

Information Flow vs. Resource Access in the Asynchronous Pi-Calculus (Extended Abstract)

Matthew Hennessy¹ and James Riely²

¹ COGS, University of Sussex, UK
email: matthewh@cogs.sussex.ac.uk,

home page: <http://www.cogs.susx.ac.uk/users/matthewh/>

² DePaul University, US

email: jriley@cs.depaul.edu

home page: <http://www.depaul.edu/~jriely/>

ABSTRACT. We propose an extension of the asynchronous π -calculus in which a variety of security properties may be captured using types. These are an extension of the Input/Output types for the π -calculus in which I/O capabilities are assigned specific security levels.

We define a typing system which ensures that processes running at security level σ cannot access resources with a security level higher than σ . The notion of access control guaranteed by this system is formalized in terms of a Type Safety theorem.

We then show that, for a certain class of processes, our system prohibits implicit information flow from high-level to low-level processes. We prove that low-level behaviour can not be influenced by changes to high-level behaviour. This is formalized as a Non-Interference Theorem with respect to may testing.

1 Introduction

The problem of protecting information and resources in systems with multiple sensitivity or security levels, [2], has been studied extensively. Flow analysis techniques have been used in [3, 4], axiomatic logic in [13] while in [27, 15] type systems have been developed for a number of prototypical programming languages. In this paper, we explore the extent to which type systems for ensuring various forms of security can also be developed for the asynchronous π -calculus [5, 17]. We discuss two security issues: resource access control and information control. The former is described in terms of runtime errors, the latter in terms of non-interference [27, 11].

The (asynchronous) π -calculus is a very expressive language for describing distributed systems, [5, 23, 12], in which processes intercommunicate using channels. Thus $n?(x)P$ is a process which receives some value on the channel named n , binds it to the variable x and executes the code P . Corresponding to this input command is the asynchronous output command $n!\langle v \rangle$ which outputs the value v on n . The set of values which may be transmitted on channels includes channel names themselves; this, together with the ability to dynamically create new channel names, gives the language its descriptive power.

Within the setting of the π -calculus we wish to investigate the use of types to enforce security policies. To facilitate the discussion we extend the syntax with a new construct to represent a process running at a given security clearance, $\sigma[P]$. Here σ is some security level taken from a complete lattice of security levels SL and P is the code of the process. Further, we associate with each channel, the *resources* in our language, a set of input/output capabilities [22, 24], each decorated with a specific security level. Intuitively, if channel n has a read capability at level σ , then only processes running at security level σ or higher may be read from n . This leads to the notion of a *security policy* Σ , which associates a set of capabilities with each channel in the system. The question then is to design a typing system which ensures that processes do not violate the given security policy.

Of course this depends on when we consider such a violation to take place. For example if Σ assigns the channel or resource n the highest security level top then it is reasonable to say that a violation will eventually occur in

$$c!\langle n \rangle \mid \text{bot}[[c?(x) x?(y) P]]$$

as after the communication on c , a low level process, $\text{bot}[[n?(y) P]]$ has gained access to the high level resource n . Underlying this example is the principle that processes at a given security level σ should have access to resources at security level *at most* σ . We formalize this principle in terms of a relation $P \xrightarrow{\Sigma} \text{err}$, indicating that P violates the security policy Σ .

To prevent such errors, we restrict attention to security policies that are somehow consistent. Let Γ be such a consistent policy; consistency is defined by restricting types so that they respect a subtyping relation. We then introduce a typing system, $\Gamma \vdash P$, which ensures that P can never violate Γ :

If $\Gamma \vdash P$ then for every context $C[\]$ such that $\Gamma \vdash C[P]$ and every Q which occurs during the execution of $C[P]$, that is $C[P] \mapsto^* Q$, we have $Q \xrightarrow{\Gamma} \text{err}$.

Thus our typing system ensures that low level processes will never gain access to high level resources. The typing system implements a particular view of security, which we refer to as the *R-security policy*, as it offers protection to *resources*. Here communication is allowed between high level and low level principals, provided of course that the values involved are appropriate.

This policy does not rule out the possibility of information leaking indirectly from high security to low security principals. Suppose h is a high channel and hl is a channel with high-level write access and low-level read access in:

$$\text{top}[[h?(x) \text{ if } x = 0 \text{ then } hl!\langle 0 \rangle \text{ else } hl!\langle 1 \rangle]] \mid \text{bot}[[hl?(z) Q]] \quad (\star)$$

This system can be well-typed although there is some implicit information flow from the high security agent to the low security one; the value received on the high level channel h can be determined by the low level process Q .

It is difficult to formalize exactly what is meant by *implicit information flow* and in the literature various authors have instead relied on *non-interference*, [14,

25, 11, 26], a concept more amenable to formalization, which ensures, at least informally, the absence of implicit information flow.

To obtain such results for the π -calculus we need, as the above example shows, a stricter security policy, which we refer to as the *I-security policy*. This allows a high level principal to read from low level resources but not to write to them. Using the terminology of [2, 7]:

- *write up*: a process at level σ may only write to channels at level σ or above
- *read down*: a process at level σ may only read from channels at level σ or below.

In fact the type inference system remains the same and we only need constrain the notion of type. In this restricted type system well-typing, $\Gamma \Vdash P$, ensures a form of *non-interference*.

To formalize this non-interference result we need to develop a notion of process behaviour, relative to a given security level. Since the behaviour of processes also depends on the type environment in which they operate we need to define a relation

$$P \approx_{\Gamma}^{\sigma} Q$$

which intuitively states that, relative to Γ , there is no observable distinction between the behaviour of P and Q at security level σ ; processes running at security level σ can observe no difference in the behaviour of P and Q . Lack of information flow from high to low security levels now means that this relation is invariant under changes in high-level values; or indeed under changes in high-level behaviour.

It turns out that the extent to which this is true depends on the exact formulation of the behavioural equivalence $\approx_{\Gamma}^{\sigma}$. We show that it is *not* true if $\approx_{\Gamma}^{\sigma}$ is based on *observational* equivalence [19] or *must testing* equivalence [21]. But a result can be established if we restrict our attention to *may testing* equivalence (here written \simeq_{Γ}^{σ}). Specifically we will show that, for certain H, K :

$$\text{If } \Gamma \Vdash P, Q \text{ and } \Gamma \Vdash^{\text{top}} H, K \text{ then } P \simeq_{\Gamma}^{\sigma} Q \text{ implies } P|H \simeq_{\Gamma}^{\sigma} Q|K \quad (\star\star)$$

The remainder of the paper is organized as follows. In the next section we define the *security π -calculus*, giving a labelled transition semantics and a formal definition of runtime errors. In Section 3 we design a set of types and a typing system which implements the resource control policy. This section also contains Subject Reduction and Type Safety theorems. In Section 4 we motivate the restrictions required on types and terms in order to implement the information control policy. We also give a precise statement of our non-interference result, and give counter-examples to related conjectures based on equivalences other than *may testing*.

The proof of our main theorem depends on an analysis of *may testing* in terms of *asynchronous* sequences of actions [6] which in turn depends on a more explicit operational semantics for our language, where actions are parameterised relative to a typing environment. The details may be found in the full version of the paper, [16].

Fig. 1 Syntax

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

2 The Language

The syntax of the *security π -calculus*, given in Figure 1, uses a predefined set of *names*, ranged over by a, b, \dots, n and a set of *variables*, ranged over by x, y, z . *Identifiers* are either variables or names. *Security annotations*, ranged over by small Greek letters σ, ρ, \dots , are taken from a complete lattice $\langle SL, \preceq, \sqcap, \sqcup, \text{top}, \text{bot} \rangle$ of security levels. We also assume for each σ a set of *basic values* BV_σ ; we use bv to range over base values. We require that all syntactic sets be disjoint.

The binding constructs $u?(X : A) Q$ and $(\text{new } a : A) Q$ introduce the usual notions of free names and variables, $\text{fn}(P)$ and $\text{fv}(P)$, respectively, and associated notions of substitution; details may be found in the full version. Moreover the typing annotations on the binding constructs, which will be explained in Section 3, are omitted whenever they do not play a role.

The behaviour of a process is determined by the interactions in which it can engage. To define these, we give a labelled transition semantics (LTS) for the language. The set *Act of labels*, or *actions*, is defined as follows:

$\mu ::=$	<i>Actions</i>
τ	Internal action
$(\tilde{c} : \tilde{C})a?v$	Input of v on a learning private names \tilde{c}
$(\tilde{c} : \tilde{C})a!v$	Output of v on a revealing private names \tilde{c}

Visible actions (all except τ) are ranged over by α, β and we use $\mathcal{E}(\alpha)$ to denote the bound names in α , together with their types. $\mathcal{E}((\tilde{c} : \tilde{C})a!v) = \mathcal{E}((\tilde{c} : \tilde{C})a?v) = (\tilde{c} : \tilde{C})$. Further, let $\text{n}(\mu)$ be the set of *names* occurring in μ , whether free or bound. We say that the actions ‘ $(\tilde{c} : \tilde{C})a?v$ ’ and ‘ $(\tilde{c} : \tilde{C})a!v$ ’ are *complementary*, with $\bar{\alpha}$ denoting the complement of α .

The LTS is defined in Figure 2 and for the most part the rules are straightforward; it is based on the standard operational semantics from [20], to which the reader is referred for more motivation.

Informally a security policy associates with each input/output capability on a channel a security level. To this end, *Pre-capabilities* and *pre-types* are defined

Fig. 2 Labelled Transition Semantics

$\frac{}{a!\langle v \rangle \xrightarrow{a!v} \mathbf{0}} \quad \text{(L-OUT)}$ $\frac{}{a?(X)P \xrightarrow{(\tilde{c}:C)a?v} P\{v/X\}} \quad \text{(L-IN)}$ $\frac{}{P \xrightarrow{(\tilde{c}:C)a!v} P'} \quad \text{(L-OPEN)}$ $\frac{}{(\text{new } b : B)P \xrightarrow{(b:B)(\tilde{c}:C)a!v} P'} \quad b \neq a, b \in \text{fn}(v)$ $\frac{}{P \xrightarrow{\alpha} P', Q \xrightarrow{\bar{\alpha}} Q'} \quad \text{(L-COM)}$ $\frac{}{P Q \xrightarrow{\tau} (\text{new } \mathcal{E}(\alpha)) (P' Q')} \quad \text{(L-EQ)}$ $\frac{}{\text{if } u = u \text{ then } P \text{ else } Q \xrightarrow{\tau} P} \quad \text{(L-CTX)}$ $\frac{}{P \xrightarrow{\mu} P'} \quad \text{(L-CTXT)}$ $\frac{}{*P \xrightarrow{\mu} *P P'} \quad \text{(L-CTXT)}$ $\frac{}{\sigma[P] \xrightarrow{\mu} \sigma[P']}$ $\frac{}{P \xrightarrow{\mu} P'} \quad \text{(L-CTXT)}$ $\frac{}{P Q \xrightarrow{\mu} P' Q} \quad \text{bn}(\mu) \notin \text{fn}(Q)$ $\frac{}{Q P \xrightarrow{\mu} Q P'}$ $\frac{}{P \xrightarrow{\mu} P'} \quad \text{(L-CTXT)}$ $\frac{}{(\text{new } a : A)P \xrightarrow{\mu} (\text{new } a : A)P'} \quad a \notin \text{n}(\mu)$	$\tilde{c} \notin \text{fn}(P), \tilde{c} \in \text{fn}(v)$ $u \neq w$
--	--

as follows:

$cap ::=$ $w_\sigma \langle A \rangle$ $r_\sigma \langle A \rangle$	<i>Pre-Capability</i> σ -level process can write values with type A σ -level process can read values with type A
$A ::=$ B_σ $\{cap_1, \dots, cap_k\}$ (A_1, \dots, A_k)	<i>Pre-Type</i> Base type Resource type ($k \geq 0$) Tuple type ($k \geq 0$)

We will tend to abbreviate a singleton set of capabilities, $\{cap\}$, to cap .

A *security policy*, Σ , is a finite mapping from names to pre-types. Thus, for example, if Σ maps the channel lh to the pre-type $\{w_{\text{bot}}\langle B \rangle, r_{\text{top}}\langle A \rangle\}$, for some appropriate A, B, then low level processes may write to lh but only high level ones may read from it; this is an approximation of the security associated with a mailbox. On the other hand if Σ maps hl to $\{r_{\text{bot}}\langle A \rangle, w_{\text{top}}\langle B \rangle\}$ then hl acts more like an information channel; anybody can read from it but only high level processes may place information there.

The import of a security policy may be underlined by defining what it means to violate it. Our definition is given in Figure 3, in terms of a relation $P \xrightarrow{\Sigma} err$. For example, relative to the policy Σ defined above, after one reduction step of the process $\text{top}[\![c!\langle hl \rangle]\!] | \text{bot}[\![c?(x) x!\langle v \rangle]\!]$, there is a security error because

Fig. 3 Runtime Errors

$$\begin{array}{l}
\text{(E-RD)} \quad \rho \llbracket a?(X) P \rrbracket \xrightarrow{\Sigma} \text{err} \quad \text{if } \sigma \preceq \rho \text{ implies for all } A, r_\sigma \langle A \rangle \notin \Sigma(a) \\
\text{(E-WR}_1\text{)} \quad \rho \llbracket a!\langle v \rangle \rrbracket \xrightarrow{\Sigma} \text{err} \quad \text{if } \sigma \preceq \rho \text{ implies for all } A, w_\sigma \langle A \rangle \notin \Sigma(a) \\
\text{(E-WR}_2\text{)} \quad \rho \llbracket a!\langle v \rangle \rrbracket \xrightarrow{\Sigma} \text{err} \quad \text{if } bv \in v, bv \in \mathbf{B}_\sigma \text{ and } \sigma \not\preceq \rho \\
\text{(E-STR)} \quad \frac{P \xrightarrow{\Sigma} \text{err} \quad P \xrightarrow{\Sigma} \text{err} \quad P \equiv Q, P \xrightarrow{\Sigma} \text{err}}{P \mid Q \xrightarrow{\Sigma} \text{err} \quad \rho \llbracket P \rrbracket \xrightarrow{\Sigma} \text{err} \quad Q \xrightarrow{\Sigma} \text{err}} \\
\frac{P \xrightarrow{\Sigma, a: A} \text{err}}{(\text{new } n : A) P \xrightarrow{\Sigma} \text{err}}
\end{array}$$

$\text{bot} \llbracket \text{hl}!\langle v \rangle \rrbracket \xrightarrow{\Sigma} \text{err}$. A low security process has read access to security channel hl on which write access is reserved for high-security processes. Assuming an appropriate typing for c and v the same security error does not occur in $\text{top} \llbracket c!\langle \text{lh} \rangle \rrbracket \mid \text{bot} \llbracket c!(x) x!\langle v \rangle \rrbracket$. The low security process $\text{bot} \llbracket \text{hl}!\langle v \rangle Q \rrbracket$ has the right to write on the channel lh .

3 Resource Control

Our typing system will apply only to certain security policies, those in which the pre-types are in some sense *consistent*. Consistency is imposed using a system of kinds: the kind $R\text{Type}_\sigma$ comprises the value types accessible to processes at security level σ . These kinds are in turn defined using a subtyping relation on pre-capabilities and pre-types.

Definition 1. Let $<$: be the least preorder on pre-capabilities and pre-types such that:

$$\begin{array}{l}
\text{(U-WR)} \quad w_\sigma \langle A \rangle < w_\sigma \langle B \rangle \quad \text{if } B < A \\
\text{(U-RD)} \quad r_\sigma \langle A \rangle < r_\rho \langle B \rangle \quad \text{if } A < B \text{ and } \sigma \preceq \rho \\
\text{(U-BASE)} \quad \mathbf{B}_\sigma < \mathbf{B}_\rho \quad \text{if } \sigma \preceq \rho \\
\text{(U-RES)} \quad \{\text{cap}_i\}_{i \in I} < \{\text{cap}'_j\}_{j \in J} \quad \text{if } (\forall j)(\exists i) \text{cap}_i < \text{cap}'_j \\
\text{(U-TUP)} \quad (A_1, \dots, A_k) < (B_1, \dots, B_k) \quad \text{if } (\forall i) A_i < B_i
\end{array}$$

For each ρ , let $R\text{Type}_\rho$ be the least set that satisfies:

$$\begin{array}{l}
\text{(RT-WR)} \quad \frac{A \in R\text{Type}_\sigma}{\{w_\sigma \langle A \rangle\} \in R\text{Type}_\rho} \quad \sigma \preceq \rho \\
\text{(RT-RD)} \quad \frac{A \in R\text{Type}_\sigma}{\{r_\sigma \langle A \rangle\} \in R\text{Type}_\rho} \quad \sigma \preceq \rho \\
\text{(RT-WRRD)} \quad \frac{A \in R\text{Type}_\sigma \quad A' \in R\text{Type}_{\sigma'}}{\{w_\sigma \langle A \rangle, r_{\sigma'} \langle A' \rangle\} \in R\text{Type}_\rho} \quad \begin{array}{l} \sigma \preceq \rho \\ \sigma' \preceq \rho \\ A < A' \end{array} \\
\text{(RT-BASE)} \quad \frac{}{\mathbf{B}_\sigma \in R\text{Type}_\rho} \quad \sigma \preceq \rho \\
\text{(RT-TUP)} \quad \frac{A_i \in R\text{Type}_\rho \quad (\forall i)}{(A_1, \dots, A_k) \in R\text{Type}_\rho}
\end{array}$$

Fig. 4 Typing Rules

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

Let RType be the union of the kinds RType_ρ over all ρ . □

Note that if $\sigma \preceq \rho$ then $\text{RType}_\sigma \subseteq \text{RType}_\rho$. Intuitively, low level values are accessible to high level processes. However the converse is not true. For example $w_{\text{top}} \langle \rangle \in \text{RType}_{\text{top}}$ but $w_{\text{top}} \langle \rangle$ is not in $\text{RType}_{\text{bot}}$. The compatibility requirement between read and write capabilities in a type (RT-WRRD), in addition to the typing implications discussed in [24], also has security implications. For example suppose $r_{\text{bot}} \langle \mathbf{B}_\sigma \rangle$ and $w_{\text{top}} \langle \mathbf{B} \rangle$ are capabilities in a valid channel type. Then apriori a high level process can write to the channel while a low level process may read from it. However the only possibility for σ is bot, that is only low level values may be read. Moreover the requirement $\mathbf{B} \prec: \mathbf{B}_\sigma$ implies that \mathbf{B} must also be \mathbf{B}_{bot} . So although high level processes may write to the channel they may only write low level values.

Proposition 1. *For every ρ , RType_ρ is a preorder with respect to \prec , with both a partial meet operation \sqcap and a partial join \sqcup .* □

A *type environment* is a finite mapping from identifiers (names and variables) to types. We adopt some standard notation. For example, let $\langle \Gamma, u : A \rangle$ denote the obvious extension of Γ ; $\langle \Gamma, u : A \rangle$ 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. We will normally abbreviate the simple environment $\{u : A\}$ to $u : A$ and moreover use $v : A$ to denote its obvious generalisation to values.

The typing system is given in Figure 4 where the judgements are of the form $\langle \Gamma \Vdash P \rangle$. If $\langle \Gamma \Vdash P \rangle$ we say that P is a σ -level process. Also, let $\langle \Gamma \vdash P \rangle$ abbreviate $\langle \Gamma \Vdash^{\text{top}} P \rangle$.

Intuitively $\langle \Gamma \Vdash P \rangle$ indicates that the process P will not cause any security errors if executed with security clearance σ . The rules are very similar to those used in papers such as [24, 22] for the standard IO typing of the π -calculus. Indeed the only significant use of the security levels is in the (T-IN) and (T-OUT)

rules, where the channels are required to have a specific security level. This is inferred using auxiliary value judgements, of the form $\Gamma \vdash v : A$. It is interesting to note that security levels play no direct role in their derivation.

Theorem 1 (Subject Reduction). *Suppose $\Gamma \Vdash P$. Then*

- $P \xrightarrow{\tau} Q$ implies $\Gamma \Vdash Q$
- $P \xrightarrow{(\tilde{c}:\tilde{C})a?v} Q$ implies there exists a type A such that $\Gamma \vdash a : r_\delta \langle A \rangle$ for some $\delta \preceq \sigma$, and if $\Gamma \sqcap v : A$ is well-defined then $\Gamma \sqcap v : A \Vdash Q$.
- $P \xrightarrow{(\tilde{c}:\tilde{C})a!v} Q$ implies there exists a type A such that $\Gamma \vdash a : w_\delta \langle A \rangle$ for some $\delta \preceq \sigma$, $\Gamma, \tilde{c} : \tilde{C} \vdash v : A$ and $\Gamma, \tilde{c} : \tilde{C} \Vdash Q$.

□

We can now state the first main result:

Theorem 2 (Type Safety). *If $\Gamma \vdash P$ then for every closed context $C[\]$ such that $\Gamma \vdash C[P]$ and every Q such that $C[P] \xrightarrow{\tau}^* Q$ we have $Q \xrightarrow{F} \text{err}$*

□

Having defined our typing system we may now view $\sigma \llbracket P \rrbracket$ simply as notation for the fact that, relative to the current typing environment Γ , the process P is well-typed at level σ , i.e. $\Gamma \Vdash P$. Technically we can view $\sigma \llbracket P \rrbracket$ to be *structurally equivalent* to P , assuming we are working in an environment Γ such that $\Gamma \Vdash P$.

4 Information Flow

We have shown in the previous sections that, in well-typed systems, processes running at a given security level can only access resources appropriate to that level. However, as pointed out in the Introduction this does not rule out (implicit) information flow between levels. One way of formalizing this notion of flow of information is to consider the behaviour of processes and how it can be influenced. If the behaviour of low-level processes is independent of any high-level values in its environment then we can say that there can be no implicit flow of information from high-level to low-level. This is *not* the case in the example considered in the Introduction, (\star) . Suppose, for example, that Q is the code fragment ‘if $z = 0$ then $l_1! \langle \rangle$ else $l_2! \langle \rangle$ ’. If (\star) were placed in an environment with ‘ $\text{top} \llbracket h! \langle 0 \rangle \rrbracket$ ’, then the resource l_1 would be called. If, instead, (\star) were placed in an environment with ‘ $\text{top} \llbracket h! \langle 42 \rangle \rrbracket$ ’, then l_2 would be called. In other words the behaviour of the low-level process can be influenced by high-level changes; there is a possibility of information flow downwards.

This is not surprising in view of the type associated with the channel hl ; in the terminology of [2] it allows a *write down* from a high-level process to a low-level process. Thus if we are to eliminate implicit information flow between levels in well-typed processes we need to restrict further the allowed types; types such as $\{\text{w}_{\text{top}} \langle \rangle, \text{r}_{\text{bot}} \langle \rangle\}$ clearly contradict the spirit of secrecy. Thus, for the rest of the paper we work with the more restrictive set *IType*, the *Information types*. In order for $\{\text{w}_\sigma \langle A \rangle, \text{r}_{\sigma'} \langle A' \rangle\}$ to be in *IType*, it must be that $\sigma \preceq \sigma'$; this is not necessarily true for types in *RType*.

Definition 2. For each ρ , let IType_ρ , be the least set that satisfies the rules in Definition 1, with (RT-WRRD) replaced by:

$$\begin{array}{c} \text{(IT-WRRD)} \\ \frac{\begin{array}{l} A \in \text{IType}_\sigma \\ A' \in \text{IType}_{\sigma'} \end{array}}{\{\mathbf{w}_\sigma \langle A \rangle, r_{\sigma'} \langle A' \rangle\} \in \text{IType}_\rho} \quad \begin{array}{l} \sigma \preceq \sigma' \\ \sigma' \preceq \rho \end{array} \quad A \prec: A' \end{array}$$

Let IType be the union of IType_ρ over all ρ . We write $\Gamma \Vdash^\sigma P$ if $\Gamma \Vdash P$ can be derived from the rules of Figure 4 using these more restrictive types. \square

All of the results of the previous section carry over to the stronger typing system; we leave their elaboration to the reader.

Unfortunately, due to the expressiveness of our language, the use of *I-types* still does not preclude information flow downwards, between levels. Consider the system

$$\text{top} \llbracket h?(x) \text{ if } x = 0 \text{ then bot} \llbracket l! \langle 0 \rangle \rrbracket \text{ else bot} \llbracket l! \langle 1 \rangle \rrbracket \rrbracket \mid \text{bot} \llbracket l?(z) Q \rrbracket$$

executing in an environment in which h is a top-level read/write channel and l is a bot-level read/write channel. This system can be well-typed using *I-types*, but there still appears to be some some implicit flow of information from top to bot. The problem here is that our syntax allows a high-level process, which can not write to low-level channels, to evolve into a low-level process which does have this capability; we need to place a boundary between low- and high-level processes which ensures a high-level process never gains write access to low-level channels. This is the aim of the following definition:

Definition 3. Define the security levels of a term below ρ , $\text{sl}_\rho(P)$, as follows:

$$\begin{array}{l} \text{sl}_\rho(*P) = \text{sl}_\rho(P) \quad \text{sl}_\rho(\mathbf{0}) = \{\rho\} \quad \text{sl}_\rho(\sigma \llbracket P \rrbracket) = \{\sigma \sqcap \rho\} \cup \text{sl}_{\sigma \sqcap \rho}(P) \\ \text{sl}_\rho(\text{new } a : A \ P) = \text{sl}_\rho(P) \quad \text{sl}_\rho(u! \langle v \rangle) = \emptyset \quad \text{sl}_\rho(P \mid Q) = \text{sl}_\rho(P) \cup \text{sl}_\rho(Q) \\ \text{sl}_\rho(u?(X : B) \ P) = \text{sl}_\rho(P) \quad \text{sl}_\rho(\text{if } u = v \text{ then } P \text{ else } Q) = \text{sl}_\rho(P) \cup \text{sl}_\rho(Q) \end{array}$$

A process P is σ -free if for every ρ in $\text{sl}_{\text{top}}(P)$, $\rho \not\preceq \sigma$. \square

Non-interference, as discussed in the Introduction, ($\star\star$), depends on a formulation of a behavioural equivalence, as the following example illustrates. Let A denote the type $\{\mathbf{w}_{\text{bot}} \langle \rangle, r_{\text{bot}} \langle \rangle\}$ and B denote $\{r_{\text{bot}} \langle \rangle\}$. Further, let Γ map a and b to A and B , respectively, and n to the type $\{\mathbf{w}_{\text{bot}} \langle A \rangle, r_{\text{bot}} \langle A \rangle\}$. Now consider the terms P and H defined by

$$P \Leftarrow \text{bot} \llbracket n! \langle a \rangle \mid n?(x : A) \ x! \langle \rangle \rrbracket \quad H \Leftarrow \text{top} \llbracket n?(x : B) \ b?(y) \ \mathbf{0} \rrbracket$$

It is very easy to check that $\Gamma \Vdash P, H$ and that H is bot-free. Note that in the term $P \mid H$ there is contention between the low and high-level processes for who will receive a value on the channel n . This means that if we were to base the

semantic relation \approx on any of *strong bisimulation equivalence*, *weak bisimulation equivalence*, [19], or *must testing*, [21], we would have

$$P \mid \mathbf{0} \not\approx^\sigma P \mid H$$

The essential reason is that the consumption of writes can be detected; the reduction

$$P \mid H \xrightarrow{\tau} \text{bot}[[n?(x : A) \ x! \langle \rangle]] \mid \text{top}[[b?(y) \cdot \mathbf{0}]]$$

cannot be matched by $P \mid \mathbf{0}$. Using the terminology of [21], $P \mid \mathbf{0}$ *guarantees* the test $\text{bot}[[a?(x) \ \omega! \langle \rangle]]$ whereas $P \mid H$ does not.

May equivalence is defined in terms of tests. A *test* is a process with an occurrence of a new reserved resource name ω . We use T to range over tests, with the typing rule $\Gamma \Vdash \omega! \langle \rangle$ 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 $T \Downarrow$ if $T \xrightarrow{\tau}^* T'$, where T' has the form $(\text{new } \tilde{c}) (\omega! \langle \rangle \mid T'')$ for some T'' and \tilde{c} ; that is T can eventually report success.

We wish to capture the behaviour of processes at a given level of security. Consequently we only compare their ability to pass tests that are well-typed at that level. The definition must also take into account the environment in which the processes are used, as this determines the security level associated with resources.

Definition 4. We write $P \simeq_T^\sigma Q$ if for every test T such that $\Gamma \Vdash T$:

$$(P \mid T) \Downarrow \text{ if and only if } (Q \mid T) \Downarrow$$

□

We can now state the main result of the paper.

Theorem 3 (Non-Interference). If $\Gamma \Vdash P, Q$ and $\Gamma \Vdash^{\text{top}} H, K$ where H and K are σ -free processes, then $P \simeq_T^\sigma Q$ implies $P \mid H \simeq_T^\sigma Q \mid K$. □

The proof of the theorem relies on a constructing sufficient condition to guarantee that two processes are may equivalent. This condition involves the *asynchronous* sequences of actions which processes can perform in the type environment Γ . The details may be found in the full version of the paper, [16], which also contains the subsequent proof of the non-interference result.

Finally let us remark that if we allowed *synchronous* tests then this result would no longer hold. For an appropriate Γ would have:

$$P \mid \mathbf{0} \simeq_T^\sigma P \mid H$$

Let T be the test $\text{bot}[[b! \langle \rangle \ \omega! \langle \rangle]]$. Then $P \mid H \mid T$ may eventually produce an output on ω whereas $P \mid \mathbf{0} \mid T$ cannot. However, since our language is asynchronous, such tests are not allowed.

5 Conclusions and Related Work

Methods for controlling information flow are a central research issue in computer security [7, 14, 27] and in the Introduction we have indicated a number of different approaches to its formalisation. Non-interference has emerged as a useful concept and is widely used to infer (indirectly) the absence of information flow. In publications such as [25, 9] it has been pointed out that process algebras may be fruitfully used to formalise and investigate this concept; for example in [8] process algebra based methods are suggested for investigating security protocols, essentially using a formalisation of non-interference for CCS.

However in these publications the *non-interference* is always defined behaviourally, as a condition on the possible traces of CCS or CSP processes; useful surveys of trace based non-interference may be found in [9, 26]. Here, we work with the more expressive π -calculus, which allows dynamic process creation and network reconfiguration. Our approach to *non-interference* is also more extensional in that it is expressed in terms of how processes effect their environments, relative to a particular behavioural equivalence. However the proof of our main result, Theorem 3, describes may equivalence in terms of (typed) traces; presumably a trace based definition of *non-interference*, similar in style to those in [9, 26] could be extracted from this proof.

More importantly our approach differs from much of the recent process calculus based security research in that we develop purely *static* methods for ensuring security. Processes are shown to be secure not by demonstrating some property of trace sets, using a tool as such as that in [10], but by type-checking. Types have also been used in this manner in [1] for an extension of the π -calculus called the *spi-calculus*. But there the structure of the types are very straightforward; the type *Secret* representing a secret channel, the type *Public* representing a public one, and *Any* which could be either. However the main interest is in the type rules for the encryption/decryption primitives of the *spi-calculus*. The non-interference result also has a different formulation to ours; it states that the behaviour of well-typed processes is invariant, relative to *may* testing, under certain value-substitutions. Intuitively, it means that the encryption/decryption primitives preserve values of type *Secret* from certain kinds of attackers. It would be interesting to add these primitives to the our *security π -calculus* and to try to adapt the associated type rules to the set of *I-Types*.

An extension of the π -calculus is also considered in [18], where a sophisticated type system is used to control information flow. The judgements in their system take the form

$$\Gamma \vdash_s P \triangleright A$$

where s is a security level, P is a process term, A is a poset of so-called *action nodes* and Γ is a type environment. Their environments are quite similar to ours, essentially associating with channels a version of input/output types annotated with, among other things, security levels. However their intuition, and much of the technical development, is quite different from ours. In summary it appears that our type system addresses information flow within the core π -calculus while the more sophisticated one of [18] controls the flow allowed via the extra syntactic

constructs of their language. However a more thorough comparison between the two systems deserves to be made.

Acknowledgements: The research was partially funded by EPSRC grant GR/L93056, and ESPRT Working Group Confer2. The authors would like to thank I. Castellani for a careful reading of a draft version of the paper.

References

1. Martín Abadi. Secrecy by typing in security protocols. In *Proceedings of TACS'97*, volume 1281 of *Lecture Notes in Computer Science*, pages 611–637. Springer Verlag, 1997.
2. D. E. Bell and L. J. LaPadula. Secure computer system: Unified exposition and multics interpretation. Technical report MTR-2997, MITRE Corporation, 1975.
3. C. Bodei, P. Degano, F. Nielson, and H. R. Nielson. Control flow analysis for the π -calculus. In *Proc. CONCUR'98*, number 1466 in *Lecture Notes in Computer Science*, pages 84–98. Springer-Verlag, 1998.
4. C. Bodei, P. Degano, F. Nielson, and H. R. Nielson. Static analysis of processes for no read-up and no write-down. In *Proc. FOSSACS'99*, number 1578 in *Lecture Notes in Computer Science*, pages 120–134. Springer-Verlag, 1999.
5. G. Boudol. Asynchrony and the π -calculus. Technical Report 1702, INRIA-Sophia Antipolis, 1992.
6. 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.
7. D. Denning. Certification of programs for secure information flow. *Communications of the ACM*, 20:504–513, 1977.
8. Riccardo Focardi, Anna Ghelli, and Roberto Gorrieri. Using non interference for the analysis of security protocols. In *Proceedings of DIMACS Workshop on Design and Formal Verification of Security Protocols*, 1997.
9. Riccardo Focardi and Roberto Gorrieri. A classification of security properties for process algebras. *Journal of Computer Security*, 3(1), 1995.
10. Riccardo Focardi and Roberto Gorrieri. The compositional security checker: A tool for the verification of information flow security properties. *IEEE Transactions on Software Engineering*, 23, 1997.
11. Riccardo Focardi and Roberto Gorrieri. Non interference: Past, present and future. In *Proceedings of DARPA Workshop on Foundations for Secure Mobile Code*, 1997.
12. C. Fournet, G. Gonthier, J.J. Levy, L. Marganet, and D. Remy. A calculus of mobile agents. In U. Montanari and V. Sassone, editors, *CONCUR: Proceedings of the International Conference on Concurrency Theory*, volume 1119 of *Lecture Notes in Computer Science*, pages 406–421, Pisa, August 1996. Springer-Verlag.
13. R. Reitmas G. Andrews. An axiomatic approach to information flow in programs. *ACM Transactions on Programming Languages and Systems*, 2(1):56–76, 1980.
14. J. A. Goguen and J. Meseguer. Security policies and security models. In *IEEE Symposium on Security and privacy*, 1992.
15. Nevin Heintz and Jon G. Riecke. The SLam calculus: Programming with secrecy and integrity. In *Conference Record of the ACM Symposium on Principles of Programming Languages*, San Diego, January 1998.

16. Matthew Hennessy and James Riely. Information flow vs. resource access in the asynchronous pi-calculus. Technical report 2000:03, University of Sussex, 2000. Available from <http://www.cogs.susx.ac.uk/>.
17. 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.
18. Kohei Honda, Vasco Vasconcelos, and Nobuko Yoshida Honda. Secure information flow as typed process behaviour. In *Proceedings of European Symposium on Programming (ESOP) 2000*. Springer-Verlag, 2000.
19. R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
20. R. Milner, J. Parrow, and D. Walker. Mobile logics for mobile processes. *Theoretical Computer Science*, 114:149–171, 1993.
21. R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 24:83–113, 1984.
22. 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.
23. Benjamin C. Pierce and David N. Turner. Pict: A programming language based on the pi-calculus. Technical Report CSCI 476, Computer Science Department, Indiana University, 1997. To appear in *Proof, Language and Interaction: Essays in Honour of Robin Milner*, Gordon Plotkin, Colin Stirling, and Mads Tofte, editors, MIT Press.
24. James Riely and Matthew Hennessy. Resource access control in systems of mobile agents (extended abstract). In *Proceedings of 3rd International Workshop on High-Level Concurrent Languages*, Nice, France, September 1998. Full version available as Computer Science Technical Report 2/98, University of Sussex, 1997. Available from <http://www.cogs.susx.ac.uk/>.
25. 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.
26. P.Y.A. Ryan and S.A. Schneider. Process algebra and non-interference. In *CSFW 12*. IEEE, 1997.
27. 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.