Noninterference for Concurrent Programs *

Gérard Boudol and Ilaria Castellani
INRIA
2004 Route des Lucioles, BP 93
06902 Sophia-Antipolis – France

Abstract

We propose a type system to ensure the property of *noninterference* in a system of concurrent programs, described in a standard imperative language extended with parallelism. Our proposal is in the line of some recent work by Irvine, Volpano and Smith. Our type system, as well as our semantics for concurrent programs, seem more natural and less restrictive than those originally presented by these authors. Moreover, we show how to extend the language in order to formalise scheduling policies. The type system is extended to the new constructs, and we show that noninterference still holds, while remaining in a nonprobabilistic setting.

1 Introduction

The aim of this paper is to study the notion of secure information flow, and more specifically of *noninterference* (a notion first introduced by Goguen and Meseguer in [3]) in the setting of concurrency. Our starting point is the paper [14] by Volpano, Smith and Irvine, and the subsequent paper [11] by Smith and Volpano, where noninterference is enforced by means of a simple type system in an imperative language with security levels. The language considered in [14] is purely sequential, and is extended in [11] with asynchronous parallelism (interleaving). In this introduction, and in the examples given in the paper, the security levels will simply be high (H) and low (L). High-level variables are supposed to contain *secret* information, while low-level variables contain *public* information. However all results will be given for an arbitrary lattice of security levels.

In Volpano et al.'s work, *noninterference* means that variables of a given level do not interfere with those of lower levels: more precisely, the values of low-level variables are not dependent on the values of high-level variables.

*Research partially funded by the EU Working Group CONFER II and by the french RNRT Project MARVEL.

Noninterference is meant to model the absence of information flow from high level to low level, and more generally from any security level σ_0 to another one σ_1, except when this is explicitly allowed, that is $\sigma_0 \leq \sigma_1$. Such information flow is considered *insecure*, as it amounts to the disclosure of secret information into the public domain. Insecure flow can be *explicit*, when assigning the value of a high variable to a low variable, or *implicit*, when testing the value of a high variable and then assigning a value depending on the result of the test to a low variable, for instance. In the approach of [14, 11], these situations are prevented by means of a type system. More precisely, explicit flow is prevented by requiring that the level of the assigned variable be at least as high as that of the source variable, while implicit flow is prevented by asking that the level of the commands in the branches of a conditional (the level of a command being that of its lowest assigned variables) be at least as high as that of the tested variable. Implicit flow can also arise in while-loops, and is prevented by a similar condition on the type of the body of the loop.

In fact, because of while-loops, the definition of noninterference is more precise than what is stated above: it says that no change in the values of low-level variables should be observed as a consequence of a change in high-level variables, *provided that the program terminates successfully*. Using subscripts to explicitly indicate the security level of a variable, let us consider the following program, that terminates if $x_H \neq 0$ and loops forever (doing nothing) otherwise:

$$\texttt{while } x_H = 0 \texttt{ do nil } ;\ y_L := 1 \tag{1}$$

Should this program be accepted, that is, should it be typable? According to the above definition of noninterference the answer is "yes", since whenever the program terminates it produces the same value $y_L = 1$ for its low-level variable. Indeed, this program is typable in Volpano and Smith's type system, since the loop is typable and the sequential composition of typable programs is always typable.

However, accepting such a program leads to problems when parallelism is introduced in the language. These problems can be concisely described as "disguising information flow as control flow". Let us illustrate the problem by means of an example, which is a simplified version of the *PIN* example given by Smith and Volpano in [11]. In this example, given in Figure 1, three threads α, β and γ are run (asynchronously) in parallel. There are four variables, a high-level variable *PIN* tested by thread γ, two high-level variables t_α and t_β serving as "triggers" for threads α and β, and a low-level variable r written by α and β. As can be easily seen, with initial values $t_\alpha = t_\beta = \mathit{ff}$ the effect of this program is to copy the value of the secret variable *PIN* into the public variable r. The illicit information flow from *PIN* to r is implemented through the *control flow* from γ to α or β. However,

$$\gamma:\ \texttt{if}\ PIN = 0\ \texttt{then}\ t_\beta := tt\ \texttt{else}\ t_\alpha := tt$$
$$\alpha:\ \texttt{while}\ t_\alpha \neq tt\ \texttt{do nil}\ ;\ r := 0\ ;\ t_\beta := tt$$
$$\beta:\ \texttt{while}\ t_\beta \neq tt\ \texttt{do nil}\ ;\ r := 1\ ;\ t_\alpha := tt$$

PIN, t_α, t_β : boolean variables of type H

r : boolean variable of type L

γ : thread of type H, α, β : threads of type L

Figure 1: Information Flow through Control Flow

---

if we assume that a system of concurrent threads is typable provided each component is typable, as would be natural, this particular system is to be accepted.

To circumvent this problem, Smith and Volpano propose in [11] to forbid the use of high-level variables as guards in while-loops, and more specifically[1] to accept only while-loops of low level. While ruling out the program in (1), and also the threads α and β of the *PIN* example, this solution seems a bit drastic. It excludes inoffensive programs such as $\texttt{while}\ x_H = 0\ \texttt{do nil}$. We shall propose here a different solution to the problem raised by while-loops in the presence of parallelism, which allows this program to be typed, while ruling out the programs of example (1) and Figure 1. Our solution is based on the observation that a program such as

$$\texttt{while}\ x_H = 0\ \texttt{do nil}$$

should indeed be considered with some care in a concurrent setting, but only as a "guard", that is, as regards what may follow it. In the context of concurrent threads, if the control comes back to this `while` loop, this may be with a value for x_H different from 0, contrarily to what happens in a sequential setting. In other words, this program may observe the behaviour of other, concurrent components, in the course of their execution, and influence accordingly the behaviour of the thread in which it participates. Technically, this means that we will abandon the *big-step semantics* which is the basis of Volpano et al.'s analysis in favour of a *small step semantics* for programs, which is the approach usually adopted in dealing with parallelism. Our aim is then to ensure a stronger form of noninterference, where the course of values – not just the final value – of a low-level variable does

[1] assuming, in the case of an arbitrary lattice of security levels, that there is a lowest level $\bot$.

not depend upon the value of high variables. Typically, the program (1) is no longer interference-free in this stronger sense. In order to reject it, we introduce a refinement of the type system, where the level of loop guards – the expressions tested by `while` loops – is taken into account in sequential composition. We shall see that also conditional guards – the expressions tested by conditional statements – need to be taken into account.

We will also examine the situation where a scheduling policy is in force in a thread system: we will introduce a couple of new programming primitives to describe formally such a situation, and show how to extend the type system for this new setting, where new interference phenomena arise. We show that our type system still ensures the noninterference property in this case, while imposing a less severe restriction than that prefigured in [11].

The rest of the paper is organised as follows. In Section 2 we introduce the language, its operational semantics and its type system. Section 3 presents the properties of typed programs, including subject reduction and noninterference. Finally, in Section 4 we consider the extended language with scheduling policies.

2 The Language and Type System

The language we consider is essentially that of [11] (where e stands for a boolean or arithmetic expression, whose syntax we do not detail here). We use the following two-level syntax, where U, V denote sequential programs, while P, Q denote general (concurrent) programs:

$$\begin{array}{rcl} U,\ V\ldots & := & \texttt{nil} \mid x := e \mid U;V \mid \texttt{if}\ e\ \texttt{then}\ U\ \texttt{else}\ V \\ & \mid & \texttt{while}\ e\ \texttt{do}\ U \\ P,\ Q\ldots & := & U \mid U;P \mid \texttt{if}\ e\ \texttt{then}\ P\ \texttt{else}\ Q \mid \texttt{while}\ e\ \texttt{do}\ P \\ & \mid & P \parallel Q \end{array}$$

Note that on the left of a sequential composition, we must have a sequential program. Thus programs of the form $(P \parallel Q); R$ are not allowed. With this restriction, our language is still more general than that of [11], which describes concurrent systems as collections of threads, thus allowing only top-level parallelism, while we allow the dynamic spawning of new threads.

The operational semantics of the language is given in terms of transitions between configurations $(P, \mu) \rightarrow (P', \mu')$ where P, P' are programs and μ, μ' stand for *memories*, that is mappings from variables to values. These mappings are extended in the obvious way to expressions, whose evaluation is assumed to be atomic as in [11]. We use the notation $\mu[v/x]$ for memory update. The rules specifying the operational semantics of programs are presented in Figure 2. As pointed out already in the introduction, the semantics used here is a *small step semantics*, as opposed to the *big step*

$$\text{(Assign-Op)}\quad \frac{}{(x := e, \mu) \to (\mathtt{nil}, \mu[\mu(e)/x])}$$

$$\text{(Seq-Op1)}\quad \frac{(U, \mu) \to (U', \mu')}{(U; P, \mu) \to (U'; P, \mu')}$$

$$\text{(Seq-Op2)}\quad \frac{(P, \mu) \to (P', \mu')}{(\mathtt{nil}; P, \mu) \to (P', \mu')}$$

$$\text{(Cond-Op1)}\quad \frac{\mu(e) = tt}{(\mathtt{if}\ e\ \mathtt{then}\ P\ \mathtt{else}\ Q, \mu) \to (P, \mu)}$$

$$\text{(Cond-Op2)}\quad \frac{\mu(e) \neq tt}{(\mathtt{if}\ e\ \mathtt{then}\ P\ \mathtt{else}\ Q, \mu) \to (Q, \mu)}$$

$$\text{(While-Op1)}\quad \frac{\mu(e) = tt}{(\mathtt{while}\ e\ \mathtt{do}\ P, \mu) \to (P; \mathtt{while}\ e\ \mathtt{do}\ P, \mu)}$$

$$\text{(While-Op2)}\quad \frac{\mu(e) \neq tt}{(\mathtt{while}\ e\ \mathtt{do}\ P, \mu) \to (\mathtt{nil}, \mu)}$$

$$\text{(Par-Op1)}\quad \frac{(P, \mu) \to (P', \mu')}{(P \parallel Q, \mu) \to (P' \parallel Q, \mu')}$$

$$\text{(Par-Op2)}\quad \frac{(Q, \mu) \to (Q', \mu')}{(P \parallel Q, \mu) \to (P \parallel Q', \mu')}$$

Figure 2: Operational Semantics for Parallel Programs

semantics of [11][2]. The rules are fairly standard, and we shall not comment on them.

In the introduction we argued that, in a small-steps semantics, the program (1) should be treated as another case of implicit information flow. Intuitively, when exiting a loop one gets some information about its guard; it seems then appropriate to require that what follows the loop – its "continuation" – have level at least as high as that of the loop guard. This will be the basic idea of our new type system, which is closely inspired by that given by Volpano et al. in [14] – however as suggested by the above example it will be more restrictive than that of [14] on the sequential sublanguage,

[2]In fact, the semantics of [11] is a mixture of small and big step semantics: transitions are given between configurations but there are two kinds of configurations, intermediate and final ones, suggesting that termination should be observed.

$$\text{(Nil)}\quad \frac{}{\Gamma \vdash \texttt{nil} : (\tau, \sigma)\, cmd}$$

$$\text{(Assign)}\quad \frac{\Gamma \vdash e : \tau, \quad \Gamma(x) = \tau\, var}{\Gamma \vdash x := e : (\tau, \sigma)\, cmd}$$

$$\text{(Seq)}\quad \frac{\Gamma \vdash U : (\tau, \sigma)\, cmd, \quad \Gamma \vdash P : (\tau', \sigma')\, cmd, \quad \sigma \leq \tau'}{\Gamma \vdash U\,;\, P : (\tau \sqcap \tau', \sigma \sqcup \sigma')\, cmd}$$

$$\text{(Cond)}\quad \frac{\Gamma \vdash e : \theta, \quad \Gamma \vdash P_i : (\tau, \sigma)\, cmd, \quad \theta \leq \tau}{\Gamma \vdash \texttt{if}\ e\ \texttt{then}\ P_0\ \texttt{else}\ P_1 : (\tau, \theta \sqcup \sigma)\, cmd}$$

$$\text{(While)}\quad \frac{\Gamma \vdash e : \theta, \quad \Gamma \vdash P : (\tau, \sigma)\, cmd, \quad \theta \sqcup \sigma \leq \tau}{\Gamma \vdash \texttt{while}\ e\ \texttt{do}\ P : (\tau, \theta \sqcup \sigma)\, cmd}$$

$$\text{(Par)}\quad \frac{\Gamma \vdash P_i : (\tau, \sigma)\, cmd}{\Gamma \vdash P_0 \parallel P_1 : (\tau, \sigma)\, cmd}$$

$$\text{(Subtyping)}\quad \frac{\Gamma \vdash P : (\tau, \sigma)\, cmd, \quad \tau' \leq \tau, \quad \sigma \leq \sigma'}{\Gamma \vdash P : (\tau', \sigma')\, cmd}$$

Figure 3: Typing Rules for Concurrent Programs

---

because of our more detailed observation of programs.

The types of data and expressions are *security levels*, that is elements of a lattice $(\mathcal{T}, \leq)$. We denote the operations of meet and join respectively by $\sqcap$ and $\sqcup$. These types are ranged over by $\tau,\ \sigma \ldots$ In the examples, the lattice of security levels will simply be $\{L, H\}$, with $L < H$. The types of variables (when used in the left-hand side of an assignment) are of the form $\tau\, var$. Our first point of departure from [14] concerns the types for programs. Type judgements in [14] are of the form $\Gamma \vdash P : \tau\, cmd$, where Γ is a mapping from variables to types of variables, i.e. elements of $\{\tau\, var \mid \tau \in \mathcal{T}\}$. The meaning of $\Gamma \vdash P : \tau\ cmd$ is that in the type environment Γ, the type τ is a *lower bound* for the level of the *assigned variables* of P. In line with this intuition, subtyping for programs is contravariant, that is $\tau\, cmd \leq \tau'\, cmd$ if $\tau' \leq \tau$. Thus for instance any program of type $H\ cmd$ can be downgraded to type $L cmd$. A program of type $H cmd$ is guaranteed not to contain any assignment to a low variable.

As we said previously, we have to take into account the level of loop guards in typing. To deal with possible divergence (and, as we shall see, with the scheduling of threads) we also have to record the level of the tested

expression in conditional branching, which we regard as a guard as well. Thus we shall use here more refined types $(\tau, \sigma)\ cmd$, where the first component τ plays the same rôle as in the type $\tau\ cmd$, while the second component σ is the *guard type*, an *upper bound* on the level of the loop and conditional guards occurring in a program. Accordingly, the subtyping for programs is contravariant in its first component and covariant in the second:

$$(\tau, \sigma)\ cmd \leq (\tau', \sigma')\ cmd \ \text{ if } \ \tau' \leq \tau \ \text{ and } \ \sigma \leq \sigma'$$

The guard type will be set up by while-loops and conditional branchings, and looked up by sequential composition. The complete type system for programs is shown in Figure 3, where we assume given a system allowing us to infer judgements $\Gamma \vdash e : \tau$ about the security level of an expression. These are used in the (Assign), (Cond) and (While) rules. Notice that the guard type plays no particular rôle in rules (Nil) and (Assign), which are plain adaptations of the ones in [14]. As in this paper, the side condition $\theta \leq \tau$ in the rules (Cond) and (while) is meant to avoid implicit insecure flow. Let us comment a little on the rules for conditional branching, while-loops and sequential composition, which are the main novelty w.r.t. [14, 11]. As explained, the guard type is at least θ for a conditional branching or a while-loop testing an expression of level θ, and from then onwards it should remain at least θ to prevent concatenation with programs of a lower or incomparable level. Rule (Seq) is precisely designed to avoid sequencing "low" assignments after a program with "high" guards. This rules out the kind of implicit flow exhibited by the program (1). It also rules out the (semantically equivalent) program

$$\begin{array}{l} \texttt{if}\ x_H = 0\ \ \texttt{then}\ \ \texttt{while}\ tt\ \texttt{do}\ \texttt{nil} \\ \qquad\qquad\ \ \texttt{else}\ \ \texttt{nil}\,; \\ y_L := 1 \end{array} \tag{2}$$

which is not interference-free in our sense. Indeed, this example shows why the level of conditional guards should be recorded as a guard type. The side condition $\sigma \leq \tau$ in the rule (While) is needed for the preservation of types during execution, as we shall see in the next section.

3 Properties of Typed Programs

In this section we prove some desired properties of our type system. The first property, *subject reduction*, states that types are preserved along execution.

Theorem 3.1 (Subject Reduction)
If $\Gamma \vdash P : (\tau, \sigma)\ cmd$ *and* $(P, \mu) \to (P', \mu')$, *then* $\Gamma \vdash P' : (\tau, \sigma)\ cmd$.

Proof: By induction on the inference of $\Gamma \vdash P : (\tau, \sigma)\ cmd$, and then case analysis on the last rule used in this inference. If this rule is the subtyping rule, we simply use the induction hypothesis. We examine the main cases:

(ASSIGN). Here we have $P = x := e$ with $\Gamma \vdash e : \tau$ and $\Gamma(x) = \tau\, var$. Moreover, it must be that $P' = \texttt{nil}$ and $\mu' = \mu[\mu(e)/x]$. The theorem is trivial in this case, since $\Gamma \vdash \texttt{nil} : (\tau, \sigma)\, cmd$.

(SEQ). Here $P = U; Q$ with $\Gamma \vdash U : (\tau_1, \sigma_1)\, cmd$, $\Gamma \vdash Q : (\tau_2, \sigma_2)\, cmd$, $\tau = \tau_1 \sqcap \tau_2$ and $\sigma = \sigma_1 \sqcup \sigma_2$ with $\sigma_1 \leq \tau_2$. There are two possibilities for a transition $(P, \mu) \to (P', \mu')$. If this results from $(U, \mu) \to (U', \mu')$, by means of (SEQ-OP1), we have $P' = U'; Q$, where $\Gamma \vdash U' : (\tau_1, \sigma_1)\, cmd$ by induction, and we conclude using the rule (SEQ). If (SEQ-OP2) is used, we have $U = \texttt{nil}$ with $(Q, \mu) \to (P', \mu')$. By induction $\Gamma \vdash P' : (\tau_2, \sigma_2)\, cmd$, and we use the subtyping rule to conclude.

(WHILE). Here $P = \texttt{while}\ e\ \texttt{do}\ Q$ and $\Gamma \vdash Q : (\tau, \sigma')\, cmd$, with $\sigma = \theta \sqcup \sigma'$ and $\sigma \leq \tau$. If $\mu(e) = tt$ we have $P' = Q; P$, and since $\sigma' \leq \tau$ we may use the rule (SEQ) to conclude $\Gamma \vdash P' : (\tau, \sigma)\, cmd$. Otherwise $P' = \texttt{nil}$, and this case is trivial.

All the other cases are similar (in the case of (COND) we use subtyping). □

We shall use the following assumptions about expressions:

Assumption 3.2 (Termination of Expression Evaluation)
For any memory μ and expression e, the value $\mu(e)$ is defined.

Assumption 3.3 (Simple Security)
If $\Gamma \vdash e : \tau$, then every variable occurring in e has type $\tau'\, var$ in Γ, with $\tau' \leq \tau$.

We introduce now a notion of equality on memories, depending on a given type environment Γ and a given set $\mathcal{L} \subseteq \mathcal{T}$ of security levels, which are to be understood as "low" security levels. Two memories are $(\Gamma, \mathcal{L})$-equal if they coincide on variables to which Γ assigns a level in $\mathcal{L}$. Intuitively, such memories are indistinguishable for an observer whose level is in $\mathcal{L}$.

Definition 3.4 ($(\Gamma, \mathcal{L})$-Equality of Memories)
$\mu =_\Gamma^{\mathcal{L}} \nu \;\Leftrightarrow_{\text{def}}\; \forall x.\, \Gamma(x) = \tau\, var \;\&\; \tau \in \mathcal{L} \;\Rightarrow\; \mu(x) = \nu(x)$.

In the examples we implicitly used the particular case $=_\Gamma^{\{L\}}$, which is denoted $\sim_\Gamma$ in [11]. To define our notion of bisimulation, it is convenient to introduce a notation: we denote by $\twoheadrightarrow$ the reflexive closure of $\to$, that is:

$$(P, \mu) \twoheadrightarrow (P', \mu') \;\Leftrightarrow_{\text{def}}\; P' = P \;\&\; \mu' = \mu \quad \text{or} \quad (P, \mu) \to (P', \mu')$$

Definition 3.5 ($(\Gamma, \mathcal{L})$-Bisimulation)
A relation $\mathcal{R}$ on configurations is a $(\Gamma, \mathcal{L})$-bisimulation if $(P, \mu)\, \mathcal{R}\, (Q, \nu)$ implies

(i) $\mu =_\Gamma^{\mathcal{L}} \nu$

(ii) $(P,\mu) \rightarrow (P',\mu') \;\Rightarrow\; \exists Q',\nu'.\ (Q,\nu) \twoheadrightarrow (Q',\nu') \;\wedge\; (P',\mu')\,\mathcal{R}\,(Q',\nu')$

(iii) $(Q,\nu) \rightarrow (Q',\nu') \;\Rightarrow\; \exists P',\mu'.\ (P,\mu) \twoheadrightarrow (P',\mu') \;\wedge\; (P',\mu')\,\mathcal{R}\,(Q',\nu')$

The (Γ,ω)-*bisimulation equivalence* on configurations, noted $\simeq^{\mathcal{L}}_{\Gamma}$, is the largest (Γ,ω)-bisimulation.

To establish the noninterference result, we need some preliminary lemmas. The following two lemmas confirm the intuition, discussed earlier, behind the type judgements $\Gamma \vdash P : (\tau,\sigma)\ cmd$. Both results are easily proved by induction on the inference of type judgements.

Lemma 3.6 (Confinement)
If $\Gamma \vdash P : (\tau,\sigma)\ cmd$ *then every variable assigned to in* P *has type* $\theta\ var$ *in* Γ*, with* $\tau \leq \theta$*.*

Lemma 3.7 (Guard Safety)
If $\Gamma \vdash P : (\tau,\sigma)\ cmd$ *then every loop or conditional branching guard in* P *has type* θ *in* Γ*, with* $\theta \leq \sigma$*.*

Definition 3.8 ($\mathcal{L}$-Boundedness)
A program P is $\mathcal{L}$-*bounded* in Γ if $\Gamma \vdash P : (\tau,\sigma)\ cmd$ implies $\tau \in \mathcal{L}$.

We shall often use the complementary notion of *non-$\mathcal{L}$-boundedness*: P is *not $\mathcal{L}$-bounded* in Γ if there exist τ,σ, with $\tau \notin \mathcal{L}$, such that $\Gamma \vdash P : (\tau,\sigma)\ cmd$.

Definition 3.9 ($\mathcal{L}$-Guardedness)
A program P is $\mathcal{L}$-*guarded* in Γ if there exist $\sigma \in \mathcal{L}$ and τ such that $\Gamma \vdash P : (\tau,\sigma)\ cmd$.

In the following, we shall most often use the terminology "(non-)$\mathcal{L}$-bounded" and "$\mathcal{L}$-guarded" without any explicit reference to the current typing context Γ. As a consequence of subject reduction, both *non-$\mathcal{L}$*-boundedness and $\mathcal{L}$-guardedness are preserved by execution. Let us now introduce a definition: we say that the set $\mathcal{L}$ of low security levels is *downward-closed* if it satisfies

$$\tau \in \mathcal{L}\ \&\ \sigma \leq \tau \;\Rightarrow\; \sigma \in \mathcal{L}$$

Note that, if $\mathcal{L}$ is downward-closed, then by the Confinement Lemma a program which is *not $\mathcal{L}$-bounded* cannot write on variables of low level, that is in $\mathcal{L}$. Similarly, by the Guard Safety Lemma, a program which is $\mathcal{L}$-*guarded* does not contain loop guards of high level, that is not in $\mathcal{L}$.

Proposition 3.10 (Bisimilarity of Non-$\mathcal{L}$-Bounded Programs)
Let $\mathcal{L}$ *be a downward-closed set of security levels, and let* $\mathcal{S}^{\Gamma,\mathcal{L}}$ *be the relation consisting of the pairs* $((P,\mu),(Q,\nu))$ *such that* $\mu =^{\mathcal{L}}_{\Gamma} \nu$ *and both* P *and* Q *are non-*$\mathcal{L}$*-bounded in* Γ*. Then* $\mathcal{S}^{\Gamma,\mathcal{L}}$ *is a* $(\Gamma,\mathcal{L})$*-bisimulation.*

Proof: Let $((P,\mu),(Q,\nu)) \in \mathcal{S}^{\Gamma,\mathcal{L}}$ and $(P,\mu) \to (P',\mu')$. This can be matched by $(Q,\nu) \twoheadrightarrow (Q,\nu)$, since by the Confinement Lemma $\mu' =^{\mathcal{L}}_{\Gamma} \mu =^{\mathcal{L}}_{\Gamma} \nu$, and by the Subject Reduction Theorem $\Gamma \vdash P : (\tau,\sigma)$ implies $\Gamma \vdash P' : (\tau,\sigma)$. □

Note that $(\Gamma,\mathcal{L})$-bisimilarity does not preserve termination. For instance, for any memories μ and ν such that $\mu =^{\mathcal{L}}_{\Gamma} \nu$ we have:

$$(\texttt{nil},\mu) \simeq^{\mathcal{L}}_{\Gamma} (\texttt{while } tt \texttt{ do nil},\nu)$$

Lemma 3.11 (Behaviour of $\mathcal{L}$-Guarded Programs)
Let $\mathcal{L}$ be a downward-closed set of security levels. If P is $\mathcal{L}$-guarded in Γ and $\mu =^{\mathcal{L}}_{\Gamma} \nu$, then $(P,\mu) \to (P',\mu')$ implies $(P,\nu) \to (P',\nu')$, with $\mu' =^{\mathcal{L}}_{\Gamma} \nu'$.

Proof: By induction on the inference of $\Gamma \vdash P : (\tau,\sigma)\, cmd$, and by case analysis on the last rule used in this inference. We examine the main cases.

(Assign). Here $P = x := e$ with $\Gamma \vdash e : \tau$, $\Gamma(x) = \tau\, var$, $P' = \texttt{nil}$ and $\mu' = \mu[\mu(e)/x]$. By Assumption 3.2, $\nu(e)$ is defined. Let $\nu' = \nu[\nu(e)/x]$. It is easy to see that $\mu' =^{\mathcal{L}}_{\Gamma} \nu'$, since if $\tau \notin \mathcal{L}$ then $\mu' =^{\mathcal{L}}_{\Gamma} \nu'$ by the Confinement Lemma, while if $\tau \in \mathcal{L}$ then $\mu(e) = \nu(e)$ by Assumption 3.3.

(Seq). Here $P = U;Q$, $\Gamma \vdash U : (\tau_1,\sigma_1)\, cmd$ and $\Gamma \vdash Q : (\tau_2,\sigma_2)\, cmd$ with $\sigma = \sigma_1 \sqcup \sigma_2$ and $\sigma_1 \leq \tau_2$. If the transition $(P,\mu) \to (P',\mu')$ is proved by means of (Seq-Op1), that is $P' = U';Q$ with $(U,\mu) \to (U',\mu')$, then we use the induction hypothesis, since $\sigma_1 \sqcup \sigma_2 \in \mathcal{L} \;\Rightarrow\; \sigma_1 \in \mathcal{L}$. This will give us $(U,\nu) \to (U',\nu')$, with $\mu' =^{\mathcal{L}}_{\Gamma} \nu'$, and by (Seq-Op1) we conclude that $(P,\nu) \to (U';Q,\nu')$, as required. In the case where (Seq-Op2) is used, with $U = \texttt{nil}$ and $(Q,\mu) \to (P',\mu')$, we simply use the induction hypothesis.

(Cond). Here $P = \texttt{if } e \texttt{ then } P_1 \texttt{ else } P_2$ with $\Gamma \vdash e : \theta$, $\Gamma \vdash P_i : (\tau,\sigma')\, cmd$, with $\theta \leq \tau$ and $\sigma = \theta \sqcup \sigma'$. Since $\sigma \in \mathcal{L}$, also $\theta \in \mathcal{L}$, and thus $\mu(e) = \nu(e)$ by Assumption 3.3. Hence a transition $(P,\mu) \to (P_i,\mu)$ is matched by $(P,\nu) \to (P_i,\nu)$.

(While). Here $P = \texttt{while } e \texttt{ do } Q$ and $\Gamma \vdash e : \theta$, $\Gamma \vdash Q : (\tau,\sigma')\, cmd$, with $\sigma \leq \tau$ and $\sigma = \theta \sqcup \sigma'$. Clearly $\sigma \in \mathcal{L}$ implies $\theta \in \mathcal{L}$, therefore $\mu(e) = \nu(e)$ by Assumption 3.3, and thus a transition of the loop, which is necessarily of the form $(P,\mu) \to (P',\mu)$, will be matched by $(P,\nu) \to (P',\nu)$.

The case of (Par) is easy, since the components have the same type as the whole program, and the case of (Subtyping) is immediate by induction. □

Theorem 3.12 (Noninterference)
If P is typable in Γ then $(P,\mu) \simeq^{\mathcal{L}}_{\Gamma} (P,\nu)$ for any downward-closed set $\mathcal{L}$ of security levels, and for any μ,ν such that $\mu =^{\mathcal{L}}_{\Gamma} \nu$.

Proof: We define inductively the relation $\mathcal{R}_0^{\Gamma,\mathcal{L}}$ on configurations as follows: $(P,\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(Q,\nu)$ if and only if P and Q are typable in Γ, $\mu =_\Gamma^{\mathcal{L}} \nu$ and one of the following holds:

1. $(P,\mu)\,\mathcal{S}^{\Gamma,\mathcal{L}}\,(Q,\nu)$, where $\mathcal{S}^{\Gamma,\mathcal{L}}$ is the relation defined in Proposition 3.10
2. $P = Q$ and P is $\mathcal{L}$-bounded
3. $P = U;R$ and $Q = V;R$, where $(U,\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(V,\nu)$ and R is not $\mathcal{L}$-bounded
4. P is not $\mathcal{L}$-bounded and $Q = V;R$, where $(\mathtt{nil},\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(V,\nu)$ and R is not $\mathcal{L}$-bounded
5. $P = P_1 \parallel P_2$ and $Q = Q_1 \parallel Q_2$ with $(P_i,\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(Q_i,\nu)$.

We show that $\mathcal{R}_0^{\Gamma,\mathcal{L}}$ is a $(\Gamma,\mathcal{L})$-bisimulation. The theorem will be a consequence of this fact, since if P is typable, then $(P,\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(P,\nu)$, either by Clause 1 or by Clause 2. By construction, $(P,\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(Q,\nu)$ implies $\mu =_\Gamma^{\mathcal{L}} \nu$. We show now, by induction on the definition of $\mathcal{R}_0^{\Gamma,\mathcal{L}}$, that if $(P,\mu) \to (P',\mu')$ and $(P,\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(Q,\nu)$, then there exist Q' and ν' such that $(Q,\nu) \twoheadrightarrow (Q',\nu')$, with $(P',\mu')\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(Q',\nu')$.

Clause 1: here we use the Proposition 3.10.

Clause 2: we proceed by induction on the structure of P.

(Assignment). Here $P = x := e$, $P' = \mathtt{nil}$ and $\mu' = \mu[\mu(e)/x]$. Then the matching move is $(P,\nu) \to (\mathtt{nil},\nu')$, where $\nu' = \nu[\nu(e)/x]$. Indeed from $\mu =_\Gamma^{\mathcal{L}} \nu$ we deduce $\mu' =_\Gamma^{\mathcal{L}} \nu'$ since P is $\mathcal{L}$-bounded and thus by Lemma 3.6 $\mu(e) = \nu(e)$.

(Sequence). Here $P = U;Q$. Consider a transition $(P,\mu) \to (P',\mu')$. We distinguish two cases, according to which of the rules (Seq-Op1) or (Seq-Op2) is used to deduce this transition:

1. If $(P,\mu) \to (P',\mu')$ is proved by means of (Seq-Op1), that is $P' = U';Q$ with $(U,\mu) \to (U',\mu')$, then there are two possibilities:
   1.1. Q is $\mathcal{L}$-bounded: then $\Gamma \vdash Q : (\tau_2,\sigma_2)\ cmd$ implies $\tau_2 \in \mathcal{L}$, therefore since P is typable there exist τ_1,σ_1 such that $\Gamma \vdash U : (\tau_1,\sigma_1)\ cmd$, with $\sigma_1 \in \mathcal{L}$. Thus U is $\mathcal{L}$-guarded. Then by Lemma 3.11 we have $(U,\nu) \to (U',\nu')$, with $\mu' =_\Gamma^{\mathcal{L}} \nu'$. Therefore $(P,\nu) \to (U';Q,\nu')$ by (Seq-Op1), which is the required matching move by Clause 2 again since $U';Q$ is $\mathcal{L}$-bounded (because $\tau_2 \in \mathcal{L}$ implies $\tau_1 \sqcap \tau_2 \in \mathcal{L}$).

1.2. Q is not $\mathcal{L}$-bounded. There exist U'' and ν' such that $(U,\nu) \twoheadrightarrow (U'',\nu')$ with $\mu' =^{\mathcal{L}}_{\Gamma} \nu'$ and $(U',\mu')\mathcal{R}_0^{\Gamma,\mathcal{L}}(U'',\nu')$, either by Proposition 3.10 if U is not $\mathcal{L}$-bounded, or by the induction hypothesis on the structure of terms if U is $\mathcal{L}$-bounded. We conclude using Clause 3.

2. Let $(P,\mu) \to (P',\mu')$ be proved by means of (Seq-Op2), with $U = \texttt{nil}$ and $(Q,\mu) \to (P',\mu')$. Then, as in case 1.2, we deduce $(Q,\nu) \twoheadrightarrow (P'',\nu')$, with $\mu' =^{\mathcal{L}}_{\Gamma} \nu'$ and $(P',\mu')\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(P'',\nu')$, using either Proposition 3.10 or the induction hypothesis.

(Conditional). Here $P = \texttt{if}\ e\ \texttt{then}\ P_1\ \texttt{else}\ P_2$ with $\Gamma \vdash e : \theta$, $\Gamma \vdash P_i : (\tau,\sigma')\ cmd$, with $\theta \leq \tau$ and $\sigma = \theta \sqcup \sigma'$. Then $\tau \in \mathcal{L}$ implies $\theta \in \mathcal{L}$, and by Assumption 3.3 we have $\mu(e) = \nu(e)$. Thus the move $(P,\mu) \to (P_i,\mu)$ is matched by $(P,\nu) \to (P_i,\nu)$, where $(P_i,\mu)\mathcal{R}_0^{\Gamma,\mathcal{L}}(P_i,\nu)$ by Clause 1 or Clause 2.

(While-loop). Here $P = \texttt{while}\ e\ \texttt{do}\ Q$. Since P is $\mathcal{L}$-bounded, $\mu(e) = \nu(e)$ by Assumption 3.3, and we proceed as in the case of (Cond) to obtain related derivatives.

(Parallel composition). This case is trivial, using the induction hypothesis, and brings us to Clause 5.

Clause 3: Here $P = U;R$ and $Q = V;R$, where $(U,\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(V,\nu)$ and R is not $\mathcal{L}$-bounded. Then if $(P,\mu) \to (P',\mu')$ is proved by means of (Seq-Op1), that is $(U,\mu) \to (U',\mu')$ with $P' = U';R$, we use induction on the definition of $\mathcal{R}_0^{\Gamma,\mathcal{L}}$ to deduce $(V,\nu) \twoheadrightarrow (V',\nu')$, with $\mu' =^{\mathcal{L}}_{\Gamma} \nu'$ and $(U',\mu')\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(V',\nu')$, and we stay in Clause 3. If $(P,\mu) \to (P',\mu')$ is proved by means of (Seq-Op2), that is $U = \texttt{nil}$ and $(R,\mu) \to (P',\mu')$, then we know that $\mu' =^{\mathcal{L}}_{\Gamma} \mu$ by Lemma 3.6. Then we take as matching move $(Q,\nu) \twoheadrightarrow (Q,\nu)$, and we fall in Clause 4 since by Subject Reduction P' is not $\mathcal{L}$-bounded.

Clause 4: Here P is not $\mathcal{L}$-bounded, $Q = V;R$, where $(\texttt{nil},\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(V,\nu)$ and R is not $\mathcal{L}$-bounded. Then a transition $(P,\mu) \to (P',\mu')$ is matched simply by $(Q,\nu) \twoheadrightarrow (Q,\nu)$, since $\mu' =^{\mathcal{L}}_{\Gamma} \mu$ by Lemma 3.6 and thus also $\mu' =^{\mathcal{L}}_{\Gamma} \nu$. Suppose now that (Q,ν) moves. If (Seq-Op1) is used, that is $(V,\nu) \to (V',\nu')$, by induction on the definition of $\mathcal{R}_0^{\Gamma,\mathcal{L}}$ we have $(\texttt{nil},\mu) \twoheadrightarrow (U,\mu')$ with $(U,\mu')\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(V',\nu')$. Since the only derivative of $(\texttt{nil},\mu)$ is itself, we have $U = \texttt{nil}$ and $\mu' = \mu$. Thus we can take $(P,\mu) \twoheadrightarrow (P,\mu)$ as matching move for $(Q,\nu) \to (V';R,\nu')$, and we stay in $\mathcal{R}_0^{\Gamma,\mathcal{L}}$ by Clause 4. If instead (Seq-Op2) is used, we have $V = \texttt{nil}$ and $(R,\nu) \to (R',\nu')$. Since R is not $\mathcal{L}$-bounded we know that $\nu =^{\mathcal{L}}_{\Gamma} \nu'$ and R' is not $\mathcal{L}$-bounded. Then $(P,\mu) \twoheadrightarrow (P,\mu)$ is again the matching move because $(P,\mu)\,\mathcal{R}_0^{\Gamma,\mathcal{L}}\,(R',\nu')$ by Clause 1.

Clause 5: trivial, using the induction hypothesis. From here we always remain in Clause 5. □

Observe that if τ_1 and τ_2 are incomparable levels in $(\mathcal{T}, \leq)$, they may be separated by a downward-closed set $\mathcal{L}$, namely $\mathcal{L} = \{\tau \mid \tau \leq \tau_1\}$, and therefore a typable program cannot yield any flow from level τ_2 to τ_1 (and conversely). In other words, information flow arising from typable programs is indeed of the kind allowed by the security lattice, that is from τ to σ where $\tau \leq \sigma$.

As is well-known, bisimulation equivalence is a strong relation, and therefore our noninterference result is quite sharp. In particular, bisimulation equivalence is usually stronger than *trace semantics*, and therefore we can draw some consequences of our result regarding the trace behaviour of programs. To this end, let us introduce a notation: we write $(P, \mu) \rightsquigarrow^{\mathcal{L}}_{\Gamma} (P', \mu')$ if and only if $\mu' \neq^{\mathcal{L}}_{\Gamma} \mu$ and there exist $P_0, \mu_0, \ldots, P_n, \mu_n$ (possibly with $n = 0$) such that $\mu_i =^{\mathcal{L}}_{\Gamma} \mu$ for all i and

$$(P, \mu) = (P_0, \mu_0) \rightarrow (P_1, \mu_1) \rightarrow \cdots (P_n, \mu_n) \rightarrow (P', \mu')$$

Given an initial memory μ_0, let us define the set $\mathsf{Trace}^{\mathcal{L}}_{\Gamma}(P, \mu_0)$ of $(\mathcal{L}, \Gamma)$-*traces* of the program P as the set of finite or infinite sequences $\mu_0, \mu_1, \ldots \mu_n \ldots$ such that there exist $P_1, \ldots P_n \ldots$ with

$$(P, \mu_0) \rightsquigarrow^{\mathcal{L}}_{\Gamma} (P_1, \mu_1) \rightsquigarrow^{\mathcal{L}}_{\Gamma} \cdots (P_n, \mu_n) \rightsquigarrow^{\mathcal{L}}_{\Gamma} \cdots$$

Notice that, by definition, $\mu_{i+1} \neq^{\mathcal{L}}_{\Gamma} \mu_i$ for all i. The $(\Gamma, \mathcal{L})$-equality of memories is easily extended into a $(\Gamma, \mathcal{L})$-equality of traces of the same length, as follows:

$$\mu_0, \mu_1, \ldots \mu_n \ldots =^{\mathcal{L}}_{\Gamma} \nu_0, \nu_1, \ldots \nu_n \ldots \;\Leftrightarrow_{\text{def}}\; \forall i.\ \mu_i =^{\mathcal{L}}_{\Gamma} \nu_i$$

Then we can show that a typable program has the same traces (up to $=^{\mathcal{L}}_{\Gamma}$) when execution starts from $(\Gamma, \mathcal{L})$-equal memories:

Corollary 3.13 (Trace Equivalence)
Let $\mathcal{L}$ be a downward-closed set of security levels. If P is typable in Γ and $\mu =^{\mathcal{L}}_{\Gamma} \nu$ then for any trace $t \in \mathsf{Trace}^{\mathcal{L}}_{\Gamma}(P, \mu)$ there exists $t' \in \mathsf{Trace}^{\mathcal{L}}_{\Gamma}(P, \nu)$ such that $t' =^{\mathcal{L}}_{\Gamma} t$.

Proof: By the Theorem 3.12, we have $(P, \mu) \simeq^{\mathcal{L}}_{\Gamma} (P, \nu)$, and it is easy to see that $(P, \mu) \rightsquigarrow^{\mathcal{L}}_{\Gamma} (P', \mu')$ and $(P, \mu) \simeq^{\mathcal{L}}_{\Gamma} (Q, \nu)$ implies that there exist Q' and ν' such that $(Q, \nu) \rightsquigarrow^{\mathcal{L}}_{\Gamma} (Q', \nu')$. □

4 Adding a Scheduler

As pointed out by Smith and Volpano in [11], their noninterference result for concurrent threads relies on the hypothesis of a purely nondeterministic execution. This result breaks down if particular scheduling policies are enforced. We recall the example given in [11]. Assume a *round robin time slicing* scheduler, with a time slice of t steps, $t \geq 2$, and consider the composition $P = P_1 \parallel P_2$ of the following two threads:

$$\begin{array}{rcl} P_1 & = & \texttt{if } x_H = 0 \texttt{ then } Q \texttt{ else nil}; \\ & & y_L := 0 \\ P_2 & = & y_L := 1 \end{array} \tag{3}$$

Then, supposing that Q is a convergent program that takes at least $t-1$ steps to execute, and that the scheduler gives precedence to P_1, the value of y_L will depend on that of x_H. The solution proposed in [11] to preserve noninterference in the presence of an arbitrary scheduler consists in forbidding conditionals with high guards[3], that is, again assuming that there is a lowest security level, to accept only conditional branching on low level expressions. This condition, combined with the exclusion of loops with high guards, required for multi-threading, resulted in [11] in a very severe limitation: the impossibility for any program to test a variable, except at the lowest level.

One may observe that the program P_1 of Example (3) is actually ruled out by our typing rule for conditional branching, which ensures that low assignments cannot be performed after a high test (*cf.* Example 2). To show that our typing is indeed adequate for dealing with scheduling, we first formalise what it means for a system of concurrent programs to be controlled by a scheduler. Essentially, this means running the system in lockstep with a program that implements the scheduling policy. To describe controlled execution, we use a construction $P[Q]$, which makes P and Q move hand in hand, but allows the controller, P, to proceed by itself whenever Q is unable to move. Then a system consisting of n parallel programs P_i controlled by a scheduler *Sched* will be described as:

$$\mathit{Sched}\ [P'_1 \parallel \cdots \parallel P'_n]$$

where the P'_i are adaptations of the P_i, so that the threads can be triggered and suspended by the scheduler. To this end we introduce a new construct `when` e `do` P, whose semantics is that P is allowed to proceed, for one step, when the condition e holds. It is technically convenient to introduce another level in the syntax: besides the programs P, written according to the

[3]It is also suggested there that a better approach to scheduling would be *probabilistic*. Indeed a whole line of research on probabilistic noninterference has been developed, but this will not be our concern here, where we stick to a *possibilistic* setting.

$$\text{(CONTROL-OP1)} \quad \frac{(S,\mu) \to (S',\mu'),\ (T,\mu) \to (T',\mu'')}{(S[T],\mu) \to (S'[T'],\mu' \sqcup_\mu \mu'')}$$

$$\text{(CONTROL-OP2)} \quad \frac{(S,\mu) \to (S',\mu'), \quad (T,\mu) \not\to}{(S[T],\mu) \to (S'[T],\mu')}$$

$$\text{(WHEN-OP)} \quad \frac{\mu(e) = tt, \qquad (S,\mu) \to (S',\mu')}{(\texttt{when}\ e\ \texttt{do}\ S,\mu) \to (\texttt{when}\ e\ \texttt{do}\ S',\mu')}$$

Figure 4: Additional Operational Rules for Systems

---

grammar given in Section 2, there is a set of "systems" S, T built as follows:

$$S \ ::= \ P \ \mid \ S \parallel T \ \mid \ S[T] \ \mid \ \texttt{when}\ e\ \texttt{do}\ S$$

Letting $w(S)$ denote the set of variables written (assigned to) by S, the construct $S[T]$ is only legal under the condition $w(S) \cap w(T) = \emptyset$.

Notation 4.1 *We use* $(S,\mu) \to$ *to mean* $\exists S', \mu'$ *such that* $(S,\mu) \to (S',\mu')$, *and* $(S,\mu) \not\to$ *for the negation of* $(S,\mu) \to$.

The semantics of the new constructs is given in Figure 4, where $\mu' \sqcup_\mu \mu''$ represents the memory μ with the conjunction of the updates operated by S and by T, that is $\mu' \backslash \mu \cup \mu'' \backslash \mu \cup (\mu' \cap \mu'')$. Then for instance the scheduled programs may be written $P_i' = \texttt{when}\ s_i\ \texttt{do}\ P_i$ where s_i is the "*proceed*" signal for program P_i, set up by the scheduler. The following program

$$\begin{array}{lcl} \mathit{Sched}\ {}^{t}_{n} & = & i := 0\,; \\ & & \texttt{while}\ tt\ \texttt{do} \quad i := [i+1]_{\mathrm{mod}\ n}; k := 0; \\ & & \qquad\qquad \texttt{while}\ k < t\ \texttt{do}\ s_i := tt; s_i := \mathit{ff}; k := k+1 \end{array}$$

describes a scheduler for a system of n threads, implementing round robin with time slice t, provided that all the s_i's are initially false. It is easy to imagine how to program other scheduling policies in a similar style.

The typing rules for the new operators are given in Figure 5. The side-conditions in rules (CONTROL) and (WHEN) need some comments. First, note that a `when` statement can induce an implicit flow, just like the conditional and `while` statements, as for instance in the system:

$$\texttt{when}\ x_H = 0\ \texttt{do}\ y_L := y_L + 1$$

$$\text{(CONTROL)} \quad \frac{\Gamma \vdash S : (\tau, \sigma)\, cmd, \quad \Gamma \vdash T : (\tau, \sigma)\, cmd, \quad \tau \geq \sigma}{\Gamma \vdash S[T] : (\tau, \sigma)\, cmd}$$

$$\text{(WHEN)} \quad \frac{\Gamma \vdash e : \theta, \quad \Gamma \vdash S : (\tau, \sigma)\, cmd, \quad \theta \leq \tau}{\Gamma \vdash \texttt{when}\ e\ \texttt{do}\ S : (\tau, \theta \sqcup \sigma)\, cmd}$$

Figure 5: Additional Typing Rules for Systems

This explains the requirement $\theta \leq \tau$ in rule (WHEN). On the other hand, the condition that the guard of the `when` statement should affect its guard type may seem superfluous at first sight, since a `when` statement can never be followed (in sequential composition) by any other system. The reason for this condition is that in a controlled system $S[T]$, a *blocked* behaviour of the controller S can create interferences if the controlled system T is low, and this blocked behaviour of S may be due to a `when` statement. Consider for instance the system $S[T]$ where P is a high program that does at least one step:

$$\begin{aligned} S &= \texttt{when}\ x_H = 0\ \texttt{do}\ P \\ T &= y_L := y_L + 1 \end{aligned}$$

We let the reader check that this system can lead to interference. Now if the `when` statement S were allowed to have type $(L, L)\, cmd$, the whole system $S[T]$ would be typable.

As regards the rule (CONTROL), the condition $\tau \geq \sigma$ excludes for instance – if the security levels are L and H – systems $S[T]$ whose unique type is $(L, H)\, cmd$. Consider for instance the controlled system $S'[T']$, where:

$$\begin{aligned} S' &= \texttt{while}\ x_H = 0\ \texttt{do}\ \texttt{nil} \\ T' &= y_L := 0\ ;\ y_L := y_L + 1 \end{aligned}$$

Here again there is a possible interference due to the blocking of the controller after one step if the guard of the loop is false. Note that the only possible type of $S'[T']$ would be indeed $(L, H)\, cmd$, since it affects a low variable and has a high loop guard.

It is easy to check that the Subject Reduction Theorem and the Confinement Lemma, and therefore also the Proposition 3.10, extend to the new language. Similarly, the definitions of $(\Gamma, \mathcal{L})$-bisimulation and $\mathcal{L}$-boundedness remain formally the same as those for the base language (modulo the replacement of programs by systems). Obviously, the Guard Safety Lemma may now be strengthened into:

Lemma 4.2 (Strong Guard Safety)
If $\Gamma \vdash S : (\tau, \sigma)\, cmd$ *then every loop, conditional or when statement guard in* S *has type* θ *in* Γ*, with* $\theta \leq \sigma$*.*

It is also easy to see that the Lemma 3.11 still holds in the extended language, thanks to the condition $\theta \leq \tau$ in the rule (WHEN). We are now able to generalise our noninterference result.

Theorem 4.3 (Extended noninterference)
If S *is typable in* Γ*, then* $(S, \mu) \simeq_\Gamma^{\mathcal{L}} (S, \nu)$ *for any downward-closed set* $\mathcal{L}$ *of security levels, and for any* μ, ν *such that* $\mu =_\Gamma \nu$*.*

Proof: We define inductively the relation $\mathcal{R}_1^{\Gamma,\mathcal{L}}$ as follows: $(S, \mu)\, \mathcal{R}_1^{\Gamma,\mathcal{L}}\, (T, \nu)$ if and only if S and T are typable in Γ, $\mu =_\Gamma^{\mathcal{L}} \nu$ and one of the following holds:

1. $(S, \mu)\, \mathcal{R}_0^{\Gamma,\mathcal{L}}\, (T, \nu)$, where $\mathcal{R}_0^{\Gamma,\mathcal{L}}$ is the relation considered in the proof of Theorem 3.12
2. S and T are non-$\mathcal{L}$-bounded
3. $S = T$ and S is $\mathcal{L}$-bounded
4. $S = S_0 \parallel S_1$, $T = T_0 \parallel T_1$ and $(S_i, \mu) \mathcal{R}_1^{\Gamma,\mathcal{L}}\, (T_i, \nu)$
5. $S = \texttt{when}\ e\ \texttt{do}\ S_1$, $T = \texttt{when}\ e\ \texttt{do}\ T_1$, $\Gamma(e) \leq \mathcal{L}$ and $(S_1, \mu) \mathcal{R}_1^{\Gamma,\mathcal{L}}\, (T_1, \nu)$

Then we show that $\mathcal{R}_1^{\Gamma,\mathcal{L}}$ is a $(\Gamma, \mathcal{L})$-bisimulation. As in the proof of Theorem 3.12, it is easy to see that this is enough to show the intended result. We proceed by induction on the definition of $\mathcal{R}_1^{\Gamma,\mathcal{L}}$ to show that, if $(S, \mu) \mathcal{R}_1^{\Gamma,\mathcal{L}} (T, \nu)$ and $(S, \mu) \to (S', \mu')$ then there is a matching move, that is $(T, \nu) \twoheadrightarrow (T', \nu')$ with $(S', \mu') \mathcal{R}_1^{\Gamma,\mathcal{L}} (T', \nu')$. The proof in the case of Clause 1 is just the proof of Theorem 3.12. The proof in the case of Clause 2 is the same as the one of Proposition 3.10. We consider now the case of Clause 3. Here the pair is of the form $((S, \mu), (S, \nu))$, and we proceed by induction on the structure of S. If S is a program P, the pair is already in $\mathcal{R}_0^{\Gamma,\mathcal{L}}$. The case where S is a parallel composition of systems is handled just as for programs. We examine the remaining two cases.

(i) $S = S_1[S_2]$. Suppose $(S, \mu) \to (S', \mu')$ is derived by the rule

$$\frac{(S_1, \mu) \to (S_1', \mu'),\ (S_2, \mu) \to (S_2', \mu'')}{(S_1[S_2], \mu) \to (S_1'[S_2'], \mu' \sqcup_\mu \mu'')}$$

Let $\Gamma \vdash S : (\tau, \sigma)\, cmd$. Since S is $\mathcal{L}$-bounded, $\tau \in \mathcal{L}$, thus by rule (CONTROL) we deduce $\Gamma \vdash S_i : (\tau, \sigma)\, cmd$ with $\sigma \in \mathcal{L}$. Then by the (extended) Lemma 3.11 we know that $(S_1, \nu) \to (S_1', \nu')$ with $\mu' =_\Gamma \nu'$, and $(S_2, \nu) \to$

(S_2', ν'') with $\mu'' =_\Gamma \nu''$. Thus recomposing with rule (CONTROL-OP1) we obtain

$$(S_1[S_2], \nu) \to (S_1'[S_2'], \nu' \sqcup_\nu \nu'')$$

where $(S_1'[S_2'], \mu' \sqcup_\mu \mu'')\, \mathcal{R}_1^{\Gamma,\mathcal{L}}\, (S_1'[S_2'], \nu' \sqcup_\nu \nu'')$, by Clause 3 if $S_1'[S_2']$ is $\mathcal{L}$-bounded, and by Clause 2 otherwise.

Consider now the case where only the controller moves, that is when rule (CONTROL-OP2) is used. We have here:

$$\frac{(S_1, \mu) \to (S_1', \mu'), \quad (S_2, \mu) \not\to}{(S_1[S_2], \mu) \to (S_1'[S_2], \mu')}$$

As in the previous case, Lemma 3.11 gives us $(S_1, \nu) \to (S_1', \nu')$, with $\mu' =_\Gamma \nu'$, and $(S_2, \nu) \not\to$. Then we can apply rule (CONTROL-OP2) to get the matching move $(S_1[S_2], \nu) \to (S_1'[S_2], \nu')$, and conclude as in the previous case.

(ii) $S = \texttt{when}\ e\ \texttt{do}\ T$. In this case $(S, \mu) \to (S', \mu')$ is deduced by rule (WHEN-OP), that is

$$\frac{\mu(e) = tt, \quad (T, \mu) \to (T', \mu')}{(\texttt{when}\ e\ \texttt{do}\ T, \mu) \to (\texttt{when}\ e\ \texttt{do}\ T', \mu')}$$

Since S is $\mathcal{L}$-bounded, by rule (WHEN) we know that T is $\mathcal{L}$-bounded and $\Gamma(e) \in \mathcal{L}$. From the first fact, using induction we deduce $(T, \nu) \twoheadrightarrow (T'', \nu')$ with $(T', \mu')\mathcal{R}_1^{\Gamma,\mathcal{L}}(T'', \nu)$. From the second fact we deduce $\mu(e) = \nu(e)$. Then we can use rule (WHEN-OP) to get

$$(\texttt{when}\ e\ \texttt{do}\ T,\ \nu) \twoheadrightarrow (\texttt{when}\ e\ \texttt{do}\ T'', \nu')$$

which is the required move since $(\texttt{when}\ e\ \texttt{do}\ T', \mu')\, \mathcal{R}_1^{\Gamma,\mathcal{L}}\, (\texttt{when}\ e\ \texttt{do}\ T'', \nu')$ by Clause 5.

The remaining cases of the proof, Clauses 4 and 5, are easy to prove, by a direct use of the induction hypothesis. Notice that in the case of Clause 5, we use, as above, the fact that $\Gamma(e) \in \mathcal{L}$ to deduce $\mu(e) = \nu(e)$. □

5 Conclusion and related work

We have addressed the question of secure information flow in systems of concurrent programs. This covers one of the security problems that can arise, for instance, when a mobile program visits different sites, namely that of preserving the *confidentiality* of the visited sites' private data. In fact, in [2], it is shown how a form of noninterference called *non deducibility on composition* may be used to model also other security properties like *authenticity*, *non repudiation* and *fairness*. Noninterference thus appears as a rather interesting notion to study when security is concerned. On the other hand,

it may be argued [7] that *covert channels*, that is implicit information flows of the kind considered here, are unavoidable in practice, as they can arise also at the hardware level. Thus the aim of statically ensuring the absence of covert channels might be a hard one to realise. We certainly do not claim here to cover the whole range of possible attacks from a hostile party.

The issue of noninterference has been largely studied in the literature, using different models, and it is not our intention here to review the various approaches. We focussed on the approach of Volpano et al., as it applies to a fairly standard language, which can be assumed to be the kernel of more sophisticated practical languages.

The question of secure flow and noninterference has also started to be investigated in the setting of process calculi, and in particular in mobile process calculi [5], [6], [9] and [4]. The treatment in the first two papers seems however overly restrictive: it amounts (at least in the core calculus) to forbid all control flow from actions on high channels to actions on low channels. In [6], the core calculus is extended with more sophisticated constructs; in the extended calculus some actions may be classified as "innocuous", and the restriction on control flow may be relaxed when these actions are involved. The last two papers, [9] and [4], are less restrictive and closer in spirit to our approach, as they try to distinguish the dangerous control flow (implementing information flow) from the harmless control flow which should not be restricted. Another related paper is [1], which studies secrecy properties in security protocols expressed in the *spi*-calculus.

As concerns noninterference in the presence of scheduling policies, the most popular approach has been so far the probabilistic one, taken for instance in [13] and [10]. Our stand here was to handle scheduling within a possibilistic setting.

An issue which has not been addressed here, but is planned for future work, is the feasibility of checking noninterference using a type inference algorithm, in the line of [12]. Current work is also oriented towards the treatment of more realistic languages, as advocated for instance in [8], including exceptions and some form of higher-order.

References

[1] M. Abadi. Secrecy by typing in security protocols. *Journal of the ACM*, 46(5):749–786, 1999.

[2] R. Focardi, R. Gorrieri, and F. Martinelli. Non interference for the analysis of cryptographic protocols. In *Proceedings ICALP'00*, number 1853 in LNCS, 2000.

[3] J. A. Goguen and J. Meseguer. Security policies and security models. In *Proceedings 1982 IEEE Symposium on Security and Privacy*, pages 11–20, 1982.

[4] M. Hennessy. The security π-calculus and noninterference. Computer Science Technical Report 2000:05, University of Sussex, 2000.

[5] M. Hennessy and J. Riely. Information flow vs resource access in the asynchronous pi-calculus (extended abstract). In *Proceedings ICALP'00*, number 1853 in LNCS, 2000.

[6] K. Honda, V. Vasconcelos, and N. Yoshida. Secure information flow as typed process behaviour. In *Proceedings ESOP'00*, number 1782 in LNCS, 2000.

[7] J. Millen. 20 years of covert channel modeling and analysis. In *IEEE Symposium on Security and Privacy*, 1999.

[8] A. Myers. Jflow: Practical mostly-static information flow control. In *26th ACM Symposium on Principles of Programming Languages (POPL)*, 1999.

[9] F. Pottier and S. Conchon. Information flow inference for free. In *Proceedings ICFP'00*, 2000.

[10] A. Sabelfeld and D. Sands. Probabilistic noninterference for multi-threaded programs. In *IEEE Computer Security Foundations Workshop*, 2000.

[11] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In ACM, editor, *Proceedings POPL '98*, pages 355–364. ACM Press, 1998.

[12] D. Volpano and G. Smith. A type-based approach to program security. In *TAPSOFT'97*, number 1214 in LNCS, pages 607–621, 1997.

[13] D. Volpano and G. Smith. Probabilistic noninterference in a concurrent language. *Journal of Computer Security*, 7(2-3), 1999.

[14] D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):167–187, 1996.