

# A Simple View of Type-Secure Information Flow in the $\pi$ -Calculus

François Pottier

INRIA

E-mail: Francois.Pottier@inria.fr

## Abstract

*One way of enforcing an information flow control policy is to use a static type system capable of guaranteeing a noninterference property. Noninterference requires that two processes with distinct “high”-level components, but common “low”-level structure, cannot be distinguished by “low”-level observers. We state this property in terms of a rather strict notion of process equivalence, namely weak barbed reduction congruence.*

*Because noninterference is not a safety property, it is often regarded as more difficult to establish than a conventional type safety result. This paper aims to provide an elementary noninterference proof in the setting of the  $\pi$ -calculus. This is done by reducing the problem to subject reduction – a safety property – for a nonstandard, but fairly natural, extension of the  $\pi$ -calculus, baptized the  $\langle\pi\rangle$ -calculus.*

## 1 Introduction

*Information flow analysis* consists in analyzing a program so as to determine how its output *depends* on the inputs it is given; or, more generally, how its observable behavior depends on the stimuli provided by its environment. Such an analysis allows static enforcement of so-called “information flow control” security policies, which prevent secret data from being leaked on public communication channels, or unreliable information from affecting critical decisions. More generally, dependency analysis is at the heart of several program transformation techniques [1].

Information flow analysis, reformulated as a type inference problem, has been heavily studied in the past few years, especially in the area of high-level, sequential languages; for references, see e.g. [20]. Recently, researchers [9, 11, 22, 10, 12, 27] have begun further extending this study to low-level, concurrent calculi, such as the  $\pi$ -calculus. This is made clearly worthwhile by the fact that these calculi allow modeling distributed computing systems

and protocols, which today are at the heart of many real-world security concerns.

An information flow analysis usually comes in the form of a type system, together with a soundness proof. The former is typically (although not necessarily) derived from an existing system, whose types are augmented with annotations taken from a fixed *security lattice*  $\mathcal{L}$  [3]. The soundness theorem gives a *noninterference* statement: it asserts that, if an appropriate type judgement holds, then no change in the high-security inputs of a process can affect a low-security observer. (Here, the usual *subject reduction* theorem becomes of rather secondary importance; nevertheless, it is often useful as a tool in the noninterference proof.) Such a statement must rely on some (preferably standard) notion of *process equivalence*, giving rise to a variety of choices.

Examining the theoretical results offered by existing works, we find some of them to be somewhat lacking in strength or simplicity. Hepburn and Wright [10] propose a type system which enjoys subject reduction, but do not establish any kind of security property. Hennessy and Riely [9] and Hennessy [8] base their noninterference statements on *may*- or *must*-testing equivalence; a statement based on a stronger notion of process equivalence would seem preferable. Sewell and Vitek [22] do not prove noninterference. They instead define a so-called “causal flow property” concerning the traces of a process in a “colored” labelled transition semantics. Because both the property and the colored semantics are *ad hoc*, it is difficult to determine exactly what is being guaranteed. Honda, Yoshida, Vasconcelos and Berger, in a number of papers [11, 12, 27] propose several advanced type systems, which allow exploiting linearity and deadlock-freedom properties to refine the information flow analysis. Their noninterference result, which is stated in terms of a weak bisimulation, requires a rather involved proof [26], due to the need to keep track of liveness properties. Furthermore, the type systems presented in [12, 27] have a rigid type structure (due to so-called “IO alternation” and “sequentiality” constraints), which is meant to allow encoding *sequential* computation only. They would need to be extended if they were to ac-

cept every  $\pi$ -calculus process that is well-typed in a simple type discipline.

In this paper, we wish to define a simple type-based information flow analysis, roughly equivalent in expressive power to Hennessy and Riely’s [9], to Sewell and Vitek’s [22], or to the nonlinear fragment of Honda *et al.*’s [11], and to give an *elementary* noninterference proof for it. By showing how straightforward it is to establish noninterference in this simple setting, we hope to convey some useful insights, and to later facilitate the understanding of more advanced type systems.

Our noninterference result will be stated in terms of bisimulation equivalence, rather than (say) may-testing equivalence, so as to identify more processes as insecure. (See Focardi and Gorrieri’s work [7] for a discussion of this issue.) In short, relying on a bisimulation equivalence allows detecting information leaks caused by *contention* between “low”-level and “high”-level processes waiting on a single channel  $x$ . This, we argue, sounds desirable, because, in practice (that is, under a reasonable fairness assumption), a “low”-level process may detect the presence of a “high”-level one by noticing that a message that was sent to it on channel  $x$  was not received after a certain amount of time.

Of course, one may object that this phenomenon is only a particular kind of *timing leak*. Our type system does not, in general, detect such leaks. Neither does it detect *probabilistic* information leaks, because we use notions of process equivalence based on *possible*, rather than *likely*, behaviors. Still, we argue that weak bisimulation yields a stronger noninterference result than may-testing equivalence, and we choose to rely on it.

## 2 Overview

We wish to establish a noninterference result, i.e. prove that two (well-typed) processes which differ only in some “high”-level components behave identically when observed through “low”-level channels. Note that this is *not* a so-called *safety* property, because it requires examining the traces of *two* processes, rather than of a single one. (See e.g. [15] for more details.) This explains why noninterference is considered more difficult to establish than a conventional type safety result. Our approach consists in reducing noninterference to a simple *subject reduction* property (i.e. to a safety property) for a *nonstandard* extension of the  $\pi$ -calculus. It is inspired by previous joint work with Vincent Simonet [21].

Our calculus, baptized the  $\langle\pi\rangle$ -calculus, describes the independent execution of a *pair* of processes (which we arbitrarily index by  $\{1, 2\}$ ), while keeping track of their *shared* sub-processes. For instance, the  $\langle\pi\rangle$ -calculus process  $\mathbf{P} \mid \langle\mathbf{Q}\rangle_1 \mid \langle\mathbf{R}\rangle_2$  represents the pair  $(\mathbf{P} \mid \mathbf{Q}, \mathbf{P} \mid \mathbf{R})$ , while explicitly recording the fact that the sub-process  $\mathbf{P}$

is shared. A sub-term of the form  $\langle\mathbf{P}\rangle_i$  indicates that  $\mathbf{P}$  is present only in the process of index  $i$ , and absent in the one of index  $j$ , where  $\{i, j\} = \{1, 2\}$ . Brackets cannot be nested. This would not make any sense, since we only wish to encode two standard processes, not more.

Because we are interested in comparing two processes that differ only in their “high”-level components, and because we encode these differences using “bracket” constructors, any sub-process that appears within brackets will be considered a “high”-level process. Conversely, processes that are shared will be considered “low”-level.

One may wonder why it is not sufficient to employ only one kind of brackets, i.e. a single construction  $\langle\cdot\rangle$ . For instance, the process  $\bar{x} \mid \langle x.\mathbf{P} \rangle$  would represent the pair of processes  $(\bar{x}, \bar{x} \mid x.\mathbf{P})$ , while encoding the fact that  $x.\mathbf{P}$  is considered a “high”-level process. However, such a language would not be rich enough to express reduction; indeed, the pair of processes above evolves to  $(\bar{x}, \mathbf{P})$ , which cannot be represented in the single-bracket syntax – hence the need for two distinct bracket constructs.

In Sect. 3, we equip the  $\langle\pi\rangle$ -calculus with an operational semantics which reflects that of the standard  $\pi$ -calculus, but preserves sharing information. Reduction may take place outside brackets (meaning that a common reduction step is performed), inside brackets (meaning that some process performs an independent, “high”-level step), or at bracket boundaries. In the latter case, we discard sharing information by reducing  $\mathbf{P}$  to  $\langle\mathbf{P}\rangle_1 \mid \langle\mathbf{P}\rangle_2$ . However, this is allowed only if some communication step cannot otherwise take place; gratuitous loss of sharing is forbidden. (This is a crucial point; because we consider any process that appears under brackets as “high”-level, the semantics should not cause brackets to appear needlessly, lest the precision of the information flow analysis be compromised.) We begin our formal development by showing how the  $\langle\pi\rangle$ -calculus relates to the standard  $\pi$ -calculus.

In Sect. 4, we equip the  $\langle\pi\rangle$ -calculus with a type system. It is very similar to Pierce and Sangiorgi’s type system for the  $\pi$ -calculus [18], but is extended with *security annotations*, a standard notion in information flow analyses; see e.g. [11, 22]. We prove that it enjoys a *subject reduction* property.

Sect. 5 shows how these results combine to yield a weak-bisimulation-based noninterference property. In short, the bisimulation diagram naturally arises by design of the  $\langle\pi\rangle$ -calculus, while preservation of “low”-level bars is a simple consequence of the typing hypothesis. Then, for convenience, Sect. 6 rephrases our results using “colored” (rather than “bracket”) notation for processes.

### 3 The $\langle\pi\rangle$ -Calculus

#### 3.1 Presentation

The  $\langle\pi\rangle$ -calculus is an extension of the synchronous polyadic  $\pi$ -calculus [16], whose semantics describes the independent execution of a *pair* of processes, while keeping track of their *shared* sub-processes. (We consider a synchronous variant, because it contains the asynchronous variant as a fragment; see Sect. 7.1 for comments about the asynchronous case.)

**Definition 1** Let  $x, y, z, \dots$  range over a denumerable set  $\mathcal{N}$  of names. Let  $\tilde{y}, \tilde{z}, \dots$  denote vectors of such names. Let  $i$  range over  $\{1, 2\}$ . Normal processes  $M, N, \dots$  and processes  $P, Q, R, \dots$  are given by

$$\begin{aligned} N &::= x(\tilde{y}).P \mid \bar{x}(\tilde{z}).P \mid \mathbf{0} \mid N + N \\ P &::= N \mid (P \mid P) \mid !P \mid \nu x.P \mid \langle P \rangle_i \end{aligned}$$

$x(\tilde{y}).P$  binds  $\tilde{y}$  in  $P$ , and  $\nu x.P$  binds  $x$  in  $P$ . From this information, the usual notions of free names, capture-free substitution, and  $\alpha$ -convertibility are deduced. We write  $\text{fn}(P)$  for the set of free names of a process  $P$ . We identify processes up to  $\alpha$ -conversion.

In the following, **bold** meta-variables denote standard processes, i.e. processes which have no sub-term of the form  $\langle P \rangle_i$ . Throughout the paper, we restrict our attention to processes where every sub-term of the form  $\langle P \rangle_i$  is in fact of the form  $\langle \mathbf{P} \rangle_i$ , i.e. we never nest brackets.

A single  $\langle\pi\rangle$ -calculus process is meant to represent a *pair* of standard  $\pi$ -calculus processes. In particular,  $\langle \mathbf{P} \rangle_1$  stands for the pair  $(\mathbf{P}, \mathbf{0})$ , while  $\langle \mathbf{P} \rangle_2$  stands for  $(\mathbf{0}, \mathbf{P})$ . A  $\langle\pi\rangle$ -calculus process of the form  $\mathbf{P}$  represents *shared* structure: it stands for the pair  $(\mathbf{P}, \mathbf{P})$ . More generally, an arbitrary process  $P$  stands for the pair  $(\pi_1(P), \pi_2(P))$ , where the projection functions  $\pi_1$  and  $\pi_2$  are defined as follows.

**Definition 2** Let  $\{i, j\} = \{1, 2\}$ . The  $i^{\text{th}}$  projection function, written  $\pi_i$ , satisfies the laws  $\pi_i(\langle \mathbf{P} \rangle_i) = \mathbf{P}$  and  $\pi_i(\langle \mathbf{P} \rangle_j) = \mathbf{0}$  and is a homomorphism on other (i.e. standard) process forms. We often write  $\pi_i P$  for  $\pi_i(P)$ . For all  $P$ ,  $\pi_i P$  is a standard process.

Projection may create superfluous null processes, which it is convenient to disregard. We introduce an auxiliary relation for this purpose.

**Definition 3** Let  $\leq_0$  be the smallest reflexive, compatible relation over processes which satisfies the law  $P \leq_0 P \mid \mathbf{0}$ .

(By *compatible*, we mean closed under all contexts.) Viewing  $\pi_i$  as a relation, we will write  $P \pi_i \cdot \geq_0^* \mathbf{P}$  or  $\mathbf{P} \leq_0^* \cdot \pi_i^{-1} P$  to denote  $\pi_i P \geq_0^* \mathbf{P}$ . (The superscript  $*$  denotes the reflexive, transitive closure of a relation; the infix operator  $\cdot$  denotes the composition of relations.)

#### 3.2 Semantics

We now define an operational semantics for the  $\langle\pi\rangle$ -calculus. Its restriction to standard processes will coincide with the standard synchronous polyadic  $\pi$ -calculus [16].

For purely technical reasons, we choose a slightly altered presentation of the latter as a starting point: we turn three structural congruence laws, namely scope extrusion, replication, and spontaneous creation of new bound names, into reduction rules, thus making them irreversible. Proving that these changes do not affect the semantics of the  $\pi$ -calculus is an orthogonal issue, which we do not address here.

For technical convenience, structural congruence is not made transitive. Similarly, evaluation contexts are not allowed to be nested. These presentational choices eliminate some redundancy from our definitions, thus simplifying proofs.

**Definition 4** Structural congruence  $\equiv$  is the smallest reflexive, symmetric, compatible relation over processes such that the following laws hold:

1.  $N + \mathbf{0} \equiv N$ ,  $N_1 + N_2 \equiv N_2 + N_1$ ,  $(N_1 + N_2) + N_3 \equiv N_1 + (N_2 + N_3)$ ;
2.  $P \mid \mathbf{0} \equiv P$ ,  $P_1 \mid P_2 \equiv P_2 \mid P_1$ ,  $(P_1 \mid P_2) \mid P_3 \equiv P_1 \mid (P_2 \mid P_3)$ ;
3.  $\nu x.\nu y.P \equiv \nu y.\nu x.P$ .

**Definition 5** An evaluation context  $E$  is one of  $(\square \mid P)$ ,  $\nu x.\square$  or  $\langle \square \rangle_i$ .

All structural congruence laws are standard. All evaluation contexts are standard as well, except  $\langle \square \rangle_i$ , which allows reduction under brackets.

**Definition 6** The raw one-step reduction relation  $\mapsto$  is given by Fig. 1. We write  $M \# N$  (read:  $M$  and  $N$  may communicate) for  $(M \mid N \mapsto) \vee (N \mid M \mapsto)$ , where  $(P \mapsto)$  is itself a short-hand for  $(\exists P' \quad P \mapsto P')$ . Weak reduction, written  $\Rightarrow$ , is defined as  $(\equiv \cup \mapsto)^*$ .

Rules COMM and CONTEXT are standard. Rules EXTR, REPL and NEW are directed versions of standard congruence rules. The crucial rule is SPLIT, which allows discarding sharing information if required by further reductions. SPLIT can be viewed as a restriction of the following, more liberal rule:

$$\begin{array}{c} \text{SPLIT}' \\ M \mapsto \langle \pi_1 M \rangle_1 \mid \langle \pi_2 M \rangle_2 \end{array}$$

SPLIT' explicitly replaces a shared process with its projections. Thus, it implements our intuition that a  $\langle\pi\rangle$ -calculus process stands for a pair of standard processes. However,

<b>COMM</b> $(M + x(\tilde{y}).P) \mid (N + \tilde{x}(\tilde{z}).Q) \mapsto P[\tilde{z}/\tilde{y}] \mid Q$	
<b>EXTR</b> $\frac{x \notin \text{fn}(Q)}{(\nu x.P) \mid Q \mapsto \nu x.(P \mid Q)}$	<b>REPL</b> $!P \mapsto P \mid !P$
<b>NEW</b> $\mathbf{0} \mapsto \nu x.\mathbf{0}$	<b>CONTEXT</b> $\frac{P \mapsto P'}{E[P] \mapsto E[P']}$
<b>SPLIT</b> $\frac{\pi_i M \# N \quad \{i, j\} = \{1, 2\}}{M \mid \langle N \rangle_i \mapsto \langle \pi_i M \mid N \rangle_i \mid \langle \pi_j M \rangle_j}$	
<b>GLUE</b> $\frac{M \# N}{\langle M \rangle_i \mid \langle N \rangle_i \mapsto \langle M \mid N \rangle_i}$	<b>BREAK</b> $\langle P \mid Q \rangle_i \mapsto \langle P \rangle_i \mid \langle Q \rangle_i$
<b>PUSH</b> $\langle \nu x.P \rangle_i \mapsto \nu x.\langle P \rangle_i$	

**Figure 1. Semantics of the  $\langle \pi \rangle$ -calculus**

because it has no side-condition, it allows sharing information to be discarded at will. Although such a behavior would be perfectly valid as far as the untyped semantics of the  $\langle \pi \rangle$ -calculus is concerned, it would lead to a useless type system. Indeed, as we have said, every process which appears under brackets must remain invisible to “low”-level observers. However, SPLIT’ potentially causes *every* process to appear under brackets. Because the type system must have subject reduction, adopting this rule would force it to typecheck every process under the most restrictive security assumption.

As a result, we must replace SPLIT’ with a restricted version that preserves as much sharing information as possible, i.e. that allows it to be discarded only if some communication step cannot otherwise take place. Thus, we obtain SPLIT, where  $M$  can be split into  $\langle \pi_i M \rangle_i \mid \langle \pi_j M \rangle_j$  only if one of its projections, say  $\pi_i M$ , is able to communicate with some term  $N$  which already appears under a bracket  $\langle N \rangle_i$ .

Rules GLUE and BREAK can be understood as restrictions of a (hypothetical) structural congruence rule:

$$\langle P \mid Q \rangle_i \equiv \langle P \rangle_i \mid \langle Q \rangle_i$$

Again, allowing this equivalence to hold would be correct as far as the untyped semantics is concerned, but would pose a slight technical typing problem. The premise in rule GLUE works around it by allowing brackets to be merged only if

required to allow some communication step. (Have a look at the subject reduction proof for GLUE for more details.)

Lastly, rule PUSH allows pushing brackets down inside  $\nu$  binders. In conjunction with BREAK, this lets them move down to the level of normal processes, where they can be dealt with by SPLIT.

### 3.3 Relating the $\langle \pi \rangle$ -Calculus to the $\pi$ -Calculus

We now relate the  $\langle \pi \rangle$ -calculus to the  $\pi$ -calculus. (Because the latter is a fragment of the former, we do not define it separately.) Our aim is to show that the  $\langle \pi \rangle$ -calculus allows reasoning about the execution of a *pair* of standard  $\pi$ -calculus processes. That is, every reduction of a  $\langle \pi \rangle$ -calculus process represents correct reductions of its projections; conversely, every reduction of a projection can be emulated by reductions of the whole. In other words, the two semantics are in a weak bisimulation relation.

We begin with the easier part, i.e. proving that the projection of every reduction step is a correct reduction step as well. This is immediate; we omit the proof. (All missing proofs can be found in the full version of this paper [19].)

**Lemma 1** *Let  $i \in \{1, 2\}$ . If  $P \Rightarrow P'$ , then  $\pi_i P \Rightarrow \pi_i P'$ .*

We continue with the subtler part, i.e. proving that every correct (congruence or reduction) step performed by a projection  $\pi_i P$  can be emulated by the term  $P$ . There is a slight technical twist: if  $\pi_i P$  is  $P \mid \mathbf{0}$ , which reduces to  $P$  via a structural congruence step, then  $P$  may be of the form  $P \mid \langle Q \rangle_j$ , where  $\{i, j\} = \{1, 2\}$ . In that case,  $P$  will be unable to perform the same step. To account for this, we use the pre-order  $\leq_0^*$  introduced in Definition 3.

**Lemma 2** *Let  $i \in \{1, 2\}$ . If  $P \Rightarrow P'$  and  $P \leq_0^* \cdot \pi_i^{-1} P$ , then there exists some process  $P'$  such that  $P \Rightarrow P'$  and  $P' \leq_0^* \cdot \pi_i^{-1} P'$ .*

The proof, which requires a number of auxiliary lemmas, is given in the full version of this paper [19].

## 4 Typing the $\langle \pi \rangle$ -Calculus

### 4.1 Presentation

We now introduce a type system for the  $\langle \pi \rangle$ -calculus. It extends an existing type system for the  $\pi$ -calculus – namely Pierce and Sangiorgi’s [18], which we choose for its simplicity – with security annotations. Its typing judgements are of the form  $\Gamma \vdash_{\text{pc}} P$ , where  $\text{pc}$  is a *security level*, i.e. a member of a fixed *security lattice*  $\mathcal{L}$ . Such a judgement may be read: *under assumptions  $\Gamma$ ,  $P$  is well-typed and will affect only observers of security clearance  $\text{pc}$  or higher*. It may also be read: *under assumptions  $\Gamma$ , assuming  $P$  gains*

information of level  $pc$  by being executed,  $P$  is well-typed. This formulation explains why this meta-variable is historically named  $pc$  [6]: it is the security level which  $P$  attains simply by virtue of being executed, i.e. the security level associated with its “program counter”.

Even though the security lattice  $\mathcal{L}$  is arbitrary, it is desirable to establish a simple dichotomy between “low” and “high” security levels. Such a distinction allows simple proofs; full generality will be recovered in Sect. 6. To this end, in the present section, we assume given a fixed, downward-closed set  $L \subseteq \mathcal{L}$ . We will view levels within (resp. outside)  $L$  as “low” (resp. “high”).

Noninterference states that two processes which differ only in some high-level sub-terms cannot be distinguished by low-level observers. To achieve this, our type system will guarantee that processes of the form  $\langle \mathbf{P} \rangle_i$  – which we use to encode the differences between two processes – can affect only high-security-level observers. In other words, for  $\Gamma \vdash_{pc} \langle \mathbf{P} \rangle_i$  to hold, we will require  $pc \notin L$ . (See rule T-BRACKET in Fig. 2.) This will be our only use of  $L$  in this section.

As in [18], every channel type  $t$  carries a polarity  $p$  which tells whether the channel may be used for input, output, or both. It is further annotated with a security level  $l \in \mathcal{L}$ , which tells how much information may be obtained by successfully reading from or writing to the channel.

**Definition 7** A polarity  $p$  is one of  $\{-, +, \pm\}$ . Types  $t$  are of the form  $\langle \tilde{t} \rangle_l^p$ .

**Definition 8** Polarities are ordered by  $\pm \leq -, \pm \leq +$ . Types are ordered by

$$\frac{p \leq p' \quad (p' \leq - \Rightarrow \tilde{t} \leq \tilde{t}') \quad (p' \leq + \Rightarrow \tilde{t}' \leq \tilde{t})}{\langle \tilde{t} \rangle_l^p \leq \langle \tilde{t}' \rangle_l^{p'}}$$

(The ordering is extended point-wise to vectors of types.) These definitions can be understood either inductively or co-inductively, yielding finite or infinite types. The choice of one or the other is orthogonal to our concerns; indeed, the subject reduction and noninterference proofs are entirely independent of this issue. We leave it open, to be settled by the analysis designer at a later stage.

Note that our definition of subtyping does *not* allow a channel’s security level to be modified, be it covariantly or contravariantly. In other words, two channel types that are in a subtyping relationship must have the same security level. This (admittedly strong) requirement reflects the fact that information flows *both ways* along a channel, regardless of the direction of messages (i.e. regardless of the channel’s polarity). This property will be used in the proof of Lemma 6, which itself plays a key role in establishing subject reduction.

$\frac{\text{T-RECV} \quad \Gamma(x) \leq \langle \tilde{t} \rangle_{pc}^- \quad \Gamma; \tilde{y} : \tilde{t} \vdash_{pc} P}{\Gamma \vdash_{(pc)} x(\tilde{y}).P}$		
$\frac{\text{T-SEND} \quad \Gamma(x) \leq \langle \tilde{t} \rangle_{pc}^+ \quad \Gamma(\tilde{y}) \leq \tilde{t} \quad \Gamma \vdash_{pc} P}{\Gamma \vdash_{(pc)} \bar{x}(\tilde{y}).P}$		$\text{T-NULL} \quad \Gamma \vdash_{(pc)} \mathbf{0}$
$\frac{\text{T-SUM} \quad \Gamma \vdash_{(pc)} M \quad \Gamma \vdash_{(pc)} N}{\Gamma \vdash_{(pc)} M + N}$		$\frac{\text{T-NORMAL} \quad \Gamma \vdash_{(pc)} N}{\Gamma \vdash_{pc} N}$
$\frac{\text{T-PAR} \quad \Gamma \vdash_{pc} P \quad \Gamma \vdash_{pc} Q}{\Gamma \vdash_{pc} P \mid Q}$	$\frac{\text{T-REPL} \quad \Gamma \vdash_{pc} P}{\Gamma \vdash_{pc} !P}$	$\frac{\text{T-NEW} \quad \Gamma; x : t \vdash_{pc} P}{\Gamma \vdash_{pc} \nu x.P}$
$\frac{\text{T-BRACKET} \quad \Gamma \vdash_{pc} \mathbf{P} \quad pc \notin L}{\Gamma \vdash_{pc} \langle \mathbf{P} \rangle_i}$		$\frac{\text{T-SUB} \quad \Gamma \vdash_{pc} P \quad pc' \leq pc}{\Gamma \vdash_{pc'} P}$

**Figure 2.** Typing rules of the  $\langle \pi \rangle$ -calculus

**Definition 9** The type system of the  $\langle \pi \rangle$ -calculus is defined by Fig. 2. It involves two separate judgement forms. Judgements of the form  $\Gamma \vdash_{(pc)} N$  concern normal processes. Judgements of the form  $\Gamma \vdash_{pc} P$  concern arbitrary processes.

Rules T-RECV and T-SEND require the channel’s security level to match the level attained by the process, namely  $pc$ . Furthermore, T-SUM requires all components of a sum to have matching levels. As a result, if  $\Gamma \vdash_{(pc)} N$  holds, then all channels liable to be read or written by  $N$  must have the same security level. This reflects the fact that information may flow arbitrarily between these channels. As a simple illustration of this fact, consider  $N = x.P + y.Q$ . If, in the presence of  $N$ , a message sent on channel  $x$  is *not* consumed after a while, then its sender knows that some message was available on channel  $y$ , so information flows from  $y$  to  $x$ . (Recall that relying on weak bisimulation amounts, in practice, to formulating a fairness hypothesis.) By symmetry, the converse is also true.

Rule T-SUB allows strengthening the security requirements bearing on a process. This rule applies only to judgements of the form  $\Gamma \vdash_{pc} P$ ; applying it to judgements of the form  $\Gamma \vdash_{(pc)} N$  would break the property that all components of a sum have a common security level.

In T-RECV and T-SEND, the continuation process  $P$  is typed at a security level equal to that of the channel  $x$  (or greater, thanks to T-SUB). This reflects the fact that a successful synchronization at  $x$  yields information which  $P$

may exploit.

## 4.2 Type Preservation

We begin with three easy lemmas, stating that typing is preserved by projection, by structural congruence, and by substitution of names for names. Proofs are omitted.

**Lemma 3** For  $i \in \{1, 2\}$ ,  $\Gamma \vdash_{\text{pc}} P$  implies  $\Gamma \vdash_{\text{pc}} \pi_i P$ .

**Lemma 4**  $\Gamma \vdash_{\text{pc}} P$  and  $P \equiv P'$  imply  $\Gamma \vdash_{\text{pc}} P'$ .

**Lemma 5**  $\Gamma; \tilde{y} : \tilde{t} \vdash_{\text{pc}} P$  and  $\Gamma(\tilde{z}) \leq \tilde{t}$  imply  $\Gamma \vdash_{\text{pc}} P[\tilde{z}/\tilde{y}]$ .

Next comes an important auxiliary lemma, stating that, if two normal processes  $M$  and  $N$  are able to communicate with each other, and if they are typed under a common environment  $\Gamma$ , then they must be typed at the same security level. (Additionally, the lemma states that the reduced process  $R$  is well-typed at that level.) Indeed,  $M$  and  $N$  must share some channel  $x$ , so they must both be typed at  $x$ 's security level.

**Lemma 6** Assume  $\Gamma \vdash_{(\text{pc}_1)} M$  and  $\Gamma \vdash_{(\text{pc}_2)} N$ . If  $M \mid N \mapsto R$ , then  $\Gamma \vdash_{\text{pc}} R$  holds, where  $\text{pc} = \text{pc}_1 = \text{pc}_2$ .

*Proof.* Because  $M$  and  $N$  are normal processes, they are irreducible; thus, the reduction  $M \mid N \mapsto R$  must be an instance of rule COMM. Thus,  $M$  and  $N$  must be of the form  $M' + x(\tilde{y}).P$  and  $N' + x(\tilde{z}).Q$ , respectively.

By T-SUM and T-RECV, the first hypothesis yields  $\Gamma \vdash_{(\text{pc}_1)} M'$ ,  $\Gamma(x) \leq \langle \tilde{t}_1 \rangle_{\text{pc}_1}^-$  and  $\Gamma; \tilde{y} : \tilde{t}_1 \vdash_{\text{pc}_1} P$ . By T-SUM and T-SEND, the second one yields  $\Gamma \vdash_{(\text{pc}_2)} N'$ ,  $\Gamma(x) \leq \langle \tilde{t}_2 \rangle_{\text{pc}_2}^+$ ,  $\Gamma(\tilde{z}) \leq \tilde{t}_2$  and  $\Gamma \vdash_{\text{pc}_2} Q$ .

From  $\Gamma(x) \leq \langle \tilde{t}_1 \rangle_{\text{pc}_1}^-$  and  $\Gamma(x) \leq \langle \tilde{t}_2 \rangle_{\text{pc}_2}^+$ , we deduce  $\text{pc}_1 = \text{pc}_2$  and  $\tilde{t}_2 \leq \tilde{t}_1$ . Take  $\text{pc} = \text{pc}_1 = \text{pc}_2$ . By transitivity of subtyping,  $\Gamma(\tilde{z}) \leq \tilde{t}_1$  holds. Lemma 5 then yields  $\Gamma \vdash_{\text{pc}_1} P[\tilde{z}/\tilde{y}]$ . Using  $\text{pc} = \text{pc}_1 = \text{pc}_2$  and T-PAR, we obtain  $\Gamma \vdash_{\text{pc}} P[\tilde{z}/\tilde{y}] \mid Q$ , that is,  $\Gamma \vdash_{\text{pc}} R$ .  $\square$

These results allow us to establish that typing is preserved by reduction.

**Lemma 7 (Subject Reduction)**  $\Gamma \vdash_{\text{pc}} P$  and  $P \mapsto P'$  imply  $\Gamma \vdash_{\text{pc}} P'$ .

*Proof.* By induction on the derivation of  $P \mapsto P'$ .

Case COMM.  $P$  is of the form  $M \mid N$ . By T-SUB, T-PAR and T-NORMAL, this yields  $\Gamma \vdash_{(\text{pc}_1)} M$  and  $\Gamma \vdash_{(\text{pc}_2)} N$ , where  $\text{pc} \leq \text{pc}_1$  and  $\text{pc} \leq \text{pc}_2$  hold. By Lemma 6 and by T-SUB,  $\Gamma \vdash_{\text{pc}} P'$  holds.

Case SPLIT. The reduction is  $M \mid \langle \mathbf{N} \rangle_i \mapsto \langle \pi_i M \mid \mathbf{N} \rangle_i \mid \langle \pi_j M \rangle_j$ , where  $\{i, j\} = \{1, 2\}$  and  $\pi_i M \# \mathbf{N}$ . Our hypothesis is  $\Gamma \vdash_{\text{pc}} M \mid \langle \mathbf{N} \rangle_i$ . By T-SUB, T-PAR, T-NORMAL and T-BRACKET, this yields  $\Gamma \vdash_{(\text{pc}_1)} M$  and

$\Gamma \vdash_{\text{pc}_2} \mathbf{N}$ , where  $\text{pc} \leq \text{pc}_1$ ,  $\text{pc} \leq \text{pc}_2$ , and  $\text{pc}_2 \notin L$  hold. Because  $L$  is downward-closed, we may assume, without loss of generality, that T-SUB is never used immediately above T-BRACKET; as a result,  $\Gamma \vdash_{(\text{pc}_2)} \mathbf{N}$  holds.  $\pi_i M \# \mathbf{N}$  implies  $M \# \mathbf{N}$ . As a result, we may apply Lemma 6 (either to  $M$  and  $\mathbf{N}$  or to  $\mathbf{N}$  and  $M$ ), yielding  $\text{pc}_1 = \text{pc}_2$ .

It follows that  $\Gamma \vdash_{\text{pc}_2} M$  holds. By Lemma 3, both  $\pi_i M$  and  $\pi_j M$  are well-typed under  $\Gamma$  and  $\text{pc}_2$ . By hypothesis, so is  $\mathbf{N}$ . Thus, considering  $\text{pc}_2 \notin L$ , so is  $\langle \pi_i M \mid \mathbf{N} \rangle_i \mid \langle \pi_j M \rangle_j$ . Because  $\text{pc}_2 \geq \text{pc}$ , the result follows by T-SUB.

Case GLUE. The reduction is  $\langle \mathbf{M} \rangle_i \mid \langle \mathbf{N} \rangle_i \mapsto \langle \mathbf{M} \mid \mathbf{N} \rangle_i$ , where  $\mathbf{M} \# \mathbf{N}$ . Our hypothesis is  $\Gamma \vdash_{\text{pc}} \langle \mathbf{M} \rangle_i \mid \langle \mathbf{N} \rangle_i$ . As above, this yields  $\Gamma \vdash_{(\text{pc}_1)} \mathbf{M}$  and  $\Gamma \vdash_{(\text{pc}_2)} \mathbf{N}$ , where  $\text{pc} \leq \text{pc}_k$  and  $\text{pc}_k \notin L$  hold for  $k \in \{1, 2\}$ . Again, Lemma 6 yields  $\text{pc}_1 = \text{pc}_2$ . Thus,  $\Gamma \vdash_{\text{pc}_2} \langle \mathbf{M} \mid \mathbf{N} \rangle_i$  holds. The result follows by T-SUB.

Case EXTR, REPL, CONTEXT, BREAK, PUSH. Immediate.  $\square$

## 5 Noninterference

Combining the results of Sect. 3.3 and 4.2, we will now derive a noninterference property, expressed in terms of weak barbed reduction congruence [17].

**Definition 10** Let  $\alpha$  range over names and co-names  $(x, \bar{x}, \dots)$ . If  $\alpha$  is  $x$  (resp.  $\bar{x}$ ), then  $\bar{\alpha}$  stands for  $\bar{x}$  (resp.  $x$ ). If  $\alpha$  is  $x$  or  $\bar{x}$ , then  $|\alpha|$  is  $x$ . The predicate  $P \downarrow_\alpha$  (read: the process  $P$  is observable at  $\alpha$ ) is defined as follows:

$$(M + x(\tilde{y}).P) \downarrow_x \qquad (M + \bar{x}(\tilde{y}).P) \downarrow_{\bar{x}}$$

$$\frac{P \downarrow_\alpha \quad E \text{ does not bind } |\alpha|}{E[P] \downarrow_\alpha}$$

$P \downarrow_\alpha$  stands for  $(\exists P' \quad P \Rightarrow P' \wedge P' \downarrow_\alpha)$ .

**Definition 11** Let  $B$  be an arbitrary set of names. A binary relation  $\mathcal{R}$  over processes is a weak  $B$ -simulation if and only if

- $P \mathcal{R} Q \wedge P \Rightarrow P'$  implies  $\exists Q' \quad Q \Rightarrow Q' \wedge P' \mathcal{R} Q'$ ; and
- $|\alpha| \in B$ ,  $P \mathcal{R} Q$  and  $P \downarrow_\alpha$  imply  $Q \downarrow_\alpha$ .

$\mathcal{R}$  is a weak  $B$ -bisimulation if and only if  $\mathcal{R}$  and  $\mathcal{R}^{-1}$  are weak  $B$ -simulations. Two processes are weakly  $B$ -bisimilar if they are related by some weak  $B$ -bisimulation. They are weakly bisimilar if they are  $\mathcal{N}$ -bisimilar.

In this section,  $L$  is fixed, as in Sect. 4. The set of channels which, according to a type environment  $\Gamma$ , do not leak any high-level information is referred to as  $\text{low}(\Gamma)$ . It is defined as follows.

**Definition 12** Given a type environment  $\Gamma$ ,  $\text{low}(\Gamma)$  denotes the largest set  $B \subseteq \mathcal{N}$  such that  $x \in B$  and  $\Gamma(x) = \langle \tilde{i} \rangle_l^p$  imply  $l \in L$ .  $\Gamma$  is said to be an  $L$ -environment if and only if  $\text{low}(\Gamma) = \mathcal{N}$ .

If  $x \in \text{low}(\Gamma)$  holds, then, by rule T-BRACKET, no well-typed process observable at  $x$  can appear under brackets. So, in that case, observability at  $x$  must be preserved by projection. This is expressed by the following simple lemma, whose proof appears in the full version of this paper. (We write  $\Gamma \vdash R$  to indicate that  $\Gamma \vdash_{\text{pc}} R$  holds for some  $\text{pc} \in \mathcal{L}$ .)

**Lemma 8 (Barb Preservation)** Assume  $\Gamma \vdash R$ ,  $R \downarrow_\alpha$  and  $|\alpha| \in \text{low}(\Gamma)$ . Let  $i \in \{1, 2\}$ . Then,  $(\pi_i R) \downarrow_\alpha$  holds.

Then, it is easy to establish that any two processes which are respectively the left- and right-hand projections of a single, well-typed  $\langle \pi \rangle$ -calculus process are barbed bisimilar, provided only low-security barbs are observed.

**Lemma 9 (Barbed Bisimulation)** Let  $\mathbf{P} \mathcal{R}_\Gamma \mathbf{Q}$  hold if and only if, for some  $\langle \pi \rangle$ -calculus process  $R$ , both  $\Gamma \vdash R$  and  $\mathbf{P} \leq_0^* \cdot \pi_1^{-1} R \pi_2 \cdot \geq_0^* \mathbf{Q}$  hold. Then,  $\mathcal{R}_\Gamma$  is a weak  $\text{low}(\Gamma)$ -bisimulation.

*Proof.* We prove that  $\mathcal{R}_\Gamma$  is a weak  $\text{low}(\Gamma)$ -simulation. (To deal with  $\mathcal{R}_\Gamma^{-1}$ , simply swap  $\pi_1$  and  $\pi_2$ .) First, assume  $\mathbf{P} \mathcal{R}_\Gamma \mathbf{Q}$  and  $\mathbf{P} \Rightarrow \mathbf{P}'$ . Then, we have

$$\begin{array}{ccccc} \mathbf{P} & \leq_0^* \cdot \pi_1^{-1} R & \pi_2 \cdot & \geq_0^* \mathbf{Q} \\ \Downarrow & & \Downarrow & \searrow & \\ \mathbf{P}' & \leq_0^* \cdot \pi_1^{-1} R' & \pi_2 \cdot & \mathbf{Q}' & \end{array}$$

The leftmost commutative diagram is given by Lemma 2, the middle one by Lemma 1, and the rightmost one merely stems from the fact that  $\leq_0$  is contained within  $\equiv$ . Furthermore, Lemmas 4 and 7 yield  $\Gamma \vdash R'$ , which shows that  $\mathbf{P}' \mathcal{R}_\Gamma \mathbf{Q}'$  holds.

Next, we check that  $\mathcal{R}_\Gamma$  preserves *strong* barbs, which, given the above, implies that it also preserves weak ones, meeting the second condition of Definition 11. Assume  $\mathbf{P} \mathcal{R}_\Gamma \mathbf{Q}$ ,  $\mathbf{P} \downarrow_\alpha$  and  $|\alpha| \in \text{low}(\Gamma)$ . For some process  $R$ ,  $\Gamma \vdash R$  and  $\mathbf{P} \leq_0^* \cdot \pi_1^{-1} R \pi_2 \cdot \geq_0^* \mathbf{Q}$  hold. As a result of the latter, we must have  $R \downarrow_\alpha$ . By Lemma 8, this implies  $(\pi_2 R) \downarrow_\alpha$ . Because  $\pi_2 R \geq_0^* \mathbf{Q}$  holds,  $\mathbf{Q} \downarrow_\alpha$  follows.  $\square$

As a corollary, we show that, for any process  $R$ , the projections  $\pi_1 R$  and  $\pi_2 R$  are weakly barbed-congruent [17], provided we admit only contexts which, according to the type system, emit only low-security barbs.

**Theorem 1** For any process  $R$ , for any (standard) context  $\mathbf{C}$  such that  $\Gamma \vdash \mathbf{C}[R]$  holds in some  $L$ -environment  $\Gamma$ ,  $\mathbf{C}[\pi_1 R]$  and  $\mathbf{C}[\pi_2 R]$  are weakly bisimilar.

*Proof.* Apply Lemma 9. For  $i \in \{1, 2\}$ ,  $\pi_i \mathbf{C}[R]$  is  $\mathbf{C}[\pi_i R]$ . As a result,  $\mathbf{C}[\pi_1 R]$  and  $\mathbf{C}[\pi_2 R]$  are related by  $\mathcal{R}_\Gamma$ . Thus, they are weakly  $\text{low}(\Gamma)$ -bisimilar. The result follows from  $\text{low}(\Gamma) = \mathcal{N}$ .  $\square$

Our second corollary is in the style of Honda *et al.*'s non-interference claim [11]:

**Theorem 2** Assume  $\Gamma_0 \vdash_{\text{pc}_0} \mathbf{P}_i$ , where  $\text{pc}_0 \notin L$ , holds for  $i \in \{1, 2\}$ . Then, for any context  $\mathbf{C}$  and for any environment  $\Gamma$  such that  $\Gamma \vdash \mathbf{C}[\ ]$  holds under the assumption  $\Gamma_0 \vdash_{\text{pc}_0} [\ ]$ ,  $\mathbf{C}[\mathbf{P}_1]$  and  $\mathbf{C}[\mathbf{P}_2]$  are weakly  $\text{low}(\Gamma)$ -bisimilar.

*Proof.* Because  $\text{pc}_0 \notin L$ , T-BRACKET yields  $\Gamma_0 \vdash_{\text{pc}_0} \langle \mathbf{P}_i \rangle_i$  for  $i \in \{1, 2\}$ . Define  $R$  as  $\mathbf{C}[\langle \mathbf{P}_1 \rangle_1 \mid \langle \mathbf{P}_2 \rangle_2]$ . According to T-PAR and to our hypothesis about  $\mathbf{C}$ ,  $\Gamma \vdash R$  holds. Apply Lemma 9.  $\mathbf{C}[\mathbf{P}_1]$  and  $\mathbf{C}[\mathbf{P}_2]$  are related by  $\mathcal{R}_\Gamma$ . As a result, they are weakly  $\text{low}(\Gamma)$ -bisimilar.  $\square$

Theorem 1 is, in our opinion, more directly useful than Theorem 2. Indeed, the former allows precise reasoning about two processes which have some common structure, while the latter treats  $\mathbf{P}_1$  and  $\mathbf{P}_2$  as entirely separate. However, Theorem 1 is expressed in terms of a nonstandard theoretical tool, namely the  $\langle \pi \rangle$ -calculus. For this reason, we will now reformulate it, with a practical aim: provide a specification of the security guarantees offered by the type system that can be read and written by the programmer.

## 6 Programmer-Oriented Specification

We suggest allowing the programmer to color any subprocess with a security level  $l \in \mathcal{L}$ , indicating that its presence should not affect observers whose security clearance is not at least  $l$ . Thus, we introduce a ‘‘colored’’  $\pi$ -calculus. It is reminiscent of Abadi, Lamson, and Lévy labelled  $\lambda$ -calculus [2] and of Sewell and Vitek’s colored box- $\pi$ -calculus [22]. However, we do not need to define a semantics for it; here, we view it only as a programming notation.

$$\begin{aligned} N &::= x(\tilde{x}).P \mid \bar{x}(\tilde{x}).P \mid \mathbf{0} \mid N + N \\ P &::= N \mid (P \mid P) \mid !P \mid \nu x.P \mid l : P \end{aligned}$$

The only variation with respect to the  $\langle \pi \rangle$ -calculus is that the construct  $\langle \mathbf{P} \rangle_i$  is replaced with the coloring construct  $l : P$ . (Unlike brackets, coloring constructs can be nested.) The type system is modified by replacing rule T-BRACKET with

$$\frac{\text{T-COLOR} \quad \Gamma \Vdash_{\text{pc}} P \quad l \leq \text{pc}}{\Gamma \Vdash_{\text{pc}} l : P}$$

Note that the new type system is no longer parameterized by a set  $L$ . We use  $\Gamma \Vdash_{\text{pc}} P$  to denote its judgements.

Rule T-COLOR simply forces the sub-process  $P$  to be typechecked at a level that equals or exceeds  $l$ . Thus, “colors” in the source program cause the typechecker to enforce additional constraints, which will then guarantee a certain security property, as we will now explain.

Given  $L \subseteq \mathcal{L}$ , we define an *erasure* function  $[\cdot]_L$ , which drops all sub-terms whose color is not a member of  $L$ , and produces a term in the standard  $\pi$ -calculus. The function which strips off all colors, namely  $[\cdot]_{\mathcal{L}}$ , is written  $[\cdot]$ .

**Definition 13** *Let  $[\cdot]_L$  satisfy  $[l : P]_L = [P]_L$  when  $l \in L$ ,  $[l : P]_L = \mathbf{0}$  when  $l \notin L$ , and be a homomorphism on standard process forms.*

Then, our final noninterference theorem states that pruning sub-terms which carry “high” colors does not alter the behavior of a process under a certain typed barbed congruence, whereby a context is allowable only if its type states that it will affect “low” channels only. Note that the meaning of “low” and “high” is parameterized by the choice of  $L$ , of which the type system is now independent.

**Theorem 3 (Noninterference)** *Let  $L$  be a downward-closed subset of  $\mathcal{L}$ . For any process  $R$  of the colored  $\pi$ -calculus, for any (standard) context  $\mathbf{C}$  such that  $\Gamma \Vdash \mathbf{C}[R]$  holds in some  $L$ -environment  $\Gamma$ ,  $\mathbf{C}[[R]]$  and  $\mathbf{C}[[R]_L]$  are weakly bisimilar.*

*Proof.* Define a mapping  $\cdot^*$  from the colored  $\pi$ -calculus into the  $\langle \pi \rangle$ -calculus, which satisfies  $(l : P)^* = P^*$  when  $l \in L$ ,  $(l : P)^* = \langle [P] \rangle_1$  when  $l \notin L$ , and is a homomorphism on standard process forms. Then, the identities  $\pi_1(P^*) = [P]$  and  $\pi_2(P^*) = [P]_L$  hold for any colored process  $P$ . Furthermore, it is easy to check that  $\Gamma \Vdash_{\text{pc}} P$  implies  $\Gamma \vdash_{\text{pc}} P^*$ . As a consequence,  $\Gamma \Vdash \mathbf{C}[R]$  yields  $\Gamma \vdash (\mathbf{C}[R])^*$ . Furthermore,  $(\mathbf{C}[R])^*$  is  $\mathbf{C}[R^*]$ . By Theorem 1,  $\mathbf{C}[\pi_1(R^*)]$  and  $\mathbf{C}[\pi_2(R^*)]$  are weakly bisimilar. By the identities above, these are none other than  $\mathbf{C}[[R]]$  and  $\mathbf{C}[[R]_L]$ .  $\square$

In particular, if  $R$  itself is well-typed within some  $L$ -environment, then every context  $\mathbf{C}$  such that  $\mathbf{C}[R]$  is well-typed will do.

## 7 Discussion

### 7.1 Asynchrony

Our type system allows input (resp. output) channel types  $\langle \tilde{t} \rangle_l^-$  (resp.  $\langle \tilde{t} \rangle_l^+$ ) to be covariant (resp. contravariant) in their *parameters*  $\tilde{t}$ , but both are *invariant* in their security level  $l$  (Definition 8). Intuitively, this is because synchronization causes information to flow not only from the sender of a message to its receiver, but also in the reverse direction. As a result, both processes must have the

*same* notion of “pc”. This is enforced by the invariance of security annotations.

Let us now restrict our interest to the asynchronous fragment of the  $\pi$ -calculus, where every sender is of the form  $\bar{x}\langle \tilde{z} \rangle.\mathbf{0}$ . It may seem that, under this restriction, senders can no longer observe the reception of their messages. This would suggest that making input (resp. output) channel types covariant (resp. contravariant) in their security level is safe. This, however, is not the case; in fact, message reception can still be observed by exploiting contention between receivers. This is illustrated by the following example. Take  $\mathcal{L} = \{\mathbf{L}, \mathbf{H}\}$ , with  $\mathbf{L} \leq \mathbf{H}$ . Take  $L = \{\mathbf{L}\}$ . Consider the typing judgement

$$x : \langle \rangle_{\mathbf{H}}^{\pm}, y : \langle \rangle_{\mathbf{L}}^{\pm}, z : \langle \rangle_{\mathbf{L}}^+ \vdash_{\mathbf{L}} \langle \bar{x} \rangle_2 \mid \bar{y} \mid x.y.\mathbf{0} \mid y.\bar{z}$$

This judgement is *incorrect* in our type system, because the sub-term  $x.y.\mathbf{0}$  is ill-typed. Indeed, listening on channel  $x$ , which has type  $\langle \rangle_{\mathbf{H}}^{\pm}$ , causes “pc” to become  $\mathbf{H}$ . Subsequently, listening on channel  $y$  becomes illegal, since  $y$  has type  $\langle \rangle_{\mathbf{L}}^{\pm}$ . On the other hand, if input channel types were covariant, then  $\langle \rangle_{\mathbf{L}}^{\pm}$  would be a subtype of  $\langle \rangle_{\mathbf{H}}^{\pm}$ , and the judgement would become correct. Yet the two projections of this process, namely

$$\bar{y} \mid x.y.\mathbf{0} \mid y.\bar{z} \quad \text{and} \quad \bar{x} \mid \bar{y} \mid x.y.\mathbf{0} \mid y.\bar{z}$$

are *not* weakly  $\{y, z\}$ -bisimilar. Indeed, the left-hand process must send a message on channel  $z$ , while the right-hand one may choose to never do so. In other words, the presence of  $\bar{x}$  causes contention between two receivers on  $y$ , which can be detected by observing  $z$ . This example shows that channel types must remain invariant in their security annotations, even in the asynchronous fragment of the  $\pi$ -calculus.

### 7.2 Exploiting Linearity

In the example above, the information leak is caused by contention. If the channel  $y$  was linear [14], then no contention would be possible, and it would be safe for the “high”-level process  $x.y.\mathbf{0}$  to receive a signal through the “low”-level channel  $y$ . In fact, under a strong notion of linearity, every communication action on a linear channel must eventually succeed (see e.g. [25]), so its success alone does not carry any information. This fact is pointed out and exploited by Honda *et al.* [11, 12, 27]. It is indeed crucial, in practice, to take advantage of it, because the  $\pi$ -calculus is a low-level programming language, where continuation-passing style is ubiquitous: control is encoded through the use of (often linear) communications, rather than evident in the program’s syntax. Zdancewic and Myers [28] address this issue in the case of a low-level, sequential calculus.

Can our proof approach be extended to deal with linearity information? Let us give a rough sketch of how we envision such an extension. The semantics of the  $\langle\pi\rangle$ -calculus must be modified to disallow linear communications from taking place under brackets. Instead, one should introduce reduction rules akin to the following, which re-discovers sharing:

$$\text{JOIN} \\ \langle\bar{x}.\mathbf{P}\rangle_1 \mid \langle\bar{x}.\mathbf{Q}\rangle_2 \mapsto \bar{x}.\langle\mathbf{P}\rangle_1 \mid \langle\mathbf{Q}\rangle_2 \quad \text{if } x \text{ is linear}$$

These changes allow a continuation  $x.P$  to be triggered without splitting  $P$ . In turn, this will allow the type system to view  $P$  as a “low”-level process, even though the trigger  $\bar{x}$  is sent from a “high”-level process. The anticipated difficulty is in establishing the completeness lemma, i.e. the analogue of Lemma 2. Indeed, proving that JOIN is always applicable requires proving that every linear communication action will eventually succeed. This requires using typing information, whereas, in our current development, the results in Sect. 3.3 were (pleasantly) independent of types. We view this research direction as most promising.

### 7.3 Type Inference

Type inference is no more difficult for this system than for Pierce and Sangiorgi’s original type system [18]. Type inference for (an extended version of) the latter has been studied by Igarashi and Kobayashi [13]. To adapt their algorithm, one must generate and solve extra inequalities, as required by rules T-SUB and T-COLOR. These are atomic, i.e. only involve variables and constants taken in  $\mathcal{L}$ .

### 7.4 Related Work

In previous work with Sylvain Conchon [20], we proposed a generic approach to information flow analysis, based on a suitable *colored semantics*, where terms are annotated with *security colors* taken from  $\mathcal{L}$ . (See e.g. Sewell and Vitek’s colored box- $\pi$ -calculus [22].) We suggested, at the time, that this approach should be applicable to the  $\pi$ -calculus. Later experiments confirmed our intuition, but showed that it naturally leads to a *may-testing*-based noninterference result. Our attempts to adapt it to a bisimulation setting lead us to the present formulation, where brackets replace colors, allowing a simple bisimulation proof.

Hennessy and Riely’s system [9] shares several basic mechanisms, such as the use of security annotations on processes, channels and judgements, with ours. Instead of using channel types annotated with polarities à la Pierce and Sangiorgi, they use sets of so-called read or write “capabilities”; this seems only a superficial difference. More importantly, they study the *asynchronous*  $\pi$ -calculus under *may-testing* equivalence, which means that information

only flows from senders to receivers and allows a channel to be read at a higher security level than it was written. (Contrast this with our discussion of Sect. 7.1.) Lastly, their boxing construct  $\rho[P]$  is *not* analogous to our coloring construct  $l : P$ . (Compare the corresponding typing rules.) It is used with different meanings in their “resource control” system and in their “information flow” system, which creates a tension and seems to make their noninterference statement a bit awkward.

Sewell and Vitek [22] develop a type system similar to ours. (Their annotations are sets of principals, whereas we employ a slightly more abstract notion of security level.) They do not prove a noninterference result; instead, they state a so-called “causal flow” property. We view this as a serious shortcoming: it is difficult to determine exactly which notion of causality the property reflects. The same criticism can be held against Bodei *et al.*’s “no read up/no write down” property [4].

The restriction of Honda *et al.*’s system [11] to nonlinear types seems essentially identical to our system: judgements are annotated with a security level similar to ours (compare  $(\text{Deg}_s)$  with T-SUB), and nonlinear channel types are invariant in their security level. The noninterference results stated in [11, 12, 27] also rely on a form of weak bisimulation. They propose the most advanced systems to date; whether our technique can be modified to establish their soundness is an interesting issue.

The problem of noninterference in multi-threaded imperative languages is similar to the one studied here, and has been investigated by several researchers [23, 24, 5]. We believe that our proof technique could be re-used in their setting, allowing the candidate bisimulation relation to be defined implicitly in terms of a type system for a “bracket” calculus – as done here – rather than explicitly, which is clumsier (see e.g. [5]).

### Acknowledgements

The paper owes much to numerous discussions with Sylvain Conchon. Alan Schmitt’s diagram-chasing skills were also most helpful.

### References

- [1] M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke. A core calculus of dependency. In *Conference Record of the 26th ACM Symposium on Principles of Programming Languages*, pages 147–160, San Antonio, Texas, Jan. 1999. ACM Press. URL: <http://www.soe.ucsc.edu/~abadi/Papers/flowpopl.ps>.
- [2] M. Abadi, B. Lampson, and J.-J. Lévy. Analysis and caching of dependencies. In *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming*, pages 83–91, Philadelphia, Pennsylvania, May 1996.

- ACM Press. URL: <http://www.soe.ucsc.edu/~abadi/Papers/make-preprint.ps>.
- [3] D. E. Bell and L. J. LaPadula. Secure computer systems: Unified exposition and Multics interpretation. Technical Report MTR-2997, The MITRE Corp., Bedford, Massachusetts, July 1975. URL: <http://www.mitre.org/resources/centers/infosec/infosec.html>.
- [4] C. Bodei, P. Degano, F. Nielson, and H. R. Nielson. Static analysis of processes for no read-up and no write-down. In W. Thomas, editor, *Proceedings of FoSSaCS'99*, volume 1578 of *Lecture Notes in Computer Science*, pages 120–134. Springer, Mar. 1999. URL: <http://www.di.unipi.it/~chiara/publ-40/BDNN99.ps>.
- [5] G. Boudol and I. Castellani. Non-interference for concurrent programs and thread systems. To appear. URL: <ftp://ftp-sop.inria.fr/mimosa/personnel/gbo/non-interf-threads.ps.gz>, Sept. 2001.
- [6] D. E. Denning. *Cryptography and Data Security*. Addison-Wesley, Reading, Massachusetts, 1982.
- [7] R. Focardi and R. Gorrieri. A classification of security properties for process algebras. *Journal of Computer Security*, 3(1):5–33, 1995. URL: <http://www.cs.unibo.it/~gorrieri/Papers/jcsfinal.ps.gz>.
- [8] M. Hennessy. The security picalculus and non-interference. Technical Report 2000:05, University of Sussex, Nov. 2000. URL: <ftp://ftp.cogs.susx.ac.uk/pub/reports/compsci/cs052000.ps.z>.
- [9] M. Hennessy and J. Riely. Information flow vs. resource access in the asynchronous pi-calculus. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming*, Lecture Notes in Computer Science. Springer-Verlag, July 2000. URL: <http://www.depaul.edu/~jriely/papers/00icalp.ps.gz>.
- [10] M. Hepburn and D. Wright. Trust in the pi-calculus. In *Third International Conference on Principles and Practice of Declarative Programming (PPDP'01)*, Sept. 2001.
- [11] K. Honda, V. Vasconcelos, and N. Yoshida. Secure information flow as typed process behaviour. In G. Smolka, editor, *Proceedings of the 2000 European Symposium on Programming (ESOP'00)*, volume 1782 of *Lecture Notes in Computer Science*, pages 180–199. Springer Verlag, Mar. 2000. URL: <ftp://ftp.dcs.qmw.ac.uk/lfp/kohei/siftp-esop00.ps.gz>.
- [12] K. Honda and N. Yoshida. A uniform type structure for secure information flow. In *Proceedings of the 29th ACM Symposium on Principles of Programming Languages (POPL'02)*, pages 81–92, Portland, Oregon, Jan. 2002. URL: <http://www.mcs.le.ac.uk/~nyoshida/paper/ifal.ps.gz>.
- [13] A. Igarashi and N. Kobayashi. Type reconstruction for linear  $\pi$ -calculus with I/O subtyping. *Information & Computation*, 161:1–44, Aug. 2000. URL: <http://www.graco.c.u-tokyo.ac.jp/~igarashi/papers/LinearPi.ps.gz>.
- [14] N. Kobayashi, B. C. Pierce, and D. N. Turner. Linearity and the Pi-Calculus. *ACM Transactions on Programming Languages and Systems*, 21(5):914–947, Sept. 1999. URL: <http://www.acm.org/pubs/citations/journals/toplas/1999-21-5/p914-kobayashi/>.
- [15] J. McLean. A general theory of composition for trace sets closed under selective interleaving functions. In *Proceedings of the 1994 IEEE Symposium on Research in Security and Privacy*. IEEE Press, 1994. URL: <http://chacs.nrl.navy.mil/publications/CHACS/1994/1994mclean-sp.ps>.
- [16] R. Milner. The polyadic  $\pi$ -calculus: a tutorial. Technical Report ECS-LFCS-91-180, Laboratory for Foundations of Computer Science, Department of Computer Science, University of Edinburgh, Oct. 1991. URL: <ftp://ftp.cl.cam.ac.uk/users/rml35/ppi.ps.z>.
- [17] R. Milner and D. Sangiorgi. Barbed bisimulation. In W. Kuich, editor, *International Colloquium on Automata, Languages and Programming*, volume 623 of *Lecture Notes in Computer Science*, pages 685–695, Vienna, Austria, July 1992. Springer-Verlag. URL: <ftp://ftp-sop.inria.fr/meije/theorie-par/davides/bn.ps.gz>.
- [18] B. Pierce and D. Sangiorgi. Typing and subtyping for mobile processes. In *Proceedings of the Eighth Annual IEEE Symposium on Logic in Computer Science*, pages 376–385, June 1993. URL: <http://www.cis.upenn.edu/~bcpierce/papers/pi-lics.ps>.
- [19] F. Pottier. A simple view of type-secure information flow in the  $\pi$ -calculus. Full version. URL: <http://pauillac.inria.fr/~fpottier/publis/fpottier-csfw15-long.ps.gz>, Feb. 2002.
- [20] F. Pottier and S. Conchon. Information flow inference for free. In *Proceedings of the 5th ACM SIGPLAN International Conference on Functional Programming (ICFP'00)*, pages 46–57, Montréal, Canada, Sept. 2000. ACM Press. URL: <http://pauillac.inria.fr/~fpottier/publis/fpottier-conchon-icfp00.ps.gz>.
- [21] F. Pottier and V. Simonet. Information flow inference for ML. In *Proceedings of the 29th ACM Symposium on Principles of Programming Languages (POPL'02)*, pages 319–330, Portland, Oregon, Jan. 2002. ACM Press. URL: <http://pauillac.inria.fr/~fpottier/publis/fpottier-simonet-popl02.ps.gz>.
- [22] P. Sewell and J. Vitek. Secure composition of untrusted code: Wrappers and causality types. In *Proceedings of the 13th Computer Security Foundations Workshop*. IEEE Computer Society Press, July 2000. URL: <http://www.cl.cam.ac.uk/users/pes20/wraptypes.ps.gz>.
- [23] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In *Conference Record of the 25th ACM Symposium on Principles of Programming Languages*, pages 355–364, Jan. 1998. URL: <http://www.cs.nps.navy.mil/people/faculty/volpano/papers/popl98.ps.z>.
- [24] G. S. Smith. A new type system for secure information flow. In *Proc. 14th IEEE Computer Security Foundations Workshop*, pages 115–125, Cape Breton, Nova Scotia, June 2001. URL: <http://www.cs.fiu.edu/~smithg/papers/csfw01.ps.gz>.
- [25] N. Yoshida. Graph types for monadic mobile processes. In V. Chandru and V. Vinay, editors, *Proceedings of the 16th Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 1180 of *Lecture Notes in Computer Science*, pages 371–386.

Springer-Verlag, 1996. URL: [http://www.mcs.le.ac.uk/~nyoshida/paper/graph1\\_short.ps.gz](http://www.mcs.le.ac.uk/~nyoshida/paper/graph1_short.ps.gz).

- [26] N. Yoshida, K. Honda, and M. Berger. Linearity and bisimulation. Technical Report MSC-2001/48, University of Leicester, Dec. 2001. URL: <http://www.mcs.le.ac.uk/~nyoshida/paper/lb.ps.gz>.
- [27] N. Yoshida, K. Honda, and M. Berger. Linearity and bisimulation. In *Proceedings of 5th International Conference of Foundations of Software Science and Computer Structures (FoSSaCs 2002)*, Lecture Notes in Computer Science. Springer Verlag, Apr. 2002. URL: [http://www.mcs.le.ac.uk/~nyoshida/paper/fossacs\\_ca\\_final.ps.gz](http://www.mcs.le.ac.uk/~nyoshida/paper/fossacs_ca_final.ps.gz).
- [28] S. Zdancewic and A. C. Myers. Secure information flow and CPS. In D. Sands, editor, *Proceedings of the 2001 European Symposium on Programming (ESOP'01)*, Lecture Notes in Computer Science, Genova, Italy, Apr. 2001. Springer Verlag. URL: <http://www.cs.cornell.edu/zdance/lincont.ps>.