculus with read and write capabilities on channels. Sewell [23] generalizes the type system of [22] to distinguish between *local* communication, which can be efficiently implemented, and *non-local* communication. Fournet *et al.* [10] have developed an *ML*-style typing system for the join calculus where channels are allowed a certain amount of polymorphism. Amadio [2] has presented a type system that guarantees that channel names are defined at exactly one location, whereas the type system of Kobayashi *et al.* [17] ensures that some channels are used linearly.

The work closest to ours is that of de Nicola, Ferrari and Pugliese [7]. Their goals are the same as ours, but the specifics of their solution are quite different. They work with a variant of Linda [6] with multiple “tuple spaces”. Tuple spaces correspond to locations in our setting, and tuples (named data) correspond to resources. The type system of [7] controls access to tuple spaces, rather than to specific tuples, and thus provides coarser-grained control of resource access than that provided by our typing system.

Static analyses for proving various security properties of programs have also been proposed by several authors; two recent references are [16,13].

## Acknowledgements

We would like to thank INRIA Sophia Antipolis for their hospitality while conducting this research. We have benefited from conversations with Alan Jeffrey, Peter Sewell and Luca Cardelli, among others.

## References

- [1] R. Amadio and S. Prasad. Modelling IP mobility. Internal Report 244, Laboratoire d'Informatique de Marseille, 1997.
- [2] R. Amadio. An asynchronous model of locality, failure, and process mobility. In *COORDINATION '97*, volume 1282 of *Lecture Notes in Computer Science*. Springer-Verlag, 1997.
- [3] L. Cardelli and A. D. Gordon. Mobile ambients. In Maurice Nivat, editor, *Proc. FOSSACS'98, International Conference on Foundations of Software Science and Computation Structures*, volume 1378 of *Lecture Notes in Computer Science*, pages 140–155. Springer-Verlag, 1998.
- [4] L. Cardelli. A language with distributed scope. *Computing Systems*, 8(1):27–59, January 1995. A preliminary version appeared in Proceedings of the 22nd ACM Symposium on Principles of Programming.
- [5] N. Carriero and D. Gelernter. Linda in context. *Communications of the ACM*, 32(4):444–458, 1989.
- [6] N. Carriero, D. Gelernter, and L. Zuck. Bauhaus Linda. In *Object-Based Models and Languages for Concurrent Systems*, number 924 in *Lecture Notes in Computer Science*, pages 66–76. Springer-Verlag, 1995.
- [7] R. De Nicola, G. Ferrari, and R. Pugliese. Coordinating mobile agents via blackboards and access rights. In *COORDINATION '97*, volume 1282 of *Lecture Notes in Computer Science*. Springer-Verlag, 1997.
- [8] K.M. Chandy *et al.* A world-wide distributed system using java and the internet. In *IEEE International Symposium on High Performance Distributed Computing*. IEEE, August 1996.
- [9] C. Fournet, G. Gonthier, J.J. Levy, L. Marganet, and D. Rémy. A calculus of mobile agents. In U. Montanari and V. Sassone, editors, *CONCUR: Proceedings of the International Conference on Concurrency Theory*, volume 1119 of *Lecture Notes in Computer Science*, pages 406–421, Pisa, August 1996. Springer-Verlag.
- [10] C. Fournet, C. Laneve, L. Maranget, and D. Rémy. Implicit typing la ml for the join-calculus. In *CONCUR: Proceedings of the International Conference on Concurrency Theory*, *Lecture Notes in Computer Science*, Warsaw, August 1997. Springer-Verlag.
- [11] General Magic Inc. Agent technology. [http://www.genmagic.com/html/agent\\_overview.html](http://www.genmagic.com/html/agent_overview.html), 1997.

- [12] A. Giacalone, P. Mishra, and S. Prasad. A symmetric integration of concurrent and functional programming. *International Journal of Parallel Programming*, 18(2):121–160, 1989.
- [13] N. Heintz and J.G. Riecke. The SLam calculus: Programming with secrecy and integrity. In *Conference Record of the ACM Symposium on Principles of Programming Languages*, San Diego, January 1998. ACM Press.
- [14] M. Hennessy and J. Riely. Resource access control in systems of mobile agents. Computer Science Technical Report 2/98, University of Sussex, 1998. Available from <http://www.cogs.susx.ac.uk/>.
- [15] IBM Corp. The IBM aglets workbench. <http://www.trl.ibm.co.jp/aglets/>, 1996.
- [16] G. Karjoth, D.B. Lange, and M. Oshima. A security model for aglets. *IEEE Internet Computing*, 1(4), 1997.
- [17] N. Kobayashi, B.C. Pierce, and D.N. Turner. Linearity and the pi-calculus. In *Conference Record of the ACM Symposium on Principles of Programming Languages*, Paris, January 1996. ACM Press.
- [18] R. Milner. The polyadic  $\pi$ -calculus: a tutorial. Technical Report ECS-LFCS-91-180, Laboratory for Foundations of Computer Science, Department of Computer Science, University of Edinburgh, UK, October 1991. Also in *Logic and Algebra of Specification*, ed. F. L. Bauer, W. Brauer and H. Schwichtenberg, Springer-Verlag, 1993.
- [19] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, Parts I and II. *Information and Computation*, 100:1–77, September 1992.
- [20] ObjectSpace Inc. Objectspace voyager. <http://www.objectspace.com/voyager>, 1997.
- [21] C. Perkins. IP mobility support. RFC 2002, 1996.
- [22] B. Pierce and D. Sangiorgi. Typing and subtyping for mobile processes. *Mathematical Structures in Computer Science*, 6(5):409–454, 1996. Extended abstract in LICS '93.
- [23] P. Sewell. Global/local subtyping for a distributed  $\pi$ -calculus. Technical Report 435, Computer Laboratory, University of Cambridge, August 1997.
- [24] Sun Microsystems Inc. Java home page. <http://www.javasoft.com/>, 1995.