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, $\sigma[\![$

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:

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

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-interfer*

FIGURE 2 Labelled Transition Semantics

(L-OUT)   (L-IN)

$$\frac{}{a!\langle v\rangle \xrightarrow{a!v} \mathbf{0}} \qquad \frac{}{a?(X)\, P \xrightarrow{(\tilde c:\tilde C)a?v} P\{v/X\}} \ \ \tilde c \notin \mathsf{fn}(P)$$

(L-OPEN)

$$\frac{P \xrightarrow{(\tilde c:\tilde C)a!v} P'}{(\mathsf{new}\, b:B)\ P \xrightarrow{(b:B)(\tilde c:\tilde C)a!v} P'}\ \ \begin{array}{l} b \neq a \\ b \in \mathsf{fn}(v)\end{array}$$

(L-COM)

$$\frac{P \xrightarrow{\alpha} P',\ \ Q \xrightarrow{\overline{\alpha}} Q'}{P \mid Q \xrightarrow{\tau} (\mathsf{new}\,\mathcal{E}(\alpha))\ (P' \mid Q')}$$

(L-EQ)

$$\frac{}{\mathsf{if}\ u = u\ \mathsf{then}\ P\ \mathsf{else}\ Q \xrightarrow{\tau} P} \qquad\qquad \frac{}{\mathsf{if}\ u = w\ \mathsf{then}\ P\ \mathsf{else}\ Q \xrightarrow{\tau} Q}\ \ u \neq w$$

(L-CTXT)

$$\frac{P \xrightarrow{\mu} P'}{\begin{array}{l} *P \xrightarrow{\mu} *P \mid P' \\ \sigma[\![P]\!] \xrightarrow{\mu} \sigma[\![P']\!]\end{array}} \qquad\qquad \frac{P \xrightarrow{\mu} P'}{\begin{array}{l} P \mid Q \xrightarrow{\mu} P' \mid Q \\ Q \mid P \xrightarrow{\mu} Q \mid P'\end{array}}\ \ \mathsf{bn}(\mu) \notin \mathsf{fn}(Q)$$

$$\frac{P \xrightarrow{\mu} P'}{(\mathsf{new}\, a:A)\ P \xrightarrow{\mu} (\mathsf{new}\, a:A)\ P'}\ \ a \notin \mathsf{n}(\mu)$$

a set of *basic values* BV_σ; we use bv to range over base values. We require that all syntactic sets be disjoint.

The input construct '$u?(X:A)\,P$' binds all variables in the pattern X while the construct '$(\mathsf{new}\, a:A)\,P$' binds names and associated with these. We have the usual notions of free and bound names and variables, α-equivalence and substitution. We identify terms up to α-equivalence. Let $\mathsf{fn}(P)$ and $\mathsf{fv}(P)$ denote the set of free names and variables, respectively, of the term P. We use '$P\{v/X\}$' to denote the substitution of the identifiers occurring in the value v for the variables occurring in the pattern X. For '$P\{v/X\}$' to be well-defined X and v must have the same structure; to avoid unnecessary complications we assume that a variable can occur at most once in a pattern. The binding constructs have types associated with them; these will be explained in Section 3 but are ignored for the moment. In general these types (and the various security annotations) will be omitted from terms unless they are relevant to the discussion at hand.

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

this end, *Pre-capabilities* and *pre-types* are defined as follows:

cap	$::=$	*Pre-Capability*
	$w_\sigma \langle A \rangle$	σ-level process can write values with type A
	$r_\sigma \langle A \rangle$	σ-level process can read values with type A

A	$::=$	*Pre-Type*
	$\mathbf{B}_\sigma$	Base type
	$\{cap_1, \dots, cap_k\}$	Resource type $(k \geq 0)$
	$(A_1, \dots, A_k)$	Tuple type $(k \geq 0)$

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

FIGURE 3 Runtime Errors

(E-RD) $\rho[\![a?(X)\,P]\!] \xmapsto{\Sigma} err$ if $\sigma \preceq \rho$ implies for all A, $r_\sigma\langle A\rangle \notin \Sigma(a)$

(E-WR$_1$) $\rho[\![a!\langle v\rangle]\!] \xmapsto{\Sigma} err$ if $\sigma \preceq \rho$ implies for all A, $w_\sigma\langle A\rangle \notin \Sigma(a)$

(E-WR$_2$) $\rho[\![a!\langle v\rangle]\!] \xmapsto{\Sigma} err$ if $bv \in v$, $bv \in \mathbf{B}_\sigma$ and $\sigma \npreceq \rho$

(E-STR) $\dfrac{P \xmapsto{\Sigma} err}{P \mid Q \xmapsto{\Sigma} err}$ $\dfrac{P \xmapsto{\Sigma} err}{\rho[\![P]\!] \xmapsto{\Sigma} err}$ $\dfrac{P \equiv Q,\ P \xmapsto{\Sigma} err}{Q \xmapsto{\Sigma} err}$

$\dfrac{P \xmapsto{\Sigma, a\,:\,A} err}{(\mathsf{new}\,n:A)\ P \xmapsto{\Sigma} err}$

$c!\langle lh\rangle$ although intuitively it involves a security leak; a low security agent can read from c a channel which has at least some capability which should only be accessible to high security principals. However it is straightforward to place it in a context in which a security leak occurs: $c!\langle lh\rangle \mid bot[\![c?(x)\,x!\langle v\rangle]\!]$. Thus our typing system will also be required to rule out such processes.   □

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

For each ρ, let $RType_\rho$ be the least set that satisfies:

(RT-WR)
$$\frac{\mathrm{A} \in RType_\sigma}{\{\mathsf{w}_\sigma\langle\mathrm{A}\rangle\} \in RType_\rho} \ \sigma \preceq \rho$$

(RT-RD)
$$\frac{\mathrm{A} \in RType_\sigma}{\{\mathsf{r}_\sigma\langle\mathrm{A}\rangle\} \in RType_\rho} \ \sigma \preceq \rho$$

(RT-BASE)
$$\frac{}{\mathbf{B}_\sigma \in RType_\rho} \ \sigma \preceq \rho$$

(RT-WRRD)
$$\frac{\mathrm{A} \ \in RType_\sigma \quad \mathrm{A}' \in RType_{\sigma'}}{\{\mathsf{w}_\sigma\langle\mathrm{A}\rangle, \mathsf{r}_{\sigma'}\langle\mathrm{A}'\rangle\} \in RType_\rho} \ \begin{matrix}\sigma \ \preceq \rho \\ \sigma' \preceq \rho \\ \mathrm{A} <: \mathrm{A}'\end{matrix}$$

(RT-TUP)
$$\frac{\mathrm{A}_i \in RType_\rho \quad (\forall i)}{(\mathrm{A}_1, \dots, \mathrm{A}_k) \in RType_\rho}$$

Let $RType$ be the union of the kinds $RType_\rho$ over all ρ. $\square$

Note that if $\sigma \preceq \rho$ then $RType_\sigma \subseteq RType_\rho$. Intuitively, low level values are accessible to high level processes. However the converse is not true. For example $\mathsf{w}_{\mathsf{top}}\langle\rangle \in RType_{\mathsf{top}}$ but $\mathsf{w}_{\mathsf{top}}\langle\rangle$ is not in $RType_{\mathsf{bot}}$. Note also that there is no relation between subtyping and accessibility at a given security level. For example:

$\mathsf{w}_{\mathsf{bot}}\langle\rangle \in RType_{\mathsf{bot}}$ and $\{\mathsf{w}_{\mathsf{bot}}\langle\rangle, \mathsf{r}_{\mathsf{top}}\langle\rangle\} <: \mathsf{r}_{\mathsf{bot}}\langle\rangle$ but $\{\mathsf{w}_{\mathsf{bot}}\langle\rangle, \mathsf{w}_{\mathsf{top}}\langle\rangle\} \notin RType_{\mathsf{bot}}$
$\mathsf{r}_{\mathsf{bot}}\langle\rangle \in RType_{\mathsf{bot}}$ and $\mathsf{r}_{\mathsf{bot}}\langle\rangle <: \mathsf{r}_{\mathsf{top}}\langle\rangle$ but $\mathsf{r}_{\mathsf{top}}\langle\rangle \notin RType_{\mathsf{bot}}$

The compatibility requirement between read and write capabilities in a type (RT-WRRD), in addition to the typing implications discussed in [23], also has security implications. For example suppose $\mathsf{r}_{\mathsf{bot}}\langle\mathbf{B}_\sigma\rangle$ and $\mathsf{w}_{\mathsf{top}}\langle\mathrm{B}\rangle$ 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 $\mathrm{B} <: \mathbf{B}_\sigma$ implies that B must also be $\mathbf{B}_{\mathsf{bot}}$. So although high level processes may write to the channel they may only write low level values.

Remark. Most of the restrictions imposed on types are essential to achieving Subject Reduction, but a few are not. First, Subject Reduction still holds if we weaken (U-WR) to: $\mathsf{w}_\sigma\langle\mathrm{A}\rangle <: \mathsf{w}_\rho\langle\mathrm{B}\rangle$ if $\mathrm{B} <: \mathrm{A}$ and $\sigma \preceq \rho$. Were we to adopt this rule, it would be true that every process typable at level σ would also be typable at level ρ, for $\sigma \preceq \rho$. Given our definition, this is not true. Nonetheless, every process typable at σ can be trivially rewritten so that it is typable at ρ given our definition (one must simply surround output actions with explicit security restrictions). We have

FIGURE 4 Typing Rules

(T-ID)	(T-BASE)	(T-TUP)
$\dfrac{\Gamma(u) <: A}{\Gamma \vdash u : A}$	$\dfrac{bv \in \mathbf{B}_\sigma}{\Gamma \vdash bv : \mathbf{B}_\sigma}$	$\dfrac{\Gamma \vdash v_i : A_i \quad (\forall i)}{\Gamma \vdash (v_1, \ldots, v_k) : (A_1, \ldots, A_k)}$

$$
\frac{\Gamma, X : A \vdash^\sigma P \qquad \Gamma \vdash u : r_\sigma\langle A \rangle}{\Gamma \vdash^\sigma u?(X : A)\, P} \text{(T-IN)}
\qquad
\frac{\Gamma \vdash u : w_\sigma\langle A \rangle \qquad \Gamma \vdash v : A}{\Gamma \vdash^\sigma u!\langle v \rangle} \text{(T-OUT)}
$$

(T-EQ)
$$
\frac{\Gamma \vdash u : A, v : B \qquad \Gamma \vdash^\sigma Q \qquad \Gamma \sqcap \{u : B, v : A\} \vdash^\sigma P}{\Gamma \vdash^\sigma \text{if } u = v \text{ then } P \text{ else } Q}
$$

(T-SR)	(T-NEW)	(T-STR)
$\dfrac{\Gamma \vdash^{\sigma \sqcap \rho} P}{\Gamma \vdash^\sigma \rho[\![P]\!]}$	$\dfrac{\Gamma, a : A \vdash^\sigma P}{\Gamma \vdash^\sigma (\text{new } a : A)\, P}$	$\dfrac{\Gamma \vdash^\sigma P, Q}{\Gamma \vdash^\sigma P \mid Q, \; *P, \; \mathbf{0}}$

adopted the stronger rule because it is necessary in the next section and results in no substantive loss of expressivity.

Second, we have limited types to contain at most one read and one write capability. We have done so to simplify the proofs, particularly in the next section. This clearly results in a loss of expressiveness. We have yet to find, however, a compelling example that requires a resource to have more than one read or one write capability. It is usually sensible to simply take the meet.   □

PROPOSITION 3.2. *For every ρ, RType$_\rho$ is a preorder with respect to $<:$, with both a partial meet operation $\sqcap$ and a partial join $\sqcup$.*

Proof. Straightforward adaptation of Proposition 6.2 of [23]. The partial operations $\sqcap$ and $\sqcup$ are first defined by structural induction on types. Typical clauses are

$$
\begin{aligned}
r_\sigma\langle A \rangle \sqcap r_{\sigma'}\langle A' \rangle &= r_{\sigma \sqcap \sigma'}\langle A \sqcap A' \rangle \\
w_\sigma\langle A \rangle \sqcap w_\sigma\langle A' \rangle &= w_\sigma\langle A \sqcup A' \rangle \\
r_\sigma\langle A \rangle \sqcup r_{\sigma'}\langle A \rangle &= r_{\sigma \sqcup \sigma'}\langle A \sqcup A' \rangle \\
w_\sigma\langle A \rangle \sqcup w_\sigma\langle A' \rangle &= w_\sigma\langle A \sqcap A' \rangle
\end{aligned}
$$

One can then show, by induction on the definitions, that:

A $\in$ *RType$_\rho$* and A $\in$ *RType$_{\rho'}$* implies A $\sqcap$ B $\in$ *RType$_{\rho \sqcap \rho'}$* and A $\sqcup$ B $\in$ *RType$_{\rho \sqcup \rho'}$*.

Finally it is straightforward to show that $\sqcap$ and $\sqcup$, defined in this manner, are indeed partial meet and partial join operators.   □

We now discuss the typing system, which is defined using restricted security policies, called type environments. A *type environment* is a finite mapping from identifiers (names and variables) to types. We adopt some standard notation. For example, 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 $<:$ together with the partial operators $\sqcap$ and $\sqcup$ may also be extended to environments. For example $\Gamma <: \Delta$ if for all u in the domain of Δ, $\Gamma(u) <: \Delta(u)$. The partial meet enables us to define more subtle extensions. For example $\Gamma \sqcap \{u : A\}$ may be defined even if u is already in the domain of Γ. It is well defined when $\Gamma(u) \sqcap A$ exists, in which case it maps u to this type. We will normally abbreviate the simple environment $\{u : A\}$ to $u : A$ and moreover use $v : A$ to denote its obvious generalisation to values; this is only well-defined when the value v has the same structure as the type A.

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

Intuitively '$\Gamma \vdash^{\sigma} P$' 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 [23, 21] 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. One might expect that the judgements for values would need to ensure that a value written to a channel be accessible at the appropriate security level. This job, however, is already handled by our definition of types. For example, in order for $\mathsf{w}_{\sigma}\langle A \rangle$ to be a type, A must be a type accessible to σ.

The typing system enjoys many expected properties, the proof of which we leave to the reader.

PROPOSITION 3.3.

- (SPECIALIZATION) $\Gamma \vdash v : A$ *and* $A <: B$ *then* $\Gamma \vdash v : B$
- (WEAKENING) $\Gamma \vdash^{\sigma} P$ *and* $\Delta <: \Gamma$ *then* $\Delta \vdash^{\sigma} P$
- (RESTRICTION) $\Gamma, u : A \vdash^{\sigma} P$ *and* $u \notin \mathsf{fv}(P) \cup \mathsf{fn}(P)$ *implies* $\Gamma \vdash^{\sigma} P$.   □

The main technical tool required for Subject Reduction is, as usual, a substitution result.

LEMMA 3.4 (SUBSTITUTION). *If $\Gamma \vdash v : A$ then*

- $\Gamma \vdash$

(T-STR), followed by (T-NEW), gives the required $\Gamma \vDash^\sigma (\mathsf{new}\,\mathcal{E}(\alpha))\,(P'\,|\,Q')$. $\qquad\square$

We can now prove the first main result:

THEOREM 3.6 (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* $Q \not\xmapsto{\Gamma}\;err$

Proof. By Subject Reduction we know that $\Gamma \vDash^{\mathsf{top}} Q$ and therefore it is sufficient to prove that $\Gamma \vDash^{\mathsf{top}} Q$ implies $Q \not\xmapsto{\Gamma}\;err$. In fact we prove the contrapositive, $Q \xmapsto{\Gamma}\;err$ implies $\Gamma \not\vDash^{\mathsf{top}} Q$ by induction on the definition of $Q \xmapsto{\Gamma}\;err$.

This is a straightforward inductive proof on the derivation of $Q \xmapsto{\Gamma}\;err$. For example consider the case (E-RD). Suppose that $\rho[\![a?(X)\,P]\!] \xmapsto{\Gamma}\;err$ because $\sigma \preceq \rho$ implies for all A, $\mathsf{r}_\sigma\langle A\rangle \notin \Sigma(a)$. By supposition, we have that $\Gamma(a)$ either has no read capability or it has a read capability at level δ, where $\delta \not\preceq \rho$. In either case, the judgement $\Gamma \not\vDash^\rho a?(X)\,P$ cannot be derived, and therefore $\Gamma \not\vDash^{\mathsf{top}} \rho[\![a?(X)\,P]\!]$ is also underivable. $\qquad\square$

We end this section with a brief discussion on the use of the syntax $\sigma[\![P]\!]$ in our language. We have primarily introduced it in order to discuss typing issues. Having defined our typing system we may now view $\sigma[\![P]\!]$ 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. $\Gamma \vDash^\sigma P$. Technically we can view $\sigma[\![P]\!]$ to be *structurally equivalent* to P, assuming we are working in an environment Γ such that $\Gamma \vDash^\sigma P$. This will be formalised in Section 5.

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. Consider the following system

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

executing in an environment in which h is a top-level read/write channel and hl is a top-level write and bot-level read channel. This system can be well-typed, using *R-types*, so the processes only access resources appropriate to their security level. Nevertheless there is some implicit flow of

top $[\![h?(x)$ if $x = 0$ then bot

Here $\approx^\sigma$ is some form of behavioural equivalence that is sensitive only to behaviour of processes that are σ-level or lower. It turns out that such a result is very dependent on the exact formulation used, as the following example illustrates.

Let A denote the type $\{\mathsf{w}_{\mathsf{bot}}\langle\rangle, \mathsf{r}_{\mathsf{bot}}\langle\rangle\}$ and B denote $\{\mathsf{r}_{\mathsf{bot}}\langle\rangle\}$. Further, let Γ map a and b to A and B, respectively, and n to the type $\{\mathsf{w}_{\mathsf{bot}}\langle\mathrm{A}\rangle, \mathsf{r}_{\mathsf{bot}}\langle\mathrm{A}\rangle\}$. Now consider the terms P and H defined by

$$P \Leftarrow \mathsf{bot}[\![n!\langle a\rangle \mid n?(x:\mathrm{A})\ x!\langle\rangle]\!] \qquad H \Leftarrow \mathsf{top}[\![n?(x:\mathrm{B})\ b?(y)\ \mathbf{0}]\!]$$

It is very easy to check that $\Gamma \Vdash P, H$ and that H is bot-free. Note that in the term $P\,|\,H$ 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*, [18], or *must testing*, [20], 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

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

cannot be matched by $P\,|\,\mathbf{0}$. Using the terminology of [20], $P\,|\,\mathbf{0}$ *guarantees* the test $\mathsf{bot}[\![a?(x)\ \omega!\langle\rangle]\!]$ whereas $P \mid H$ does not.

Even obtaining results with respect to *may testing*, defined in Section 5, is delicate. If we allowed *synchronous* tests then we would also have:

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

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

In the following section, we prove a non-interference result using may testing on processes typable using *I-types*.

5   Noninterference up to May Testing

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 $\Gamma \Vdash^\sigma \omega!\langle\rangle$ 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.

DEFINITION 5.1. We write $T\Downarrow$ if $T \xrightarrow{\tau}^* T'$, where T' has the form $(\mathsf{new}\ \tilde{c})\,(\omega!\langle\rangle \mid T'')$ for some T'' and $\tilde{c}$.   □

Information

FIGURE 5 Context LTS

(C-RED)

$$\frac{P \xrightarrow{\tau} P'}{\Delta \rhd P \xrightarrow{\tau}_\sigma \Delta \rhd P'}$$

(C-OUT)

$$\frac{\Delta \Vdash a : \mathsf{r}_\delta \langle \mathrm{B} \rangle}{\Delta \rhd a!\langle v \rangle \xrightarrow{a!v}_\sigma \Delta \rhd \mathbf{0}} \quad \delta \preceq \sigma$$

(C-IN)

$$\frac{\Delta \Vdash a : \mathsf{w}_\delta \langle \mathrm{B} \rangle \qquad \Delta, \tilde{c} : \tilde{\mathrm{C}} \Vdash v : \mathrm{B}}{\Delta \rhd a?(X : \mathrm{A})\, P \xrightarrow{(\tilde{c}\,:\,\tilde{\mathrm{C}})a?v}_\sigma \Delta, \tilde{c} : \tilde{\mathrm{C}} \rhd P\{v/X\}} \quad \begin{array}{l} \delta \preceq \sigma \\ \tilde{c} \notin \mathsf{fn}(P) \end{array}$$

(C-OPEN)

$$\frac{\Delta \rhd P \xrightarrow{(\tilde{c}\,:\,\tilde{\mathrm{C}})a!v}_\sigma \Delta' \rhd P'}{\Delta \rhd (\mathsf{new}\, b : \mathrm{B})\, P \xrightarrow{(b\,:\,\mathrm{B})(\tilde{c}\,:\,\tilde{\mathrm{C}})a!v}_\sigma \Delta',\ b : \mathrm{B} \rhd P'} \quad \begin{array}{l} b \neq a \\ b \in \mathsf{fn}(v) \end{array}$$

(C-CTXT)

$$\frac{\Delta \rhd P \xrightarrow{\mu}_\sigma \Delta' \rhd P'}{\begin{array}{l} \Delta \rhd *P \xrightarrow{\mu}_\sigma \Delta' \rhd *P \mid P' \\ \Delta \rhd \rho[\![P]\!] \xrightarrow{\mu}_\sigma \Delta' \rhd \rho[\![P']\!] \end{array}}$$

$$\frac{\Delta \rhd P \xrightarrow{\mu}_\sigma \Delta' \rhd P'}{\begin{array}{l} \Delta \rhd P \mid Q \xrightarrow{\mu}_\sigma \Delta' \rhd P' \mid Q \\ \Delta \rhd Q \mid P \xrightarrow{\mu}_\sigma \Delta' \rhd Q \mid P' \end{array}} \quad \mathsf{bn}(\mu) \notin \mathsf{fn}(Q)$$

$$\frac{\Delta \rhd P \xrightarrow{\mu}_\sigma \Delta' \rhd P'}{\Delta \rhd (\mathsf{new}\, a : \mathrm{A})\, P \xrightarrow{\mu}_\sigma \Delta' \rhd (\mathsf{new}\, a : \mathrm{A})\, P'} \quad a \notin \mathsf{n}(\mu)$$

Let T be a test such that $\Gamma \Vdash^\sigma T$. Then P can interact with T by performing the action μ and evolving to P'. As a result of this interaction, the capabilities of the context may be increased, as reflected in Γ'.

The modified LTS is defined in Figure 5 and the rules are straightforward. However note that in the rule (C-OUT) it is understood that the environment already knows the value v being output; it is only in the rule (C-OPEN) where the environment learns new information.

Some properties of this modified LTS are easy to establish. For example in $\Gamma \rhd P \xrightarrow{\mu}_\sigma \Gamma' \rhd P'$ the new environment Γ' is completely determined by Γ and the action μ. If μ is τ then Γ' coincides with Γ; otherwise it is Γ augmented with the type environment $\mathcal{E}(\mu)$, the bound names together with their declared types. For this reason the following Lemma is easily established:

LEMMA 5.4. $\Gamma \rhd P \xrightarrow{\mu}_\sigma \Gamma' \rhd P'$ *and* $\Gamma \Vdash P$ *implies* $\Gamma' \Vdash P'$.

Information Flow vs. Resource

Note that in this Lemma the requirement $\Gamma \Vdash P$ is essential to ensure that if T receives a value v then that value is compatible with the type environment Γ.

May testing is determined by the *traces*, s, t, in $VAct^*$ which processes can perform. Let ϵ represent the empty trace. The notion of complementary actions lifts element-wise to traces, $\bar{s}$. The names in a trace $\mathsf{n}(s)$ is defined as the union of the names in the individual actions; likewise the bound names in a trace $\mathsf{bn}(s)$ is defined as the union of the bound names in the individual actions.

DEFINITION 5.7 (TRACES). Let $\Gamma \rhd P \overset{s}{\Longrightarrow}_\sigma \Gamma' \rhd P'$ be the least relation such that:

(TR-ϵ)

$$\overline{\Gamma \rhd P \overset{\epsilon}{\Longrightarrow}_\sigma}$$

DEFINITION 5.9. (Asynchronous traces) Let $\Gamma \triangleright P \xRightarrow{s}{}^a_\sigma \Gamma' \triangleright Q$ be the least relation which, in addition to the clauses in Definition 5.7, satisfies

$$
\frac{
\begin{array}{l}
(\text{C-AIN}) \\
\Gamma \Vdash a : \mathsf{w}_\delta \langle \mathrm{B} \rangle, \\
\Gamma, \tilde{c} : \tilde{\mathrm{C}} \Vdash v : \mathrm{B}, \\
\Gamma, \tilde{c} : \tilde{\mathrm{C}} \triangleright P \mid \delta[\![a!\langle v \rangle]\!] \xRightarrow{s}{}^a_\sigma \Gamma \triangleright Q
\end{array}
}{
\Gamma \triangleright P \xRightarrow{(\tilde{c}:\tilde{\mathrm{C}})a?v.s}{}^a_\sigma \Gamma' \triangleright Q
} \quad
\begin{array}{l}
\delta \preceq \sigma \\
\tilde{c} \notin \mathsf{fn}(P)
\end{array}
$$

$\square$

The ability to compose asynchronous traces depends on the fact that

- $\rho[\![P]\!] \xrightarrow{(\tilde{c}\,:\,\tilde{C}a!v)} \rho[\![P']\!]$ because $P \xrightarrow{(\tilde{c}\,:\,\tilde{C}a!v)} P'$.

  $\Gamma \Vdash^\sigma \rho[\![P]\!]$ implies $\Gamma \Vdash^{\sigma \sqcap \rho} P$ and so by induction

$$P \equiv_\Gamma (\mathsf{new}\ \tilde{c}\,:\,\tilde{C})\ (\delta[\![a!\langle v\rangle]\!] \mid P')$$

  for some $\delta \preceq \sigma \sqcap \rho$. Using the rules (S-SRNEW)(S-SRSR) and (S-SRPAR) we can then show $\rho[\![P]\!] \equiv_\Gamma (\mathsf{new}\ \tilde{c}\,:\,\tilde{C})\ (\rho \sqcap \delta[\![a!\langle v\rangle]\!] \mid \rho[\![P']\!])$.

$\square$

PROPOSITION 5.12 (TRACE COMPOSITION). *Suppose* $\Gamma \Vdash^\sigma T$.

Proof. The proof is by induction on the derivation of $\Gamma \rhd P \mid H \overset{s}{\Longrightarrow}{}^a_\sigma$. We examine the most interesting cases.

- $\Gamma \rhd P \mid H \overset{\tau}{\longrightarrow}_\sigma \Gamma \rhd R \overset{s}{\Longrightarrow}{}^a_\sigma$.

  The most important case here is when there is communication between P and H. Here $P \overset{\alpha}{\longrightarrow} P'$, $H \overset{\overline{\alpha}}{\longrightarrow} H'$, R is $(\mathsf{new}\,\tilde{c}:\tilde{C})\,(P' \mid H')$, where $\tilde{c}$ are the bound variables in α. There are two possibilities.

  - Output from P to H; α has the form $(\tilde{c}:\tilde{C})a!v$. Let us examine the trace $\Gamma \rhd (\mathsf{new}\,\tilde{c}:\tilde{C})\,(P' \mid H') \overset{s}{\Longrightarrow}{}^a_\sigma$. Somewhere in s the names in $\tilde{c}$ may be exported. In general we can construct a related trace s_c such that $\Gamma, \tilde{c}:\tilde{C} \rhd (P' \mid H) \overset{s_c}{\Longrightarrow}{}^a_\sigma$, with the property that for any Q, $\Gamma, \tilde{c}:\tilde{C} \rhd Q \overset{s_c}{\Longrightarrow}{}^a_\sigma$ implies $\Gamma \rhd Q \overset{s_c}{\Longrightarrow}{}^a_\sigma$; s_c is obtained from s by omitting any bounds $(c:C)$ found on its output actions.

    Now we may apply induction to $\Gamma, \tilde{c}:\tilde{C}(P' \mid H') \rhd \overset{s_c}{\Longrightarrow}{}^a_\sigma$, since $\Gamma \Vdash^{\mathsf{g}} P'$ by Subject Reduction and   to $\Gamma\not\!\!c$

$\Gamma, \tilde{c} : \tilde{C} \triangleright (P \mid \delta[\![a!\langle v \rangle]\!]) \stackrel{s'}{\Longrightarrow}_{\sigma}^{a}$. Again we may now Definition 5.9 to obtain the required $\Gamma, \tilde{c} : \tilde{C} \triangleright P \stackrel{s'}{\Longrightarrow}_{\sigma}^{a}$.

$\square$

Given this technical result, we can now prove the Non-Interference Theorem.

THEOREM (5.3). *If $\Gamma \Vdash^{\sigma} P, Q$ and $\Gamma \Vdash^{\text{top}} H, K$ where H, K are σ-free processes, then:*

$$P \simeq_{\Gamma}^{\sigma} Q \ implies \ P \mid H \simeq_{\Gamma}^{\sigma} Q \mid K.$$

Proof. To establish the result, it is sufficient to show that $P \simeq_{\Gamma}^{\sigma} P \mid H$. In fact by Theorem 5.14 it is sufficient to show $\Gamma \triangleright P \stackrel{s}{\Longrightarrow}_{\sigma}^{a}$ implies $\Gamma \triangleright P \mid H \stackrel{s}{\Longrightarrow}_{\sigma}^{a}$, which is immediate, and $\Gamma \triangleright P \mid H \stackrel{s}{\Longrightarrow}_{\sigma}^{a}$ implies $\Gamma \triangleright P \stackrel{s}{\Longrightarrow}_{\sigma}^{a}$; this follows from the previous Proposition. $\square$

Note that the requirement that P, Q be well-typed processes at level σ is necessary for this result to be true. For example consider the process P defined by $h?(x)\,l?y.\,\mathbf{0}$ in an environment Γ in which h, l are high-level and low-level resources respectively. Then $P \simeq_{\Gamma}^{\text{bot}} \mathbf{0}$. However $P \mid H \not\simeq_{\Gamma}^{\text{bot}} H$, where H is the high-level process $h!\langle\rangle$.

6 Conclusions and Related Work

In this paper we have proposed simple typing systems for enforcing a variety of security properties for the *security π-calculus*. The types are obtained by adding security levels to the standard input/output types of the π-calculus, [21, 23]. The first typing system, based on *R-Types*, is designed with resource access control in mind; the security level of a resource (or more formally a capability on a resource) dictates the security clearanceP

a high-level process reads a value from a low-level channel

ticated type system is used to control information flow. The judgements
in their system take

[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. Marganget, 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

retical Computer Science, 114:149–171, 1993.

[20] R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 24:83–113, 1984.

[21] 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.

[22] 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.

[23] 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/`.

[24] 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.

[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.

[28] Nobuko Yoshida. Graph types for monadic mobile processes. In *FSTTCS*, volume 1180, pages 371–386. Springer-Verlag, 1996.