

Global/Local Subtyping and Capability Inference for a Distributed π -calculus

Peter Sewell

University of Cambridge

Peter.Sewell@cl.cam.ac.uk

Abstract

This paper considers how locality restrictions on the use of capabilities can be enforced by a static type system. A distributed π -calculus with a simple reduction semantics is introduced, integrating location and migration primitives from the Distributed Join Calculus with asynchronous π communication. It is given a type system in which the input and output capabilities of channels may be either global, local or absent. This allows compile-time optimization where possible but retains the expressiveness of channel communication. Subtyping allows all communications to be invoked uniformly. We show that the most local possible capabilities for internal channels can be inferred automatically.

1 Introduction


A central theme in programming language and system design is that of restricting access to resources to be in some sense *local*. This can support clean design, allow efficient implementation, and provide robustness against accidental errors and malicious attacks. The development of ubiquitous networking, particularly of systems in which executing agents or executable code are communicable, brings new kinds of resource and locality to the fore. One has resources such as the capability to input or output on a communication channel, to read or write a distributed reference cell, and to decrypt or encrypt with a cryptographic key pair. It may be desirable to restrict the use of these capabilities to, for example, a single agent, a group of trusted agents, a region of the network or a machine address space.

In this paper we consider how such restrictions, particularly the first, can be enforced by a static type system. We introduce a *distributed π -calculus*, with primitives for location, migration and π -calculus style channel communication, as an idealisation of a mobile agent programming language. It is given a type system in which the input and output capabilities for a communication channel can be either *global*, and therefore usable within any location, restricted to be *local*, and therefore usable only within the location where the channel is declared, or *absent*. The type system allows local communication to be implemented efficiently, while subtyping and subsumption ensure that, from the programmer's point of view, it is not unduly restrictive. The constructs for input and output along a channel are independent of whether its capabilities are global or local, thus facilitating programming. At the same time the programmer can distinguish between local and (potentially expensive) global communications via the typing of channel declarations. For the type system to be pragmatically usable it is essential that capability annotations can often be inferred automatically – one would like internal channels to be given the most local capabilities possible, to allow optimisation. We show that this can be done, by showing that typing is preserved by least upper bounds in a modified subtype order.

The distributed π -calculus used is introduced in Section 2. It is designed to allow the global/local type system to be presented clearly; it does not address other issues, such as name services, failure and administrative domains, that arise in mobile agent programming. It builds on the π -calculus of Milner, Parrow and Walker [MPW92]. The π -calculus is often described as a calculus of mobile processes. Strictly, however, this refers to the mobility of the scopes of channel declarations – channels are statically scoped, but their scopes may change over time as channel names are sent outside their current scopes. There is no other notion of locality or of identity of processes, so to directly model distributed phenomena, such as migration of agents, failure of machines or knowledge of agents, one must add primitives for grouping π -calculus

processes, into units of migration, failure or shared knowledge respectively. This was done by Amadio and Prasad [AP94] in order to model an abstraction of the failure semantics of *Facile* [TLK96], an extension of ML with distribution primitives. More recently, Fournet et al have proposed the *Distributed Join Calculus* and a related programming language [FG96, FGL⁺96]. To model migrating agents one needs at least a two level hierarchy, with agents, each containing some processes, located at named virtual machines. The Distributed Join Calculus takes a generalisation of this, with primitive tree-structured *locations*. Immediate sublocations of the root model virtual machines; descendants of these represent hierarchically structured mobile agents. Locations are named; new locations can be created and their names can be scope-extruded by communication just as π -calculus channel names can be. Locations that are not immediate sublocations of the root may migrate, changing their parent location. The sublocations of a location thereby migrate with it. We adopt similar location and migration primitives (this choice gives a reasonably clean calculus, but it is not critical for the type system). For communication between agents a wide variety of primitives may be useful in practice. In this paper we adopt those of an asynchronous π -calculus [Bou92, HT92]. They are rather high level, operating independently of the physical locations of agents, and expressive, allowing any number of readers and writers on a channel. As in [AP94] channels have locations, giving the agents in which their queues of blocked writers or readers are stored.

The type system is introduced in Section 3. Generalising the Input/Output type system of Pierce and Sangiorgi [PS96], it has channel types annotated with capabilities. Here they are of the form $\uparrow_{io} T$, with input and output capabilities i and o each taken from $\{-, L, G\}$. The intuition is that a G capability may be used at any location, an L capability may be used only at the location of the channel concerned and a $-$ capability may not be used at all. The types $\uparrow_{--} T$ are replaced by a single top type $\top$, leaving the capabilities below.



For example, consider a channel x of type $\uparrow_{io} T$, which is located at a location named k . If $io = GG$ then x is usable at any location, both for input and for output. If $io = LL$ then x is usable only within location k , but still for both input and output. If $io = LG$ then x can be used for output anywhere, but for input only within location k . Such a channel might be used for sending requests to a server located at k . Conversely, if $io = GL$ then x can be used for input anywhere, but for output only at location k . Such a channel might be used for receiving results from servers, or for ‘pushed’ data from an information source. The tags $\pm, -, +$ of [PS96] correspond to the capabilities $GG, G-, -G$, for global communication, and to $LL, L-, -L$, for local communication.

A subtyping order is lifted from the tag ordering above (with $i-$ and $-o$ covariant and contravariant respectively), allowing subcapabilities to be communicated. For example if $x: \uparrow_{LG} T$ then x may be transmitted globally, along channels of type $\uparrow_{GG} \uparrow_{-G} T$, to readers that are guaranteed to use it only at type $\uparrow_{-G} T$.

The expressiveness of the system should aid programmers by detecting errors at compile time, including communications that inadvertently potentially involve network communication. In an implementation, channels with tags $LL, L-$ and $-L$ can be implemented with data structures that are local to a single agent, and so always on the same (albeit possibly changing) machine. Their names need not be globally unique, but only unique within their location (note that this implies that equality testing of channel names should not be available) and need not

be registered with global name services. Channels with tags **GL** (resp. **LG**) are subject to fewer optimizations, but still allow the references to the channel data structures by writers (resp. readers) to be local pointers.

The typing rules involve two novel features — the formation of certain types must be forbidden by kinding rules and the capabilities of channels must be compared with capabilities at which they can be used by readers. The main soundness result is subject reduction (Theorem 1); in addition one can see by examination of the typing rules that no well-typed process can immediately use a capability that it does not have.

In Section 4 we show that the most local possible capabilities for channels can be inferred (Theorem 2). Some related work and possible generalisations are discussed in Section 5. Proofs are omitted for lack of space. Further discussion of alternative calculi, and the proof of soundness for an extension of the type system with type variables and recursive types, may be found in [Sew97a].

2 A distributed π -calculus

In this section the syntax and operational semantics of our distributed π -calculus (dpi for short) are given. The operational semantics is a rather mild extension of that for the asynchronous π -calculus. It is a reduction semantics, defining reductions over process terms (no additional notion of configuration is required) using a structural congruence. It differs from an asynchronous π semantics in only two respects — there is a reduction rule for migration and the standard structural congruence and reduction rules are adapted to terms containing location information. Location and channel names are both subject to scope extrusion, just as π -calculus channel names are. The semantics and type system can therefore be simplified by treating new location and channel declarations similarly, taking a single binder (**new** $x : @_l T$), which declares x to be a sublocation of l (respectively a channel located at l) if the type T is the type loc of location names (respectively a channel type).

Types To the types of channels and locations introduced above we add *base types*, ranged over by B , for example a unit type 1 and Int, pairs, as a first step towards more interesting datatypes, and a type to be the top of the subtype order. The *pre-types*, ranged over by S, T, U, V , are given by

$$T ::= B \mid T \times T \mid \downarrow_{io} T \mid \text{loc} \mid \top$$

Only some pre-types will be considered well-formed. The syntax of processes involves types, and hence the reduction semantics does also. Its definition does not depend on them in any interesting way, however, so we defer the type formation rules to Section 3.

Processes We take an infinite set $\mathcal{X}$ of *names*, ranged over by a, j, k, l, x, y, z and containing a distinguished name top. We let m, n, p, q range over $\mathbb{N}$. *Values*, ranged over by u, v, w , are

$$v ::= b \mid x \mid \langle v, v \rangle$$

where b ranges over elements of the base types. There are two extremal possibilities for adding location information to terms. In one a locator applies to the largest possible unit, with all co-located subterms gathered into a single subterm. This is adopted, for example, in the *Ambient Calculus* of Cardelli and Gordon [CG98]. For dpi, however, communication is possible *across* the location tree structure, so to give a reduction semantics (in which writers and readers at different locations must be brought syntactically adjacent by a structural congruence) the other extreme is adopted, with each elementary subterm explicitly located. Accordingly, *processes*,

ranged over by P, Q, R , are:

$P ::=$	$@_u \bar{v} w$	at location u , output value w on channel v
	$@_u v(y).P$	at u , input a value from channel v and bind it to y in P
	$@_u ! v(y).P$	replicated input
	$@_u \mathbf{migrate_to} \ v \ \mathbf{then} \ P$	migrate location u to become a sublocation of v
	$@_u \mathbf{let} \ \langle y:T, y':T' \rangle = w \ \mathbf{in} \ P$	at u , bind the halves of the pair w to y and y' in P
	$(\mathbf{new} \ y : @_u T)P$	declare a new channel or location named y , of type T , located at u and binding in P
	0	the null process
	$P \mid P$	parallel composition

The names y and y' above, which must be distinct in the **let** case, bind in the respective subterms P (in particular, in $(\mathbf{new} \ y : @_u T)P$ the scope of y does not include u); we work up to alpha conversion of bound names. The free names of a value v and process P will be denoted by $\text{fn}(v)$ and $\text{fn}(P)$ respectively. The substitution of a value v for a name x in P will be written $\{v/x\}P$. Output values and input binders of type 1 will often be elided.

The syntax of processes includes some nonsensical terms, which the reduction semantics gives nonsensical reductions to. They will be formally excluded by the typing rules but two points are worth mentioning now. Firstly, in well-typed processes the u and v appearing in the grammar will always be names. They are allowed to be arbitrary values in the syntax so that substitution of values for names is always defined. Secondly, the syntax includes terms which can teleport after a prefix, e.g.

$$@_k x(y).@_l x(z).0 \quad @_k x(y).@_y x(z).0$$

For conceptual simplicity we would like migration to be the *only* way in which processes may move, and so want locators $@_l$ to describe the locations of processes rather than cause them to move. Teleporting terms are excluded by considering the set of free inhabited locations $\text{fl}(P)$ of a process P . This is the set of the free location names that are inhabited in P by outputs, inputs, migrates, pair splits, channels or locations. It is the set of names which occur free in P within the argument u of a subterm $@_u$. For example, $\text{fl}((\mathbf{new} \ j : @_l \text{loc})(@_j \bar{x} \mid @_l \bar{x} \mid @_k \bar{x})) = \{k, l\}$. The type system will require, in every prefix located at l with continuation P , that $\text{fl}(P) \subseteq \{l\}$.

Reduction semantics Structural congruence $\equiv$ is the least congruence relation over processes satisfying the following.

$$P \mid 0 \equiv P \tag{1}$$

$$P \mid Q \equiv Q \mid P \tag{2}$$

$$P \mid (Q \mid R) \equiv (P \mid Q) \mid R \tag{3}$$

$$(\mathbf{new} \ x : @_u S)(\mathbf{new} \ y : @_v T)P \equiv (\mathbf{new} \ y : @_v T)(\mathbf{new} \ x : @_u S)P \quad x \notin \text{fn}(v), y \wedge y \notin \text{fn}(u) \tag{4}$$

$$P \mid (\mathbf{new} \ x : @_v T)Q \equiv (\mathbf{new} \ x : @_v T)(P \mid Q) \quad x \notin \text{fn}(P) \tag{5}$$

The first three equations are standard, allowing parallel compositions to be treated as multisets. Equation 5 allows scope extrusion, both of channel names and of location names. Equation 4 allows new-binders to be permuted; the side condition ensures that the location tree structure, and the locations of channels, are preserved.

The reduction relation $\longrightarrow$ over processes is the least relation satisfying the following.

$$\begin{aligned} @_k \bar{x} v \mid @_l x(y).P &\longrightarrow \{v/y\}P \\ @_k \bar{x} v \mid @_l ! x(y).P &\longrightarrow \{v/y\}P \mid @_l ! x(y).P \\ @_l \mathbf{let} \ \langle y_1:T_1, y_2:T_2 \rangle = \langle v_1, v_2 \rangle \ \mathbf{in} \ P &\longrightarrow \{v_1/y_1\}\{v_2/y_2\}P \\ (\mathbf{new} \ l : @_j T)(\mathbf{new} \ \Delta)(Q \mid @_l \mathbf{migrate_to} \ k \ \mathbf{then} \ P) &\longrightarrow (\mathbf{new} \ l : @_k T)(\mathbf{new} \ \Delta)(Q \mid P) \\ &\text{if } \{k, l\} \cap \text{dom}(\Delta) = \{ \} \wedge k \neq l \end{aligned}$$

$$\frac{P \longrightarrow Q}{P \mid R \longrightarrow Q \mid R} \quad \frac{P \longrightarrow Q}{(\mathbf{new} \ x : @_l T)P \longrightarrow (\mathbf{new} \ x : @_l T)Q} \quad \frac{P \equiv P' \quad P' \longrightarrow Q' \quad Q' \equiv Q}{P \longrightarrow Q}$$

where we define $(\mathbf{new} \Delta)P$, for lists of the grammar $\Delta ::= \bullet \mid \Delta, x : @_l T$, by $(\mathbf{new} \bullet)P = P$ and $(\mathbf{new} \Delta, x : @_l T)P = (\mathbf{new} \Delta)(\mathbf{new} x : @_l T)P$. The first two reduction rules are the standard communication rules for an asynchronous π -calculus; note that the communications can take place irrespective of the locations of the writer, reader and channel. The third is an unproblematic pair splitting reduction. The fourth is the only substantially new reduction rule. It allows location l to migrate from being a sublocation of j to become a sublocation of k . After the migration the continuation P is released. The additional context $(\mathbf{new} \Delta)(Q \mid _)$, which is preserved by the reduction, is required as the scope of l may contain other location and channel declarations, and processes, that mention l . In particular, note that Q may contain other subterms $@_l \dots$ that remain located at l as it migrates. Note also that the side condition means that the rule is not applicable if k is a sublocation of l . Such migrations, which would introduce a cycle into the location tree, are blocked, although later migrations may unblock them. The last three rules are standard.

The sublocation tree of a migrating location is unchanged, and so migrates with it. The unit of migration is thus a subtree of locations with all their processes and channels. The largest unit that is guaranteed to stay together (and so always be on the same machine), however, is not a subtree but just the processes and channels at a single location — its sublocations may migrate away. The tree structure is therefore essentially orthogonal to global/local typing.

Examples We give some simple example processes that will be well-typed in the empty context. First, a server that returns the result of some computation (in this trivial example it simply pairs the argument with itself):

$$\begin{aligned} & (\mathbf{new} \text{pairServer} : @_{\text{toploc}})(\mathbf{new} \text{client} : @_{\text{toploc}}) \\ & \quad (\mathbf{new} \text{pair} : @_{\text{pairServer}} \downarrow_{\text{LG}}(\text{Int} \times \uparrow_{-\text{G}}(\text{Int} \times \text{Int}))) \\ & \quad \quad @_{\text{pairServer}} ! \text{pair}(y).@_{\text{pairServer}} \mathbf{let} \langle n : \text{Int}, c : \uparrow_{-\text{G}}(\text{Int} \times \text{Int}) \rangle = y \mathbf{in} @_{\text{pairServer}} \overline{c}\langle n, n \rangle \\ & \quad \mid \\ & \quad (\mathbf{new} c : @_{\text{client}} \downarrow_{\text{LG}}(\text{Int} \times \text{Int}))@_{\text{client}} \overline{\text{pair}}\langle 7, c \rangle \mid @_{\text{client}} c(x). \dots \end{aligned}$$

Note the use of subsumption for typing the output $@_{\text{client}} \overline{\text{pair}}\langle 7, c \rangle$. Secondly, a rudimentary tracker, that receives location names on a channel move (perhaps provided by an active badge system controller) and migrates to them:

$$\begin{aligned} & (\mathbf{new} l_1 : @_{\text{toploc}}) \dots (\mathbf{new} l_3 : @_{\text{toploc}}) \\ & (\mathbf{new} \text{controller} : @_{\text{toploc}}) \\ & \quad (\mathbf{new} \text{move} : @_{\text{controller}} \downarrow_{\text{GL}} \text{loc}) \\ & \quad \quad @_{\text{controller}} \overline{\text{move}} l_1 \mid \dots \mid @_{\text{controller}} \overline{\text{move}} l_3 \\ & \quad \mid \\ & \quad (\mathbf{new} \text{follower} : @_{\text{controllerloc}}) \\ & \quad \quad @_{\text{follower}} ! \text{move}(l).@_{\text{follower}} \mathbf{migrate_to} \ l \ \mathbf{then} \ \dots \\ & \quad \mid @_{\text{follower}} \dots \end{aligned}$$

A more realistic tracker would have additional communications so that the moves could be sequentialised. As these examples show, the syntax of processes contains redundant location information. The design of a less verbose representation, allowing co-located processes to be gathered together at compile time, is discussed in [Sew97a]. Also discussed there is a calculus that is better suited to use by programmers, allowing locators to occur more freely. Both of these require a more complex operational semantics.

Action calculus semantics There is a rather large space of possible calculi with reduction semantics. One way of understanding it, particularly for comparing different calculi, is to put them into a common framework, such as the *Action Calculi* of Milner [Mil96]. This provides a well-understood structural congruence, with a graphical intuition, that has clarified the design of dpi. As an illustration, we give an action calculus presentation of the fragment of dpi

without pairs or base type values. We take the arity monoid $(\mathbb{N}, +, 0)$, the names of arity 1 to be $\mathcal{X}$ and controls:

$$\begin{array}{c} \mathbf{new}_T : 1 \rightarrow 1 \quad \mathbf{out} : 3 \rightarrow 0 \quad \frac{a : 1 \rightarrow 0}{\mathbf{in}(a) : 2 \rightarrow 0} \quad \frac{a : 0 \rightarrow 0}{\mathbf{mig}(a) : 2 \rightarrow 0} \\ \mathbf{rep}(a) : 2 \rightarrow 0 \end{array}$$

Comparing with the action calculus for the π -calculus $\mathbf{AC}(\nu, \mathbf{out}, \mathbf{box}, \mathbf{rep})$, from [Mil96, §5.4], the arities of **new**, **out**, **in** and **rep** are obtained by adding one to the source of their corresponding arities; the name binding the new port on a control giving the location of that control. There is an obvious mapping taking processes in the fragment of dpi considered to actions of arity $0 \rightarrow 0$. Taking the reaction rules of the action calculus to be the translation of the first, second and fourth dpi rules, this is a bijection, up to structural congruence, that preserves one-step reaction.

3 Global/local subtyping

This section gives the global/local type system. It defines a judgement $\Gamma \vdash P : \text{process}$ which should be read as ‘under assumptions Γ the process P is well-formed’. As usual these contexts Γ contain assumptions on the types of names that may occur free in P . They must also contain assumptions on the locations of such names. *Pre-contexts* are therefore lists:

$$\begin{array}{ll} \Gamma ::= \bullet & \text{the empty context} \\ \Gamma, x : @_l T & \Gamma \text{ extended with name } x, \text{ located at } l, \text{ of type } T \end{array}$$

We now illustrate the three main phenomena that the type system must address. Firstly, a channel name must only be used (for input or output) if it has the appropriate capability, i.e. **L** or **G** for usages at its location; **G** for usages at other locations. For example, with respect to the context

$$\Gamma \stackrel{\text{def}}{=} k : @_{\text{top}} \text{loc}, l : @_{\text{top}} \text{loc}, w : @_l \uparrow_{-G} 1, z : @_l \uparrow_{-L} 1$$

we should have

$$\begin{array}{ll} \Gamma \vdash @_l \bar{w} : \text{process} & \Gamma \vdash @_l \bar{z} : \text{process} \\ \Gamma \vdash @_k \bar{w} : \text{process} & \Gamma \not\vdash @_k \bar{z} : \text{process} \end{array}$$

Secondly, local capabilities must not be sent outside their locations. Consider the context

$$\begin{array}{ll} \Gamma \stackrel{\text{def}}{=} k : @_{\text{top}} \text{loc}, & \text{top level location} \\ l : @_{\text{top}} \text{loc}, & \text{top level location} \\ z : @_l \uparrow_{LL} 1, & \text{local channel carrying 1, at } l \\ x : @_l \uparrow_{GG} \uparrow_{LL} 1 & \text{global channel carrying names of local channels carrying 1, at } l \end{array}$$

and the process $P \stackrel{\text{def}}{=} @_l \bar{x}z \mid @_k x(y).@_k \bar{y}$. At first sight one might expect $\Gamma \vdash P : \text{process}$, but the reduction

$$@_l \bar{x}z \mid @_k x(y).@_k \bar{y} \longrightarrow @_k \bar{z}$$

can send both **L** capabilities of z out of l — it is clear that $\Gamma \vdash @_k \bar{z} : \text{process}$ should not hold, and hence that $\Gamma \vdash P : \text{process}$ should not. It is prevented by restricting type formation, ruling out channel types, such as $\uparrow_{GG} \uparrow_{LL} 1$, that can be used to communicate local capabilities globally.

Thirdly, there must be a restriction on the mention of names outside their locations. This is a little delicate, as one cannot simply forbid all such mentions of the names of channels that are declared with some local capability. Suppose $x : \uparrow_{LL} \uparrow_{-o} 1$ is located at l and $z : \uparrow_{-o'} 1$, and consider when $@_l \bar{x}z$ should be well typed. If $o' = \mathbf{G}$ then z may be located anywhere, as its output capability is global. If $o = \mathbf{G}$ and $o' = \mathbf{L}$ then z is not a subtype of the expected value, so $@_l \bar{x}z$ should never be well typed. On the other hand, if $o = o' = \mathbf{L}$ then the output capability of z is local and may be used by a reader on x , so z must be located at l also. The essential point is whether the capabilities of z and the capabilities at which it can be used

by readers (as determined by the type of x) share a local capability (either for input or for output). This will be captured by a relation of *colocality* over types. Note that if $x: \downarrow_{LL} \downarrow_{-L} 1$ and $z: \downarrow_{LG}$ was not located at l the output should still be well typed, despite the fact that both types have a local capability.

Kinds, Contexts, Types and Values We first define four mutually recursive judgements, $\vdash T:K$, read as ‘type T is well-formed and has kind K ’, $\vdash \Gamma ok$, read as ‘context Γ is well-formed’, $\Gamma \vdash v:T$, read as ‘value v has type T ’, and $\Gamma \vdash x@l$, read as ‘name x is located at l ’.

The *kinds*, ranged over by K , are $\text{Type}_{\gamma\varepsilon}$ where γ and ε range over the 2-point lattices $G \leq -$ and $E \leq -$ respectively. They are ordered by the product order. The intuition is that types that have a kind $\text{Type}_{G\varepsilon}$ are *global*, with values of such types being freely communicable between locations. Types that have a kind $\text{Type}_{\gamma E}$ are *extensible*; new names at these types may be created by new-binders. We write $\sqcup$ for the least upper bounds in these lattices. The kinding rules for *types* are:

$$\begin{array}{c} \vdash B: \text{Type}_{G-} \\ \vdash T: \text{Type}_{\gamma\varepsilon} \\ \vdash T': \text{Type}_{\gamma'\varepsilon'} \\ \hline \vdash T \times T': \text{Type}_{(\gamma \sqcup \gamma') -} \end{array} \quad \begin{array}{c} \vdash \text{loc}: \text{Type}_{GE} \\ \vdash \top: \text{Type}_{GE} \end{array} \quad \begin{array}{c} \vdash T: \text{Type}_{G\varepsilon} \\ io \in \{GG, G-, -G\} \\ \hline \vdash \downarrow_{io} T: \text{Type}_{GE} \end{array} \quad \begin{array}{c} \vdash T: \text{Type}_{G\varepsilon} \\ io \in \{GL, LG\} \\ \hline \vdash \downarrow_{io} T: \text{Type}_{-E} \end{array} \quad \begin{array}{c} \vdash T: \text{Type}_{\gamma\varepsilon} \\ io \in \{LL, -L, L-\} \\ \hline \vdash \downarrow_{io} T: \text{Type}_{-E} \end{array} \quad \begin{array}{c} \vdash T: K \\ K \leq K' \\ \hline \vdash T: K' \end{array}$$

The rules for channel types prevent the formation of types that could be used to carry local capabilities between locations. For example, we have:

$$\begin{array}{c} \vdash \downarrow_{LL} \downarrow_{LL} 1: \text{Type}_{-E} \quad \not\vdash \downarrow_{GG} \downarrow_{LL} 1: \text{Type}_{--} \\ \vdash \downarrow_{LL} \downarrow_{GG} 1: \text{Type}_{-E} \quad \vdash \downarrow_{GG} \downarrow_{GG} 1: \text{Type}_{GE} \end{array}$$

and $\vdash \downarrow_{io} \downarrow_{i'o'} 1: \text{Type}_{-E}$ iff $io \in \{GG, G-, -G, GL, LG\} \Rightarrow i'o' \in \{GG, G-, -G\}$, i.e. if io is at all global then $i'o'$ must be not at all local. Products are global only if both their components are global. Base types and $\top$ are global, as is loc , so location names may be communicated freely. For illustration, the types (in boxes) and global types (in double boxes) of the form $\downarrow_{io} \downarrow_{i'o'} 1$ are shown in Figure 1. The only extensible types are channel types, loc , and $\top$. The formation rules for *contexts* are:

$$\frac{\vdash T: K \quad (\Gamma \vdash l: \text{loc}) \vee (l = \text{top}) \quad x \notin \text{dom}(\Gamma) \cup \{\text{top}\}}{\vdash \Gamma, x: @_l T ok} \quad \frac{}{\vdash \bullet ok}$$

Contexts thus contain location and type assumptions on free names. The rules ensure that locations are tree structured, with root top . The typing rules for values, and the rule for the location of names, are straightforward.

$$\frac{\vdash \Gamma, x: @_l T, \Delta ok}{\Gamma, x: @_l T, \Delta \vdash x: T} \quad \frac{\vdash \Gamma ok \quad b \in |B|}{\Gamma \vdash b: B} \quad \frac{\Gamma \vdash v: T \quad \Gamma \vdash v': T'}{\Gamma \vdash \langle v, v' \rangle: T \times T'} \quad \frac{\vdash \Gamma, x: @_l T, \Delta ok}{\Gamma, x: @_l T, \Delta \vdash x@l}$$

Subtyping The ordering on tags induces a subtype order on types — if $io \leq i'o'$ then a channel of type $\downarrow_{io} T$ may be used as if it were a channel of type $\downarrow_{i'o'} T$, which has weaker capabilities. As in [PS96], a tag io is covariant iff $o = -$, contravariant if $i = -$ and non-variant otherwise. The subtype order $\leq$ is the least binary relation over the pre-types such that

$$\begin{array}{c} \overline{B \leq B} \quad \frac{S_1 \leq T_1 \quad S_2 \leq T_2}{S_1 \times S_2 \leq T_1 \times T_2} \quad \frac{io \leq i'o' \quad i' \neq - \Rightarrow S \leq T \quad o' \neq - \Rightarrow T \leq S}{\downarrow_{io} S \leq \downarrow_{i'o'} T} \quad \text{loc} \leq \text{loc} \quad S \leq \top \end{array}$$




Figure 1: The subtype order over well-formed types $\downarrow_{io} \downarrow_{i''o''} 1$. They are ordered by $\downarrow_{io} \downarrow_{i''o''} 1 \leq \downarrow_{i'o'} \downarrow_{i'''o'''} 1$ iff $io \leq i'o'$ and $o' = - \Rightarrow i'o' \leq i'''o'''$, $i' = - \Rightarrow i'''o''' \leq i'o'$ and $o' \neq - \wedge i' \neq - \Rightarrow i'o' = i'''o'''$.

The replacement of the tag $--$ by a single top type ensures that subtyping is a partial order – otherwise names of types $\downarrow_{--} T$ would be communicable but not usable, so we would have $\downarrow_{--} S \leq \downarrow_{--} T$ for all S and T . The subtype order over well-formed types $\downarrow_{io} \downarrow_{i'o'} 1$ is

illustrated in Figure 1. Note that the well-formed types are not up, down or convex-closed under the subtype order on pre-types.

Colocality We say that a tag is *local* if it contains an $\mathbf{L}$ capability and that two tags are *colocal* if they share a common $\mathbf{L}$ capability, i.e. $\text{local}(io) \stackrel{\text{def}}{\iff} i = \mathbf{L} \vee o = \mathbf{L}$ and $\text{colocal}(io, i'o') \stackrel{\text{def}}{\iff} (i = \mathbf{L} \wedge i' = \mathbf{L}) \vee (o = \mathbf{L} \wedge o' = \mathbf{L})$. The key properties of these definitions are that $\text{colocal}(io, io) \iff \text{local}(io)$ and that, if $io \leq i'o' \leq i''o''$ and $\text{colocal}(io, i''o'')$, then $\text{colocal}(io, i'o')$ and $\text{colocal}(i'o', i''o'')$. Note that the local tags are neither up, down or convex closed in the tag ordering. Further, colocal is a symmetric relation but is not reflexive or transitive, or closed under relational composition with the tag ordering. It does satisfy $\text{colocal}(io, i'o') \Rightarrow (io \leq i'o' \vee i'o' \leq io)$. Colocality is lifted from tags to a relation on well-formed types that are in the subtype relation as follows.

$$\frac{\begin{array}{l} (i = i' = \mathbf{L}) \vee (o = o' = \mathbf{L}) \\ \vdash \downarrow_{io} S : \text{Type}_{--} \\ \vdash \downarrow_{i'o'} T : \text{Type}_{--} \\ \downarrow_{io} S \leq \downarrow_{i'o'} T \end{array}}{\text{colocal}(\downarrow_{io} S, \downarrow_{i'o'} T)} \quad \frac{\begin{array}{l} \text{colocal}(S_i, T_i) \\ \vdash S_{1-i} : \text{Type}_{--} \\ \vdash T_{1-i} : \text{Type}_{--} \\ S_{1-i} \leq T_{1-i} \end{array}}{\text{colocal}(S_0 \times S_1, T_0 \times T_1)} i \in \{0, 1\}$$

We define the colocal names of a value with respect to two types that are in the subtype relation:

$$\frac{\begin{array}{l} \Gamma \vdash x : S \\ \text{colocal}(S, T) \end{array}}{\Gamma \vdash x \in \text{colocaln}(x, S, T)} \quad \frac{\begin{array}{l} \Gamma \vdash x \in \text{colocaln}(v_i, S_i, T_i) \\ \Gamma \vdash v_{1-i} : S_{1-i} \\ \vdash T_{1-i} : \text{Type}_{--} \\ S_{1-i} \leq T_{1-i} \end{array}}{\Gamma \vdash x \in \text{colocaln}(\langle v_0, v_1 \rangle, S_0 \times S_1, T_0 \times T_1)} i \in \{0, 1\}$$

The key properties lift as follows. If a value has any colocal names with respect to two types then those types are colocal, the types of the colocal names of a value are themselves local and the set of colocal names of a value varies contravariantly with the upper type.

Lemma 1 *If $\Gamma \vdash x \in \text{colocaln}(v, V, T)$ then $\text{colocal}(V, T)$ and there exists U such that $\Gamma \vdash x : U$ and $\text{colocal}(U, U)$. If in addition $\vdash S : \text{Type}_{--}$ and $V \leq S \leq T$ then $\Gamma \vdash x \in \text{colocaln}(v, V, S)$.*

Processes Finally the typing rules for processes can be given.

$$\begin{array}{ll} \text{OUT} \frac{\begin{array}{l} \Gamma \vdash l : \text{loc} \\ \Gamma \vdash x : \downarrow_{io} T \\ \Gamma \vdash v : T' \\ T' \leq T \\ o \leq \mathbf{L} \\ o = \mathbf{L} \Rightarrow \Gamma \vdash x @ l \\ \forall a. \Gamma \vdash a \in \text{colocaln}(v, T', T) \Rightarrow \Gamma \vdash a @ l \end{array}}{\Gamma \vdash @_l \overline{x} v : \text{process}} & (\text{REP-IN}) \frac{\begin{array}{l} \Gamma \vdash l : \text{loc} \\ \Gamma \vdash x : \downarrow_{io} T \\ \Gamma, y : @_l T \vdash P : \text{process} \\ i \leq \mathbf{L} \\ i = \mathbf{L} \Rightarrow \Gamma \vdash x @ l \\ \text{fl}(P) \subseteq \{l\} \end{array}}{\Gamma \vdash @_l x(y).P : \text{process}} \\ & \Gamma \vdash @_l !x(y).P : \text{process} \\ \\ \text{LET} \frac{\begin{array}{l} \Gamma \vdash l : \text{loc} \\ \Gamma \vdash v : T' \\ \Gamma, y_1 : @_l T_1, y_2 : @_l T_2 \vdash P : \text{process} \\ T' \leq T_1 \times T_2 \\ \forall a. \Gamma \vdash a \in \text{colocaln}(v, T', T_1 \times T_2) \Rightarrow \Gamma \vdash a @ l \\ \text{fl}(P) \subseteq \{l\} \end{array}}{\Gamma \vdash @_l \text{let } \langle y_1 : T_1, y_2 : T_2 \rangle = v \text{ in } P : \text{process}} & \text{MIG} \frac{\begin{array}{l} \Gamma \vdash l : \text{loc} \\ \Gamma \vdash v : \text{loc} \\ \Gamma \vdash P : \text{process} \\ \text{fl}(P) \subseteq \{l\} \end{array}}{\Gamma \vdash @_l \text{migrate.to } v \text{ then } P : \text{process}} \\ \\ \text{NEW} \frac{\vdash T : \text{Type}_{-E} \quad \Gamma, x : @_l T \vdash P : \text{process}}{\Gamma \vdash (\text{new } x : @_l T) P : \text{process}} & \text{NIL} \frac{\vdash \Gamma \text{ ok}}{\Gamma \vdash \emptyset : \text{process}} \quad \text{PAR} \frac{\Gamma \vdash P : \text{process} \quad \Gamma \vdash Q : \text{process}}{\Gamma \vdash P \mid Q : \text{process}} \end{array}$$

Most of the premises of these rules are routine; we discuss the others briefly.

Out The first premise ensures that l is a location. The second through fifth premises are analogous to those of the **OUT** rule of [PS96]. Name x must be a channel, value v must be of a subtype of the type carried by the channel, and the channel must have an output capability (either **G** or **L**). The fourth and fifth premises could be replaced by $\uparrow_{io} T \leq \uparrow_{-l} T'$. The penultimate premise addresses the first phenomenon discussed at the beginning of this section, ensuring that if x has only a local output capability then it can only be used at its own location. The last premise addresses the third such phenomenon, ensuring that any transmitted channel names that have a local capability which can be used by receivers on x are located at l .

(Rep-)In This is very similar to **OUT** except for the premise $\text{fl}(P) \subseteq \{l\}$, which prevents teleportation after the input. Note that for typing P it is assumed that y is located at l .

New This allows new-binding of names at channel types, **loc**, and $\top$.

Let This is similar to a combination of **OUT** and **(Rep-)IN** (as, indeed, the reduction rule for **LET** is).

A few remarks: (1) The rules allow locations and channels, but not processes, to be located at top. This is consistent with the intuition that immediate sublocations of top model virtual machines. For other applications of the calculus different treatments of top are appropriate and should be straightforward. (2) Local channels can be sent outside their location (with reduced capabilities) and then back inside. Their local capabilities cannot then be used, however. (3) A name may be assumed to have a local type in a process P and still, if P is placed in a process context, engage in cross-location communication. (4) The **let** construct includes an explicit type for its pattern, which may be a supertype of the type of its value. Without this the set of typable processes would be unduly restricted. In the input construct the type of the pattern can be left implicit, as it is bounded by the type of the channel. (5) To add recursive types contexts must contain kind assumptions on type variables, type formation rules must be relativised to contexts, and enforce guardedness, subtyping must be defined coinductively, and type unfolding must be allowed in the value typing rules and definitions of subtyping and colocality. The details can be found in [Sew97a].

Soundness The main soundness result is that typing is preserved by reduction.

Theorem 1 (Subject reduction) *If $\Gamma \vdash P : \text{process}$ and $P \rightarrow Q$ then $\Gamma \vdash Q : \text{process}$.*

To prove this it is necessary to show that typing is preserved by legitimate *context permutations*, by *relocation* (changes of location assumptions for names of non-local types), by *narrowing* (taking type assumptions of lower types, while keeping location assumptions constant) and by *substitution*. For reasons of space we state only the substitution lemma for processes.

Lemma 2 (Substitution — Processes) *The rule below is admissible.*

$$\frac{\begin{array}{l} \Gamma, z : @_j V \vdash P : \text{process} \\ \Gamma \vdash u : U \\ U \leq V \\ \forall a . \Gamma \vdash a \in \text{colocaln}(u, U, V) \Rightarrow \Gamma \vdash a @ j \\ z \notin \text{fl}(P) \end{array}}{\Gamma \vdash \{u/z\}P : \text{process}}$$

The first three premises are standard. The fourth ensures that any names of the substituted value u are located at the same place as the substituted variable z was assumed to be at, if their actual and assumed types are colocal. The last premise ensures that no locators in P can be affected by the substitution.

In addition, it is easy to see from the typing rules that no well-typed process can immediately use a local capability outside its location. This can be made precise by immediate-soundness results such as the following.

Proposition 1 *If $\Gamma \vdash (\mathbf{new} \Delta)(@_l \bar{x}v \mid Q) : \text{process}$, $\Gamma, \Delta \vdash x : \downarrow_{iL} T$ and $\Gamma, \Delta \vdash x @ k$ then $k = l$.*

4 Capability Inference

One would like to be able to automatically infer the most local possible type for new-bound channels, to allow compile-time optimisation. Unfortunately, this is not possible in any straightforward sense based on the subtype order. Consider for example $k : @_{\text{toploc}}, z : @_k \downarrow_{LL} 1 \vdash (\mathbf{new} x : @_k T)(@_k \bar{x}z \mid @_k x(y).@_k \bar{y})$. This holds iff T is either $\downarrow_{LL} \downarrow_{-L} 1$ or $\downarrow_{LL} \downarrow_{LL} 1$; these types are not related by subtyping. We can, however, infer the most local possible top-level capabilities for T . Take the modified ‘subtype’ order $\sqsubseteq$ (with all channel type constructors covariant) to be the least relation over the pre-types such that

$$\frac{}{B \sqsubseteq B} \quad \frac{S_1 \sqsubseteq T_1 \quad S_2 \sqsubseteq T_2}{S_1 \times S_2 \sqsubseteq T_1 \times T_2} \quad \frac{T \sqsubseteq T' \quad io \leq i'o'}{\downarrow_{io} T \sqsubseteq \downarrow_{i'o'} T'} \quad \frac{}{\text{loc} \sqsubseteq \text{loc}} \quad \frac{}{\top \sqsubseteq \top} \quad \frac{}{\downarrow_{io} S \sqsubseteq \top}$$

and define $\simeq$ to be the least relation over the pre-types such that

$$\frac{}{B \simeq B} \quad \frac{S_1 \simeq T_1 \quad S_2 \simeq T_2}{S_1 \times S_2 \simeq T_1 \times T_2} \quad \frac{T \simeq T'}{\downarrow_{io} T \simeq \downarrow_{i'o'} T'} \quad \frac{}{\text{loc} \simeq \text{loc}} \quad \frac{}{\top \simeq \top} \quad \frac{}{\downarrow_{io} S \simeq \top} \quad \frac{}{\top \simeq \downarrow_{io} S}$$

relating any two types that have essentially the same shape, neglecting capabilities. Say a set of types $\tilde{T} = \{T_n \mid n \in N\}$ is *compatible* if it is non-empty and $\forall m, n \in N. T_m \simeq T_n$. One can show that any compatible $\tilde{T}$ has a least upper bound, written $\sqcup \tilde{T}$, with respect to $\sqsubseteq$. Lifting $\sqsubseteq, \simeq$, compatibility and $\sqcup$ pointwise to pre-contexts and processes, one can show that the typing judgements are preserved by taking least upper bounds with respect to $\sqsubseteq$.

Theorem 2 (Capability Inference)

1. *If $\tilde{T}$ is compatible and $\forall n \in N. \vdash T_n : K_n$ then $\vdash \sqcup \tilde{T} : \sqcup \tilde{K}$.*
2. *If $\tilde{S}$ is compatible, $\tilde{T}$ is compatible and $\forall n. S_n \leq T_n$ then $\sqcup \tilde{S} \leq \sqcup \tilde{T}$.*
3. *If $\tilde{\Gamma}$ is compatible, $\tilde{S}$ is compatible, $\tilde{T}$ is compatible, $\forall n. \Gamma_n \vdash v : S_n, \forall n. \vdash T_n : \text{Type}_{--}, \forall n. S_n \leq T_n$ and $\sqcup \tilde{\Gamma} \vdash x \in \text{colocaln}(v, \sqcup \tilde{S}, \sqcup \tilde{T})$ then $\exists n. \Gamma_n \vdash x \in \text{colocaln}(v, S_n, T_n)$.*
4. *If $\tilde{\Gamma}$ is compatible, $\tilde{P}$ is compatible and $\forall n \in N. \Gamma_n \vdash P_n : \text{process}$ then $\sqcup \tilde{\Gamma} \vdash \sqcup \tilde{P} : \text{process}$.*

For any pre-type S the set $\{T \mid S \sqsubseteq T\}$ is finite. Given some $\Gamma \vdash P : \text{process}$ (perhaps with types containing only GG capabilities, inferred by an algorithm along the lines of [Gay93, VH93, Tur96]) one can therefore compute the least upper bound of $\{P' \mid P \sqsubseteq P' \wedge \Gamma \vdash P' : \text{process}\}$. For the example above this gives $T = (\downarrow_{LL} \downarrow_{-L} 1) \sqcup (\downarrow_{LL} \downarrow_{LL} 1) = \downarrow_{LL} \downarrow_{-L} 1$. A more efficient algorithm will clearly be required in practice.

5 Conclusion

We conclude by briefly mentioning some related type systems and some possible future work. Capability-based type systems for process calculi have been given by De Nicola, Ferrari and Pugliese [DFP97], for a variant of Linda with localities, and by Riely and Hennessy [RH98], for a distributed π -calculus with site failure. Several authors have given type systems that enforce information flow properties, e.g. [HR98, SV98]. A type system that enforces secrecy and freshness for the Spi Calculus [AG97] has been proposed by Abadi in [Aba97]. In [Ste96] Steckler has given a static analysis technique for distributed Poly/ML with similar motivation to ours — to detect when channels are guaranteed to be local to a single processor. It incorporates also some reachability analysis, but does not separate input and output capabilities. Finally, Nielson and Nielson have studied static analysis techniques for CML, focussing on the number of usages of capabilities, in [NN95].

Special cases Three special cases of the type system may be of interest. In the Join Calculus the names introduced by a definition **def** D **in** P can only be used in P for output (to a first approximation D declares a single replicated reader on these names). For typing P , therefore, they are analogous to channels with capability $-G$. One could allow the output capability to be local, taking the suborder of tags $-G \leq -L$. In some circumstances it may not be necessary to allow the input and output capabilities of channels to vary separately, cutting down to the suborder of tags $GG \leq LL$. This greatly reduces the complexity (although also the expressiveness) of the type system as all channel type constructors become nonvariant. It can be used to prevent the extrusion of local references from agents. A milder simplification is to omit the tags GL and LG , i.e. to take the product of the tags $\pm, -, +$ of [PS96] with the two-point lattice $G \leq L$. For such tags, if $io \leq i'o'$ then $\text{colocal}(io, i'o') \iff \text{local}(io) \wedge \text{local}(i'o')$.

Linearity and Location Types In a distributed application one would expect many channels to be in some sense *linear*; in particular many servers will have a single replicated receiver (this observation motivates the introduction of join patterns in [FG96]). The integration of global/local typing with some form of linearity or receptiveness [Ama97, KPT96, San97] would allow more precise typing, and hence further optimizations, while retaining the expressiveness of general channel communication. One might also refine the system to allow location names to be local, with types loc_G and loc_L , enabling migration to locations to be restricted, and allow locations to be *immobile* or *mobile*, restricting the migration of locations. Linearity would again be useful — a common case is that of *one-hop* locations (c.f. Java Applets).

Behavioural equivalences In order to reason about dpi processes a labelled transition system and behavioural congruence are required, perhaps building on the bisimulation congruence results of Riely and Hennessy [RH97, RH98], together with an understanding of the appropriate extensional equivalence for a mobile agent programming language, building on [Sew97b].

Typing for secrecy properties The focus of this paper has been on locality information that can be used for implementation optimization. Very similar type systems should be applicable to the enforcement of secrecy properties for cryptographic keys or nonces. For this it would be desirable to take capabilities not just from $\{G, L, -\}$ but from the lattice of arbitrary sets of location names, lifted above a bottom element G . These (dependent) types would allow new names (modelling keys, for example, as in the Spi Calculus) to be created that are restricted to a dynamically calculated set of individuals. One would want a rather strong soundness result — the analogue of Theorem 1 would only show that secrecy is preserved by well-typed processes, whereas an attacker may perform some ill-typed computation.

Acknowledgements The author would like to thank Cédric Fournet, Robin Milner, Benjamin Pierce, Paweł Wojciechowski, and the Thursday group, for interesting discussions about this work, and to acknowledge support from EPSRC grant GR/K 38403 and Esprit Working group 21836 (CONFER-2).

References

- [Aba97] Martín Abadi. Secrecy by typing in security protocols. In *TACS '97 (open lecture), LNCS 1281*, pages 611–638, September 1997.
- [AG97] Martín Abadi and Andrew D. Gordon. A calculus for cryptographic protocols: The spi calculus. In *Proceedings of the Fourth ACM Conference on Computer and Communications Security, Zürich*, pages 36–47. ACM Press, April 1997.
- [Ama97] R. M. Amadio. An asynchronous model of locality, failure, and process mobility. In *Proc. COORDINATION 97, LNCS 1282*, 1997.
- [AP94] R. M. Amadio and S. Prasad. Localities and failures. In P. S. Thiagarajan, editor, *Proceedings of 14th FST and TCS Conference, FST-TCS'94. LNCS 880*, pages 205–216. Springer-Verlag, 1994.

- [Bou92] Gérard Boudol. Asynchrony and the π -calculus (note). Rapport de Recherche 1702, INRIA Sofia-Antipolis, May 1992.
- [CG98] Luca Cardelli and Andrew D. Gordon. Mobile ambients. In *Proc. of Foundations of Software Science and Computation Structures (FoSSaCS), ETAPS'98*, March 1998.
- [DFP97] Rocco De Nicola, GianLuigi Ferrari, and Rosario Pugliese. Coordinating mobile agents via blackboards and access rights. In *Proc. COORDINATION '97, LNCS 1282*, 1997.
- [FG96] Cédric Fournet and Georges Gonthier. The reflexive CHAM and the join-calculus. In *Proceedings of the 23rd POPL*, pages 372–385. ACM press, January 1996.
- [FGL⁺96] Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A calculus of mobile agents. In *Proceedings of CONCUR '96. LNCS 1119*, pages 406–421. Springer-Verlag, August 1996.
- [Gay93] Simon J. Gay. A sort inference algorithm for the polyadic π -calculus. In *Proceedings of the 20th POPL*. ACM Press, 1993.
- [HR98] Nevin Heintze and Jon G. Riecke. The SLam calculus: Programming with secrecy and integrity. In *Proceedings of the 25th POPL*, January 1998.
- [HT92] Kohei Honda and Mario Tokoro. On asynchronous communication semantics. In M. Tokoro, O. Nierstrasz, and P. Wegner, editors, *Object-Based Concurrent Computing. LNCS 612*, pages 21–51, 1992.
- [KPT96] Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. Linearity and the pi-calculus. In *Proceedings of the 23rd POPL*, pages 358–371. ACM press, January 1996.
- [Mil96] Robin Milner. Calculi for interaction. *Acta Informatica*, 33:707–737, 1996.
- [MPW92] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, Parts I + II. *Information and Computation*, 100(1):1–77, 1992.
- [NN95] Hanne Riis Nielson and Flemming Nielson. Static and dynamic processor allocation for higher-order concurrent languages. In *Proceedings of TAPSOFT 95 (FASE). LNCS 915*, 1995.
- [PS96] Benjamin Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. *Mathematical Structures in Computer Science*, 6(5):409–454, 1996.
- [RH97] James Riely and Matthew Hennessy. Distributed processes and location failures. In *Proceedings of ICALP '97. LNCS 1256*, pages 471–481. Springer-Verlag, July 1997.
- [RH98] James Riely and Matthew Hennessy. A typed language for distributed mobile processes. In *Proceedings of the 25th POPL*, January 1998.
- [San97] Davide Sangiorgi. The name discipline of uniform receptiveness. In *Proceedings of ICALP '97. LNCS 1256*, pages 303–313, 1997.
- [Sew97a] Peter Sewell. Global/local subtyping for a distributed π -calculus. Technical Report 435, University of Cambridge, August 1997. Available from <http://www.cl.cam.ac.uk/users/pes20/>.
- [Sew97b] Peter Sewell. On implementations and semantics of a concurrent programming language. In *Proceedings of CONCUR '97. LNCS 1243*, pages 391–405. Springer-Verlag, 1997.
- [Ste96] Paul Steckler. Detecting local channels in distributed Poly/ML. Technical Report ECS-LFCS-96-340, University of Edinburgh, January 1996.
- [SV98] Geoffrey Smith and Dennis Volpano. Secure information flow in a multi-threaded imperative language. In *Proceedings of the 25th POPL*, January 1998.
- [TLK96] Bent Thomsen, Lone Leth, and Tsung-Min Kuo. A Facile tutorial. In *Proceedings of CONCUR '96. LNCS 1119*, pages 278–298. Springer-Verlag, August 1996.
- [Tur96] David N. Turner. *The Polymorphic Pi-calculus: Theory and Implementation*. PhD thesis, University of Edinburgh, 1996.
- [VH93] Vasco Thudichum Vasconcelos and Kohei Honda. Principal typing schemes in a polyadic π -calculus. In *Proceedings of CONCUR '93. LNCS 715*, pages 524–538, 1993.