MIKADO
Mobile Calculi based on Domains

MIKADO Global Computing Project
IST-2001-32222

Type Systems for Mikado inspired mobility and security programming concepts

MIKADO Deliverable D2.1.2

Editor : M. HENNESSY (SUSSEX)

Authors : M HENNESSY, C LHOUSSAINE, J RATHKE, V SASSONE (SUSSEX)

Contributors : M HENNESSY, C LHOUSSAINE, J RATHKE, V SASSONE (SUSSEX)

Classification : Public
Deliverable no. : D2.1.2
Reference : RR/WP2/3
Date : February 2004

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

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

Abstract

This document reports on specific developments on type systems for languages inspired by the Mikado programming model. The main scientific contents are contained in two technical annexes but references are also made to related advances, often associated with other Global Computing projects.

Research pertinent to this delivery continues at the Mikado sites, Firenze, INRIA and Sussex, often as joint collaborations. However this deliverable consists of two specific technical papers, relating novel developments on type systems for distributed systems, which we believe are applicable to languages using the general Mikado programming model.

The essential point is the perception that *dependent types*, which have long been applied in Computer Science, can be used to give detailed, yet flexible, control over agent behaviour. The detail allows hosts to implement security policies protecting their resources, for example by explicitly naming the agents which are allowed to perform specific actions. On the other hand flexibility means that servers can still offer personalised services to particular client, since these policies need not necessarily constrain the behaviour of well-meaning clients.

A good example of an implicit use of dependent types is the use of *groups* in various typing systems for Ambients, [CGG00]. Intuitively the space of Ambients is divided in possibly overlapping sets of *groups* and types defining ambient capabilities, such as where an ambient may reside, or where they may move to, are then expressed at the level of these *groups*. Rather than using these static classifications we could instead use the actual names of ambients themselves when defining capabilities. This would lead to a type system capable of expressing very detailed security policies, but would be very inflexible. For example servers could only be designed to serve specifically named clients. But by parameterising these types on the names used, effectively defining dependent ambient types, we re-introduce flexibility. In particular we allow typed servers to be defined, which on the one hand guarantee high-levels of security, and on the other allow flexible services to be offered to particular clients.

In this deliverable this idea is pursued in two directions. In the first it is applied to a higher-order version of Dpi, a language within the framework of Mikado-inspired mobility principles, although the distribution space is flat. In the second we examine the language of Ambients, where of course the distribution space is hierarchical. In both cases the details of type-checking are complex. However this is inevitable, as the kinds of security policies we envisage are necessarily subtle, and even being able to express them in terms of a type system is in itself an achievement. Moreover we do not envisage applying type reconstruction in this scenario, a topic which is well-known to be difficult for dependent types. Instead we envisage these languages to be explicitly typed, as the types associated with the various entities are to be designed to implement desired security policies.

Finally although we have pursued these ideas for two specific Mikado-related languages, we believe that the concepts are more generally applicable to arbitrary languages within the Mikado framework.

Dpi: In last years deliverable, [HMR03], a type system was developed for a distributed version of the picalculus, Dpi, which enables sites to selectively distribute and manage access rights to resources. In [HRY03], the first half of which should be considered part of this deliverable, this approach is extended to a higher-order version of Dpi, which gives much more control to hosts over general access and use of their own internal computation space. The general idea is to replace the unsecure migration command

$$\mathsf{goto}\, l.C$$

with the more restrictive

$$\mathsf{goto}_p\, l.F$$

The former allows any client to run any code C at a host site, in this case l. In the latter the client must name an entry *port*, here p, at the host's site; this can be viewed as a higher-order channel, which can be

programmed by the host site in a manner similar to the standard channels of the picalculus. So for example the host may not make available input on the port p, thereby denying access to the client. If input is available at p then the *parametric* code F is received by the host, and can be applied to whatever arguments it deems fit.

However the real expressive power in the language is in the types which can be associated with these higher-order ports. First a fine-grained system of types for agents is designed in which definite limits can be put on the local, and non-local resources which can be used by the agent. For example the record type

$$[\mathsf{info} : \mathsf{r}\langle \mathit{string} \rangle @\mathsf{here},\ \mathsf{put} : \mathsf{w}\langle \mathit{string} \rangle @\mathsf{cl}]$$

restricts an agent to *reading* some information from the server site, referred to as here, and writing to a resource called put at a client site cl. Effectively a server with a port at this type is servicing only one specific client called cl.

However this focus on a specific client site called cl can be eliminated by the use of *dependent* types. For example consider the type expression:

$$\mathsf{TDep}(x : \mathsf{loc})\,[\mathsf{info} : \mathsf{r}\langle \mathit{string} \rangle @\mathsf{here},\ \mathsf{put} : \mathsf{w}\langle \mathit{string} \rangle @\mathsf{x}]$$

Here the site at which the information is to be written, using put, is now parameterised with reject to an arbitrary location x. A port with this type would expect to receive both the name of a location say k , which is bound to x, and code whose type is guaranteed to be

$$[\mathsf{info} : \mathsf{r}\langle \mathit{string} \rangle @\mathsf{here},\ \mathsf{put} : \mathsf{w}\langle \mathit{string} \rangle @\mathsf{k}]$$

The resulting server would therefore be able to offer a generalised service to any client willing to supply code at the required type.

In [HRY03] these dependent types are investigated in detail, and it is shown how they can be used on the one hand to define general but secure servers, and on the other how clients may also use them to protect themselves in turn from abusive servers.

Ambients: In [LS03], which should be considered as part of this deliverable, similar ideas on dependent types are applied to ambients. It provides an alternative to the commonly used *group-based* type systems dealing with mobility control in the framework of the Ambient Calculus. The objectives are both to simplify the specification of security policies by means of types (i.e. to make programmer's life easier) and to improve the expressivity and the flexibility of the typed language. This inevitably results in an increasing complexity of the typing rules, which require the new pivotal notion of *abstract names*. Those names allow the *internalisation* of (a dynamic version of) group names, which move from being part of the language to being a inner mechanism of the type system.

Assuming $\mathcal{P}$ and $\mathcal{C}$ are sets of ambient names, an ambient n with type $\mathsf{mob}[\mathcal{P}, \mathcal{C}]$ is allowed to be a sub-ambient of all ambients whose name occurs in $\mathcal{P}$ (the possible parents of n), and is allowed to wrap any ambient whose name occurs in $\mathcal{C}$ (the set of possible children of n). The main example illustrating the power of this new system is the one of a specialised *taxi server* which provides a service for a specific client c:

$$T(c) = !(\nu \mathit{taxi} : A)(\mathit{taxi}[] \mid c[\mathsf{in}\mathit{taxi}])$$

By defining the type A of *taxi* as $\mathsf{mob}[\emptyset, \{c\}]$, we specify that only ambients named c are allowed to use (i.e. here to access) the freshly created taxi. To guarantee the same restricted access to *taxi* using groups we would need a more elaborate, and thus more complex, specification. Still, the main improvement is the dynamicity allowed by our dependent types. It is indeed easy to generalise the previous example to get a *real* server, i.e. which provides a service usable only by any client whose name only becomes known dynamically:

$$\mathit{TaxiServ} = !(x)(\nu \mathit{taxi} : A')(\mathit{taxi}[] \mid x[\mathsf{in}\mathit{taxi}])$$

Here A' is $\mathsf{mob}[\emptyset, \{x\}]$ which refers to the received name bound to x.

References

[CGG00] Luca Cardelli, Giorgio Ghelli, and Andrew D. Gordon. Ambient groups and mobility types. In *IFIP TCS*, pages 333–347, 2000.

[HMR03] Matthew Hennessy, Massimo Merro, and Julian Rathke. Towards a behavioural theory of access and mobility control in distributed systems. *Theoretical Computer Science*, 2003. To Appear.

[HRY03] Matthew Hennessy, Julian Rathke, and Nobuko Yoshida. Safedpi: a language for controlling mobile code. Technical Report 2003:02, University of Sussex, 2003. Extended abstract to be presented at Fossacs04.

[LS03] C. Lhoussaine and V. Sassone. A dependently typed ambient calculus. In *Proceedings of European Symposium on Programming, ESOP'03*, LNCS. Springer, 2003. To appear.

SAFEDPI: a language for controlling mobile code

Matthew Hennessy Julian Rathke Nobuko Yoshida

JANUARY 2004

Abstract

SAFEDPI is a distributed version of the PICALCULUS, in which processes are located at dynamically created *sites*. Parametrised code may be sent between sites using so-called *ports*, which are essentially higher-order versions of PICALCULUS communication channels. A host location may protect itself by only accepting code which conforms to a given type associated to the incoming port.

We define a sophisticated static type system for these ports, which restrict the capabilities and access rights of any processes launched by incoming code. Dependent and existential types are used to add flexibility, allowing the behaviour of these launched processes, encoded as *process types*, to depend on the host's instantiation of the incoming code.

We also show that a natural contextually defined behavioural equivalence can be characterised coinductively, using bisimulations based on *typed actions*. The characterisation is based on the idea of knowledge acquisition by a testing environment and makes explicit some of the subtleties of determining equivalence in this language of highly constrained distributed code.

1 Introduction

In this paper we elaborate a theory of distributed systems which incorporates resource policies. Our main results are:

- a language for distributed systems in which access to hosts by mobile code is controlled using capability-based types
- a fine-grained type system using novel forms of dependent and existential types which gives hosts considerable flexibility in determining the allowed behaviour of incoming code
- a coinductive characterisation of a natural contextual equivalence, based on the notion of typed actions.

This is developed in terms of an extension of the language DPI, [10, 8, 20, 14], a version of the PICALCULUS, [21], in which processes may migrate

between locations, which in turn can be dynamically created. In DPI a typical system takes the form

$$l[\![P]\!] \mid (\mathsf{new}\, e : \mathsf{E})(k[\![Q]\!] \mid l[\![R]\!])$$

where there are two threads P and R running at l and one, Q, running at k. The threads Q and R share the private name e at type E. The threads P , Q, R are similar to processes in the PICALCULUS in that they can receive and send values on local channels; the types of these channels indicate the kind of values which may be transmitted. Locations may be dynamically created. For example in

$$l[\![(\mathsf{newloc}\, k : \mathsf{K})\, \mathsf{with}\, Q \ \mathsf{in}\ \mathsf{xpt}_1!\langle k\rangle \mid \mathsf{xpt}_2!\langle k\rangle]\!]$$

a new location k is created at type K, the code Q is installed at k and the name of the new location is exported via the channels xpt_i. Location types are similar to record types, their form being

$$\mathsf{loc}[c_1 : \mathsf{C}_1, \ldots c_n : \mathsf{C}_n]$$

This indicates that the channels, or resources, c_i at types C_i are available at the location. So for example K above could be

$$\mathsf{loc}[\mathsf{ping} : \mathsf{rw}\langle \mathsf{P}\rangle, \mathsf{fing} : \mathsf{rw}\langle \mathsf{F}\rangle]$$

indicating that the services ping and fing(er) are supported at k; r indicates the permission to *read from* a channel, while w indicates the permission to *write to* the channel. However the types at which k becomes known depends on the types of the exporting channels. Suppose for example these had the types

$$\mathsf{xpt}_1 : \mathsf{w}\langle \mathsf{loc}[\mathsf{ping} : \mathsf{w}\langle \mathsf{P}\rangle]\rangle$$
$$\mathsf{xpt}_2 : \mathsf{w}\langle \mathsf{loc}[\mathsf{fing} : \mathsf{w}\langle \mathsf{F}\rangle]\rangle$$

Then processes receiving the name k from the source xpt_1 would only be able to write to the ping service at k, i.e. send messages to that service, while the source xpt_2 only allows similar restricted access to the finger service. It is in this way, by selectively distributing names at particular subtypes, that resource access policies are implemented in DPI.

In this paper we make two extensions to DPI. The first allows more control to locations over code which wishes to access their computation space. In DPI the migration rule is given by

$$k[\![\mathsf{goto}\, l.P]\!] \longrightarrow l[\![P]\!];$$

any thread is allowed to migrate to the site l. In SAFEDPI, the language of this paper, migration is represented by

$$k[\![\mathsf{goto}_p\, l.F]\!] \longrightarrow l[\![p!\langle F\rangle]\!]$$

A thread must designate a *port* p at l in order to migrate. It then reduces to the system $l[\![p!\langle F\rangle]\!]$, which apriori represents a thread running at location l. However this thread will have no effect until the site l makes available a corresponding thread of the form $l[\![p?(x)\,Q]\!]$; using standard communication this will now allow the effective entry of F. In this manner, by programming the presence or absence of ports, the site l can control the immigration of code.

Effectively we have replaced unconstrained spawning of processes at arbitrary sites by *higher-order communication*. Moreover these ports, higher-order channels, have types associated with them. The types on ports are the second major extension to the language. In general we allow *scripts*, parameterised code, to be sent via ports. These take the form

$$\lambda(\tilde{x} : \tilde{\mathsf{T}})P$$

where each x_i can be matched by arbitrary *transmittable values*; it is the types T_i which determine the nature of the abstraction. But when such a script is transmitted it may be instantiated at the receiving site by values of the appropriate type. This gives added security to sites by controlling the type at which scripts will be accepted. This of course depends on the granularity of the type structure for scripts.

The most straightforward form of type for scripts is

$$(\tilde{x} : \tilde{\mathsf{T}}) \to \mathbf{proc}$$

stating that, whenever a script of this type is instantiated with appropriate parameters, the result is guaranteed to be a well-typed process. But a priori there is no constraint on the resources it can use. To limit the access of incoming code to resources we introduce fine-grained *process types*, [26]. These dictate the capabilities, on both local and third-party channels, which the code is allowed to access, and take the form of a record:

$$\mathsf{pr}[c_1 : \mathsf{C}_1@k_1, \ldots, c_n : \mathsf{C}_n@k_n]$$

A process of this type can use *at most* the set of channels c_i, located respectively at the locations k_i, with the capabilities C_i; in these process types the use of a local channel c is indicated by an entry of the form $c : \mathsf{C}@\mathsf{here}$.

When these process types are incorporated into script types a host location can have much more effective control over the behaviour of incoming code, particularly when we use a form of *dependent* function type. For example suppose a port only accepts scripts at the type

$$\mathsf{Fdep}(x : \mathsf{r}\langle\mathsf{T}\rangle \to \mathsf{pr}[x : \mathsf{r}\langle\mathsf{T}\rangle@\mathsf{here},\ \mathsf{reply} : \mathsf{w}\langle\mathsf{T}\rangle@k])$$

Then an incoming script can only be instantiated by a local channel, with read capability at type T. Moreover the resulting running code is now only

allowed to read from this local channel and write to the third-party channel called reply located at the specific location k. With a port at the type

$$\mathsf{Fdep}(y : \mathsf{w}\langle\mathsf{T}\rangle@k \to \mathsf{pr}[\mathsf{info} : \mathsf{r}\langle\mathsf{T}\rangle@\mathsf{here},\ y : \mathsf{w}\langle\mathsf{T}\rangle@k])$$

the host can instantiate the incoming script with some channel located at the site k, on which it has write permission, and the running code is restricted to writing there, and reading from a local channel called info.

Note that in both these examples the location k is built into the script types. Thus a server with an access port at this type would only allow entry to scripts which guarantee to write only at k. However *dependent types* can be used to allow this target site to be parameterised. Consider the simple example

$$\mathsf{Tdep}(z : \mathsf{L})\,\mathsf{Fdep}(y : \mathsf{w}\langle\mathsf{T}\rangle@z \to \mathsf{pr}[\mathsf{info} : \mathsf{r}\langle\mathsf{T}\rangle@\mathsf{here},\ y : \mathsf{w}\langle\mathsf{T}\rangle@z])$$

where the script type is now parameterised by locations of some type L. This allows the server to accepts scripts which can write the information at sites determined by the client.

Although these dependent types add considerable flexibility to the interaction between clients and servers, they have potential drawbacks; as we will see the client has to send with the script the actual objects on which their type is parameterised. In principle this opens up the possibility of (rogue) servers abusing this extra information. However *existential types* provide extra protection to clients, because, as we will see, this extra information is not required as part of the communication.

The language SAFEDPI is formally defined in Section 2, together with a reduction semantics. In Section 3 we define the set of types and the type inference system; the formal development relies heavily on the type systems already given in [8, 19]. In Section 4 we develop a series of example systems. These are designed on the one hand, to explain the intricacies of the the type inference rules, and on the other to demonstrate the power and flexibility of the types. This is followed by a section devoted to establishing the expected properties of the type system, in particular Subject Reduction.

We now turn to the second topic of the paper, typed behavioural equivalences. In untyped languages, these are normally defined coinductively, as the largest equivalences over processes which preserve, in some sense, actions of the form

$$M \xrightarrow{\mu} M' \tag{1}$$

Typically these actions describe the possible forms of interactions between a process and its environment. In a typed setting many of these actions will not be possible, because the environment will not have the power to participate in them. As a simple example consider the system

$$l[\![(\mathsf{newc}\, c : \mathsf{C})\ (\mathsf{xpt!}\langle c\rangle\ \mid c?(x)\,Q)]\!]$$

in an environment in which the export channel xpt can only send channels with the read capability. The environment will receive c along xpt but will not be able to transmit on c. Consequently the potential input actions on c by the process above will not be possible.

Following [9, 8] we replace the untyped actions in (1) with typed actions of the form

$$\mathcal{I} \rhd M \xrightarrow{\mu} \mathcal{I}' \rhd M'$$

where M is the system being observed while $\mathcal{I}$ is a constraint on the observing environment representing its knowledge of the system M. Actions change both the processes and the environment in which they are being observed. This will lead, in the standard manner, to a coinductively defined, bisimulation-based, relation between systems, which we denote by

$$\mathcal{I} \models M \approx_{bis} N$$

In our second main result of the paper, we prove that this coinductive relation coincides with a naturally defined contextual equivalence. One of the features of our approach is the explicit representation of the information which the environment can obtain from systems through testing with contexts. In such a highly constrained setting as this, this becomes a genuine aid in understanding the equivalence. This is the topic of Section 6.

The paper ends, in Section 7, with some conclusions and a brief survey of related work.

2 The language SAFEDPI

Syntax: The syntax, given in Figure 1, is a slight extension of that of DPI from [8]. It is explicitly typed, but for expository purposes we defer the description of types until Section 3. The syntax also presupposes a general set of channel names NAMES, ranged over by n, m, and a set of variables VARS, ranged over by x, y. *Identifiers*, ranged over by u, w, may come from either of these sets. NAMES is partitioned into two sets, LOCS ranged over by $k, l, \ldots$ for locations, and CHANS ranged over by $a, b, c, \ldots$ for channels. There is also a distinguished subset of channels called *ports*, and ranged over by $p, q, \ldots$, which are used to handle higher-order values. Similarly we will sometimes use ξ, ξ' for variables which will be instantiated by higher-order values.

The syntax for systems, ranged over by M, N, O, is the same as in DPI, allowing the parallel composition of located processes $l[\![P]\!]$, which may share defined names, using the construct $(\mathsf{new}\, e : \mathsf{E})\, -$.

The syntax for processes, ranged over by P, Q is an extension of the PICALCULUS, [21], with primitives for migration between locations. Parallelism is allowed, we have the terminated process stop, and we also allow

M, N ::=	*Systems*
$l[\![P]\!]$	Located Process
$M \mid N$	Composition
$(\mathsf{new}\, e : \mathsf{E})\, M$	Name Creation
$\mathbf{0}$	Termination

P, Q ::=	*Processes*
$u!\langle V \rangle$	Output
$u?(X : \mathsf{T})\, P$	Input
$\mathsf{goto}_u\, v.P$	Migration
$\mathsf{if}\ u_1 = u_2\ \mathsf{then}\ P\ \mathsf{else}\ Q$	Matching
$(\mathsf{newc}\, c : \mathsf{C})\ P$	Channel creation
$(\mathsf{newreg}\, n : \mathsf{N})\ P$	Global name creation
$(\mathsf{newloc}\, k : \mathsf{K})\, \mathsf{with}\, Q\ \mathsf{in}\ P$	Location creation
$P \mid Q$	Composition
$F\,(\tilde{v})$	Application
$*P$	Iteration
stop	Termination

U, V, W ::=	*Values*
$(\tilde{v})$	Tuples
v ::=	*Value components*
$\lambda\,(\tilde{x} : \tilde{\mathsf{T}}).\ P$	Scripts
u	Identifiers

Figure 1: Syntax of SAFEDPI

matching and mismatching, with the construct $\mathsf{if}\ u_1 = u_2\ \mathsf{then}\ P\ \mathsf{else}\ Q$, and a form of iteration $*P$.

In the input construct $u?(X : \mathsf{T})\, P$ we take X to be a pattern which is used to deconstruct incoming values; this is a value which only contains distinct occurrences of variables. In our somewhat restricted format for values this means that X has the form $(\tilde{x})$, with each x_i being distinct. The output construct is asynchronous, $u!\langle V \rangle$. Here V is a tuple consisting of either identifiers or higher-order values. The latter can take the form of scripts, $\lambda\,(\tilde{x} : \tilde{\mathsf{T}}).\ P$, where P is an arbitrary process term; we will often use F to indicate an arbitrary script, whereas v will be reserved for the individual components in a tuple V; thus it will represent either an identifier or a script. Of particular interest to us will be tuples of the form $(\tilde{v}, F)$ which will be interpreted as *dependent values*; intuitively the script F depends on the

values $\tilde{v}$.

At the risk of being verbose, the syntax has explicit notations for the various forms of names which can be declared. In $(\mathsf{newc}\, c : \mathsf{C})\ P$ a new local channel named c is declared, while $(\mathsf{newreg}\, n : \mathsf{N})\ P$ represents the generation of a new globally registered name n for channels; see [8] for motivation. When a new location is declared, in $(\mathsf{newloc}\, k : \mathsf{K})\, \mathsf{with}\, Q\ \mathsf{in}\ P$, its declaration type K can only involve channel names which have been registered. This construct generates the new location k, sets the code Q running there, and in parallel continues with the execution of P. This specific construct for new locations is required since code may only be executed at a location once entry has been gained via a port; so here Q represents the code with which the location is initialised.

The main novelty in SAFEDPI, over DPI, is the construct

$$\mathsf{goto}_p\, k.F$$

Intuitively this means: migrate to location k via the port p with the code F. Our type system will ensure that F is in fact a script with a type appropriate to the port p; moreover entry will only be gained if at the location k the port p is currently active.

The various binding structures, for names and variables, give rise to the standard notions of free and bound occurrences of identifiers, α-conversion, and (capture-avoiding) substitution of values for identifiers in terms, $P\{\!|^{v}/_{u}|\!\}$; this is extended to patterns, $P\{\!|^{V}/_{X}|\!\}$ in the standard manner. We omit the details but three points are worth emphasising. The first is that many such substitutions may give rise to badly formed process terms but our typing system will ensure that this will never occur in well-typed terms. The second is that identifiers may occur in certain types and therefore we require a notion of substitution into types; this will be explained in Section 3. Finally terms will be identified up to α-equivalence, and bound identifiers will always be chosen to be distinct, and different from any free identifiers.

In the sequel we use *system* to refer to a *closed* system term, that is a system term which contain no free occurrences of variables; similarly a *process* means a closed process term.

Reduction Semantics: This is given in terms of a binary relation between systems

$$M \longrightarrow N$$

and is a mild generalisation of that given in [8, 10] for DPI.

Definition 2.1 (Contextual relations) A relation $\mathcal{R}$ over systems is said to be *contextual* if it preserves all the system constructors of the language; that is $M\,\mathcal{R}\,N$ implies

- $M \mid O\,\mathcal{R}\,N \mid O$ and $O \mid M\,\mathcal{R}\,O \mid N$

(R-COMM)
$$k[\![c!\langle V\rangle]\!] \mid k[\![c?(X:\mathsf{T})\,P]\!] \longrightarrow k[\![P\{\!|^{V}/_{X}|\!\}]\!]$$

(R-SPLIT)
$$k[\![P \mid Q]\!] \longrightarrow k[\![P]\!] \mid k[\![Q]\!]$$

(R-N.CREATE)
$$k[\![(\mathsf{newreg}\, n : \mathsf{N})\ P]\!] \longrightarrow (\mathsf{new}\, n : \mathsf{N})\, k[\![P]\!]$$

(R-MOVE)
$$k[\![\mathsf{goto}_p\, l.F]\!] \longrightarrow l[\![p!\langle F\rangle]\!]$$

(R-L.CREATE)
$$k[\![(\mathsf{newloc}\, l : \mathsf{L})\, \mathsf{with}\, C\ \mathsf{in}\ P]\!] \longrightarrow (\mathsf{new}\, l : \mathsf{L})(k[\![P]\!] \mid l[\![C]\!])$$

(R-C.CREATE)
$$k[\![(\mathsf{newc}\, c : \mathsf{C})\ P]\!] \longrightarrow (\mathsf{new}\, c : \mathsf{C}@k)\, k[\![P]\!]$$

(R-UNWIND)
$$\frac{k[\![P]\!] \mid M \longrightarrow M'}{k[\![*P]\!] \mid M \longrightarrow k[\![*P]\!] \mid M'}$$

(R-EQ)
$$k[\![\mathsf{if}\ u = u\ \mathsf{then}\ P\ \mathsf{else}\ Q]\!] \longrightarrow k[\![P]\!]$$

(R-BETA)
$$k[\![(\lambda\,(\widetilde{x} : \widetilde{\mathsf{T}}).\ P)(\widetilde{v})]\!] \longrightarrow k[\![P\{\!|^{\widetilde{v}}/_{\widetilde{x}}|\!\}]\!]$$

(R-NEQ)
$$k[\![\mathsf{if}\ u = v\ \mathsf{then}\ P\ \mathsf{else}\ Q]\!] \longrightarrow k[\![Q]\!] \quad u \neq v$$

(R-STR)
$$\frac{M \equiv N,\ M \longrightarrow M',\ M' \equiv N'}{N \longrightarrow N'}$$

Figure 2: Reduction semantics for SAFEDPI

- $(\mathsf{new}\, e : \mathsf{E})\, M\,\mathcal{R}\,(\mathsf{new}\, e : \mathsf{E})\, N$. ■

The reduction relation is defined to be the least contextual relation which satisfies the axioms and rules in Figure 2. The rule (R-STR) merely says that we are working up to a structural equivalence, $\equiv$, which abstracts from inessential details in the terms representing systems. Formally structural equivalence is defined to be the least contextual relation between (*closed*) systems which satisfies the axioms which are given in Figure 3; these are the natural adaptations of the usual axioms for structural equivalence in the PICALCULUS.

The main reduction involves *local communication*, governed by the rule (R-COMM), taken directly from DPI. However here the value V may be a script; in other words this rule encompasses higher-order communication. Higher-order output commands are generated by (R-MOVE), which has already been explained in the introduction.

Migration to a site l must designate a *port* p at which the migrating code is to be received. The rule

$$k[\![\mathsf{goto}_p\, l.F]\!] \longrightarrow l[\![p!\langle F\rangle]\!]$$

then translates the migration command into the system $l[\![p!\langle F\rangle]\!]$, which apriori represents a thread running at the target location l. However this

(S-EXTR)	$(\mathsf{new}\, n : \mathsf{E})(M \mid N)$	$=$	$M \mid (\mathsf{new}\, n : \mathsf{E})\, N$ if $n \notin \mathsf{fn}(M)$
(S-COM)	$M \mid N$	$=$	$N \mid M$
(S-ASSOC)	$(M \mid N) \mid O$	$=$	$M \mid (N \mid O)$
(S-ZERO)	$M \mid \mathbf{0}$	$=$	M
(S-STOP)	$k[\![\mathsf{stop}]\!]$	$=$	$\mathbf{0}$
(S-FLIP)	$(\mathsf{new}\, n : \mathsf{E})\,(\mathsf{new}\, n' : \mathsf{E}')\, M$	$=$	$(\mathsf{new}\, n' : \mathsf{E}')\,(\mathsf{new}\, n : \mathsf{E})\, M$ if $n' \notin \mathsf{E}, n \notin \mathsf{E}'$

Figure 3: Structural equivalence for SAFEDPI

will have no effect until the site l makes available a corresponding thread of the form $l[\![p?(\xi)\, Q]\!]$; using the rule (R-COMM) this will now allow the effective entry of F. In this manner the site l can control the immigration of code.

The rule (R-C.CREATE) exports the new channel name c generated by a process at k to the system level, where it is tagged with the declaration type $\mathsf{C}@k$; this records the location of the new channel. There is a corresponding rule for registered names, (R-N.CREATE); but such names are global and therefore there is no need to record where they were declared. The generation of new locations is governed by (R-L.CREATE):

$$k[\![(\mathsf{newloc}\, l : \mathsf{L})\, \mathsf{with}\, C \ \mathsf{in}\ P]\!] \longrightarrow (\mathsf{new}\, l : \mathsf{L})(k[\![P]\!] \mid l[\![C]\!])$$

The code C is set to run at the new location l, and note that this name is known to the continuation thread P running at the initiating location k.

The remaining axioms are self-explanatory; there is testing of simple identifiers in (R-EQ), (R-NEQ), β-reduction in the rule (R-BETA) for instantiating scripts and a standard rule for iterated processes.

For examples of reductions see Section 4.

3 Typing

In this section we discuss the types and type inference for SAFEDPI. There are three subsections. The first discusses informally the types used, which builds on those in [24, 10, 8, 26], while the second describes the type environments required to infer that systems are well-typed. Because the details are heavily syntactic, on first reading it may be better to skip directly to the final subsection which deals with the type inference rules, referring to the first two sections only on a call-by-need basis.

Basic types:	$\mathbf{base} ::= \mathbf{int} \mid \mathbf{string} \mid \mathbf{unit} \mid \top \mid \mathbf{proc} \mid \ldots$
Local Channels:	$\mathsf{C}, \mathsf{D} ::= \mathsf{r}\langle \mathsf{T} \rangle \mid \mathsf{w}\langle \mathsf{T} \rangle \mid \mathsf{rw}\langle \mathsf{T}, \mathsf{U} \rangle$
Locations:	$\mathsf{L}, \mathsf{K} ::= \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n],\ n \geq 0$ provided $u_i = u_j$ implies $i = j$
Global resources:	$\mathsf{N} ::= \mathsf{rc}\langle \mathsf{C} \rangle$
First-order:	$\mathsf{A} ::= \mathbf{base} \mid \mathsf{C} \mid \mathsf{L} \mid \mathsf{N} \mid \mathsf{C}@w$
Processes:	$\pi ::= \mathbf{proc} \mid \mathsf{pr}[u_1 : \mathsf{C}_1@w_1, \ldots, u_n : \mathsf{C}_n@w_n]$ provided $u_i = u_j, w_i = w_j$ implies $i = j$
Scripts:	$\mathsf{S} ::= \mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi)$
Values:	$\mathsf{T}, \mathsf{U} ::= \mathsf{A} \mid \mathsf{S} \mid \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\, \mathsf{T} \mid \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{T}})\, \mathsf{T}$

Figure 4: Type expressions - informal

3.1 The Types

The collection of types is an extension of those used in [8, 10], to which the reader is referred for more background and motivation. They are described informally in Figure 4 and intuitively they may be classified as follows:

Base types, ranged over by **base**: We include some predefined collection of types such as **int**, **unit**, **bool**, etc. for various constants in the language. The association of a particular type with a particular constant will be *global*, that is not dependent on a particular location. We also include **proc**, to indicate that a process is well-typed, and a *top* type $\top$, which can be associated with any identifier.

Local channel types, ranged over by C, D: These take the form

$$\mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \rangle$$

where T_r, T_w are *transmission* or *value types*; that is types of values which may be transmitted along channels. If an agent has a name at this type then it can transmit values of *at most* type T_w along it and receive from it values which have *at least* type T_r. In the formal description of types there will be a subtyping constraint, that T_w must be a subtype of T_r, explained in detail in [19]. When the transmit and receive types coincide we abbreviate this type by $\mathsf{rw}\langle \mathsf{T} \rangle$. We also allow the types $\mathsf{w}\langle \mathsf{T}_w \rangle$ and $\mathsf{r}\langle \mathsf{T}_r \rangle$, which only allow the transmission, reception respectively, of values.

Global resource name types, ranged over by N: These take the form $\mathsf{rc}\langle \mathsf{C} \rangle$, where C is a channel type. Intuitively these are the types of

names which are available to be used in the declaration of new locations. They allow an individual resource name, such as print, to be used in multiple locations, resulting in a form of *dynamic typing*.

Location types, ranged over by K, L: The standard form for these is

$$\mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n]$$

where C_i are channel types, and the identifiers u_i are distinct. An agent possessing a location name k with this type may use the channels/resources u_i located there at the types C_i; from the point of view of the agent, k is a site which offers the services $u_1, \ldots u_n$ at the corresponding types. In the formal definition we will require each u_i to be already declared as a global resource name. If n is zero then the agent knows of the existence of k but has no right to use resources there. We abbreviate this trivial type from $\mathsf{loc}[]$ to loc. We also identify location types up to re-orderings.

Process types, ranged over by π. The simplest process type is **proc**, which can be assigned to any well-typed process. More fine-grained process types take the form

$$\mathsf{pr}[u_1 : \mathsf{C}_1@w_1, \ldots, u_n : \mathsf{C}_n@w_n]$$

where the pairs (u_i, w_i) are assumed to be distinct. A process of this type can use *at most* the resource names u_i at the location w_i with their specified types C_i; these types determine the locations at which the channels u_i may be used.

Script types, ranged over by S: The general form here is

$$\mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi)$$

Scripts of this type require parameters $(\tilde{v})$ of type $(\tilde{\mathsf{T}})$; when these are supplied the resulting process will be of type $\pi\{\!|\tilde{v}/\tilde{x}|\!\}$. In other words the type of the resulting process may in general depend on the parameters. In these types we allow π to contain occurrences of a special location constant here to denote the current location.
These types will be abbreviated to $(\tilde{\mathsf{T}} \to \pi)$ whenever the variables $(\tilde{x})$ do not appear in the process type π, that is when the type of result is in fact independent of the parameters.

Script types, a generalisation of those used in [26], are one major innovation of the current paper; they allow parameterised processes, or scripts to be transmitted. Examples of such types include

$\mathsf{w}\langle\mathsf{T}\rangle \to$ **proc:** the type of a script which is parameterised on a local channel name, on which write permission at type T is needed.

$(\mathsf{r}\langle\mathsf{R}\rangle, \mathsf{w}\langle\mathsf{W}\rangle@k) \to$ **proc:** a value of this type will be applied to a pair, the first element will be a local channel with read capability at type R and the second a channel located at k with write capability at type W.

More importantly by using fine-grained process types, access to resources by incoming code can be restricted. Here are two examples:

$$\mathsf{Fdep}(x : \mathsf{r}\langle\mathsf{T}\rangle \to \mathsf{pr}[x : \mathsf{r}\langle\mathsf{T}\rangle@\mathsf{here},\ \mathsf{reply} : \mathsf{w}\langle\mathsf{T}\rangle@k])$$

Incoming code received at this type, can be instantiated by any local channel, say c, from which values can be read at type T. The resulting process is then only allowed access to two channels, namely the local channel c, from which it can read, and a channel named reply at the location k, to which it can write. This process will have the type $\mathsf{pr}[c : \mathsf{r}\langle\mathsf{T}\rangle@\mathsf{here},\ \mathsf{reply} : \mathsf{w}\langle\mathsf{T}\rangle@k]$. Code at the type

$$\mathsf{Fdep}((x, y, z) : (\mathsf{loc}, \mathsf{r}\langle\mathsf{T}\rangle@x, \mathsf{w}\langle\mathsf{T}\rangle) \to \mathsf{pr}[y : \mathsf{r}\langle\mathsf{T}\rangle@x,\ z : \mathsf{w}\langle\mathsf{T}\rangle@\mathsf{here}])$$

needs to be instantiated by a location, a channel at that location, and a local channel. For example the location could be called source, the channel located there info, from which values can be read at the type T, and the local channel record, at which values can be written at type T. The resulting process will then have type

$$\mathsf{pr}[\mathsf{info} : \mathsf{r}\langle\mathsf{T}\rangle@\mathsf{source},\ \mathsf{record} : \mathsf{w}\langle\mathsf{T}\rangle@\mathsf{here}]$$

It can download information from the third-party source site source via the channel info there.

Finally Transmission or V value types dictate the kind of values which can be transmitted over channels. These may be first order values, or scripts. We also allow dependent and existential types to be used. For example inputting a value of the dependent type $\mathsf{Tdep}(x : \mathsf{K})\,\mathsf{S}$ will result in the reception of a pair (k, F), where F is guaranteed to be of type $\mathsf{S}\{\!|k/x|\!\}$; k is the witness that the script F has the required type, and is received with the script. On the other hand inputting at the corresponding existential type $\mathsf{Edep}(x : \mathsf{K})\,\mathsf{S}$ will only result in the reception of the value F, although, as we will see, when the overall system is typechecked the witness v must be produced, to verify that F is indeed well-typed.

Notation 3.1 [Globalising types] It is worth noting that there is a crucial distinction between local channel types C and, for example, location types. The former only make sense relative to a specific location, whereas the latter are location independent, or *global types*. We can convert the local channel type C to a global type by appending a location, $\mathsf{C}@w$; this is the type of a channel of type C located at w. In various contexts it will be convenient to apply this globalisation operation to an arbitrary type, $(\mathsf{T})@w$; this will

only have an effect on any components of T which are local channel or script types. The operation is defined by induction on T:

$$(\mathsf{C})@w = \mathsf{C}@w, \quad (\mathsf{S})@w = \mathsf{S}$$
$$(\mathsf{K})@w = \mathsf{K}, \quad (\mathsf{C}@w')@w = \mathsf{C}@w'$$
$$(\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T})@w = \mathsf{Tdep}(\tilde{x} : (\tilde{\mathsf{T}})@w\,)\,(\mathsf{T})@w$$
$$(\mathsf{Edep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T})@w = \mathsf{Edep}(\tilde{x} : (\tilde{\mathsf{T}})@w\,)\,(\mathsf{T})@w$$

Note that in the last two clauses we have used the obvious notation $(\tilde{\mathsf{T}})@w$, for the list $\mathsf{T}_1@w, \ldots, \mathsf{T}_n@w$. ■

There are numerous constraints on the formation rules for types, well-documented in [10, 8]. The description given in Figure 4 should be viewed as defining *pre-types*; those which satisfy the formation constraints will then be considered to be types. It is best to describe these constraints relative to a *type environment*.

3.2 Type environments

A type judgement will take the form

$$\Gamma \vdash M$$

where Γ is a *type environment*, a list of assumptions about the types to be associated with the identifiers in the system M. These can take the form

- $u : \mathsf{loc}$, meaning that u is a location
- $u : \mathsf{C}@w$, meaning the channel u located at w has type C; for technical reasons C must be a *closed* type, that is containing no free ocurrences of variables.
- $u : \mathsf{rc}\langle\mathsf{C}\rangle$, meaning u is a global resource name, which may be installed at any new location.
- $x : \mathsf{S}$, meaning x can be instantiated by any script which can be inferred to have type S
- $x : \langle \mathsf{T}\ \mathsf{with}\ \tilde{y} : \tilde{\mathsf{E}}\rangle$. This represents a *package*, which will be used to handle existential types. Intuitively this defines the association $x : \mathsf{T}$ but the type T may depend on the auxiliary associations $\tilde{y} : \tilde{\mathsf{E}}$.

Lists of assumptions are created dynamically during typechecking, typically by augmenting a current environment with new assumptions on bound variables. It is convenient to introduce a particular notation for this operation:

Definition 3.2 (Forming environments) Let $\{V : \mathsf{T}\}$ be a list of type assumptions defined by

- $\{v : \mathsf{C}@w\} = v : \mathsf{C}@w$
- $\{x : \mathsf{S}\} = x : \mathsf{S}$
- $\{v : \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots u_n : \mathsf{C}_n]\} = v : \mathsf{loc}, u_1 : \mathsf{C}_1@v, \ldots u_n : \mathsf{C}_n@v$
- $\{(\tilde{y}, x) : \mathsf{Tdep}(\tilde{y} : \tilde{\mathsf{E}})\,\mathsf{T}\} = \{y_1 : \mathsf{E}_1\} \ldots, \{y_n : \mathsf{E}_n\},\ \{x : \mathsf{T}\}$
- $\{x : \mathsf{Edep}(\tilde{y} : \tilde{\mathsf{E}})\,\mathsf{T}\} = x : \langle \mathsf{T}\ \mathsf{with}\ \{y_1 : \mathsf{E}_1\} \ldots, \{y_n : \mathsf{E}_n\}\rangle$ ■

Of course there a lots of other possibilities for V and T but only those mentioned give rise to lists of assumptions. Moreover even those given may give rise to lists which are not consistent. For example we should not be able to introduce an assumption $u : \mathsf{loc}$ if u is already designated a channel, or introduce $u : \mathsf{C}@w$ unless w is known to be a location. Since type expressions also use identifiers, before introducing this assumption we would need to ensure that C is a properly formed type; for example it should only use identifiers which are already known. In order to describe the set of valid environments we introduce judgements of the form

$$\Gamma \vdash \mathbf{env}$$

The inference rules are straightforward and consequently are relegated to the appendix, in Figure 10. We also relegate to there the definition of subtyping judgements, of the form

$$\Gamma \vdash \mathsf{T} <: \mathsf{U},$$

given in Figure 11. Again the rules are straightforward, and mostly inherited from [8]. However it is worth noting that process types are ordered differently than location types. For example we have

$$\Gamma \vdash \mathsf{pr}[u_1 : \mathsf{C}_1@k] <: \mathsf{pr}[u_1 : \mathsf{C}_1@k,\ u_2 : \mathsf{C}_2@l]$$

but

$$\Gamma \vdash \mathsf{loc}[u_1 : \mathsf{C}_1, u_2 : \mathsf{C}_2] <: \mathsf{loc}[u_1 : \mathsf{C}_1]$$

assuming, of course, that the various types used, $\mathsf{C}_i, \mathsf{C}_j$ are well-defined relative to Γ.

These rules have been formulated so that they can also be used to say what is a valid type relative to a type expression.

Definition 3.3 (Valid types) We say the type expression T is a valid type relative to Γ, written $\Gamma \vdash \mathsf{T} : \mathbf{ty}$, whenever we can derive the judgement $\Gamma \vdash \mathsf{T} <: \mathsf{T}$. ■

Types can be viewed intuitively as *sets of capabilities* and *unioning* these sets corresponds to performing a *meet* operation with respect to subtyping. This we now explain. Let $(D, \leq)$ be a preorder. We say a subset $E \subseteq D$ is lower-bounded by $d \in D$ if $d \leq e$ for every e in E. Upper bounds are defined in a similar manner.

Definition 3.4 (partial meets and joins) We say that the preorder $(D, \leq)$ has *partial meets* if every pair of elements in D which has a lower bound also has a greatest lower bound. This means that for every pair of elements d_1, d_2 in D which has some lower bound, that is there is some element in $d \in D$ such that $d \leq d_1$, $d \leq d_2$, there is a particular lower bound, denoted $d_1 \sqcap d_2$ which is greater than or equal to every lower bound. The upper bound of pairs of elements, $d_1 \sqcup d_2$ is defined in an analogous manner. ■

Let $\mathbf{Types}_\Gamma$ denote the set of all type expressions T such that $\Gamma \vdash \mathsf{T} : \mathbf{ty}$.

Theorem 3.5 *For every Γ, the set $\mathbf{Types}_\Gamma$, ordered by $<:$, has partial meets and partial joins.*

Proof: See Proposition A.2 in Appendix A ■

Intuitively the existence of $\mathsf{T} \sqcap \mathsf{U}$ means that T and U are *compatible*, in that they allow compatible capabilities on values at these types. Moreover the type $\mathsf{T} \sqcap \mathsf{U}$ may be viewed as a *unioning* of the capabilities allowed by the individual types.

It is worth pointing out that with our type expressions set $\mathbf{Types}_\Gamma$ turns out to be not only a preorder but also a partial order. However this would no longer be the case if we allowed recursive types; nevertheless with this extension our results would still apply. Note also that because of the existence of the top type $\top$, useful in Section 6, joins of types are always guaranteed to exist.

3.3 Type Inference

We are now ready to describe the type inference system for ensuring that systems are well-typed. There are three forms of judgements, for systems, processes and values. The type inference rules for the first,

$$\Gamma \vdash M,$$

meaning that M is a well-typed system relative to Γ, are given in Figure 5. The intention is that whenever such a judgement can be inferred it will follow that Γ is a well-formed environment.

The main inference rule is (TY-PROC). In order to ensure that $k[\![P]\!]$ is a well-typed system we must show that the process P is well-typed to run at k; at the system level it is sufficient to be able to associate *any* process type

$$\frac{\Gamma, n : \mathsf{rc}\langle \mathsf{C} \rangle \vdash M}{\Gamma \vdash (\mathsf{new}\, n : \mathsf{rc}\langle \mathsf{C} \rangle)\, M}\ (\text{TY-GNEW}) \qquad \frac{\Gamma, c : \mathsf{C}@k \vdash M}{\Gamma \vdash (\mathsf{new}\, c : \mathsf{C}@k)\, M}\ (\text{TY-CNEW})$$

$$\frac{\Gamma \vdash \mathbf{env}}{\Gamma \vdash \mathbf{0}}\ (\text{TY-NIL}) \qquad \frac{\Gamma \vdash M \quad \Gamma \vdash N}{\Gamma \vdash M \mid N}\ (\text{TY-PAR})$$

$$\frac{\Gamma \vdash_k P : \mathbf{proc}}{\Gamma \vdash k[\![P]\!]}\ (\text{TY-PROC}) \qquad \frac{\Gamma, \{k : \mathsf{K}\} \vdash M \quad \Gamma, \{k : \mathsf{K}\} \vdash k : \mathsf{K}}{\Gamma \vdash (\mathsf{new}\, k : \mathsf{K})\, M}\ (\text{TY-LNEW})$$

Figure 5: Typing Systems

with P. The typing of processes must be relative to a location because it may use *local* channels which are required to exist at k; it also turns out that typing of scripts will depend on their location. There is also a subtlety in the typing of name creation. First note that in these, and all subsequent rules, we assume that all bound names in a judgement must be different than any free names used as part of the judgement. Thus in (TY-CNEW) we know that c is actually fresh to Γ. However we are still not guaranteed that $\Gamma, c : \mathsf{C}@k$ is a well-defined environment even when Γ is. From the type environment rules it will only be so when C is a well-defined type expression relative to Γ, and k is known as a location. There is a further complication in (TY-LNEW), the rule for new location creation. Deriving $\Gamma, \{k : \mathsf{K}\} \vdash M$ will ensure that $\Gamma, \{k : \mathsf{K}\}$ is a well-defined environment, but we must also ensure that all of the channels used in the location type K have already been declared, in Γ, as global resource names. This is enforced by the second requirement, $\Gamma, \{k : \mathsf{K}\} \vdash k : \mathsf{K}$.

The typing rules for the judgements on processes

$$\Gamma \vdash_w P : \pi$$

are given in Figure 7, and are defined simultaneously with the judgements for values, in Figure 6,

$$\Gamma \vdash V : \mathsf{T}$$

Let us first examine those for values. The rule (TY-LOOKUP) simply looks up the type of the identifier v relative to w in Γ, whereas (TY-BASE) allows base values to be typed for free. Note that the rule (TY-LOC) ensures that the judgement $\Gamma \vdash_w v : \mathsf{K}$, where K is a location type, can only be made when each channel used in K is already known to Γ, at a suitable type, as

(TY-LOOKUP)
$$\frac{\Gamma, v : (\mathsf{E})@w, \Gamma' \vdash \mathbf{env}}{\Gamma, v : (\mathsf{E})@w, \Gamma' \vdash_w v : \mathsf{E}}$$

(TY-BASE)
$$\frac{\Gamma \vdash \mathbf{env}}{\Gamma \vdash_w b : \mathbf{base}} \quad b \in \mathbf{base}$$

(TY-SUBVAL)
$$\frac{\Gamma \vdash_w V : \mathsf{T} \qquad \Gamma \vdash \mathsf{T} <: \mathsf{T}'}{\Gamma \vdash_w V : \mathsf{T}'}$$

(TY-MEET)
$$\frac{\Gamma \vdash_w u : \mathsf{T}_1 \qquad \Gamma \vdash_w u : \mathsf{T}_2}{\Gamma \vdash_w u : \mathsf{T}_1 \sqcap \mathsf{T}_2}$$

(TY-LOC)
$$\frac{\Gamma \vdash_v u_i : \mathsf{C}_i \qquad \Gamma \vdash_w u_i : \mathsf{rc}\langle \mathsf{D}_i \rangle \qquad \Gamma \vdash \mathsf{D}_i <: \mathsf{C}_i}{\Gamma \vdash_w v : \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n]}$$

(TY-TUDEP)
$$\frac{\Gamma \vdash_w v_i : \mathsf{E}_i\{\tilde{v}/\tilde{x}\} \qquad \Gamma \vdash_w v : \mathsf{T}\{\tilde{v}/\tilde{x}\}}{\Gamma \vdash_w (\tilde{v}, v) : \mathsf{Tdep}(\widetilde{x} : \widetilde{\mathsf{E}})\,\mathsf{T}}$$

(TY-EDEP)
$$\frac{\Gamma \vdash_w v_i : \mathsf{E}_i\{\tilde{v}/\tilde{x}\} \qquad \Gamma \vdash_w v : \mathsf{T}\{\tilde{v}/\tilde{x}\}}{\Gamma \vdash_w \langle \tilde{v}, v \rangle : \mathsf{Edep}(\widetilde{x} : \widetilde{\mathsf{E}})\,\mathsf{T}}$$

(TY-ELOOKUP)
$$\frac{\Gamma, y : \langle (\mathsf{T})@w \text{ with } \tilde{x} : \tilde{\mathsf{E}} \rangle, \Gamma' \vdash \mathbf{env}}{\Gamma, y : \langle \mathsf{T}@w \text{ with } \widetilde{x} : \widetilde{\mathsf{E}} \rangle, \Gamma' \vdash_w y : \mathsf{T}}$$

(TY-UNPACK)
$$\frac{\Gamma \vdash_w \langle \tilde{v}, v \rangle : \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{T}}{\Gamma \vdash_w v : \mathsf{T}\{\tilde{v}/\tilde{x}\}}$$

Figure 6: Typing Values

a global resource name. The rule (TY-MEET) is required because in certain circumstances we allow multiple associations with identifiers in valid environments; of course it can only be applied for types $\mathsf{T}_1, \mathsf{T}_2$ for which $\mathsf{T}_1 \sqcap \mathsf{T}_2$ exists. Dependent tuple values are typed with (TY-TUDEP). The value $(\tilde{v}, v)$ can be assigned the type $\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{T}$ provided each v_i can be assigned the type $\mathsf{E}_i\{\tilde{v}/\tilde{x}\}$ and v the type $\mathsf{T}\{\tilde{v}/\tilde{x}\}$. For existential types we need to invent a new kind of value $\langle \tilde{v}, v \rangle$; these do not occur in the language SAFEDPI, and are only used by the type inference system; intuitively $\langle \tilde{v}, v \rangle$ is a *package* consisting of the value v together with the witnesses $\tilde{v}$, which provide evidence (for the type inference system) that v has it's required type. The rule (TY-EDEP), which might also be called (TY-PACK), allows us to construct such values. It is similar to the rule for dependent tuples. The package $\langle \tilde{v}, v \rangle$ can be assigned the type $\mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{T}$ provided we can establish that v_i can be assigned the type $v_i : \mathsf{E}_i\{\tilde{v}/\tilde{x}\}$ and v the type $\mathsf{T}\{\tilde{v}/\tilde{x}\}$. Dependent tuples can be deconstructed and their components accessed in the standard manner; see the fourth clause of Definition 3.2. However the corresponding deconstruction for existential types only allows access to the final component, and not the witnesses; (TY-UNPACK) allows the value, rather than the witnesses, to be extracted at the appropriate type

(TY-OUT)
$$\frac{\Gamma \vdash_w V : \mathsf{T} \qquad \Gamma \vdash \mathsf{pr}[u : \mathsf{w}\langle \mathsf{T} \rangle@w] <: \pi \qquad \Gamma \vdash \mathsf{pr}_{\mathsf{ch}}[V : (\mathsf{T})@w] <: \pi}{\Gamma \vdash_w u!\langle V \rangle : \pi}$$

(TY-OUTE)
$$\frac{\Gamma \vdash_w \langle \tilde{v}, v \rangle : \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{T} \qquad \Gamma \vdash \mathsf{pr}[u : \mathsf{w}\langle \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{T} \rangle@w] <: \pi \qquad \Gamma \vdash \mathsf{pr}_{\mathsf{ch}}[\tilde{v} : (\tilde{\mathsf{E}})@w] <: \pi}{\Gamma \vdash_w u!\langle v \rangle : \pi}$$

(TY-IN)
$$\frac{\Gamma \vdash \mathsf{pr}[u : \mathsf{r}\langle \mathsf{T} \rangle@w] <: \pi \qquad \Gamma,\ \{X : (\mathsf{T})@w\} \vdash_w P : \pi \sqcup \mathsf{pr}_{\mathsf{ch}}[X : (\mathsf{T})@w]}{\Gamma \vdash_w u?(X : \mathsf{T})\,P : \pi}$$

(TY-SUBPROC)
$$\frac{\Gamma \vdash_w P : \pi \qquad \Gamma \vdash \pi <: \pi'}{\Gamma \vdash_w P : \pi'}$$

(TY-GO)
$$\frac{\Gamma \vdash_u v!\langle F \rangle : \pi}{\Gamma \vdash_w \mathsf{goto}_v\, u.F : \pi} \quad v \text{ a port}$$

(TY-STOP)
$$\frac{\Gamma \vdash \pi : \mathsf{ty}}{\Gamma \vdash_w \mathsf{stop} : \pi}$$

(TY-NEWLOC)
$$\frac{\Gamma, \{k : \mathsf{K}\} \vdash_k C : \pi \qquad \Gamma, \{k : \mathsf{K}\} \vdash_w P : \pi \qquad \Gamma, \{k : \mathsf{K}\} \vdash_w k : \mathsf{K}}{\Gamma \vdash_w (\mathsf{newloc}\, k : \mathsf{K})\,\mathsf{with}\, C\ \mathsf{in}\ P : \pi}$$

(TY-NEWCHAN)
$$\frac{\Gamma, c : \mathsf{C}@w \vdash_w P : \pi \sqcup \mathsf{pr}[c@w : \mathsf{C}]}{\Gamma \vdash_w (\mathsf{newc}\, c : \mathsf{C})\ P : \pi}$$

(TY-EQ)
$$\frac{\Gamma \vdash_w u_1 : \mathsf{A}_1, u_2 : \mathsf{A}_2 \qquad \Gamma \vdash_w Q : \pi \qquad \Gamma, \{u_1 : \mathsf{A}_2\}, \{u_2 : \mathsf{A}_1\} \vdash_w P : \pi \qquad (\text{provided } \Gamma, \{u_1 : \mathsf{A}_2\}, \{u_2 : \mathsf{A}_1\} \vdash \mathbf{env})}{\Gamma \vdash_w \mathsf{if}\ u_1 = u_2\ \mathsf{then}\ P\ \mathsf{else}\ Q : \pi}$$

(TY-NEWREG)
$$\frac{\Gamma, n : \mathsf{N} \vdash_w P : \pi}{\Gamma \vdash_w (\mathsf{newreg}\, n : \mathsf{N})\ P : \pi}$$

(TY-ABS)
$$\frac{\Gamma, \{\tilde{x} : (\tilde{\mathsf{T}})@w\} \vdash_w P : \pi\{w/\mathsf{here}\}}{\Gamma \vdash_w \lambda\,(\widetilde{x} : \widetilde{\mathsf{T}}).\ P : \mathsf{Fdep}(\widetilde{x} : \widetilde{\mathsf{T}} \to \pi)}$$

(TY-BETA)
$$\frac{\Gamma \vdash_w F : \mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi) \qquad \Gamma \vdash_w v_i : \mathsf{T}_i}{\Gamma \vdash_w F\ (\tilde{v}) : \pi\{\tilde{v}/\tilde{x}\}\{w/\mathsf{here}\}}$$

(TY-ITER)
$$\frac{\Gamma \vdash_w P : \pi}{\Gamma \vdash_w * P : \pi}$$

(TY-PAR)
$$\frac{\Gamma \vdash_w P : \pi \qquad \Gamma \vdash_w Q : \pi}{\Gamma \vdash_w P \mid Q : \pi}$$

Figure 7: Typing Processes

from the package. Similarly (TY-ELOOKUP) only allows knowledge of the value, and not the witnesses, to be deduced from an existential assumption.

In Figure 7 the rules for name generation, (TY-NEWCHAN),(TY-NEWLOC) and (TY-NEWREG), are simple adaptations of the corresponding rules at the system level; note that in (TY-NEWLOC) we are guaranteed that the

new name k does not occur in the type π, because of our convention on bound names; similarly for c in (TY-NEWCHAN) and n in (TY-NEWREG). (TY-STOP), (TY-ITER) and (TY-PAR) need no commentary. The rule (TY-EQ) is adapted from the analogous rule (TY-MATCH) in [10, 8]; recall that A_i range over first-order types and therefore this rule only allows the testing of first-order identifiers. (TY-ABS) and (TY-BETA) are standard rules for abstraction and application, adapted to dependent function types. But note the use of $\{\tilde{x} : (\tilde{\mathsf{T}})@w\}$ in the premise of the former; the arguments in an abstraction are relativised to the current location w. The rule for migration, (TY-GO), is justified by the reduction semantics, although we could easily have phrased it in terms of the premises of the output rule.

However the real interest is in the typing of the input and output processes. To ensure $u!\langle V\rangle$ has a process type π relative to Γ, (TY-OUT), we have to ensure that u has the output capability at some type appropriate to V. Thus we need to find some type T such that $\Gamma \vdash_w V : \mathsf{T}$ and u has the output capability on T. But we must *also* check that this capability is allowed by π. Both of these requirements are encapsulated in the second premise of the rule

$$\Gamma \vdash \mathsf{pr}[u : \mathsf{w}\langle\mathsf{T}\rangle@w] <: \pi$$

But there is a further complication. If the value being sent, V, contains channels, or more precisely capabilities on channels, then these must also be allowed by π. This is the intent of the third premise

$$\Gamma \vdash \mathsf{pr}_{\mathsf{ch}}[V : \mathsf{T}@w] <: \pi$$

which uses a (partial) function which constructs a process type from a value V and its type; it essentially extracts out any channels which may be in V. To define this we use $\sqcup$ which is a *join* operator on types, relative to $<:$ the subtyping order; when applied to process types it effectively takes the union of the capabilities of the individual types. It is worth noting that $\mathsf{pr}_{\mathsf{ch}}[v : \mathsf{T}]$ is the trivial process type $\mathsf{pr}[]$ when T is a script type.

$$\begin{aligned}
\mathsf{pr}_{\mathsf{ch}}[v : \mathsf{C}@w] &= \mathsf{pr}[v : \mathsf{C}@w] \\
\mathsf{pr}_{\mathsf{ch}}[v : \mathsf{K}] &= \mathsf{pr}[c_1 : \mathsf{C}_1@v, \ldots, c_n : \mathsf{C}_n@v] \\
&\qquad \text{where } \mathsf{K} = \mathsf{loc}[c_1 : \mathsf{C}_1, \ldots, c_n : \mathsf{C}_n] \\
\mathsf{pr}_{\mathsf{ch}}[\tilde{v} : \tilde{\mathsf{T}}] &= \mathsf{pr}_{\mathsf{ch}}[v_1 : \mathsf{T}_1] \sqcup \ldots \sqcup \mathsf{pr}_{\mathsf{ch}}[v_1 : \mathsf{T}_1] \\
\mathsf{pr}_{\mathsf{ch}}[(\tilde{v}, v) : \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{T}] &= \mathsf{pr}_{\mathsf{ch}}[\tilde{v} : \tilde{\mathsf{E}}] \sqcup \mathsf{pr}_{\mathsf{ch}}[v : \mathsf{T}] \\
\mathsf{pr}_{\mathsf{ch}}[\langle\tilde{v}, v\rangle : \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{T}] &= \mathsf{pr}_{\mathsf{ch}}[\tilde{v} : \tilde{\mathsf{E}}] \sqcup \mathsf{pr}_{\mathsf{ch}}[v : \mathsf{T}] \\
\mathsf{pr}_{\mathsf{ch}}[v : \mathsf{T}] &= \mathsf{pr}[] \qquad \text{otherwise}
\end{aligned}$$

The rule for transmitting existential values, (TY-OUTE) is a slight variation. We must establish a *package* $\langle\tilde{v}, v\rangle$ of the correct outgoing type, but only the (unpacked) value v is actually transmitted. Finally to ensure $u?(X : \mathsf{T})\,P$ has the type π, we need to check that u has the appropriate read capability, which also is allowed by π,

$$\Gamma \vdash \mathsf{pr}[u : \mathsf{r}\langle\mathsf{T}\rangle@w] <: \pi$$

and that the capabilities exercised by the residual P are either allowed by π or inherited by values which are input and bound to X:

$$\Gamma,\ \{X : (\mathsf{T})@w\} \vdash_w P : \pi \sqcup \mathsf{pr}_{\mathsf{ch}}[X : \mathsf{T}@w]$$

It is worth noting that the typing rules for input and output degenerate to the more standard form, for example as in [10], when we wish to establish that the processes are simply well-typed, that is have the type **proc**. For example we have the derived instances:

$$\frac{\Gamma \vdash_w V : \mathsf{T} \quad \Gamma \vdash_w u : \mathsf{w}\langle\mathsf{T}\rangle}{\Gamma \vdash_w u!\langle V\rangle : \mathbf{proc}}\ (\text{TY-OUT}) \qquad \frac{\Gamma \vdash_w u : \mathsf{r}\langle\mathsf{T}\rangle \quad \Gamma,\ \{X : (\mathsf{T})@w\} \vdash_w P : \mathbf{proc}}{\Gamma \vdash_w u?(X : \mathsf{T})\,P : \mathbf{proc}}\ (\text{TY-IN})$$

4 Examples

In this section we demonstrate the usefulness of the type system by a series of examples of increasing sophistication.

To make the examples more readable let us introduce some convenient notation. First we will abbreviate the transmission type $\mathbf{unit} \to \mathbf{proc}$, for thunked processes, simply to thunk. Then we use run as an abbreviation for the term $\lambda\xi\ \xi()$, where $()$ is the only value of type unit. So the type of run is $\mathsf{thunk} \to \mathbf{proc}$; it takes a thunked process and runs it. Thunked processes, which we often refer to as *thunks*, take the form $\lambda\ ().\,P$ but in the context of $\mathsf{goto}\,p.\ldots$ and port outputs $p!\langle\ldots\rangle$ we will omit the λ abstraction; thus $\mathsf{goto}_{in}\,l.\lambda\ ().\,P$ is abbreviated to $\mathsf{goto}_{in}\,l.P$. Finally we mimic the notation of process types for thunks, by letting $\mathsf{th}[....]$ denote the type $\mathbf{unit} \to \mathsf{pr}[....]$.

4.1 Installing and broadcasting services

Suppose there are two globally defined channel names ping and fing and a port name in; that is we are working in a type environment Γ with the property that

$$\Gamma \vdash \mathsf{ping} : \mathsf{rc}\langle\mathsf{D}_p\rangle,\ \mathsf{fing} : \mathsf{rc}\langle\mathsf{D}_f\rangle, \mathsf{in} : \mathsf{rc}\langle\mathsf{D}_i\rangle \qquad (2)$$

for some types $\mathsf{D}_p, \mathsf{D}_f$ and D_i. Let L be a location type such that

$$\mathsf{L} <: \mathsf{loc}[\mathsf{in} : \mathsf{C}_i, \mathsf{ping} : \mathsf{C}_p, \mathsf{fing} : \mathsf{C}_f]. \tag{3}$$

Then in the system

$$r[\![(\mathsf{newloc}\, l : \mathsf{L})\, \mathsf{with}\, Q \;\; \mathsf{in} \;\; P]\!]$$

the site r generates a new location l with declaration type L; it evolves to the new system

$$(\mathsf{new}\, l : \mathsf{L})(r[\![P]\!] \mid l[\![Q]\!])$$

To be well-typed with respect to Γ we need that

- L is a proper declaration type for locations, that is $\Gamma, \{l : \mathsf{L}\} \vdash l : \mathsf{L}$. This means that all the resource names in L should be globally defined in Γ with a type which supports their use in L. For example this would require $\mathsf{D}_p <: \mathsf{C}_p$, $\mathsf{D}_f <: \mathsf{C}_f$ and $\mathsf{D}_i <: \mathsf{C}_i$ with respect to Γ.
- the residual P is well-typed to run at r, that is

$$\Gamma, \{l : \mathsf{L}\} \vdash_r P : \mathbf{proc}$$

- the installed code is well-typed to run at the new location l, that is

$$\Gamma, \{l : \mathsf{L}\} \vdash_l Q : \mathbf{proc}.$$

The residual P running at r now knows the location l and its type, and may make it known to other agents. Suppose P has the form

$$*\mathsf{dist}_1!\langle l\rangle \mid *\mathsf{dist}_2!\langle l\rangle \mid R$$

where dist_i are distribution channels at r for broadcasting information. Agents with access to these channels can find out about l. More importantly the type at which they receive l depends on the types of dist_i at the site r. For example suppose Γ contains

$$\begin{array}{c} \mathsf{dist}_1 : \mathsf{w}\langle \mathsf{loc}[\mathsf{in} : \mathsf{w}\langle \mathsf{l}\rangle, \mathsf{ping} : \mathsf{w}\langle \mathsf{V}_p\rangle]\rangle @ r, \\ \mathsf{dist}_2 : \mathsf{w}\langle \mathsf{loc}[\mathsf{in} : \mathsf{w}\langle \mathsf{l}\rangle, \mathsf{fing} : \mathsf{w}\langle \mathsf{V}_f\rangle]\rangle @ r \end{array} \tag{4}$$

for some types $\mathsf{l}, \mathsf{V}_p, \mathsf{V}_f$. Then agents finding out about l from the source dist_1 only know about the resource ping there (in addition to the port in), while if the source of information is dist_2 only fing may be used. Of course an agent may have access to both sources. If that is the case then they can eventually come to know l at the type $\mathsf{loc}[\mathsf{in} : \mathsf{w}\langle \mathsf{l}\rangle, \mathsf{ping} : \mathsf{w}\langle \mathsf{V}_p\rangle, \mathsf{fing} : \mathsf{w}\langle \mathsf{V}_f\rangle]$, thereby obtaining knowledge of both resources. But more generally access to l will be governed by ports such as in and their programming via the installed code Q.

4.2 Servicing resources

The installed code Q determines, at least initially, who has access to the newly created site l. A typical example of the installed code C may take the form

$$*\mathsf{in}?(\xi : \mathsf{thunk}) \;\; (\mathsf{run}\, \xi) \;\mid\; S$$

Entry will be allowed to any well-typed thread at the port in, and the thread can subsequently interact with the servicing code S. This will only be well-typed if the original declaration type for the global name in allows values of type thunk to be received at that port. For example it will be well-typed if $\Gamma \vdash \mathsf{in} : \mathsf{rc}\langle \mathsf{rw}\langle \mathsf{thunk}\rangle\rangle$, that is setting the declaration type D_i in (2) above to be thunk, and the type l in the typing for the sources at r, in (4), to be thunk also.

Note that there is some choice in the type at which in is declared at l, in (3) above. If C_i is set to $\mathsf{rw}\langle \mathsf{thunk}\rangle$ then the servicing code S at l can both read and write at in, but the type $\mathsf{r}\langle \mathsf{thunk}\rangle$ is sufficient for well-typing, if S never writes to that port.

Consider a thread running at r such as

$$r[\![\mathsf{dist}_1?(x : \mathsf{L}_p)\, \mathsf{goto}_{\mathsf{in}}\, x.\mathsf{ping!}\langle v\rangle]\!] \tag{5}$$

which gains knowledge of the newly created location l via the source dist_1. Here we use L_p to be an abbreviation for an instance of the type used in (4) above, $\mathsf{loc}[\mathsf{in} : \mathsf{w}\langle \mathsf{thunk}\rangle, \mathsf{ping} : \mathsf{w}\langle \mathsf{V}_p\rangle]$. This thread is well-typed,

$$\Gamma \vdash r[\![\mathsf{dist}_1?(x : \mathsf{L}_p)\, \mathsf{goto}_{\mathsf{in}}\, x.\mathsf{ping!}\langle v\rangle]\!]$$

provided the value v can be assigned the proper type for ping namely V_p. This follows from the fact that for such a Γ we can establish

$$\Gamma \vdash_r \mathsf{dist}_1?(x : \mathsf{L}_p)\, \mathsf{goto}_{\mathsf{in}}\, x.\mathsf{ping!}\langle v\rangle \;\; : \mathbf{proc}$$

which in turn follows from

$$\Gamma, \{x : \mathsf{L}_p\} \vdash_r \mathsf{goto}_{\mathsf{in}}\, x.\mathsf{ping!}\langle v\rangle \;\; : \mathbf{proc}$$

This is a consequence of applying the typing rule (TY-GO) to the judgement

$$\Gamma, \{x : \mathsf{L}_p\} \vdash_x \mathsf{in!}\langle \mathsf{ping!}\langle v\rangle\rangle \;\; : \mathbf{proc} \tag{6}$$

The type environment $\Gamma, \{x : \mathsf{L}_p\}$ takes the form

$$\Gamma, x : \mathsf{loc}, \mathsf{in} : \mathsf{w}\langle \mathsf{thunk}\rangle @ x, \mathsf{ping} : \mathsf{w}\langle \mathsf{V}_p\rangle @ x$$

Therefore (6) follows from an application of the simple form of the output rule (TY-OUT), provided we can establish

$$\Gamma, x : \mathsf{loc}, \mathsf{in} : \mathsf{w}\langle \mathsf{thunk}\rangle @ x, \mathsf{ping} : \mathsf{w}\langle \mathsf{V}_p\rangle @ x \vdash_x \lambda\,().\ \mathsf{ping!}\langle v\rangle : \mathsf{thunk},$$

that is

$$\Gamma, x : \mathsf{loc}, \mathsf{in} : \mathsf{w}\langle \mathsf{thunk}\rangle @ x, \mathsf{ping} : \mathsf{w}\langle \mathsf{V}_p\rangle @ x \vdash_x \mathsf{ping}!\langle v\rangle : \mathbf{proc}$$

Finally this requires the judgement

$$\Gamma, x : \mathsf{loc}, \mathsf{in} : \mathsf{w}\langle \mathsf{thunk}\rangle @ x, \mathsf{ping} : \mathsf{w}\langle \mathsf{V}_p\rangle @ x \vdash_x v : \mathsf{V}_p \qquad (7)$$

Note that this checking of v is carried out relative to the variable location x; so the type V_p needs to be some *global type*, whose meaning is independent of the current location. This could be a base type such as **int**, although we will see more interesting examples, such as *return channels*, later.

4.3 Site protection

A simple infrastructure for a typical site could take the form

$$h[\![\mathsf{in}?(\xi : \mathsf{I}) \ \ *\mathsf{run}\,\xi \ \mid \ S]\!]$$

The on-site code S could provide various services for incoming agents, repeatedly accepted at the input port in. In a relaxed computing environment the type I could simply be thunk indicating that any well-typed code will be allowed to immigrate. In the sequel we will always assume that when the type of the port in is not discussed it has this liberal type.

However constraints can be imposed on incoming code by only publicising ports which have associated with them more restrictive *guardian* types. In such cases it is important that read capabilities on the these ports be retained by the host. This point will be ignored in the ensuring discussion, which instead concentrates on the forms the guardian types can take.

Consider a system consisting of a server and client, defined below, running in parallel.

$$\begin{array}{ll} \text{Server:} & s[\![\mathsf{req}?(\xi : \mathsf{S})\,\mathsf{run}\,\xi \ \mid \ *\mathsf{news}!\langle \mathit{scandal}\rangle]\!] \\ \text{Client:} & c[\![\mathsf{goto}_{\mathsf{req}}\, s.\mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report}!\langle x\rangle \\ & \quad \mid \mathsf{in}?(\xi : \mathsf{R})\,\mathsf{run}\,\xi \ \mid \mathsf{report}?(y)\ldots]\!] \end{array} \qquad (8)$$

The server is straightforward; it accepts incoming code at the port req and runs it. The only service it provides is some information on a channel called news. The client, who knows of the req port at the server sends code there to collect the news and report it back to it's own channel report; the type at which it inputs from news, which obviously must be **string**, is elided. This code migrates twice, once via the port req from the client to the server, and once via the port in, from the server to the client.

The server protects its site using the guardian type S while the client protects its site using R. What should these be? Let us assume that both

sites have the required channels at appropriate types; suppose in Γ we have the entries

$$\mathsf{news} : \mathsf{rw}\langle \mathbf{string}\rangle @ s, \ \ \mathsf{req} : \mathsf{rw}\langle \mathsf{S}\rangle @ s,$$
$$\mathsf{report} : \mathsf{rw}\langle \mathbf{string}\rangle @ c, \ \ \mathsf{in} : \mathsf{rw}\langle \mathsf{R}\rangle @ c$$

The first possibility is for the client to be relaxed but the server vigilant:

$$\begin{array}{ll} \mathsf{R}: & \mathsf{thunk} \\ \mathsf{S}: & \mathsf{th}[\mathsf{news} : \mathsf{r}\langle \mathbf{string}\rangle @ s, \ \mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle @ c] \end{array}$$

Here the client allows in any well-typed process, whereas the server will only accept at the port req processes which use at most the local channel news and the port in at the site c; moreover the local channel news can only be used in read mode.

With these types one can show that the overall system is well-typed. Typing the server is straightforward but to type the client we need to establish, among other requirements,

$$\Gamma \vdash_c \mathsf{goto}_{\mathsf{req}}\, s.\mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report}!\langle x\rangle \ : \mathbf{proc}$$

As usual this follows by an application of (TY-GO) from

$$\Gamma \vdash_s \mathsf{req}!\langle\, \mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report}!\langle x\rangle\ \rangle \ : \mathbf{proc}$$

which in turn requires establishing

$$\Gamma \vdash_s \lambda\,().\ \mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report}!\langle x\rangle \ : \mathsf{S}$$

In other words the incoming code should match the guardian type of the server, S. By *dethunking* we get the requirement

$$\Gamma \vdash_s \mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report}!\langle x\rangle \ : \mathsf{pr}[\mathsf{news} : \mathsf{r}\langle \mathbf{string}\rangle @ s, \ \mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle @ c]$$

This is established via an application of the rule (TY-IN). The first premise is immediate since we assume $\Gamma \vdash_s \mathsf{news} : \mathsf{rw}\langle \mathbf{string}\rangle$. Moreover the second amounts to

$$\Gamma, x : \mathbf{string} \vdash_s \mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report}!\langle x\rangle \ : \mathsf{pr}[\mathsf{news} : \mathsf{r}\langle \mathbf{string}\rangle @ s, \ \mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle @ c]$$

because the value being received is a string; that is $\mathsf{pr}_{\mathsf{ch}}[x : \mathbf{string} @ s]$ is the trivial process type pr[].

The significant step in establishing this second premise is to check that the code returning to the client satisfies its guardian type R:

$$\Gamma, x : \mathbf{string} \vdash_c \mathsf{in}!\langle\, \mathsf{report}!\langle x\rangle\rangle \ : \mathsf{pr}[\mathsf{news} : \mathsf{r}\langle \mathbf{string}\rangle @ s, \ \mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle @ c] \qquad (9)$$

However this is straightforward since R is the liberal guardian thunk. It follows by an application of the output rule (TY-OUT) in Figure 7. But it is important to note that in the application the third premise is vacuous, as $\mathsf{pr}_{\mathsf{ch}}[\lambda\,().\ \mathsf{report!}\langle x\rangle : \mathbf{proc}]$ is the trivial type $\mathsf{pr}[]$.

The current type $\mathsf{R} = \mathsf{thunk}$ leaves the client site open to abuse but it is easy to check that the above reasoning is still valid if the guardians are changed to

$$\begin{array}{ll} \mathsf{R}: & \mathsf{th}[\mathsf{report} : \mathsf{w}\langle\mathbf{string}\rangle@c] \\ \mathsf{S}: & \mathsf{th}[\mathsf{news} : \mathsf{r}\langle\mathbf{string}\rangle@s,\ \mathsf{in} : \mathsf{w}\langle\mathsf{R}\rangle@c] \end{array}$$

Here the guardian for the client only allows in agents which write to the local port report; note that this change requires that the guardian at the server site also uses this more restrictive type in its annotation for the port in at c.

One can check that with these new restrictive guardians the system is still well-typed. The only extra work required is in providing a proof for the judgement (9) above, ensuring that the code returning to the client satisfies the more demanding guardian. By an application of (TY-GO) and (TY-OUT) this reduces to the judgement

$$\Gamma, x : \mathbf{string} \vdash_c \lambda\,().\ \mathsf{report!}\langle x\rangle\ : \mathsf{th}[\mathsf{report} : \mathsf{w}\langle\mathbf{string}\rangle@c]$$

which is a straightforward consequence of (TY-OUT).

It might be tempting to define the guardians by

$$\begin{array}{ll} \mathsf{R}: & \mathsf{th}[\mathsf{report} : \mathsf{w}\langle\mathbf{string}\rangle@c] \\ \mathsf{S}: & \mathsf{th}[\mathsf{news} : \mathsf{r}\langle\mathbf{string}\rangle@s,\ \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@c] \end{array}$$

Here both server and client protect their own resources but the server is uninterested in what happens at the client site, by allowing code with arbitrary capabilities on the client port in. However there is an intuitive inconsistency here. The client has read capability at its port, at the restrictive type R, while somehow the server has obtained a more liberal write capability there, namely thunk.

In fact the system can not be typed with these revised guardians. In particular

$$\Gamma \not\vdash s[\![\mathsf{req}?(\xi : \mathsf{S})\,\mathsf{run}\,\xi]\!]$$

Any derivation of this judgement would require the judgement

$$\Gamma, \xi : \mathsf{S} \vdash_s \mathsf{run}\,\xi$$

which in turn would require

$$\Gamma \vdash \mathsf{S} : \mathbf{ty}$$

or more formally

$$\Gamma \vdash \mathsf{S} <: \mathsf{S}$$

But as we will see this can not be inferred; that is S is not a valid type, relative to Γ.

To see why let us suppose, for simplicity, that the port in has been declared at the site c with a type of the form $\mathsf{rw}\langle\mathsf{R},\mathsf{W}\rangle$ for some type W. One constraint in the type formation rules, (see (TY-CHAN) in Figure 11) is that the write capabilities on a channel are always a subtype of the read capabilities; in our setting this means that $\Gamma \vdash \mathsf{W} <: \mathsf{R}$. Our rules also ensure that $\Gamma \vdash_c \mathsf{in} : \mathsf{w}\langle\mathsf{T}_w\rangle$ implies $\Gamma \vdash \mathsf{T}_w <: \mathsf{W}$ and consequently $\Gamma \vdash \mathsf{T}_w <: \mathsf{R}$.

However the structure of R ensures that $\Gamma' \vdash \mathsf{thunk} <: \mathsf{R}$ for no Γ', from which we can conclude that $\Gamma \not\vdash_c \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@c$. But this is one of the requirements, in the formation rules in Figure 11, to establish $\Gamma \vdash \mathsf{S} : \mathbf{ty}$.

4.4 Anonymous channels

Consider the following variation on the server/client system:

$$\begin{array}{ll} \text{Server:} & s[\![\mathsf{req}?(\xi : \mathsf{S})\,\mathsf{run}\,\xi \mid \mathsf{where}?((y,z) : \mathsf{T})\,\mathsf{goto}_{\mathsf{in}}\,y.z!\langle \mathit{scandal}\rangle]\!] \\ \text{Client:} & c[\![\mathsf{goto}_{\mathsf{req}}\,s.\mathsf{where}!\langle c, \mathsf{report}\rangle \mid \mathsf{in}?(\xi : \mathsf{R})\,\mathsf{run}\,\xi \mid \mathsf{report}?(y)\ldots]\!] \end{array} \tag{10}$$

Here the protocol is somewhat different; the client simply supplies to the server, via the channel where, the address of a channel on which to supply the news; this consists of the pair of a location and a channel on which to report. The server then launches a thread which migrates to the relevant location, which is assumed to have an in port, to deliver the news.

Defining guardians is straightforward. For example these could be

$$\begin{array}{ll} \mathsf{R}: & \mathsf{thunk} \\ \mathsf{S}: & \mathsf{th}[\mathsf{where} : \mathsf{w}\langle\mathsf{T}\rangle@s,\ \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@c] \end{array}$$

However the difficulty is in ascertaining the required type T for the pair of values. One possibility is to set

$$\mathsf{T} = (\mathsf{l},\ \mathsf{w}\langle\mathbf{string}\rangle)$$

where l is the location type $\mathsf{loc}[\mathsf{in} : \mathsf{w}\langle\mathsf{R}\rangle]$, allowing the first component to be a location with an in port at the appropriate type and the second to be a channel for sending strings.

Unfortunately the server can not be typed with such a T. The problem arises when we try to establish

$$\Gamma, \{(y,z) : (\mathsf{T})@s\} \vdash_s \mathsf{goto}_{\mathsf{in}}\,y.z!\langle \mathit{scandal}\rangle : \mathbf{proc} \tag{11}$$

Unravelling the extended environment this means establishing

$$\Gamma, y : \mathsf{loc}, \mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle @ y, z : \mathsf{w}\langle \mathbf{string}\rangle @ s \vdash_y z!\langle \mathit{scandal}\rangle : \mathbf{proc}$$

which is not possible; the output rule (TY-OUT) demands that z be a channel at the location y.

So to be able to statically type this example we need to be able to use the first component in the pair (y, z) as part of the type of the second component; we need a dependent type.

Let

$$\mathsf{T} = \mathsf{Tdep}(x : \mathsf{l})\, \mathsf{w}\langle \mathbf{string}\rangle @ x$$

Note that (SUB-TUDEP) from Figure 11 ensures that this is a well-defined type:

$$\Gamma \vdash \mathsf{T} : \mathbf{ty}$$

because

$$\Gamma, \{x : \mathsf{l}\} \vdash \mathsf{w}\langle \mathbf{string}\rangle @ x : \mathbf{ty}$$

So this type can be safely used as part of a process. Moreover it is now easy to establish (11) above as the extended environment $\Gamma, \{(y, z) : \mathsf{T}\}$ unravels to $\Gamma, y : \mathsf{loc},\ \mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle @ y,\ z : \mathsf{w}\langle \mathbf{string}\rangle @ y$.

These location dependent types were introduced in [10], where they are shown to be very useful for typing migrating code, as they allow the transmission of *anonymous* channels between sites. In our example the server does not need to know, apriori, the name of the report channel at the client site. In the sequel we will borrow the notation used in [10] for these dependent types; we use $(u @ w)$ to denote any pair of identifiers (u, w) which is expected to have a dependent type of the form $\mathsf{Tdep}(x : \mathsf{l})\, \mathsf{C}$. In a similar vein we abbreviate this type to $\mathsf{C} @ \mathsf{l}$. Thus we can reformulate the example (10) above as:

$$\begin{array}{ll} \text{Server} & s[\![\mathsf{req}?(\xi : \mathsf{S})\, \mathsf{run}\, \xi \mid \mathsf{where}?((z @ y) : \mathsf{w}\langle \mathbf{string}\rangle @ \mathsf{l})\, \mathsf{goto}_{\mathsf{in}}\, y.z!\langle \mathit{scandal}\rangle]\!] \\ \text{Client} & c[\![(\mathsf{new}\, \mathsf{c}\, \mathsf{report}) \\ & \qquad \mathsf{goto}_{\mathsf{req}}\, s.\mathsf{where}!\langle \mathsf{report} @ c\rangle \mid \\ & \qquad \mathsf{in}?(\xi : \mathsf{R})\, \mathsf{run}\, \xi \mid \mathsf{report}?(y) \ldots]\!] \end{array}$$

Here, as a form of self-protection, the client generates a new return channel, also called report and whose obvious type is elided, which it sends to the server. The client's self-protection consists of reading this channel exactly once, which it knows will be a response to its request to the server.

Note that these location dependent types are exactly what is required to type the example (5) above. In the type judgement (7) we need to find an appropriate type V_p for values transmitted on the channel ping. We can now let V_p be the dependent type $\mathsf{w}\langle \mathbf{string}\rangle @ \mathsf{loc}$, consisting of a return address; that is a location, and a write capability at some channel at that location.

4.5 Dependent process types

There remains a major difficulty with the server in (10) and (8) above. The guardian type of the server S uses the name of the client c, and therefore it can only be used by that client. To overcome this difficulty we need to allow process types to depend on locations and channels. Here the general form will be

$$\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{E}})\, \mathsf{S}$$

where S is a script type which may depend on the variables $\tilde{x}$. A value of this type takes the form

$$(\tilde{v}, v)$$

where v is some script. But again to emphasise the occurrence of these types we will use the more descriptive syntax

$$v \text{ with } \tilde{v}$$

An example of the use of such types is in the following variation of the client server from (8) above:

$$\begin{array}{ll} \text{Server:} & s[\![\mathsf{req}?(\xi \text{ with } y : \mathsf{S}_d)\, \mathsf{run}\, \xi \mid *\mathsf{news}!\langle \mathit{scandal}\rangle]\!] \\ \text{Client:} & c[\![(\mathsf{new}\, \mathsf{c}\, \mathsf{report}) \\ & \qquad \mathsf{goto}_{\mathsf{req}}\, s.\mathsf{news}?(x)\, \mathsf{goto}_{\mathsf{in}}\, c.\, \mathsf{report}!\langle x\rangle \text{ with } c \mid \\ & \qquad \mathsf{in}?(\xi : \mathsf{R})\, \mathsf{run}\, \xi \mid \mathsf{report}?(y) \ldots]\!] \end{array} \qquad (12)$$

with the types

$$\begin{array}{rl} \mathsf{R} : & \mathsf{thunk} \\ \mathsf{S}_d : & \mathsf{Tdep}(y : \mathsf{l})\, \mathsf{th}[\mathsf{news} : \mathsf{r}\langle \mathbf{string}\rangle @ s,\ \mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle @ y] \\ \mathsf{l} : & \mathsf{loc}[\mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle] \end{array}$$

Here the important point to notice is the server's guardian type at the port req, S_d, no longer mentions any clients name; it can be used by any client which satisfies the types requirements. The server accepts a thunk, of type $\mathsf{th}[\mathsf{news} : \mathsf{r}\langle \mathbf{string}\rangle @ s,\ \mathsf{in} : \mathsf{w}\langle \mathsf{R}\rangle @ y]$ which must be accompanied by a location of type l to be used in place of the variable y in S_d. A typical client c can generate a new reply channel report and send to the server

- the thunk $\mathsf{news}?(x)\, \mathsf{goto}_{\mathsf{in}}\, c.\, \mathsf{report}!\langle x\rangle$
- accompanied by a required location, in this case c.

Let us now see how the overall system typechecks, assuming as usual an environment in which the channel news and ports req, in, have the appropriate types, and that the declaration type of report is $\mathsf{rw}\langle \mathbf{string}\rangle$. At the server let us concentrate on establishing

$$\Gamma \vdash_s \mathsf{req}?(\xi \text{ with } y : \mathsf{S}_d)\, \mathsf{run}\, \xi : \mathbf{proc}$$

This follows by an application of the simple form of (TY-IN) to

$$\Gamma, \{(y,\xi) : (\mathsf{S}_d)@s\} \vdash_s \mathsf{run}\,\xi : \mathbf{proc}$$

Noting that $(\mathsf{S}_d)@s$ is the same as S_d, unravelling the extended environment gives the requirement

$$\Gamma, y : \mathsf{loc},\ \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@y,\ \xi : \mathsf{th}_y \vdash_s \mathsf{run}\,\xi$$

where th_y is an abbreviation for the type $\mathsf{th}[\mathsf{news} : \mathsf{r}\langle\mathbf{string}\rangle@s,\ \mathsf{in} : \mathsf{w}\langle\mathsf{R}\rangle@y]$. Apriori typing the process $\mathsf{run}\,\xi$ should be straightforward with respect to this environment. But there is a subtlety; at some point in establishing this judgement we need to apply (TY-BASE) from Figure 6 to conclude

$$\Gamma, y : \mathsf{loc},\ \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@y,\ \xi : \mathsf{th}_y \vdash_s () : \mathbf{unit}$$

and this requires the premise

$$\Gamma, y : \mathsf{loc},\ \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@y,\ \xi : \mathsf{th}_y \vdash \mathbf{env}$$

which in turn requires the premise

$$\Gamma, y : \mathsf{loc},\ \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@y \vdash \mathsf{th}_y : \mathbf{ty} \tag{13}$$

In other words we have to check that th_y is a well-defined type, relative to the extended environment. However this is now straightforward using the rule (SUB-PROC) from Figure 11, in the presence of the new associations involving y in the extended environment.

Let us now turn our attention to typechecking the client in (12) above, where we concentrate on ensuring that the process sent to the port req satisfies the type S_d. We have to ensure

$$\Gamma, \mathsf{report} : \mathsf{rw}\langle\mathbf{string}\rangle@c \vdash_s \lambda\,().\ \mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\,c.\,\mathsf{report}!\langle x\rangle\ \mathsf{with}\ c\ \ : \mathsf{S}_d$$

The rule (TY-TUDEP) in Figure 6 reduces this to two premises:

$$\Gamma, \mathsf{report} : \mathsf{rw}\langle\mathbf{string}\rangle@c \vdash_s c : \mathsf{I}$$
$$\Gamma, \mathsf{report} : \mathsf{rw}\langle\mathbf{string}\rangle@c \vdash_s \mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\,c.\,\mathsf{report}!\langle x\rangle\ : \mathsf{th}[\mathsf{news} : \mathsf{r}\langle\mathbf{string}\rangle@s,\ \mathsf{in} : \mathsf{w}\langle\mathsf{R}\rangle@c]$$

The first is immediate from our assumptions about Γ and the second is essentially the same as a derivation we have already seen on page 25.

Thus using dependent process types we can define general purpose servers which can be used by multiple clients. The example we have just considered, (12), apriori leaves the clients insecure because of the use of the liberal type thunk for the clients guardian type R. But it can be generalised so that this

guardian is strengthened, allowing in only threads which are going to write to the locally declared reporting channel. Here is one possible formulation:

$$\begin{array}{ll}\text{Server:} & s[\![\mathsf{req}?(\xi\ \mathsf{with}\ (y,z,x) : \mathsf{S}_d)\,\mathsf{run}\,\xi\ \mid\ *\mathsf{news}!\langle \mathit{scandal}\rangle]\!] \\ \text{Client:} & c[\![(\mathsf{newc}\,\mathsf{report}) \\ & \quad (\mathsf{newc}\,\mathsf{in} : \mathsf{rw}\langle\mathsf{R}\rangle) \\ & \quad\quad \mathsf{goto}_{\mathsf{req}}\,s.\mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\,c.\,\mathsf{report}!\langle x\rangle\ \mathsf{with}\ (c,\mathsf{report},\mathsf{in})\ \mid \\ & \quad\quad \mathsf{in}?(\xi : \mathsf{R})\,\mathsf{run}\,\xi\ \mid \mathsf{report}?(y)\ldots]\!]\end{array} \tag{14}$$

Here a client generates a local channel report, whose type $\mathsf{rw}\langle\mathbf{string}\rangle$ we have elided, and a local port in whose declaration type is $\mathsf{rw}\langle\mathsf{R}\rangle$, where R is the more restrictive guardian type $\mathsf{th}[\mathsf{report} : \mathsf{w}\langle\mathbf{string}\rangle@c]$. In other words in has been specially created to restrict entry to processes which will only write on the newly created channel report. The client then sends the usual process to the server but now accompanies it with the triple $(c, \mathsf{report}, \mathsf{in})$

The code for the server is the same except that accompanying the incoming thread it expects three values. Its guardian type S_d however is changed to

$$\begin{array}{ll}\mathsf{S}_d : & \mathsf{Tdep}(y : \mathsf{loc},\ z : \mathsf{w}\langle\mathbf{string}\rangle@y,\ x : \mathsf{w}\langle\mathsf{th}[z : \mathsf{w}\langle\mathbf{string}\rangle@y]\rangle@y) \\ & \qquad \mathsf{th}[\mathsf{news} : \mathsf{r}\langle\mathbf{string}\rangle@s,\ x : \mathsf{w}\langle\mathsf{th}[z : \mathsf{w}\langle\mathbf{string}\rangle@y]\rangle@y]\end{array}$$

Here, once more, this guardian type does not mention any client names, but it allows clients to employ much more restrictive guardian types at their own sites. We leave the reader to check that this revised system can still be typechecked.

4.6 Existential process types

The use of dependent script types, as in the previous subsection, has certain disadvantages from the point of view of the clients. For example in (14) above the client sends to the server, in addition to the script to be executed, the triple $(c, \mathsf{report}, \mathsf{in})$. Although these are not used by the server we have defined other than as part of the received script clients are in principle able to use them in any way they seem fit. An alternative server could be given by

$$\text{badServer:} \qquad s[\![\mathsf{req}?(\xi\ \mathsf{with}\ (y,z,x) : \mathsf{S}_d)\,\mathsf{goto}_x\,y.z!\langle \mathit{boring}\rangle]\!] \tag{15}$$

This rogue server does not run the incoming script to obtain the latest news; instead it uses the incoming accompanying values and sends directly to the client some boring data.

Existential types allow the client to hide from the server the data which accompanies the incoming scripts. These type take the form

$$\mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{S}$$

where, as with dependent types, the type of the script S may depend on the parameters $\tilde{x}$. Intuitively a value of this type is once more a form of tuple $(\tilde{v}, v)$, although access to the accompanying parameters is restricted. That is reading a value of this type from a port only results in the script being obtained, although that script itself may use these parameters. This new form of tuple, often called a *package*, is denoted by

$$\langle \tilde{v}, v \rangle$$

The important point about such a package is that it only gives access to the script v and not the internal parameters $\tilde{v}$. In our formulation to send such a value on a channel the sender must have the package $\langle \tilde{v}, v \rangle$, although only the script v is emitted. For this reason we need a special output rule for existential types; see (TY-OUTE) in Figure 7, which has already been explained in Section 3.3.

Let us now reformulate (14) above using existential types:

$$\begin{array}{ll}
\text{Server:} & s[\![\mathsf{req}?(\xi : \mathsf{S}_e)\,\mathsf{run}\,\xi \;\mid\; *\mathsf{news!}\langle \mathit{scandal} \rangle]\!] \\
\text{Client:} & c[\![(\mathsf{new}\,c\,\mathsf{report}) \\
& \quad (\mathsf{new}\,\mathsf{in} : \mathsf{rw}\langle \mathsf{R} \rangle) \\
& \qquad \mathsf{goto}_{\mathsf{req}}\, s.\mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report!}\langle x \rangle \;\mid \\
& \qquad \mathsf{in}?(\xi : \mathsf{R})\,\mathsf{run}\,\xi \;\mid \mathsf{report}?(y) \ldots]\!]
\end{array} \tag{16}$$

Here the guardian type S_e is

$$\begin{array}{l}
\mathsf{Edep}(y : \mathsf{loc}, z : \mathsf{w}\langle \mathbf{string} \rangle @ y, x : \mathsf{w}\langle \mathsf{th}[z : \mathsf{w}\langle \mathbf{string} \rangle @ y] \rangle @ y) \\
\qquad \mathsf{th}[\mathsf{info} : \mathsf{r}\langle \mathbf{string} \rangle @ s,\; x : \mathsf{w}\langle \mathsf{th}[z : \mathsf{w}\langle \mathbf{string} \rangle @ y] \rangle @ y]
\end{array}$$

The server is much the same as before except that it does not receive any parameters with the incoming script. Similarly the client only sends the script.

Let us now see that this example typechecks. Establishing that the server is well-typed is a little more complicated than with dependent type S_d. The interest centres on establishing

$$\Gamma, \{\xi : (\mathsf{S}_e) @ s\} \vdash_s \mathsf{run}\,\xi : \mathbf{proc}$$

and there are two essential steps. Note that, as with S_d, $(\mathsf{S}_e) @ s$ is the same as S_e, and so in the sequel we will omit the $(-) @ s$. The first step is deriving

$$\Gamma, \{\xi : \mathsf{S}_e\} \vdash_s () : \mathbf{unit}$$

and proceeds as with the use of S_d on page 29; unravelling the environment this amounts to establishing

$$\Gamma, \xi : \langle \mathsf{th}_y \text{ with } y : \mathsf{loc},\; z : \mathsf{w}\langle \mathbf{string} \rangle @ y, x : \mathsf{w}\langle \mathsf{th}[z : \mathsf{w}\langle \mathbf{string} \rangle] \rangle @ y \rangle \vdash \mathbf{env} \tag{17}$$

where now th_y represents $\mathsf{th}[\mathsf{info} : \mathsf{r}\langle \mathbf{string} \rangle @ s,\; \mathsf{in} : \mathsf{w}\langle \mathsf{th}[z : \mathsf{w}\langle \mathbf{string} \rangle] \rangle @ y]$. Here the relevant type formation rule is (E-EDEP) from Figure 10, which requires the premise

$$\Gamma, y : \mathsf{loc},\; z : \mathsf{w}\langle \mathbf{string} \rangle @ y, x : \mathsf{w}\langle \mathsf{th}[z : \mathsf{w}\langle \mathbf{string} \rangle] \rangle @ y \vdash \mathsf{th}_y : \mathbf{ty}$$

However this is easily established from (SUB-SCRIPT) of the same Figure.

The second essential step in typechecking the server is

$$\Gamma, \{\xi : \mathsf{S}_e\} \vdash_s \xi : \mathbf{proc} \tag{18}$$

This is necessary in order to ensure that run can be applied to ξ. Here we use an application of (TY-ELOOKUP) from Figure 6 to obtain

$$\Gamma, \{\xi : \mathsf{S}_e\} \vdash_s \xi : \mathsf{th}_y$$

One can also establish, using the subtyping rules,

$$\Gamma, \{\xi : \mathsf{S}_e\} \vdash \mathsf{th}_y <: \mathbf{proc}$$

and therefore by (TY-SUBVAL) from Figure 6 we obtain the required judgement (18) above.

Now let us examine the client. The central point is to ensure that the $\mathsf{goto}_{\mathsf{req}}\, s \ldots.$ command is well-typed, which amounts to establishing the judgement:

$$\Gamma, \mathsf{report} : \mathsf{rw}\langle \mathbf{string} \rangle @ c \vdash_s \mathsf{req!}\langle \mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report!}\langle x \rangle \rangle \;\; : \mathbf{proc}$$

Here the relevant rule is (TY-OUTE) from Figure 7. The second premise follows from our assumption about the type of req at s while the third is vacuous as π is instantiated to $\mathbf{proc}$. However the first premise requires us to find some $\tilde{v}$ such that

$$\Gamma, \mathsf{report} : \mathsf{rw}\langle \mathbf{string} \rangle @ c \vdash_s \langle \tilde{v}, \mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report!}\langle x \rangle \rangle \;\; : \mathsf{S}_e \tag{19}$$

In fact the required $\tilde{v}$ is obviously going to be $(c, \mathsf{report}, \mathsf{in})$.

With these values the judgement (19) can be established using the rule (TY-EDEP) from Figure 6. This requires the following four premises, where for convenience we use Γ_e as an abbreviation for the extended environment $\Gamma, \mathsf{report} : \mathsf{rw}\langle \mathbf{string} \rangle @ c, \mathsf{in} : \mathsf{rw}\langle \mathsf{R} \rangle$. Recall that R is the type $\mathsf{th}[\mathsf{report} :$

$\mathsf{w}\langle\mathbf{string}\rangle]$.

$$\begin{array}{l}\Gamma_e \vdash_s c : \mathsf{loc}\\ \Gamma_e \vdash_s \mathsf{report} : \mathsf{w}\langle\mathbf{string}\rangle@c\\ \Gamma_e \vdash_s \mathsf{in} : \mathsf{w}\langle\mathsf{R}\rangle@c\\ \Gamma_e \vdash_s \mathsf{news}?(x)\,\mathsf{goto}_{\mathsf{in}}\,c.\,\mathsf{report}!\langle x\rangle\\ \qquad\qquad : \mathsf{th}[\mathsf{news} : \mathsf{r}\langle\mathbf{string}\rangle@s,\ \mathsf{in} : \mathsf{w}\langle\mathsf{R}\rangle@c]\end{array}$$

The first three are simple value judgements and we have already seen a derivation of the last.

This ends our consideration of the client/server in (16) above. But let us reconsider the badServer from (15) above. Using existential types this example might be written

$$\text{badServer:}\qquad s[\![\mathsf{req}?(\xi : \mathsf{S}_e)\,\mathsf{goto}_x\,y.z!\langle \mathit{boring}\rangle]\!]$$

But one can show that this no longer typechecks. The problem arises when trying to establish

$$\Gamma, \{\xi : \mathsf{S}_e\} \vdash_y x!\langle z!\langle \mathit{boring}\rangle\rangle \tag{20}$$

We have already seen the expanded environment in (17) above, which is

$$\Gamma, \xi : \langle \mathsf{th}_y \text{ with } y : \mathsf{loc},\ z : \mathsf{w}\langle\mathbf{string}\rangle@y, x : \mathsf{w}\langle\mathsf{th}[z : \mathsf{w}\langle\mathbf{string}\rangle]\rangle@y\rangle$$

However the only way to get information from the package

$$\xi : \langle \mathsf{th}_y \text{ with } y : \ldots, z : \ldots, x : \ldots\rangle$$

in this environment is to use the rule (TY-UNPACK) from Figure 6. This will only give information on the variable ξ whereas the judgement (20) requires information on the other components of the package y, z, x which are inaccessible.

4.7 Script types

In all of the examples so far servers react to data furnished directly from clients. The general form of script types,

$$\mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi),$$

allow servers to accept *parameterised* scripts, which can be instantiated by data owned, or trusted, by the server itself. Consider the following variation on the client used in (8):

$$\begin{array}{rl}\text{Client:} & c[\![\mathsf{goto}_{\mathsf{req}}\,s.F \ \mid \mathsf{in}?(\xi : \mathsf{R})\,\mathsf{run}\,\xi \ \mid \mathsf{report}?(y)\ldots]\!]\\ F = & \lambda\, y : \mathsf{w}\langle\mathbf{string}\rangle.\ y?(x)\,\mathsf{goto}_{\mathsf{in}}\,c.\,\mathsf{report}!\langle x\rangle\end{array}$$

It does not know the source of the news at the server; so it sends the script F there, a script which uses the pre-existing port and channel $\mathsf{in}, \mathsf{report}$, but is parameterised on an information channel local to the server. The server inputs the script and is now free to apply it to whatever information source it deems fit. A simple server, with the same functionality as that in (8), is given by

$$\text{Server:}\quad s[\![\mathsf{req}?(\xi : \mathsf{S}_s)\,\xi(\mathsf{news}) \ \mid\ *\mathsf{news}!\langle \mathit{scandal}\rangle]\!]$$

It simply applies the incoming script to the local channel news. However it could also dynamically generate the local news channel, along the lines of

$$\text{ServerDy:}\quad s[\![\mathsf{req}?(\xi : \mathsf{S}_s)\,\mathsf{latest}?(z)\,(\xi\ z)]\!]$$

Note that when F is received by the server and instantiated, the type of the resulting process is dependent on that of the channel to which F is applied. Under the assumptions in place during the discussion of (8), and assuming that foo is a local channel, one would expect the process $(F\ \mathsf{foo})$, running at s, to behave in accordance with the type

$$\mathsf{pr}[\mathsf{foo} : \mathsf{r}\langle\mathbf{string}\rangle@s, \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@c]$$

This is indeed the case as F can be assigned the parameterised type

$$\mathsf{Fdep}(y : \mathsf{r}\langle\mathbf{string}\rangle \to \mathsf{pr}[y : \mathsf{r}\langle\mathbf{string}\rangle@\mathsf{here}, \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@c]) \tag{21}$$

To see this let Γ be as described on 24. Then, using a simple variation on the inference described there, we can infer

$$\Gamma, y : \mathsf{r}\langle\mathbf{string}\rangle@s \vdash_s y?(x)\,\mathsf{goto}_{\mathsf{in}}\,c.\,\mathsf{report}!\langle x\rangle \ :\ \mathsf{pr}[y : \mathsf{r}\langle\mathbf{string}\rangle@\mathsf{here}, \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@c]$$

An application of (TY-ABS) from Figure 7 gives the required

$$\Gamma \vdash_s F \ :\ \mathsf{Fdep}(y : \mathsf{r}\langle\mathbf{string}\rangle \to \mathsf{pr}[y : \mathsf{r}\langle\mathbf{string}\rangle@\mathsf{here}, \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@c])$$

Under the further assumption that $\Gamma \vdash_s \mathsf{foo} : \mathsf{r}\langle\mathbf{string}\rangle$ an application of (TY-BETA) gives

$$\Gamma \vdash_s (F\ \mathsf{foo}) \ :\ \mathsf{pr}[\mathsf{foo} : \mathsf{r}\langle\mathbf{string}\rangle@s, \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle@c]$$

Following this discussion it should be apparent that to ensure that the overall system is well-typed it is sufficient to use the dependent type (21) above for the guardian type S_s. Then it is easy to check

$$\Gamma \vdash \text{Client} \mid \text{Server}$$

For example typing the server involves establishing

$$\Gamma, \xi : \mathsf{S}_s \vdash_s (\xi\ \mathsf{news}) : \mathbf{proc} \tag{22}$$

Assuming that $\Gamma \vdash_s \mathsf{news} : \mathsf{r}\langle\mathbf{string}\rangle$, we have already seen that an application of (TY-BETA) gives

$$\Gamma, \xi : \mathsf{S}_s \vdash_s (\xi\ \mathsf{news}) : \mathsf{pr}[\mathsf{news} : \mathsf{r}\langle\mathbf{string}\rangle @ s, \mathsf{in} : \mathsf{w}\langle\mathsf{thunk}\rangle @ c]$$

and the required (22) follows by subtyping.

These parameterised functional types can be used in conjunction with the other constructions we have considered, dependent and existential types, to give a very sophisticated language for guardian types which on the one hand allows non-trivial interaction between types, and on the other enables sites to protect their local resources by implementing powerful dynamic access policies. As a final example, to indicate the potential of these types, consider the the following variation on the client in (16), which is in turn an elaboration of the example we have just considered:

$$\begin{array}{rl}
\text{Server:} & s[\![\mathsf{req}?(\xi : \mathsf{S}_{se})\,(\xi\ \mathsf{news})\ \mid\ *\mathsf{news}!\langle \mathit{scandal}\rangle]\!] \\
\text{Client:} & c[\![(\mathsf{new}\, c\ \mathsf{report}) \\
& \quad (\mathsf{new}\, c\ \mathsf{in} : \mathsf{rw}\langle \mathsf{R}\rangle) \\
& \quad\quad \mathsf{goto}_{\mathsf{req}}\, s.F\ \mid \\
& \quad\quad \mathsf{in}?(\xi : \mathsf{R})\,\mathsf{run}\,\xi\ \mid \mathsf{report}?(y)\ldots \\
F = & \quad \lambda\, y : \mathsf{w}\langle\mathbf{string}\rangle.\ y?(x)\,\mathsf{goto}_{\mathsf{in}}\, c.\,\mathsf{report}!\langle x\rangle]\!]
\end{array}$$

Here the client does not know the source of the news at the server, and at the same time the server is not aware of the reply mechanisms in place at the client; indeed these are generated dynamically by the client and used to construct the script F to be sent to the server. One can show that this system is well-typed if we let the guardian type for the client and server to be

$$\begin{array}{ll}
\mathsf{R} : & \mathsf{th}[\mathsf{report} : \mathsf{w}\langle\mathbf{string}\rangle @ c] \\
\mathsf{S}_{se} & \mathsf{Fdep}(w : \mathsf{r}\langle\mathbf{string}\rangle \to \mathsf{S}_e^w)
\end{array}$$

respectively, where S_e^w is the existential type

$$\begin{array}{l}
\mathsf{Edep}(y : \mathsf{loc}, z : \mathsf{w}\langle\mathbf{string}\rangle @ y, x : \mathsf{w}\langle\mathsf{th}[z : \mathsf{w}\langle\mathbf{string}\rangle @ y]\rangle @ y) \\
\qquad \mathsf{th}[w : \mathsf{r}\langle\mathbf{string}\rangle @ s,\ x : \mathsf{w}\langle\mathsf{th}[z : \mathsf{w}\langle\mathbf{string}\rangle @ y]\rangle @ y]
\end{array}$$

5 Subject Reduction

Many of the expected properties can be derived for our type inference system. To state these succinctly it will useful to use

$$\Gamma \vdash_w J : \mathsf{T}$$

to denote either a value judgement $\Gamma \vdash_w v : \mathsf{T}$ or a process judgement $\Gamma \vdash_w P : \mathsf{T}$. We will confine our attention to judgements in which Γ contains no occurrences of the special symbol here; thus they will only occur as part of dependent types $\mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi)$ and note that in applications of the rule (TY-ABS) from Figure 7 they are eliminated.

Proposition 5.1 (Sanity Checks)

- $\Gamma \vdash_w J : \mathsf{T}$ *implies* $\Gamma \vdash \mathbf{env}$.
- $\Gamma \vdash_w P : \pi$ *implies* $\Gamma \vdash \pi : \mathbf{ty}$

Proof: The first is proved by induction on the inference of $\Gamma \vdash_w J : \mathsf{T}$ while the second is on that of the inference of $\Gamma \vdash_w P : \pi$. It is required by the base case (TY-STOP) while in the cases (TY-OUT), (TY-OUTE), (TY-IN) and (TY-SUB) it follows from the corresponding result for subtyping, Proposition A.1. All other cases follow by induction except (TY-BETA). There we have $\Gamma \vdash F\,(\tilde{v}) : \pi\{\!|\tilde{v}/\tilde{x}|\!\}\{\!|w/\mathsf{here}|\!\}$ because

(i) $\Gamma \vdash_w v_i : \mathsf{T}_i$

(ii) $\Gamma \vdash_w F : \mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi)$

The latter can only be inferred by (TY-ABS) from which we know that $\Gamma, \{\tilde{x} : (\tilde{\mathsf{T}}) @ w\} \vdash_w P : \pi\{\!|w/\mathsf{here}|\!\}$. By the induction hypothesis we have that $\Gamma, \{\tilde{x} : (\tilde{\mathsf{T}}) @ w\} \vdash \pi\{\!|w/\mathsf{here}|\!\} : \mathbf{ty}$. It follows by the substitution result, Proposition A.5, applied to (i), that $\Gamma \vdash \pi\{\!|w/\mathsf{here}|\!\}\{\!|\tilde{v}/\tilde{x}|\!\} : \mathbf{ty}$. However since we know that w is different than each x_i this type is $\pi\{\!|\tilde{v}/\tilde{x}|\!\}\{\!|w/\mathsf{here}|\!\}$. ■

In a similar vein we can show that well-typed processes can only use well-defined types. For example if $\Gamma \vdash_w u?(X : \mathsf{T})\,P : \mathbf{proc}$ then $\Gamma \vdash \mathsf{T} : \mathbf{ty}$.

Environments can be ordered by their ability to assign types to identifiers: $\Gamma_1 <: \Gamma_2$ if for every identifier u, $\Gamma_2 \vdash_w u : \mathsf{T}$ implies $\Gamma_1 \vdash_w u : \mathsf{T}$. We will write $\Gamma_1 \equiv \Gamma_2$ whenever $\Gamma_1 <: \Gamma_2$ and $\Gamma_2 <: \Gamma_1$.

Proposition 5.2

- *(Weakening) Suppose* $\Gamma_2 \vdash_w J : \mathsf{T}$ *and* $\Gamma_1 <: \Gamma_2$ *for some* Γ_1 *such that* $\Gamma_1 \vdash \mathbf{env}$. *Then* $\Gamma_1 \vdash_w J : \mathsf{T}$.
- *(Strengthening) Suppose* $\Gamma, u : \mathsf{T} \vdash_w J : \mathbf{proc}$, *where* u *does not occur free in* J. *Then* $\Gamma \vdash_w J : \mathbf{proc}$.
- *(Subtyping) Suppose* $\Gamma \vdash_w J : \mathsf{T}$. *Then* $\Gamma \vdash \mathsf{T} <: \mathsf{T}'$ *implies* $\Gamma \vdash_w J : \mathsf{T}'$

Proof: The first two statements are proved by induction on the inferences. The third follows immediately from (TY-SUBPROC) in Figure 7 and (TY-SUBVAL) in Figure 6. ■

Multiple occurrences of an identifier is governed by the following result:

Proposition 5.3 $\Gamma \vdash_w u : \mathsf{C}_1@w_1$ *and* $\Gamma \vdash_w u : \mathsf{C}_2@w_2$ *implies* $\Gamma \vdash_w u : \mathsf{rc}\langle\mathsf{D}\rangle$ *for some* D *such that* $\Gamma \vdash \mathsf{D} <: \mathsf{C}_1,\ \mathsf{D} <: \mathsf{C}_2$.

Proof: This property is essentially enforced by the formation rules for well-defined environments. These ensure that if $\Gamma_1, u : \mathsf{C}_1@w_1, \ldots, u : \mathsf{C}_2@w_2, \ldots$ is a well-defined environment then Γ_1 must contain an entry $u : \mathsf{rc}\langle\mathsf{D}\rangle$, where $\Gamma_1 \vdash \mathsf{D} <: \mathsf{C}_1$ and $\Gamma_1, u : \mathsf{C}_1@w_1, \ldots \vdash \mathsf{D} <: \mathsf{C}_2$.

The formal proof is by induction on the inferences of $\Gamma \vdash_w u : \mathsf{C}_1@w_1$ and $\Gamma \vdash_w u : \mathsf{C}_2@w_2$. The base case, when both are inferred from (TY-LOOKUP), depends on this property of well-defined environments. ■

An interesting consequence of this result is that whenever the conditions of the proposition hold $\mathsf{C}_1 \sqcap \mathsf{C}_2$ is guaranteed to exist. This is spelled out in detail in Proposition A.2 in the Appendix.

As usual the proof of Subject Reduction relies on the fact that, in a suitable sense, type inference is preserved under substitutions. We require two such results, one for standard values, and one for the existential values used in type inference.

Lemma 5.4 (Substitution) *Suppose* $\Gamma \vdash_{w_1} v : \mathsf{T}$ *with* x *not in* Γ. *Then* $\Gamma, x : (\mathsf{T})@w_1, \Delta \vdash_{w_2} J : \mathsf{T}$ *implies* $\Gamma, \Delta\{\!|v/x|\!\} \vdash_{w_2\{v/x\}} J\{\!|v/x|\!\} : \mathsf{T}\{\!|v/x|\!\}$

Proof: First note that the entry $x : (\mathsf{T})@w_1$ can only take one of the four forms, a channel registration, $x : \mathsf{rc}\langle\mathsf{D}\rangle$, a location declaration $x : \mathsf{loc}$, a channel declaration, $x : \mathsf{C}@u'$ or a script declaration $x : \mathsf{S}$. The proof is by induction on the inference of $\Gamma, x : (\mathsf{T})@w_1, \Delta \vdash_{w_2} J : \mathsf{T}$, which can use the rules from Figure 6 or Figure 7. For convenience we use α' to denote $\alpha\{\!|v/x|\!\}$ for the various syntactic categories α. Also we use Γ_e as an abbreviation for the environment $\Gamma, x : (\mathsf{T})@w_1, \Delta$. First let us look at some cases from Figure 6.

- Suppose (TY-LOOKUP) is used. So $\Gamma_e \vdash_{w_2} u : \mathsf{E}$ because

  (i) $\Gamma_e \vdash$ **env**

  (ii) Γ_e has the form $\Gamma_1, u : (\mathsf{E})@w_2, \ldots$.

  The substitution result for well-defined environments, Proposition A.5 in the appendix, ensures that

  (i') $\Gamma, \Delta' \vdash$ **env**

  To obtain the corresponding

  (ii') Γ, Δ' has the form $\Delta_1, u' : (\mathsf{E}')@w_2', \ldots$

  we perform a case analysis on where $u : (\mathsf{E})@w_2$ occurs in Γ_e; with (i') and (ii') an application of the rule (TY-LOOKUP) gives the required $\Gamma \vdash_{w_2'} u' : \mathsf{E}'$.

  If it occurs in Γ then (ii') is immediate since the substitutions have no effect. If it occurs in Δ then $u' : (\mathsf{E}')@w_2'$ occurs in Δ' and so (ii') holds. Finally $u : (\mathsf{E})@w_2$ could coincide with $x : (\mathsf{T})@w_1$. There are now a number of cases, depending on the form of $(\mathsf{T})@w_1$. As an example suppose it is $\mathsf{C}@w_1$. Then w_1 and w_2 coincide and x can not appear in C, w_1. Therefore the hypothesis $\Gamma \vdash_{w_1} v : \mathsf{C}$ gives the required result, $\Gamma, \Delta' \vdash_{w_2} v : \mathsf{C}$, by Weakening.

- The case (TY-ELOOKUP) is very similar, although there are only two rather than three possibilities for the occurrence of the association in Γ_e.

- Suppose (TY-LOC) is used. So $\Gamma_e \vdash_{w_2} w : \mathsf{K}$, where K is the type $\mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n]$ because

  (i) $\Gamma_e \vdash_w u_i : \mathsf{C}_i$

  (ii) $\Gamma_e \vdash_{w_2} u_i : \mathsf{rc}\langle\mathsf{D}_i\rangle$

  (iii) $\Gamma_e \vdash \mathsf{D}_i <: \mathsf{C}_i$

  Induction, and the substitution result for subtyping, Proposition A.5 in the Appendix, can be applied to these to obtain

  (i') $\Gamma, \Delta' \vdash_{w'} u_i' : \mathsf{C}_i'$

  (ii') $\Gamma, \Delta' \vdash_{w_2'} u_i' : \mathsf{rc}\langle\mathsf{D}_i'\rangle$

  (iii') $\Gamma, \Delta' \vdash \mathsf{D}_i' <: \mathsf{C}_i'$

  The interesting case is when both v and x occur in $u_1, \ldots u_n$; without loss of generality suppose these are u_1, u_2 respectively, in which case $u_1' = u_2' = v$. Then we know, by Proposition 5.3, that $\mathsf{C}_1' \sqcap \mathsf{C}_2'$ exists and K' is $\mathsf{loc}[u_2' : (\mathsf{C}_1' \sqcap \mathsf{C}_2'), \ldots]$. Applying the rule (TY-MEET) to (i') above gives $\Gamma, \Delta' \vdash_{w'} u_2' : (\mathsf{C}_1' \sqcap \mathsf{C}_2')$ and therefore we can apply (TY-LOC) to this, together with (i'), (ii') and (iii') to obtain the required $\Gamma, \Delta' \vdash_{w_2'} w' : \mathsf{K}'$.

The other cases from Figure 6 are similar, mostly following by induction. Now let us look at some cases from Figure 7.

- Suppose (TY-NEWLOC) is used so $\Gamma_e \vdash_{w_2} (\mathsf{newloc}\, k : \mathsf{K})\, \mathsf{with}\, C\ \mathsf{in}\ P : \pi$ because

  (i) $\Gamma_e, \{k : \mathsf{K}\} \vdash_k C : \pi$

  (ii) $\Gamma_e, \{k : \mathsf{K}\} \vdash_{w_2} P : \pi$

(iii) $\Gamma_e, \{k : \mathsf{K}\} \vdash_{w_2} k : \mathsf{K}$

Induction can be applied to each of these, to obtain

(i') $\Gamma, \Delta', (\{k : \mathsf{K}\})' \vdash_k C' : \pi'$
(ii') $\Gamma, \Delta', (\{k : \mathsf{K}\})' \vdash_{w_2'} P' : \pi'$
(iii') $\Gamma, \Delta', (\{k : \mathsf{K}\})' \vdash_{w_2'} k : \mathsf{K}'$

Unfortunately it is not true in general that $\Gamma, \Delta', (\{k : \mathsf{K}\})'$ is the same as $\Gamma, \Delta', (\{k : \mathsf{K}'\})$. For example if K is $\mathsf{loc}[x : \mathsf{C}_1', v : \mathsf{C}_2', \ldots]$ then the former contains the entries $\ldots k : \mathsf{loc}, v : \mathsf{C}_1' @ k, v : \mathsf{C}_2' @ k, \ldots$ whereas the latter contains $\ldots k : \mathsf{loc}, v : (\mathsf{C}_1' \sqcap \mathsf{C}_2') @ k, \ldots$. Nevertheless it will always be the case that

$$\Gamma, \Delta', (\{k : \mathsf{K}\})' \equiv \Gamma, \Delta', (\{k : \mathsf{K}'\})$$

and therefore by Weakening (i'),(ii') and (iii') apply also to the latter. So (TY-LOC) can be applied to these to obtain the required

$$\Gamma, \Delta' \vdash_{w_2'} (\mathsf{new}\mathsf{loc}\, k : \mathsf{K}')\, \mathsf{with}\, C' \ \mathsf{in}\ P' : \pi'$$

- Suppose (TY-IN) is used. So $\Gamma \vdash_{w_2} u?(X : \mathsf{U})\, P : \pi$ because

  (i) $\Gamma_e \vdash \mathsf{pr}[u : \mathsf{r}\langle \mathsf{U} \rangle @ w_2] <: \pi$
  (ii) $\Gamma_e, \{X : (\mathsf{U}) @ w_2\} \vdash_{w_2} P : (\pi \sqcup \mathsf{pr}_{\mathsf{ch}}[X : (\mathsf{U}) @ w_2])$

  Applying the substitution result for subtyping, Proposition A.5, we get

  (i') $\Gamma, \Delta' \vdash \mathsf{pr}[u' : \mathsf{r}\langle \mathsf{U}' \rangle @ w_2'] <: \pi'$

  since $(\mathsf{pr}[u : \mathsf{r}\langle \mathsf{U} \rangle @ w_2])'$ is $\mathsf{pr}[u' : \mathsf{r}\langle \mathsf{U}' \rangle @ w_2']$. Applying induction to (ii) gives

  (ii') $\Gamma, \Delta', (\{X : (\mathsf{U}) @ w_2\})' \vdash_{w_2'} P' : (\pi \sqcup \mathsf{pr}_{\mathsf{ch}}[X : (\mathsf{U}) @ w_2])'$

  Now substitutions distribute over $\sqcup$ (see Proposition A.3 in the Appendix), and also over the channel extraction function (See Proposition A.4). So this may be rewritten

  (ii') $\Gamma, \Delta', (\{X : (\mathsf{U}) @ w_2\})' \vdash_{w_2'} P' : (\pi' \sqcup \mathsf{pr}_{\mathsf{ch}}[X : (\mathsf{U}') @ w_2'])$

  as x is guaranteed not to be in the pattern X. As in the previous case, we can show that

  $$\Gamma, \Delta', (\{X : (\mathsf{U}) @ w_2\})' \equiv \Gamma, \Delta', \{X : (\mathsf{U}') @ w_2'\}$$

  although because of location types they may not be identical. Nevertheless this is sufficient to be able to apply (TY-IN) to (i'),(ii') to obtain the required $\Gamma, \Delta' \vdash_{w_2'} u?(X : \mathsf{U}')\, P' : \pi$ ■

This substitution result can be generalised to arbitrary patterns, but we only require it in a special case:

Corollary 5.5 *Let X be a pattern and suppose $\Gamma \vdash_{w_1} V : \mathsf{T}$ where T is not an existential type. Then $\Gamma, \{X : (\mathsf{T}) @ w_1\} \vdash_{w_2} J : \mathsf{T}$ implies $\Gamma \vdash_{w_2\{V/X\}} J\{V/X\} : \mathsf{T}\{V/X\}$*

Proof: By induction on the structure of T. The base cases are covered by the previous lemma. There are two other cases, when T is a location type and when it is a dependent type. As an example we consider the former, when it has the form $\mathsf{K} = \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots u_n : \mathsf{C}_n]$; in this case X must be a variable x and V and identifier, say v.

So $\Gamma, \{X : (\mathsf{K}) @ w\}$ is $\Gamma, x : \mathsf{loc}, u_1 : \mathsf{C}_1 @ x, \ldots, u_n : \mathsf{C}_n @ x$ which can be written as

$$\Gamma,\ x : \mathsf{loc},\ (u_1 : \mathsf{C}_1 @ x, \ldots, u_n : \mathsf{C}_n @ x)$$

So applying the previous lemma we obtain

$$\Gamma, u_1 : \mathsf{C}_1 @ v, \ldots u_n : \mathsf{C}_n @ v \vdash_{w_2\{v/x\}} J\{v/x\} : \mathsf{T}\{v/x\}$$

But $\Gamma \vdash_{w_2} v : \mathsf{K}$ means that $\Gamma \vdash_v u_i : \mathsf{C}_i$ for each i. So we see that $\Gamma \equiv \Gamma, u_1 : \mathsf{C}_1 @ v, \ldots u_n : \mathsf{C}_n @ v$ from which the required

$$\Gamma \vdash_{w_2\{v/x\}} J\{v/x\} : \mathsf{T}\{v/x\}$$

follows. ■

The corresponding result for existential types uses different substitutions into processes and types. The crucial property of existential values is that the use of their *witnesses* is very limited:

Proposition 5.6 *Suppose $\Gamma, y : \langle \mathsf{T}\ \mathsf{with}\ \tilde{x} : \tilde{\mathsf{E}} \rangle, \Gamma' \vdash_w J : \mathsf{T}$. Then $x_i \notin \mathsf{fv}(J)$ and x_i does not occur in Γ', w.*

Proof: By induction on the inference. Intuitively the result follows from the fact that the only information available, via (TY-ELOOKUP), from the entry $y : \langle \mathsf{T}\ \mathsf{with}\ \tilde{x} : \tilde{\mathsf{E}} \rangle$ is that y has the type T; no information on x_i is available. The proof relies on the corresponding result for well-defined environments and subtyping, Proposition A.6 ■

This result provides the central property underlying the substitution result for existential values.

Lemma 5.7 (ESubstitution) *Suppose $\Gamma \vdash_{w_1} \langle \tilde{v}, v \rangle : \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\, \mathsf{T}$. Then $\Gamma, y : \langle (\mathsf{T}) @ w_1\ \mathsf{with}\ \tilde{x} : \tilde{\mathsf{E}} \rangle, \Delta \vdash_{w_2} J : \mathsf{T},\ w_2 : \mathsf{loc}$ implies $\Gamma, \Delta\{\tilde{v}/\tilde{y}\} \vdash_{w_2\{\tilde{v}/\tilde{y}\}} J\{\tilde{v}/\tilde{y}\} : \mathsf{T}\{\tilde{v}/\tilde{x}\}$*

Proof: The proof follows the lines of that of Lemma 5.4, with frequent applications of the previous proposition, Proposition 5.6, to ensure that only the substitution of v for x is applied to process terms and names. As usual certain cases depend on the corresponding result for well-typed environments and subtyping judgements, Proposition A.7 in the Appendix. ■

Theorem 5.8 (Subject Reduction)

Suppose $\Gamma \vdash M$. *Then* $M \longrightarrow N$ *implies* $\Gamma \vdash N$.

Proof: It is a question of examining each of the rules in Figure 2 in turn. Note that (R-STR) requires that typing is preserved by the structural equivalence; we leave the proof of this fact to the reader, as it follows the standard approach.

Consider the rule (R-COMM):

$$k[\![c!\langle V\rangle]\!] \mid k[\![c?(X:\mathsf{T})\,P]\!] \longrightarrow k[\![P\{^{V}\!/_{X}\}]\!]$$

and suppose $\Gamma \vdash k[\![c!\langle V\rangle]\!] \mid k[\![c?(X:\mathsf{T})\,P]\!]$. Because **proc** is a top type for processes this means that

(i) $\Gamma \vdash_k c!\langle V\rangle$: **proc**

(ii) $\Gamma \vdash_k k[\![c?(X:\mathsf{T})\,P]\!]$: **proc**

We need to show $\Gamma \vdash k[\![P\{^{V}\!/_{X}\}]\!]$ which follows easily if we can establish $\Gamma \vdash_k P\{^{V}\!/_{X}\}$: **proc**.

From (i),(ii), we can show that $\Gamma \vdash_k c : \mathsf{rw}\langle \mathsf{T}, \mathsf{T}\rangle$ and $\Gamma \vdash_k V : \mathsf{T}$. There are now two cases, depending on the structure of T. First suppose it is an existential type $\mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{U}$, in which case the pattern X is a single variable, say y. Here (i) above can only be inferred by using (TY-OUTE), which means that V is a singleton, say v and there must be some vector $\tilde{v}$ of witnesses such that $\Gamma \vdash_k \langle \tilde{v}, v\rangle$: $\mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\,\mathsf{U}$. Deconstructing (ii) we know that $\Gamma, y : \langle \mathsf{U} \text{ with } \tilde{x} : \tilde{\mathsf{E}}\rangle \vdash_k P$: **proc**. We may now apply Lemma 5.7 to obtain the required $\Gamma \vdash_k P\{^{v}\!/_{y}\}$.

When T is not an existential type the proof is similar but uses an application of Corollary 5.5 in place of Lemma 5.7.

We leave the proof for the other rules to the reader. ■

6 The behaviour of SAFEDPI systems

In this section we investigate what might be an appropriate notion of semantic equivalence between SAFEDPI systems. We first propose what we believe to be a natural notion of contextual equivalence. Then, in the following

sections, we give a coinductive characterisation using actions between configurations, consisting of SAFEDPI systems together with the environment's current knowledge of the system.

For notational convenience we limit ourselves to the case when the only transmission types allowed are of the form

$$\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{A}})\,\mathsf{A} \qquad \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{A}})\,\mathsf{S} \qquad \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{A}})\,\mathsf{S}$$

Effectively this means that the values transmitted must either be of the form

- $(\tilde{u})$, a tuple of first-order values, of type $\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{A}})\,\mathsf{A}$
- $(\tilde{u}, F)$ a tuple in which the last value F, a script, may depend on the first-order values $(\tilde{u})$. These have a type of the form $\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{A}})\,\mathsf{S}$.
- F a script, the final component of an existential value $\langle \tilde{u}, F\rangle$ with a type of the form $\mathsf{Edep}(\tilde{x} : \tilde{\mathsf{A}})\,\mathsf{S}$.

Simple scripts may be simulated via the empty dependent type $\mathsf{Tdep}()\,\mathsf{S}$, as can simple first-order values, via the type $\mathsf{Tdep}()\,\mathsf{A}$. Our results extend to the full language, although the proofs require the development of more complicated notations.

6.1 A contextual equivalence

We intend to use a context based equivalence in which systems are asked to be deemed equivalent in all *reasonable* SAFEDPI contexts. What is perhaps not so clear here is the notion of reasonable context. In previous work on mobile calculi, [9, 8, 1], the equivalence took the form

$$\Gamma \models M \approx_{cxt} N$$

meaning, intuitively, that M and N are indistinguishable in any context typeable by the typing environment Γ. Although one is primarily interested in such judgements in which Γ has sufficient knowledge to type M and N, one is lead to consider more general judgements where Γ only contains a subset of that knowledge. Such equivalences, for PICALCULUS and DPI, can be characterised inductively using actions of the form

$$(\Gamma \rhd M) \xrightarrow{\mu} (\Gamma' \rhd M')$$

where $(\Gamma \rhd M)$, $(\Gamma \rhd M')$ are configurations, consisting of systems M, M' and type environments Γ, Γ', representing the current knowlege of the testing context. In general such actions change not only the systems, M to M' but also the current knowledge, from Γ to Γ', typically by adding new information.

However, there are further subtleties which need to be considered in the current setting. We discuss this with a motivating example.

Example 6.1 Consider

$$M = (\mathsf{new}\, k : \mathsf{loc}[b : \mathsf{rw}\langle \mathbf{unit}\rangle])\, l[\![a!\langle k\rangle]\!] \mid k[\![b!\langle\rangle]\!]$$
$$N = (\mathsf{new}\, k : \mathsf{loc}[b : \mathsf{rw}\langle \mathbf{unit}\rangle])\, l[\![a!\langle k\rangle]\!] \mid k[\![\mathsf{stop}]\!]$$
and
$$\Gamma = l : \mathsf{loc}, b : \mathsf{rc}\langle \mathsf{rw}\langle \mathbf{unit}\rangle\rangle, a : \mathsf{rw}\langle \mathsf{loc}[b : \mathsf{rw}\langle \mathbf{unit}\rangle]\rangle @ l$$

These two systems are well-typed with respect to Γ and should be considered equivalent under most reasonable notions of behavioural equivalence; it is impossible for a testing process to interact with M on b at k, even after the interaction on a at l. Indeed, consider what form a test which could achieve this must take:

$$- \mid l[\![a?(x)\, \mathsf{goto}_?\, x.b?()]\!]$$

It is clear that there is no port for the testing process to enter the location k on. Moreover, tests cannot be placed directly at k as k is only discovered through interaction.

To sum up we would expect

$$\Gamma \models M \approx_{cxt} N$$

to hold, for an appropriate formulation of contextual equivalence for SAFEDPI. But a naive labelled transition system of the form discussed above would distinguish them. For example a naive system might yield actions such as

$$(\Gamma \rhd M) \xrightarrow{\text{outputs } k \text{ on } a \text{ at } l} (\Gamma' \rhd l[\![\mathsf{stop}]\!] \mid k[\![b!\langle\rangle]\!])$$

where Γ' is the environment Γ updated with the knowledge about the new location $k : \mathsf{loc}[b : \mathsf{rw}\langle \mathbf{unit}\rangle]$. However, in such a system, a subsequent interaction at this newly discovered k would be possible. This interaction would suffice to distinguish M and N.

In other words we need to consider more sophisticated notions of actions in order to capture contextual equivalences for SAFEDPI. ■

It should be clear from this discussion then that in modelling behavioural equivalence in this setting, we must be aware of those locations at which we can, and can not, perform tests. And this is not simply a question of which locations the environment has immigration rights for, via some port.

Example 6.2 Consider the following scenario:

$$M = k[\![(\mathsf{newc}\, b : \mathsf{rw}\langle \mathbf{unit}\rangle)\ a!\langle b\rangle \mid b!\langle\rangle]\!]$$
$$N = k[\![(\mathsf{newc}\, b : \mathsf{rw}\langle \mathbf{unit}\rangle)\ a!\langle b\rangle \mid \mathsf{stop}]\!]$$
and
$$\Gamma = k : \mathsf{loc}, a : \mathsf{rw}\langle \mathsf{rw}\langle \mathbf{unit}\rangle\rangle @ k$$

Here the testing environment already knows about k but does not have any immigration rights there. Nevertheless M and N can be distinguished by a reasonable test, one which is typeable by Γ:

$$- \mid k[\![a?(x)\, x?\langle\rangle\, \mathsf{eureka}!\langle\rangle]\!]$$

■

Thus, in representing the environment's knowledge of the system we must also represent the information about which locations are available for direct testing. This motivates the following definition.

Definition 6.3 (Knowledge structures) A *knowledge structure* is a pair $(\Gamma, \mathcal{T})$, where

- Γ is a type environment such that $\Gamma \vdash \mathbf{env}$
- $\mathcal{T}$ is a subset of Locs such that if $k \in \mathcal{T}$ then $k : \mathsf{loc} \in \Gamma$

We use $\mathcal{I}$ to range over knowledge structures and write $\mathcal{I}_\Gamma$ and $\mathcal{I}_\mathcal{T}$ to refer to the respective components of the structure. We sometimes refer to the locations in $\mathcal{I}_\mathcal{T}$ as those to which the information structure allows *access rights*. We often abuse notation by writing $\mathcal{I}, \Gamma$ to mean the knowledge structure $((\mathcal{I}_\Gamma, \Gamma), \mathcal{I}_\mathcal{T})$. ■

Definition 6.4 (Configurations) We write $\mathcal{I} \rhd M$ for a *configuration* where

- $\mathcal{I}$ is a knowledge structure
- there exists some Δ such that $\Delta \vdash M$, $\Delta <: \mathcal{I}_\Gamma$, and $\mathsf{dom}(\Delta) = \mathsf{dom}(\mathcal{I}_\Gamma)$. ■

Definition 6.5 (Knowledge-indexed relations) We call a family of binary relations between systems indexed by knowledge structures a *knowledge-indexed relation* over systems. We write $\mathcal{I} \models M\, \mathcal{R}\, N$ to mean that systems M and N are related by $\mathcal{R}$ at index $\mathcal{I}$ and moreover, $\mathcal{I} \rhd M$ and $\mathcal{I} \rhd N$ are both configurations. ■

We will use knowledge-indexed relations to propose a notion of behavioural equivalence appropriate to this setting. We do this in an established manner [11, 6, 9] by proposing that we consider the largest equivalence closed under certain natural properties listed below.

Reduction closure: We say that a knowledge-indexed relation is *reduction closed* if whenever $\mathcal{I} \models M\, \mathcal{R}\, N$ and $M \longrightarrow M'$ there exists some N' such that $N \longrightarrow^* N'$ and $\mathcal{I} \models M'\, \mathcal{R}\, N'$.

Context closure: We say that a knowledge-indexed relation is *contextual* if

(1) $\mathcal{I} \models M \mathcal{R} N$ and $\mathcal{I}_\Gamma, k : \mathsf{loc} \vdash \mathbf{env}$ implies $\mathcal{I}' \models M \mathcal{R} N$ where $\mathcal{I}'$ is $((\mathcal{I}_\Gamma, k : \mathsf{loc}), \mathcal{I}_\mathcal{T} + k)$

(2) $\mathcal{I} \models M \mathcal{R} N$ and $\mathcal{I}_\Gamma, \Gamma' \vdash \mathbf{env}$ implies $\mathcal{I}, \Gamma' \models M \mathcal{R} N$

(3) $\mathcal{I} \models M \mathcal{R} N$ and $\mathcal{I}_\Gamma \vdash k[\![P]\!]$ with $k \in \mathcal{I}_\mathcal{T}$ implies $\mathcal{I} \models (M \mid k[\![P]\!]) \mathcal{R} (N \mid k[\![P]\!])$

(4) $\mathcal{I}, \{n : E\} \models M \mathcal{R} N$ implies $\mathcal{I} \models (\mathsf{new}\, n : \mathsf{E})\, M \mathcal{R} (\mathsf{new}\, n : \mathsf{E})\, N$ ■

In the first condition we are assured that k is a fresh location; therefore this form of weakening allows the environment to create for itself fresh locations at which it may deploy code. The second form of weakening, in (2), allows it to invent new names with which to program processes. Condition (3) allows it to place well-typed code at sites to which it has access rights, while (4) is the standard mechanism for handling names which are private to the systems being investigated.

Barb Preservation: For any given location k and any given channel a such that $k \in \mathcal{I}_\mathcal{T}$ and $\mathcal{I}_\Gamma \vdash_k a : \mathsf{rw}\langle \mathbf{unit} \rangle$ we write $\mathcal{I} \vdash M \Downarrow^{\mathsf{barb}} a@k$ if there exists some M' such that $M \longrightarrow^* M' \mid k[\![a!\langle\rangle]\!]$. We say that a knowledge-indexed relation is *barb preserving* if $\mathcal{I} \models M \mathcal{R} N$ and $\mathcal{I} \vdash M \Downarrow^{\mathsf{barb}} a@k$ implies $\mathcal{I} \vdash N \Downarrow^{\mathsf{barb}} a@k$.

Definition 6.6 (Reduction barbed congruence) We let $\approx_{cxt}$ be the largest knowledge-indexed relation over systems which is

- pointwise symmetric (that is $\mathcal{I} \models M \approx_{cxt} N$ implies $\mathcal{I} \models N \approx_{cxt} M$)
- reduction closed
- contextual
- barb preserving ■

We take reduction barbed congruence to be our touchstone equivalence for SAFEDPI as it is based on simple observable behaviour respected in all contexts. The definition above is stated relative to choice of the knowledge structure $\mathcal{I}$. We should point out however that, for any given systems M, N and type environment Γ such that $\Gamma \vdash M$ and $\Gamma \vdash N$ then there is a canonical choice of knowledge structure $\mathcal{I}$, namely, $(\Gamma, \mathcal{T}_\Gamma)$ where we let $\mathcal{T}_\Gamma = \{ k \mid k : \mathsf{loc} \in \Gamma \}$. This choice of knowledge structure gives rise to what we feel to be a natural and intuitive notion of equivalence for well-typed SAFEDPI systems.

Of course, the quantification over all contexts makes reasoning about the equivalence virtually intractable. However it is common practice, [19, 21, 1, 9, 8], to provide some sort of model or alternative characterisation in terms of labelled transition systems, which makes the behaviour of systems much more accessible. In particular if the actions in the labelled transition system are sufficiently simple this can lead to automatic, or semi-automatic verification methods.

In the next section we show that this contextual equivalence for SAFEDPI can be characterised in a similar manner, as a bisimulation equivalence over a suitably defined labelled transition system.

6.2 A bisimulation equivalence

We first discuss the labels, or *actions*, to be used in the labelled transition system. They are given by the following grammar:

$$\alpha ::= \tau \mid (\tilde{n} : \tilde{\mathsf{E}})\mathsf{go}_p k.F \mid (\tilde{n} : \tilde{\mathsf{E}})(\tilde{m})k.a.\beta$$
$$\beta ::= V? \mid V!$$

where it is assumed that $k, a, p \notin \tilde{n}, \tilde{m}$. These are intended to be read as follows:

- τ represents *internal* communication in which no interaction with the environment takes place
- $\mathsf{go}_p k.F$ represents an attempt by the environment to enter location k on port p. The code to be deployed, if this attempt succeeds, is given by the script F.
- $k.a.V!$ represents a communication between the system and the environment in which the system exports on channel a at k. The value V in this action depends on the type of the channel. First order values can be recognised by the environment and so they are recorded in the action label. Scripts, on the other hand, can not necessarily be identified. So instead the environment provides a suitable receiving context for a script. For example suppose the system exports some script F on a channel a of script type S. To test F the environment can supply any abstraction G of type $G : S \rightarrow \mathbf{proc}$, with which F can be investigated; see rule (M-SEND − SCRIPT) in Figure 8.
- $k.a.V?$ represents a communication between system and environment in which the system imports on channel a at k. The value V is always provided by the environment.
- $(n)\alpha$ represents an action α in which the new name n has been exported from the system; it is new in the sense that it has not previously been

encountered by the testing environment. The type of n is not recorded since it can be inferred from the type of the channel on which it is exported.

- $(n : E)\alpha$ represents an action α in which the fresh name n is being provided by the environment.

The following notation is useful in defining the labelled transition system. Firstly, the *subject labels*, $\mathsf{subj}(\alpha)$ of an action are given by:

$$\begin{array}{l}\mathsf{subj}(\tau) = \emptyset\\ \mathsf{subj}((\tilde{n} : \tilde{\mathsf{E}})(\tilde{m})k.a.\beta) = \{k, a\}\\ \mathsf{subj}((\tilde{n} : \tilde{\mathsf{E}})\mathsf{go}_p k.V) = \{k, p\}\end{array}$$

Next, we define the *object labels* of an action. These are divided into both input and output object labels using the two functions $\mathsf{obj}_?(\alpha)$ and $\mathsf{obj}_!(\alpha)$ in order to identify whether the names returned are being provided by the environment or exported from the system. We use input object labels to identify the former and output object labels the latter.

$$\begin{array}{ll}\mathsf{obj}_?(\tau) = \emptyset & \mathsf{obj}_!(\tau) = \emptyset\\ \mathsf{obj}_?(\tilde{u}!) = \emptyset & \mathsf{obj}_!(\tilde{u}!) = \mathsf{fn}(\tilde{u})\\ \mathsf{obj}_?(\tilde{V}?) = \mathsf{fn}(V) & \mathsf{obj}_!(\tilde{V}?) = \emptyset\\ \mathsf{obj}_?((\tilde{u}, G)!) = \mathsf{fn}(G) & \mathsf{obj}_!((\tilde{u}, G)!) = \mathsf{fn}(\tilde{u})\end{array}$$

$$\mathsf{obj}_?((\tilde{n} : \tilde{\mathsf{E}})\mathsf{go}_p k.V) = \mathsf{fn}(V) \setminus \tilde{n} \qquad \mathsf{obj}_!((\tilde{n} : \tilde{\mathsf{E}})\mathsf{go}_p k.V) = \emptyset$$

$$\mathsf{obj}_?((\tilde{n} : \tilde{\mathsf{E}})(\tilde{m})k.a.\beta) = \mathsf{obj}_?(\beta) \setminus \tilde{n} \quad \mathsf{obj}_!((\tilde{n} : \tilde{\mathsf{E}})(\tilde{m})k.a.\beta) = \mathsf{obj}_!(\beta) \setminus \tilde{m}$$

The interesting case here is $\beta = (\tilde{u}, G)!$, which represents the export from the system to the environment a higher-order script, dependent on the first-order values $(\tilde{u})$. This exported script is not represented in the label; instead G, which is supplied by the environment, is applied to it. So $\mathsf{obj}_?(\beta)$ is all the free names in G, since these are supplied by the environment, while $\mathsf{obj}_!(\beta)$ are all the idenfiers in $\tilde{u}$, since these are supplied by the system.

With this notation we define judgements of the form

$$(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M') \tag{23}$$

representing the effect of the system M performing the action labelled α, in an environment whose knowlege is $\mathcal{I}$. This action changes the system, from M to M', and the knowledge, from $\mathcal{I}$ to $\mathcal{I}'$. Typically this is an *increase* in knowledge of the testing environment of the system, represented as the change from the type environment, $\mathcal{I}_\Gamma$ to $\mathcal{I}'_\Gamma$.

The axioms for the judgements (23) are given in Figure 8; these are based on the rules in Figure 10 of [8]. We make use of the following notation

(M-RECEIVE)
$$\frac{k \in \mathcal{I}_\mathcal{T} \quad \mathsf{T} = \bigsqcap \mathcal{I}^w_\Gamma(a,k) \quad \mathcal{I}^w_\Gamma(a,k) \neq \emptyset \quad \mathcal{I}_\Gamma \vdash_k V : \mathsf{T}}{(\mathcal{I} \rhd k[\![a?(X : \mathsf{U})\, P]\!]) \xrightarrow{k.a.V?} (\mathcal{I} \rhd k[\![P\{\!|V/X|\!\}]\!])}$$

(M-DELIVER)
$$\frac{k \in \mathcal{I}_\mathcal{T} \quad \mathsf{T} = \bigsqcap \mathcal{I}^w_\Gamma(a,k) \quad \mathcal{I}^w_\Gamma(a,k) \neq \emptyset \quad \mathcal{I}_\Gamma \vdash_k V : \mathsf{T}}{(\mathcal{I} \rhd M) \xrightarrow{k.a.V?} (\mathcal{I} \rhd M \mid k[\![a!\langle V\rangle]\!])}$$

(M-SEND.VAL)
$$\frac{k \in \mathcal{I}_\mathcal{T} \quad \mathsf{T} \text{ a first-order type} \quad \mathsf{T} = \bigsqcap \mathcal{I}^r_\Gamma(a,k) \quad \mathcal{I}^r_\Gamma(a,k) \neq \emptyset \quad \mathcal{I}_\Gamma, \{\tilde{u} : (\mathsf{T})@k\} \vdash \mathbf{env}}{(\mathcal{I} \rhd k[\![a!\langle\tilde{u}\rangle]\!]) \xrightarrow{k.a.\tilde{u}!} (\mathcal{I}, \{\tilde{u} : (\mathsf{T})@k\} \rhd k[\![\mathsf{stop}]\!])}$$

(M-SEND.SCRIPT)
$$\frac{k \in \mathcal{I}_\mathcal{T} \quad \mathsf{T} \text{ of the form } \mathsf{Edep}(\tilde{x} : \tilde{T})\, S \quad \mathsf{T} = \bigsqcap \mathcal{I}^r_\Gamma(a,k) \quad \mathcal{I}^r_\Gamma(a,k) \neq \emptyset \quad \mathcal{I}_\Gamma \vdash_k G : \mathsf{T} \to \mathbf{proc}}{(\mathcal{I} \rhd k[\![a!\langle F\rangle]\!]) \xrightarrow{k.a.G!} (\mathcal{I} \rhd k[\![G\ (F)]\!])}$$

(M-SEND.DEP.SCRIPT)
$$\frac{k \in \mathcal{I}_\mathcal{T} \quad \mathsf{T} \text{ of the form } \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{E}})\, \mathsf{S} \quad \mathsf{T} = \bigsqcap \mathcal{I}^r_\Gamma(a,k) \quad \mathcal{I}^r_\Gamma(a,k) \neq \emptyset \quad \mathcal{I}_\Gamma, \{\tilde{u} : (\tilde{\mathsf{E}})@k\} \vdash \mathbf{env} \quad \mathcal{I}_\Gamma \vdash_k G : T \to \mathbf{proc}}{(\mathcal{I} \rhd k[\![a!\langle(\tilde{u}, F)\rangle]\!]) \xrightarrow{k.a.(\tilde{u},G)!} (\mathcal{I}, \{\tilde{u} : (\tilde{\mathsf{E}})@k\} \rhd k[\![G\ (\tilde{u}, F)]\!])}$$

(M-GOTO)
$$\frac{k \notin \mathcal{I}_\mathcal{T} \quad \mathcal{I}_\Gamma \vdash_k p!\langle V\rangle : \mathbf{proc}}{(\mathcal{I} \rhd M) \xrightarrow{\mathsf{go}_p k.V} (\mathcal{I} \rhd M \mid k[\![p!\langle V\rangle]\!])}$$

Figure 8: Labelled Transition System Axioms

in the presentation of the rules: for a type environment $\mathcal{I}_\Gamma$ we write

$$\begin{array}{l}\mathcal{I}^r_\Gamma(a,k) = \{\mathsf{T} \mid a : \mathsf{r}\langle\mathsf{T}\rangle@k \in \mathcal{I}_\Gamma \text{ or } a : \mathsf{rw}\langle\mathsf{T},\mathsf{U}\rangle@k \in \mathcal{I}_\Gamma\}\\ \mathcal{I}^w_\Gamma(a,k) = \{\mathsf{U} \mid a : \mathsf{w}\langle\mathsf{U}\rangle@k \in \mathcal{I}_\Gamma \text{ or } a : \mathsf{rw}\langle\mathsf{T},\mathsf{U}\rangle@k \in \mathcal{I}_\Gamma\}\end{array}$$

The input rule (M-RECEIVE) is a mild generalisation of the corresponding rule in [8], given there as (LTS-IN). Note that the action is only possible if the environment has access rights to its location k, that is if k is in $\mathcal{I}_\mathcal{T}$. Because SAFEDPI is asynchronous there are two forms of output actions. The rule

(M-RED)
$$\frac{M \longrightarrow M'}{(\mathcal{I} \rhd M) \xrightarrow{\tau} (\mathcal{I} \rhd M')}$$
(M-PAR)
$$\frac{(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')}{(\mathcal{I} \rhd M \mid N) \xrightarrow{\alpha} (\mathcal{I}' \rhd M' \mid N)} \qquad \frac{(\mathcal{I} \rhd N \mid M) \xrightarrow{\alpha} (\mathcal{I}' \rhd N \mid M')}{}$$
(M-NEW)
$$\frac{(\mathcal{I}, n : \top \rhd M) \xrightarrow{\alpha} (\mathcal{I}', n : \top \rhd M')}{(\mathcal{I} \rhd (\mathsf{new}\, n : \mathsf{E})\, M) \xrightarrow{\alpha} (\mathcal{I}' \rhd (\mathsf{new}\, n : \mathsf{E})\, M')} \quad n \notin \mathsf{n}(\alpha)$$
(M-OPEN)
$$\frac{(\mathcal{I}, m : \top \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')}{(\mathcal{I} \rhd (\mathsf{new}\, m : \mathsf{E})\, M) \xrightarrow{(m)\alpha} (\mathcal{I}' \rhd M')} \quad m \notin \mathsf{subj}(\alpha), m \in \mathsf{obj}_{\mathsf{I}}(\alpha)$$
(M-WEAK)
$$\frac{(\mathcal{I}, \{n : \mathsf{E}\} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')}{(\mathcal{I} \rhd M) \xrightarrow{(n:\mathsf{E})\alpha} (\mathcal{I}' \rhd M')} \quad n \notin \mathsf{subj}(\alpha), n \in \mathsf{obj}_{\mathsf{I}}(\alpha)$$
(M-τWEAK)
$$\frac{((\mathcal{I}_{\mathrm{T}}, \{k : \mathsf{K}\}), \mathcal{I}_{\mathcal{T}} + k) \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')}{(\mathcal{I} \rhd M) \xrightarrow{(k:\mathsf{K})\alpha} (\mathcal{I}' \rhd M')} \quad k \notin \mathsf{subj}(\alpha), k \in \mathsf{obj}_{\mathsf{I}}(\alpha)$$

Figure 9: Labelled Transition System Rules

(M-DELIVER) represents the delivery of a value to a channel, although it may not necessarily be consumed; note again that access rights are required to the channels' location.

There are three versions of the second form of output rule, in which the value is consumed by the channel; the variation depends on the type of the channel, but all require access rights. The first, (M-SEND.VAL), for first-order values, is an extension of the corresponding rule, (LTS-OUT), from [8]; note that here the environment's knowledge is increased, by adding the information contained in $\{\tilde{u} : (\mathsf{T})_{\omega} k\}$. Output of scripts is handled by (M-SEND.SCRIPT), where the environment supplies an appropriate G for further investigation of the script F. Dependent scripts, $(\tilde{u}, F)$ are handled by (M-SEND.DEP.SCRIPT); here the values $(\tilde{u})$ are exported from the system to the environment, while G, used for further investigation of F is imported to the system from the environment.

The final rule in Figure 8, (M-GOTO), is novel. It allows the environment to place arbitrary (well-typed) code at a site k, even if it does not have access rights there, provided it knows of a port p at k. Of course, in accordance with our operational semantics, k is free to ignore this code, by not proffering an input at the port p.

The inference rules for the action judgements (23) are given in Figure 9, and again they are informed by the corresponding rules in Figure 10 of [8]. Here we abuse notation a little by writing $(m)\alpha$ to mean $(\tilde{n} : \tilde{\mathsf{E}})(m, \tilde{m})\alpha'$ whenever α is $(\tilde{n} : \tilde{\mathsf{E}})(\tilde{m})\alpha'$. Note that, unlike in [8], we have two weakening rules; the new one, (M-τWEAK), allows the environment to invent a new location k at which it has access rights.

As a sanity check on these judgements we give a precise description of the possible forms the actions can take; to aid readability we will use G to represent a script furnished by the environment and F to represent one furnished by the system:

Proposition 6.7 *Suppose that $\mathcal{I} \rhd M$ is a configuration from which $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd N)$, where α is not τ. Then α takes one of the following forms:*

First-order: *input $(\tilde{n} : \tilde{\mathsf{E}})k.a.(\tilde{u})?$, where $(\tilde{n}) \subseteq (\tilde{u})$, or output $(\tilde{m})k.a.(\tilde{u})!$, where $(\tilde{m}) \subseteq (\tilde{u})$*

Script: *input $(\tilde{n} : \tilde{\mathsf{E}})k.a.F?$, where $(\tilde{n}) \subset \mathsf{fn}(F)$, or output $(\tilde{n} : \tilde{\mathsf{E}})k.a.G!$ where $(\tilde{n}) \subset \mathsf{fn}(G)$*

Dependent script: *input $(\tilde{n} : \tilde{\mathsf{E}})k.a.(\tilde{u}, F)?$, where $(\tilde{n}) \subseteq (\tilde{u}) \cup \mathsf{fn}(F)$, or output $(\tilde{n} : \tilde{\mathsf{E}})(\tilde{m})k.a.(\tilde{u}, G)!$, where $(\tilde{n}) \subset \mathsf{fn}(G)$ and $(\tilde{m}) \subset (\tilde{u})$*

Aynchronous-goto: $(\tilde{n} : \tilde{\mathsf{E}})\mathit{go}_p k.F$, *where* $(\tilde{n}) \subseteq \mathsf{fn}(F)$.

Proof: By induction on the inference of $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd N)$. ■

Proposition 6.8 (Well-definedness) *Suppose $\mathcal{I} \rhd M$ is a configuration. Then $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd N)$ implies $\mathcal{I}' \rhd N$ is also a configuration.*

Proof: By induction on the inference of $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd N)$, and an analysis of the last rule used; the details are similar to the corresponding result, Proposition 4.4 of [8]; the access rights component of $\mathcal{I}, \mathcal{I}_{\mathcal{T}}$ only plays a role in one rule, (M-τWEAK), and even there it is a minor role.

The axiom (M-RECEIVE) requires an application of the substitution results, Corollary 5.5 or Lemma 5.7 depending on the transmission type involved. The remaining axioms are straightforward, as their premises contain sufficient typing information to guarantee that the residual is indeed a configuration.

The proof for the rule (M-RED) depends on Subject Reduction, Theorem 5.8, while that for (M-NEW) relies on Weakening; the remaining rules follow immediately by induction. ■

With this result we now have a labelled transition system for SAFEDPI, the nodes being configurations and the actions all judgements (23) which can be inferred from Figure 8 and Figure 9. The standard definition of bisimulation therefore gives a co-inductive relation over configurations:

Definition 6.9 (Bisimulations) We say the binary relation between configurations $\mathcal{R}$ is a *typed bisimulation* if $\mathcal{C}\,\mathcal{R}\,\mathcal{D}$ implies

- $\mathcal{C} \xrightarrow{\alpha} \mathcal{C}'$ implies $\mathcal{D} \xRightarrow{\hat{\alpha}} \mathcal{D}'$ for $\mathcal{D}'$ such that $\mathcal{C}'\,\mathcal{R}\,\mathcal{D}'$
- $\mathcal{D} \xrightarrow{\alpha} \mathcal{D}'$ implies $\mathcal{C} \xRightarrow{\hat{\alpha}} \mathcal{C}'$ for $\mathcal{C}'$ such that $\mathcal{C}'\,\mathcal{R}\,\mathcal{D}'$

where $\xRightarrow{\hat{\alpha}}$ is the standard notation, meaning $\xrightarrow{\tau}{}^* \xrightarrow{\alpha} \xrightarrow{\tau}{}^*$ for α not equal to τ and $\xrightarrow{\tau}{}^*$ otherwise.

We write $\mathcal{I} \models M \approx_{bis} N$ whenever there exists some bisimulation $\mathcal{R}$ such that $(\mathcal{I} \rhd M)\,\mathcal{R}\,(\mathcal{I} \rhd N)$. ■

With this notation, that is by viewing the knowledge-structure $\mathcal{I}$ as a parameter, we construe $\approx_{bis}$ to be a knowledge-indexed relation over systems. This enables us to compare it directly with the touchstone behavioural equivalence $\approx_{cxt}$. The main technical property we require of $\approx_{bis}$ is given in the following result:

Proposition 6.10 *The knowledge-indexed relaton $\approx_{bis}$ is contextual.*

Proof: This follows similar lines to the equivalent statement in [8]. For this reason we only show that $\approx_{bis}$ is preserved by parallel composition here. Let $\mathcal{R}$ be defined by

$$(\mathcal{I} \rhd (\mathsf{new}\,\tilde{n} : \tilde{\mathsf{T}}_1)\,M \mid \prod_{i\in I} k_i[\![P_i]\!])\;\mathcal{R}\;(\mathcal{I} \rhd (\mathsf{new}\,\tilde{n} : \tilde{\mathsf{T}}_2)\,N \mid \prod_{i\in I} k_i[\![P_i]\!])$$

if and only if there exists some $\mathcal{I}'_\Gamma$, $(\tilde{\mathsf{T}})$ and $\mathcal{T}'$ such that

$$\begin{array}{l}
\mathcal{I}'_\Gamma <: \mathcal{I}_\Gamma \\
(\tilde{\mathsf{T}}_1) <: (\tilde{\mathsf{T}}) \qquad \text{and} \qquad (\tilde{\mathsf{T}}_2) <: (\tilde{\mathsf{T}}) \\
\mathcal{T}' \subseteq \tilde{n} \\
\mathcal{I}'_\Gamma \vdash k_i[\![P_i]\!] \text{ and } k_i \in \mathcal{I}_\mathcal{T} + \mathcal{T}' \text{ for each } i \in I \\
(\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T} + \mathcal{T}'), \{\tilde{n} : \mathsf{T}\} \models M \approx_{bis} N
\end{array}$$

We aim to show that $\mathcal{R}$ is a bisimulation from which the result follows immediately. For the purposes of this exposition we will assume that $\tilde{n}$ is empty and that the indexing set I is a singleton. We take any

$$(\mathcal{I} \rhd M \mid k[\![P]\!])\;\mathcal{R}\;(\mathcal{I} \rhd N \mid k[\![P]\!])$$

so we have some $\mathcal{I}'_\Gamma$ such that

$$(\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \models M \approx_{bis} N \tag{24}$$

with $\mathcal{I}'_\Gamma \vdash k[\![P]\!]$ and $k \in \mathcal{I}_\mathcal{T}$. We suppose that $(\mathcal{I} \rhd M \mid k[\![P]\!]) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')$ and now must show that there is a corresponding matching move from $(\mathcal{I} \rhd N \mid k[\![P]\!])$. In cases in which α is not τ this is easily done by appealing to (24). For $\alpha = \tau$ we know that $\mathcal{I}' = \mathcal{I}$ also. By an analogue of the Decomposition Lemma of [8] we can obtain five possibilities:

1. $(\mathcal{I} \rhd M) \xrightarrow{\tau} (\mathcal{I} \rhd M'')$ such that $M' \equiv M'' \mid k[\![P]\!]$
2. $k[\![P]\!] \longrightarrow M''$ such that $M' \equiv M \mid M''$
3. for first order T, $(\mathcal{I} \rhd M) \xrightarrow{(\tilde{m})k.a.\tilde{v}!} (\mathcal{I}'' \rhd M'')$ with
   - $k[\![P]\!] \equiv k[\![a?(\tilde{x} : \mathsf{T})\,Q]\!]$
   - $M' \equiv (\mathsf{new}\,\tilde{m} : \mathsf{U})\,M'' \mid k[\![Q[\{\!|\tilde{v}/\tilde{x}|\!\}]]\!]$
4. for other T, $(\mathcal{I} \rhd M) \xrightarrow{(\tilde{m})k.a.V!} (\mathcal{I}'' \rhd M'')$ with
   - $k[\![P]\!] \equiv k[\![a?(\tilde{x} : \mathsf{T})\,Q]\!]$
   - $V = (\tilde{v}, \lambda\,\tilde{x} : \mathsf{T}.\,Q)$
   - $(\mathsf{new}\,\tilde{m} : \mathsf{U})\,M'' \longrightarrow M'$ derived from (R-BETA)
5. $(\mathcal{I} \rhd M) \xrightarrow{(\tilde{n}:\mathsf{T})k.a.V?} (\mathcal{I}' \rhd M'')$ with
   - $k[\![P]\!] \equiv k[\![a!\langle V\rangle]\!]$
   - $M' \equiv M'' \mid k[\![\mathsf{stop}]\!]$

For each case we show that these conditions lead to the desired matching transition. We deal with each of them in turn.

- For (1) we appeal directly to (24).
- More interesting is case (2), particularly when the reduction is generated by use of the rule (R-L.CREATE) or (R-MOVE). We examine each of these: suppose $k[\![P]\!] \longrightarrow M''$ is derived from a use of (R-L.CREATE) so that
  $$\begin{array}{l} P = (\mathsf{newloc}\,l : \mathsf{L})\,\mathsf{with}\,C\;\mathsf{in}\;Q \\ M'' \equiv (\mathsf{new}\,l : \mathsf{L})\,l[\![C]\!] \mid k[\![Q]\!] \end{array}$$
  We know by (24) that $(\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \models M \approx_{bis} N$ and hence, by weakening to introduce a new testable location, we have
  $$(\mathcal{I}'_\Gamma, l : \mathsf{loc}, \mathcal{I}_\mathcal{T} + l) \models M \approx_{bis} N$$
  and by further weakening we obtain,
  $$(\mathcal{I}'_\Gamma, \{l : \mathsf{L}\}, \mathcal{I}_\mathcal{T} + l) \models M \approx_{bis} N$$
  Call the knowledge structure above, $\mathcal{I}''$. We know, by construction of $\mathcal{R}$, that $\mathcal{I}'_\Gamma \vdash k[\![P]\!]$ with $k \in \mathcal{I}_\mathcal{T}$, and therefore, according to the type rules (TY-NEWLOC), (TY-SUBPROC) and (TY-PROC), we must have
  $$\mathcal{I}'' \vdash l[\![C]\!] \qquad \text{and} \qquad \mathcal{I}'' \vdash k[\![Q]\!]$$
  with $l, k \in \mathcal{I}''_\mathcal{T}$ also. Therefore, by definition of $\mathcal{R}$ again, we see that
  $$\mathcal{I} \models (\mathsf{new}\,l : \mathsf{L})(M \mid l[\![C]\!] \mid k[\![Q]\!])\;\mathcal{R}\;(\mathsf{new}\,l : \mathsf{L})(N \mid l[\![C]\!] \mid k[\![Q]\!]) \tag{25}$$

We know that $(\mathcal{I} \rhd N \mid k[\![P]\!]) \xrightarrow{\tau} (\mathcal{I} \rhd (\mathsf{new}\, l : \mathsf{L})(N \mid l[\![C]\!] \mid k[\![Q]\!]))$ and that $M' \equiv M \mid M'' \equiv (\mathsf{new}\, l : \mathsf{L})(M \mid l[\![C]\!] \mid k[\![Q]\!])$, so by (25), we have

$$\mathcal{I} \models M' \,\mathcal{R}\, (\mathsf{new}\, l : \mathsf{L})(N \mid l[\![C]\!] \mid k[\![Q]\!])$$

and our matching transition as required.

Alternatively, suppose that $k[\![P]\!] \longrightarrow M''$ is derived from an instance of (R-MOVE). We then have

$$P = \mathsf{goto}\, p.lF \qquad \text{and} \qquad M'' \equiv l[\![p!\langle F\rangle]\!]$$

for some p, l, F. It is important to note here that the location l may not be contained in $\mathcal{I}_\mathcal{T}$ and this prevents us from immediately using the definition of relation $\mathcal{R}$ to claim that

$$\mathcal{I} \models M \mid l[\![p!\langle F\rangle]\!] \,\mathcal{R}\, N \mid l[\![p!\langle F\rangle]\!]$$

However, we do know that $\mathcal{I}'_\Gamma \vdash k[\![P]\!]$ so

$$(\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \rhd M \xrightarrow{\mathsf{go}_p l.F} (\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \rhd M \mid l[\![p!\langle F\rangle]\!]$$

is a valid transition. The hypothesis (24) tells us that there is a matching transition

$$(\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \rhd N \xrightarrow{\mathsf{go}_p l.F} (\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \rhd N''$$

such that $(\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \models M \mid l[\![p!\langle F\rangle]\!] \approx_{bis} N''$. This tells us that there is some N' such that

$$N \longrightarrow^* N' \qquad \text{and} \qquad N' \mid l[\![p!\langle F\rangle]\!] \longrightarrow^* N''$$

Therefore, it is clear that $(\mathcal{I} \rhd N \mid k[\![P]\!]) \Longrightarrow (\mathcal{I} \rhd N'')$ with $\mathcal{I} \models M \mid l[\![p!F]\!] \approx_{bis} N''$ as required.

- Cases (3) and (4) are similar in nature so we only show the reasoning for the latter. We have, in this instance, that

$$(\mathcal{I} \rhd M) \xrightarrow{(\tilde{m})k.a.(\tilde{m}', G)!} (\mathcal{I}'' \rhd M'')$$

where

$G = \lambda\, \tilde{x} : \mathsf{T}.\, Q$
$P = a?(\tilde{x} : \mathsf{T})\, Q$
$M'' \longrightarrow M'''$ (from (R-BETA) such that $M' \equiv (\mathsf{new}\, \tilde{m} : \mathsf{U}')\, M'''$
$\tilde{m} \subseteq \tilde{m}'$

It is easy to check (*cf.* Lemma 4.8 of [8]) that

$$(\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \rhd M \xrightarrow{(\tilde{m})k.a.(\tilde{m}', G)!} (\mathcal{I}'_\Gamma, \{\tilde{m}' : \mathsf{U}\} \rhd M''), \mathcal{I}_\mathcal{T}) \rhd M''$$

where $\mathsf{U}' <: \mathsf{U}$. Call the target knowledge structure $\mathcal{I}'''$. This tells us, by (24) that there exists a matching transition

$$(\mathcal{I}'_\Gamma, \mathcal{I}_\mathcal{T}) \rhd N \xrightarrow{(\tilde{m})k.a.(\tilde{m}', G)?} (\mathcal{I}''' \rhd N'')$$

with $\mathcal{I}''' \models M'' \approx_{bis} N''$. Note that $M'' \longrightarrow M'''$ (derived from (R-BETA)) guarantees, by confluence properties of beta-reduction, that $\mathcal{I}''' \models M''' \approx_{bis} N''$ and we can also assume, without loss of generality that N'' is stable with respect to β–reductions. By analysing the above transition we see that there exists some N''', $\tilde{n} : \mathsf{T}'$ and V such that

$$N \longrightarrow^* (\mathsf{new}\, \tilde{m} : \mathsf{U}'')\, (\mathsf{new}\, \tilde{n} : \mathsf{T}')(N''' \mid k[\![a!\langle V\rangle]\!])$$

with

$$(\mathsf{new}\, \tilde{n} : \mathsf{T}')(N''' \mid k[\![\lambda\, \tilde{x} : \mathsf{T}.\, Q(V)]\!]) \longrightarrow^* N'' \qquad \text{and} \qquad \mathsf{U}'' <: \mathsf{U}$$

Therefore we have

$$\begin{array}{rcl} N \mid k[\![P]\!] & \longrightarrow^* & (\mathsf{new}\, \tilde{m} : \mathsf{U}'')\, (\mathsf{new}\, \tilde{n} : \mathsf{T}')(N''' \mid k[\![a!\langle V\rangle]\!] \mid k[\![a?(\tilde{x} : \mathsf{T})\, Q]\!]) \\ & \longrightarrow^* & (\mathsf{new}\, \tilde{m} : \mathsf{U}'')\, (\mathsf{new}\, \tilde{n} : \mathsf{T}')(N''' \mid k[\![Q\{\!|V/\tilde{x}|\!\}]\!]) \\ & \longrightarrow^* & (\mathsf{new}\, \tilde{m} : \mathsf{U}'')\, N'' \\ & \equiv & N' \end{array}$$

Given that $M' \equiv (\mathsf{new}\, \tilde{m} : \mathsf{U}')\, M'''$, we have enough to conclude that $\mathcal{I} \models M' \,\mathcal{R}\, N'$ as required.

- Finally, in case (5) we follow a similar argument to that in [8] with only a slight modification to account for the asynchronous nature of SAFEDPI. ■

6.3 Relating bisimulation and contextual barbed congruence

This section is devoted to showing that these equivalences, viewed as knowledge-indexed relations coincide.

Proposition 6.11 (Soundness of $\approx_{bis}$ for $\approx_{cxt}$)

$$\mathcal{I} \models M \approx_{bis} N \;\; \textit{implies} \;\; \mathcal{I} \models M \approx_{cxt} N.$$

Proof: It is evident that $\approx_{bis}$ forms a symmetric, reduction closed and barb preserving knowledge-indexed relation. Therefore, because of Proposition 6.10 $\approx_{bis}$ satisfies all the defining properties of $\approx_{cxt}$. Since $\approx_{cxt}$ is the largest such relation the result follows. ■

The force of this proposition is that any distinctions made between systems by the contextual congruence can also be made by the labelled transition system. This means that we have provided enough labels of sufficient distinguishing power. We must also check that we have not provided too much distinguishing power in the labelled transition system. This is done by relating each action defined in the labelled transition system to an actual well-typed SAFEDPI context.

Proposition 6.12 (Definability (*cf. Prop 4.4 of [9]*)) *For each label α and each knowledge structure $\mathcal{I}$ there exists a system $\mathcal{C}_\alpha^\mathcal{I}$ which uses the fresh barb name δ, port name δ_{in} and location k_0 and tests for α, in the sense that*

- *if $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')$ then $\mathcal{I}, \{k_0 : K_0\} \vdash \mathcal{C}_\alpha^\mathcal{I}$ and moreover, $\mathcal{C}_\alpha^\mathcal{I} \mid M \longrightarrow^* (\mathsf{new}\, \tilde{m} : \tilde{\mathsf{E}})(k_0[\![\delta_{\mathsf{in}}!\langle \delta!\langle\rangle\rangle]\!] \mid M'')$ with $M'' \equiv M'$*
- *if $\mathcal{C}_\alpha^\mathcal{I} \mid M \longrightarrow^* (\mathsf{new}\, \tilde{m} : \tilde{\mathsf{E}})(k_0[\![\delta_{\mathsf{in}}!\langle \delta!\langle\rangle\rangle]\!] \mid M'')$ and $\mathcal{I}, \{k_0 : K_0\} \vdash \mathcal{C}_\alpha^\mathcal{I}$ where $\tilde{m} = \mathit{obj}_\mathsf{i}(\alpha)$ then $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')$ with $M'' \equiv M'$.*

where

$$K_0 = \mathsf{loc}[\delta_{\mathsf{in}} : \mathsf{rw}\langle \mathsf{thunk}\rangle, \delta : \mathsf{rw}\langle \mathbf{unit}\rangle, \delta_{\mathsf{fail}} : \mathsf{rw}\langle \mathbf{unit}\rangle, \delta_{\mathsf{succ}} : \mathsf{rw}\langle \mathbf{unit}\rangle]$$

(the barbs δ_{fail} and δ_{succ} are to be used later).

Proof: These systems are, for the most part, straightforward, and readers familiar with the work in [8, 9] will have little trouble reconstructing them.

As an example we show the systems for $k.a.(\tilde{v}, G)!$ and $\mathsf{go}_p l.V$ actions: we define

$$\mathcal{C}^\mathcal{I}_{k.a.(\tilde{v},G)!} \stackrel{def}{=} k[\![a?(\tilde{x}, y)\, \mathsf{if}\ \tilde{x} = \tilde{v}\ \mathsf{then}\ G(\tilde{x}, y) \mid \mathsf{goto}_{\delta_{\mathsf{in}}}\, k_0.\delta!\langle\rangle\ \ \mathsf{else}\ \ \mathsf{stop}]\!]$$

and

$$\mathcal{C}^\mathcal{I}_{\mathsf{go}_p l.V} \stackrel{def}{=} k_0[\![\delta_{\mathsf{in}}!\langle \delta!\langle\rangle\rangle \mid \mathsf{goto}_p\, l.V]\!]$$

The interested reader is invited to check that, for any configuration such that $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I} \rhd M')$ for one of these actions then it is the case that $\mathcal{I}_\Gamma, \{k_0 : K_0\} \vdash \mathcal{C}_\alpha$ and moreover $\mathcal{C}_\alpha \mid M \longrightarrow^* k_0[\![\delta_{\mathsf{in}}!\langle \delta!\langle\rangle\rangle]\!] \mid M''$ where M'' is structurally equivalent to M' up to collection of terminated garbage threads $l[\![\mathsf{stop}]\!]$. ■

By providing such testing systems for each action in the lts provided above we are able to establish our second main result

Theorem 6.13 (Full abstraction of $\approx_{bis}$ for $\approx_{cxt}$)

$$\mathcal{I} \models M \approx_{cxt} N \text{ if and only if } \mathcal{I} \models M \approx_{bis} N.$$

Proof: (Sketch) One direction is given by Proposition 6.11. The converse is shown by building a bisimulation from all pairs of configurations such that $\mathcal{I} \models M \approx_{cxt} N$. Specifically, let $\mathcal{R}$ be a relation over configurations defined by

$$(\mathcal{I} \models M)\ \ \mathcal{R}\ \ (\mathcal{I} \models N)$$

if $\mathcal{I} \models M \approx_{cxt} N$. We outline the proof that $\mathcal{R}$ defines a bisimulation, from which the result follows.

To this end suppose $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')$, where $\mathcal{I} \models M\ \mathcal{R}\ N$. We must find a matching move $(\mathcal{I} \rhd N) \stackrel{\alpha}{\Longrightarrow} (\mathcal{I}' \rhd N')$, such that $\mathcal{I}' \models M'\ \mathcal{R}\ N'$. For the purposes of this sketch we assume for simplicity that $\mathcal{I} = \mathcal{I}'$. By Definability, Proposition 6.12. We know that there exists a system $\mathcal{C}_\alpha^\mathcal{I}$, typeable from $\mathcal{I}_\Gamma, \{k_0 : K_0\}$, which satisfies the conditions of contextuality for knowledge-indexed relations and moreover, induces an interaction when plugged with M. In other words,

$$\mathcal{C}_\alpha^\mathcal{I} \mid M \longrightarrow^* k_0[\![\delta_{\mathsf{in}}!\langle \delta!\langle\rangle\rangle]\!] \mid M'' \tag{26}$$

for some δ, δ_{in} at k_0 and M'' equivalent to M' up to structure and garbage collection. We make use of this property of $\mathcal{C}_\alpha^\mathcal{I}$ as follows: first for the barb names, δ_{fail} and δ_{succ} in K_0 let

$$\mathsf{Flip} \stackrel{def}{=} k_0[\![\delta_{\mathsf{fail}}!\langle\rangle \mid \delta?().\delta_{\mathsf{fail}}?().\delta_{\mathsf{succ}}!\langle\rangle]\!]$$

and let

$$\mathcal{D}_\alpha^\mathcal{I} \stackrel{def}{=} (k_0[\![\delta_{\mathsf{in}}?(X : \mathsf{thunk})X()]\!] \mid \mathsf{Flip} \mid \mathcal{C}_\alpha^\mathcal{I} \mid -)$$

It is easy to check that $\mathcal{I}_\Gamma, \{k_0 : K_0\} \vdash \mathcal{D}_\alpha^\mathcal{I}$ whenever $\mathcal{I}_\Gamma, \{k_0 : K_0\} \vdash \mathcal{C}_\alpha^\mathcal{I}$. We should note that the reductions (26) above extend so that (up to structure and garbage collection)

$$\mathcal{D}_\alpha^\mathcal{I}[M] \longrightarrow^* k_0[\![\delta_{\mathsf{succ}}!\langle\rangle]\!] \mid M''$$

The hypothesis $\mathcal{I} \models M \approx_{cxt} N$, the fact that $\mathcal{I}_\Gamma, \{k_0 : K_0\} \vdash \mathcal{C}_\alpha^\mathcal{I}$ and weakening, contextuality and barb preserving properties of $\approx_{cxt}$ together allow us to use $(\mathcal{I}_\Gamma, \{k_0 : K_0\}, \mathcal{I}_\mathcal{T} + k_0) \models \mathcal{D}_\alpha^\mathcal{I}[M] \approx_{cxt} \mathcal{D}_\alpha^\mathcal{I}[N]$ to find a matching transition

$$\mathcal{D}_\alpha^\mathcal{I}[N] \longrightarrow^* k_0[\![\delta_{\mathsf{succ}}!()]\!] \mid N''$$

with

$$(\mathcal{I}_\Gamma, \{k_0 : K_0\}, \mathcal{I}_\mathcal{T} + k_0) \models k[\![\delta!\langle\rangle]\!] \mid M'' \approx_{cxt} k[\![\delta!\langle\rangle]\!] \mid N''.$$

Note that we can guarantee this form by the absence of the δ_{fail} barb in $k_0[\![\delta_{\mathsf{succ}}!\langle\rangle]\!] \mid M''$ and the fact that, by symmetry, absence of barbs must also be preserved. The systems $\mathcal{C}_\alpha^\mathcal{I}$ are also built in such a way as to guarantee that whenever $\mathcal{D}_\alpha^\mathcal{I}[N] \longrightarrow^* k_0[\![\delta_{\mathsf{succ}}!\langle\rangle]\!] \mid N''$ then we must also have $\mathcal{I} \rhd N \stackrel{\alpha}{\Longrightarrow}$

$\mathcal{I} \rhd N'$ where, again, N'' is equivalent to N' up to structural equivalence and garbage collection. It is easy to show directly that

$$(\mathcal{I}_\Gamma, \{k_0 : K_0\}, \mathcal{I}_\mathcal{I} + k_0) \models k_0[\![\delta_{\mathsf{succ}}!\langle\rangle]\!] \mid M'' \approx_{cxt} k_0[\![\delta_{\mathsf{succ}}!\langle\rangle]\!] \mid N''$$

implies $\mathcal{I} \models M' \approx_{cxt} N'$ which is enough to conclude with $\mathcal{I} \models M' \mathcal{R} N'$. A symmetric argument establishes that $\mathcal{R}$ is a bisimulation.

The case in which $(\mathcal{I} \rhd M) \xrightarrow{\alpha} (\mathcal{I}' \rhd M')$ for $\mathcal{I}'$ not equal to $\mathcal{I}$ is slightly more complicated and is dealt with using an *Extrusion Lemma* similar to that found in [6, 9, 8]. ■

This provides an alternative characterisation of reduction barbed congruence which models the nature of knowledge acquisition possible by testing with highly constrained mobile code in an explicit way.

7 Conclusion

We have developed a sophisticated type system for controlling the behaviour of mobile code in distributed systems, and demonstrated that, at least in principle, coinductive proof principles can still be applied to investigate their behaviour.

The use of types in this manner could be considered as a particular case of the general approach of *proof-carrying code*, [18] and *typed assembly language (TAL)* [17]. Here hosts would publish their safety policies in terms of a type or logical proposition and code wishing to enter would have to arrive with a proof, which a typechecker or proofchecker can use to verify that it satisfies the published policy. Indeed we intend to use the types of the current paper in this manner, by extending the work in [20]. The work of [18] and [17] has inspired much further research into the use of type systems in higher-level languages for resource access and usage monitoring, [23], [12], for example. However the emphasis in these papers is on dynamics and counting of resource usage rather than using sophisticated types to specify fine-grained access control.

There has been much work on modelling mobility and locations using particular process calculi. Perhaps the calculus closest to SAFEDPI is the Seal Calculus, [5]. Seals are hierarchically organised computational sites in which inter-seal communication, which is channel-based, is only allowed among siblings or between parents and siblings. Seals may also be communicated, rather like the communication of higher-order processes along ports in SAFEDPI; indeed in some sense it is more general as the seal being transmitted may be computationally active. However the communication of seals is more complicated, as it involves agreement between three participants, the sender, the receiver, and the seal being transmitted. Seals are also typed using *interfaces*, similar to our fine-grained process types, π. But

these only record the *input* capabilities a seal offers to its parents, and in order to preserve interfaces under reduction the transmission of input channel capabilities is forbidden in the language. This is a severe restriction, at least in general distributed computing, if not in the more focused application area of seals. For example the generation of new servers requires the the transmission of input capabilities. We believe that our dependent and existential types can also be applied to the Seal Calculus, to obtain a more general notion of interface, which will still be preserved by reduction.

The M-calculus, [22], a higher-order extension of the distributed join calculus, is also closely related, at least conceptually, to SAFEDPI. Here, not only are locations hierarchically organised, but are *programmable*, in the sense that entry and exit policies for each location can be explicitly programmed. In addition it has an interesting operator, called *passivation*, which can freeze the contents of a site into a value. However their type system is not related to one we have developed for SAFEDPI; the latter addresses access control issues for migrating code whereas the former is concerned with unicity of locations; in a higher-order language with a passivation operator it is important to ensure that each locality has a unique name. Thus the type system for the M-Calculus draws on that presented in [25], where unicity of the location of channel names was addressed, rather than that of [26], which developed fine-grained access control types for processes.

Type systems have also been used to explicitly control mobility in distributed calculi, most notably in variants of the Ambient calculus of Cardelli and Gordon [3]. In particular, [2], [16] use subtyping to control movement of mobile processes in a hierarchically distributed system by introducing explicit types to express permission to migrate. A similar technique was used for DPI in [10], [8]. In contrast, here we control mobility only indirectly through types. Code is always permitted to migrate provided it has access to a suitable port at the target location. But by restricting the use of channels in the types this consequently restricts migration. Indeed, we decouple permission to migrate from the location name itself, affording more flexibility in the control of migration.

The coinductive characterisation presented here makes use of *higher-order actions* in the sense that, to interact with a system willing to send a script V, the environment must supply a receiving script G to which V will be applied. A similar approach is used in the characterisation theorems for various forms of ambients in [7] and [15]. Higher-order actions are also used in the bisimulation equivalence presented in [4] for the Seal calculus. However, there the three way nature of higher-order communication leads to a proliferation of such actions, some of which can not be simulated by seal contexts; see Section 4.4 of [5] for examples. As a result the bisimulation equivalence is more discriminating than the natural contextual equivalence for seals.

Such higher-order bisimulations do not directly result in automatic veri-

(E-EMPTY)
$$\frac{}{\vdash \mathbf{env}}$$

(E-GRES)
$$\frac{\Gamma \vdash \mathsf{C} : \mathbf{ty}}{\Gamma,\ u : \mathsf{rc}\langle \mathsf{C}\rangle \vdash \mathbf{env}} \quad u \notin \Gamma,\ \mathsf{C}\text{ closed}$$

(TY-LOOKUP)
$$\frac{\Gamma, u : \mathsf{T}, \Gamma' \vdash \mathbf{env}}{\Gamma, u : \mathsf{T}, \Gamma' \vdash_{lookup} u : \mathsf{T}}$$

(E-NEWLCHAN)
$$\frac{\Gamma \vdash_{lookup} w : \mathsf{loc} \quad \Gamma \vdash \mathsf{C} : \mathbf{ty}}{\Gamma,\ u : \mathsf{C}@w \vdash \mathbf{env}} \quad u \notin \Gamma,\ \mathsf{C}\text{ closed}$$

(E-EDEP)
$$\frac{\Gamma, \{x_1 : \mathsf{E}_1\}, \ldots, \{x_n : \mathsf{E}_n\} \vdash \mathsf{T} : \mathbf{ty}}{\Gamma, y : \langle \mathsf{T}\ \mathsf{with}\ \widetilde{x} : \widetilde{\mathsf{E}}\rangle \vdash \mathbf{env}} \quad x_i, y \notin \Gamma,\ y \neq x_i$$

(E-SCRIPT)
$$\frac{\Gamma \vdash \mathsf{S} : \mathbf{ty}}{\Gamma, x : \mathsf{S} \vdash \mathbf{env}} \quad x \notin \Gamma$$

(E-LOC)
$$\frac{\Gamma \vdash \mathbf{env}}{\Gamma,\ u : \mathsf{loc} \vdash \mathbf{env}} \quad u \notin \Gamma$$

(TY-ELOOKUP)
$$\frac{\Gamma, \langle \tilde{x} : \tilde{\mathsf{E}}, y : \mathsf{T}\rangle, \Gamma' \vdash \mathbf{env}}{\Gamma, \langle \widetilde{x} : \widetilde{\mathsf{E}}, y : \mathsf{T}\rangle, \Gamma' \vdash_{lookup} y : \mathsf{T}}$$

(E-LCHAN)
$$\frac{\Gamma \vdash_{lookup} w : \mathsf{loc} \quad \Gamma \vdash_{lookup} u : \mathsf{rc}\langle \mathsf{D}\rangle \quad \Gamma \vdash \mathsf{D} <: \mathsf{C}}{\Gamma,\ u : \mathsf{C}@w \vdash \mathbf{env}} \quad \mathsf{C}\text{ closed}$$

Figure 10: Well-defined Environments

fication methods for distributed systems. But they do serve to focus on the essential features of systems which determine their behaviour; for example our results for SAFEDPI have demonstrated the importance of the *goto* moves $\mathsf{go}_p k.V$. Moreover they serve as a starting point for more in-depth analyses of the behaviour of SAFEDPI systems, and more particularly of interesting sub-languages. For example is it possible to use the technique of [13] to find a fully-abstract bisimulation equivalence which only uses first-order labels? There the receiving contexts for higher-order values are replaced by symbolic representatives. Although not directly applicable due to the extra complication of distribution and mobility control, it would be of great interest to pursue those ideas in the current setting.

Acknowledgements: The first two authors would like to acknowledge the financial support of the two EU Global computing projects, **Mikado** and **Myths**.

(SUB-BASE)
$$\frac{\Gamma \vdash \mathbf{env}}{\Gamma \vdash \mathbf{base} <: \mathbf{base}}$$

(SUB-TOP)
$$\frac{\Gamma \vdash \mathsf{T} <: \mathsf{T}}{\Gamma \vdash \mathsf{T} <: \top}$$

(SUB-PROCTOP)
$$\frac{\Gamma \vdash \pi <: \pi}{\Gamma \vdash \pi <: \mathbf{proc}}$$

(SUB-CHAN)
$$\frac{\Gamma \vdash \mathsf{T}_r <: \mathsf{U}_r, \mathsf{U}_w <: \mathsf{T}_w,\ \mathsf{T}_w <: \mathsf{T}_r}{\begin{array}{l}\Gamma \vdash \mathsf{w}\langle \mathsf{T}_w\rangle <: \mathsf{w}\langle \mathsf{U}_w\rangle,\\ \Gamma \vdash \mathsf{r}\langle \mathsf{T}_r\rangle <: \mathsf{r}\langle \mathsf{U}_r\rangle\\ \Gamma \vdash \mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w\rangle <: \mathsf{rw}\langle \mathsf{U}_r, \mathsf{U}_w\rangle\end{array}} \qquad \frac{\Gamma \vdash \mathsf{T}_r <: \mathsf{U}_r, \mathsf{U}_w <: \mathsf{T}_w\ \ \mathsf{T}_w <: \mathsf{T}_r}{\begin{array}{l}\Gamma \vdash \mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w\rangle <: \mathsf{w}\langle \mathsf{U}_w\rangle\\ \Gamma \vdash \mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w\rangle <: \mathsf{r}\langle \mathsf{U}_r\rangle\end{array}}$$

(SUB-LOC)
$$\frac{\Gamma \vdash_{lookup} u_i : \mathsf{rc}\langle \mathsf{D}_i\rangle\ \ \Gamma \vdash \mathsf{D}_i <: \mathsf{C}_i,\ \mathsf{D}_j <: \mathsf{C}'_j,\ \Gamma \vdash \mathsf{C}_i <: \mathsf{C}'_i,}{\Gamma \vdash \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_m : \mathsf{C}_m] <: \mathsf{loc}[u_1 : \mathsf{C}'_1, \ldots, u_n : \mathsf{C}'_n]} \quad 0 \leq n \leq m$$

(SUB-HOM)
$$\frac{\Gamma \vdash \mathsf{C} <: \mathsf{C}' \quad \Gamma \vdash_{lookup} w : \mathsf{loc}}{\Gamma \vdash \mathsf{C}@w <: \mathsf{C}'@w}$$
$$\Gamma \vdash \mathsf{rc}\langle \mathsf{C}\rangle <: \mathsf{rc}\langle \mathsf{C}'\rangle$$

(SUB-SCRIPT)
$$\frac{\Gamma, \{x_1 : (\mathsf{T}_1)@\mathsf{here}\}, \ldots, \{x_n : (\mathsf{T}_n)@\mathsf{here}\} \vdash \pi <: \pi'}{\Gamma \vdash \mathsf{Fdep}(\widetilde{x} : \widetilde{\mathsf{T}} \to \pi) <: \mathsf{Fdep}(\widetilde{x} : \widetilde{\mathsf{T}} \to \pi')}$$

(SUB-PROC)
$$\frac{\Gamma \vdash u_i : \mathsf{C}_i@w_i,\ u_j : \mathsf{C}'_j@w_j\ \ \Gamma \vdash \mathsf{C}'_i@w'_i <: \mathsf{C}_i@w_i}{\Gamma \vdash \mathsf{pr}[u_1 : \mathsf{C}_1@w_1, \ldots, u_m : \mathsf{C}_m@w_m] <: \mathsf{pr}[u_1 : \mathsf{C}'_1@w_1, \ldots, u_n : \mathsf{C}'_n@w_n]} \quad 0 \leq m \leq n$$

(SUB-TUDEP)
$$\frac{\Gamma, \{x_1 : \mathsf{E}_1\}, \ldots, \{x_n : \mathsf{E}_n\} \vdash \mathsf{T} <: \mathsf{T}'}{\Gamma \vdash \mathsf{Tdep}(\widetilde{x} : \widetilde{\mathsf{E}})\,\mathsf{T} <: \mathsf{Tdep}(\widetilde{x} : \widetilde{\mathsf{E}})\,\mathsf{T}'}$$

(SUB-EDEP)
$$\frac{\Gamma, \{x_1 : \mathsf{E}_1\}, \ldots, \{x_n : \mathsf{E}_n\} \vdash \mathsf{T} <: \mathsf{T}'}{\Gamma \vdash \mathsf{Edep}(\widetilde{x} : \widetilde{\mathsf{E}})\,\mathsf{T} <: \mathsf{Edep}(\widetilde{x} : \widetilde{\mathsf{E}})\,\mathsf{T}'}$$

Figure 11: Subtyping

A Auxiliary Definitions and Results

Types and Type Environments: The judgements for well-defined environments, $\Gamma \vdash \mathbf{env}$, and subtyping, $\Gamma \vdash \mathsf{T} <: \mathsf{U}$, are defined simultaneously,

using the rules in Figure 10 and Figure 11. The former are a mild extension of the corresponding rules in Figure 6 of [8] to accommodate script and dependent types and rely on a predicate $\Gamma \vdash_{lookup} u : \mathsf{T}$, which simply looks up the type associated with u in Γ. The latter is an extension of the well-known subtyping rules of types in the PICALCULUS, [21], and DPI, [10, 8]; the rules for process types are similar to those used in [26]. The judgements also check that the identifiers used in T, U are actually declared appropriately in Γ.

Proposition A.1 (Sanity Checks)

- $\Gamma \vdash \mathsf{T} <: \mathsf{U}$ *implies* $\Gamma \vdash \mathbf{env}$
- $\Gamma \vdash \mathsf{T} <: \mathsf{U}$ *implies* $\Gamma \vdash \mathsf{T} : \mathbf{ty}$ *and* $\Gamma \vdash \mathsf{U} : \mathbf{ty}$
- $\Gamma \vdash \mathsf{T} <: \mathsf{U}$, $\Gamma \vdash \mathsf{U} <: \mathsf{R}$ *implies* $\Gamma \vdash \mathsf{T} <: \mathsf{R}$
- $\Gamma, u : \mathsf{T} \vdash \mathbf{env}$ *implies* $\Gamma \vdash \mathbf{env}$ *and* $\Gamma \vdash \mathsf{T} : \mathbf{ty}$

Proof: By rule induction. ■

Meets and Joins: The partial operators $\sqcap$, $\sqcup$ on type expressions are defined by extending the definitions used in [10, 8] for channel and location types. We take them to be the least reflexive and symmetric operators which satisfy a series of rules for combining together various kinds of type expressions. Those governing channel expressions are, as in [10]:

- $\mathsf{r}\langle \mathsf{T}_1 \rangle \sqcap \mathsf{r}\langle \mathsf{T}_2 \rangle = \mathsf{r}\langle \mathsf{T}_1 \sqcap \mathsf{T}_2 \rangle$, $\quad \mathsf{r}\langle \mathsf{T}_1 \rangle \sqcup \mathsf{r}\langle \mathsf{T}_2 \rangle = \mathsf{r}\langle \mathsf{T}_1 \sqcup \mathsf{T}_2 \rangle$
- $\mathsf{w}\langle \mathsf{T}_1 \rangle \sqcap \mathsf{w}\langle \mathsf{T}_2 \rangle = \mathsf{w}\langle \mathsf{T}_1 \sqcup \mathsf{T}_2 \rangle$, $\quad \mathsf{w}\langle \mathsf{T}_1 \rangle \sqcup \mathsf{w}\langle \mathsf{T}_2 \rangle = \mathsf{w}\langle \mathsf{T}_1 \sqcap \mathsf{T}_2 \rangle$
- $\mathsf{r}\langle \mathsf{T}_r \rangle \sqcap \mathsf{w}\langle \mathsf{T}_w \rangle = \mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \rangle$
- $\mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \rangle \sqcap \mathsf{r}\langle \mathsf{T}'_r \rangle = \mathsf{rw}\langle \mathsf{T}_r \sqcap \mathsf{T}'_r, \mathsf{T}_w \rangle$,
  $\mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \rangle \sqcup \mathsf{r}\langle \mathsf{T}'_r \rangle = \mathsf{rw}\langle \mathsf{T}_r \sqcup \mathsf{T}'_r, \mathsf{T}_w \rangle$,
- $\mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \rangle \sqcap \mathsf{w}\langle \mathsf{T}'_w \rangle = \mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \sqcup \mathsf{T}'_w \rangle$,
  $\mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \rangle \sqcup \mathsf{r}\langle \mathsf{T}'_w \rangle = \mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \sqcap \mathsf{T}'_w \rangle$,

To express the rules for location types we take advantage of the fact that the ordering of their components is immaterial:

- $\mathsf{loc}[u_1 : \mathsf{C}'_1] \sqcap \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n] = \mathsf{loc}[u_1 : (\mathsf{C}'_1 \sqcap \mathsf{C}_1), \ldots, u_n : \mathsf{C}_n]$,
  $\mathsf{loc}[u_1 : \mathsf{C}'_1] \sqcup \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n] = \mathsf{loc}[u_1 : (\mathsf{C}'_1 \sqcup \mathsf{C}_1)]$
- if u does not occur in $\{u_1, \ldots, u_n\}$ then
  $\mathsf{loc}[u : \mathsf{C}] \sqcap \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n] = \mathsf{loc}[u : \mathsf{C}, u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n]$,
  $\mathsf{loc}[u : \mathsf{C}] \sqcup \mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n] = \mathsf{loc}[]$

- $\mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n] \sqcap \mathsf{K} = \mathsf{loc}[u_1 : \mathsf{C}_1] \sqcap (\ldots (\mathsf{loc}[u_n : \mathsf{C}_n] \sqcap \mathsf{K}) \ldots)$,
  $\mathsf{loc}[u_1 : \mathsf{C}_1, \ldots, u_n : \mathsf{C}_n] \sqcup \mathsf{K} = (\mathsf{loc}[u_1 : \mathsf{C}_1] \sqcup \mathsf{K}) \sqcap \ldots \sqcap (\mathsf{loc}[u_n : \mathsf{C}_n] \sqcup \mathsf{K})$

We use a similar approach to defining the operations on process types, where we use GC as an arbitrary type of the form $\mathsf{C}@w$. However the process type constructor is contravariant, whereas the location constructor is covariant.

- $\mathsf{pr}[u_1 : \mathsf{C}'_1@w_1] \sqcap \mathsf{pr}[u_1 : \mathsf{C}_1@w_1, \ldots, u_n : \mathsf{GC}_n] = \mathsf{pr}[u_1 : (\mathsf{C}'_1 \sqcup \mathsf{C}_1)@w_1]$,
  $\mathsf{pr}[u_1 : \mathsf{C}'_1@w_1] \sqcup \mathsf{pr}[u_1 : \mathsf{C}_1@w_1, \ldots, u_n : \mathsf{GC}_n] =$
  $\mathsf{pr}[u_1 : (\mathsf{C}'_1 \sqcap \mathsf{C}_1)@w_1, \ldots, u_n : \mathsf{GC}_n]$
- if $u@w$ does not occur in $\{u_1@w_1, \ldots, u_n@w_n\}$ then
  $\mathsf{pr}[u : \mathsf{C}@w] \sqcap \mathsf{pr}[u_1 : \mathsf{C}_1@w, \ldots, u_n : \mathsf{C}_n@w_n] = \mathsf{pr}[]$,
  $\mathsf{pr}[u : \mathsf{C}@w] \sqcup \mathsf{pr}[u_1 : \mathsf{C}_1@w_1, \ldots, u_n : \mathsf{C}_n@w_n] =$
  $\mathsf{pr}[u : \mathsf{C}@w, u_1 : \mathsf{C}_1@w_1 \ldots, u_n : \mathsf{C}_n@w]$
- $\mathsf{pr}[u_1 : \mathsf{GC}_1, \ldots, u_n : \mathsf{GC}_n] \sqcap \pi =$
  $(\mathsf{pr}[u_1 : \mathsf{GC}_1] \sqcap \pi) \sqcup \ldots \sqcup (\mathsf{pr}[u_n : \mathsf{GC}_n] \sqcap \pi)$,
  $\mathsf{pr}[u_1 : \mathsf{GC}_1, \ldots, u_n : \mathsf{GC}_n] \sqcup \pi = \mathsf{pr}[u_1 : \mathsf{GC}_1] \sqcup (\ldots (u_n : \mathsf{GC}_n \sqcup \pi) \ldots)$
- $\mathbf{proc} \sqcap \pi = \pi$, $\quad \mathbf{proc} \sqcup \pi = \mathbf{proc}$

For the various forms of dependent types, the rules are straightforward:

- $\mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi) \sqcap \mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi') = \mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to (\pi \sqcap \pi'))$,
  $\mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi) \sqcup \mathsf{Fdep}(\tilde{x} : \tilde{\mathsf{T}} \to \pi') = \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\,(\pi \sqcup \pi')$
- $\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T} \sqcap \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T}' = \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\,(\mathsf{T} \sqcap \mathsf{T}')$,
  $\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T} \sqcup \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T}' = \mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{T}})\,(\mathsf{T} \sqcup \mathsf{T}')$
- $\mathsf{Edep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T} \sqcap \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T}' = \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{T}})\,(\mathsf{T} \sqcap \mathsf{T}')$,
  $\mathsf{Edep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T} \sqcup \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{T}})\,\mathsf{T}' = \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{T}})\,(\mathsf{T} \sqcup \mathsf{T}')$

For the remaining kinds of type expressions we merely extend the definitions homomorphically:

- $\mathsf{rc}\langle \mathsf{C} \rangle \sqcap \mathsf{rc}\langle \mathsf{C}' \rangle = \mathsf{rc}\langle \mathsf{C} \sqcap \mathsf{C}' \rangle$, $\quad \mathsf{rc}\langle \mathsf{C} \rangle \sqcup \mathsf{rc}\langle \mathsf{C}' \rangle = \mathsf{rc}\langle \mathsf{C} \sqcup \mathsf{C}' \rangle$
- $\mathsf{T}@w \sqcap \mathsf{T}'@w = (\mathsf{T} \sqcap \mathsf{T}')@w$

Proposition A.2

- *If there exists some type expression* T *such that* $\Gamma \vdash \mathsf{T} <: \mathsf{T}_1$ *and* $\Gamma \vdash \mathsf{T} <: \mathsf{T}_2$ *then* $\mathsf{T}_1 \sqcap \mathsf{T}_2$ *is well-defined*
- *When* $\mathsf{T}_1 \sqcap \mathsf{T}_2$ *is well-defined,* $\Gamma \vdash \mathsf{T}_1 \sqcap \mathsf{T}_2 <: \mathsf{T}_i$ *and* $\Gamma \vdash \mathsf{T} <: \mathsf{T}_1 \sqcap \mathsf{T}_2$*, for any type expression* T *such that* $\Gamma \vdash \mathsf{T} <: \mathsf{T}_1$ *and* $\Gamma \vdash \mathsf{T} <: \mathsf{T}_2$*.*
- *If there exists some type expression* T *such that* $\Gamma \vdash \mathsf{T}_1 <: \mathsf{T}$ *and* $\Gamma \vdash \mathsf{T}_2 <: \mathsf{T}$ *then* $\mathsf{T}_1 \sqcup \mathsf{T}_2$ *is well-defined*

- *When* $\mathsf{T}_1 \sqcup \mathsf{T}_2$ *is well-defined,* $\Gamma \vdash \mathsf{T}_i <: \mathsf{T}_1 \sqcup \mathsf{T}_2$*, and* $\Gamma \vdash \mathsf{T}_1 \sqcup \mathsf{T}_2 <: \mathsf{T}$*, for any type expression* T *such that* $\Gamma \vdash \mathsf{T}_1 <: \mathsf{T}$ *and* $\Gamma \vdash \mathsf{T}_2 <: \mathsf{T}$.

Proof: The first and third statements are proved by induction on the derivations of $\Gamma \vdash \mathsf{T}_i <: \mathsf{T}$ and $\Gamma \vdash \mathsf{T} <: \mathsf{T}_i$ respectively. The second and fourth are by induction on the construction of $\mathsf{T}_1 \sqcap \mathsf{T}_2$, $\mathsf{T}_1 \sqcup \mathsf{T}_2$ respectively. ■

Note that because of the top type $\top$ the premise of the third statement is always true; so $\mathsf{T}_1 \sqcup \mathsf{T}_2$ always exists, although in many cases it will be the uninformative type $\top$.

Substitutions: Free identifiers may occur in type expressions and therefore we need to define $\mathsf{T}\{\!|v/u|\!\}$ for an arbitrary type expression T; this is then used as part of the definition of substitution into process terms. The definition of $\mathsf{T}\{\!|v/u|\!\}$ is by induction on the structure of T. The only interesting cases are location and process types, where the definition needs to ensure that the entries remain unique:

- $\mathsf{loc}[u' : \mathsf{C}]\{\!|v/u|\!\} = \mathsf{loc}[u'\{\!|v/u|\!\} \ : \ \mathsf{C}\{\!|v/u|\!\}]$
- $\mathsf{loc}[u_1 : \mathsf{C}_1, \ldots u_n : \mathsf{C}_n]\{\!|v/u|\!\} = (\mathsf{loc}[u_1 : \mathsf{C}_1]\{\!|v/u|\!\}) \sqcap \ldots \sqcap (\mathsf{loc}[u_n : \mathsf{C}_n]\{\!|v/u|\!\})$
- $\mathsf{pr}[u' : \mathsf{C}]\{\!|v/u|\!\} = \mathsf{pr}[u'\{\!|v/u|\!\} \ : \ \mathsf{C}\{\!|v/u|\!\}]$
- $\mathsf{pr}[u_1 : \mathsf{C}_1, \ldots u_n : \mathsf{C}_n]\{\!|v/u|\!\} = (\mathsf{pr}[u_1 : \mathsf{C}_1]\{\!|v/u|\!\}) \sqcup \ldots \sqcup (\mathsf{pr}[u_n : \mathsf{C}_n]\{\!|v/u|\!\})$
- All other cases are defined homomorphically. For example
  - $\mathsf{rw}\langle \mathsf{T}_r, \mathsf{T}_w \rangle\{\!|v/u|\!\} = \mathsf{rw}\langle \mathsf{T}_r\{\!|v/u|\!\}, \mathsf{T}_w\{\!|v/u|\!\} \rangle$
  - $\mathsf{Tdep}(\tilde{x} : \tilde{\mathsf{E}})\, \mathsf{T}\{\!|v/u|\!\} = \mathsf{Tdep}(\tilde{x} : (\tilde{\mathsf{E}}\{\!|v/u|\!\}))\, (\mathsf{T}\{\!|v/u|\!\})$, where we assume v is different from each x_i

Proposition A.3 • *Suppose* $\mathsf{T} \sqcap \mathsf{U}$ *is defined. Then so is* $\mathsf{T}\{\!|v/u|\!\} \sqcap \mathsf{U}\{\!|v/u|\!\}$ *and (up to* α*-equivalence) is the same as* $(\mathsf{T} \sqcap \mathsf{U})\{\!|v/u|\!\}$

- *Similarly for* $\mathsf{T} \sqcup \mathsf{U}$.

Proof: By simultaneous induction on the definitions of $\mathsf{T} \sqcap \mathsf{U}$ and $\mathsf{T} \sqcup \mathsf{U}$. ■

Substitution of identifiers also commutes with the channel extraction function.

Proposition A.4 *For all identifiers* u, v,

$$\mathsf{pr}_{\mathsf{ch}}[V : \mathsf{T}]\{\!|v/u|\!\} = \mathsf{pr}_{\mathsf{ch}}[V\{\!|v/u|\!\} : \mathsf{T}\{\!|v/u|\!\}]$$

Proof: By induction on the definition of $\mathsf{pr}_{\mathsf{ch}}[V : \mathsf{T}]$. The only non-trivial case is when V is an identifier w and T a location type, when the proof depends on the peculiaries of the application of substitutions to location and process types. ■

Proposition A.5 (Substitution) *Suppose* $\Gamma \vdash_w v : \mathsf{T}$ *and* $x \notin \Gamma$. *Then*

- $\Gamma, x : (\mathsf{T})@w, \Delta \vdash \mathbf{env}$ *implies* $\Gamma, \Delta\{\!|v/x|\!\} \vdash \mathbf{env}$
- $\Gamma, , x : (\mathsf{T})@w, \Delta \vdash \mathsf{T} <: \mathsf{U}$ *implies* $\Gamma, \Delta\{\!|v/x|\!\} \vdash \mathsf{T}\{\!|v/x|\!\} <: \mathsf{U}\{\!|v/x|\!\}$

Proof: By simulataneous induction on the derivations. Note that there are only four possibilities for the entry $x : (\mathsf{T})@w$, namely $x : \mathsf{loc}$, $x : \mathsf{rc}\langle \mathsf{D} \rangle$, $x : \mathsf{C}@w$ or $x : \mathsf{S}$. ■

The corresponding substitution result for existential values depends on the following property of existential witnesses.

Proposition A.6 *Let* Γ_e *denote* $\Gamma, y : \langle \mathsf{T} \text{ with } \tilde{x} : \tilde{\mathsf{T}} \rangle, \Gamma'$. *Then*

- $\Gamma_e \vdash \mathbf{env}$ *implies* x_i *does not occur in* Γ'.
- $\Gamma_e \vdash \mathsf{T} <: \mathsf{U}$ *implies* x_i *does not ocur free in* T, U.

Proposition A.7 *Suppose* $\Gamma \vdash_w \langle \tilde{v}, v \rangle : \mathsf{Edep}(\tilde{x} : \tilde{\mathsf{E}})\, \mathsf{T}$. *Let* Γ_e *denote* $\Gamma, y : \langle (\mathsf{T})@w \text{ with } \tilde{x} : (\tilde{\mathsf{E}})@w \rangle, \Delta$. *Then*

- $\Gamma_e \vdash \mathbf{env}$ *implies* $\Gamma, \Delta\{\!|v/y|\!\} \vdash \mathbf{env}$
- $\Gamma \vdash \mathsf{U}_1 <: \mathsf{U}_2$ *implies* $\mathsf{U}_1\{\!|\tilde{v}, v/\tilde{x}, y|\!\} <: \mathsf{U}_2\{\!|\tilde{v}, v/\tilde{x}, y|\!\}$

Proof: By simultaneous induction on the inferences. ■

Adding knowledge to environments: Here we extend the meet operator $\sqcap$ to lists of type associations. This is used in Figure 9, in the rules (M-SEND.VAL) and (M-SEND.DEP.SCRIPT), for increasing the knowledge in a type environment. We first define the (partial) operation $\Gamma \sqcap u : \mathsf{E}$ between an arbitrary association list Γ and a singleton:

- If E is a located channel $\mathsf{A}@w$ then $\Gamma \sqcap u : \mathsf{E}$ is $\Gamma, u : \mathsf{E}$.
- Otherwise if u has no association in Γ then $\Gamma \sqcap u : \mathsf{E}$ is also $\Gamma, u : \mathsf{E}$.
- Otherwise $\Gamma \sqcap u : \mathsf{E}$ is obtained by replacing the association of u in Γ, say $u : \mathsf{E}'$, by the new association $u : (\mathsf{E} \sqcap \mathsf{E}')$; in this case the operation is only defined if $(\mathsf{E} \sqcap \mathsf{E}')$ exists.

The general definition of $\Gamma_1 \sqcap \Gamma_2$ then follows by induction on the size of Γ_2:

- $\Gamma_1 \sqcap \epsilon = \Gamma_1$
- $\Gamma_1 \sqcap (\Gamma_2', u : \mathsf{E}) = (\Gamma_1 \sqcap \Gamma_2') \sqcap u : \mathsf{E}$

References

[1] M. Boreale and D. Sangiorgi. Bisimulation in name-passing calculi without matching. In *Proc. 13th LICS Conf.* IEEE Computer Society Press, 1998.

[2] L. Cardelli, G. Ghelli, and A. Gordon. Ambient groups and mobility types. In *Proc. IFIP TCS 2000*, volume 1872 of *Lecture Notes in Computer Science*. Springer-Verlag, 2000.

[3] L. Cardelli and A. Gordon. Mobile ambients. In *Proc. FoSSaCS '98*, LNCS. Springer-Verlag, 1998.

[4] G. Castagna and F. Zappa Nardelli. The Seal calculus revisited: Contextual equivalences and bisimilarity. In *Proceedings of FSTTCS*, Lecture Notes in Computer Science, 2002.

[5] Giuseppe Castagna, Jan Vitek, and Francesco Zappa. The Seal calculus. 2003. Available from ftp://ftp.di.ens.fr/pub/users/castagna/seal.ps.gz.

[6] C. Fournet, G. Gonthier, J-J. Levy, L. Maranget, and D. Remy. A calculus of mobile agents. In *Proc. CONCUR*, volume 1119 of *Lecture notes in computer science*. Springer-Verlag, 1996.

[7] M. Hennessy and M. Merro. Bisimulation congruences in safe ambients. In *Proc. POPL 02*. ACM Press, 2002.

[8] Matthew Hennessy, Massimo Merro, and Julian Rathke. Towards a behavioural theory of access and mobility control in distributed systems. Technical Report 2002:01, COGS, University of Sussex, 2002. Extended Abstract published in the Proceedings of FoSSaCS 2003.

[9] Matthew Hennessy and Julian Rathke. Typed behavioural equivalences for processes in the presence of subtyping. In James Harland, editor, *Electronic Notes in Theoretical Computer Science*, volume 61. Elsevier Science Publishers, 2002. To appear in *Mathematical Structures in Computer Science*.

[10] Matthew Hennessy and James Riely. Resource access control in systems of mobile agents. *Information and Computation*, 173:82–120, 2002.

[11] K. Honda and N. Yoshida. On reduction-based process semantics. *Theoretical Computer Science*, 152(2):437–486, 1995.

[12] Atsushi Igarashi and Naoki Kobayashi. Resource usage analysis. In *Proceedings of ACM Symposium on Principles of Programming Languages (POPL'02)*, pages 331–342, 2002.

[13] Alan Jeffrey and Julian Rathke. Contextual equivalence for higher-order π-calculus revisited. In *Proceedings MFPS XIX, Montreal*, 2003.

[14] Cedric Lhoussaine. Type inference for a distributed pi-calculus. In *ESOP'02*, volume 2618 of *LNCS*, pages 253–269. Springer-Verlag, 2002.

[15] Massimo Merro and Francesco Zappa Nardelli. Bisimulation proof techniques for mobile ambients. In *Proc. 30^{th} International Colloquium on Automata, Languages, and Programming (ICALP 2003), Eindhoven*, Lecture Notes in Computer Science. Springer-Verlag, 2003.

[16] Massimo Merro and Vladimiro Sassone. Typing and subtyping mobility in boxed ambients. In *Proceedings CONCUR 02*, volume 1644 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.

[17] Greg Morrisett, Karl Crary, Neal Glew, and David Walker. Stack-based typed assembly language. In *Types in Compilation*, volume 1473 of *Lecture notes in Computer Science*, pages 25–35. Springer-Verlag, 1998.

[18] George C. Necula. Proof-carrying code. In *Conference Record of POPL '97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 106–119, Paris, France, jan 1997.

[19] B. Pierce and D. Sangiorgi. Behavioral equivalence in the polymorphic pi-calculus. *Journal of the ACM*, 47(3):531–584, 2000.

[20] James Riely and Matthew Hennessy. Trust and partial typing in open systems of mobile agents (extended abstract). In *Conference Record of POPL '99 The 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 93–104, 1999. To appear in the Journal of Automated Reasoning.

[21] Davide Sangiorgi and David Walker. *The π-calculus*. Cambridge University Press, 2001.

[22] A. Schmitt. and J.-B. Stefani. The M-calculus: A higher-order distributed process calculus. In *POPL2003*, January 2003.

[23] David Walker. A type system for expressive security properties. In *the twenty seventh ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Boston*, pages 254–267, 2000.

[24] Nobuko Yoshida. Channel dependent types for higher-order mobile processes. In *Conference Record of POPL '04: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Venice,Italy, jan 2004.

[25] Nobuko Yoshida and Matthew Hennessy. Subtyping and locality in distributed higher order processes. In *Proc. CONCUR*, volume 1664 of *Lecture notes in computer science*. Springer-Verlag, 1999.

[26] Nobuko Yoshida and Matthew Hennessy. Assigning types to processes. *Information and Computation*, 172:82–120, 2002.

A Dependently Typed Ambient Calculus*

Cédric Lhoussaine and Vladimiro Sassone
University of Sussex, UK

Abstract. *The Ambient calculus is a successful model of distributed, mobile computation, and has been the vehicle of new ideas for resource access control. Mobility types have been used to enforce elementary access control policies, expressed indirectly via classification of ambients in groups by means of 'group types.' The paper presents a theory of dependent types for the Ambient calculus which allows greater flexibility, while keeping the complexity away from the programmer into the type system.*

1 Introduction

The work on types for *Mobile Ambients* [6] (MA), initiated in [7], has introduced some interesting innovations. In particular, it led to natural concepts of mobility types (cf., e.g. [4, 3, 12, 8]). The common framework proposed so far in the literature for mobility control relies on *groups* and *group types*. Groups represent collections of ambients with uniform access rights, and policies like '*n CanMoveTo m*' are expressed by saying, for instance, '*n belongs to group G* and all ambients of *group G can move to m*.' These and similar simple constraints, such as '*m accepts ambients of group G*,' are then enforced statically by embodying groups in types. This approach's merit is to simplify considerably the type system. It allows to avoid the difficulties of dependent types, since it replaces (variable) ambient names with (constant) group names. The loss in expressiveness and flexibility, however, is self-evident.

This paper generalises the approaches to access control based on groups by using types explicitly dependent on ambient names, so that policies like '*n CanMoveTo m*' can be expressed directly. The task involves issues of technical complexity, and we believe that our approach contributes significantly to advance the theory of type-based resource access control in the Ambient calculus. The guiding principle behind our work is to allow for flexibility in control policies, while pushing the complexity to the type system. We elaborate further on this below.

Dependent types vs groups. At the formal level, dependent types are indeed trickier to work with than groups, as already suggested in [4]. However, they are very expressive and ultimately, at the programming level, not necessarily more difficult to use than groups. It may indeed be challenging, if not impossible, to partition ambient names into groups in order to implement the desired security policy in complex applications. By referring directly to names, instead, a programmer avoids a level of indirection, and is less prone to policy specification errors. The type systems we propose here aim

* Research supported by 'MyThS: Models and Types for Security in Mobile Distributed Systems', EU FET-GC IST-2001-32617.

at harnessing the complexity of dependent types and relegating it to inner working of the type system, so as to keep the burden off the programmer, and consequently make programs more robust.

To illustrate the point, let us consider a naive and specialised taxi server whose intention is to provide a taxi service usable only by ambient n:

$$\mathit{TaxiServ}(n) = !(\nu\, \mathit{taxi} : A)\,(\mathit{taxi}[] \mid n[\mathsf{in}\, \mathit{taxi}]).$$

The process creates a new ambient called *taxi*, with some type A, together with an ambient n which can move to it. Due to the scoping rules, an external ambient $m[P]$ can move into the taxi by first moving to n and then escaping it after being transported to *taxi*. To avoid this, A must restrict the children allowed of *taxi* to, say, only ambient n. Using groups – for instance in the setting of [12], where an ambient type $a : \mathsf{amb}[G, \mathsf{mob}[\mathcal{G}]]$ assigns group G to a and specifies that $\mathcal{G}$ is the set of groups of those ambients which a is allowed to move to – to prevent any ambient other than n to occupy the taxi, n must be the *only* ambient with a type of the form $\mathsf{amb}[H, \mathsf{mob}[\{G\} \cup \mathcal{H}]]$, where G is the group of *taxi*. This condition is clearly beyond the control of the $\mathit{TaxiServ}(n)$ application, as the environment may create new names with types that violate it. Making G private is not a solution, as it forces n to be confined too, so defeating its very purpose. Indeed, in a term such as $(\nu\, G)\,(\nu\, n : A)\, P \mid Q$, with G in A, n can not be passed to Q (not even via subtyping [12]), as the type of any suitable receiving variable would contain G.

With dependent types we would simply and directly mention n as a possible (or perhaps the only) child for *taxi*. This leads to ambient types of the form $\mathsf{amb}[\mathsf{mob}[\mathcal{P}, \mathcal{C}]]$, where $\mathcal{P}$ is the set of names of those ambients allowed to contain *taxi* and, dually, $\mathcal{C}$ is the set of names of those ambients allowed to reside within *taxi*. Type A above can be simply $\mathsf{amb}[\mathsf{mob}[\{\mathit{top}\}, \{n\}]]$, where name *top* represents the top-level ambient. Yet, the *taxi* ambient would be guaranteed to have n and only n as a possible child, with no need for n to be private. Of course, such simplicity in expressing types comes at the cost of additional complexity at the type system level. Primarily, we need to enforce consistency between the type of a parent ambient and those of its children. To exemplify, when a taxi is created, n receives a new capability to have *taxi* as its parent; the type of n – which actually predates the creation of the fresh instance of *taxi* – must then be updated accordingly. Observe that in a conventional type system, the effect of the new name would be confined inside the scope of its definition, and would certainly not affect the type of n. In our case, however, leaving n's type unchanged leads to an unsound system (cf. Example 1).

In order for n to be aware of any additional capability it may acquire over time, we introduce the pivotal notion of *abstract names*. Such names are used only in the derivation of typing judgements, to record potential name creations, either by a process or by its environment. As it will be clear later, abstract names intuitively play the role of group names in their ability to cross the scope of name restrictions.

Flexibility. The main contribution of this paper is the use of dynamic types to increase expressiveness, which can be fleshed out as the ability to deploy security policies at run time. More precisely, we want servers flexible enough to provide services specialised to each particular client, so leading towards secure, highly modular programming. In other

words, a service specifically provided for client c_1 must be unusable by client c_2. Such objectives may be achieved in several different ways, as for instance by adding special operators to the calculus. In this paper we stick to the basic primitives of the Ambient calculus and investigate the issue of personalised services at the level of types.

The taxi server discussed above is not particularly interesting, because it is hardly a server at all: it does not interact with generic clients, but provides again and again the same 'secure' service to the same client. A brief reflection proves that if we were to design a 'true' server using groups, we would need to exchange group names, yielding soon a form of (group) dependent types. The type system we propose for dynamic types is a natural extension of the one we illustrated above, for mobility only. Namely, we simply require that communications of ambient names and variables be reflected in the types. For instance, the dynamic taxi server will have the form:

$$!(x)\mathit{TaxiServ}(x) = !(x)(\nu\, \mathit{taxi} : \mathsf{amb}[\mathsf{mob}[\{\mathit{top}\}, \{x\}]])\,(\mathit{taxi}[] \mid x[\mathsf{in}\ \mathit{taxi}]).$$

Communication increases the complexity of the type system, because the rules need to track information about receivers, too. It makes little sense to type communication via the usual exchange types, because receiving different names may lead to orthogonal behavioural and mobility properties. Subtyping does not help towards a unifying treatment (cf. Example 2). We approach the problem by tracking the names that may be received at a given variable. This enriches ambient types which, in addition to a mobility component, acquire a communication one. A typing assignment $n : \mathsf{amb}[\,\mathsf{mob}[\mathcal{P}, C], \mathsf{com}[\mathcal{E}, \mathcal{L}]\,]$, where $\mathcal{E}$ and $\mathcal{L}$ are sets of ambient names, means that n may be communicated inside any ambient in $\mathcal{E}$, and dually, that any ambient in $\mathcal{L}$ may be communicated within n. It is important to remark that even if such components only carry finite sets, the volume of information exchanged is in general not bounded: the joint effect of name creation and the use of abstract names, allows to distribute new communication capabilities at run time, and communicate infinitely many names.

Plan. The paper proceeds as follows. Section 2 presents a simple type system with dependent types. We focus on mobility, that is the Ambient calculus with no communication at all. Then, in Section 3, we investigate dynamic types allowing communication of ambient names. For the sake of simplicity, we do not address the issue of communication of capabilities, even though this can easily be integrated in our approach.

2 A Simple Dependent Type System

In the paper we use the standard Ambient calculus with *asynchronous*, *monadic* communication. (We remark that our results do not depend on such choice.) Syntax and reduction semantics can be found in [6]. In this section we consider the pure Ambient calculus, that is the calculus restricted to the mobility features. Communication will be studied in Section 3.

2.1 Types and definitions

The syntax of types is given in Figure 1 where N denotes an infinite set of *abstract ambient names* and the symbol $\star$ stands for the *self*-ambient, whose meaning is explained

Ambient (concrete) names $n, m, \ldots \in \mathcal{N}$		*Abstract names* $\mathsf{n}, \mathsf{m}, \ldots \in \mathsf{N}$	
Capability types	$K ::= \mathsf{cap}[\mathcal{P}]$	*Mobility types*	$M, N ::= \mathsf{mob}[\mathcal{P}, C]$
Ambient types	$A ::= \mathsf{amb}[M]$		where $\mathcal{P}, C \subseteq \mathcal{N} \cup \mathsf{N} \cup \{\star\}$

Fig. 1. Types

below. The type $\mathsf{cap}[\mathcal{P}]$ denotes the type of a capability which can be exercised within any ambient whose name occurs in $\mathcal{P}$. An ambient type has, for the moment, just one component: a mobility type. Ambient n with mobility type $\mathsf{mob}[\mathcal{P}, C]$ is allowed to be a sub-ambient of all ambients whose name occurs in $\mathcal{P}$; i.e. $\mathcal{P}$ is the set of *possible parents* of n. Moreover, an ambient is allowed to occur as sub-ambient of n if its name occurs in C, the set of *possible children* of n. We define some useful notations on types:

$$\mathsf{amb}[M]^{\mathsf{mob}} = M \qquad \mathsf{mob}[\mathcal{P}, C]^{\uparrow} = \mathcal{P} \qquad \mathsf{mob}[\mathcal{P}, C]^{\downarrow} = C$$

We use two kinds of typing contexts: *abstract contexts* (ranged over by $\Xi, \Theta, \ldots$) and *concrete contexts* (ranged over by $\Gamma, \Delta, \ldots$), which map respectively abstract and concrete names to ambient types. We use Π to denote either an abstract or a concrete context. We note by (Π, Π') the union of disjoint contexts of the same kind. According to the informal meaning of types, type assignments in concrete typing contexts are related to each other. Consider for instance

$$\Gamma = n : \mathsf{amb}[\,\mathsf{mob}[\mathcal{P}, C]\,],\ m : \mathsf{amb}[\,\mathsf{mob}[\emptyset, \{n\}]\,],$$

where the type of m allows it to have n as a child. Coherently, the type of n should allow n to have m as a parent, i.e. $m \in \mathcal{P}$. This can be expressed by the central notion of *context update* $\Gamma^{(n:A)}$, which updates Γ with respect to a fresh type assignment $n : A$:

$$\begin{aligned} (\Gamma, \Delta)^{(n:A)} &= \Gamma^{(n:A)}, \Delta^{(n:A)} \\ (m : \mathsf{amb}[M_0])^{(n:\mathsf{amb}[N])} &= m : \mathsf{amb}[M_1], \end{aligned} \tag{1}$$

where

$$M_1^{\uparrow} = \begin{cases} M_0^{\uparrow} \cup \{n\} \text{ if } m \in N^{\downarrow}, \\ M_0^{\uparrow} \qquad \text{otherwise;} \end{cases} \qquad \text{and} \qquad M_1^{\downarrow} = \begin{cases} M_0^{\downarrow} \cup \{n\} \text{ if } m \in N^{\uparrow}, \\ M_0^{\downarrow}. \end{cases}$$

Context update is pivotal in our type systems to express the typing of name creation. We define union (resp. inclusion and equality) of types as component-wise set union (resp. inclusion and equality); the notation $nm(A)$ (resp. $an(A)$) stands for the set of (abstract) names occurring in A. It is extended to all types and typing contexts.

2.2 The type system

Our type system deals with typing judgements for values and processes, that is

$$\Gamma \vdash^{\Theta} V : T \qquad \text{and} \qquad \Gamma \vdash_a^{\Xi;\Theta} P,$$

$$\frac{A \in \mathit{Types}(\Gamma, \Theta, a)}{\Gamma \vdash^{\Theta} a : A}\ (\textsc{VAmb}) \qquad \frac{\Gamma \vdash^{\Theta} U : K, V : K}{\Gamma \vdash^{\Theta} U.V : K}\ (\textsc{VPfx})$$

$$\frac{\Gamma \vdash^{\Theta} a : A, a_i : A_i \quad A^{\mathsf{mob}\uparrow} \subseteq A_i^{\mathsf{mob}\uparrow} \quad 1 \le i \le n}{\Gamma \vdash^{\Theta} \mathsf{out}\, a : \mathsf{cap}[\{a_1, \ldots, a_n\}]}\ (\textsc{VOut}) \qquad \frac{\Gamma \vdash^{\Theta} a : A \qquad \mathcal{P} \subseteq A^{\mathsf{mob}\downarrow}}{\Gamma \vdash^{\Theta} \mathsf{in}\, a : \mathsf{cap}[\mathcal{P}]}\ (\textsc{VIn})$$

$$\frac{\Gamma \vdash^{\Theta} a : A, a_i : A_i \qquad A^{\mathsf{mob}} \subseteq A_i^{\mathsf{mob}} \qquad 1 \le i \le n}{\Gamma \vdash^{\Theta} \mathsf{open}\, a : \mathsf{cap}[\{a_1, \ldots, a_n\}]}\ (\textsc{VOpen})$$

$$\frac{\Gamma \vdash_a^{\Xi;\Theta} P \quad \Gamma \vdash^{\Xi,\Theta} a : A \quad b \in A^{\mathsf{mob}\uparrow}}{\Gamma \vdash_b^{\Xi;\Theta} a[P]}\ (\textsc{Amb}) \qquad \frac{}{\Gamma \vdash_a^{\emptyset;\Theta} \mathbf{0}}\ (\textsc{Nil}) \qquad \frac{\Gamma \vdash_a^{\Xi;\Theta,\Xi\sigma} P}{\Gamma \vdash_a^{\Xi;\Theta} !P}\ (\textsc{Rep})$$

$$\frac{\Gamma \vdash_a^{\Xi;\Theta} P \quad \Gamma \vdash^{\Theta,\Xi} V : \mathsf{cap}[a]}{\Gamma \vdash_a^{\Xi;\Theta} V.P}\ (\textsc{Pfx}) \qquad \frac{\Gamma \vdash_a^{\Xi_1;\Xi_2,\Theta} P \qquad \Gamma \vdash_a^{\Xi_2;\Xi_1,\Theta} Q}{\Gamma \vdash_a^{\Xi_1,\Xi_2;\Theta} P \mid Q}\ (\textsc{Par})$$

$$\frac{\Gamma^{(n:A)}, n : A\{n/\star\} \vdash_a^{\Xi\{n/\mathsf{n}\};\Theta} P}{\Gamma \vdash_a^{\mathsf{n}:A,\Xi;\Theta} (\nu\, n : A)\, P}\ (\textsc{Res})$$

Fig. 2. Simple Type System

where T is either an ambient or a capability type. The first judgement reads as "V has type T with respect to the concrete context Γ and the abstract context Θ." In the second one, ambient name a stands for the ambient where P is running, which we call the *current location* of P. In other words, P is guaranteed to be well-typed if run in a. Abstract contexts Ξ and Θ are used to account for potential name creations: those arising from P are registered in the *local* abstract context Ξ, those performed by the environment appear in the *external* abstract context Θ. The role of such contexts is to "internalise" (a dynamic notion of) group names, which move from being part of the language to being a inner mechanism of the type system.

Example 1. Let

$$\mathit{TaxiServ}(n) = !(\nu\, \mathit{taxi} : A)\,(\mathit{taxi}[] \mid n[\mathsf{in}\, \mathit{taxi}]), \qquad \text{for} \quad A = \mathsf{amb}[\,\mathsf{mob}[\{\mathit{top}\}, \{n\}]\,].$$

Let $Q = n[\,\mathsf{in}\, n \mid m[\mathsf{out}\, n.\mathsf{out}\, n]\,]$. We study the system $P = \mathit{TaxiServ}(n) \mid Q$, which we assume running in some ambient named *top*. The server $\mathit{TaxiServ}(n)$ creates an ambient *taxi* whose type allows it to accept only n as a child. The scope (or "visibility") of *taxi* is $(\mathit{taxi}[] \mid n[\mathsf{in}\, \mathit{taxi}])$. We assume that Q is an ambient n which may move in an ambient of the same name n and contains a child ambient m, willing to escape twice out n. Thus, m must be allowed to run at the same level of n, and a type system should ensure that any possible parent for n is also a possible one for m. This leads to an assignment like Γ below.

$$\Gamma = \begin{cases} \mathit{top} : \mathsf{amb}[\,\mathsf{mob}[\emptyset, \{n, m\}]\,] \\ \quad n : \mathsf{amb}[\,\mathsf{mob}[\{\mathit{top}, n\}, \{n, m\}]\,] \\ \quad m : \mathsf{amb}[\,\mathsf{mob}[\{\mathit{top}, n\}, \emptyset]\,] \end{cases} \qquad \Delta = \begin{cases} \mathit{top} : \mathsf{amb}[\,\mathsf{mob}[\emptyset, \{n, m, \mathit{taxi}\}]\,] \\ \quad n : \mathsf{amb}[\,\mathsf{mob}[\{\mathit{top}, n, \mathit{taxi}\}, \{n, m\}]\,] \\ \quad m : \mathsf{amb}[\,\mathsf{mob}[\{\mathit{top}, n\}, \emptyset]\,], \\ \mathit{taxi} : \mathsf{amb}[\,\mathsf{mob}[\{\mathit{top}\}, \{n\}]\,] \end{cases}$$

In ($taxi[] \mid n[\mathsf{in}\, taxi]$), ambient n gains access to *taxi* by means of name creation, viz. *taxi*, and the type assignment must evolve to Δ above (which actually is of the form $\Gamma^{(taxi:A)}, taxi : A$). However, running P may obviously lead to m becoming a child of *taxi*, thus violating the specification expressed by the type of *taxi*. So, P must be considered ill-typed. A naive approach would however try to typecheck ($taxi[] \mid n[\mathsf{in}\, taxi]$) (against Δ) and Q (against Γ) independently, and therefore accept P. Indeed, from Q's point of view, *taxi* does not exist (it is bound in $TaxiServ(n)$), and Q behaves well. From the viewpoint of $TaxiServ(n)$, although *taxi* exists, the specification given by its type is respected: ambient n is allowed to move into *taxi*. Considering $TaxiServ(n)$ and Q as stand-alone processes prevents from realising that their interaction may lead to a breach of the intended access policy. □

Example 1 motivates the use of abstract contexts. They are a means to record the new capabilities which may potentially arise from name creations performed in external processes. Since such names are bound, they cannot be referred to by name in typing contexts: we use *abstract names* to represent them. The typechecking of process Q in Example 1 must then be carried out with an *external* abstract context $\Theta = \mathsf{taxi} : A$, representing the potential creation of a name of type A, and an empty *local* abstract context, as Q does not create names. In other words, we are led to prove $\Gamma \vdash_{top}^{\emptyset;\Theta} Q$. From Γ and Θ, we can deduce a more informative type for n, viz. its *global type* $\mathsf{amb}[\,\mathsf{mob}[\{top, n, \mathsf{taxi}\}, \{n, m\}]\,]$, which states that n is allowed to run within *top*, n and some ambient taxi, to be created by its environment and whose actual name we do not known yet. (We stress that the name taxi is just a placeholder; the information that leads to the global type is carried by A.) Since m does not appear in Θ, its global type remains $\mathsf{amb}[\mathsf{mob}[\{top, n\}, \emptyset]]$, and we can conclude that Q is ill-typed. Indeed, there exists a possible parent of n which is not a possible one for m, namely taxi.

We define the *symmetric type* $\Pi[\alpha]$ of a name α with respect to a typing context Π, as $\mathsf{amb}[\mathsf{mob}[\mathcal{P}, C]]$, where $\mathcal{P} = \{\beta \mid \alpha \in \Pi(\beta)^{\mathsf{mob}\downarrow}\}$ and $C = \{\beta \mid \alpha \in \Pi(\beta)^{\mathsf{mob}\uparrow}\}$.

Definition 1. The *global type* of a name α with respect to typing contexts Γ and Θ is defined as

$$gt(\Gamma;\Theta)(\alpha) = \begin{cases} \Gamma(\alpha) \cup \Theta[\alpha] & \text{if } \alpha \in \mathcal{N} \\ \Theta(\alpha)\{\alpha/\star\} \cup \Theta[\alpha] & \text{if } \alpha \in \mathsf{N} \end{cases}$$

We define $Types(\Gamma, \Theta, a) = \{gt(\Gamma;\Theta)(a)\}$. (We use a singleton set here for uniformity with next section.) Turning back to the previous discussion, one can verify that indeed

$$gt(\Gamma;\Theta)(n) = \mathsf{amb}[\mathsf{mob}[\{top, n, \mathsf{taxi}\}, \{n, m\}]].$$

The notion of symmetric type leads to the definition of *symmetric typing context*.

Definition 2. A concrete typing context Γ is said *symmetric* if $\Gamma(n) = \Gamma[n]$, for all $n \in nm(\Gamma)$.

Symmetric typing context formalises the notion of "consistency" of a typing context. It ensures that a name n has a type allowing m as a child of n if and only if m itself has a type allowing n as a parent. It also guarantees that a typing context is "complete,"

i.e. that any name occurring in it has a type assignment. In the following we will assume all concrete typing contexts to be symmetric, and all abstract contexts to be "complete" for abstract names, that is $dom(\Theta) = an(\Theta)$.

Symmetric typing replicates parenthood information both in children and in parents types. Nevertheless, both components of mobility types are necessary to give sufficient access control in the presence of name creation. We remark that comparable mechanisms play in groups based approaches via the assignment of names to groups. Using the types of [12] to exemplify, recall that $\mathcal{G}$ in $n : \mathsf{amb}[G, \mathsf{mob}[\mathcal{G}]]$ is the set of groups of ambients allowed as parents of n. Therefore, assigning n to G to has the effect of determining the children allowed of n: namely, those ambients whose mobility component contains G.

Let us now comment on the rules of the Figure 2, starting with the judgements for values. We look at the rules from conclusions to premises. A judgement $\Gamma \vdash^{\Theta} U : T, V : T'$ stands for $\Gamma \vdash^{\Theta} U : T$ and $\Gamma \vdash^{\Theta} V : T'$. Axiom (VAMB) simply gives the global type of a name, while (VPFX) asserts that the prefix $U.V$ has the capability type K if both U and V have such type. In order to type a $\mathsf{out}\, a$ performed in some ambient a_i, (VOUT) checks that the possible parents of a are allowed for a_i (for $i \in \{1, \ldots, n\}$). Rule (VIN) typechecks the capability of ambient b (a member of $\mathcal{P}$) to move to ambient a. This is allowed if b is a possible child of a. Finally, (VOPEN) states that some ambient a_i can open a if both the allowed parents and children of a are allowed for a_i, $i \in \{1, \ldots, n\}$. This ensures that whatever happens to a_i because of the capabilities inherited from a via the open is compatible with its specification.

A process is always typed with respect to a current location. Process $a[P]$ is well-typed in ambient b (rule AMB) if b is allowed as a's parent and P is well-typed in a. Process $\mathbf{0}$ is always well-typed with local abstract context empty. Process $V.P$ is well-typed in a if P is well-typed in a and V is a capability that can be exercised in a. In order to type a parallel composition of processes (rule (PAR)), the local abstract context must be split in two parts: in the typing of each component one of these remains local, the other becomes external.

Rule (RES) deserves a close analysis. The typing of $(\nu n : A)\, P$ requires the local abstract context to assign type A to some abstract name (here n). The rule consumes such assignment, and P is then typed with its concrete typing context updated with n's typing. We apply the substitution $\{n/\star\}$ to A, where $\star$ is an "alias" for the self ambient (n here). This is necessary, since in the conclusion n cannot occur in A. Indeed, we make the usual assumption that a bound name does not occur in the typing contexts and in the current location (which here means $n \neq a$). Finally, the substitution $\{n/\mathsf{n}\}$ is applied to the local abstract context, as types in Ξ may refer to n. This would be the case of, e.g., the term

$$(\nu n : \mathsf{amb}[\,\mathsf{mob}[\{top\}, \emptyset]\,])\,(\nu m : \mathsf{amb}[\,\mathsf{mob}[\{n\}, \emptyset]\,])\,P$$

typed with the local abstract context $\mathsf{n} : \mathsf{amb}[\,\mathsf{mob}[\{top\}, \emptyset]\,],\ \mathsf{m} : \mathsf{amb}[\,\mathsf{mob}[\{\mathsf{n}\}, \emptyset]\,]$.

In Rule (REP) we assume that σ is an *abstract renaming*. By this we mean an injective substitution whose domain is the set of abstract names occurring in Ξ and whose codomain is a set of fresh abstract names. In other words, if σ is an abstract renaming, then $\Xi\sigma$ is a fresh copy of Ξ. Process $!P$ can be seen as an infinite parallel composition of copies of P. Therefore, typing $!P$ is equivalent to typing P in an environment which

can provide new capabilities as P can. These are represented by the local abstract context of $!P$, which explains why we add a copy of the local abstract context to the external one. It is easy to prove that one copy is sufficient.

It is worth remarking that there are behavioural properties which are expressible with groups but not with this system of dependent types. This is because in the present system, name creation assigns types which may only refer to names already in the scope of the created name. For instance, in

$$(\nu\, n : A)\, P \mid (\nu\, m : A')\, Q$$

m (resp. n) can not occur in A (resp. A'). This means that n and m will not be able to interact. With group types one may use group names shared by m and n, i.e. whose scope contains both names. The extension of the type system to dynamic types introduced in the next section circumvents the problem by the usual mechanisms of scope extrusion in name passing calculi. Indeed, in order for m to give (and acquire) new capabilities to n, we send n to m, which then may refer in A' to the received name.

3 Dynamic Types

As illustrated in the previous section, name creation is itself a means to dynamically change ambient capabilities. In the type system, abstract names are used to track the relevant changes outside the actual scope of the name. New names may refer directly in their types to known ambient names. For instance, the taxi server

$$(\nu\, taxi : \mathsf{amb}[\mathsf{mob}[\{top\}, \{n\}]])\, (taxi[] \mid n[\mathsf{in}\, taxi])$$

creates a name *taxi* which may have n as a child. At the same time, the type of n is enriched with the capability to have *taxi* as a parent. This process, however, is restricted in that it provides a private taxi for one and only one ambient, namely n. By enabling communication of ambient names, types in name creation constructs may also refer to a received name, so boosting the dynamic aspects of the calculus. This allows, for instance, to write a *generic* taxi server which dynamically provides private taxis for any ambient whose name is received by a communication:

$$DTaxiServ = !(x)(\nu\, taxi : \mathsf{amb}[\,\mathsf{mob}[\{top\}, \{x\}]\,])\, (taxi[] \mid x[\mathsf{in}\, taxi]),$$

Such an increase in dynamic adaptation, and therefore in expressiveness, comes with a corresponding increase in the complexity of the typing system. There are two reasons:

- firstly, types may now refer to ambients indirectly, via variables;
- secondly, upon communication a variable may become instantiated by a name or by another, thus leading to different capabilities, that is different access control policies. Moreover, such policies are not necessarily comparable, which, for instance, makes subtyping useless.

Example 2. Let *DTaxiServ* be as above, and consider $P = \langle n \rangle \mid \langle m \rangle \mid DTaxiServ$. Process P may evolve to either

$$P_1 = \langle n \rangle \mid taxi[] \mid m[\mathsf{in}\, taxi] \mid DTaxiServ \quad \text{or} \quad P_2 = \langle m \rangle \mid taxi[] \mid n[\mathsf{in}\, taxi] \mid DTaxiServ$$

$$
\begin{array}{llll}
\textit{Abstract variables } \mathsf{x}, \mathsf{y}, \ldots \in \mathsf{X} & & \textit{Capability types} & K ::= \mathsf{cap}[\mathcal{P}] \\
\textit{Mobility types} & M, N ::= \mathsf{mob}[\mathcal{P}, C] & \textit{Communication types} & C ::= \mathsf{com}[\mathcal{E}, \mathcal{L}] \\
\textit{Ambient types} & A ::= \mathsf{amb}[M, C] & \textit{Variable types} & B ::= \mathsf{var}[\mathcal{B}]
\end{array}
$$

$$
\text{where } \mathcal{P}, C \subseteq \mathcal{N} \cup \mathsf{N} \cup \mathcal{X} \cup \mathsf{X} \cup \{\star\} \quad \text{and} \quad \mathcal{B}, \mathcal{E}, \mathcal{L} \subseteq \mathcal{N} \cup \mathsf{N} \cup \{\star\}
$$

Fig. 3. New types for communication

according to whether *DTaxiServ* receives m or n. As regards to the type assigned to *taxi*, in P_1 only m – that is the received name – is allowed to go inside *taxi*, whereas in P_2 only n is. Such situations are orthogonal and thus incomparable: in the first one, m acquires a new capability ('*CanMoveTo*' *taxi*) and n remains unchanged, whereas in the second it is the other way around. This means that the type system will have to deal with all possible communications. □

Example 2 suggests that an ambient name may have different possible types depending on the communications that may happen.

3.1 New types and definitions

The needed additional types are given in Figure 3, where X is an infinite set of *abstract variables*. Ambient types have now two components: one for mobility, which remains unchanged, and one for communication. Following our "symmetric style" for types, a communication type consists of two sets of names: an *external* one, $\mathcal{E}$, which describes the ambients where the name in object may be communicated, and a *local* one, $\mathcal{L}$, which describes the names which may be communicated within the object ambient. For instance, if n has type $\mathsf{amb}[M, \mathsf{com}[\{m\}, \{o\}]]$, then n can be exchanged inside m, and o can be communicated inside n; two terms which satisfy such specification are: $m[\langle n \rangle]$ and $n[\langle o \rangle]$.

A variable type is the set of ambients where a variable may be bound. For instance, in $n[(x)P]$ the variable x may have type $\mathsf{var}[\{n\}]$. Because we deal with an ambient calculus with the open capability, a variable may of course be bound in several places. Variable x in

$$n[\langle o_1 \rangle \mid \mathsf{open}\, m \mid m[\langle o_2 \rangle \mid (x)P]$$

may be bound in n or in m, depending on whether m is opened before x is bound or not. The type of x in the term above must therefore be $\mathsf{var}[\{m, n\}]$.

We remark that, contrary to usual type systems, but as in [1], we do not use the so-called *exchange* types, which build on the notion of "topic of conversation" permitted within a given ambient. This is a particular type (modulo subtyping) which all processes within that ambient agree to exchange. Indeed, suppose that n is an ambient whose exchange type is A, and m_1 and m_2 have type A. Then, the term

$$n[\langle m_1 \rangle \mid \langle m_2 \rangle \mid (x : A)P]$$

would be well-typed. However, if an external process gives a new capabilities to m_1 via name creation, then $\langle m_1 \rangle$ may become ill-typed. An instance of this is the following term.

$$(\nu\, o : \mathsf{amb}[\mathsf{mob}[\{m_1\}, \emptyset], C])\, Q \mid n[\langle m_1 \rangle \mid \langle m_2 \rangle \mid (x : A)P]$$

Here, $\langle m_2 \rangle$ is still well-typed, but $\langle m_1 \rangle$ is manifestly not. In fact, after scope extrusion, m_1 has not anymore type A, but A enriched with the capability to admit o as a child, which leads to a type error when trying to bind x to m_1. Our type system does not feature uniform exchange types, but it allows in principle to bind a variable to names which may have completely different types. This also leads to an increase in expressiveness.

Assuming the set of communicable ambient names to appear in types may at first appear to give a very restricted form of communication, where only finitely many names can be communicated in a given ambient. Due to the dynamic nature of types, this is actually not the case: name creation can give new communication capabilities to existing ambients. For instance, if ambient n has a communication type $\mathsf{com}[\emptyset, \{o\}]$ – that is, only o may be sent in n – the creation of name m with communication type $\mathsf{com}[\{n\}, \emptyset]$ gives n the new capability to have m has communicable ambient. Combining name creation with replication leads to potentially infinite communicable names in a given ambient.

Finally, we remark that variable and communication types, contain only ambient names and no variables. This means that communication is forbidden inside ambients created using a received name, as in $(x)x[P]$. We embraced such restriction for the sake of simplicity: allowing variables in such types is possible, at the price of a more complex definition of the set of types (Definition 7) and of additional side conditions in the type system's rules.

We introduce notations for communication types corresponding to those given for mobility types in (1). We define $\mathsf{amb}[M, C]^{\mathsf{com}} = C$, and then $C^{\uparrow}$, $C^{\downarrow}$ and $(\Gamma, \Delta)^{(n:A)}$, where $(x : B)^{(n:A)} = x : B$, as obvious. The notion of *symmetric type* is extended to communication types straightforwardly.

No variable occurs in symmetric types, therefore the notion of symmetric typing context must be defined accordingly.

Definition 3. A concrete typing context Γ is *symmetric* if, for all $n \in nm(\Gamma)$, we have $\Gamma(n) - (\mathsf{X} \cup \mathcal{X}) = \Gamma[n]$.

Here $\Gamma(n) - (\mathsf{X} \cup \mathcal{X})$ denotes the type of n in Γ with all variables removed. Definition 1 of global types for ambient names remains unchanged; we complete it below for variables.

Definition 4. The *global type* of a variable α with respect to typing contexts Γ and Θ is defined by

$$gt(\Gamma; \Theta)(\alpha) = \begin{cases} \Gamma[\alpha] \cup \Theta[\alpha] & \text{if } \alpha \in \mathcal{X} \\ \Theta[\alpha] & \text{if } \alpha \in \mathsf{X} \end{cases}$$

The global type of a variable is the ambient type constructed from the capabilities the variable is given by typing contexts. In case of an abstract variable, only the abstract context provides the information, as no abstract variable occurs in the concrete context.

$$\frac{\text{for any } \alpha,\ \text{if } lt(\Gamma;\Theta)(\alpha) = B \text{ then } a \in B \text{ iff } a_i \in B \qquad \Gamma \vdash^{\Theta} a : A, a_i : A_i \qquad A^{\mathsf{mob}} \subseteq A_i^{\mathsf{mob}} \qquad A^{\mathsf{com}} = A_i^{\mathsf{com}} \qquad \forall i \in \{1,\dots n\}}{\Gamma \vdash^{\Theta} \mathsf{open}\, a : \mathsf{cap}[\{a_1,\dots,a_n\}]}\ (\text{VOpen'})$$

$$\frac{\Gamma, x : B \vdash_a^{\Xi\{x/\mathsf{x}\};\Theta} P \qquad a \in B}{\Gamma \vdash_a^{\mathsf{x}:B,\Xi;\Theta} (x)P}\ (\text{Input}) \qquad \frac{\Gamma \vdash^{\Theta} V : A \qquad a \in A^{\mathsf{com}\uparrow}}{\Gamma \vdash_a^{\emptyset;\Theta} \langle V \rangle}\ (\text{Ouput})$$

Fig. 4. Revised typing system and additional ones for communication

We finally define the *local type* of a name, as the usual notion of type provided by a typing context.

Definition 5. The *local type* of a name α with respect to typing contexts Γ and Θ is defined by

$$lt(\Gamma;\Theta)(\alpha) = \begin{cases} \Gamma(\alpha) & \text{if } \alpha \in dom(\Gamma) \\ \Theta(\alpha) & \text{if } \alpha \in dom(\Theta) \\ \text{undefined otherwise.} \end{cases}$$

For $B = \mathsf{var}[\mathcal{B}]$, we usually identify B and $\mathcal{B}$ writing, for instance, $\alpha \in B$ instead of $\alpha \in \mathcal{B}$.

3.2 Revised typing system

The type system for the calculus with communication of ambient names is obtained from Figure 2 adding rules (Input) and (Ouput) for communication in Figure 4, and replacing (VOpen) with (VOpen').

Rule (VOpen') has two new conditions. The first verifies that the communication type is the same in the opened ambient and in the enclosing one. The condition on the top line checks that variables may be bound in the opened ambient if and only if they may in the opening one. Rule (Input) is similar to (Res): we pick up an abstract variable from the local abstract context whose type contains the current location. We do not perform any context update, since such updates only concern ambient names. Finally, rule (Output) says that we can send the value V inside a if that is allowed by the type of V.

The most significant revision is in the definition of $Types(\Gamma, \Theta, a)$, which is not anymore just a singleton, since names may now have several types. This means that any typing derivation with an axiom (VAmb) as a premise has to be read as: "for all types A such that $\Gamma \vdash^{\Theta} a : A$." Before we give the formal definition of $Types(\Gamma, \Theta, a)$, let us focus on an example to gather some intuition.

Example 3. Let us consider a non-replicated version of the taxi server of Example 2, but with the new communication types:

$$DTaxiServ_1 = (x)(\nu\, taxi : \mathsf{amb}[M, C])\,(taxi[] \mid x[\mathsf{in}\, taxi]),$$

with $M = \mathsf{mob}[\{top\}, \{x\}]$ and $C = \mathsf{com}[\emptyset, \emptyset]$. Let $P = \langle n \rangle \mid \langle m \rangle \mid DTaxiServ_1$, and suppose to work with the following typing context which assigns types to the free ambient names:

$$\Gamma = \begin{cases} n : A_n, \\ m : A_m, \\ top : \mathsf{amb}[\,\mathsf{mob}[\emptyset, \{m, n\}], \mathsf{com}[\emptyset, \{m, n\}]\,], \end{cases}$$

where

$$A_n = \mathsf{amb}[\,\mathsf{mob}[\{top\}, \{m\}], \mathsf{com}[\{top\}, \emptyset]\,]$$
$$A_m = \mathsf{amb}[\,\mathsf{mob}[\{top, n\}, \emptyset], \mathsf{com}[\{top\}, \emptyset]\,].$$

Here, n and m can be communicated in top, and have top as a parent. Moreover, n may have m as child and, consequently, m may have n as parent. One can verify that Γ is symmetric. In order to type $DTaxiServ_1$ we need an abstract context which assigns variable type $\mathsf{var}[\{top\}]$ to an abstract variable, say x, because x may be bound in top. Also, it has to assign an ambient type A to an abstract ambient name, say taxi. Type A has to match $\mathsf{amb}[M, C]$, but we cannot directly use the latter, since M contains x which is initially bound. However, we can refer to the abstract name x, meant to correspond to x. Therefore the abstract context is

$$\Xi = \mathsf{x} : \mathsf{var}[\{top\}], \mathsf{taxi} : \mathsf{amb}[N, C],$$

for $N = \mathsf{mob}[\{top\}, \{\mathsf{x}\}]$. It is easy to check that the proof of $\Gamma \vdash_{top}^{\Xi;\emptyset} DTaxiServ_1$ leads to the judgement $\Delta \vdash_{top}^{0;\emptyset} (taxi[\,] \mid x[\mathsf{in}\, taxi])$, by application of the rules (Input) and (Res), where

$$\Delta = \begin{cases} n : A_n, \\ m : A_m, \\ top : \mathsf{amb}[\,\mathsf{mob}[\emptyset, \{m, n, taxi\}], \mathsf{com}[\emptyset, \{m, n\}]\,], \\ x : \mathsf{var}[\{top\}], \\ taxi : \mathsf{amb}[\,\mathsf{mob}[\{top\}, \{x\}], \mathsf{com}[\emptyset, \emptyset]\,]. \end{cases}$$

In $(taxi[\,] \mid x[taxi])$, x is an ambient whose global *ambient* type is

$$A_x = \Delta[x] = \mathsf{amb}[\,\mathsf{mob}[\{taxi\}, \emptyset], \mathsf{com}[\emptyset, \emptyset]\,].$$

This is actually the least type assignable to x, since x is intended to be instantiated by an ambient name. Indeed, x may be bound to n or m, which means that the global types it may possibly assume include

$$A_x \cup A_n = \mathsf{amb}[\,\mathsf{mob}[\{top, taxi\}, \{m\}], \mathsf{com}[\{top\}, \emptyset]\,],$$

if n is received for x, and

$$A_x \cup A_m = \mathsf{amb}[\,\mathsf{mob}[\{top, taxi, n\}, \emptyset], \mathsf{com}[\{top\}, \emptyset]\,],$$

if m is received for x. Type $A_x \cup A_n$ is possible for n in the scope of the restriction. Indeed, if n is received, then the newly created ambient $taxi$ gives to n – the received

name – the capability to be its child. Outside the scope of the restriction, a possible type for n is rather

$$\Xi[\mathsf{x}] \cup A_n = \mathsf{amb}[\,\mathsf{mob}[\{top, \mathsf{taxi}\}, \{m\}], \mathsf{com}[\{top\}, \emptyset]\,].$$

Since *taxi* is not visible, we have to refer to its abstract representative taxi. Inspecting the various possibilities, we are therefore led to consider the following sets of possible ambient types.
(1) Outside the scope of *taxi*:

$$Types(\Gamma, \Xi, n) = \{\, A_n, A_n \cup \Xi[\mathsf{x}]\,\} \quad \text{and} \quad Types(\Gamma, \Xi, m) = \{A_m, A_m \cup \Xi[\mathsf{x}]\,\};$$

(2) Inside the scope of *taxi*:

$$Types(\Delta, \emptyset, n) = \{\, A_n, A_n \cup \Delta[x]\,\}, \qquad Types(\Delta, \emptyset, m) = \{\, A_m, A_m \cup \Delta[x]\,\}$$
$$\text{and} \quad Types(\Delta, \emptyset, x) = \{\, A_x \cup A_n, A_x \cup A_m\,\}.\square$$

3.3 Estimating the set of possible types

In this section we give the formal estimate of the set of possible types for name a with respect to chosen typing contexts. As illustrated by Example 3, the possible types of an ambient name n are given by its global type merged with those of the variables n might instantiate. Consider the term

$$R = !\langle n\rangle \mid \langle m\rangle \mid (x)P \mid (y)Q,$$

and suppose that n and m have respectively global types A_n and A_m. Then, n may instantiate variables x and y with global types, say, A_x and A_y respectively. The possible types of n are, therefore, A_n, $A_n \cup A_x$, $A_n \cup A_y$ and $A_n \cup A_x \cup A_y$. The latter corresponds to the case in which n instantiates both x and y. The possible types of m are A_m, $A_m \cup A_x$, $A_m \cup A_y$, and $A_m \cup A_x \cup A_y$. (Observe that the latter is actually impossible, since m may instantiate at most one variable; it is an approximation error introduced by our estimation of possible types.) Symmetrically, the possible types for x are $A_x \cup A_n$, $A_x \cup A_m$, $A_x \cup A_y \cup A_n$ and $A_x \cup A_y \cup A_m$.

We first define the set of names which a given name may possibly instantiate or be instantiated by. To this aim we use binary relations over names. Intuitively, if a name α can instantiate a variable β, then the pair (α,β) is in one such a relation. We focus exclusively on so-called binding relations, which only depend on types and typing contexts. Basically, this amounts to saying that α is related to β if β is a variable which might be bound in some ambient where α may be communicated, as formalised by the definition below.

Definition 6. The *binding relation* with respect to typing contexts Γ and Θ is

$$Bind(\Gamma;\Theta) = \{\, (\alpha,\beta) \mid gt(\Gamma;\Theta)(\alpha) = A \text{ and } gt(\Gamma;\Theta)(\beta) = B \text{ with } A^{\mathsf{com}\uparrow}\{\alpha/\star\} \cap B \neq \emptyset\,\}.$$

Observe that $Bind(\Gamma,\Theta)$ is actually a subset of $(\mathsf{N}\cup\mathcal{N})\times(\mathsf{X}\cup\mathcal{X})$. Indeed, due to our restrictions, the global (ambient) type of a variable is always empty – more precisely, it is always $\mathsf{com}[\emptyset,\emptyset]$. It follows that a variable cannot be related to another variable.

Relying on binding relations, we can estimate which variables may be instantiated by a given ambient name, and which ambient names may instantiate a given variable.

$$Vars(\Gamma;\Theta;\alpha) = \{\beta \mid (\alpha,\beta) \in Bind(\Gamma;\Theta)\};$$
$$Ambs(\Gamma;\Theta;\alpha) = \{\beta \mid (\beta,\alpha) \in Bind(\Gamma;\Theta)\}.$$

We use the notation $A \triangleright \mathcal{T}$, for $\mathcal{T}$ a set of ambient types, to denote the set $\{A\cup B \mid B \in \mathcal{T}\}$, and $A \triangleright^\star \mathcal{T}$ for $\{A\} \cup A \triangleright \mathcal{T}$.

Definition 7. The set $Types(\Gamma;\Theta;\alpha)$ of possible types for α with respect to typing contexts Γ and Θ is defined by

$$Types(\Gamma;\Theta;\alpha) = gt(\Gamma;\Theta)(\alpha) \triangleright^\star \left\{ gt(\Gamma;\Theta)(\beta) \;\middle|\; \beta \in Vars(\Gamma;\Theta;\alpha) \right\}, \quad \text{if } \alpha \in \mathsf{N}\cup\mathcal{N}, \text{ and}$$

$$Types(\Gamma;\Theta;\alpha) = \bigcup_{\beta\in Ambs(\Gamma;\Theta;\alpha)} gt(\Gamma;\Theta)(\alpha) \triangleright Types(\Gamma;\Theta;\beta), \qquad \text{if } \alpha \in \mathsf{X}\cup\mathcal{X}.$$

Observe that the constraints on our typing rules are expressed essentially as type inclusion constraints. Therefore, we do not really need to consider all the possible types of an ambient name. In practice, as well as in proofs, it is more convenient the use notions of maximum and minimum types which are respectively the union and the intersection of all possible types. Due to space limitation, we leave to the reader to reformulate our rules in such terms. For instance, rule (VAmb) can be removed, whilst (VIn) is rewritten by replacing A (standing for all possible types for a) by the minimum type for a.

The main result of this paper is a 'Subject Reduction' theorem for the type systems we introduced. Write $\Gamma \vdash_a P$ if there exist some abstract contexts Ξ and Θ such that $\Gamma \vdash_a^{\Xi;\Theta} P$, it can be expressed as follows.

Theorem 1 (Subject Reduction). If $\Gamma \vdash_n P$ and $P \to Q$, then $\Gamma \vdash_n Q$.

4 Conclusion, related and future work

The paper proposed a type system to describe access control policies and related behavioural properties of processes in the Ambient calculus. Whilst all type systems in the literature make use of groups to describe such properties [4, 12, 8], ours is based on dependent types. Despite an additional complexity in building typing derivations, our system has simple, natural types which may make policy specifications easier, and, therefore, simplify the programming task.

The major technical device of our approach is the novel notion of *abstract names*, which mimic the role of groups internally to typing derivations. More precisely, they keep track of – and dynamically embody in an ambient's type – any new capability that the ambient may gain during its lifespan. Abstract names are, in a sense, a dynamic notion of group, made internal to the type system rather than part of the language.

We showed how to extend a basic system to deal with communication of ambient names. The resulting type system is, we believe, the main contribution of this paper. It allows modular, per-client programming and is, to the best of our knowledge, the first one dealing at such a level with dynamic specifications of security policies in the ambient calculus. Communication types come together with a noticeable increase of complexity of typing derivations in the system. Indeed, when typing a communication, one has to consider all ambient names that can possibly instantiate a given variable. However, only a finite number of names are necessary for the proofs, and the notions of maximum and minimum types simplify the matter considerably.

We plan to study in future work the question of decidability of type checking, as a positive answer would make our type system usable in practice. We conjecture that it should be possible to devise a type checking algorithm, for the following reasons.

- The use of abstract contexts is not a concern: for a typeable term P, it is easy to provide (algorithmically or "by hand"), an abstract context Θ, such that $\Gamma \vdash_n^{\Theta;\emptyset} P$ for some Γ and n. Indeed, Θ is essentially the collection of bound names in P.
- Communication is quite hard to manage "manually" because it requires the synthesis of several global types. However, given variable types and communication types, it does not seem difficult to design an algorithm that does the job.

Because processes are typed against an environment – represented by the external abstract contexts – our type system is not compositional, in that $\Gamma \vdash_a P$ and $\Gamma \vdash_a Q$ does not imply $\Gamma \vdash_a P \mid Q$. This may be problematic in the framework of open (or partially typed) systems which have to deal with the arrival of unknown agents. Thus, the following questions arises: if $\Gamma \vdash_n^{\Xi;\emptyset} P$ is provable, does a class exist of external abstract contexts Θ such that $\Gamma \vdash_n^{\Xi;\Theta} P$? Such a class would determine the environments which make P typeable. Another question concerns the design of a type inference algorithm, which is a crucial component in open systems, where information coming from external sources cannot always be trusted. As studied in [11] for $D\pi$, dependent types call for a very accurate treatment, which would certainly be required for our type system as well. Such challenging questions are topics of our current and future research.

Related work. Beside the already mentioned work on groups [4, 12, 8], which provided direct inspiration for our research, related work includes [14, 10]. These papers introduce dynamic and dependent types for a distributed π calculus, and study their impact on behavioural semantics. A difficulty arises in [14] with referring to bound names created in external processes, at all analogous to the one we tackled here with the introduction of *abstract names*. As a matter of fact, it is pointed out in *loc. cit.* as the main limitation of their work. The problem has been addressed in [13], completely independently of our work, using existential types. There appear to be analogies between Yoshida's existential types and our approach, although at this stage it is difficult to assess them in the details.

Since our type system works linearly on internal abstract contexts, and classically on external ones, the closest match appears to be with Yoshida's linear type discipline. However, while we conjecture we can reformulate local abstract names with (some form

of) existential types, it looks as though they cannot provide a notion corresponding to our external abstract contexts. But such contexts are the keystones of our work. They allow, in particular, a correct treatment of parallel composition and replication (cf. Figure 2), where it is not always the case that if P is correct, so is $P \mid P$ or, a fortiori, $!P$. We plan to investigate this matter further.

Concerning typing systems for the ambient calculus, [1] uses dependent types for communication in order to achieve advanced type polymorphism of the sort usually encountered in lambda calculi. Types in *loc. cit.* track – like ours – all messages exchanged, and not rely on topics of conversation. They are also used to bound the nesting of ambients.

We conclude by observing that, as pointed out in [5] there are intriguing connections between groups and channels and binders of the flow analysis, as in [2, 9]. Indeed, our approach has a lot to share with control flow analysis, and we believe our work can shed further light on such connections.

References

1. T. Amtoft and J. Wells. Mobile processes with dependent communication types and singleton types for names and capabilities. Technical Report 2002-3, Kansas State University, 2002.
2. C. Bodei, P. Degano, F. Nielson, and H. R. Nielson. Static analysis for the π calculus with applications to security. *Information and Computation*, 168:68–92, 2001.
3. M. Bugliesi, G. Castagna, and S. Crafa. Boxed ambients. In *TACS'01*, volume 2215 of *Lecture Notes in Computer Science*, pages 38–63. Springer, 2001.
4. L. Cardelli, G. Ghelli, and A. Gordon. Ambient groups and mobility types. In *International Conference IFIP TCS*, volume 1872 of *Lecture Notes in Computer Science*, pages 333–347. Springer, 2000.
5. L. Cardelli, G. Ghelli, and A. Gordon. Secrecy and group creation. In *CONCUR'00*, volume 1877 of *Lecture Notes in Computer Science*, pages 365–379. Springer, 2000.
6. L. Cardelli and A. Gordon. Mobile ambients. In *FOSSACS'98*, volume 1378 of *Lecture Notes in Computer Science*, pages 140–155. Springer, 1998.
7. L. Cardelli and A. Gordon. Types for mobile ambients. In *POPL'99*, pages 79–92. ACM Press, 1999.
8. M. Coppo, M. Dezani-Ciancaglini, E. Giovannetti, and I. Salvo. M3: Mobility types for mobile processes in mobile ambients. In *CAT'03*, volume 78 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 2003.
9. P. Degano, F. Levi, and C. Bodei. Safe ambients: Control flow analysis and security. In *Proceedins of ASIAN'00*, volume 1961 of *LNCS*, pages 199–214. Springer-Verlag, 2000.
10. M. Hennessy, M. Merro, and J. Rathke. Towards a behavioural theory of access and mobility control in distributed systems. In *FOSSACS'03*, Lecture Notes in Computer Science. Springer, 2003.
11. C. Lhoussaine. Type inference for a distributed π-calculus. In *ESOP'03*, volume 2618 of *Lecture Notes in Computer Science*, pages 253–268. Springer, 2003.
12. M. Merro and V. Sassone. Typing and subtyping mobility in boxed ambients. In *CONCUR'02*, volume 2421 of *Lecture Notes in Computer Science*, pages 304–320. Springer, 2002.
13. N. Yoshida. Channel dependent types for higher-order mobile processes. In *POPL'04*, 2004. To appear.
14. N. Yoshida and M. Hennessy. Assigning types to processes. In *LICS'00*, pages 334–345, 2000.