

On Declassification and the Non-Disclosure Policy (*)

Ana Almeida Matos Gérard Boudol
INRIA
06902 Sophia Antipolis – France

Abstract

We address the issue of declassification in a language-based security approach. We introduce, in a Core ML-like language with concurrent threads, a declassification mechanism that takes the form of a local flow policy declaration. The computation in the scope of such a declaration is allowed to implement information flow according to the local policy. This dynamic view of information flow policies is supported by a concrete presentation of the security lattice, where the confidentiality levels are sets of principals, similar to access control lists. To take into account declassification, and more generally dynamic flow policies, we introduce a generalization of non-interference, that we call the non-disclosure policy, and we design a type and effect system for our language that enforces this policy.

1. Introduction

This paper addresses the issue of declassification in a language-based security approach. We are therefore more generally concerned with the confidentiality aspect of security. It has often been argued (see [11, 19, 31, 38] for instance) that the standard techniques used for access control are not enough to fully protect confidential information. Ideally, one would like to have a way of controlling how this information is used by subjects having the required clearance. Indeed, it is useless to restrict access to confidential information if one does not have some guarantee that the authorized subjects will not publicly disclose a significant part of this information. In other words, one should be interested in how information flows in a computer system, especially when the “subjects” are programs roaming over the web, that are not easily amenable to non-disclosure agreement.

Since Bell and La Padula and Denning’s pioneering works [3, 11], the classical approach to secure information flow is to use a lattice of security levels (see for instance the survey [40] for the use of security lattices). The objects

of a system are then labelled by security levels, and information is allowed to flow from one object to another if the source object has a lower confidentiality level than the target one. That is, the ordering relation on security levels determines the legal flows, and a program is secure if, roughly speaking, it does not set up illegal flows from inputs to outputs. This was first formally stated via a notion of *strong dependency* by Cohen in [9], and is also referred to as *non-interference* according to the terminology used by Goguen and Meseguer in [16].

A lot of work has been devoted to the design of methods for analyzing information flow in programs (see for instance [2] for early references), even though this has not always been related to a security property like non-interference by a soundness result. Some of these methods consist in runtime checks, and have been criticized for various reasons, like for the fact that they generally suffer from the “label creep problem” (see [38]). More importantly, the failure of a run-time check can serve as a covert channel [11, 30]. As an alternative, static analysis methods have been developed for information flow. One can highlight the use of type systems, which started with the work of Volpano, Smith and Irvine [50]. Although they offer only approximate analysis, type systems have well-known advantages, like in preventing some programming errors at an early stage. Indeed, in the context of a security policy like the one provided by a flow lattice of security levels, it is an error not to comply with the policy, and a type safety result should in this case establish that well-typed programs are secure. Type systems for secure information flow have been designed for various languages (see e.g. [6, 10, 17, 33, 42, 43, 47, 50, 54], and further references in [38]), culminating with Jif (or JFlow, see [29]) and Flow CAML [41] as regards the size of the language. In this paper we shall base our study on Core ML [28, 51], a call-by-value λ -calculus extended with imperative constructs that we enrich with concurrent threads.

The classical non-interference property has been a matter of debate from various points of view (see [36]). A fundamental observation, which was made very early (e.g. in [18]), is that non-interference rules out, by its very definition, programs that deliberately *declassify* information from a confidential level to a more public one. These programs

(*) Work partially supported by the CRISS project of the ACI Sécurité Informatique. The first author is supported by the PhD scholarship POSI/SFRH/BD/7100/2001.

are quite common and very useful, making it hard to use non-interference in practice. A standard example is a password checking procedure, which delivers to any user the result of comparing a submitted password with secret information contained in a database, thus leaking a bit of confidential information. Another one is encryption, where secret information is encoded into a ciphertext that can be read by anyone. Besides these classical security issues, some new circumstances where downgrading information is required arise in the context of networked systems. A typical example is the one of a service selling electronic information, like articles in a journal for instance. The contents of an article have to be kept secret from the client of such a service until he has paid for it, or has identified himself as a subscriber. Then the designer of the electronic purchase service has to provide a procedure that dynamically declassifies information, depending on informations provided by the client. To support this kind of programming, it would be useful for the programmer to have means to check that his code implements only the intended information flows. Our goal here is to provide the programmer with such a support.

The incompatibility of information flow security (in its current setting) with declassification is a challenging problem that has motivated a lot of work. We will comment on this later. This paper intends to contribute to its solution by proposing a turn on how the problem is set. Our view is that there are two different issues to be considered:

1. How may we *justify* that a program is allowed to declassify information, i.e. that it is not actually revealing “too much”?
2. How may we *accept* such programs in a language-based security setting, while still preserving some secure information flow property?

To illustrate this distinction, let us imagine for a while a password checking procedure that returns the root password, instead of a “yes/no” answer. Although the first is most probably wrong, the two programs are no different from the point of view of information flow: they both disclose information from the security level of the secret password database. One should therefore be able to reject the first by a semantical program analysis, while the second should be accepted in a language supporting downgrading facilities.

There is clearly a quantitative aspect to the first question. Indeed, some researchers who have studied it have proposed to quantify the amount of information that a program may leak, or to use complexity-theoretic or probabilistic arguments to establish that it is not feasible to exploit the allowed information leakage (see [7, 12, 21, 22, 46, 49], to mention just a few recent papers). Justifying declassification is a very interesting research problem which is by no

means easy – in fact, justifying practically sound encryption mechanisms is a whole research domain – and seems to be beyond the reach of static analysis techniques. However, in our view, it is the *programmer* who has responsibility for solving the first question, of *what*, or *how much* information he intends to declassify in his program, whereas the (designer of the) *programming language* has to provide an answer to the second question above. This is the language design issue that we address.

Given that deliberately downgrading programs are validated by the programmer, the programming language should be as flexible as possible in expressing them. To this end, we introduce in our core language a programming construct for directly manipulating flow relations, namely a *flow declaration* construct (flow F in M) where F is a *flow policy*, i.e. a binary relation on security levels, and M is any expression of the language. The meaning is that M is executed in the context of the current flow policy *extended with* F , and after termination the current policy is restored, that is, the scope of F is M . For instance, if we have security levels $A(\textit{lice})$ and $B(\textit{ob})$, then – using the ML notation $!x$ for dereferencing, and x_A, y_B for memory locations with confidentiality level A and B – a program like:

$$(\text{flow } A \prec B \text{ in } y_B := !x_A) \quad (1)$$

is legal, since in the context of the relation $A \prec B$, information is allowed to flow from A to B . With respect to the current flow policy, this is a declassification operation – unless, obviously, the current policy already says that information may flow from A to B . Moreover, this expression appears to read at the confidentiality level B for the rest of the program. Then the expression $y_B := (\text{flow } A \prec B \text{ in } !x_A)$ is also legal, and has the same meaning as the previous one. It should be clear, on the other hand, that a statement like $y_B := (\text{flow } C \prec B \text{ in } !x_A)$ is not legal, unless the current policy allows information to flow from level A to C (or B). Another example is

$$(\text{flow } A \prec B \text{ in } M) ; (\text{flow } B \prec C \text{ in } N)$$

that shows a way to achieve a kind of non-transitive flow relation (see [4, 34, 35]).

A similar construct for introducing flow policies exists in Flow CAML, but with an important difference: there it adds the flow relation F to the global security policy, whereas in our case the declaration is *local*. Such a construct has been mentioned in [44] under the name of “delegation”, but it was not formally studied there. The operation $\text{declassify}(M, \ell)$, that is used in some languages (see [29, 39] for instance) to downgrade the value of M to the confidentiality level ℓ , may be represented as – again using ML notations

$$\text{let } x = (\text{ref}_H M) \text{ in } (\text{flow } H \prec \ell \text{ in } !x)$$

where $(\text{ref}_H M)$ creates a new memory address with security level H and contents M (that is, the value of M), and H is a security level that is higher (w.r.t the current flow policy) than any other one. This interpretation of the declassification operation as a local operator on the flow policies seems to have never been pointed out. Furthermore, the flow declaration construct allows us to express more precise ways of declassifying, by specifying the levels from which information may flow.

We should point out here that, contrarily to most studies concerning security for functional languages (with the exception of [10]), we do not regard values as having a security level – there is no “secret 0” or “public 100” for instance. Our standpoint is that confidentiality levels are associated with the objects in which information is stored (or the channels on which it is communicated), like files, databases or references in the case of ML, and that what is to be controlled is the access to the object (typically read or write). This is consistent with most studies dealing with imperative languages, and has the pleasant consequence that the security type system is just a standard *type and effect system* [26], with flow constraints, where the security levels play the role of *regions* (as noted in [10]). In the construct (flow F in M), only the effects of M , and the resulting confidentiality level of the expression are concerned with the flow relation F ; the final value of M , which has no confidentiality status, is not.

Once declassification is permitted in the language, a question is: what kind of security property do we have that takes declassification into account? And what means could we have to ensure that programs have this property? Not surprisingly, here the answer to this second question will be: a type and effect system. To answer the first one, we must find an alternative to non-interference. In the language-based security approach, and more specifically in concurrent settings, this property is often based on the small-step semantics, where one specifies transitions $(P, \mu) \rightarrow (P', \mu')$ between successive states of the program and the memory. This is well suited for our approach, where declassification is based on a dynamically evolving structure of the lattice of confidentiality levels. Indeed, the scope of a local flow policy is only a portion of the computation, and this has to be reflected in the semantics in order to state a security property. The way we do this is by decorating each transition with a label, which is the local flow policy that is in force for this particular step:

$$(P, \mu) \xrightarrow{F} (P', \mu') \quad (2)$$

The intuition is that, as regards information flow, the memory μ should be considered from the point of view of the current flow policy extended with F . That is, if F says that information is allowed to flow from level ℓ to level ℓ' , what is read at level ℓ at this step may be regarded as having level

ℓ' . Then our new confidentiality property, which is a generalization of non-interference that we call the *non-disclosure policy*, roughly says that a program P is secure if at each (small) step it satisfies non-interference *with respect to the flow policy that holds for this step*. More precisely, given that P performs the transition (2) above under the memory μ , and given that ν is a memory which only differs from μ regarding confidential information with respect to the current flow policy extended with F , then there is a transition $(P, \nu) \xrightarrow{F'} (P'', \nu')$ from P under memory ν such that ν' is again equal to μ' as regards public information. Moreover, since we have to check this at every possible step, we shall also require that the programs P' and P'' have similar behaviours, from the confidentiality point of view. For instance, program (1) satisfies this property, since the reference x_A may be considered as having the level B under the flow relation $A \prec B$.

The main technical contribution of this paper is a proof that the type and effect system we design for our core language with declassification enforces the non-disclosure policy. We thus provide a direct generalization of the standard result regarding type systems for information flow. We shall start by introducing dynamic flow policies, which provide a way to deal with declassification that, up to our knowledge, has never been formally explored before. This is supported by a specific notion of security (pre-)lattice, which is also new. These will appear in the next section, where we present the language and its operational semantics. Then, in Section 3, we introduce our generalization of non-interference, namely the non-disclosure policy, that takes into account dynamic flow policies. A sound type and effect system is given for the language in Section 4, featuring the notion of a “termination effect” that allows us to deal with a strong notion of security, while still providing a flexible type system. We then briefly discuss related work and conclude. For lack of space, the proofs are omitted; they can be found in the full version of the paper.

2. The Language

2.1 Security (pre-)lattices

As we said in the Introduction, we will use dynamically evolving flow relations for dealing with declassification. However, the security levels associated with references that appear in the expression M are the same as those in (flow F in M) – it is only the flow policy that changes. We are then faced with the issue of maintaining a (varying) lattice structure over a given set of security levels, since we shall use the meet and join operations in the type system, as usual. As a matter of fact, a “pre-lattice” structure turns out to be more convenient for our purpose. We call *pre-lattice* a pair (\mathcal{L}, \preceq)

$M, N \dots \in \mathcal{E}xpr$	$::= W \mid (\text{if } M \text{ then } N \text{ else } N') \mid (MN)$ $\mid M ; N \mid (\text{ref}_{\ell, \theta} N) \mid (!N) \mid (M := N)$ $\mid (\text{thread } M) \mid (\text{flow } F \text{ in } M)$	expressions
$W \in \mathcal{W}$	$::= V \mid \rho x W$	pseudo-values
$V \in \mathcal{V}al$	$::= x \mid u_{\ell, \theta} \mid \lambda x M \mid tt \mid ff \mid ()$	values

Figure 1: Syntax

where \preceq is a preorder on \mathcal{L} , that is a reflexive and transitive (but not necessarily anti-symmetric) relation, such that for any $x, y \in \mathcal{L}$ there exist a meet $x \wedge y$ and a join $x \vee y$ for x and y , satisfying

$$\begin{aligned} x \wedge y &\preceq x \\ x \wedge y &\preceq y \\ z \preceq x \ \& \ z \preceq y &\Rightarrow z \preceq x \wedge y \end{aligned}$$

and

$$\begin{aligned} x &\preceq x \vee y \\ y &\preceq x \vee y \\ x \preceq z \ \& \ y \preceq z &\Rightarrow x \vee y \preceq z \end{aligned}$$

Now we will define our security pre-lattices, where the set of security levels is fixed, and only the flow relations may vary. We assume given a set \mathcal{P} of *principals*, ranged over by $p, q \dots$. A *confidentiality level* is any set of principals, that is any subset ℓ of \mathcal{P} . The intuition is that whenever ℓ is the confidentiality label of an object, i.e. a reference, it represents a set of programs that are allowed to get the value of the object, i.e. to read the reference. From this point of view, a reference labelled \mathcal{P} (also denoted \perp) is a most public one – every program is allowed to read it –, whereas the label \emptyset (also denoted \top) indicates a secret reference, so secret that no one is allowed to read it. We can interpret the reverse inclusion of security levels as indicating allowed flows of information: if a reference x is labelled ℓ , and $\ell \supseteq \ell'$ then the value of x may be transferred to a reference y labelled ℓ' , since the programs allowed to read this value from y were already allowed to read it from x .

The dynamically varying information flow policies are determined by relations on principals. This is slightly different from what we informally presented in the Introduction, where, for simplicity, a flow relation was assumed to relate security levels. However, as we shall see, a relation on principals induces a preorder on security levels. A *flow policy* is a binary relation over \mathcal{P} . We let $F, G \dots$ range over such relations. A pair $(p, q) \in F$ is to be understood as “information may flow from principal p to principal q ”, that is, more precisely, “*everything that principal p is allowed to read may also be read by principal q* ”. We must point out here that, since we are dealing with confidentiality (and not integrity) a flow policy will only affect the reading capabilities of programs (and not their writing capabilities). As a

member of a flow policy, a pair (p, q) will most often be written $p \prec q$. We denote, as usual, by F^* the preorder generated by F (that is, the reflexive and transitive closure of F). Then we introduce the *preorder on confidentiality levels* determined by the flow relation F :

$$\ell \preceq_F \ell' \Leftrightarrow_{\text{def}} \forall q \in \ell'. \exists p \in \ell. p F^* q$$

which is denoted \preceq (instead of \supseteq) when $F = \emptyset$. We shall use without notice the fact that

$$G \subseteq F \ \& \ \ell \preceq_G \ell' \Rightarrow \ell \preceq_F \ell'$$

It is not difficult to see that the preorder \preceq_F induces a pre-lattice structure on the set of confidentiality levels, where a meet is simply the union, and a join of ℓ and ℓ' is

$$\{ q \mid \exists p \in \ell. \exists p' \in \ell'. p F^* q \ \& \ p' F^* q \}$$

This observation justifies the following definition.

DEFINITION (SECURITY PRE-LATTICES) 2.1.

A confidentiality level is any subset ℓ of the set \mathcal{P} of principals. Given a flow policy $F \subseteq \mathcal{P} \times \mathcal{P}$, the confidentiality levels are pre-ordered by the relation

$$\ell \preceq_F \ell' \Leftrightarrow_{\text{def}} \forall q \in \ell'. \exists p \in \ell. p F^* q$$

The meet and join, w.r.t. F , of two security levels ℓ and ℓ' are respectively given by $\ell \cup \ell'$ and

$$\ell \vee_F \ell' = \{ q \mid \exists p \in \ell. \exists p' \in \ell'. p F^* q \ \& \ p' F^* q \}$$

2.2 Syntax and operational semantics

The language is a call-by-value λ -calculus extended with the imperative constructs of ML, conditional branching and boolean values. We also introduce the possibility of dynamically creating concurrent threads, and of declassifying computations. Clearly, the latter is the main novelty, and one could probably deal with other programming paradigms in a similar way, by adding local flow declarations. The syntax is given in Figure 1, where x is any variable, F is any flow policy that is most often written as a list $p_1 \prec q_1, \dots, p_n \prec q_n$ of pairs of principals, and $u_{\ell, \theta}$ is a triple made of a memory address u – or location, or *reference* –, a type θ (see Section 4 below) and a label ℓ which is a confidentiality level. The label ℓ , most often written p_1, \dots, p_n instead of $\{p_1, \dots, p_n\}$, is similar to an access control list. We use locations explicitly decorated with types and confidentiality

$$\begin{array}{c}
((\text{if } tt \text{ then } M \text{ else } N), \mu) \xrightarrow[\emptyset]{} (M, \mu) \\
((\text{if } ff \text{ then } M \text{ else } N), \mu) \xrightarrow[\emptyset]{} (N, \mu) \\
((\lambda x.MV), \mu) \xrightarrow[\emptyset]{} (\{x \mapsto V\}M, \mu) \\
(V ; N, \mu) \xrightarrow[\emptyset]{} (N, \mu) \\
((\text{ref}_{\ell, \theta} V), \mu) \xrightarrow[\emptyset]{} (u_{\ell, \theta}, \mu \cup \{u_{\ell, \theta} \mapsto V\}) \quad u \text{ fresh for } \mu \\
((! u_{\ell, \theta}), \mu) \xrightarrow[\emptyset]{} (V, \mu) \quad \mu(u_{\ell, \theta}) = V \\
((u_{\ell, \theta} := V), \mu) \xrightarrow[\emptyset]{} (\emptyset, \mu[u_{\ell, \theta} := V]) \\
(\varrho xW, \mu) \xrightarrow[\emptyset]{} (\{x \mapsto \varrho xW\}W, \mu) \\
((\text{flow } F \text{ in } V), \mu) \xrightarrow[\emptyset]{} (V, \mu)
\end{array}$$

$$\begin{array}{c}
(M, \mu) \xrightarrow{F} (M', \mu') \\
\hline
(\mathbf{E}[M], \mu) \xrightarrow{F \cup [\mathbf{E}]} (\mathbf{E}[M'], \mu') \qquad \hline
(\mathbf{E}[(\text{thread } M)], \mu) \xrightarrow[\emptyset]{} ((\mathbf{E}[\emptyset]) \parallel (\text{flow } [\mathbf{E}] \text{ in } M)), \mu) \\
(P, \mu) \xrightarrow{F} (P', \mu') \qquad \hline
((P \parallel Q), \mu) \xrightarrow{F} ((P' \parallel Q), \mu') \qquad \hline
(P, \mu) \xrightarrow{F} (P', \mu') \qquad \hline
((Q \parallel P), \mu) \xrightarrow{F} ((Q \parallel P'), \mu')
\end{array}$$

Figure 2: Operational Semantics

labels for the purpose of the proof of type soundness. However, as we shall see, these annotations do not have any role in the operational semantics, and therefore they do not have to appear in an implementation (although in a mobile code setting, one would like to keep these annotations in order to perform security checks when loading a piece of code). We denote by $\text{loc}(M)$ the set of decorated locations occurring in M . These addresses are regarded as providing the *inputs* of the expression M .

For typing reasons, we do not regard sequential composition as a derived construct, and we do not regard the imperative constructs ref , $!$ and $:=$ as first-class functions. Indeed, the typing of $(\text{ref}_{\ell, \theta} N)$ will generally be different from the typing of $(\lambda x.(\text{ref}_{\ell, \theta} x)N)$ for instance. Applying the construct $\text{ref}_{\ell, \theta}$ to a value V creates a new reference with initial value V . Here, as we shall see, the value is assumed to be of type θ , and the confidentiality level ℓ will be assigned to the created reference. While the type θ could probably be inferred, as in ML, it seems natural for security purposes to explicitly assign a confidentiality level to the created reference. In a pure type and effect inference approach, with an unlabelled ref function, we would only get constraints that this level should satisfy. The construct ϱxW , which is a binder for the variable x in W , provides a way to deal with recursive values. As a matter of fact, given that the set of values is quite limited in our core language, the only interesting case is that of recursive functions, i.e. $\varrho f \lambda x.M$, which could be denoted $(\text{let rec } f = \lambda x.M \text{ in } f)$ in an ML-like notation. We denote by loop the expression $\varrho x.x$, and we may

use the following standard abbreviation:

$$(\text{while } M \text{ do } N) =_{\text{def}} (\varrho y \lambda x.(\text{if } M \text{ then } N ; (y \emptyset) \text{ else } \emptyset) \emptyset)$$

We let $\text{fv}(M)$ be the set of variables occurring free in M , and we denote by $\{x \mapsto W\}M$ the capture-avoiding substitution of W for the free occurrences of x in M , where $W \in \mathcal{W}$. The evaluation relation is a transition relation between configurations of the form (P, μ) where P is a *process*, written according to the following syntax:

$$P, Q \dots \in \text{Proc} ::= M \mid (P \parallel Q)$$

and μ , the *memory* (or *heap*), is a mapping from a finite set $\text{dom}(\mu)$ of decorated references to values. The operation of updating the value of a reference in the memory is denoted, as usual, $\mu[u_{\ell, \theta} := V]$. We say that the name u is *fresh* for μ if $v_{\ell, \theta} \in \text{dom}(\mu) \Rightarrow v \neq u$. To define the operational semantics, we introduce evaluation contexts:

$$\begin{array}{l}
\mathbf{E} ::= \square \mid \mathbf{F}[\mathbf{E}] \mid (\text{flow } F \text{ in } \mathbf{E}) \\
\mathbf{F} ::= (\text{if } \square \text{ then } M \text{ else } N) \mid (\square N) \mid (V \square) \\
\quad \mid \square ; N \mid (\text{ref}_{\ell, \theta} \square) \mid (! \square) \mid (\square := N) \mid (V := \square)
\end{array}$$

and we denote by $[\mathbf{E}]$ the flow policy enforced by the context \mathbf{E} . This is defined as follows:

$$\begin{array}{l}
[\square] = \emptyset \\
[\mathbf{F}[\mathbf{E}]] = [\mathbf{E}] \\
[(\text{flow } F \text{ in } \mathbf{E})] = F \cup [\mathbf{E}]
\end{array}$$

The labelled transition rules are given in Figure 2. As one can see, the transitions do not depend on the types and labels of memory locations. Observe also that the flow label F of the transitions plays no role in determining the resulting configuration. To evaluate (flow F in M), we simply evaluate M , until termination (that is, when M is a value). Then (flow F in M) is operationally the same as M , and the context (flow F in \square) should disappear at compile time. We denote somewhat abusively by \rightarrow the relation given by

$$(P, \mu) \rightarrow (P', \mu') \Leftrightarrow_{\text{def}} \exists F. (P, \mu) \xrightarrow{F} (P', \mu')$$

and we let $\xrightarrow{*}$ denote the reflexive and transitive closure of the relation \rightarrow . We shall actually only consider transitions from *well-formed* configurations: a configuration (P, μ) is well-formed if $\text{loc}(P) \subseteq \text{dom}(\mu)$ and for any $u_{\ell, \theta} \in \text{dom}(\mu)$ we have $\text{loc}(\mu(u_{\ell, \theta})) \subseteq \text{dom}(\mu)$. It is easy to see that well-formedness is preserved by transitions.

To conclude this section, we introduce another kind of transitions which is useful for the proof of type soundness. These transitions, denoted $(M, \mu) \xrightarrow{F} (M', \mu')$, should be read as follows: the expression M , in the context of the memory μ , performs a step, assuming the local flow policy F , and resulting in the new expression M' and memory μ' , while possibly spawning the expression N as a new thread to execute (with $N = \emptyset$ if M actually does not spawn any thread). This formalizes the evaluation steps of an expression M as the “main thread”. Formally, this is defined as \xrightarrow{F} , with $N = \emptyset$, except for:

$$\frac{}{(\mathbf{E}[(\text{thread } N)], \mu) \xrightarrow[\emptyset]{(\text{flow } [\mathbf{E}] \text{ in } N)} (\mathbf{E}[\emptyset], \mu)}$$

The following lemma relates these transitions with the ones that we have used to describe the operational semantics:

LEMMA 2.2.

- (i) If $(M, \mu) \xrightarrow{F} (M', \mu')$ then either $N = \emptyset$ and $(M, \mu) \xrightarrow{F} (M', \mu')$ or $(M, \mu) \xrightarrow{F} ((M' \parallel N), \mu')$.
- (ii) If $(M, \mu) \xrightarrow{F} (P, \mu')$ then either P is an expression and $(M, \mu) \xrightarrow{F} (P, \mu')$ or $(M, \mu) \xrightarrow{F} (M', \mu')$ for some M' and N such that $P = (M' \parallel N)$.

Next we make a simple but crucial observation, stating that, if the evaluation of an expression M differs in the context of two distinct memories while not creating two distinct references, this is because M is performing a dereferencing operation, which yields different results depending on the memory. Apart from non-deterministically choosing new references, this is the only way for computations of expressions to split.

LEMMA 2.3.

If $(M, \mu) \xrightarrow{F} (M', \mu')$ and $(M, \nu) \xrightarrow{F'} (M'', \nu')$ with $M' \neq M''$ and $\text{dom}(\nu' - \nu) = \text{dom}(\mu' - \mu)$, then $N = \emptyset = N'$ and there exist \mathbf{E} and $u_{\ell, \theta}$ such that $F = [\mathbf{E}] = F'$, $M = \mathbf{E}[\!(u_{\ell, \theta})]$, and $M' = \mathbf{E}[\mu(u_{\ell, \theta})]$, $M'' = \mathbf{E}[\nu(u_{\ell, \theta})]$ with $\mu' = \mu$ and $\nu' = \nu$.

3. The non-disclosure policy

In this section we introduce our security property. The definitions that follow could be formulated in an abstract way, that is for any semantic framework consisting of a given labelled transition system, where the transitions have, as above, the form

$$(P, \mu) \xrightarrow{F} (P', \mu')$$

However, we instantiate here the definitions with the notion of a process that we have introduced in the previous section.

To state the security property, we use, as it is standard, a notion of *memory equality* relative to a given confidentiality level: two memories μ and ν are equal up to level ℓ if they assign the same value to every location with security level lower than ℓ (this is sometimes referred to as “low equality” of memories). However, there is an implicit parameter in this notion, which is the flow relation used to determine that a level is lower than ℓ . Since the flow policy is not fixed in our setting, we make it explicit in the notion of memory equality. Furthermore, we will compare two memories only with respect to the references they share. The low equality of memories is thus defined:

$$\begin{aligned} \mu \simeq^{F, \ell} \nu &\Leftrightarrow_{\text{def}} \forall u_{\ell', \theta} \in \text{dom}(\mu) \cap \text{dom}(\nu). \\ &\ell' \preceq_F \ell \Rightarrow \mu(u_{\ell', \theta}) = \nu(u_{\ell', \theta}) \end{aligned}$$

This relation is not transitive, but it is reflexive and symmetric. We shall use without notice the fact that

$$G \subseteq F \ \& \ \mu \simeq^{F, \ell} \nu \Rightarrow \mu \simeq^{G, \ell} \nu$$

Our security property is defined in terms of *bisimulations* (see [6, 14, 37, 42] for the use of bisimulations in stating security properties, and [25] for a review of various other approaches). Bisimulations are relations on states of transition systems, that relate two states whenever any transition of either of these states can be “matched” by a transition of the other. Following [37], here we shall make use of this notion in a slightly non-standard way, since we shall call “bisimulation” a relation between processes P , rather than configurations (P, μ) . Moreover, such a relation is parameterized upon a flow policy G , which is the current flow relation.

DEFINITION (BISIMULATION) 3.1.

A (G, ℓ) -bisimulation is a symmetric relation \mathcal{R} on processes such that if

$$P \mathcal{R} Q \ \& \ (P, \mu) \xrightarrow{F} (P', \mu') \ \& \ \mu \simeq^{F \cup G, \ell} \nu \ \& \\ u_{\ell', \theta} \in \text{dom}(\mu' - \mu) \Rightarrow u \text{ is fresh for } \nu$$

then there exist Q' and ν' such that

$$(Q, \nu) \xrightarrow{*} (Q', \nu') \ \& \ P' \mathcal{R} Q' \ \& \ \mu' \simeq^{G, \ell} \nu' \ \& \\ \text{dom}(\nu' - \nu) \subseteq \text{dom}(\mu' - \mu)$$

It is implicit in this definition that the configurations (P, μ) and (Q, ν) are well-formed. This implies in particular that $\text{loc}(P) \subseteq \text{dom}(\mu)$ and $\text{loc}(Q) \subseteq \text{dom}(\nu)$. The definition says that the transition

$$(P, \mu) \xrightarrow{F} (P', \mu')$$

has to be matched by a transition from (Q, ν) , whenever $P \mathcal{R} Q$, and the memories μ and ν satisfy some conditions. The first one, namely that $\mu \simeq^{F \cup G, \ell} \nu$, can be interpreted as follows: since P is performing a transition within the scope of the current flow policy G extended with the local flow relation F , it is allowed to read references from the input memory according to the policy $F \cup G$. That is, in the input memories μ and ν , “low” means less than ℓ with respect to $F \cup G$. The condition “ $u_{\ell', \theta} \in \text{dom}(\mu' - \mu) \Rightarrow u$ is fresh for ν ” simply means that if P creates a new reference, then we assume that the created name does not conflict with other names under consideration. Indeed, this new name can always be chosen so that it satisfies this constraint. The conclusion is that the state (Q, ν) should be able to evolve, possibly in several steps, into a configuration (Q', ν') , satisfying the following constraints: first, $P' \mathcal{R} Q'$ means that no mismatch should occur in future computations. Next, we require $\mu' \simeq^{G, \ell} \nu'$, thus considering the output μ' of the step $(P, \mu) \xrightarrow{F} (P', \mu')$ from the point of view of the flow policy that is restored after it, that is G . This means in particular that the local flow policy F does not affect the level of references from the writing, or output point of view (recall that $p \prec q$ means that “everything that principal p is allowed to read may also be read by principal q ”). Finally $\text{dom}(\nu' - \nu) \subseteq \text{dom}(\mu' - \mu)$ ensures that Q only creates a reference to match P 's transition if P itself has created a reference, in which case we require the new location names to be the same. We need this in order to get the transitivity of the greatest bisimulation. The reader may have noticed that there is no condition on the flow policy of the matching moves for Q . This is because we wish to consider all the *pure* programs, written without $(!N)$ and $(M := N)$, to be bisimilar.

As mentioned above, the notion of bisimulation we use is stronger than the standard one, since if the transition $(P, \mu) \xrightarrow{F} (P', \mu')$ is matched by $(Q, \nu) \xrightarrow{*} (Q', \nu')$, we

restart the bisimulation game by comparing the processes P' and Q' , in the context of any new low equal memories, rather than the configurations (P', μ') and (Q', ν') . This allows us to restore a more restrictive flow relation after a local flow relation has been used, as in

$$(\text{flow } H \prec L \text{ in } v_L := !u_H) ; w_L := !u'_H \quad (3)$$

where the second assignment implements an illegal flow (denoting simply by $H, L \dots$ a singleton security level $\{H\}, \{L\} \dots$ assigned to a reference). Defining bisimulations over processes, rather than between configurations, also allows us to detect an illegal flow in the program

$$(\text{if } !w_X \text{ then } (\text{if } !w_X \text{ then } () \text{ else } v_L := !u_H) \text{ else } ())$$

This demanding definition for bisimulations seems also appropriate for dealing with a mobile code scenario, where the shared memory of a system of threads can be modified by incoming code.

REMARKS AND NOTATION 3.2.

- (i) For any G and ℓ there exists a (G, ℓ) -bisimulation, like for instance the set $\text{Val} \times \text{Val}$ of pairs of values.
- (ii) The union of a family of (G, ℓ) -bisimulations is a (G, ℓ) -bisimulation. Consequently, there is a largest (G, ℓ) -bisimulation, which we denote $\simeq^{G, \ell}$. This is the union of all such bisimulations.

One should observe that $\simeq^{G, \ell}$ is a *partial* equivalence relation. That is, this relation is not reflexive. Indeed, a process which is not bisimilar to itself, like $v_L := !u_H$ if $H \not\prec_G L$, is not secure. As in [37], our definition states that a program is secure if it is bisimilar to itself:

DEFINITION (THE NON-DISCLOSURE POLICY) 3.3.

A process P satisfies the non-disclosure policy (or is secure from the confidentiality point of view) with respect to the flow policy G if it satisfies $P \simeq^{G, \ell} P$ for any ℓ . We then write $P \in \mathcal{ND}(G)$.

It is easily seen that the set $\mathcal{ND}(G)$ is non-empty. For instance, any value is secure. An important property of our notion of security is that if an expression M does not violate the current flow policy G extended with a local policy F , then the expression $(\text{flow } F \text{ in } M)$ is secure with respect to the current flow relation:

PROPOSITION 3.4.

$$M \in \mathcal{ND}(F \cup G) \Rightarrow (\text{flow } F \text{ in } M) \in \mathcal{ND}(G)$$

Then for instance a program like the one of example (1) is secure. Another property is that security is compatible with parallel composition:

PROPOSITION (COMPOSITIONALITY) 3.5.

$$P \in \mathcal{ND}(G) \ \& \ Q \in \mathcal{ND}(G) \Rightarrow (P \parallel Q) \in \mathcal{ND}(G)$$

Our non-disclosure policy generalizes the usual non-interference property for sequential programs (without declassification). To see this point, let us first recall that the latter is based on the “big-step” semantics of programs, that is on the relation $(P, \mu) \Rightarrow \mu'$ that a program P establishes from an initial state of the memory μ to the final state μ' . Namely, P is *non-interfering* if $(P, \mu) \Rightarrow \mu'$ and $(P, \nu) \Rightarrow \nu'$, for μ and ν that differ only regarding confidential information, implies that μ' and ν' are equal as regards public information, that is:

$$(P, \mu) \Rightarrow \mu' \ \& \ (P, \nu) \Rightarrow \nu' \ \& \ \mu \simeq^{G, \ell} \nu \ \Rightarrow \ \mu' \simeq^{G, \ell} \nu'$$

Let us denote for a while by \mathcal{DExpr} the set of expressions written without using thread, flow and ref, and let us show that the expressions in \mathcal{DExpr} satisfying the non-disclosure policy with respect to a given flow policy G are non-interfering. The big-step semantics for expressions in \mathcal{DExpr} can be defined as follows:

$$(M, \mu) \Rightarrow \mu' \ \Leftrightarrow_{\text{def}} \ \exists V \in \text{Val}. \ (M, \mu) \xrightarrow{*} (V, \mu')$$

It is easy to see that the evaluation mechanism is deterministic for $M \in \mathcal{DExpr}$, and that if $(M, \mu) \Rightarrow \mu'$ then $\text{dom}(\mu') = \text{dom}(\mu)$. Now assume that $M \in \mathcal{DExpr} \cap \mathcal{ND}(G)$, $(M, \mu) \xrightarrow{*} (V, \mu')$ and $(M, \nu) \xrightarrow{*} (V', \nu')$ with $\mu \simeq^{G, \ell} \nu$. Then there exist M' and ν'' such that $(M, \nu) \xrightarrow{*} (M', \nu'')$, $V \sqsubset^{G, \ell} M'$ and $\mu' \simeq^{G, \ell} \nu''$. Since M is deterministic, we have $(M', \nu'') \xrightarrow{*} (V', \nu')$, and from (V, μ') there must be a sequence of transitions matching the move from (M', ν'') to (V', ν') . This sequence must be empty, and we then have $\mu' \simeq^{G, \ell} \nu'$.

Now let us see some examples. We assume given two principals H and L , and a current flow relation G consisting of the pair $L \prec H$. We shall denote references with security levels $\{H\}$ or $\{L\}$ simply by u_H or v_L (leaving out the type), as usual. Since, as we have just seen, the non-disclosure policy implies the standard non-interference property for expressions of \mathcal{DExpr} , it is obvious that the standard examples of explicit and implicit flow, namely:

$$v_L := !u_H \quad (4)$$

$$(\text{if } !u_H \text{ then } v_L := tt \text{ else } v_L := ff) \quad (5)$$

do not satisfy the non-disclosure policy, whereas these programs are secure in the context of the flow declaration $(\text{flow } H \prec L \text{ in } \square)$. Since we follow a bisimulation approach to security, we also reject termination leaks, like for instance

$$(\text{if } !u_H \text{ then } () \text{ else loop}) ; v_L := tt \quad (6)$$

where writing at level L depends on reading at level H (we refer to [2, 6, 17, 38, 42, 47] for discussions about this kind of leaks). Another example of a termination leak is

$$(!u_H)(); v_L := tt \quad (7)$$

Indeed, the value of the reference u could be $\lambda y \lambda x x$ or $\lambda y \text{loop}$. Similarly, there is a termination leak in

$$(\lambda x(x))(!u_H) ; v_L := tt \quad (8)$$

since the value of the reference u might be $\lambda y y$ or $\lambda y \text{loop}$. We shall put a constraint on sequential composition in the type system to rule out such programs. However, this constraint will not be as strict as “no low write after a high read”, because we would like to accept for instance the following (secure) program:

$$(w_H := !u_H) ; (v_L := tt) \quad (9)$$

Regarding the flow declaration construct, we notice for instance that the program

$$v_L := (\text{flow } H \prec L \text{ in } !u_H) \quad (10)$$

which is essentially the same as example (1), is secure, as well as

$$(\text{if } !u_H \text{ then } w_H := tt \text{ else } ())$$

whereas

$$(\text{if } !u_H \text{ then } (\text{flow } H \prec L \text{ in } v_L := tt) \text{ else } ()) \quad (11)$$

is not. The reason is that the flow declaration $H \prec L$ is a way of giving (temporarily) the same reading capabilities to the principals H and L , whereas it does not affect the writing capabilities of a program. A type system for information flow has to take this into account.

4. The type and effect system

The types we use in the type and effect system are quite standard. Namely, a reference type θref_ℓ records the type θ of values the reference contains, as well as the “region” ℓ where it is created. Here this is the *confidentiality level* of the reference, indicating who is allowed to read it. A function type records the *latent effect* [26] of a function of that type, which is the effect the function may have when applied to an argument. It also records the “latent flow relation”, which is assumed to hold when the function is applied to an argument. The syntax of types is

$$\tau, \sigma, \theta \dots ::= t \mid \text{bool} \mid \text{unit} \mid \theta \text{ref}_\ell \mid (\tau \xrightarrow[s]{F} \sigma)$$

where t is any type variable and s is any “security effect” – see below. The judgements of the type and effect system have the form

$$G; \Gamma \vdash M : s, \tau$$

where G is a flow relation, Γ is a typing context, assigning types to variables, s is a security effect, that is a triple (ℓ_0, ℓ_1, ℓ_2) of confidentiality levels, and τ is a type. The intuition is:

$$\begin{array}{c}
\frac{}{G; \Gamma \vdash u_{\ell, \theta} : \perp, \theta \text{ref}_{\ell}} \text{ (LOC)} \quad \frac{}{G; \Gamma, x : \tau \vdash x : \perp, \tau} \text{ (VAR)} \\
\frac{F; \Gamma, x : \tau \vdash M : s, \sigma}{G; \Gamma \vdash \lambda x M : \perp, (\tau \xrightarrow{F} \sigma)} \text{ (ABS)} \quad \frac{}{G; \Gamma \vdash () : \perp, \text{unit}} \text{ (NIL)} \\
\frac{}{G; \Gamma \vdash tt : \perp, \text{bool}} \text{ (BOOLT)} \quad \frac{}{G; \Gamma \vdash ff : \perp, \text{bool}} \text{ (BOOLF)} \\
\frac{G; \Gamma \vdash M : s, \text{bool} \quad G; \Gamma \vdash N_i : s_i, \tau \quad s.r \preceq_G s_0.w \cup s_1.w}{G; \Gamma \vdash (\text{if } M \text{ then } N_0 \text{ else } N_1) : s \curlywedge s_0 \curlywedge s_1 \curlywedge (\perp, \top, s.r), \tau} \text{ (COND)} \\
\frac{G; \Gamma \vdash M : s, \tau \xrightarrow{F} \sigma \quad G; \Gamma \vdash N : s'', \tau \quad s.t \preceq_G s''.w \quad s.r \curlywedge s''.r \preceq_G s'.w}{G; \Gamma \vdash (MN) : s \curlywedge s' \curlywedge s'' \curlywedge (\perp, \top, s.r \curlywedge s''.r), \sigma} \text{ (APP)} \\
\frac{G; \Gamma \vdash M : s, \tau \quad G; \Gamma \vdash N : s', \sigma \quad s.t \preceq_G s'.w}{G; \Gamma \vdash M ; N : s \curlywedge s', \sigma} \text{ (SEQ)} \\
\frac{G; \Gamma \vdash M : s, \theta}{G; \Gamma \vdash (\text{ref}_{\ell, \theta} M) : s, \theta \text{ref}_{\ell}} \text{ (REF)} \quad \frac{G; \Gamma \vdash M : s, \theta \text{ref}_{\ell}}{G; \Gamma \vdash (!M) : s \curlywedge (\ell, \top, \perp), \theta} \text{ (DEREF)} \\
\frac{G; \Gamma \vdash M : s, \theta \text{ref}_{\ell} \quad G; \Gamma \vdash N : s', \theta \quad s.t \preceq_G s'.w, s.r \curlywedge s'.r \preceq_G \ell}{G; \Gamma \vdash (M := N) : s \curlywedge s' \curlywedge (\perp, \ell, \perp), \text{unit}} \text{ (ASSIGN)} \\
\frac{G; \Gamma \vdash M : s, \text{unit}}{G; \Gamma \vdash (\text{thread } M) : (s.r, s.w, \perp), \text{unit}} \text{ (THREAD)} \\
\frac{G; \Gamma, x : \tau \vdash W : s, \tau}{G; \Gamma \vdash \rho x W : s, \tau} \text{ (REC)} \quad \frac{F, G; \Gamma \vdash M : s, \tau \quad s.r \preceq_{GUF} r \quad s.t \preceq_{GUF} t \preceq_G r}{G; \Gamma \vdash (\text{flow } F \text{ in } M) : (r, s.w, t), \tau} \text{ (FLOW)}
\end{array}$$

Figure 3: The Type and Effect System

- G is the current flow policy that is in force when evaluating M ;
- ℓ_0 , also denoted by $s.r$, is the *reading effect*, that is an upper bound (up to the current flow relation) of the security levels of the references the expression M may read. This may be regarded as the security level, or more precisely *the confidentiality level of the expression M* ;
- ℓ_1 , also denoted $s.w$, is the *writing effect*, that is a lower bound (w.r.t. the relation \preceq) of the level of references that the expression M may update;
- ℓ_2 , also denoted $s.t$, is an upper bound (w.r.t. the current flow relation) of some of the levels of dereferencing the expression M may perform. This is used to avoid termination leaks, and therefore we call this the *termination effect* – although the intention is not to guarantee termination.

With respect to the various type systems for information flow, the main novelty here is the G parameter in the typing

context, which is used to relax the constraints on how information may flow in a piece of code to type (such a flow relation in the typing context also appears, under the name of a “hierarchy”, in [44]). The termination effect is similar to the “guard level” of [6] and to the “running time level” of [42]. According to the intuition above, in the type system the reading and termination levels will be composed in a covariant way, whereas the writing level is contravariant, and not concerned with the flow relations between principals. Then we abusively denote by \perp and \top the triples (\perp, \top, \perp) and (\top, \perp, \top) respectively. In the typing rules for compound expressions, we will use the join operation on security effects:

$$s \curlywedge_G s' =_{\text{def}} (s.r \curlywedge_G s'.r, s.w \cup s'.w, s.t \curlywedge_G s'.t)$$

as well as the following convention:

CONVENTION.

In the type system, when the security effects occurring in the context of a judgement $G; \Gamma \vdash M : s, \tau$ involve the join

operation Υ , it is assumed that the join is taken w.r.t. G , i.e. it is Υ_G .

The typing system is given in Figure 3. Notice that this system is syntax-directed: there is exactly one rule per construction of the language. Let us comment on some of the rules, justifying the side conditions that constrain the typing of an expression, as well as the resulting effect of the expression. We see that the reading and writing effects are respectively introduced by the functions for dereferencing and updating the memory – rules (DEREF) and (ASSIGN). We notice that an expression (thread M) has no termination effect, since its evaluation terminates in one step. Indeed, the non-termination of M as a thread cannot influence the computations in the thread that spawned it. The constraints on information flow are implemented in the rules (COND), (APP), (SEQ) and (ASSIGN). In the former, the constraint $s.r \preceq_G s_0.w \cup s_1.w$ means that the branches N_0 and N_1 may only write at a level which is greater, with respect to the current flow relation G , than the reading level of the predicate. This is to prevent indirect flows, like in example (5). A slightly more subtle example is

$$(\text{if } !u_H \text{ then } (\text{thread } v_L := tt) \text{ else } ())$$

which also shows why we record the writing level of the body M of the thread in the effect of (thread M). In the conclusion of (COND), we record the reading level of the predicate as the termination level of the whole expression. This, combined with the condition $s.t \preceq_G s'.w$ in the (SEQ) rule, is to prevent termination leaks as in example (6). This example is essentially the same as

$$((\text{if } !u_H \text{ then } \lambda x x \text{ else loop})(v_L := tt))$$

which is ruled out by the condition $s.t \preceq_G s''.w$ in the (APP) rule. In this rule, the condition $s''.r \preceq_G s'.w$ is to prevent a direct flow, like in

$$(\lambda x(v_L := x)(!u_H))$$

The condition $s.r \preceq_G s'.w$ is meant to exclude expressions that read a secret function that writes in a public location, and unravel this effect by applying it. For instance, it rules out an expression like $((!u_{\ell,\theta})())$ where $\theta = \text{unit} \xrightarrow{(\perp, \ell', \perp)}$ unit and ℓ is not lower than ℓ' . Indeed, the value of the reference u might be $\lambda z(v_{\ell',\theta'} := V)$, with different values for V in different memories (see [52] for a similar example). Example (7) shows why the reading level of the function is recorded in the termination level of the application, and example (8) justifies that we record the reading level of the argument in the termination level of the application. We can use the typing rules for abstraction and application to derive the typing of the let construct, that is

$(\text{let } x = N \text{ in } M) = (\lambda x MN)$, namely:

$$\frac{G; \Gamma \vdash N : s, \tau \quad G; \Gamma, x : \tau \vdash M : s', \sigma \quad s.r \preceq_G s'.w}{G; \Gamma \vdash (\text{let } x = N \text{ in } M) : s \Upsilon s' \Upsilon (\perp, \top, s.r), \sigma}$$

The reader may also check that the typing of $M ; N$ is slightly more liberal than the one of $(\text{let } z = M \text{ in } N)$, where $z \notin \text{fv}(N)$ – and similarly for $(\text{ref}_{\ell,\theta} M)$, $(!M)$ and $(M := N)$ with respect to $(\lambda x(\text{ref}_{\ell,\theta} x)M)$, $(\lambda x(!x)M)$ and $((\lambda x \lambda y(x := y)M)N)$ –, since we do not have to record the reading effect of the component M in the termination effect of $M ; N$. For instance neither $(\text{let } x = (w_H := !u_H) \text{ in } (v_L := tt))$ nor $(\lambda x(w_H := x)(!u_H)) ; (v_L := tt)$ can be typed, whereas the expression of example (9) is accepted. This explains our syntax for the imperative part of the language.

The condition $s'.r \preceq_G \ell$ in the rule (ASSIGN) is to prevent a direct flow, like in example (4). With the condition $s.r \preceq_G \ell$ we rule out the expression $(!u_H) := tt$. Indeed, the value of the reference u might be different locations with level L in different memories. Finally the condition $s.t \preceq_G s'.w$ is to prevent termination leaks, as in

$$(\text{if } !u_H \text{ then } w_H \text{ else loop}) := (v_L := tt)$$

These examples show that all the constraints put on information flow in the typing rules for conditional branching, application, sequential composition and assignment are in fact necessary. Regarding recursion, the reader can check for instance that a derived typing for the while construct is

$$\frac{G; \Gamma \vdash M : s, \text{bool} \quad G; \Gamma \vdash N : s', \tau \quad s.r \Upsilon s'.t \preceq_G s.w \Upsilon s'.w}{G; \Gamma \vdash (\text{while } M \text{ do } N) : s \Upsilon s' \Upsilon (\perp, \top, s.r), \text{unit}}$$

As in [6, 42], we record the confidentiality level of the boolean guard expression in the termination level of the while construct.

To type a flow declaration (flow F in M), we have to type M in the context of the current flow policy extended with F . In the (FLOW) rule, we use a kind of subsumption for the security effect. Namely, the apparent reading and termination effects of the expression (flow F in M) are allowed to be higher, with respect to F , than the ones of M . For instance, one can check that the following is a valid proof of typing (leaving out the type annotation of references):

$$\frac{\frac{H \prec L, L \prec H; \Gamma \vdash u_H : \perp, \tau}{H \prec L, L \prec H; \Gamma \vdash (!u_H) : (H, \top, \perp), \tau}}{L \prec H; \Gamma \vdash (\text{flow } H \prec L \text{ in } (!u_H)) : (L, \top, \perp), \tau} \quad H \prec L$$

and therefore one can see that the type system accepts the expression of example (10). Another example is

$$v_L := (\text{flow } H \prec L \text{ in } \text{encrypt}(!u_H, K))$$

where encrypt is a given encryption function, and K is the encryption key. Subsumption in the (FLOW) rule will be used in the proof of the Subject Reduction property⁽¹⁾. However, one should notice that in the typing rule for flow declaration we do not allow subsumption for the writing level. Example (11) shows why it would be wrong to do so.

The way we build the termination effect of an expression allows us to accept the expression of example (9). This would not be the case if we had approximated $s.t$ as $s.r$ by dealing with security effects of the form (r, w) . The classical approach to “weak” non-interference, based on a big-step semantics, actually corresponds to a variant of our type system where $s.t = \perp$, but obviously this is too weak to ensure the non-disclosure policy. We could think of further refining our type system by building the termination level in a more clever way, since clearly secure programs such as

$$(\text{if } !u_H \text{ then } v_H := tt \text{ else } v_H := ff) ; w_L := tt$$

are still rejected by our type system. For instance, we would like to say that the termination level of $(\text{if } !u_H \text{ then } M \text{ else } N)$ may be taken as \perp if we know that both M and N terminate. However, little is known on how to ensure termination in a higher-order imperative language. Indeed, it is well-known since Landin’s work [20] that circular higher-order references introduce non-termination, like for instance in (using ref without subscripts)

$$(\text{let } x = (\text{ref } \lambda y y) \text{ in } x := \lambda y ((!x)y) ; ((!x)V))$$

which has the type of V . Nevertheless, if we know for a fact that both M and N do not cause any trouble (see [1, 37, 48] for some ways of ensuring this in a simple language), we may compensate for the inflexibility of the type system as regards $(\text{if } !u_H \text{ then } M \text{ else } N)$ for instance by writing instead:

$$(\text{flow } H \prec \perp \text{ in } (\text{if } !u_H \text{ then } M \text{ else } N))$$

Our type system rejects in the same way expressions that are regarded as involving a *timing leak* (see [18, 38, 43]), like

$$(\text{if } !u_H \text{ then } M \text{ else } ()) ; v_L := ff \quad (12)$$

where M is an expression that takes some time to compute, like $() ; \dots ; ()$, although program (12) is generally secure in the sense of Definition 3.3. Such a program should be regarded as unsecure if some specific scheduling discipline were to be taken into account (cf. [6, 37, 43]). We leave for further investigations the question as to whether our type

⁽¹⁾ The proof we have for our Soundness Theorem uses all the features of the type system, apart from the fact that the reading level of the body M of a thread is recorded in the effect of (thread M).

system is also adequate to deal with scheduling disciplines (we plan to do this for the cooperative concurrency of ULM [5]), but we observe that it does not restrict the confidentiality level of the predicate in a conditional branching to be \perp , as suggested in [43, 47]. In this respect, our system is close to the ones of [6, 42].

As a last example, let us examine a program showing that the flow policy under which the value of a reference will be put cannot be predicted statically. Let M be the following expression:

$$\begin{aligned} \text{let } f = \lambda x \lambda y \text{ if } x \text{ then } (\text{flow } p \prec q \text{ in } v_{q,\theta} := !y) \\ \text{else } (\text{flow } p \prec r \text{ in } w_{r,\theta} := !y) \\ \text{in } ((fN)u_{p,\theta}) \end{aligned}$$

Then one can see that the following typing is admissible:

$$\frac{G; \Gamma \vdash N : s, \text{bool} \quad s.r \preceq_G \{q, r\}}{G; \Gamma \vdash M : s \curlywedge (p, \{q, r\}, s.r), \text{unit}}$$

As we explained above, the constraint $s.r \preceq_G \{q, r\}$ is to prevent termination leaks, since the evaluation of N could terminate or not, depending on the values read in the memory. One should notice that there is no constraint relating p , the confidentiality level of the reference u , with q or r . This means that the value of this reference can be downgraded to either level q or r , depending on the value computed for the boolean N .

As previously announced, the main technical result of our paper is a type soundness property. To prove it, we rely, as usual, on a Subject Reduction property, which states that the type of an expression is preserved by reduction. Regarding the effects, some may be performed, by reading or updating a reference, and some may be discarded, when a branch in a conditional expression is taken. Then the effects of an expression “decrease” along the computations, and, in particular, its confidentiality level becomes less critical. The Subject Reduction property is formulated using the reduction of the “main thread” in an expression:

PROPOSITION (SUBJECT REDUCTION) 4.1.

If $G; \Gamma \vdash M : s, \tau$ and $(M, \mu) \xrightarrow[F]{N} (M', \mu')$ with $u_{\ell, \theta} \in \text{dom}(\mu) \Rightarrow G; \Gamma \vdash \mu(u_{\ell, \theta}) : \perp, \theta$ then $G; \Gamma \vdash M' : s', \tau$ and $G; \Gamma \vdash N : s'', \text{unit}$ for some s' and s'' such that $s'.r \curlywedge_G s''.r \preceq_G s.r$, $s.w \preceq s'.w \cup s''.w$ and $s'.t \preceq_G s.t$.

Finally our main result is that the type and effect system provides a way to guarantee that confidentiality is preserved by programs:

THEOREM (SOUNDNESS).

If M is typable in the context of a flow policy G , that is if for some Γ , s and τ we have $G; \Gamma \vdash M : s, \tau$, then M satisfies the non-disclosure policy with respect to G , that is $M \in \mathcal{ND}(G)$.

To prove this result, for any security level ℓ we exhibit a (G, ℓ) -bisimulation that contains the pair (M, M) for any G -typable expression M . Such a relation is built by examining the possible cases for pairs (P, Q) such that $(M, \mu) \rightarrow (P, \mu')$ and $(M, \nu) \rightarrow (Q, \nu')$, where μ and ν satisfy the condition of Definition 3.1 of bisimulations. Lemma 2.3 shows that, starting from a given expression M in the context of two different memories, the evaluation process may only branch (to P and Q) when the expression comes to read a reference. This is the basis for building our bisimulations. The details are given in the full version of the paper.

To conclude this section, we notice that we could add to the language a *reading clearance* construct $(M : \ell)$, with the intention that it is accepted only if the expression M reads references with confidentiality level below ℓ , with respect to the current flow policy. The typing of such a construct is obvious:

$$\frac{G; \Gamma \vdash M : s, \tau \quad s.r \preceq_G \ell}{G; \Gamma \vdash (M : \ell) : s, \tau}$$

This is sound because, if $G; \Gamma \vdash M : s, \tau$, then $s.r$ is an upper bound of the confidentiality level of the references that the expression M may read (thanks to the (THREAD) rule that records the reading level of the threads M may spawn).

5. Conclusion and related work

We have proposed a way to face the “*challenge [of] determining what the nature of a downgrading mechanism should be and what kinds of security guarantees it permits*” [53]. Taking the view that one should distinguish the questions of *what* or *how much* can be revealed, from that of *how* it can be revealed, we addressed the second question by proposing a simple and powerful construct for declassification based on dynamically varying flow policies. Although this idea has already been mentioned in the literature (see [44]), it does not seem to have been previously studied in a formal way (in [45] it is shown that if the downgrading relations used in a program do not modify the security lattice under some level ℓ , then the program is secure up to this level). Our main achievement is the design of a security property that is a natural generalization of classical non-interference, based on dynamic flow policies. We notice that the idea of stating the security property in a way that reflects the *local* nature of declassification, that is, using a decorated small-step semantics, could perhaps be used in other settings. Moreover, we have shown that, for programs written in an expressive, higher-order imperative core language, this property can be enforced by static analysis. The idea of using a construct for dynamically introducing flow policies can certainly be applied to various other programming paradigms.

The language-based security approach is now well established – though not widely used –, and there is no question that lattices of information flow provide a good basis for ensuring end-to-end confidentiality properties. Indeed, checking that a program does not violate a given flow policy by means of a type system which enforces the extensional property of non-interference, is regarded as providing a reasonable security guarantee. Therefore we believe that checking, piece by piece, that a program does not violate *local* flow policies should provide a similar guarantee, while allowing us to deal with declassification. Borrowing an example from [52], we may then write a piece of code to release precious information held in a safe place A by *Alice*, to *Bob* who wishes to purchase it, provided a payment has indeed been done:

if *paid* then (flow *Alice* \prec *Bob* in $B := !A$) else \dots

without checking further that this intended leak of information is justified. In some other cases, where a program has to be certified against high security standards for instance, one would have to provide a formal proof that a declassified portion of code satisfies some specification, like that of not releasing too much information. This justification is left to the programmer, whereas our programming language design provides him with a flexible programming construct for declassification, together with a static checking technique to prevent some errors.

In this respect, our approach contrasts with most previous works on declassification in a language-based security setting that aimed at imposing constraints, at a linguistic level, on this operation – sometimes without justifying such constraints, for lack of an extensional notion of security. For instance, in [46, 49], Volpano and Smith restrict downgrading to occur by means of specific “hard” functions. This is certainly relevant for some applications, especially those involving cryptography, but is less appropriate for applications where the programmer intends to let information leak in some places (like in the example above). Another example of constrained downgrading is *robust declassification* which was proposed, and then studied in a series of papers [29, 30, 31, 32, 44, 52] by Myers and colleagues. The idea of robust declassification is to allow this operation only for extending the reading clearances assigned by the owner of an object, and to control it by requiring that this operation runs under appropriate authority. This was first conceived as a run-time constraint, and was later approximated in a type system by means of integrity levels. Compared to our approach, robust declassification is obviously more restrictive (for instance one can only set up flows from lower to higher levels). However, it would be interesting to see whether we can accommodate our setting to deal with it (even though it is not very clear that the “robustness property” of [32] is related to our non-disclosure property). An

obvious idea would be to restrict the use of (flow F in M) to the case where F only allows to declassify the contents of references that have been created by the thread executing this piece of code, and then to see how our security property can be made more accurate for this case.

In another paper [39], Sabelfeld and Myers introduce a different way of restricting declassification, with the idea that downgrading is acceptable provided that the program does not modify data if that could influence the value of declassified expressions, therefore addressing the question of *what* is downgraded. For instance, the program $u_H := \text{ff} ; v_L := \text{declassify}(u_H, L)$ is not regarded as safe according to the definition of *delimited release*. On the other hand, a program like $v_L := \text{declassify}(u_H, L) ; w_L := u_H$, which is similar to the one of example (3), is considered safe, but is ruled out by the type system. The type system, based on the idea that “*variables under declassification may not be updated prior to declassification*”, might be difficult to extend to a more sophisticated language, with a less predictable order of execution than the one considered in [39].

More recently, Chong and Myers introduced *declassification policies* [8], that specify the levels through which a value can be downgraded. This also involves conditions, which are supposed to be satisfied in order to perform the declassification steps. These are used in the definition of a generalized noninterference property to mark the steps where declassification occurs. This bears some resemblance to our transitions labelled by a local flow relation, although conditions are rather used to single out sequences of steps that do not involve downgrading operations. The declassification policies of [8] look a bit inflexible since, as far as we can see, there is no possibility for a value to be used in another way than the one prescribed by the specific policy assigned to it. Therefore it seems that, with these policies, the programmer must accurately anticipate the run-time behaviour of the declassified values. By contrast, in our setting a reference can be involved in various declassification scenarios, and this does not have to be reflected in its type.

Closer to ours is the work by Ferrari & al. [13], who proposed to attach “*waivers*” to methods in an object-oriented language to provide a way of making information flow from objects to users. Although the authors claim that “*only privileged methods*” have associated waivers, there seems to be actually no constraint on the flow of information they allow. This idea of a waiver is therefore similar to a local flow relation, though it is not clear whether the notion of “*safe information flow*” that the authors define is similar to our non-disclosure property (as far as we can see, this definition does not treat waivers as having a local scope). A work that is also close to ours, at least as regards the motivations, is the one by Li and Zdancewic [24]. After having made the initial decision that “*instead of studying who can downgrade the data* [like in the work on robust declassification],

we take an orthogonal direction and study how data can be downgraded”, they intend to offer the programmer a way of specifying sophisticated *downgrading policies*. Therefore we can say we share the same motivations. However, the ways we take from this starting point differ considerably. Li and Zdancewic introduce a very sophisticated notion of a downgrading policy (an expression in a typed λ -calculus), where we use flow relations between principals, which look easy to use in practice. Our non-disclosure property also looks simpler than the notion of *relaxed noninterference*, which is based on program equivalence (in the language of downgrading policies). Their main result is again very close in its spirit to ours, since “*the security guarantee [provided by relaxed noninterference] only assures that the program respects the user’s security policies*”. Therefore it would be interesting to compare in greater details the two approaches, especially from the point of view of expressiveness.

Finally, the work that is the closest to ours is the one by Mantel and Sands [27]. In addition to a given lattice structure of security levels, they consider an extra relation on these levels, that can be used in specific instructions – namely, assignments of the form $v_{\ell'} := (!u_{\ell})$ – to downgrade information: in such an instruction, the flow from ℓ to ℓ' should be allowed by the “exceptional” flow relation. In our syntax, we would write this as (flow $\ell \prec \ell'$ in $v_{\ell'} := (!u_{\ell})$) (where $\ell \prec \ell'$ means $\{p \prec q \mid p \in \ell \ \& \ q \in \ell'\}$). Then Mantel and Sands introduce a security property generalizing classical non-interference, defined by means of a notion of bisimulation with respect to transitions annotated by a flow relation, and they show a type soundness result. Therefore one can see that this is very close to what we did in this paper (the two works were done independently, and a precise comparison of our security properties remains to be made). There are some differences, however. A first difference is that Mantel and Sands choose to restrict declassification to very specific instructions, whereas we allow any computation to be declassified. From a pragmatic point of view, the main difference we see is that in their work, declassification is governed by a specific global flow policy – the “exceptional” flow relation – that cannot be manipulated by programs. We think that it could be useful in practice to have the ability of choosing various ways of downgrading, depending on the point in the program where this is performed, without necessarily complying with a predetermined, global downgrading policy. Moreover, such a dynamic view seems to be needed in order to deal with mobile code, where agents migrate with their own flow policy. Another noticeable – though not related to declassification – difference is that we are using a higher-order imperative language, whereas Mantel and Sands consider a simple while language with threads, where there is no interaction between commands and expressions. Moreover, our type

system appears to be less restrictive (as regards the while construct for instance).

Some obvious topics for further investigations are polymorphism and type inference [29, 33], dynamic labels [44, 55], and more generally first-class security levels. One could also wish to deal with a richer set of effects, including for instance the creation and deletion of references, the creation of threads, and more generally any action that modifies the context of an expression in the (abstract) machine evaluating it. We are currently working on using the idea of local flow policies in a mobile code setting, and more precisely in the ULM language [5]. Indeed, a mobile agent may carry its own flow policy, and run in various sites, each having their own, local flow policies, and therefore this is a scenario where one has to deal with various flow relations. Regarding declassification, one may think our approach is too permissive, since it allows any program to declassify anything, provided that no other flow than the declared ones is implemented. Therefore it would be interesting to see how we could restrict the usage of the flow declaration construct in some sensible ways, and adapt the non-disclosure policy accordingly. We have mentioned a possible way of doing this in discussing the work on robust declassification. It would also be interesting to find a simple notion of “security error”, that could be used as a basis for designing error messages in a type inference approach. Finally, we observe that, following Biba’s remark that integrity is dual to confidentiality in some sense (see [23, 31]), we may design a framework for the integrity aspect of security in a similar way to what we did for confidentiality. It could support, in particular, downgrading facilities like the “endorse” construct of [23] (which is also considered in [32], but with a different semantics). Although the expectations are even stronger regarding integrity than confidentiality (see [23] for instance), it would be interesting, from a practical point of view (*cf.* [53]), to have in a programming language such flexible downgrading facilities with respect to integrity.

References

- [1] J. AGAT, *Transforming out timing leaks*, POPL’00 (2000) 40-53.
- [2] G.R. ANDREWS, R.P. REITMAN, *An axiomatic approach to information flow in programs*, ACM TOPLAS, Vol. 2 No. 1 (1980) 56-76.
- [3] D.E. BELL, L.J. LA PADULA, *Secure computer system: unified exposition and Multics interpretation*, Mitre Corp. Rep. MTR-2997 Rev. 1 (1976).
- [4] A. BOSSI, C. PIAZZA, S. ROSSI, *Modelling downgrading in information flow security*, CSFW’04 (2004).
- [5] G. BOUDOL, ULM, *a core programming model for global computing*, ESOP’04, Lecture Notes in Comput. Sci. 2986 (2004) 234-248.
- [6] G. BOUDOL, I. CASTELLANI, *Non-interference for concurrent programs and thread systems*, Theoretical Comput. Sci. Vol. 281, No. 1 (2002) 109-130.
- [7] D. CLARK, S. HUNT, P. MALACARIA, *Quantified interference: information theory and information flow*, WITS’04 (2004).
- [8] S. CHONG, A.C. MYERS, *Security policies for downgrading*, 11th ACM Conf. on Computer and Communications Security (2004).
- [9] E. COHEN, *Information transmission in computational systems*, 6th ACM Symp. on Operating Systems Principles (1977) 133-139.
- [10] K. CRARY, A. KLIGER, F. PFENNING, *A monadic analysis of information flow security with mutable state*, J. of Functional Programming, Vol. 15 No. 2 (2005) 249-291.
- [11] D.E. DENNING, *A lattice model of secure information flow*, CACM Vol. 19 No. 5 (1976) 236-243.
- [12] A. DI PIERRO, C. HANKIN, H. WIKLICKY, *Approximate non-interference*, CSFW’02 (2002) 1-15.
- [13] E. FERRARI, P. SAMARATI, E. BERTINO, S. JAJODIA, *Providing flexibility in information flow control for object-oriented systems*, IEEE Symp. on Security and Privacy (1997) 130-140.
- [14] R. FOCARDI, R. GORRIERI, *A classification of security properties for process algebras*, J. of Computer Security, Vol. 3 No. 1 (1995) 5-33.
- [15] R. FOCARDI, S. ROSSI, *Information flow security in dynamic contexts*, CSFW’01 (2001) 307-319.
- [16] J.A. GOGUEN, J. MESEGUER, *Security policies and security models*, IEEE Symp. on Security and Privacy (1982) 11-20.
- [17] N. HEINTZE, J. RIECKE, *The SLam calculus: programming with secrecy and integrity*, POPL’98 (1998) 365-377.
- [18] A.K. JONES, R.J. LIPTON, *The enforcement of security policies for computation*, 5th ACM Symp. on Operating Systems Principles (1975) 197-206.
- [19] B.W. LAMPSON, *A note on the confinement problem*, CACM Vol. 16 No. 10 (1973) 613-615.
- [20] P.J. LANDIN, *The mechanical evaluation of expressions*, Computer Journal Vol. 6 (1964) 308-320.
- [21] P. LAUD, *Semantics and program analysis of computationally secure information flow*, ESOP’01, Lecture Notes in Comput. Sci. 2028 (2001) 77-91.
- [22] P. LAUD, *Handling encryption in an analysis for secure information flow*, ESOP’03, Lecture Notes in Comput. Sci. 2618 (2003) 159-173.
- [23] P. LI, Y. MAO, S. ZDANCEWIC, *Information integrity policies*, Formal Aspects of Security and Trust Workshop (2003).
- [24] P. LI, S. ZDANCEWIC, *Downgrading policies and relaxed noninterference*, POPL’05 (2005) 158-170.
- [25] G. LOWE, *Semantic models of information flow*, Theoretical Comput. Sci. 315 (2004) 209-256.
- [26] J.M. LUCASSEN, D.K. GIFFORD, *Polymorphic effect systems*, POPL’88 (1988) 47-57.

- [27] H. MANTEL, D. SANDS, *Controlled declassification based on intransitive noninterference*, APLAS'04, Lecture Notes in Comput. Sci. 3302 (2004) 129-145.
- [28] R. MILNER, M. TOFTE, R. HARPER, D. MACQUEEN, *The definition of Standard ML (Revised)*, The MIT Press (1997).
- [29] A. MYERS, *JFlow: practical mostly-static information flow control*, POPL'99 (1999).
- [30] A. C. MYERS, B. LISKOV, *A decentralized model for information flow control*, ACM Symp. on Operating Systems Principles (1997) 129-142.
- [31] A. C. MYERS, B. LISKOV, *Protecting privacy using the decentralized label model*, ACM Trans. on Soft. Eng. and Methodology, Vol. 9 No. 4 (2000) 410-442.
- [32] A. C. MYERS, A. SABELFELD, S. ZDANCEWIC, *Enforcing robust declassification*, CSFW'04 (2004).
- [33] F. POTTIER, V. SIMONET, *Information flow inference for ML*, ACM TOPLAS Vol. 25 No. 1 (2003) 117-158.
- [34] A. W. ROSCOE, M. H. GOLDSMITH, *What is intransitive noninterference?*, CSFW'99 (1999).
- [35] J. RUSHBY, *Noninterference, transitivity, and channel-control security policies*, Comput. Sci. Lab. SRI International, Tech. Rep. CSL-92-02 (1992).
- [36] P. RYAN, J. MCLEAN, J. MILLEN, V. GLIGOR, *Noninterference, who needs it?*, CSFW'01 (2001).
- [37] A. SABELFELD, D. SANDS, *Probabilistic noninterference for multi-threaded programs*, CSFW'00 (2000).
- [38] A. SABELFELD, A. C. MYERS, *Language-based information-flow security*, IEEE J. on Selected Areas in Communications Vol. 21 No. 1 (2003) 5-19.
- [39] A. SABELFELD, A. C. MYERS, *A model for delimited information release*, Intern. Symp. on Software Security, Lecture Notes in Comput. Sci. to appear (2003).
- [40] R. S. SANDHU, *Lattice-based access control models*, IEEE Computer Vol. 26 No. 11 (1993) 9-19.
- [41] V. SIMONET, *The Flow Caml system: documentation and user's manual*, INRIA Tech. Rep. 0282 (2003).
- [42] G. SMITH, *A new type system for secure information flow*, CSFW'01 (2001).
- [43] G. SMITH, D. VOLPANO, *Secure information flow in a multi-threaded imperative language*, POPL'98 (1998).
- [44] S. TSE, S. ZDANCEWIC, *Run-time principals in information-flow type systems*, IEEE Symp. on Security and Privacy (2004).
- [45] S. TSE, S. ZDANCEWIC, *A design for a security-typed language with certificate-based declassification*, ESOP'05, Lecture Notes in Comput. Sci. to appear (2005).
- [46] D. VOLPANO, *Secure introduction of one-way functions*, CSFW'00 (2000) 246-254.
- [47] D. VOLPANO, G. SMITH, *Eliminating covert flows with minimum typings*, CSFW'97 (1997) 156-168.
- [48] D. VOLPANO, G. SMITH, *Probabilistic noninterference in a concurrent language*, CSFW'98 (1998) 34-43.
- [49] D. VOLPANO, G. SMITH, *Verifying secrets and relative secrecy*, POPL'00 (2000) 268-276.
- [50] D. VOLPANO, G. SMITH, C. IRVINE, *A sound type system for secure flow analysis*, J. of Computer Security, Vol. 4, No 3 (1996) 167-187.
- [51] A. WRIGHT, M. FELLEISEN, *A syntactic approach to type soundness*, Information and Computation Vol. 115 No. 1 (1994) 38-94.
- [52] S. ZDANCEWIC, *A type system for robust declassification*, MFPS'03, ENTCS Vol. 83 (2003).
- [53] S. ZDANCEWIC, *Challenges for information-flow security*, PLID'04 (2004).
- [54] S. ZDANCEWIC, A. C. MYERS, *Secure information flow via linear continuations*, HOSC Vol. 15 No. 2-3 (2002) 209-234.
- [55] L. ZHENG, A. C. MYERS, *Dynamic security labels and noninterference*, Formal Aspects of Security and Trust Workshop (2004).