



Compile-Time Security Certification of Imperative Programming Languages

Sandip Ghosal¹(✉), R. K. Shyamasundar¹, and N. V. Narendra Kumar²

¹ Department of Computer Science and Engineering,
Indian Institute of Technology Bombay, Mumbai 400076, India
sandipsmit@gmail.com

² Institute for Development and Research in Banking Technology, Hyderabad, India

Abstract. With the ever increase in the demand of building secure systems, recent years are witnessing a plethora of research on information flow control (IFC) techniques in programming languages to enforce a finer-grained restriction on the propagation of information among untrusted objects. In this paper, we introduce a *dynamic labelling* (DL) algorithm (This paper is an extended version of the article [1] presented in SECUREPT'18.) for security certification of imperative programming languages that follows a combination of mutable and immutable labelling referred to as *hybrid* labelling approach. First, we study the possible methods of binding security labels with the subjects and objects of the program which include *program counter* that represent implicit flow within a program and compare the precision achieved by the applications of methods on benchmark programs. Next, we describe our labelling algorithm that generates labels for intermediate subjects/objects of a program from the given set of initial labels (some of which could be immutable throughout the computation) adhering to the constraints defined in [2] for a program to be *information-flow secure*. Apart from the usual control statements found in the imperative languages, we also present the labelling approach for a procedure call highlighting subtleties of different parameter passing mechanisms adopted in modern languages. Further, we discuss a variant of the algorithm for concurrent programs. It is shown that our algorithm always terminates after a finite number of iterations, also establish the soundness concerning non-interference as given by [3]. We compare the labelling precision realizable by our approach with the existing approaches in the literature.

1 Introduction

The seminal work of Denning [4] on security certification of programs built on *information-flow security* led to a firm foundation for language-based security. The extension of such a theory through the proposal of the Decentralized Label Model (DLM) [5] provided a momentum for language-based security. Since then there has been an enormous amount of literature on language-based security [6–9]. Various well articulate assessments of the status of language-based security

have been discussed in [10–12]. With the need for security everywhere including IoT, language-based security is becoming prominent as it deals with security at various levels, and also brings out various points of leaks and attacks.

Non-interference was developed after Denning’s security certification as a more semantic characterization of security [13], followed by many extensions. Informally the non-interference property says the impact on the program due to changes in *high* inputs should not be observable by the *low* outputs. [3] have used a purely value-based interpretation of non-interference as a semantic characterization of information-flow, and derived sound typing rules to capture Denning’s certification semantics effectively. Volpano *et al.*’s notion of non-interference, and its extensions have become the de-facto standard for the semantics of information-flow in the literature on language-based security. Usually, the objective of IFC is to enforce non-interference property to ensure end-to-end flow security.

A broad spectrum of information-flow security mechanisms varies from fully dynamic ones, e.g., in the form of execution-monitors [14, 15] to static ones, e.g., in the form of type systems [3]. While one would prefer static labels as that leads to certification of programs at compile-time, it has the problem of classifying programs that would not leak any information under the information-flow policy at execution time due to the underlying inputs that arrive at run-time. Real-world web applications often require to interact with the external environment that cannot be predicted during compile-time which motivates researchers to enforce security at run-time. For instance, security settings of files and database records are updated frequently, and these changes might affect the information flow control which cannot be handled by static mechanisms. Dynamic labels are essential to capture the changes in security label and accordingly labels are changed at run-time. However, unlike static or immutable labelling implementation, a dynamic mechanism has a cost that user has to pay in the form of significant run-time overhead, and also *implicit flows* introduced due to uncovered flow paths not considered at run-time. Hence, an ideal label-checking mechanism should have a *hybrid labelling* that would have an excellent trade-off for mutable and immutable labels to realize acceptable precision and performance.

In this paper, we first discuss various labelling approaches that use a combination of attributes like mutable, immutable, static, compile-time, run-time, etc., for the security certification along with the corresponding realizable precision of security. Having assessed the shortcomings [1], we propose a new hybrid (mutable and immutable) labelling approach for certifying programs for information-flow security using the standard certification of constraints as elaborated by Denning. Our dynamic labelling algorithm is sound with respect to non-interference, and we further prove the termination of the labelling algorithm. Our proposed labelling algorithm leads to certification that is more security precise than other labelling approaches in the literature [7, 8, 16–20]. It may be pointed out that the labels are generated succinctly without unnecessarily blowing up the label space. As the method is not tied to any particular security model, it provides a sound basis for the security certification of programs for information-flow security. We

further compare the precision realizable by our approach with those in the literature. The comparison of our approach also brings to light, an intrinsic property of our labelling algorithm that could be effectively used for non-deterministic or concurrent programs illustrated in Sect. 4.

Structure of the Paper. Section 2 presents different labelling schemes, and assess the limitations of them. Section 3 describes the proposed dynamic labelling algorithm along with illustrative examples, and proofs of characteristic properties as well as soundness with respect to *non-interference*. Section 4 presents an extended version of the labelling algorithm for concurrent programs. Section 5 provides a comparison with earlier approaches. Finally, Sect. 6 summarizes the contributions along with the ongoing work.

2 Certification of Programs

According to Denning’s Information Flow Model (DFM) the necessary and sufficient condition for the flow security of a program P is: if there is an information flow from object x to object y , denoted by $x \rightarrow y$, the flow is secured by P only if $\lambda(x) \leq \lambda(y)$. ‘ λ ’ is a labelling function, responsible for binding subject/object of the program to a security class (either statically or dynamically depending on the application) from the lattice of security classes as described in Denning’s lattice model [4]. ‘ \leq ’ is a binary relation on security classes that specifies permissible information flows. ‘ \oplus ’ is a binary class-combining operator evaluates least upper bound (LUB) of two security classes in the lattice. The above condition is usually referred to as the *Information-flow policy* (IFP). A program is certified for IFP if there are no violations of the policy during program execution. The *Information Flow Secure* policy forms a basis for certifying programs for security. The crux of certification lies in assuring that all the information flows over legitimate channels or storage channels follow the specified flow-policy. The outcome of approaches could be measured in terms of *precision* defined below.

Let us suppose that F is the set of possible flows in an information flow system, and let A be the subset of F authorized by a given flow policy, and let E be the subset of F “executable” given the flow control mechanisms in operation. The system is said to be **secure** if $E \subseteq A$; that is all executable flows are authorized. A secure system is **precise** if $E = A$.

The method of binding security classes/labels to objects play an essential role in the analysis of programs. Each of the static certification and runtime enforcement algorithms proposed in the literature chooses an object labelling method, some of which may also use the label of the implicit variable, usually referred to as *Program Counter* (PC) which may be reset after every statement or keeps updating monotonically. First, let us consider the following three broad object labelling schemes:

Scheme 1: fixed labels for all the variables,

Scheme 2: labels of all the variables can be modified; for example, based on the information contained in them at any given program point, and

Table 1. There is implicit flow from x to y while there is no direct flow (Cf. [1]).

```

1 void test(int x){
2   int y=0;
3   int z=0;
4   if(x==0)
5     z=1;
6   if(z==0)
7     y=1;
8 }

```

Table 2. Need to label local variables dynamically (Cf. [1]).

```

1 void test(){
2   int a;
3   a=x;
4   y=a;
5   a=z;
6 }

```

Scheme 3: labels of some variables are fixed while the labels of the other variables could be modified.

In this section, we argue that from the perspective of capturing the notion of security,

(i) Scheme 1 is inappropriate as it is too stringent, and results in secure programs being incorrectly rejected as insecure,

(ii) Scheme 2 is also inappropriate as it allows all programs as secure programs, and

(iii) Scheme 3 is the ideal candidate, if the variables whose labels can be allowed to be modified are carefully chosen.

A majority of literature on language-based security follows Scheme 1, as it has advantages in certifying tricky programs such as the one given in Table 1.

As the program in Table 1 executes, the following information flow is observed: if the value of x is 0, the value of z becomes 1, and the second *if* block will not execute, thus the value of y remains 0. On the other hand, if the value of x is 1, the second *if* block executes, and y is initialized to 1. In either case, the value of y is the same as the value of x although there is no such explicit assignment, e.g., $x = y$. If we consider the security labels of x and y are \underline{x} and \underline{y} and $\underline{x} \not\leq \underline{y}$ then this is an example of the insecure program as there is an implicit flow from x to y even though they belong to different security classes where an explicit flow is not allowed. Table 3 analyzes two different labelling approaches, i.e., Scheme 1 and Scheme 2 in respect of the program in Table 1.

However, purely static labelling is too restrictive and rejects secure programs as insecure. This is illustrated by considering the program fragment shown in Table 2, where x, y , and z are global variables, and a is a local variable (we do not consider pointer variables).

If the program in Table 2 is analyzed under Scheme 1, it generates the following set of flow-constraints to be satisfied for the program to be secure: $\underline{x} \leq \underline{a}$, $\underline{a} \leq \underline{y}$, and $\underline{z} \leq \underline{a}$. These constraints will be satisfied only if $\underline{z} \leq \underline{y}$, which implies that there is an information-flow from z to y . However, from an intuitive perspective, the program never causes an information flow from z to y and must be

Table 3. Analyzing information flow at each line from example in Table 1 according to static and dynamic labelling (Cf. [1]).

Line No.	Static labelling (Scheme 1)	Dynamic labelling (Scheme 2)
3	Label of z would be inferred in such a way that all the flows to and from z should be secured. If z is assigned to \underline{x} , the flow from $z \rightarrow y$ is not permitted. There is no way to label z so that all the flows are safe. For this reason the program is insecure	Label of z is initialized to least confidential security class e.g. public (\perp)
4		As there is a flow from $x \rightarrow z$ z is updated to \underline{x} such that $\underline{x} \leq \underline{z}$
5		
6		There is a flow from $z \rightarrow y$. But \underline{x} cannot flow into \underline{y} . Hence the program is insecure
7		

considered secure if $\underline{x} \leq \underline{y}$. Table 4 analyzes two different labelling approaches, i.e., Scheme 1 and Scheme 2 in respect of the program in Table 2.

Table 4. Analyzing information flow at each line of example shown in Table 2 according to static and dynamic labelling (Cf. [1]).

Line No	Flow direction	Static labelling (Scheme 1)	Dynamic labelling (Scheme 2)
5		Label of a is automatically inferred in such a way that all the flows to and from a is secure. If a is labelled as $\underline{x} \oplus \underline{z}$ due to explicit flows from x and y to a , the constraint $\underline{x} \oplus \underline{z} \leq \underline{y}$ will not be satisfied because $\underline{z} \not\leq \underline{y}$. As static labelling fails to label the local variable a , the program is insecure	Label of a is initialized to least confidential security class e.g. public \perp
6	$x \rightarrow a$		Label of a is updated to label of x i.e. \underline{x} so that x can flow to a
7	$a \rightarrow y$		Flow is allowed as the constraint $\underline{x} \leq \underline{y}$ is satisfied.
8	$z \rightarrow a$		Label of a is updated to $\underline{x} \oplus \underline{z}$ so that flow is allowed as $\underline{z} \leq \underline{x} \oplus \underline{z}$. Hence the program is flow-safe

From the above examples, it follows that the use of purely static labelling is too conservative, and misses several secure programs.

2.1 Refinement via PC Labels

Information flow is quite tricky to capture and can happen even if a statement does not get executed. Such flows are called ‘‘implicit’’ [21], and are possible due to conditional and iteration statements. To keep track of such impact, the notion of the program counter (pc) label is introduced that denotes the sensitivity of the current context. Traditionally, once the control exits the conditional or iteration statements, the pc label is reset to its previous value, thus denoting that the variables in the condition expression no longer impact the current context. Subsequently, a sequential composition $S_1; S_2$ is deemed secure if both S_1 and S_2 are individually secure.

Examples `copy3` and `copy4` in [21] have highlighted that certain subtle flows cannot be captured unless the `pc` label is updated monotonically, and tracks the influence of all the information the program has accessed. It was noted that this might lead to a phenomenon called “label creep” [11] wherein the `pc` label rises too high resulting in rejecting any further flows. To avoid label creeping, the current literature on language-based security takes the route of resetting `pc` label after exiting from a control structure.

Tracking PC Labels [18]. The method described for Haskell envisaged in [18] uses a label for current control without reinitializing every time the control exits a statement. A labeled IO Haskell library unit `LIO` is built to track a single mutable *current label* (similar to *pc*) at run-time and allows access to IO operations. The unit is responsible for ensuring that the current label keeps track of all the observed data and regulates label modifications. At each computation, `LIO` keeps tracks of the *current label* and allows access to IO functionality, e.g., labeled file systems. The current label is evaluated as an *upper bound* of all the labels observed during program execution.

2.2 Labelling Schemes: A Summary

Table 5 summarizes possible ways of binding labels with objects.

Table 5. Binding labels with objects for certification (Cf. [1]).

pc Label→	Reset	Monotonic
Labelling↓		
Static	P_1	P_2
Hybrid	P_3	P_4

Some programs where ignoring the label of program point leads to incorrect certification are given in Tables 6 and 7. Table 6 shows an example with information leaks due to abnormal termination of a program. The value of x can be calculated from the value of *sum* (maximum possible integer value) and y on terminating the program due to integer overflow. There is an implicit flow from x to y although the assignment to y is conditioned on the *sum*.

A non-terminating flow insecure program is shown in Table 7. Let us consider the given label of the global variables x and y are given as \underline{x} and \underline{y} such that $\underline{x} \not\leq \underline{y}$. It can be observed that the variable y holds the value equal to x depending on the termination of the program.

Although `LIO` follows the label binding P_2 that keeps track of program point, the run-time monitor fails to stop an adversary from obtaining the *high* values by observing the termination of programs. Later in Sect. 3, we illustrate examples that manifest P_4 also covers P_2 . Considering all the limitations discussed above, a compile-time monitor based on the labelling Scheme 3 and binding mechanism P_4

Table 6. Information-flow through abnormal termination (cf. Copy6 from [21]).

Program	Label constraints	
y=0;	$\underline{y} = high$	
int sum=0;	$\underline{x} = high$	
while(true){		
sum=sum+x;	$sum \oplus \underline{x} \leq sum$	$sum = \underline{x}$
y=y+1;	$\underline{y} \oplus Low \leq \underline{y}$	
}		

Table 7. Information-flow through non-termination (Cf. Copy5 from [21]).

Program	Label constraints
y=0;	$\underline{y} = high$
while(x==0)	$\underline{x} = high$
skip;	
y=1;	$y \oplus Low \leq y$

would be a better candidate for secure certification of programs. A comparative study is given in Sect. 5 where we categorize the existing prominent IFC tools and platforms based on the labelling mechanisms.

From the above studies, we can infer:

1. While static labelling has advantages for the security certification of programs, over-approximately annotated static labels of local variables may lead to imprecision, and adversely impacts the soundness of these approaches.
2. For certifying iterative programs (terminating, non-terminating, abnormally terminating including exceptions), annotating local variables with improper static labels, and following the scheme that resets the label of the program counter, often miss to capture both the forward label propagation, and impact on static labels due to repeated backward information flow and leads to a loss of precision and soundness.
3. While the certification approach of Denning generates the relevant constraints, it fails to assert the existence of possible labelling for local variables/objects that satisfy the initial labels of global variables/objects.

If we can compute labels (or policies) for local variables such that the flow security is satisfied at all the program points, then we could consider such programs to be secure. Thus the question will be: is there a dynamic labelling procedure that realizes the same? A sound dynamic labelling approach that overcomes the limitations of the current techniques listed above shall be presented in Sect. 3.

3 Our Approach to Certification

Our approach of certification is based on a hybrid labelling of objects in the program, that could be unrolled a finite number of times. Possibility (or otherwise)

of labelling the objects of the program leads to certification (or otherwise) of the program; the same could be used in the execution monitor for checking flow security at run-time.

In the following, we describe our Dynamic Labelling (DL) algorithm. The algorithm finds its basis in Denning's certification semantics.

3.1 Dynamic Labelling (DL) Algorithm for Sequential Programs

Notation: Let G be the set of global variables/objects, L the set of local variables of a program, $(var)(e)$ the set of variables appearing in expression e , pc the program counter, and $\lambda, \lambda_0, \lambda_1, \dots$ the labelling functions that give the security label/sensitivity-level of variables and pc .

SV is a function that takes a statement/command as input, and returns the set of source variables appearing in it as output. TV is a function that takes a statement/command as input and returns the set of target variables appearing in it as output. DL is a dynamic labelling procedure/function that takes a command, clearance level of the subject executing the program, and a labelling function, as inputs, and returns either a labelling or UNABLE TO LABEL as output.

λ_{init} denotes the initial labelling. $\forall x \in G : \lambda_{init}(x)$ is given, $\forall x \in L : \lambda_{init}(x) = \perp$, and $\lambda_{init}(pc) = \perp$, where ' \perp ' is the least restrictive security class or *public*. Let P be a given program together with initial labelling for the global objects. Let s denote the subject trying to execute the program, and cl denote his clearance. If $DL(P, cl, \lambda_{init})$ returns a valid labelling, then the program preserves information-flow security when executed by the subject s . If $DL(P, cl, \lambda_{init})$ returns 'UNABLE TO LABEL', then information-flow security will be violated if the program is executed by the subject s , and the algorithm exits at that point without proceeding further.

Algorithm DL is described in Table 8. It is illustrated through examples followed by its' soundness in the sequel.

Illustrative Examples

We illustrate the advantages of our dynamic labelling procedure by analyzing the example programs from Sect. 2. Examples clearly highlight the advantages of the dynamic labelling procedure in capturing subtle information-flows through control flow path, non-termination/abnormal termination channels, etc. For each example, initial labels of the global variables are provided. We assume that the subject executing the program has the highest security label, and thus ignore the clearance field; dynamic labelling is shown in a tabular form.

We apply the proposed algorithm to the example shown in Table 2. It can be observed that the algorithm successfully labels the intermediate variable a as shown in Table 9.

Example 1. *Initial labels for global variables:* $\lambda(x) = \lambda(y) = \underline{x}$, $\lambda(z) = \underline{z}$.

Table 8. Description of algorithm DL for sequential programs (Cf. [1]).

<p>1. $S : \text{skip}:: \text{SV}(S)=\{\emptyset\}; \text{TV}(S)=\{\emptyset\}; \text{DL}(S, \text{cl}, \lambda) : \text{return } \lambda$</p> <p>2. $S : x := e:: \text{SV}(S)=\text{var}(e); \text{TV}(S)=\{x\};$ $\text{DL}(S, \text{cl}, \lambda):$ $\text{tmp} = \bigoplus_{v \in \text{var}(e) \cap G} \lambda(v)$ if (tmp $\not\leq$ cl) then exit ‘UNABLE TO LABEL’ $\lambda_1 = \lambda$ $\lambda_1(\text{pc}) = \lambda(\text{pc}) \oplus \text{tmp}$ if $x \in L$: $\lambda_1(x) = \lambda(x) \oplus \lambda(\text{pc}) \oplus \text{tmp}$ return λ_1 if $x \in G$: if ($[\lambda(\text{pc}) \oplus \text{tmp} \oplus \text{cl}] \leq \lambda(x)$) then return λ_1 else exit ‘UNABLE TO LABEL’</p>	<p>3. $S : \text{if } e \text{ then } S_1 [\text{else } S_2]::$ $\text{SV}(S)=\text{SV}(S_1) \cup \text{SV}(S_2) \cup \text{var}(e);$ $\text{TV}(S)=\text{TV}(S_1) \cup \text{TV}(S_2)$ $\text{DL}(S, \text{cl}, \lambda):$ $\text{tmp} = \bigoplus_{v \in \text{var}(e) \cap G} \lambda(v)$ if (tmp $\not\leq$ cl) then exit ‘UNABLE TO LABEL’ $\lambda' = \lambda$ $\lambda'(\text{pc}) = \lambda(\text{pc}) \oplus \text{tmp}$ $\lambda_1 = \text{DL}(S_1, \text{cl}, \lambda')$ $\lambda_2 = \text{DL}(S_2, \text{cl}, \lambda')$ $\lambda_3(\text{pc}) = \lambda_1(\text{pc}) \oplus \lambda_2(\text{pc})$ $\forall x \in L : \lambda_3(x) = \lambda_1(x) \oplus \lambda_2(x)$ return λ_3</p>
<p>4. $S : \text{while } e \text{ then } S_1::$ $\text{SV}(S)=\text{SV}(S_1) \cup \text{var}(e);$ $\text{TV}(S)=\text{TV}(S_1);$ $\text{DL}(S, \text{cl}, \lambda):$ $\text{tmp} = \bigoplus_{v \in \text{var}(e) \cap G} \lambda(v)$ if (tmp $\not\leq$ cl) then exit ‘UNABLE TO LABEL’ $\lambda_1 = \lambda$ $\lambda_1(\text{pc}) = \lambda(\text{pc}) \oplus \text{tmp}$ $\lambda_2 = \text{DL}(S_1, \text{cl}, \lambda_1)$ if ($\lambda_2 \neq \lambda_1$) $\lambda_1 = \lambda_2$ $\lambda_2 = \text{DL}(\text{while } e \text{ do } S_1,$ $\text{cl}, \lambda_1)$ return λ_2</p>	<p>5. $S : S_1; S_2::$ $\text{SV}(S)=\text{SV}(S_1) \cup \text{SV}(S_2);$ $\text{TV}(S)=\text{TV}(S_1) \cup \text{TV}(S_2);$ $\text{DL}(S, \text{cl}, \lambda):$ return $\text{DL}(S_2, \text{cl}, \text{DL}(S_1, \text{cl}, \lambda));$</p> <p>Here, problem of “insecurity” will be indicated by one of the recursive calls.</p>
<p>Note that “UNABLE TO LABEL” yields the control point where a certain object fails to satisfy the information flow policy.</p>	

Consider the example shown in Table 6 where x and y are global variables having labels \underline{x} and \underline{y} respectively. At the point $\mathbf{y=y+1}$ the program fails to satisfy the constraint $\underline{y} \oplus \underline{pc} \leq \underline{y}$ because the label of pc is updated to \underline{x} and $\underline{x} \not\leq \underline{x}$. Therefore the algorithm declares the program as flow-insecure.

The algorithm if applied to the program in Table 7 identifies the flow violation as shown in Table 10: the label of pc is updated to \underline{x} while testing the predicate; detects the flow violation at the statement $\mathbf{y=1}$ because the constraint $\underline{pc} \leq \underline{y}$ does not satisfy. Hence the algorithm fails to proceed further and declares the program as flow-insecure, thus detects the control point where a particular object can leak information.

Table 9. DL successfully labels that was not possible by static labelling (Cf. [1]).

Statement	<i>pc</i> Label	Label of local variable(<i>a</i>)
	\perp	\perp
int a=x;	\underline{x}	\underline{x}
y=a;	\underline{x}	\underline{x}
a=z;	$\underline{x} \oplus \underline{z}$	$\underline{x} \oplus \underline{z}$

Table 10. DL detects “insecurity” by failing to label a non-terminating flow (Cf. [1]).

Statement	<i>pc</i> Label
	\perp
y=0;	\perp
while x==0 skip;	\underline{x}
y=1;	UNABLE TO LABEL

Example 2. *Global variable(s): x, y; No local variables.*
Initial labels for global variables: $\lambda(x) = \underline{x}$, $\lambda(y) = \underline{y}$.

3.2 DL Algorithm for Procedure Call

A procedure declaration in an imperative language contains: procedure identifier and name of the formal parameters along with the respective mode of binding with the actual arguments like *in*, *out*, *in out* etc., declaration of variables local to the procedure and a body of the procedure. Let us consider a call to a procedure, say $p(a_1, \dots, a_m; b_1, \dots, b_n)$, where a_1, \dots, a_m are the actual input arguments and b_1, \dots, b_n are the actual input/output arguments corresponding to formal input parameters x_1, \dots, x_m and input/output parameters y_1, \dots, y_n . According to Denning’s security certification [21], execution of a procedure call $p(a_1, \dots, a_m; b_1, \dots, b_n)$ is secure iff

1. Body of the procedure p is flow secure,
2. $\underline{a_i} \leq \underline{b_j}$ if $\underline{x_i} \leq \underline{y_j}$ ($1 \leq i \leq m$, $1 \leq j \leq n$), and
 $\underline{b_i} \leq \underline{b_j}$ if $\underline{y_i} \leq \underline{y_j}$ ($1 \leq i \leq n$, $1 \leq j \leq n$).

Note that, the specification of input/output arguments and corresponding parameters in a procedure call and definition respectively, are often deprecated in modern programming languages and implicitly identified by the information passing mechanisms between arguments and parameters. Therefore, these languages are more lenient on the rigid specifications compared to their legacy counterparts, and often distinguish a *function* from a *procedure* by the mere

presence of a **return** construct. We consider the classic definition of the procedure as found in legacy languages and present the dynamic labelling algorithm for a procedure call adhering the security constraints as given by Denning. The dynamic labelling algorithm for a procedure call is presented in Table 11. Our algorithm evaluates the procedure as soon as it encounters a procedure call and returns the control at the point of invocation once it completes the evaluation.

Table 11. DL algorithm for a procedure call.

<pre> 6. $S : p(a_1, \dots, a_m; b_1, \dots, b_n) ::$ DL=(S, cl, λ): //Initialize the label of the parameters $\lambda' = \lambda_{\text{init}}$ forall $i \in 1 \dots m, \lambda'(x_i) = \lambda(a_i)$ // Evaluate the body of the procedure $\lambda_1 = \text{DL}(p - \text{body}, cl, \lambda')$ return λ_1 </pre>

The dynamic labelling algorithm performs the following operations at the time of entry & exit from the procedure: it initializes the labels of the formal input parameters with the corresponding labels of the actual input arguments as per the order they appear in the list of the arguments; then, the label of the program counter (pc) is initialized with \perp ; next, the algorithm evaluates the procedure body and finally, resets the pc to its initial label on exiting the procedure and returns the final labels to the caller. Note that, the algorithm adhere to the parameters transmission mechanisms during transferring the control from caller to the callee procedure and vice-versa. This will be clear from the subsequent example.

Consider the procedure **Add** written in Ada (shown in Fig. 1) that receives two actual input arguments by the formal input parameters X and Y , add them and store the result into formal output parameter Z . Then for a procedure call **Add**(A, B, C), the dynamic labelling algorithm shall transfer the label of actual input arguments A and B to formal input parameters X and Y respectively. The algorithm then evaluates the statement $Z := X + Y$ and generates the label of Z as LUB of the labels of X and Y . Finally, the algorithm returns the labelling function containing the mapping from Z to its new label to the caller. Note that, the DL algorithm by itself enforce the constraints given by Denning thus certify the procedure as flow secure.

We discuss the possible variants of our labelling algorithm for programming languages that have different information passing mechanisms for subprograms such as pass-by-value (e.g., Java, C, C++), pass-by-reference (e.g., C, C++) and pass-by-object reference (e.g., Python).

```

1 procedure Add(X, Y: in Integer; Z: out Integer) is
2 begin
3   Z := X + Y;
4 end

```

Fig. 1. A procedure `Add` written in Ada.

Case 1: Consider a language where the actual input parameters act as local variables to the subprogram (like `in` parameters in Ada). Then, the labels of these local variables are initialized with the labels of the corresponding input arguments. Since the parameters are purely local variables, therefore, any changes in the labels during the evaluation process do not affect the corresponding arguments. However, it is necessary to have the construct `return` at the end of the subprogram to transfer the changes to the caller. Such subprograms are referred to as *function* and invoked within an expression. In such a case, we could refine the dynamic labelling algorithm which might treat the function return parameters as the output parameters.

Case 2: In case of a language that follows pass-by-reference (like C, C++) for passing the information to a subprogram, any reference to the label of the parameter is considered to be a reference to the label of the argument. Therefore, changes in the labels of the parameters are *direct* changes in the labels of corresponding arguments as they appear in the list of the arguments. Therefore, the dynamic labelling algorithm could be modified to consider the input parameters that are passed by reference, as the output parameters also and treat accordingly. Note that, in presence of global variables in the list of input arguments the corresponding formal parameters shall be considered as global within the scope of the procedure.

Case 3: A programming languages Python follows a complex pass-by-object reference mechanism – a combination of pass-by-value and pass-by-reference depending on the data type of the argument, i.e., *immutable* or *mutable* respectively. The original references of the immutable objects, e.g., integers, string, tuples etc. that are passed to subprogram cannot be changed in-place, therefore treated as local to the subprogram. Whereas, object references for mutable objects such as `list` might be changed in-place in the subprogram depending on the operation performed on it. E.g., a compound assignment, i.e., ‘+=’ performs an in-place assignment for a `list` argument. Therefore, in the presence of such operations we might consider to refine the dynamic labelling algorithm and treat the mutable object references in the parameter’s list as the formal output parameters and follow the approach similar to Case 2.

3.3 Soundness of Algorithm DL

We shall establish the termination of our dynamic labelling algorithm, and its soundness w.r.t. non-interference.

Clearly, the procedure in Table 8 always terminates and is efficient. In fact, it is linear in the size of the program. This fact is formally established through the propositions below.

Proposition 1. *$DL(S, cl, \lambda)$ always terminates for any program S not containing iteration, any label cl , and any labelling λ .*

Proof. The proof is by structural induction. For the base case, it is trivial to observe that the proposition holds for `skip`, and $x := e$.

For the inductive step, it is easy to prove that if the proposition holds for S_1 and S_2 , then it also holds for `if e then S_1 else S_2` , and for $S_1; S_2$.

Proposition 2. *For any program S not containing iteration, any label cl , and any labelling λ , if $DL(S, cl, \lambda)$ returns a valid labelling λ_1 , then $\lambda_1(pc) = \lambda(pc) \oplus_{v \in SV(S) \cap G} \lambda(v)$.*

The proof of this proposition is by structural induction and is omitted for brevity.

Proposition 3. *$DL(\text{while } e \text{ do } S, cl, \lambda)$ always terminates for any program S not containing iteration, any label cl , and any labelling λ .*

Proof. From the definition of DL , we note that the only case in which the evaluation of $DL(\text{while } e \text{ do } S, cl, \lambda)$ does not terminate is when either the evaluation of $DL(S, cl, \lambda_1)$ does not terminate, or $DL(\text{while } e \text{ do } S, cl, \bullet)$ goes into an infinite recursion.

The former is impossible due to Proposition 1. Impossibility of the latter is shown by considering the evaluation of $DL(\text{while } e \text{ do } S, cl, \lambda)$:

1. $\lambda_1 = \lambda$, $\lambda_1(pc) = \lambda(pc) \oplus_{v \in \text{var}(e) \cap G} \lambda(v)$
2. $\lambda_2 = DL(S, cl, \lambda_1)$
3. If $\lambda_2 == \lambda_1$ the evaluation terminates and there is nothing to prove. So we assume that $\lambda_2 \neq \lambda_1$. In this case $DL(\text{while } e \text{ do } S, cl, \lambda_2)$ is invoked which proceeds as follows.
4. $\lambda_3 = \lambda_2$, $\lambda_3(pc) = \lambda_2(pc) \oplus_{v \in \text{var}(e) \cap G} \lambda_2(v)$
5. $\lambda_4 = DL(S, cl, \lambda_3)$
6. If $\lambda_4 == \lambda_3$ the evaluation terminates and there is nothing to prove. So we assume that $\lambda_4 \neq \lambda_3$. In this case $DL(\text{while } e \text{ do } S, cl, \lambda_4)$ is invoked which proceeds as follows.
7. $\lambda_5 = \lambda_4$, $\lambda_5(pc) = \lambda_4(pc) \oplus_{v \in \text{var}(e) \cap G} \lambda_4(v)$
8. $\lambda_6 = DL(S, cl, \lambda_5)$

We claim that $\lambda_6 == \lambda_5$. The proof is given below.

1. First iteration:

$$\lambda_1(pc) = \lambda(pc) \oplus_{v \in \text{var}(e) \cap G} \lambda(v)$$

$$\begin{aligned} \lambda_2(pc) &= \lambda_1(pc) \oplus_{v \in SV(S) \cap G} \lambda(v) \\ &= \lambda(pc) \oplus_{v \in \text{var}(e) \cap G} \lambda(v) \oplus_{v \in SV(S) \cap G} \lambda(v) \\ &= \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v) \end{aligned} \tag{1}$$

2. Second iteration:

$$\begin{aligned}
\lambda_3(pc) &= \lambda_2(pc) \oplus_{v \in \text{var}(e) \cap G} \lambda(v) \\
&= \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v) \\
&\quad \oplus_{v \in \text{var}(e) \cap G} \lambda(v) \\
&= \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v)
\end{aligned} \tag{2}$$

$$\begin{aligned}
\lambda_4(pc) &= \lambda_3(pc) \oplus_{v \in SV(S) \cap G} \lambda(v) \\
&= \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v) \\
\lambda_4(x) &= \lambda_3(x) \oplus \lambda_3(pc) \oplus_{v \in SV(S) \cap G} \lambda(v) \\
&= \lambda_3(x) \oplus \lambda_3(pc)
\end{aligned} \tag{3}$$

This is because the label of PC is already influenced by all the global variables in S .

$$= \lambda_2(x) \oplus \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v)$$

3. Third iteration:

$$\begin{aligned}
\lambda_5(pc) &= \lambda_4(pc) \oplus_{v \in \text{var}(e) \cap G} \lambda(v) \\
&= \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \oplus_{v \in \text{var}(e) \cap G} \lambda(v) \\
&= \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v)
\end{aligned} \tag{4}$$

$$\begin{aligned}
\lambda_5(x) &= \lambda_4(x) \\
&= \lambda_2(x) \oplus \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v)
\end{aligned}$$

$$\begin{aligned}
\lambda_6(pc) &= \lambda_5(pc) \oplus_{v \in SV(S) \cap G} \lambda(v) \\
&= \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v) \\
\lambda_6(x) &= \lambda_5(x) \oplus \lambda_5(pc) \oplus_{v \in SV(S) \cap G} \lambda(v) \\
&= \lambda_2(x) \oplus \lambda(pc) \oplus_{v \in (\text{var}(e) \cup SV(S)) \cap G} \lambda(v)
\end{aligned} \tag{5}$$

It can be observed that $\lambda_6 = \lambda_5$. Thus, we can conclude that, for the iteration statement, the dynamic labelling procedure terminates after a maximum of three iterations.

Combining Propositions 1 and 3 leads to the following proposition.

Proposition 4. *DL(while e do S, cl, λ) always terminates for any program S , any label cl , and any labelling λ .*

Proposition 4 immediately establishes termination of DL as formalized below.

Proposition 5. *DL(S, cl, λ) always terminates for any sequential program S , any label cl , and any labelling λ .*

Next, we prove some results that highlight the important characteristics of our dynamic labelling procedure.

Proposition 6. *During the dynamic labelling of any program S with any clearance cl , i.e. during the evaluation of $DL(S, cl, \lambda_{init})$, $\lambda(pc) \leq cl$ always holds.*

Proof. In the initial state $\lambda_{init}(pc) = \perp \leq cl$. From the definition of DL , we note that the label of pc gets updated by taking its LUB with tmp only when $tmp \leq cl$. Therefore $\lambda(pc) \leq cl$ always holds due to simple lattice properties.

Proposition 7. *An assignment to a global variable x is deemed safe by the dynamic labelling algorithm for a program executing with clearance cl if and only if $cl \leq \lambda(x)$.*

Proof. (Necessity of $cl \leq \lambda(x)$) From the definition of DL for an assignment statement, it is immediately clear that if the operation is deemed safe then it must be the case that $cl \leq \lambda(x)$.

(Sufficiency of $cl \leq \lambda(x)$) Note that we have $\lambda(pc) \leq cl$ from Proposition 6, and for the control to reach the point, we need $tmp \leq cl$, thus reducing the check $(\lambda(pc) \oplus tmp \oplus cl) \leq \lambda(x)$ to $cl \leq \lambda(x)$.

Proposition 8. *During the dynamic labelling of any program S , $\lambda(pc)$ is monotonically non-decreasing.*

The proof of the above is trivially obtained by structural induction and is omitted for brevity.

Next, we prove a generalization of the result in Proposition 2.

Proposition 9. *For any program S , any label cl , and any labelling λ , if $DL(S, cl, \lambda)$ returns a valid labelling λ_1 , then $\lambda_1(pc) = \lambda(pc) \bigoplus_{v \in SV(S) \cap G} \lambda(v)$.*

The proof of this proposition is by structural induction and is omitted for brevity.

Proposition 10. *During the dynamic labelling of any program S , for all $x \in L$, $\lambda(x)$ is monotonically non-decreasing.*

Proof. For $x \in L$, the label of x is updated by the dynamic labelling procedure only in the case of explicit assignment. In this case the label of x changes to $\lambda(x) \oplus \lambda(pc) \oplus tmp$. Monotonicity of $\lambda(pc)$ immediately gives us that $\lambda(x)$ is also monotonically non-decreasing.

Proposition 11. *During the dynamic labelling of any program S i.e. $DL(S, cl, \lambda_{init})$, $\forall x \in L \lambda(x) \leq \lambda(pc)$ always holds.*

Proof. In the initial state we have $\lambda_{init}(x) = \lambda_{init}(pc) = \perp$. We will show that every time the label of x is updated, the property holds in the new state also.

– $S :: x := e$

$$\begin{aligned}
 \lambda_1(x) &\leq \lambda(x) \oplus \lambda(pc) \oplus tmp \\
 \lambda_1(x) &\leq \lambda(pc) \oplus tmp \text{ [by hypothesis } \lambda(x) \leq \lambda(pc) \text{]} \\
 \lambda_1(x) &\leq \lambda_1(pc)
 \end{aligned} \tag{6}$$

– $S :: \text{if } e \text{ then } S_1 \text{ else } S_2$

$$\begin{aligned} \lambda_3(x) &\leq \lambda_1(x) \oplus \lambda_2(x) \\ \lambda_3(x) &\leq \lambda_1(pc) \oplus \lambda_2(pc) \text{ [by hypothesis]} \\ \lambda_3(x) &\leq \lambda_3(pc) \end{aligned} \tag{7}$$

Relation with Non-interference

In this section, we establish that the dynamic labelling algorithm is sound w.r.t. non-interference [3].

A simple example illustrates the relation with non-interference. Consider the program $P_1 : x := y;$ - where x, y are global objects, and the labelling $\lambda(y) = l_2$, $\lambda(x) = l_3$, where l_2 and l_3 come from a total order $l_1 \leq l_2 \leq l_3 \leq l_4$. Consider four subjects s_1, s_2, s_3 , and s_4 , with clearances l_1, l_2, l_3 , and l_4 respectively.

P_1 is non-interfering. Dynamic labelling of P_1 succeeds only for subjects s_2 and s_3 . Dynamic labelling fails for s_1 because his clearance is below y , and therefore should not be allowed to access y . Similarly, dynamic labelling fails for s_4 because his clearance is above x and therefore should not be allowed to update x .

In the following, we shall formally establish the soundness of the dynamic labelling procedure w.r.t. non-interference. Note that globals are the only observables in this setting.

Theorem 1 (Soundness). *If there exists a subject for which a program is declared secure by the dynamic labelling procedure in Table 8, then the program is non-interfering.*

Proof. For reasoning about value based non-interference, it suffices to work with the last update to a low labelled variable. From the definition of DL , we observe that the only place where a global variable is potentially updated is guarded by the condition $\lambda(pc) \oplus tmp \oplus cl \leq \lambda(x)$. In particular, since we are dealing with $\lambda(x) = \text{low}$, and the program is declared secure by the dynamic labelling procedure, we can immediately infer that $\lambda(pc) = tmp = cl = \text{low}$. This guarantees that no high labelled variable could have been accessed by this time in the execution.

Traditional methods for the security certification of programs do not consider the subject labels. Let $DL_1(S, \lambda)$ be a modified dynamic labelling algorithm obtained by ignoring cl from the algorithm given in Table 8. We now prove that even this algorithm is sound w.r.t non-interference.

Theorem 2. *If a program S is declared secure by procedure DL_1 i.e., $DL_1(S, \lambda)$ returns a valid labelling, then the program is non-interfering.*

Proof of this theorem is exactly the same as the proof of the previous theorem, and is omitted.

Finally, the set of programs declared secure by traditional certification methods that reset PC and use static labels for variables (Jif is a prominent representative of this class) is incomparable to the set of programs declared secure by our dynamical labelling algorithm as shown below.

Proposition 12. *There are programs declared secure by static labelling that cannot be dynamically labelled (insecure by our definition), and there are programs declared secure by our approach that static labelling rejects as insecure.*

Proof: Programs in Tables 6 and 7 provide an example for the former, while the program in Table 2 provides an example for the latter.

4 DL Algorithm for Concurrent Programs

The dynamic labelling algorithm we propose for a sequential imperative programming language can be easily extended for concurrent programs, where shared variables among the concurrent threads are potential threats for leaking information. In our approach apart from the set of global variables, the dynamic labelling algorithm shall be given a set of variables that are shared among the threads during the execution. The labels of the shared variables might be defined globally and immutable or local to the system and mutable. In case of a shared variable with an immutable label, the label shall not be changed throughout the computation of the algorithm, whereas, the shared variables with mutable labels are changed to accommodate the flows between threads.

Table 12. DL algorithm for concurrent programs.

<pre> 7. $S : T_1 \dots T_n ::$ $DL=(S, cl, \lambda):$ $\lambda_{final} = \lambda$ do $\lambda_{start} = \lambda_{final}$ for each $i \in 1 \dots n$ $\lambda_i(v) = \lambda_{final}(v), \forall v \in H$ $\lambda_i = DL(T_i, cl, \lambda_i)$ $\lambda_{final}(v) = \lambda_i(v), \forall v \in H$ while $\exists v \in H(\lambda_{start}(v) \neq \lambda_{final}(v))$ return λ_{final} </pre>
--

Let us consider n number of threads T_1, \dots, T_n that are executing simultaneously, where each thread is considered as a sequential program comprise of the statements discussed in Sect. 3. Further, assume that a set of shared variables H is provided to the algorithm. The dynamic labelling algorithm for concurrent programs is shown in Table 12. The algorithm, first computes the label of shared variables by executing each thread independently. Then, it compares the final labels of shared variables with their initial labels. The algorithm repeats the process again if there exist a single shared variable with unequal label. The process is iterated until the label of all the shared variables converge in the lattice.

Proposition 13. *$DL(T_1 || T_2 || \dots || T_n, cl, \lambda)$ always terminates for any n number of concurrent threads given a priori, and any $cl \notin \lambda$.*

Proof. The algorithm DL for the concurrent programs might not terminate only for the cases in which either evaluation of any thread T_i , i.e., $DL(T_i, cl, \lambda_i)$ does not terminate, or $DL(T_1 || T_2 || \dots || T_n, cl, \bullet)$ goes into an infinite recursion. The former case is not possible since each thread is a sequential program and Proposition 5 holds good for the sequential programs. The latter case is impossible due to the reason shown below by considering the evaluation of $DL(T_1 || T_2 || \dots || T_n, cl, \lambda)$:

1. $\lambda_1 = \lambda$
2. For all $v \in H$, $\lambda_2(v) = \lambda_1(v) \oplus \lambda^1(v) \oplus \dots \oplus \lambda^n(v)$, where λ^i denotes the labelling function received after evaluating the thread T_i , i.e., $\lambda^i = DL(T_i, cl, \bullet)$, $1 \leq i \leq n$
3. If $\lambda_1(v) == \lambda_2(v)$ for all $v \in H$ then the evaluation terminates. Therefore, let us assume there exist at least a v for which $\lambda_1(v) \neq \lambda_2(v)$
4. $\lambda_3 = \lambda_2$
5. For all $v \in H$, $\lambda_4(v) = \lambda_3(v) \oplus \lambda^1(v) \oplus \dots \oplus \lambda^n(v)$
6. If $\lambda_3(v) == \lambda_4(v)$ for all $v \in H$ then the evaluation terminates. Therefore, let us assume there exist at least a v for which $\lambda_1(v) \neq \lambda_2(v)$
7. $\lambda_5 = \lambda_4$
8. For all $v \in H$, $\lambda_6(v) = \lambda_5(v) \oplus \lambda^1(v) \oplus \dots \oplus \lambda^n(v)$

We claim that $\forall v \in H$, $\lambda_6(v) == \lambda_5(v)$. The proof is similar to the proof for the Proposition 3 as likewise *pc*, the algorithm performs LUB for each shared variable while evaluating the threads individually to accommodate the changes. Therefore, we can conclude that, for the concurrent threads, the dynamic labelling algorithm terminates after maximum of three iterations.

The proof for the Propositions 5 and 13 for the concurrent context shows that the procedure requires a finite number of unrolling rather than a full termination to converge the labels.

Illustrative Example

An example of information leakage due to concurrent access is presented by [22]. The order of the assignments to x and y variables in VIP program depends on the secret value of h . The program in Newsmonger runs simultaneously and prints the values of x and y inside an infinite loop that are shared variables having no explicit label. If Newsmonger runs in between any of the assignments that exist in line 3 and 5 of VIP, it could reveal the value of h . In our approach the DL procedure would evaluate the labels for the threads VIP and Newsmonger separately along with the shared variables x and y until the label of the variables converge in the lattice. Therefore, it would compute the labels for x and y as equivalent to h and identify the possible flow security risk while executing the statement “output y ” in Newsmonger as label of h cannot flow to public.

5 Comparison with Related Work

In this section, we first briefly describe tools and platforms that enforce rich information flow policies and then compare our approach with the existing

Table 13. (a) VIP (left) and (b) Newsmonger (right) (Cf. [1]).

<pre> 1 x:=0; y:=0 2 if h then 3 x:=1; y:=1; 4 else 5 y:=1; x:=1; 6 end;</pre>	<pre> 1 while true do 2 output x; 3 output y; 4 done;</pre>
--	---

approaches. Further, we discuss the applicability and limitations for certifying different classes of programs like (i) *termination-sensitive* programs, (ii) concurrent/non-deterministic programs, etc.

Information Flow Tools

In the last decade a large number of information flow secure tools have been developed to enforce rigorous flow security policies through prevailing programming languages. For example, Jif [7], JOANA [17], Paragon [19] for Java, FlowFox [23], JSFlow [24], IFC4BC [25] for JavaScript, FlowCaml [8] for Caml, λ_{DSec} [16] for lambda calculus, LIO [18], HLIO [20] for Haskell and SPARK flow analysis [26] for SPARK. Also flow secure platforms for instance, Jif/split [27], Asbestos [28], HiStar [29], Flume [30], Aeolus [31] and flow checking systems that implements *sparse information labeling* [15], permissive-upgrade strategy [32] and identifies public labels and delayed exception [33] incorporate different label mechanisms shown in Table 14. While we have omitted some similar prominent platforms for lack of space, it may be noted that DL realizes the needed characteristics required for IFC.

Table 14. Comparison of IFC tools and platforms (Cf. [1]).

Tools and Platforms	Labelling mechanism	Flow-sensitive	Termination-sensitive
Jif	P_1	✗	✗
Paragon	P_1	✗	✗
FlowCaml	P_1	✗	✗
λ_{DSec}	P_3	✓	✗
LIO	P_2	✗	✗
$\lambda_l^{l_{IO}}$	P_4	✓	✓
Aeolus	P_1	✗	✗
DL	P_4	✓	✓

The earliest attempt to capture flow-sensitive labels at run-time was observed in work on λ_{DSec} . This was the first to propose general dynamic labels whose type system was proved to enforce non-interference. The core language λ_{DSec} is

a security-typed lambda calculus that supports first-class dynamic labels where labels can be checked and manipulated at run-time. Also, labels can be used as statically analyzed type annotations. The type system of λ_{DSec} prevents illegal information flows and guarantees that any well-typed program satisfies the non-interference property. In this language, the label of the pc is a lower bound on the memory effects of the function, and an upper bound on the pc label of the caller, but unlike DL, pc is not updated dynamically after executing each statement hence the language falls in the category P_3 . The non-interference property discussed in λ_{DSec} is termination-insensitive, and also does not deal with timing channels. In the following, we provide a detail discussion on **LI0** that shares a common paradigm and also subsumes the results of λ_{Dsec} .

Comparison with [18]

Here, the authors have built a labelled **IO** Haskell library, called **LI0**, for certifying Haskell programs. **LI0** tracks a single mutable *current label* (like program-counter) at run-time and allows access to IO operations. The unit is responsible for ensuring that the current label keeps track of all the observed data and regulate label modification. A type constructor **Labeled** is used to hold the restriction only value and is mutable during run-time. At each computation, **LI0** keeps tracks of the *current label* and allow access to IO functionality, e.g., labeled file systems. Current label is evaluated as the *upper bound* of all the label observed during the program execution.

Consider reading a secret reference: $a \leftarrow \text{readLIORef secret}$, where the value “secret” is labeled as L_S . Now to satisfy the information flow check i.e., $(L_S \text{ canFlowTo } L_C)$ the *current label* shall rise to $(L_C \text{ join } L_S)$ to read the secret value. Note that the value a is not labeled explicitly. Now let us take an example that wish to write the value of an object (a) to an output channel: $\text{writeLIORef output } a$, where the output channel is labeled as L_O (set dynamically according to the user executing the command). It is only permissible to modify or write data into the output channel when $(L_C \text{ canFlowTo } L_O)$ is satisfied. A second label *current clearance* (L_{cl}) provides an upper bound to *current label*. Hence, the computation cannot create, read or write to objects labeled L if $L \text{ canFlowTo } L_{cl} == \text{False}$.

Although our approach overlaps with that of **LI0**, there are subtle differences that are briefed below:

- Unlike statically evaluating the labels in our DL algorithm, the approach in **LI0** is based on run-time *floating-label* system.
- Compared to DL algorithm, the **LI0** library provides IO actions that perform *termination-insensitive* flow analysis.
- Due to *flow-insensitive* labelling **LI0** does not provide sensitivity level of each intermediate object precisely.

We illustrate each of these points in the following.

Comparison of Labelling Mechanism

The characterization of security labels, when associated with objects, is an essential aspect of IFC analysis [34]. Security labels of subjects/objects can be muta-

ble or immutable. *Flow-sensitive* IFC monitors allow changing the security labels throughout the computation thus increase the permissiveness, and also alleviate the burden of explicit label annotations. Note that these monitors perform the flow analysis during execution-time or compile-time. Mutable label flow analysis during execution-time helps to determine the flow-sensitivity of objects at run-time precisely, but compile-time analysis reduces the incident of *false-alarms* and allows more programs as secure.

LIO keeps track of a single mutable *current* floating-label that is elevated (e.g., from *low* to *high*) at run-time to accommodate sensitive reading; hence, **LIO** is *flow-sensitive* in the *current label*. However, **LIO** is *flow-insensitive* in intermediate object labels. To allow more programs by the run-time monitor, an extension of **LIO** presented by [35] that safely manipulates a label on the reference label. A *label on the label* describes the confidentiality of the **LIO** reference label itself. The run-time monitor upgrades a label of a reference only if that *label on the label can flow to floating-label*. Note that **LIO** follows the labelling mechanism P_2 whereas the proposed extension incorporates P_4 . Another extension of **LIO** monad, i.e., **HLIO** that provides programmers the flexibility to defer flow check of part of the program to run-time (like **LIO**) or static-time (unlike **LIO**), boosts permissiveness of the monitor.

Algorithm DL, is *flow-sensitive* in the absence of method-calls and a compile-time monitor built upon it would satisfy P_4 , and hence it is more permissive than the other approaches highlighted above. DL follows a hybrid labelling approach where the labels of global variables are assumed to be fixed, and label of each intermediate object is allowed to vary, thus, *flow-sensitive* dynamic labels are obtained.

Termination-insensitive Flow Analysis

Information leak depending on the termination of the program may remain undetected by the run-time monitor that extends the **LIO** library unit. A program that exploits `toLabeled` function as shown by [36] may lead to information leak through the termination channel. `toLabeled` l m executes the **LIO** operation m and encapsulates the returned value with label l . However, the function does not increase the *current label*. Hence one can write an iterative program that executes a `toLabeled` function depending on a secret value or diverges otherwise. Assuming the initial *current label* as *low*, and as it remains unchanged even after executing `toLabeled`, an adversary can determine the secret value by observing termination of the program through standard output.

The DL algorithm keeps track of the sensitive labels observed by the *pc* at compile-time. Therefore, a program that tries to pass termination information to standard output shall abide by the information flow policy. Hence the proposed labelling approach performs an exemplary *termination-sensitive* flow analysis.

Applicability of LIO in Concurrent Context

As initially **LIO** was not considered for dynamic flow-sensitive concurrent settings, an extension is proposed in [36] that mitigates and eliminates termination and timing channels in concurrent programs. In that article, a separate *current label* for each thread is mentioned that keeps track of the sensitivity of the

data it has observed, and restrict the locations to which the thread can write. Hence, while termination and timing of these threads that may expose the secret values, the thread requires to raise its *current label*. This prevents lower security threads from observing confidential information written in shared locations. Another extension of LIO proposed by [35] provides the primitive of automatic upgrade that safely updates *flow-sensitive* label references. Both the extensions are shown to be equally applicable for concurrent context. However, the extensions may not stop the concurrent programs from revealing the secret value. Let us assume the assignments in the program shown in Table 13 are LIO operations and the labels of the *current label* and h are L_C and L_h respectively. Then, the label of y is evaluated as $L_C \text{ join } L_h$ when h is false. Now, before labelling of x is done by LIO, Newsmonger might disclose the value in x which in turn would reveal the secret value h .

Determining the Label of Intermediate Variables

As early as 1975, Dorothy Denning proposed in her thesis [37] a run-time source-to-source transformation to guarantee flow security of programs having the selection or iteration statements. The method introduces additional code for checking possible flow violations at run-time. In a sense, her method, simulates possible information flows for each variable that has to lie between possible highest and lowest levels. As against this, our method succinctly captures security labels of variables explicitly without introducing additional code in the program.

6 Conclusions

We have presented a dynamic labelling algorithm, i.e., DL for flow security certification of imperative programs. Our labelling algorithm is appropriate for not only classic actual-formal parameter passing mechanism but also for mechanisms like pass-by-object reference used in Python as illustrated in this paper. Another characteristic of our approach allows us to capture the labels for termination-, progress-sensitive programs and has shown to be more security precise compared to existing approaches. Also, we have established the soundness of our approach with respect to non-interference. We have extended the DL algorithm to evaluate a concurrent context consisting of a finite number of sequential programs referred to as threads, sharing a given set of variables and executing concurrently. The algorithm DL is shown to be always terminating after a finite number of iterations. So far, we have built a platform for certifying sequential Python programs. The platform is enriched with the novel features such as *declassification* that helps to build a multi-level secure system that follows a decentralized labelling model. We have illustrated the efficacy of DL to concurrent programs and currently extending our platform for certification of concurrent programs as well.

Acknowledgements. The authors thank the Ministry of Electronics and Information Technology (MeitY), Govt. of India, for the generous support to the Information Security Research & Development Center (ISRDC) at IIT Bombay.

References

1. Ghosal, S., Shyamasundar, R.K., Kumar, N.V.N.: Static security certification of programs via dynamic labelling. In: Proceedings of the 15th International Joint Conference on e-Business and Telecommunications, ICETE 2018 - Volume 2: SECRIPT, Porto, Portugal, 26–28 July 2018, pp. 400–411 (2018)
2. Denning, D.E., Denning, P.J.: Certification of programs for secure information flow. *CACM* **20**(7), 504–513 (1977)
3. Volpano, D.M., Irvine, C.E., Smith, G.: A sound type system for secure flow analysis. *J. Comput. Secur.* **4**(2/3), 167–188 (1996)
4. Denning, D.E.: A lattice model of secure information flow. *CACM* **19**(5), 236–243 (1976)
5. Myers, A.C., Liskov, B.: Protecting privacy using the decentralized label model. *ACM TOSEM* **9**(4), 410–442 (2000)
6. Myers, A.C.: JFlow: practical mostly-static information flow control. In: Proceedings of 26th ACM Symposium on POPL, pp. 228–241 (1999)
7. Myers, A.C., Zheng, L., Zdancewic, S., Chong, S., Nystrom, N.: Jif: Java information flow (2001). <http://www.cs.cornell.edu/jif>
8. Simonet, V., Rocquencourt, I.: Flow Caml in a nutshell. In: Proceedings of 1st APPSEM-II Workshop, pp. 152–165 (2003)
9. Stefan, D., Russo, A., Mitchell, J.C., Mazières, D.: Flexible dynamic information flow control in the presence of exceptions. *CoRR* abs/1207.1457 (2012)
10. Ryan, P., McLean, J., Millen, J., Gligor, V.: Non-interference, who needs it? In: Proceedings of 14th IEEE CSF Workshop, pp. 237–238 (2001)
11. Sabelfeld, A., Myers, A.C.: Language-based information-flow security. *IEEE J. Sel. Areas Commun.* **21**(1), 5–19 (2003)
12. Hicks, B., King, D., McDaniel, P.: Jifclipse: development tools for security-typed languages. In: Proceedings of Workshop on PLAS, pp. 1–10 (2007)
13. Goguen, J.A., Meseguer, J.: Security policies and security models. In: IEEE Symposium on SP, p. 11 (1982)
14. Askarov, A., Sabelfeld, A.: Tight enforcement of information-release policies for dynamic languages. In: Proceedings of 22nd IEEE CSF Symposium, pp. 43–59 (2009)
15. Austin, T.H., Flanagan, C.: Efficient purely-dynamic information flow analysis. In: Proceedings of the ACM SIGPLAN 4th Workshop on PLAS, pp. 113–124 (2009)
16. Zheng, L., Myers, A.C.: Dynamic security labels and static information flow control. *Int. J. Inf. Secur.* **6**(2–3), 67–84 (2007)
17. Hammer, C., Snelling, G.: Flow-sensitive, context-sensitive, and object-sensitive information flow control based on program dependence graphs. *Int. J. Inf. Secur.* **8**(6), 399–422 (2009)
18. Stefan, D., Russo, A., Mitchell, J.C., Mazières, D.: Flexible dynamic information flow control in haskell. *ACM SIGPLAN Not.* **46**, 95–106 (2011)
19. Broberg, N., van Delft, B., Sands, D.: Paragon for practical programming with information-flow control. In: Shan, C. (ed.) *APLAS 2013*. LNCS, vol. 8301, pp. 217–232. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-03542-0_16
20. Buiras, P., Vytiniotis, D., Russo, A.: HLIO: mixing static and dynamic typing for information-flow control in Haskell. In: *ACM SIGPLAN Notices*, vol. 50, pp. 289–301. ACM (2015)
21. Robling Denning, D.E.: *Cryptography and Data Security*. Addison-Wesley Longman Publishing Co., Boston (1982)

22. Le Guernic, G.: Automaton-based confidentiality monitoring of concurrent programs. In: Proceedings of 20th IEEE CSF Symposium, pp. 218–232 (2007)
23. De Groef, W., Devriese, D., Nikiforakis, N., Piessens, F.: FlowFox: a web browser with flexible and precise information flow control. In: Proceedings of ACM CCS, pp. 748–759 (2012)
24. Hedin, D., Birgisson, A., Bello, L., Sabelfeld, A.: JSFlow: tracking information flow in Javascript and its APIs. In: Proceedings of 29th Annual ACM SAC, pp. 1663–1671 (2014)
25. Bichhawat, A., Rajani, V., Garg, D., Hammer, C.: Information flow control in WebKit’s JavaScript bytecode. In: Abadi, M., Kremer, S. (eds.) POST 2014. LNCS, vol. 8414, pp. 159–178. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54792-8_9
26. Barnes, J.G.P.: High Integrity Software: The Spark Approach to Safety and Security. Pearson Education, London (2003)
27. Zdancewic, S., Zheng, L., Nystrom, N., Myers, A.C.: Secure program partitioning. *ACM Trans. Comput. Syst. (TOCS)* **20**(3), 283–328 (2002)
28. Efstathopoulos, P., et al.: Labels and event processes in the asbestos operating system. In: Proceedings of 20th ACM SOSP, vol. 39, pp. 17–30 (2005)
29. Zeldovich, N., Boyd-Wickizer, S., Kohler, E., Mazières, D.: Making information flow explicit in HiStar. In: Proceedings of 7th Symposium on OSDI, pp. 263–278 (2006)
30. Krohn, M.N., et al.: Information flow control for standard OS abstractions. In: Proceedings of 21st ACM SOSP, pp. 321–334 (2007)
31. Cheng, W., et al.: Abstractions for usable information flow control in Aeolus. In: USENIX Annual Technical Conference, pp. 139–151 (2012)
32. Austin, T.H., Flanagan, C.: Permissive dynamic information flow analysis. In: Proceedings of the 5th ACM SIGPLAN Workshop on PLAS, p. 3 (2010)
33. Hritcu, C., Greenberg, M., Karel, B., Pierce, B.C., Morrisett, G.: All your IFCException are belong to us. In: IEEE Symposium on SP, pp. 3–17 (2013)
34. Hunt, S., Sands, D.: On flow-sensitive security types. *ACM SIGPLAN Not.* **41**, 79–90 (2006)
35. Buiras, P., Stefan, D., Russo, A.: On dynamic flow-sensitive floating-label systems. In: Proceedings of IEEE 27th CSF Symposium, pp. 65–79 (2014)
36. Stefan, D., Russo, A., Buiras, P., Levy, A., Mitchell, J.C., Mazières, D.: Addressing covert termination and timing channels in concurrent information flow systems. *ACM SIGPLAN Not.* **47**, 201–214 (2012)
37. Denning, D.E.R.: Secure information flow in computer systems (1975)