

Type-Safe Execution of Mobile Agents in Anonymous Networks

Matthew Hennessy¹ and James Riely²

¹ Univ. of Sussex. matthewh@sussex.ac.uk

² North Carolina State Univ. riely@csc.ncsu.edu

Abstract. We study type-safety properties of *open* distributed systems of mobile agents, where not all sites are known to be well-typed. We adopt the underlying model of an *anonymous network*, allowing that code may be corrupted on transmission and that the source of incoming code is unknowable. Nonetheless, we are able to guarantee a weak form of type-safety at “good” sites using a mix of static and dynamic typing.

1 Introduction

In [6] we presented a type system for controlling the use of resources in a distributed system. The type system guarantees that resource access is always *safe*, in the sense that, for example, integer channels are always used with integers and boolean channels are always used with booleans. The type system of [6], however, requires that all agents in the system be well-typed. In open systems, such as the internet, such global properties are impossible to verify. In this paper, we present a type system for *partially typed* networks, where only a subset of agents are assumed to be well typed.

This notion of partial typing is presented using the language $D\pi$, from [6]. In $D\pi$ mobile agents are modeled as *threads*, using a thread language based on the π -calculus. Threads are *located*, carrying out computations at *locations* or *sites*. Located threads, or *agents* interact using local *channels*, or *resources*.

In an open system, not all sites are necessarily benign. Some sites may harbor malicious agents that do not respect the typing rules laid down for the use of resources. For example, consider the system

$$k \llbracket (vc:\text{chan}\langle\text{int}\rangle) \text{gom}.a!\langle k[c]\rangle \rrbracket \\ | m \llbracket a?(z[x]) \text{go}z.x!\langle t \rangle \rrbracket$$

consisting of two sites k and m . The first generates a new local channel c for transmitting integers and makes it known to the second site m , by sending it along the channel a local to m . Here, the value $k[c]$ and the pattern $z[x]$ are *dependent pairs*, where the first element represents a location and the second element represents a resource *at that location*. In a benign world k could assume that any mobile agent that subsequently migrates from m to k would only use this new channel c to transmit integers at k . However in an insecure world m may not play according to the rules; in our example it sends an agent to k which misuses the new resource by sending the boolean value t along it.

In this paper we formalize one strategy that sites can use to protect themselves from such attacks. The strategy makes no assumptions about the security of the underlying network. For example, it is not assumed that the source of a message (or agent) can be reliably determined. We refer to such networks as *anonymous networks*.

In the presence of anonymous networks a reasonable strategy for sites is based on *paranoia*. Since the source of messages cannot be determined it is impossible to distinguish messages from potentially “trusted” sites; thus no site can be trusted. To protect itself, a site must bar entry of any mobile agent that cannot be proven to use local resources as intended.

The requirement that incoming agents be checked entails that the runtime semantics is of necessity more complicated than that of [6]. Each location must retain information about its local resources against which the incoming agents are checked. This information is encapsulated as a location type, giving the names of the resources available locally, together with their types; agent checking then amounts to a form of *local runtime typechecking*.

The paper proceeds as follows. In the next section we review the language $D\pi$ of [6] and recall its *standard* semantics. While the paper is self contained, we rely on [6] for examples and motivation for $D\pi$. This is followed, in Section 3, by a formal description of the modified run-time semantics, where sites retain local information against which incoming agents are checked.

In Section 4, we discuss the effectiveness of this runtime strategy; we show that *good sites* cannot be harmed by *bad* sites, in the sense that local resources can not be misused at good sites, even in systems which may contain malicious sites. This is formalized in terms of a static typing system, for which we prove Subject Reduction and Type Safety theorems. The paper ends with a critique of both the run-time strategy and the typing system.

2 The Language

We present the syntax and standard semantics of $D\pi$. For a full treatment of the language, including many examples, see [6]. The language is a simplification and refinement of that introduced in [11].

2.1 Syntax

The syntax is given in Table 1. We defer the discussion of types, T , to Section 2.3. The syntax is parameterized with respect to the following syntactic sets, which we assume to be disjoint:

- *Var*, of *variables*, ranged over by $x-z$,
- *Name*, of *names*, ranged over by $a-m$,
- *Int*, of *integers*, ranged over by i , and
- *Bool* = $\{t, f\}$, of *booleans*, ranged over by bv .

We use $u-w$ to range over the set of *identifiers*, $Id = Var \cup Name$. We typically use the names $a-d$ to refer to channels and $k-m$ to refer to locations, although the distinction is

Table 1 Syntax

$$\begin{array}{l}
\text{Id: } u, v, w ::= e \mid x \\
\text{Val: } U, V ::= i \mid \text{bv} \mid u \mid w[u_1, \dots, u_n] \mid (U_1, \dots, U_n) \\
\text{Pat: } X, Y ::= x \mid z[x_1, \dots, x_n] \mid (X_1, \dots, X_n) \\
\\
\text{Thread: } P, Q, R ::= \text{stop} \mid u!\langle V \rangle P \mid u?(X:T) Q \mid *P \\
\quad \quad \quad \mid \text{go } u.P \mid \text{if } U = V \text{ then } P \text{ else } Q \\
\quad \quad \quad \mid P \mid Q \mid (\nu e:T) P \\
\\
\text{System: } M, N ::= \mathbf{0} \mid \ell[[P]] \mid M \mid N \mid (\nu \ell e:T) N
\end{array}$$

formally imposed by the type system. We use e to refer to names that might be of either type. The main syntactic categories of the language are as follows:

- *Threads*, P, R , are (almost) terms of the ordinary polyadic π -calculus [8]. The thread language includes the static combinators for composition ‘|’ and typed restriction ‘ $(\nu e:T)$ ’, as well as constructs for movement ‘ $\text{go } \ell.P$ ’, output ‘ $u!\langle V \rangle P$ ’, typed input ‘ $u?(X:T) Q$ ’, (mis)matching ‘if $U = V$ then P else Q ’ and iteration ‘ $*P$ ’.
- *Agents*, $\ell[[P]]$, are located threads.
- *Systems*, M, N , are collections of agents combined using the static combinators ‘|’ and ‘ $(\nu \ell e:T)$ ’.

The output and input constructs make use of syntactic categories for *values*, U, V , and *patterns*, X, Y , respectively. Values include variables, names, base values, and tuples of these. A value of the form $w[u_1, \dots, u_n]$ includes a location w and a collection of channels u_1, \dots, u_n allocated at w . Patterns, X, Y , provide destructors for each type. To be well-formed, we require that patterns be linear, *i.e.* each variable appear at most once.

As an example of a system, consider the term:

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

This system contains three agents, $\ell[[P]]$, $\ell[[Q]]$ and $k[[R]]$. The first two agents are running at location ℓ , the third at location k . Moreover Q and R share a private channel a of type A , allocated at ℓ and unknown to P .

Unlike [3, 5], agents are relatively lightweight in $\text{D}\pi$. They are single-threaded and can be freely split and merged using structural rules and communication. As such, they are unnamed.

NOTATION. We adopt several notational conventions, as in [6].

- In the concrete syntax, “move” has greater binding power than composition. Thus $\text{go } \ell.P \mid Q$ should be read $(\text{go } \ell.P) \mid Q$. We adopt several standard abbreviations. For example, we routinely drop type annotations when they are not of interest. We omit trailing occurrences of stop and often denote tuples and other groups using a tilde. For example, we write \tilde{a} instead of (a_1, \dots, a_n) and $(\nu \tilde{e}:\tilde{T})P$ instead of $(\nu e_1:T_1) \dots (\nu e_n:T_n)P$. We also write ‘if $U = V$ then P ’ instead of ‘if $U = V$ then P else stop’ and ‘if $U \neq V$ then Q ’ instead of ‘if $U = V$ then stop else Q ’.

Table 2 Standard Reduction

Structural equivalence:

$$\begin{array}{ll}
\text{(s-stop)} & \ell[\text{stop}] \equiv \mathbf{0} \\
\text{(s-split)} & \ell[P \mid Q] \equiv \ell[P] \mid \ell[Q] \\
\text{(s-itr)} & \ell[*P] \equiv \ell[P] \mid \ell[*P] \\
\text{(s-new)} & \ell[(\nu e:\mathbf{T})P] \equiv (\nu e:\mathbf{T}) \ell[P] \\
\text{(s-extr)} & M \mid (\nu e:\mathbf{T})N \equiv (\nu e:\mathbf{T})(M \mid N) \quad \text{if } e \notin \text{fn}(M)
\end{array}$$

Reduction:

$$\begin{array}{ll}
\text{(r-move)} & \ell[\text{gok}.P] \longrightarrow k[P] \\
\text{(r-comm)} & \ell[a!(V)P] \mid \ell[a?(X)Q] \longrightarrow \ell[P] \mid \ell[Q\{V/x\}] \\
\text{(r-eq}_1\text{)} & \ell[\text{if } U = U \text{ then } P \text{ else } Q] \longrightarrow \ell[P] \\
\text{(r-eq}_2\text{)} & \ell[\text{if } U = V \text{ then } P \text{ else } Q] \longrightarrow \ell[Q] \quad \text{if } U \neq V \\
\text{(r-new)} & \frac{N \longrightarrow N'}{(\nu e:\mathbf{T})N \longrightarrow (\nu e:\mathbf{T})N'} \\
\text{(r-str)} & \frac{N \longrightarrow N'}{M \mid N \longrightarrow M \mid N'} \quad \frac{N \equiv M \longrightarrow M' \equiv N'}{N \longrightarrow N'}
\end{array}$$

- We assume the standard notion of *free* and *bound* occurrences of variables and names in systems and threads. The variables in the pattern X are bound by the input construct $u?(X)Q$, the scope is Q . The name e is bound by the restrictions $(\nu e:\mathbf{T})N$ and $(\nu e:\mathbf{T})P$, the scopes are N and P , respectively. A term with no free variables is *closed*. The functions $\text{fn}(N)$ and $\text{fid}(N)$ return respectively the sets of free names and free identifiers occurring in N .
- We also assume a standard notion *substitution*, where $N\{u/x\}$ denotes the capture-avoiding substitution of u for x in N . The notation $N\{V/x\}$ generalizes this in an obvious way as a sequence of substitutions. For example, $N\{\ell[a]/z[x]\} = N\{\ell/z\}\{a/x\}$.
- In the sequel we identify terms up to renaming of bound identifiers. \square

2.2 Standard Reduction

The standard reduction semantics is given in Table 2. The structural equivalence ($N \equiv M$) and reduction relation ($N \longrightarrow N'$) both relate *closed* system terms. The structural equivalence is defined to be the least equivalence relation that is closed under composition and restriction, satisfies the monoid laws for composition,¹ and in addition satisfies the axioms given in Table 2.

The structural rules allow for agent splitting $\ell[P \mid Q] \equiv \ell[P] \mid \ell[Q]$ and garbage collection $\ell[\text{stop}] \equiv \mathbf{0}$. Note that when a new name is lifted out of an agent, using the structural rule (s-new), the system-level restriction records the name of the location which allocated the name $\ell[(\nu e:\mathbf{T})P] \equiv (\nu e:\mathbf{T}) \ell[P]$; these location tags are used only for static typing.

¹ The monoid laws are: $N \mid \mathbf{0} \equiv N$, $N \mid M \equiv M \mid N$, and $N \mid (M \mid O) \equiv (N \mid M) \mid O$.

Most of the rules of the reduction relation are taken directly from the π -calculus, with a few changes to accommodate the fact that agents are explicitly located. The main rule of interest here is (r-move),

$$\ell \llbracket \text{gok}.P \rrbracket \longrightarrow k \llbracket P \rrbracket$$

which states that an agent located at ℓ can move to k using the move operator $\text{gok}.P$. This is the only rule that varies significantly between the standard semantics and the semantics for open systems defined later. Note that in (r-comm), communication is purely local:

$$\ell \llbracket a! \langle V \rangle P \rrbracket \mid \ell \llbracket a? \langle X \rangle Q \rrbracket \longrightarrow \ell \llbracket P \rrbracket \mid \ell \llbracket Q \{V/X\} \rrbracket$$

As an example, suppose that we wish to write a system of two agents, one at k and one at ℓ . The agent at k wishes to send a fresh integer channel a , located at k , to the other agent using the channel b , located at ℓ . This system could be written:

$$\begin{aligned} & \ell \llbracket b?(z[x]) Q \rrbracket \mid k \llbracket (\text{va}) (P \mid \text{go} \ell. b! \langle k[a] \rangle) \rrbracket \\ \equiv & \ell \llbracket b?(z[x]) Q \rrbracket \mid (\nu_k a) (k \llbracket P \rrbracket \mid k \llbracket \text{go} \ell. b! \langle k[a] \rangle \rrbracket) && \text{(s-new), (s-split)} \\ \longrightarrow & \ell \llbracket b?(z[x]) Q \rrbracket \mid (\nu_k a) (k \llbracket P \rrbracket \mid \ell \llbracket b! \langle k[a] \rangle \rrbracket) && \text{(r-move)} \\ \longrightarrow & (\nu_k a) (\ell \llbracket Q \{k[a]/z[x]\} \rrbracket \mid k \llbracket P \rrbracket) && \text{(s-extr), (r-comm), (s-stop)} \end{aligned}$$

Beside each reduction, we have written the rules used to infer it, omitting (r-str) and (r-new), which are almost always used. Note that after arriving at ℓ , the agent sends the value $k[a]$ rather than simply a . In the type system, this identifies the resource a as *non-local* at ℓ . If a simple resource value, such as a , had been communicated to Q , it would have had to have been local to ℓ , rather than k . An example of a process Q that uses the received value $z[x]$ is $\text{go} z.x! \langle 1 \rangle$, which after the communication would become $\text{go} k.a! \langle 1 \rangle$. This can move to the location k and send the integer 1 on the newly received channel a .

2.3 Types and Subtyping

The purpose of the type system is to ensure proper use of base types, channels and locations. In this paper we use the simple type languages from [6, §4], extended with base types for integers and booleans. We use uppercase Roman letters to range over types, whose syntax is as follows:

$$\begin{aligned} TChan: & A, B, C ::= \text{chan} \langle T \rangle \\ TLoc: & K, L, M ::= \text{loc} \{a_1:A_1, \dots, a_n:A_n\}, \quad a_i \text{ distinct} \\ TVal: & S, T ::= \text{bool} \mid \text{int} \mid A \mid K[B_1, \dots, B_h], \mid (T_1, \dots, T_n) \end{aligned}$$

Types are divided into the following syntactic groups:

- $TChan$ of channel types, which specify the type of values communicated over a channel, $\text{chan} \langle T \rangle$.
- $TLoc$ of simple location types, which specify the set of typed channels available at a location, $\text{loc} \{\tilde{a}:\tilde{A}\}$.

- *TVal* of value types, which include types for base values, channels, locations and tuples.

In value types, location types have the form $\text{loc}\{a_1:A_1, \dots, a_n:A_n\}[\mathbf{B}_1, \dots, \mathbf{B}_n]$. The extended form allows for a certain amount of first-order existential polymorphism. Informally, $\text{loc}\{\tilde{a}:\tilde{A}\}[\tilde{\mathbf{B}}]$ may be read “ $\exists \tilde{x}: \text{loc}\{\tilde{a}:\tilde{A}, \tilde{x}:\tilde{\mathbf{B}}\}$ ”, *i.e.* the type of a location which has channels \tilde{a} of types \tilde{A} and some (unnamed) channels of types $\tilde{\mathbf{B}}$.

Location types are essentially the same as standard record types, and we identify location types up to reordering of their “fields”. Thus $\text{loc}\{a:A, b:\mathbf{B}\}[\mathbf{C}] = \text{loc}\{b:\mathbf{B}, a:A\}[\mathbf{C}]$. But reordering is not allowed on “abstract” fields. Thus if \mathbf{B} and \mathbf{C} are different, then $\text{loc}\{a:A\}[\mathbf{B}, \mathbf{C}] \neq \text{loc}\{a:A\}[\mathbf{C}, \mathbf{B}]$.

Throughout the text, we drop empty braces when clear from context, writing ‘loc’ instead of ‘loc{ }[]’, ‘K’ instead of ‘K[]’, and ‘u’ instead of ‘u[]’.

The subtyping relation ($T \leq S$) is discussed at length in [6]. On base types and channel types there is no nontrivial subtyping; for example, $\text{chan}\langle T \rangle \leq \text{chan}\langle T' \rangle$ if and only if $T = T'$. On location types (both simple and “existential”), the subtyping relation is similar to that traditionally defined for record or object types:

$$\begin{aligned} \text{loc}\{\tilde{a}:\tilde{A}, b:\mathbf{B}\} &\leq \text{loc}\{\tilde{a}:\tilde{A}\} \\ \text{loc}\{\tilde{a}:\tilde{A}, b:\mathbf{B}\}[\tilde{\mathbf{C}}] &\leq \text{loc}\{\tilde{a}:\tilde{A}\}[\tilde{\mathbf{C}}] \end{aligned}$$

On tuples, the definition is by homomorphic extension: $\tilde{\mathbf{S}} \leq \tilde{\mathbf{T}}$ if $\forall i: S_i \leq T_i$

Both subtyping and existential dependent types are inherited from our earlier work. Whereas subtyping is not essential to the work at hand, dependent types are, as will be seen in Section 3.1.

2.4 Type Environments

Location types contain the names of the channels known to be defined at a location. To present typing systems for the language in later sections, it is useful to generalize location types to allow the inclusion of *variables* as well as names. Variables are allowed at types `int` and `bool`, in addition to channel types \mathbf{A} . The resulting types are called *v-open* location types, \mathbb{K} :

$$\begin{aligned} T\text{Simple}: \quad \mathbf{H}, \mathbf{G} &::= \text{int} \mid \text{bool} \mid \mathbf{A} \\ T\text{Open}: \quad \mathbb{K}, \mathbb{L}, \mathbb{M} &::= \text{loc}\{\tilde{u}:\tilde{\mathbf{H}}\}, \quad u_i \text{ distinct} \end{aligned}$$

To be well-formed, we require that every *name* in a v-open location type be associated with a channel type.

The subtyping relation extends directly to v-open location types:

$$\text{loc}\{\tilde{u}:\tilde{\mathbf{H}}, \tilde{v}:\tilde{\mathbf{G}}\} \leq \text{loc}\{\tilde{u}:\tilde{\mathbf{H}}\}$$

A type environment, Γ , maps identifiers to v-open location types. An example of a type environment is:

$$\Gamma = \left\{ \begin{array}{l} k : \text{loc}\{a:\text{chan}\langle \text{int} \rangle, x:\text{int}\} \\ z : \text{loc}\left\{ \begin{array}{l} a:\text{chan}\langle \text{loc}[\text{chan}\langle \text{int} \rangle] \rangle \\ y:\text{chan}\langle \text{loc}[\text{chan}\langle \text{bool} \rangle] \rangle \end{array} \right\} \end{array} \right\}$$

Here we have two locations, k and z . The first has an integer channel named a and an integer variable x . The second has two channels: a , which communicates (potentially remote) integer channels and y which communicates (potentially remote) boolean channels.

If a type environment contains no variables, we say that it is *closed*. Closed type environments map names to (closed) location types K .

In the typing system of the next section we need some notation for extending (v -open) location types and for convenience we explain this notation here. The location type \mathbb{L} may be extended using the notation ' $\mathbb{L}, V:T$ ', defined by induction on V :

$$\begin{aligned} (\text{loc}\{\tilde{u}:\tilde{H}\}), v:G &= \text{loc}\{\tilde{u}:\tilde{H}, v:G\} \quad \text{if } \forall i: u_i \neq v \\ \mathbb{L}, w[\tilde{u}]:K[\tilde{A}] &= \mathbb{L} \quad \text{if } \{w, \tilde{u}\} \text{ disjoint } \text{dom}(\mathbb{L}) \\ \mathbb{L}, \text{bv:bool} &= \mathbb{L} \\ \mathbb{L}, i:\text{int} &= \mathbb{L} \\ \mathbb{L}, \tilde{U}:\tilde{T} &= ((\dots (\mathbb{L}, U_1:T_1), \dots), U_n:T_n) \end{aligned}$$

Therefore the extension of \mathbb{L} by $V:T$ adds only information about *local identifiers* to \mathbb{L} . For example:

$$\text{loc}\{d:D\}, (0, x, z[y]):(\text{int}, A, \text{loc}\{c:C\}[B]) = \text{loc}\{d:D, x:A\}$$

Note that every location type L is also a v -open location type, and thus this definition applies to ‘‘closed’’ location types as well. In the same way the definition applies to patterns, as well as values, since syntactically every pattern is also a value.

We use a similar notation for extending type environments: If u is not in the domain of Γ then $\Gamma, u:\mathbb{L}$ denotes the new type environment, which is similar to Γ but in addition maps u to type \mathbb{L} .

3 Semantics for Open Systems

As we have explained in the introduction, using the standard run-time semantics local resources may be intentionally misused by malicious sites; in the example given in the Introduction, we have seen that the local resource a at location k is deliberately misused by a mobile agent originating at location m . In this section we modify the run-time semantics in a manner which offers protection against such misuse.

We are assuming that agents are working in an anonymous network and the run-time strategy suggested in the introduction is for every site to keep a record of local resources against which all incoming agents are checked. In our formulation this distributed information is collected together as closed type environment and thus reductions of the runtime semantics take the form

$$\Delta \triangleright N \mapsto N'$$

where Δ is a closed type environment and N and N' are systems.

The reduction semantics is given in Table 3. Most of the rules are simple adaptations of the corresponding rules in Table 2. For example, the rules for local communication

Table 3 Reduction for Open Systems

(r-move)	$\Delta \triangleright \ell \llbracket \text{go } k. P \rrbracket \longrightarrow k \llbracket P \rrbracket$	if $\Delta(k) \vdash P$
(r-comm)	$\Delta \triangleright \ell \llbracket a!(V) P \rrbracket \mid \ell \llbracket a?(X) Q \rrbracket \longrightarrow \ell \llbracket P \rrbracket \mid \ell \llbracket Q \{V/x\} \rrbracket$	
(r-eq ₁)	$\Delta \triangleright \ell \llbracket \text{if } U = U \text{ then } P \text{ else } Q \rrbracket \longrightarrow \ell \llbracket P \rrbracket$	
(r-eq ₂)	$\Delta \triangleright \ell \llbracket \text{if } U = V \text{ then } P \text{ else } Q \rrbracket \longrightarrow \ell \llbracket Q \rrbracket$	if $U \neq V$
(r-new)	$\frac{\Delta, \ell: (\mathbb{L}, a:A) \triangleright N \longrightarrow N'}{\Delta, \ell: \mathbb{L} \triangleright (\nu_{\ell} a:A) N \longrightarrow (\nu_{\ell} a:A) N'}$	$\frac{\Delta, \ell: \mathbb{L}, k: \mathbb{K} \triangleright N \longrightarrow N'}{\Delta, \ell: \mathbb{L} \triangleright (\nu_{\ell} k: \mathbb{K}) N \longrightarrow (\nu_{\ell} k: \mathbb{K}) N'}$
(r-str)	$\frac{\Delta \triangleright N \longrightarrow N'}{\Delta \triangleright M \mid N \longrightarrow M \mid N'}$	$\frac{N \equiv M \quad \Delta \triangleright M \longrightarrow M' \quad M' \equiv N'}{\Delta \triangleright N \longrightarrow N'}$

and matching of values are essentially as before, as Δ is not consulted for these reductions. There is a minor change in the rule for the restriction operator, because Δ must be augmented to reflect the addition of the new name.

The only significant change from the standard run-time semantics is in the rule for code movement:

$$\Delta \triangleright \ell \llbracket \text{go } k. P \rrbracket \longmapsto k \llbracket P \rrbracket \quad \text{if } \Delta(k) \vdash P$$

This says that the agent P can move from location ℓ to location k only if P is guaranteed not to misuse the local resources of k , *i.e.* $\Delta(k) \vdash P$. Here P is type-checked dynamically against $\Delta(k)$, which gives the names and types of the resources available at k .

3.1 Runtime Typing

The definition of this runtime local type-checking is given in Table 4. This is a *light weight* typing in that the incoming code is only checked to the extent of its references to local resources. Thus judgments are of the form

$$\mathbb{L} \vdash P$$

indicating that P can safely run at a location that provides resources as defined in \mathbb{L} .

Perhaps the most surprising rule in this light weight type checking is (t-move), which involves no type checking whatsoever. However this is reasonable as an agent such as $\text{go } \ell. P$ running at k uses no local resources; it moves immediately to the site ℓ . As a result of this rule notice that reductions of the form

$$\Delta \triangleright m \llbracket \text{go } k. \text{go } \ell. P \rrbracket \longmapsto k \llbracket \text{go } \ell. P \rrbracket$$

are always allowed, regardless of the information in Δ .

The only significant local checking is carried out by the two rules (t-r), (t-w) which we examine in some detail. The subtlety in the read rule (t-r) is to some extent hidden in the rule for updating location types and this is best explained by example. The rule dictates, for example, that the agent $a?(x:\text{chan}\langle T \rangle) Q$ can migrate to a location with local resources \mathbb{L} provided:

Table 4 Typing for Values and Threads

Values:

$$\begin{array}{l}
\text{(t-sit)} \frac{\mathbb{L} \leq \text{loc}\{u:H\}}{\mathbb{L} \vdash u:H} \qquad \text{(t-base)} \frac{}{\mathbb{L} \vdash n:\text{int}, \text{bv}:\text{bool}} \\
\text{(t-loc)} \frac{}{\mathbb{L} \vdash u[v_1, \dots, v_n]:\mathbf{K}[A_1, \dots, A_n]} \qquad \text{(t-tup)} \frac{\forall i: \mathbb{L} \vdash U_i:T_i}{\mathbb{L} \vdash (U_1, \dots, U_n):(T_1, \dots, T_n)}
\end{array}$$

Threads:

$$\begin{array}{l}
\text{(t-move)} \frac{}{\mathbb{L} \vdash \text{gou}.P} \qquad \text{(t-newl)} \frac{\mathbb{L} \vdash P}{\mathbb{L} \vdash (\nu k:\mathbf{K})P} \\
\text{(t-r)} \frac{\mathbb{L} \vdash u:\text{chan}\langle T \rangle \quad \mathbb{L}, X:T \vdash Q}{\mathbb{L} \vdash u?(X:T)Q} \qquad \text{(t-newc)} \frac{\mathbb{L}, a:A \vdash P}{\mathbb{L} \vdash (\nu a:A)P} \\
\text{(t-w)} \frac{\mathbb{L} \vdash u:\text{chan}\langle T \rangle, V:T, P}{\mathbb{L} \vdash u!\langle V \rangle P} \qquad \text{(t-str)} \frac{\mathbb{L} \vdash P, Q}{\mathbb{L} \vdash \text{stop}, *P, P|Q} \\
\text{(t-eq)} \frac{\mathbb{L} \vdash U:T, V:T, P, Q}{\mathbb{L} \vdash \text{if } U = V \text{ then } P \text{ else } Q}
\end{array}$$

- a is a local channel of the appropriate type, in this case a channel for communicating values of type $\text{chan}\langle T \rangle$
- Q is locally well-typed with respect to an augmented set of resources, $\mathbb{L}, x:\text{chan}\langle T \rangle$.

However suppose the channel a communicates non-local information; *e.g.* when is $a?(z[x]:\mathbf{K}[A])Q$ locally well-typed? The rule (t-r) simply demands that, in addition to a having the appropriate local type, Q is well-typed with respect to the same set of local resources \mathbb{L} . Formally this is because according to the definitions given in Section 2.4 $\mathbb{L}, z[x]:\mathbf{K}[A]$ is simply \mathbb{L} . Intuitively this is reasonable since any non-local information received on a will not be used locally and thus it may be ignored.

The rule (t-w) states that agent $a!\langle V \rangle P$ is locally well-typed provided

- the continuation P is locally well-typed
- the channel a has an appropriate local type, say $\text{chan}\langle T \rangle$
- the value to be transmitted V is locally well-typed to be transmitted on a , $\mathbb{L} \vdash V:T$; this means that it must be possible to assign to V the local object type of the channel a , namely T .

Once more there is a subtlety, this time in the local type checking of values. If the value V to be transmitted is a local resource, say a channel name b , then according to the rule (t-sit) b must have the local type T . If, on the other hand, V is a non-local value, say $k[b]$, then locally this is of no interest; according to (t-loc) $k[b]$ can be assigned *any* location type which in effect means that when it is transmitted locally on a its validity is not checked.

This ends our discussion of runtime local type checking, and of the runtime semantics.

3.2 An Example

As an example consider a system of three locations, k , ℓ and m , with the following distributed type environment, Δ .

$$\Delta = \left\{ \begin{array}{l} k : \text{loc}\{a : \text{chan}\langle \text{int} \rangle\} \\ \ell : \text{loc}\{b : \text{chan}\langle \text{loc}[\text{chan}\langle \text{bool} \rangle]\rangle\} \\ m : \text{loc}\{d : \text{chan}\langle \text{loc}[\text{chan}\langle \text{bool} \rangle]\rangle\} \end{array} \right\}$$

Let N be the following system:

$$\begin{array}{l} k \llbracket \text{go } m. d! \langle k[a] \rangle \rrbracket \\ | m \llbracket d?(z[x]) \text{ go } \ell. b! \langle z[x] \rangle \rrbracket \\ | \ell \llbracket b?(z[x]) \text{ go } z. x! \langle t \rangle \rrbracket \end{array}$$

Here k communicates the name of its integer channel a to m , using the channel d local to m . Then m misinforms ℓ about the type of a at k : the communication along b fools ℓ into believing that a is a boolean channel. Subsequently ℓ attempts to send an agent to k that violates the type of local resource a , by sending a boolean value where an integer is expected.

The reader can check that according to our runtime semantics the first code movement between k and m is allowed:

$$\Delta \triangleright k \llbracket \text{go } m. d! \langle k[a] \rangle \rrbracket \mapsto m \llbracket d! \langle k[a] \rangle \rrbracket$$

as local type checking of the migrating agent succeeds, $\Delta(m) \vdash d! \langle k[a] \rangle$. The local channel d is used correctly and since the value transmitted, $k[a]$, is non-local it is essentially not examined (only the number of names is checked, not their types).

The local communication at m on channel d now occurs and the second code movement between m and ℓ is also allowed,

$$\Delta \triangleright m \llbracket d! \langle k[a] \rangle \rrbracket | m \llbracket d?(z[x]) \text{ go } \ell. b! \langle z[x] \rangle \rrbracket \mapsto \cdot \mapsto \ell \llbracket b! \langle k[a] \rangle \rrbracket$$

because the migrating thread, $b! \langle k[a] \rangle$, is also successful in its type check against local resources, $\Delta(\ell)$. The local communication along b now occurs

$$\Delta \triangleright \ell \llbracket b! \langle k[a] \rangle \rrbracket | \ell \llbracket b?(z[x]) \text{ go } z. x! \langle t \rangle \rrbracket \mapsto \ell \llbracket \text{go } k. a! \langle t \rangle \rrbracket$$

However the next potential move, the migration of the thread $a! \langle t \rangle$ from ℓ to k , is disallowed by the rule (r-move); the thread is locally type checked against the resources at k , where a is known to be an integer channel, and its potential misuse is discovered.

Table 5 Static Typing

Values and threads: As in Table 4

Systems:

$$\begin{array}{c}
\text{(t-rung)} \quad \frac{\mathbb{L} \vdash P}{\Gamma, \ell : \mathbb{L} \vdash \ell \llbracket P \rrbracket} \qquad \text{(t-runb)} \quad \frac{\ell \notin \text{dom}(\Gamma)}{\Gamma \vdash \ell \llbracket P \rrbracket} \\
\text{(t-newcg)} \quad \frac{\Gamma, \ell : (\mathbb{L}, a : \mathbf{A}) \vdash N}{\Gamma, \ell : \mathbb{L} \vdash (\nu_{\ell} a : \mathbf{A}) N} \qquad \text{(t-newcb)} \quad \frac{\ell \notin \text{dom}(\Gamma) \quad \Gamma \vdash N}{\Gamma \vdash (\nu_{\ell} a : \mathbf{A}) N} \\
\text{(t-newlg)} \quad \frac{\Gamma, \ell : \mathbb{L}, k : \mathbf{K} \vdash N}{\Gamma, \ell : \mathbb{L} \vdash (\nu_{\ell} k : \mathbf{K}) N} \qquad \text{(t-newlb)} \quad \frac{\ell \notin \text{dom}(\Gamma) \quad \Gamma \vdash N}{\Gamma \vdash (\nu_{\ell} k : \mathbf{K}) N} \\
\text{(t-str)} \quad \frac{\Gamma \vdash M, N}{\Gamma \vdash \mathbf{0}, M \mid N}
\end{array}$$

Configurations:

$$\text{(t-config)} \quad \frac{\Gamma \vdash N}{\Gamma \vdash \Delta \triangleright N} \quad \forall \ell \in \text{dom}(\Gamma) : \Gamma(\ell) \leq \Delta(k)$$

4 Static typing

In the runtime semantics misuse of local resources can certainly occur, since there is no requirement that values have the object type specified by the transmitting channel. For example the reduction

$$\Delta \triangleright k \llbracket a! \langle t \rangle \rrbracket \mid k \llbracket a?(x) Q \rrbracket \mapsto k \llbracket \text{stop} \rrbracket \mid k \llbracket Q \{t/x\} \rrbracket$$

is allowed even if the object type of the channel a at k is int , *i.e.* $\Delta(k) \leq \text{loc}\{a : \text{chan}\langle \text{int} \rangle\}$.

We do not assume that all sites respect the typing constraints on their channels. For convenience let us call sites that violate the typing constraints *bad* sites, as they do not play according to the rules. A typical example is the site m , described at the end of the previous section; it receives an integer channel a from k but then attempts to use a to send a boolean value. In contrast a *good* site is one where typing constraints are enforced.

In this section we present a static type system that guarantees that:

good sites cannot be harmed by *bad* sites.

That is, local resources at a good site cannot be misused despite the existence of, and interaction with, bad sites. We prove Subject Reduction and Type Safety theorems for the type system. Intuitively, Subject Reduction can be interpreted as saying that the integrity of good sites is maintained as computation proceeds, while Type Safety demonstrates that local resources at good sites cannot be misused.

The static typing relation for anonymous networks is defined in Table 5. Judgments are of the form

$$\Gamma \vdash N$$

where Γ is a (v-open) type environment and N a system. The type environment only records the types of good locations; thus $k \in \text{dom}(\Gamma)$ is to be read “ k is good” and $m \notin \text{dom}(\Gamma)$ may be read as “ m is bad.” If $\Gamma \vdash N$, then those agents in N that are located at sites in the domain of Γ are guaranteed to be “well behaved”.

For threads and values, the static typing relation is the same as runtime typing relation given in Table 4. The typing of a located agent $\ell[[P]]$ depends on whether ℓ is good, i.e. $\ell \in \text{dom}(\Gamma)$. If ℓ is good, then according to (t-rung), to infer $\Gamma \vdash \ell[[P]]$ we need to establish that the P is well-typed to run at ℓ , $\Gamma(\ell) \vdash P$. But if ℓ is not good, then it doesn’t matter whether P is well-typed; thus (t-runb) says that if $\ell \notin \text{dom}(\Gamma)$, then $\ell[[P]]$ is well-typed for any thread P .

The typing of new resources at the systems level also depends on whether they are introduced at a good site. For example to infer $\Gamma, \ell:\mathbb{L} \vdash (v_\ell a:A) N$, where because of the notation we know ℓ is a good site, (t-newcg) says that we need to establish $\Gamma, \ell:(\mathbb{L}, a:A) \vdash N$, i.e. that N is well-typed under the assumption that there is a new resource a at site ℓ . This is in contrast to the case when ℓ is not good, $\ell \notin \text{dom}(\Gamma)$, where according to (t-newcb) it is sufficient to establish that the system N is well-typed relative to the same typing environment. Note that implicit in these two rules is the assumption that new good sites can only be generated by good sites.

Intuitively $\Gamma \vdash N$ means that in N the good sites, those in the domain of Γ use their local resources in accordance with Γ , whereas the behavior of bad sites are unconstrained. As an example consider the configuration $\Delta \triangleright N$ from Section 3.2. If we let Γ be the typing environment defined by:

$$\Gamma = \left\{ \begin{array}{l} k : \text{loc}\{a : \text{chan}\langle \text{int} \rangle\} \\ \ell : \text{loc}\{b : \text{chan}\langle \text{loc}[\text{chan}\langle \text{bool} \rangle]\} \end{array} \right\} \quad (*)$$

Then one can easily check that $\Gamma \vdash \Delta \triangleright N$, that is Δ is consistent with Γ and $\Gamma \vdash N$. Intuitively here we are saying that k and ℓ are good sites which use their local resources correctly whereas no guarantee is made about the local behavior at m . Note, however, that if one considers static typing under Δ , which includes m , then Δ does not type N .

The static typing system satisfies several important properties, such as type specialization and narrowing, which are stated and proved in Appendix A. These properties are used to establish the following result, which states that well-typing at good sites is preserved by reduction.

THEOREM 1 (SUBJECT REDUCTION).

If $\Gamma \vdash \Delta \triangleright N$ and $\Delta \triangleright N \mapsto N'$ then $\Gamma \vdash \Delta \triangleright N'$.

Proof. See Appendix A. □

We now discuss the extent to which in the type system precludes the misuse of local resources at good sites. This can be formalized using a notion of runtime error. In this paper, we confine our attention to runtime errors based on arity mismatching. Intuitively, an arity mismatch occurs when the value sent on a channel does not match the type that the recipient expects, or when two structurally dissimilar values are compared using the match construct. To formalize this notion, we define a compatibility relation $U \simeq T$ between (closed) values and types, and a compatibility relation $U \simeq V$ between (closed) values.

Table 6 Compatibility and Runtime Error

Compatibility:	
$i \asymp \text{int}$	$i \asymp i'$
$\text{bv} \asymp \text{bool}$	$\text{bv} \asymp \text{bv}'$
$e \asymp A$	$e \asymp e'$
$e[d_1, \dots, d_n] \asymp K[A_1, \dots, A_n]$	$e[d_1, \dots, d_n] \asymp e'[d'_1, \dots, d'_n]$
$(V_1, \dots, V_n) \asymp (T_1, \dots, T_n)$	$(V_1, \dots, V_n) \asymp (V'_1, \dots, V'_n)$
if $\forall i: V_i \asymp T_i$	if $\forall i: V_i \asymp V'_i$
Runtime error:	
(e-comm) $\ell \llbracket a!(V)P \rrbracket \mid \ell \llbracket a?(X:T)Q \rrbracket \xrightarrow{\text{err}\ell}$	if $V \not\asymp T$
(e-eq) $\ell \llbracket \text{if } U = V \text{ then } P \text{ else } Q \rrbracket \xrightarrow{\text{err}\ell}$	if $U \not\asymp V$
(e-new) $\frac{N \xrightarrow{\text{err}k}}{(\nu_\ell e:E)N \xrightarrow{\text{err}k/\ell}}$	(e-str) $\frac{N \xrightarrow{\text{err}k}}{M \mid N \xrightarrow{\text{err}k}} \quad \frac{N \equiv M \quad M \xrightarrow{\text{err}k}}{N \xrightarrow{\text{err}k}}$

The definitions of compatibility and of runtime error are given in Table 6. Runtime error is defined as a predicate $N \xrightarrow{\text{err}\ell}$, indicating that N is capable of a runtime error at location ℓ . Essentially an error occurs at location ℓ if either two incompatible values are compared at ℓ or an attempt is made to communicate a value along a local channel which is incompatible with the type of the channel. The only non-trivial rule is (e-new) which may effect a change in the name of the location where the error occurs. For example it will report an error at ℓ in $(\nu_\ell k:K)N$ if there is a runtime error in N at location k .

We should point out that this definition of run-time error is considerably weaker than that employed in [6]; in that paper, the notion of run-time error took into account not only to arity mismatches but also *access violations*.

THEOREM 2 (TYPE SAFETY). *If $\Gamma \vdash N$ and $\ell \in \text{dom}(\Gamma)$ then $N \not\xrightarrow{\text{err}\ell}$.*

Proof. See Appendix A. □

This theorem, together with Subject Reduction, can be interpreted informally as saying that as reductions proceed local resources cannot be locally misused at good sites, even in systems where not all sites which are necessarily well-behaved.

5 Discussion

Here we make three points which demonstrate the limitations of both the runtime strategy and the typing system.

First the static typing is very weak as it is designed only to eliminate *local* misuse of *local* resources. Let

$$\Gamma = \left\{ \begin{array}{l} \ell : \text{loc} \left\{ \begin{array}{l} b : \text{chan} \langle \text{loc} \{ d : \text{chan} \langle \text{bool} \rangle \} \rangle \\ c : \text{chan} \langle \text{loc} \{ d : \text{chan} \langle \text{int} \rangle \} \rangle \end{array} \right\} \\ k : \text{loc} \{ d : \text{chan} \langle \text{int} \rangle \} \end{array} \right\}$$

Then $\Gamma \vdash \ell \llbracket c?(z[x]) b!(z[x]) \rrbracket$ although obviously here there is, at least informally, a misuse of local channels. However formally the values exchanged on c and b are *potentially remote* and therefore these values are not considered local. For example, consider a companion site k which has one local channel d of type $\text{chan}\langle \text{int} \rangle$. Then no runtime error occurs in the system

$$\Gamma \triangleright \left(\ell \llbracket c?(z:\text{loc}\{d:\text{chan}\langle \text{int} \rangle\}) b!(z) \rrbracket \mid k \llbracket \text{go } \ell. (c!\langle k \rangle b?(z) \text{go } z. d!\langle t \rangle) \rrbracket \right)$$

although one could argue that at location ℓ there is a misuse of the channel b . A runtime error is avoided by the dynamic type-checking; after the communications on c and b the potential move

$$\ell \llbracket \text{go } k. d!\langle t \rangle \rrbracket \mapsto k \llbracket d!\langle t \rangle \rrbracket$$

is blocked because the agent attempting to move to k , $d!\langle t \rangle$ does not type check against the local resources at k .

Second, the requirement to dynamically type-check *all* incoming threads is very inefficient; however, the weak typing system makes it essential. Purely local type-checking makes it very difficult to introduce *trust* into the system. A site cannot even trust itself! For example suppose we revised the reduction rule (r -move) to read as follows:

$$\Delta \triangleright \ell \llbracket \text{go } k. P \rrbracket \mapsto \Delta \triangleright k \llbracket P \rrbracket \quad \text{if } k = \ell \text{ or } \Delta(k) \vdash P$$

Here the site k trusts itself and therefore does not type check the thread P . However this rule is not safe; Subject Reduction fails and runtime errors may be introduced. As an example, consider the following configuration, which uses the typing environment Γ given at the beginning of this section:

$$\Gamma \triangleright \ell \llbracket c?(z) \text{go } z. c!\langle t \rangle \rrbracket \mid \ell \llbracket c!\langle \ell \rangle \rrbracket$$

The configuration can be typed with respect to Γ itself, but after the communication the result is $\Gamma \triangleright \ell \llbracket c!\langle t \rangle \rrbracket$, which fails to type under Γ . Moreover $\Gamma \triangleright \ell \llbracket c!\langle t \rangle \rrbracket$ induces a runtime error due to the potential misuse of channel c . A related phenomenon is the potential misuse of channel names as locations; *e.g.* the thread $(\text{va}:\text{chan}\langle \text{int} \rangle) \text{go } a. P$ is typable in our system.

Third, the system relies heavily on dynamic type-checking to avoid misuse of resources. This can be highlighted by considering another desirable property of a runtime semantics for open systems:

Movement between good sites should always be allowed, even in the presence of badly behaved sites.

It is obvious from the example discussed in Section 3.2 that our run-time semantics does not satisfy this property. If we use the static environment Γ defined in (*) (Section 4) then k and ℓ are *good* sites. But as we have seen in Section 3.2 the intervention of m eventually prevents a movement from k and ℓ .

In a companion paper [12] we address these issues by strengthening the typing system. The second concern raised here (the inability to trust oneself or to ensure that

channels are not used as locations) can be addressed by reformalizing the typing system while retaining its basic “local” character; such a reformalization is given in the appendix of [12].

In [12], we go further, however, extending the notion of *trust* to arbitrary collections of sites. Dynamic typechecking is strengthened by making a site record information about all sites in the system, not only itself. While typechecking in this system is computationally more expensive, not all incoming threads need be checked; those originating at trusted sites are allowed through unchecked. This new semantics addresses the first two of the above concerns fully, and the third partially: a stronger notion of type safety is guaranteed, some incoming threads are not typechecked, and movement between mutually trusted sites is always allowed, although movement between good sites may not be.

6 Related Work

In this paper we have outlined a strategy for ensuring that the integrity of well-behaved sites is not compromised by the presence of potentially malicious mobile agents. Moreover we have formalized the correctness of this strategy in terms of Subject Reduction and Type Safety theorems for a *partial* type system.

In this study we used the language $D\pi$ [6], one of a number of distributed versions of the π -calculus [8]. For other variations see [5, 13]. The languages in [3, 4] are thematically similar although based on somewhat different principles. We have taken advantage of a rich type system for $D\pi$, originally presented in [6], where not only do channels have the types originally proposed in [10] for the π -calculus, but locations have types broadly similar to those of objects. An even richer type system is also proposed in [6] in which types correspond to *capabilities*, as in [4], and an interesting topic for future research would be the extension of partial typing to these richer types.

Our research is related to proposals for *proof-carrying code* outlined in [9]: code consumers, which in our case are locations, demand of code producers, in our case incoming threads, that their code is accompanied by a proof of correctness. This proof is checked by the consumer before the code is allowed to execute. The correctness is expressed in terms of a public safety policy announced by the consumer and the producer must provide code along with a proof that it satisfies this policy. In our case this safety policy is determined by the location type which records the types of the consumer’s resources, and proof checking corresponds to type checking the incoming code against this record. Our work is different in that the correctness proof can be reconstructed efficiently, and therefore the producer need not supply an explicit proof.

For other examples of related work within this framework see [7, 14]. For example the former contains a number of schemes for typechecking incoming code for access violations to local private resources. However the language is very different from ours, namely a sequential higher-order functional language, and there is no direct formalization of the fact that distributed systems which employ these schemes are well-behaved.

A very different approach to system security is based on the use of cryptography and signatures. For example [1] presents a π -calculus based language which contain

cryptographic constructs which ensure the exchange of data between trusted agents, while [2] contains a description of the application of this approach in a practical setting.

A Proofs

A.1 Properties of the static type system

First we prove two important properties of type systems with subtyping: Type Specialization and Weakening.

PROPOSITION 3 (TYPE SPECIALIZATION). *If $\mathbb{L} \vdash V:T$ and $T \leq S$ then $\mathbb{L} \vdash V:S$.*

Proof. By induction on the judgment $\mathbb{L} \vdash V:T$. If $V:T$ takes the form $V:H$ then S must coincide with H , since there is no non-trivial subtyping on channel types or base types. If $V:T$ has the form $w[\tilde{u}]:L[\tilde{A}]$ then the result is trivial, using (t-loc). Finally, the case for tuples follows by induction. \square

PROPOSITION 4 (WEAKENING).

- If $\mathbb{L} \vdash V:T$ and $\mathbb{K} \leq \mathbb{L}$ then $\mathbb{K} \vdash V:T$
- If $\mathbb{L} \vdash P$ and $\mathbb{K} \leq \mathbb{L}$ then $\mathbb{K} \vdash P$
- If $\Gamma, w:\mathbb{L} \vdash N$ and $\mathbb{K} \leq \mathbb{L}$ then $\Gamma, w:\mathbb{K} \vdash N$

Proof. In each case the proof is by induction on the type inference. We examine two examples of proof on threads:

- (t-r). Here $\mathbb{L} \vdash u?(X:T)Q$ because $\mathbb{L} \vdash u:\text{chan}\langle T \rangle$ and $\mathbb{L}, X:T \vdash Q$. We can apply the first statement in the proposition to the former, to obtain $\mathbb{K} \vdash u:\text{chan}\langle T \rangle$, while induction to the latter gives $\mathbb{K}, X:T \vdash Q$. An application of (t-r) now gives the required $\mathbb{K} \vdash u?(X:T)Q$.
- (t-newc). Here $\mathbb{L} \vdash (va:A)P$ because $\mathbb{L}, a:A \vdash P$. By α -conversion we can choose a so that it does not appear in \mathbb{K} and therefore by induction we have $\mathbb{K}, a:A \vdash P$. Now an application of (t-newc) gives the required $\mathbb{K} \vdash (va:A)P$.

We present four cases for the proof on systems.

- (t-rung). Here $\Gamma, w:\mathbb{L} \vdash m[P]$ because $\mathbb{M} \vdash P$, where $\mathbb{M} \stackrel{\text{def}}{=} (\Gamma, w:\mathbb{L})(m)$. If m and w are different then we also have $\mathbb{M} = (\Gamma, w:\mathbb{K})(m)$ and therefore an application of (t-rung) gives the required $\Gamma, w:\mathbb{K} \vdash m[P]$. On the other hand if m is the same as w then $\mathbb{M} = \mathbb{L}$. So we can apply the second part of the proposition to \mathbb{M} , obtaining $\mathbb{K} \vdash P$. Now (t-rung) also gives the required $\Gamma, w:\mathbb{K} \vdash m[P]$.
- (t-runb). This case is trivial.
- (t-newlg). Here $\Gamma, w:\mathbb{L} \vdash (v_\ell m:M)N$ because $\ell \in \text{dom}(\Gamma, w:\mathbb{L})$ and $\Gamma, w:\mathbb{L}, m:M \vdash N$. Applying induction we obtain $\Gamma, w:\mathbb{K}, m:M \vdash N$. Now (t-newlg) can be applied since $\ell \in \text{dom}(\Gamma, w:\mathbb{K})$, to obtain the required $\Gamma, w:\mathbb{K} \vdash (v_\ell m:M)N$.
- (t-newcb). Here $\Gamma, w:\mathbb{L} \vdash (v_\ell a:A)N$ because $\ell \notin \text{dom}(\Gamma, w:\mathbb{L})$ and $\Gamma, w:\mathbb{L} \vdash N$. However we also have $\ell \notin \text{dom}(\Gamma, w:\mathbb{K})$ and therefore (t-newcb) can also be applied to obtain the required $\Gamma, w:\mathbb{K} \vdash (v_\ell a:A)N$. \square

The following Restriction Lemma states that if $\Gamma \vdash N$ and some identifier u does not occur free in N then N can also be typed in an environment obtained from Γ by removing all occurrences of u . For any identifier u let $\Gamma \setminus u$ denote the result of removing all occurrences of u from Γ . For example $(\Gamma, u:\mathbb{L}) \setminus u$ denotes Γ while $(\Gamma, w:(\mathbb{L}, u:\mathbb{A})) \setminus u$ is the same as $(\Gamma \setminus u), w:\mathbb{L}$.

LEMMA 5 (RESTRICTION).

- If $\mathbb{L}, v:\mathbb{H} \vdash U:\mathbb{T}$ and $v \notin \text{fid}(U)$ then $\mathbb{L} \vdash U:\mathbb{T}$
- If $\mathbb{L}, v:\mathbb{H} \vdash P$ and $v \notin \text{fid}(P)$ then $\mathbb{L} \vdash P$
- If $\Gamma \vdash N$ and $v \notin \text{fid}(N)$ then $\Gamma \setminus v \vdash N$.

Proof. By induction on the proof of the typing judgment. □

The following corollary follows by an easy induction on V .

COROLLARY 6.

- If $\mathbb{L}, V:\mathbb{S} \vdash U:\mathbb{T}$ and $\text{fid}(V)$ disjoint $\text{fid}(U)$ then $\mathbb{L} \vdash U:\mathbb{T}$
- If $\mathbb{L}, V:\mathbb{S} \vdash P$ and $\text{fid}(V)$ disjoint $\text{fid}(P)$ then $\mathbb{L} \vdash P$ □

A.2 Subject Reduction

We first show that typing is preserved by the structural equivalence. The most complicated case is already covered by the Restriction proposition.

LEMMA 7 (SCOPE EXTRUSION).

If $e \notin \text{fn}(M)$ then $\Gamma \vdash M \mid (v_\ell e:\mathbb{T})$ if and only if $\Gamma \vdash (v_\ell e:\mathbb{T})(M \mid N)$

Proof. There are two cases. If $\ell \notin \text{dom}(\Gamma)$, then we can reason as follows:

$$\begin{aligned} \Gamma \vdash (v_\ell e:\mathbb{T})(M \mid N) &\iff \Gamma \vdash N \text{ and } \Gamma \vdash M \\ &\iff \Gamma \vdash N \mid (v_\ell e:\mathbb{T})M \end{aligned}$$

In the case that $\ell \in \text{dom}(\Gamma)$, the argument is slightly different depending on whether e is a channel or a location. As an example we consider the former, and we assume Γ has the form $\Delta, \ell:\mathbb{L}$.

$$\begin{aligned} \Gamma \vdash (v_\ell a:\mathbb{A})(M \mid N) &\iff \Delta, \ell:(\mathbb{L}, a:\mathbb{A}) \vdash (M \mid N) \\ &\iff \Delta, \ell:(\mathbb{L}, a:\mathbb{A}) \vdash M \text{ and } \Delta, \ell:(\mathbb{L}, a:\mathbb{A}) \vdash N \\ &\iff \Delta, \ell:\mathbb{L} \vdash M \text{ and } \Delta, \ell:(\mathbb{L}, a:\mathbb{A}) \vdash N \text{ by Restriction} \\ &\iff \Gamma \vdash M \text{ and } \Gamma \vdash (v_\ell a:\mathbb{A})N \\ &\iff \Gamma \vdash M \mid (v_\ell a:\mathbb{A})N \quad \square \end{aligned}$$

PROPOSITION 8. If $N \equiv M$ then $\Gamma \vdash N$ implies $\Gamma \vdash M$

Proof. By induction on the proof that $N \equiv M$. All of the rules and most of the axioms are very easy to handle; the most difficult case is (s-extr), which follows by the previous lemma. As an example, we consider the rule (s-new): $\ell \llbracket (v_\ell e:\mathbb{T})P \rrbracket \equiv (v_\ell e:\mathbb{T})\ell \llbracket P \rrbracket$.

If $\ell \notin \text{dom}(\Gamma)$ then it is immediate as both sides trivially type with respect to Γ . So suppose $\ell \in \text{dom}(\Gamma)$, and, as an example, that $e:T$ is $a:A$. Let Γ have the form $\Delta, \ell:\mathbb{L}$. Then $\Gamma \vdash \ell[(\text{va}:A)P]$ if and only if $\mathbb{L}, a:A \vdash P$ while

$$\begin{aligned} \Gamma \vdash (\text{v}_\ell a:A) \ell[P] &\iff \Delta, \ell:(\mathbb{L}, a:A) \vdash \ell[P] \\ &\iff \mathbb{L}, a:A \vdash P \end{aligned} \quad \square$$

As is normally the case the proof of Subject Reduction depends on the fact that, in some sense, typing is preserved by substitution. To prove this fact the following lemma will be useful:

LEMMA 9. *If $\mathbb{L} \vdash V:S$ and $\mathbb{L}, X:S \vdash U:T$ then $\mathbb{L} \vdash U\{V/x\}:T$.*

Proof. By induction on the structure of X .

Suppose X is a variable, *i.e.* $X = x$ for some x . We then proceed by induction on U . If U is a base value or has the form $u[\tilde{v}]$, then the result is trivial. If U is a tuple we can apply the inner induction hypothesis. If U is x , then $U\{V/x\} = V$ and, by the typing rules, $S \leq T$; using type specialization, we can therefore conclude $\mathbb{L} \vdash U\{V/x\}:T$, as required. If U is an identifier u that is different from x then $x \notin \text{fid}(U)$ and $U\{V/x\} = U$; therefore by the Restriction Lemma $\mathbb{L} \vdash U\{V/x\}:T$, as required.

Suppose X has the form $z[\tilde{x}]$, so $\mathbb{L}, X:S = \mathbb{L}$ and $\text{fid}(X)$ disjoint $\text{dom}(\mathbb{L})$ and therefore $U\{V/x\} = U$. By Corollary 6, $\mathbb{L} \vdash U:T$.

If X is a tuple, the result follows by induction. □

PROPOSITION 10 (SUBSTITUTION). *If $\mathbb{L} \vdash V:T$ and $\mathbb{L}, X:T \vdash P$ then $\mathbb{L} \vdash P\{V/x\}$.*

Proof. For convenience we use P' to denote $P\{V/x\}$ (and employ similar notation for other syntactic categories).

As before, the proof is by induction on the structure of X . The cases for tuples \tilde{Y} and structured values $z[\tilde{x}]$ are as before. We present the case where $X = x$, for some variable x . In this case we proceed by a second induction on the inference $\mathbb{L}, x:T \vdash P$. We present the cases that involve binders.

(t-r). Here $\mathbb{L}, x:T \vdash u?(Y:S)Q$ because $\mathbb{L}, x:T \vdash u:\text{chan}\langle T \rangle$ and $\mathbb{L}, x:T, Y:S \vdash Q$. Applying the previous lemma to $\mathbb{L}, x:T \vdash u:\text{chan}\langle T \rangle$, we obtain $\mathbb{L} \vdash u':\text{chan}\langle T \rangle$. In addition, $\mathbb{L}, x:T, Y:S \vdash Q$ can be rewritten as $\mathbb{L}, Y:S, x:T \vdash Q$. By the internal induction, we obtain $\mathbb{L}, Y:S \vdash Q'$. Now (t-r) gives the required $\mathbb{L} \vdash u'?(Y:S)Q'$.

(t-newc). Here $\mathbb{L}, x:T \vdash (\text{va}:A)P$ because $\mathbb{L}, x:T, a:A \vdash P$. The environment can be rewritten as $\mathbb{L}, a:A, x:T$, and therefore induction can be applied to obtain $\mathbb{L}, a: \vdash P'$. The result follows by an application of (t-newc). □

THEOREM (1, SUBJECT REDUCTION).

If $\Gamma \vdash \Delta \triangleright N$ and $\Delta \triangleright N \mapsto N'$ then $\Gamma \vdash \Delta \triangleright N'$.

Proof. From the hypothesis $\Gamma \vdash \Delta \triangleright N$ we know that $\Gamma \vdash N$ and $\Gamma(\ell) \leq \Delta(\ell)$ for every ℓ in $\text{dom}(\Gamma)$. So it is sufficient to show $\Gamma \vdash N'$, which we do by induction on the derivation of $\Delta \triangleright N \mapsto N'$. We present a number of representative cases.

(r-move). Here we have $\Delta \triangleright \ell \llbracket \text{gok}.P \rrbracket \longrightarrow k \llbracket P \rrbracket$ and $\Delta(k) \vdash P$. If $k \notin \text{dom}(\Gamma)$ then the required $\Gamma \vdash k \llbracket P \rrbracket$ follows trivially from (t-runb). If $k \in \text{dom}(\Gamma)$ then we know that $\Gamma(k) \leq \Delta(k)$ and therefore, by Weakening, $\Gamma(k) \vdash P$. Now (t-rung) can be applied to obtain $\Gamma \vdash k \llbracket P \rrbracket$.

(r-comm). Here we have

$$\Delta \triangleright \ell \llbracket a!(V)P \rrbracket \mid \ell \llbracket a?(X:T)Q \rrbracket \longrightarrow \ell \llbracket P \rrbracket \mid \ell \llbracket Q\{V/X\} \rrbracket$$

If $\ell \notin \text{dom}(\Gamma)$ then the result is trivial from (t-runb). Otherwise $\ell \in \text{dom}(\Gamma)$. From $\Gamma \vdash \ell \llbracket a!(V)P \rrbracket \mid \ell \llbracket a?(X:T)Q \rrbracket$ we know $\Gamma \vdash \ell \llbracket a!(V)P \rrbracket$ and therefore $\Gamma(\ell) \vdash P$. It follows that $\Gamma \vdash \ell \llbracket P \rrbracket$.

It remains to show that $\Gamma \vdash \ell \llbracket Q\{V/X\} \rrbracket$, that is $\Gamma(\ell) \vdash Q\{V/X\}$. Again from the hypothesis we know $\Gamma \vdash \ell \llbracket a?(X:T)Q \rrbracket$ from which we can conclude that $\Gamma(\ell) \vdash u:\text{chan}\langle T \rangle$ and $\mathbb{L}, X:T \vdash Q$. From $\Gamma \vdash \ell \llbracket a!(V)P \rrbracket$ we know that $\Gamma(\ell) \vdash V:S$ for some S for which we also have $\Gamma(\ell) \vdash u:\text{chan}\langle S \rangle$. In our typing system this must mean that S and T coincide. We may therefore apply the Substitution lemma to obtain the required $\Gamma(\ell) \vdash Q\{V/X\}$.

(r-new). We consider the case:

$$\Delta, \ell:\mathbb{L} \triangleright (v_\ell a:A)N \longrightarrow (v_\ell a:A)N' \quad \text{because} \quad \Delta, \ell:(\mathbb{L}, a:A) \triangleright N \longrightarrow N'$$

First suppose $\ell \in \text{dom}(\Gamma)$. Since $\Gamma \vdash \Delta, \ell:\mathbb{L} \triangleright (v_\ell a:A)N$ we know Γ can be written as $\Gamma', \ell:\mathbb{L}'$, where $\mathbb{L}' \leq \mathbb{L}$ and therefore $\Gamma', \ell:(\mathbb{L}', a:A) \vdash N$. We can now apply induction to obtain $\Gamma', \ell:(\mathbb{L}', a:A) \vdash N'$, to which (t-newcg) can be applied to obtain the required $\Gamma \vdash (v_\ell a:A)N'$.

If $\ell \notin \text{dom}(\Gamma)$ then by (t-newcb) it is sufficient to prove $\Gamma \vdash N'$. In this case $\Gamma \vdash \Delta, \ell:\mathbb{L} \triangleright (v_\ell a:A)N$ yields $\Gamma \vdash \Delta, \ell:(\mathbb{L}, a:A) \triangleright N$ to which induction can be applied to give the required $\Gamma \vdash N'$.

(r-str). This case follows using induction and Proposition 8. \square

A.3 Type Safety

We first show that the typing system is “compatible” with the compatibility relation \asymp .

LEMMA 11.

- $\mathbb{L} \vdash V:T$ implies $V \asymp T$
- $\mathbb{L} \vdash V:T$ and $\mathbb{L} \vdash U:T$ implies $V \asymp U$

Proof. A straightforward inductive argument, in the first case on the derivation of $\mathbb{L} \vdash V:T$ and in the second on the structure of the type T . \square

THEOREM (2, TYPE SAFETY). *If $\Gamma \vdash N$ and $\ell \in \text{dom}(\Gamma)$ then $N \xrightarrow{\text{err}\ell}$.*

Proof. By induction on the proof that $N \xrightarrow{\text{err}\ell}$, we show that if $\ell \in \text{dom}(\Gamma)$ and $N \xrightarrow{\text{err}\ell}$ then $\Gamma \not\vdash N$, which is sufficient to establish the theorem. Let \mathbb{L} denote $\Gamma(\ell)$.

(e-comm). In this case we have $\ell \llbracket a!(V)P \rrbracket \mid \ell \llbracket a?(X:T)Q \rrbracket \xrightarrow{\text{err}\ell}$ because $V \not\asymp T$. Now suppose, for a contradiction, that $\Gamma \vdash \ell \llbracket a!(V)P \rrbracket \mid \ell \llbracket a?(X:T)Q \rrbracket$. Because there is no subtyping on channel types we know $\mathbb{L} \vdash \text{chan}\langle T \rangle$ and $\mathbb{L} \vdash \text{chan}\langle S \rangle$ implies $T = S$. From the alleged typing we can therefore conclude that $\mathbb{L} \vdash V:T$. By the first part of Lemma 11, we have $V \asymp T$ which contradicts $V \not\asymp T$.

(e-eq). Here we have $\ell[\text{if } U = V \text{ then } P \text{ else } Q] \xrightarrow{\text{err}\ell}$ because $U \not\approx V$. If we assume $\Gamma \vdash \ell[\text{if } U = V \text{ then } P \text{ else } Q]$ then we must have $\Gamma \vdash U:T$ and $\Gamma \vdash V:T$ for some T . Now applying the second part of Lemma 11 we obtain a contradiction to $U \not\approx V$.

(e-new). First suppose that $(\nu_k e:E)N \xrightarrow{\text{err}\ell}$ because $N \xrightarrow{\text{err}\ell}$ and $e \neq \ell$. If $k \notin \text{dom}(\Gamma)$ then by induction we have $\Gamma \not\vdash N$ and by either (t-newcb) or (t-newlb) we can conclude $\Gamma \not\vdash (\nu_k e:E)N$.

So suppose $k \in \text{dom}(\Gamma)$. As an example suppose $e:E$ is a location; the case when it is a channel is similar. Applying induction we have $\Gamma, e:E \not\vdash N$ and therefore by (t-newlg) we can conclude $\Gamma \not\vdash (\nu_k e:E)N$.

Finally suppose $(\nu_\ell k:E)N \xrightarrow{\text{err}\ell}$ because $N \xrightarrow{\text{err}k}$. By α -conversion we can assume that k does not appear in Γ and so we can apply induction to obtain $\Gamma, k:E \not\vdash N$. We know that $\ell \in \text{dom}(\Gamma)$ and therefore by (t-newlg) we can conclude $\Gamma \not\vdash (\nu_\ell k:E)N$.

(e-str). Straightforward, using induction and Proposition 8. \square

Acknowledgements

We thank the referees for several comments that sharpened the presentation. Matthew Hennessy funded by CONFER II and EPSRC project GR/K60701. James Riely funded by NSF grant EIA-9805604.

References

- [1] M. Abadi and A.D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, To appear. Available as Compaq SRC Research Report 149 (1998).
- [2] B. Bershad, B. Savage, P. Paradyak, E. Sirer, D. Becker, M. Fiuczynski, C. Chambers, and S. Eggers. Extensibility, safety and performance in the SPIN operating system. In *Symposium on Operating Systems Principles*, pages 267–284, 1997.
- [3] L. Cardelli and A.D. Gordon. Mobile ambients. In *Foundations of Software Science and Computational Structures*, volume 1378 of *Lecture Notes in Computer Science*, pages 140–155. Berlin: Springer-Verlag, 1998.
- [4] R. DeNicola, G. Ferrari, and R. Pugliese. Types as specifications of access policies. In J. Vitek and C. Jensen, editors, *Secure Internet Programming: Security Issues for Distributed and Mobile Objects*, Lecture Notes in Computer Science. Springer-Verlag, 1999.
- [5] C. Fournet, G. Gonthier, J.J. Levy, L. Marganet, and D. Remy. 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. Berlin: Springer-Verlag.
- [6] M. Hennessy and J. Riely. Resource access control in systems of mobile agents. In U. Nestmann and B. Pierce, editors, *3rd International Workshop on High-Level Concurrent Languages (HLCL'98)*, volume 16(3) of *Electronic Notes in Theoretical Computer Science*, Nice, September 1998. Elsevier. Available from <http://www.elsevier.nl/locate/entcs>. Full version available as Sussex CSTR 98/02, 1998. Available from <http://www.informatics.sussex.ac.uk/>.

- [7] X. Leroy and F. Rouaix. Security properties of typed applets. In *Conference Record of the ACM Symposium on Principles of Programming Languages*, San Diego, January 1998. ACM Press.
- [8] R. Milner. The polyadic π -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.
- [9] G. Necula. Proof-carrying code. In *Conference Record of the ACM Symposium on Principles of Programming Languages*, Paris, January 1997. ACM Press.
- [10] 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.
- [11] J. Riely and M. Hennessy. A typed language for distributed mobile processes. In *Conference Record of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, San Diego, January 1998. ACM Press.
- [12] J. Riely and M. Hennessy. Trust and partial typing in open systems of mobile agents. In *Conference Record of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, San Antonio, January 1999. ACM Press. Full version available as Sussex CSTR 98/04, 1998. Available from <http://www.informatics.sussex.ac.uk/>.
- [13] P. Sewell. Global/local subtyping and capability inference for a distributed π -calculus. In *Proceedings of the International Colloquium on Automata, Languages and Programming*, volume 1433 of *Lecture Notes in Computer Science*. Berlin: Springer-Verlag, July 1998.
- [14] R. Stata and M. Abadi. A type system for java bytecode subroutines. In *Conference Record of the ACM Symposium on Principles of Programming Languages*, San Diego, January 1998. ACM Press.