

The Security Picalculus and Non-interference (Extended Abstract)

Matthew Hennessy^{1,2}

University of Sussex, UK

Abstract

The security π -calculus is a typed version of the asynchronous π -calculus in which the types, in addition to constraining the input/output behaviour of processes, have *security levels* associated with them. This enables us to introduce a range of typing disciplines which allow input or output behaviour, or both, to be bounded above or below by a given security level.

We define typed versions of *may* and *must* equivalences for the security π -calculus, where the tests are parameterised relative to a security level. We provide alternative characterisations of these equivalences in terms of *actions in context*; these describe the actions a process may perform, assuming the observer is constrained by a given typing environment.

Using these alternative characterisations we prove non-interference results with respect to *may* and *must* testing. These show that information flow between security levels can be controlled using our typing systems.

Keywords: Distributed Systems, picalculus, security types, non-interference, testing equivalences.

1 Introduction

The asynchronous π -calculus, [2,13], is a simple formalism for describing distributed processes. It presupposes a set of channel names through which processes communicate. Thus $\mathbf{a}?(X)P$ is a process which inputs some value v on the channel \mathbf{a} , and executes the body P in which X has been substituted by the value v , while output on the same channel is denoted by $\mathbf{a}!\langle v \rangle$. These two primitives, together with operators for parallelism, $|$, repetition, $*$, and channel scoping, $(\mathbf{new} n)$, make the π -calculus a very powerful language. For

¹ This research has been partially funded by the EU Global Computing projects *Mikado* and *Myths*

² Email: matthewh@cogs.sussex.ac.uk

example the term P ,

$$* \text{req?}(x, y) \text{ (new } r) \text{ s!}\langle x, r \rangle \mid r?(z) y!\langle z \rangle$$

describes a process which repeatedly receives a request on the channel req , consisting of a value, bound to x , and a return channel, bound to y . This value is in turn sent along the channel s , presumably serviced by some independent server, together with a private return channel r , generated specifically for this purpose. A response is awaited from the service, on the reply channel r , which is then forwarded on the original return channel y .

Numerous typing systems have been developed for this language, [17,23,24]. Most are based on judgements of the form $\Gamma \vdash P$ indicating that the process P is well-typed with respect to the channel environment Γ , which associates capabilities with the free channel names of P . Usually these capabilities are some elaboration of

read capabilities $r\langle T \rangle$: the ability to read values at type T from a channel
write capabilities $w\langle T \rangle$: the ability to write values at type T to a channel

For example let A denote the tuple type $(\text{int}, w\langle \text{int} \rangle)$; a value of this type will consist of a pair, the first element of which is an integer, and the second a channel on which integers may be written. If Γ associates the type $r\langle A \rangle$ with the channel req and the type $w\langle A \rangle$ with s , we would expect the above term, P , to be well-typed with respect to Γ . However for this to be true the local channel r needs to be generated with the write capability $w\langle \text{int} \rangle$, to be sent along the channel s , and the read capability $r\langle \text{int} \rangle$, which is used by the process itself. Thus if we were to annotate all bound names and variables with their required types we would obtain the annotated term

$$* \text{req?}(x, y) : A \text{ (new } r : R) \text{ s!}\langle x, r \rangle \mid r?(z) y!\langle z \rangle$$

where R is the type $\{w\langle \text{int} \rangle, r\langle \text{int} \rangle\}$, more usually written as $rw\langle \text{int} \rangle$. This term is well-typed with respect to the above mentioned environment Γ .

Intuitively the use of types constrains the behaviour of processes, ensuring no misuse of channels. By defining sophisticated forms of types process behaviour can be more or less constrained, while at the same time the advantages of well-typing can be preserved. For example a form of polymorphism is investigated in [17], while in [11] security levels are associated with capabilities, to obtain so-called *security types*. Suppose we have two security levels, high, denoted by top , and low, denoted by bot . Then we would have capabilities of the form $r_{\text{top}}\langle T \rangle$, $r_{\text{bot}}\langle T \rangle$, $w_{\text{top}}\langle T \rangle$, $w_{\text{bot}}\langle T \rangle$, where T in turn is a security type. By varying the precise definition of a security type we can either implement *resource access control* methodologies, or ensure forms of *non-interference*, [1,7,6]. In this paper we will be concerned with the latter, using a mild variation of the I -types of [11]; essentially types are sets of read/write capabilities, where in addition each capability is annotated by a security level taken from

some complete lattice . We will refer to the asynchronous π -calculus, augmented with these types, as the *security π -calculus*.

The statement of non-interference results requires some definition of *process behaviour*; intuitively a system is interference-free if its low level behaviour is independent of changes to high-level behaviour. The main topic of this paper is an investigation of the notion of behaviour of process, relative to a security level, for the security π -calculus.

Process behaviour is relative to some typing environment Γ and therefore we wish to develop a relation of the form $\Gamma \triangleright^\sigma P \simeq Q$ meaning, intuitively, that in the typing environment Γ , both P and Q exhibit the same σ -level behaviour. By this we mean that a σ -level observer will be unable to discern a difference between P and Q . For example low-level observers will be unable to see any high-level actions performed by P , Q . But more importantly we assume that these observers are constrained by the typing environment Γ and therefore actions disallowed by this environment will also be invisible to observers.

In this paper we investigate *may* and *must* testing equivalences, [15,8], for the security π -calculus. In particular we give an alternative characterisation of these behavioural equivalences which, as might be expected from [15,8], are based on the (asynchronous) sequences of actions that a process can perform. But here these sequences are relative to both a security level and a typing environment.

These actions, which we call *actions in context*, take the form

$$\Gamma \triangleright P \xrightarrow{\mu}_\sigma \Gamma' \triangleright P'$$

indicating that the process P can perform the action μ to interact with some σ -level observer constrained by the environment Γ ; this action affects the process P but may also change the the typing environment of the observing process, from Γ to Γ' .

The remainder of this extended abstract is organised as follows. In Section 2 we formally define the types we use, together with the syntax of the security π -calculus. This is followed with a range of typing systems. In the most straightforward we have the judgements $\Gamma \vdash P$ where Γ is a type environment, associating types to channel names and variables. This means that relative to Γ , P uses its channels correctly as input/output devices, ignoring their security annotations. We also have judgements of the form $\Gamma \vdash_\sigma P$ which indicates that in addition P uses channels with security level *at most* σ . Similarly we have a typing relation $\Gamma \vDash P$ meaning that P uses channels with *at least* security level σ . Indeed we can go further, designing relations such as $\Gamma \vdash_{r,\sigma} P$ or $\Gamma \vdash_{w,\sigma} P$ where the *read capabilities* or the *write capabilities* of processes are independently constrained. For all of these typing relations Subject Reduction is easily established.

Section 3 is the heart of the paper. First the behavioural preorders and equivalences are defined, by adapting the standard framework, [15,8], to the

security π -calculus. We obtain the relations

$$\Gamma \triangleright_{\sigma} P \simeq_{may} Q \quad \text{and} \quad \Gamma \triangleright_{\sigma} P \simeq_{must} Q$$

indicating that P and Q can not be distinguished, relative to *may/must* experiments respectively, by any testing process T such that $\Gamma \vdash_{\sigma} T$; that is any test running at security level *at most* σ , relative to the type environment. This is followed by an exposition of the actions in context, sequences of which are used the alternative characterisation of \simeq_{may} ; we also state a much more complicated characterisation of \simeq_{must} .

One benefit of having behavioural equivalences relativised to security levels is that non-interference results can be stated succinctly. Section 4 contains two such statements. The first gives conditions ensure that

$$\Gamma \triangleright_{\sigma} P \simeq_{may} Q \text{ implies } \Gamma \triangleright_{\sigma} P \mid H \simeq_{may} Q \mid K.$$

It turns out to be sufficient to require that the *read capabilities* of P and Q be bounded above by σ , that is $\Gamma \vdash_{\sigma} P, Q$, and that the *write capabilities* of H and K be bounded below by some $\delta \not\leq \sigma$, that is $\Gamma \vdash^{\omega\delta} H, K$.

This is quite a general non-interference result. For example in the case where Q is P and K is the empty process $\mathbf{0}$ we obtain

$$\Gamma \triangleright_{\sigma} P \simeq_{may} P \mid H$$

indicating that, under the conditions of the theorem, the process H can not interfere with the behaviour of P .

This result is not true for the *must* equivalence. As explained in Section 4, this is because our types allow contention between processes running at different security levels over read access to channels. However by restricting the type system, to confine *read capabilities* to a single security level, we show that a similar result holds for \simeq_{must} .

In this extended abstract all proofs are omitted. They may be found in the full version of the paper, [9].

2 The Language

We presuppose a complete lattice $\langle SL, \preceq, \sqcap, \sqcup, \mathbf{top}, \mathbf{bot} \rangle$ of security annotations, ranged over by σ, ρ, \dots . For each σ we assume a set of *basic types* at that level, of the form \mathbf{B}_{σ} . If the security annotation is omitted, as in \mathbf{int} , then we assume it has security level \mathbf{bot} ; as we shall see values of these types are available to all processes. Also, as explained in the Introduction, a σ -level channel type, for channels accessible to processes with security clearance at level σ , consists of a set of σ -level capabilities, i.e. a subset of Cap_{σ} . These may either be a read capability, of the form $r_{\rho}\langle T \rangle$, for some appropriate ρ and T , or a write capability, of the form $w_{\sigma}\langle T \rangle$. These capabilities are constrained by

Fig. 1 Types

(RT-BASE)		
$\frac{}{\mathbf{B}_\rho \in \text{Type}_\sigma} \quad \rho \preceq \sigma$		
(RT-WR)	$\frac{A \in \text{Type}_\sigma}{\mathbf{w}_\sigma \langle A \rangle \in \text{Cap}_\sigma}$	(RT-RD)
		$\frac{A \in \text{Type}_\rho}{\mathbf{r}_\rho \langle A \rangle \in \text{Cap}_\sigma} \quad \sigma \preceq \rho$
(RT-WRRD)	$\frac{S \subseteq_{fin} \text{Cap}_\sigma}{S \in \text{Type}_\sigma} \quad S \text{ consistent}$	(RT-TUP)
		$\frac{A_i \in \text{Type}_\sigma \quad (\forall i)}{(A_1, \dots, A_k) \in \text{Type}_\sigma}$
(U-WR)	$\mathbf{w}_\sigma \langle A \rangle <: \mathbf{w}_\sigma \langle B \rangle$	if $B <: A$
(U-RD)	$\mathbf{r}_\sigma \langle A \rangle <: \mathbf{r}_\sigma \langle B \rangle$	if $A <: B$
(U-BASE)	$\mathbf{B}_\sigma <: \mathbf{B}_\rho$	if $\sigma \preceq \rho$
(U-RES)	$\{cap_i\}_{i \in I} <: \{cap'_j\}_{j \in J}$	if $(\forall j)(\exists i) cap_i <: cap'_j$
(U-TUP)	$(A_1, \dots, A_k) <: (B_1, \dots, B_k)$	if $(\forall i) A_i <: B_i$

The set of capabilities Cap is *consistent* if

- $\mathbf{w}_\sigma \langle A \rangle, \mathbf{w}_\rho \langle B \rangle \in \text{Cap}$ implies $\sigma = \rho$ and A is B
 - $\mathbf{r}_\sigma \langle A \rangle, \mathbf{r}_\rho \langle B \rangle \in \text{Cap}$ implies A is B
 - $\mathbf{w}_\sigma \langle A \rangle, \mathbf{r}_\rho \langle B \rangle \in \text{Cap}$ implies $A <: B$.
-

consistency requirements. For example since values with the capability $\mathbf{w}_\sigma \langle T \rangle$ are written to by σ -level processes we require that T in turn be a σ -level type.

Types, i.e. sets of capabilities, are also constrained. For simplicity in a given type we only allow at most one write capability, and for each level σ at most one read capability at that level. More importantly we ensure that, relative to security levels, only *write-ups*, [7,1], are allowed by requiring that if $\mathbf{w}_\rho \langle T \rangle$ and $\mathbf{r}_{\rho'} \langle S \rangle$ are in a type then $\rho \preceq \rho'$; the additional constraint that T be a sub-type of S is well-known [17,12]. The formal definitions of types and capabilities, together with their *subtyping* relation are given in Figure 1. These types correspond very closely to the *I-types* of [11]; the rule (RT-RD) ensures that only *write-ups* are allowed, from low-level processes to high-level processes. But we allow multiple read capabilities, which will enable us to be more flexible with respect to allowing/disallowing reading from a channel at different security levels. However subtyping is more restrictive; unlike [11] they can only be sub-typed at the same security level; $\mathbf{r}_\sigma \langle A \rangle <: \mathbf{r}_\rho \langle B \rangle$ only if $\sigma = \rho$. Nevertheless this is compensated for in the existence of multiple read capabilities.

Fig. 2 Syntax

$P, Q ::=$	<i>Terms</i>	$X, Y ::=$	<i>Patterns</i>
$u!\langle v \rangle$	Output	x	Variable
$u?(X : A) P$	Input	(X_1, \dots, X_k)	Tuple
if $u = v$ then P else Q	Matching		
$(\text{new } a : A) P$	Name creation	$u, v, w ::=$	<i>Values</i>
$P \mid Q$	Composition	bv	Base Value
$*P$	Replication	a	Name
$\mathbf{0}$	Termination	x	Variable
		(u_1, \dots, u_k)	Tuple

Example 2.1

- The set $\{\mathbf{w}_{\text{bot}}\langle \mathbf{int} \rangle, \mathbf{r}_{\text{bot}}\langle \mathbf{int} \rangle, \mathbf{r}_{\text{top}}\langle \mathbf{int} \rangle\}$ is a **bot**-level channel type, an element of $Type_{\text{bot}}$; that is channels of this type may be transmitted on **bot**-level channels. In turn these channels may be written to by a **bot**-level process or read by either a **bot**-level or a **top**-level process.
- The type $\{\mathbf{w}_{\text{bot}}\langle \mathbf{int} \rangle, \mathbf{r}_{\text{top}}\langle \mathbf{int} \rangle\}$ restricts reading from the channel to **top**-level processes, although **bot**-level ones can write to it.
- The set $\{\mathbf{w}_{\text{top}}\langle \mathbf{int} \rangle, \mathbf{r}_{\text{bot}}\langle \mathbf{int} \rangle, \mathbf{r}_{\text{top}}\langle \mathbf{int} \rangle\}$ is not a valid type as it contains a read capability at a lower level than its write capability.
- The set $\{\mathbf{w}_{\text{top}}\langle \mathbf{int} \rangle, \mathbf{r}_{\text{top}}\langle \mathbf{int} \rangle\}$ is a **top**-level type but not a **bot**-level one; that is, it is in $Type_{\text{top}}$ but not in $Type_{\text{bot}}$.

Proposition 2.2 For every σ , $Type_{\sigma}$ is a preorder with respect to \leq , with both a partial meet operation \sqcap and a partial join \sqcup .

Multiple read capabilities in a type, such as $\{\mathbf{w}_{\text{bot}}\langle \mathbf{int} \rangle, \mathbf{r}_{\text{bot}}\langle \mathbf{int} \rangle, \mathbf{r}_{\text{top}}\langle \mathbf{int} \rangle\}$, allows processes at different security levels to read from the same channel. We can eliminate such contention by using a restricted set of types.

Definition 2.3 [Single-level Types] Let $SlType$ be the least set of types obtained by changing the condition on read capabilities in the definition Figure 1 to read:

$$\mathbf{r}_{\rho}\langle A \rangle, \mathbf{r}_{\sigma}\langle B \rangle \in \text{Cap} \text{ implies } \rho = \sigma \text{ and } A \text{ is } B.$$

Note that these types still allow communication from low-level processes to high-level processes. We leave the reader to check that these types, ordered by \leq also has both partial meet and join operations.

The syntax of the (asynchronous) π -calculus, given in Figure 2, uses a predefined set of *names*, ranged over by a, b, \dots and a set of *variables*, ranged over by x, y, z . *Identifiers* are either variables or names. We also assume a set

Fig. 3 Typing Rules

$\frac{(\text{T-ID}) \quad \Gamma(u) \prec: A}{\Gamma \vdash u : A}$	$\frac{(\text{T-BASE}) \quad bv \in \mathbf{B}_\sigma}{\Gamma \vdash bv : \mathbf{B}_\sigma}$	$\frac{(\text{T-TUP}) \quad \Gamma \vdash v_i : A_i \quad (\forall i)}{\Gamma \vdash (v_1, \dots, v_k) : (A_1, \dots, A_k)}$
$\frac{(\text{T-IN}) \quad \Gamma, X : A \vdash P}{\Gamma \vdash u : r_\sigma \langle A \rangle}$	$\frac{(\text{T-OUT}) \quad \Gamma \vdash u : w_\sigma \langle A \rangle}{\Gamma \vdash v : A}$	$\frac{(\text{T-EQ}) \quad \Gamma \vdash u : A, v : B}{\Gamma \vdash Q}$
$\frac{\Gamma \vdash u : r_\sigma \langle A \rangle}{\Gamma \vdash u ? (X : A) P}$	$\frac{\Gamma \vdash v : A}{\Gamma \vdash u ! \langle v \rangle}$	$\frac{\Gamma \sqcap \{u : B, v : A\} \vdash P}{\Gamma \vdash \text{if } u = v \text{ then } P \text{ else } Q}$
	$\frac{(\text{T-NEW}) \quad \Gamma, a : A \vdash P}{\Gamma \vdash (\text{new } a : A) P}$	$\frac{(\text{T-STR}) \quad \Gamma \vdash P, Q}{\Gamma \vdash P \mid Q, *P, \mathbf{0}}$

of *basic values*, ranged over by bv , each of which belong to a given basic type. We assume the standard notions of free names and variables, $\text{fn}(P)$ and $\text{fv}(P)$, respectively, and associated notions of substitution and α -equivalence, \equiv_α , are defined as usual. Moreover the typing annotations on the binding constructs are omitted whenever they do not play a role, as will most occurrences of the empty process $\mathbf{0}$. We also assume the standard reduction semantics for *closed* terms, usually referred to as *processes*. This is expressed in terms of a binary relation over processes, defined inductively via judgements of the form

$$P \xrightarrow{\tau} Q$$

See the full version [9] for details, or indeed any standard reference to the *piccalculus* such as [21].

A *type environment* is a finite mapping from identifiers to types. We adopt some standard notation for describing them. For example, for any identifier u let $\Gamma, u : A$ denote the obvious extension of Γ ; $\Gamma, u : A$ is only defined if u is not in the domain of Γ . The subtyping relation \prec : together with the partial operators \sqcap and \sqcup may also be extended to environments. For example $\Gamma \prec: \Delta$ if for all u in the domain of Δ , $\Gamma(u) \prec: \Delta(u)$. We will normally abbreviate the simple environment $\{u : A\}$ to $u : A$ and moreover use $v : A$ to denote its obvious generalisation to arbitrary values v ; this is only well-defined when v has the same structure as the type A .

The first typing system is given in Figure 3, where the judgements take the form

$$\Gamma \vdash P$$

Intuitively this means that the process P uses all channels as input/output devices in accordance with their types, as given in Γ . It is the standard typing system for the π -calculus, [17], adapted to our types; note that the security levels on the capabilities do not play any role.

Fig. 4 Security Typing Rules

$\frac{\text{(LT-IN)} \quad \Gamma, X : A \vdash_{\sigma} P \quad \Gamma \vdash u : r_{\delta}\langle A \rangle}{\Gamma \vdash_{\sigma} u?(X : A) P} \quad \delta \preceq \sigma$	$\text{(LT-OUT)} \quad \frac{\Gamma \vdash v : A \quad \Gamma \vdash u : w_{\delta}\langle A \rangle}{\Gamma \vdash_{\sigma} u!\langle v \rangle} \quad \delta \preceq \sigma$	$\text{(LT-EQ)} \quad \frac{\Gamma \vdash u : A, v : B \quad \Gamma \vdash_{\sigma} Q}{\Gamma \sqcap \{u : B, v : A\} \vdash_{\sigma} P} \quad \Gamma \vdash_{\sigma} \text{if } u = v \text{ then } P \text{ else } Q$
$\text{(LT-NEW)} \quad \frac{\Gamma, a : A \vdash_{\sigma} P}{\Gamma \vdash_{\sigma} (\text{new } a : A) P}$	$\text{(LT-STR)} \quad \frac{\Gamma \vdash_{\sigma} P, Q}{\Gamma \vdash_{\sigma} P \mid Q, *P, \mathbf{0}}$	

We can also design a type inference system which not only ensures that channels are used according to their types but also controls the security levels of the channels used. One such system is given in Figure 4, where the judgements now take the form

$$\Gamma \vdash_{\sigma} P$$

This indicates that not only is P well-typed as before but in addition it uses channels with security level *at most* σ . (This corresponds to the typing system used in [11].) The only difference is in the input/output rules, where the security level of the channels used are checked. For example $\Gamma \vdash_{\sigma} a!\langle v \rangle$ only if in Γ the channel a can be assigned a security level $\delta \preceq \sigma$, in addition to having the appropriate output capability in Γ .

We can also design a typing system

$$\Gamma \vdash^{\sigma} P$$

which ensures that P uses channels with security level *at least* σ . The only change is to demand in the input/output rules that $\sigma \preceq \delta$:

$$\begin{array}{c} \text{(HL-IN)} \\ \Gamma, X : A \vdash^{\sigma} P \\ \Gamma \vdash u : r_{\delta}\langle A \rangle \\ \hline \Gamma \vdash^{\sigma} u?(X : A) P \quad \sigma \preceq \delta \end{array} \qquad \begin{array}{c} \text{(HL-OUT)} \\ \Gamma \vdash u : w_{\sigma}\langle A \rangle \\ \Gamma \vdash v : A \\ \hline \Gamma \vdash^{\sigma} u!\langle v \rangle \quad \sigma \preceq \delta \end{array}$$

We can provide further mix and matches. For example the type system

$$\Gamma \vdash_{r\sigma} P$$

ensures that all channels from which values are *read* have a read capability of *at most* σ ; the security level of the output channels is unexamined. This system is obtained by using the rules in the original Figure 3 but with the rule (T-IN) replaced with (LT-IN); the output rule is left unchanged. In a similar manner we can define relations $\Gamma \vdash_{w\sigma} P$, $\Gamma \vdash^{\sigma} P$ and $\Gamma \vdash^{w\sigma} P$.

Theorem 2.4 (Subject Reduction) *Let \Vdash denote any of the relations \vdash , \vdash_σ , $\vdash_{\tau\sigma}$, \vdash^σ , $\vdash_{w\sigma}$, $\vdash^{w\sigma}$, and suppose $\Delta \Vdash P$. Then $P \xrightarrow{\tau} Q$ implies $\Delta \Vdash Q$.*

3 Behavioural Theories

A *test* or *observer* is a process with an occurrence of a new reserved resource name ω , used to report success. We let T to range over tests, with the typing rule $\Gamma \vdash_\sigma \omega!\langle \rangle$ for all Γ . When placed in parallel with a process P , a test may interact with P , producing an output on ω if some desired behaviour of P has been observed. We write

$$P \text{ may } T$$

if $T \mid P \xrightarrow{\tau}^* R$ for some R such that R can *report success*, i.e. $R \xrightarrow{\omega!\langle \rangle}$. The stronger relation

$$P \text{ must } T$$

holds when in every computation

$$T \mid P \xrightarrow{\tau} R_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} R_n \xrightarrow{\tau} \dots$$

there is some R_k , $k \geq 0$, which can report success.

We can obtain a testing based behavioural preorder between processes by demanding that they react in a similar manner to a given class of tests. Here we choose the class of tests which are well-typed and use channels from *most* a given security level σ ; that is we require that processes react in the same manner to all tests T such that $\Gamma \vdash_\sigma T$.

Definition 3.1 [Testing Preorders] We write $\Gamma \triangleright_\sigma P \sqsubseteq_{\text{may}} Q$ if for every test T such that $\Gamma \vdash_\sigma T$, $P \text{ may } T$ implies $Q \text{ may } T$.

Similarly $\Gamma \triangleright_\sigma P \sqsubseteq_{\text{must}} Q$ means that for every such T , $P \text{ must } T$ implies $Q \text{ must } T$.

We use \simeq_{may} and \simeq_{must} denote the related equivalence relations.

So for example setting σ to be **bot**, $\Gamma \triangleright_{\text{bot}} P \simeq_{\text{may}} Q$ means that in the type environment Γ , P and Q are indistinguishable by low-level observers, from a may testing point of view.

The aim of this section is to outline alternative characterisations of these equivalences. It is well-known, [15,8], that testing equivalences are closely related to the ability of processes to perform sequences of actions. We have explained in the Introduction that here we need to relativise these sequences to security levels and a typing environments for the observers.

The rules for the Context LTS are given in Figure 5. The judgements take the form

$$\Gamma \triangleright P \xrightarrow{\mu}_\sigma \Gamma' \triangleright P' \tag{1}$$

This judgement should be understood as expressing the fact that:

Fig. 5 Context LTS

$$\begin{array}{c}
\text{(C-OUT)} \\
\frac{\mathbf{r}_\delta \langle A \rangle \in \Gamma(a)}{\Gamma \triangleright a! \langle v \rangle \xrightarrow{a!v}_\sigma \Gamma \sqcap v : A \triangleright \mathbf{0}} \quad \delta \preceq \sigma \\
\\
\text{(C-IN)} \\
\frac{\mathbf{w}_\delta \langle A \rangle \in \Gamma \\ \Gamma \vdash v : A}{\Gamma \triangleright a?(X : B) P \xrightarrow{a?v}_\sigma P \{v/X\}} \quad \delta \preceq \sigma \\
\\
\text{(C-OPEN)} \\
\frac{\Gamma, b : \top \triangleright P \xrightarrow{(\tilde{c})a!v}_\sigma \Gamma' \triangleright P'}{\Gamma \triangleright (\mathbf{new} b : B) P \xrightarrow{(b)(\tilde{c})a!v}_\sigma \Gamma' \triangleright P'} \quad \begin{array}{l} b \neq a, \tilde{c} \\ b \in \text{fn}(v) \end{array} \\
\\
\text{(C-WEAK)} \\
\frac{\Gamma, b : B \triangleright P \xrightarrow{(\tilde{c} : \tilde{C})a?v}_\sigma \Gamma' \triangleright P'}{\Gamma \triangleright P \xrightarrow{(b : B \tilde{c} : \tilde{C})a?v}_\sigma \Gamma' \triangleright P'} \quad b \neq a, \tilde{c} \\
\\
\text{(C-RED)} \\
\frac{P \xrightarrow{\tau} P'}{\Gamma \triangleright P \xrightarrow{\tau}_\sigma \Gamma \triangleright P'} \\
\\
\text{(C-CTXT)} \\
\frac{\Gamma \triangleright P \xrightarrow{\mu}_\sigma \Gamma' \triangleright P'}{\Gamma \triangleright *P \xrightarrow{\mu}_\sigma \Gamma' \triangleright *P \mid P'} \\
\\
\frac{\Gamma \triangleright P \xrightarrow{\mu}_\sigma \Gamma' \triangleright P'}{\Gamma \triangleright P \mid Q \xrightarrow{\mu}_\sigma \Gamma' \triangleright P' \mid Q} \quad \text{bn}(\mu) \not\subseteq \text{fn}(Q) \\
\Gamma \triangleright Q \mid P \xrightarrow{\mu}_\sigma \Gamma' \triangleright Q \mid P' \\
\\
\frac{\Gamma, a : A \triangleright P \xrightarrow{\mu}_\sigma \Gamma', a : A \triangleright P'}{\Gamma \triangleright (\mathbf{new} a : A) P \xrightarrow{\mu}_\sigma \Gamma' \triangleright (\mathbf{new} a : A) P'} \quad a \notin \text{n}(\mu)
\end{array}$$

The process P , when run concurrently with any observing process T such that $\Gamma \vdash_\sigma T$, can perform the action μ . This will transform P into P' and may also transform the type environment of the observing process to Γ' .

These actions can take three forms:

internal move: $\Gamma \triangleright P \xrightarrow{\tau}_\sigma \Gamma \triangleright P'$ This corresponds to an internal move by P , which does not depend on its environment. These moves are completely determined by the reduction semantics.

input move: $\Gamma \triangleright P \xrightarrow{(\tilde{c} : \tilde{C})a?v}_\sigma \Gamma' \triangleright P'$ Here the observing process sends a value v , possibly constructed using the new values \tilde{c} at type \tilde{C} , to P ; the type environment of the observing process will be augmented with these new values. An appropriate write capability on a is required of the observing

Fig. 6 (asynchronous) Traces in Context

$$\begin{array}{c}
\text{(TR-}\tau\text{)} \\
\frac{\Gamma \triangleright P \xrightarrow{\tau}_\sigma \Gamma' \triangleright P' \quad \Gamma' \triangleright P' \xrightarrow{s}_\sigma \Gamma'' \triangleright P''}{\Gamma \triangleright P \xrightarrow{s}_\sigma \Gamma'' \triangleright P''} \\
\text{(TR-}\epsilon\text{)} \\
\frac{}{\Gamma \triangleright P \xrightarrow{\epsilon}_\sigma \Gamma \triangleright P} \\
\text{(TR-}\alpha\text{)} \\
\frac{\Gamma \triangleright P \xrightarrow{\alpha}_\sigma \Gamma' \triangleright P' \quad \Gamma' \triangleright P' \xrightarrow{s}_\sigma \Gamma'' \triangleright P''}{\Gamma \triangleright P \xrightarrow{\alpha \cdot s}_\sigma \Gamma'' \triangleright P''} \\
\text{(TR-ASYNC-IN)} \\
\frac{\mathbf{w}_\delta \langle A \rangle \in \Gamma \quad \Gamma \vdash v : A}{\Gamma \triangleright P \xrightarrow{a?v}_\sigma \Gamma \triangleright P \mid a! \langle v \rangle} \delta \preceq \sigma \\
\text{(TR-ASYNC-WEAK)} \\
\frac{\Gamma, b : B \triangleright P \xrightarrow{(\tilde{c} : \tilde{C})a?v}_\sigma \Gamma' \triangleright P'}{\Gamma \triangleright P \xrightarrow{(b : B \tilde{c} : \tilde{C})a?v}_\sigma \Gamma' \triangleright P'} \quad b \neq a, \tilde{c}
\end{array}$$

process for the action to take place; see the rule (C-IN).

output move: $\Gamma \triangleright P \xrightarrow{(\tilde{c})a!v}_\sigma \Gamma' \triangleright P'$ Here P sends a value v , constructed using new values \tilde{c} , along the channel a to the observing process; typically the observers type environment Γ will be augmented with knowledge of v . This is implemented via the rules (C-OUT) and (C-OPEN), which uses the *top* type \top to denote the type consisting of the empty set of capabilities; note that this dominates all channel types in the subtyping relation.

In the judgements (1) above it should not be assumed that P can be typed in Γ ; indeed in the rules of Figure 5 it is not even assumed that the process terms are even typeable. Intuitively we expect the process P to be well typed in some environment Δ and that Γ represents that part of Δ which is known to the user. This motivates the following definition.

Definition 3.2 The pair $\Gamma \triangleright P$ is said to be a *configuration* if there exists an environment Δ such that

- $\Delta <: \Gamma$
- $\text{domain}(\Delta) = \text{domain}(\Gamma)$
- $\Delta \vdash P$

One can check that configurations are preserved by the relations defined in Figure 5; in the sequel we will refer to such judgements, applied to configurations, as *actions in context*.

It is worthwhile noting that in actions in context, (1) above, the new observers environment Γ' is not in general determined by the initial configuration $\Gamma \triangleright P$, unlike in [10]. This arises because of the rule (L-OUT) in Figure 5. In

general $\Gamma(a)$ may contain two read capabilities, $r_{\delta_1}\langle A_1 \rangle$ and $r_{\delta_2}\langle A_2 \rangle$, in which case Γ' may take either of the forms $\Gamma \sqcap v : A_1$ or $\Gamma \sqcap v : A_2$. However by restricting ourselves to single-level types this problem does not arise.

We say Γ is a *single-level* environment if it only uses single-level types. For such environments we can define the partial predicate $\Gamma \text{ after}_\sigma s$ with the property that in every action in context, (1) above, if Γ is *single-level* then Γ' must be $\Gamma \text{ after}_\sigma s$; see the full version [9] for details.

It is well-known that *may testing* is related to the ability of processes to perform sequences of actions, but in an asynchronous language we must consider some form of asynchronous sequences. In the present context this means generalising *actions in context* to (*asynchronous*) *traces in context* :

Definition 3.3 [Traces] Let $\Gamma \triangleright P \xrightarrow{s}_\sigma \Gamma' \triangleright P'$ be the least relation which satisfies the rules given in Figure 6.

Theorem 3.4 (Alternative Characterisation of May Testing) *Suppose $\Gamma \triangleright P$ and $\Gamma \triangleright Q$ are configurations. Then $\Gamma \triangleright_\sigma P \sqsubseteq_{\text{may}} Q$ if and only if $\Gamma \triangleright P \xrightarrow{s}_\sigma$ implies $\Gamma \triangleright Q \xrightarrow{s}_\sigma$.*

The extra ingredients required to capture must testing, in addition to traces, are well-known from [15,8]; they include a *convergence predicate*, indicating that a process has no internal infinite computations, and *acceptance sets*, indicating the next possible actions in which a process can engage. Here these need to be generalised from processes to configurations; they must also be relativised to security levels.

Definition 3.5 [Convergence] We say the configuration \mathcal{C} *converges*, written $\mathcal{C} \Downarrow$, if there is no infinite sequence of derivations

$$\mathcal{C} \xrightarrow{\tau} \mathcal{C}_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} \mathcal{C}_k \xrightarrow{\tau}$$

This relation is then parameterised to *sequences in context* and security levels by

$$\begin{aligned} \varepsilon : \mathcal{C} \Downarrow^\sigma \varepsilon & \text{ if } \mathcal{C} \Downarrow \\ s = \alpha \cdot s' : \mathcal{C} \Downarrow^\sigma s & \text{ if } \mathcal{C} \Downarrow \text{ and whenever } \mathcal{C} \xrightarrow{\alpha}_\sigma \mathcal{C}', \mathcal{C}' \Downarrow^\sigma s'. \end{aligned}$$

Note that for a configuration $\Gamma \triangleright P$ whether or not it converges is actually independent of the typing environment Γ . However convergence relative to a *sequence in context* is in general dependent on these environments.

We now adapt the definition of *Acceptance sets*, [8], to the security π -calculus. First let

$$\mathcal{O}^\sigma(\mathcal{C}) = \{ a! \mid \exists v. \mathcal{C} \xrightarrow{a!v}_\sigma \}$$

and

$$\mathcal{R}^\sigma(\mathcal{C}) = \{ a? \mid \exists v. \mathcal{C} \xrightarrow{a?v}_\sigma \} \cup \mathcal{O}^\sigma(\mathcal{C}).$$

Definition 3.6 [Acceptance sets] For a configuration \mathcal{C} , let $A^\sigma(\mathcal{C}, s)$, its σ -

level acceptance set after s , be defined by

$$\{ \mathcal{R}^\sigma(\mathcal{C}') \mid \mathcal{C} \xrightarrow{s}_\sigma \mathcal{C}' \not\rightarrow \}$$

Similarly let its *output* acceptance set after s be given by

$$\{ \mathcal{O}^\sigma(\mathcal{C}') \mid \mathcal{C} \xrightarrow{s}_\sigma \mathcal{C}' \not\rightarrow \}$$

Note only acceptance sets from *stable* configurations, configurations \mathcal{C}' such that $\mathcal{C}' \not\rightarrow$, are used.

The security π -calculus is *asynchronous* and therefore, as explained in [3], acceptance sets are too discriminating, when used to characterise *must* testing; to see this it is sufficient to consider the simple example

$$a?(x) \mathbf{0} \sqsubseteq_{\text{must}} \mathbf{0}.$$

The same reference goes on to explain that the use of *output* acceptance sets must also be relativised to sets of input actions, which we now explain.

Input Completions. We use $I^\sigma(\mathcal{C})$ to denote the set of input actions which the configuration \mathcal{C} can perform at level σ , $\{a?v \mid \mathcal{C} \xrightarrow{a?v}_\sigma\}$. More generally we use I to denote an arbitrary *multi-set* of input actions, $c(I)$ to denote $\{a? \mid a?v \in I\}$ and $\bar{c}(I)$ its converse, $\{a! \mid a?v \in I\}$.

Then \searrow_I^σ is defined to be the least relation which satisfies

- $\mathcal{C} \not\rightarrow$ and $I^\sigma(\mathcal{C}) \cap I = \emptyset$ implies $\mathcal{C} \searrow_I^\sigma \mathcal{C}$
- $\mathcal{C} \xrightarrow{i}_\sigma \mathcal{C}'$ and $\mathcal{C}' \searrow_I^\sigma \mathcal{C}''$ implies $\mathcal{C} \searrow_{I \uplus \{i\}}^\sigma \mathcal{C}''$.

Intuitively $\mathcal{C}' \searrow_I^\sigma \mathcal{C}''$ means that \mathcal{C} can evolve to a stable configuration \mathcal{C}' by performing a subset of the input actions in the multi-set I ; moreover this subset is maximal in the sense that \mathcal{C}' can not perform any of the remaining actions.

Definition 3.7 [Asynchronous acceptance sets] For a configuration \mathcal{C} , let $\mathcal{O}_I^\sigma(\mathcal{C}, s)$, its σ -level asynchronous acceptance set after s , relative to the multi-set of input actions I , be defined by

$$\{ \mathcal{O}_I^\sigma(\mathcal{C}'') \mid \mathcal{C} \xrightarrow{c'}_\sigma \mathcal{C}' \searrow_I^\sigma \mathcal{C}'' \}.$$

With one final notational convention we can mimic the alternative characterisation of *must* testing from [3]. We write $\Gamma \text{ allows}_\sigma a?v$ if $\Gamma \vdash_\sigma a!\langle v \rangle$; this is generalised to sets of actions in the normal manner.

Definition 3.8 Let \mathcal{C}, \mathcal{D} be configurations of the form $\Gamma \triangleright P, \Gamma \triangleright Q$ respectively. Then $\mathcal{C} \ll^\sigma \mathcal{D}$ if for every s ,

- $\mathcal{C} \Downarrow s$ implies a) $\mathcal{D} \Downarrow s$
 b) $\forall D \in \mathcal{A}^\sigma(\mathcal{D}, s), \forall I$ such that $c(I) \cap D = \emptyset$ and
 $(\Gamma \text{ after}_\sigma s)$ allows $_\sigma I$,
 $\exists O \in \mathcal{O}_I^\sigma(\mathcal{C}, s)$ such that $O - \bar{c}(I) \subseteq D$.

Theorem 3.9 *Let Γ be a single-level environment such that both $\Gamma \triangleright P$ and $\Gamma \triangleright Q$ are configurations. Then $\Gamma \triangleright_\sigma P \sqsubseteq_{\text{must}} Q$ if and only if $\Gamma \triangleright P \ll^\sigma \Gamma \triangleright Q$.*

4 Non-Interference Results

In this section we reconsider the approach taken to non-interference in Section 4 of [11]. The essential idea is that if a process is well-typed at a given level σ then its behaviour at that level is independent of processes “running at higher security levels”; or more generally “running at security levels independent to σ ”. A particular formulation of such a result was given in Theorem 5.3 of [11]:

Theorem 4.1 *If $\Gamma \vdash_\sigma P, Q$ and $\Gamma \vdash_{\text{top}} H, K$, where H, K are σ -free processes, then $\Gamma \triangleright_\sigma P \simeq_{\text{may}} Q$ implies $\Gamma \triangleright_\sigma P \mid H \simeq_{\text{may}} Q \mid K$.*

Here, because of our more refined notions of well-typing, and the characterisation result Theorem 3.4 we can give offer a significant improvement on this theorem, and moreover the formulation is actually easier.

Let us say that the security level δ is independent of σ if $\delta \not\leq \sigma$. We can ensure that a process H is “running at a security level independent to σ ” by demanding that $\Gamma \vdash^\delta H$, for some δ independent of σ . In fact we will only require the weaker typing relation $\Delta \vdash^{\omega^\delta} H$. This ensures that all the output actions of H are at a level independent of σ , as can be deduced from the following property:

Lemma 4.2 *Suppose $\Gamma \vdash^{\omega^\delta} H$. Then $\Gamma \triangleright H \xrightarrow{\mu}_\rho$, where μ is an output action, implies $\delta \leq \rho$.*

Definition 4.3 We say a process H is σ -high with respect to Γ if $\Gamma \vdash^{\omega^\delta} H$ for some δ independent of σ .

We can now state our first non-interference result. Note that it applies to processes such that $\Gamma \vdash_{\tau\sigma} P, Q$ rather than the more restrictive $\Gamma \vdash_\sigma P, Q$, as in Theorem 4.1; only their input actions need to be at level at most σ .

Theorem 4.4 (Non-Interference 1) *Suppose $\Gamma \vdash_{\tau\sigma} P, Q$. Then*

$$\Gamma \triangleright^\sigma P \sqsubseteq_{\text{may}} Q \text{ implies } \Gamma \triangleright^\sigma P \mid H \sqsubseteq_{\text{may}} Q \mid K$$

for all σ -high processes H, K .

Non-interference with respect to *may* testing equivalence gives a certain level of assurance that there is no information flow from high-level processes to low-level processes. But it has been argued in [5] that the stronger the

equivalence used in the formulation of interference-freeness, the more we can be assured that high-level information can not be leaked. An example may be found in [4] where it is argued that the ability of a low-level user to observe the potential *absence* of actions may result in an undesirable flow of information. Our second non-interference result, which relies on the characterisation theorem Theorem 3.9, addresses this problem as it is expressed in terms of *must* testing equivalence; this is sensitive to potential deadlocks, and therefore the potential absence of actions.

First note that Theorem 4.4 is no longer true when \sqsubseteq_{may} is replaced by \sqsubseteq_{must} , as the following example shows.

Example 4.5 Let A, B denote the types $\{\mathbf{w}_{\text{bot}}\langle \rangle, \mathbf{r}_{\text{bot}}\langle \rangle, \mathbf{r}_{\text{top}}\langle \rangle\}$ and $\{\mathbf{r}_{\text{top}}\langle \rangle\}$ respectively. Further, let Γ map a to A and n to the type $\{\mathbf{w}_{\text{bot}}\langle A \rangle, \mathbf{r}_{\text{bot}}\langle A \rangle, \mathbf{r}_{\text{top}}\langle B \rangle\}$. Now consider the processes P and H defined by

$$P \Leftarrow n!\langle a \rangle \mid n?(x : A) \ x!\langle \rangle \qquad H \Leftarrow n?(x : B) \ \mathbf{0}$$

It is very easy to check that $\Gamma \vdash_{\text{rbot}} P$ and $\Gamma \vdash^{\text{wtop}} H$. However

$$\Gamma; \Gamma \triangleright^{\text{bot}} P \mid \mathbf{0} \not\sqsubseteq_{must} P \mid H$$

because of the **bot** level test $a?() \ \omega!\langle \rangle$.

The presence or absence of H determines whether or not there is read contention on the channel n , which in turn influences the deadlock capabilities of P with respect to the channel a .

Here the problem is the type of the channel n ; it may be read at both level **bot** and **top**. Note that such examples, where there is contention between reads at different levels, can not be expressed in the *join calculus*, [4].

A not unreasonable restriction would be to require that the read capability of channels be confined to a particular security level, using single-level types. This would not rule out inter-level communication, but simply control it more tightly. This restriction can be enforced requiring the type-checking to use single-level types and forbidding high-level processes to read from low-level channels.

Definition 4.6 We say a process H is *strong* σ -high with respect to Γ if $\Gamma \vdash^{\delta} H$ for some δ independent of σ .

Theorem 4.7 (Non-Interference 2) *Suppose $\Gamma \vdash_{\text{r}\sigma} P, Q$, where Γ is a single-level environment. Then*

$$\Gamma \triangleright^{\sigma} P \sqsubseteq_{must} Q \text{ implies } \Gamma \triangleright^{\sigma} P \mid H \sqsubseteq_{must} Q \mid K$$

for all finite strong σ -high processes H, K .

Note that we must restrict our attention to finite H and K since *must* testing is sensitive to divergence; if H is a divergent term then we could not

expect $\Gamma \triangleright^\sigma P \mid \mathbf{0} \simeq_{\text{must}} P \mid H$ to hold when P is a convergent term. This problem is avoided by restricting attention to finite terms, which can never diverge.

5 Related Work

A general overview of the use of static analysis techniques to enforce information-flow policies may be found in [20]. Useful surveys of research into non-interference in process languages are given in [5,19]³. Much of this work is behaviour based; systems are deemed to be interference-free if their trace sets, sequences of actions labelled *high* or *low*, satisfy certain properties. Here we use a more extensional approach, saying that a system is interference-free if low-level observers are unable to discern the presence or absence of high-level components. There must, of course be some connection between our definition and at least one of the behavioural definitions in the literature. However the comparison is not straightforward. The definitions, in papers such as [18,5] are for very simple untyped versions of CCS or CSP, while much of the power of our approach comes from the use of types for the more sophisticated π -calculus.

In [14] a type system is given which guarantees non-interference with respect to an extension of the π -calculus; moreover non-interference is expressed with respect to a barbed congruence. However the language used is a considerable extension of the π -calculus, with operators for selection based input/output, based on disjunctive patterns, and it is the behaviour of these operators which are mainly constrained by the type system. The types used are also very sophisticated. Unlike ours, which are simply annotated versions of the standard Pierce/Sangiorgi types, [16], they track the use of channels, using annotated affine and linear types, and capture causal relationships between actions by a partial composition on these types, using ideas based on the *graph types* of [24].

Finally [4], which uses security labels attached to messages in the join calculus to formulate non-interference, argues via an example for the use of a behavioural equivalence stronger than *may* testing. The formulation uses *weak barbed congruence* but could have equally well used *must* testing equivalence; indeed it is difficult to envisage a practical scenario in which there is something to be gained from assuming attackers have the extra power associated with the former rather than the latter. We have also already pointed out (in Example 4.5) that the join calculus can not be used to express situations in which there is read contention between different security levels. Nevertheless the approach used to develop a type system for the join calculus for detecting information flow seems to be quite general and may be applicable to the asynchronous π -calculus.

³ For the use of types for other languages see [22].

References

- [1] D. E. Bell and L. J. LaPadula. Secure computer system: Unified exposition and multics interpretation. Technical report MTR-2997, MITRE Corporation, 1975.
- [2] G. Boudol. Asynchrony and the π -calculus. Technical Report 1702, INRIA-Sophia Antipolis, 1992.
- [3] Ilaria Castellani and Matthew Hennessy. Testing theories for asynchronous languages. In V Arvind and R Ramanujam, editors, *18th Conference on Foundations of Software Technology and Theoretical Computer Science (Chennai, India, December 17–19, 1998)*, LNCS 1530. Springer-Verlag, December 1998.
- [4] Sylvain Conchon. Modular information flow analysis for process calculi. In Iliano Cervesato, editor, *Proceedings of the Foundations of Computer Security Workshop (FCS 2002)*, Copenhagen, Denmark, JULY 2002.
- [5] Riccardo Focardi and Roberto Gorrieri. A classification of security properties for process algebras. *Journal of Computer Security*, 3(1), 1995.
- [6] Riccardo Focardi and Roberto Gorrieri. Non interference: Past, present and future. In *Proceedings of DARPA Workshop on Foundations for Secure Mobile Code*, 1997.
- [7] J. A. Goguen and J. Meseguer. Security policies and security models. In *IEEE Symposium on Security and privacy*, 1992.
- [8] M. Hennessy. *An Algebraic Theory of Processes*. MIT Press, 1988.
- [9] Matthew Hennessy. The security picalculus and non-interference. Technical report 2000:05, University of Sussex, 2000. Available from <http://www.cogs.susx.ac.uk/>.
- [10] 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.
- [11] Matthew Hennessy and James Riely. Information flow vs resource access in the asynchronous pi-calculus (extended abstract). In U. Montanari, J. Rolim, and E. Welzl, editors, *Automata, Languages and Programming, 27th International Colloquium*, volume 1853 of *Lecture Notes in Computer Science*, pages 415–427, Geneva, Switzerland, July 2000. Springer-Verlag. Full version to appear in *ACM Transactions on Programming Languages and Systems*.
- [12] Matthew Hennessy and James Riely. Resource access control in systems of mobile agents. *Information and Computation*, 173:82–120, 2002.
- [13] Kohei Honda and Mario Tokoro. On asynchronous communication semantics. In P. Wegner M. Tokoro, O. Nierstrasz, editor, *Proceedings of the ECOOP '91*

- Workshop on Object-Based Concurrent Computing*, volume 612 of *LNCS 612*. Springer-Verlag, 1992.
- [14] Kohei Honda and Nobuko Yoshida. A uniform type structure for secure information flow. In *29th Annual Symposium on Principles of Programming Languages*. ACM, January 2002.
 - [15] R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 24:83–113, 1984.
 - [16] B. Pierce and D. Sangiorgi. Behavioral equivalence in the polymorphic pi-calculus. *Journal of the ACM*, 47(3):531–584, 2000.
 - [17] Benjamin Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. *Mathematical Structures in Computer Science*, 6(5):409–454, 1996. Extended abstract in LICS '93.
 - [18] A.W. Roscoe, J.C.P. Woodcock, and L. Wulf. Non-interference through determinism. In *European Symposium on Research in Computer Security*, volume 875 of *LNCS*, 1994.
 - [19] P.Y.A. Ryan and S.A. Schneider. Process algebra and non-interference. In *CSFW 12*. IEEE, 1997.
 - [20] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE J. Selected Areas in Communication*, 2002. To appear.
 - [21] D. Sangiorgi and D. Walker. *The Picalculus*. Cambridge University Press, 2001.
 - [22] Geoffrey Smith and Dennis Volpano. Secure information flow in a multi-threaded imperative language. In *Conference Record of the ACM Symposium on Principles of Programming Languages*, San Diego, January 1998.
 - [23] David Turner. *The Polymorphic Pi-Calculus: Theory and Implementation*. Ph.d. thesis, Edinburgh University, 1995.
 - [24] Nobuko Yoshida. Graph types for monadic mobile processes. In *FSTTCS*, volume 1180, pages 371–386. Springer-Verlag, 1996.