

Instruction-level security typing by abstract interpretation

Nicoletta De Francesco · Luca Martini

Published online: 13 February 2007
© Springer-Verlag 2007

Abstract We present a method based on abstract interpretation to check secure information flow in programs with dynamic structures where input and output channels are associated with security levels. In the concrete operational semantics each value is annotated by a security level dynamically taking into account both the explicit and the implicit information flows. We define a collecting semantics which associates with each program point the set of concrete states of the machine when the point is reached. The abstract domains are obtained from the concrete ones by keeping the security levels and forgetting the actual values. Using this framework, we define an abstract semantics, called instruction-level security typing, that allows us to certify a larger set of programs with respect to the typing approaches to check secure information flow. An efficient implementation is shown, operating a fixpoint iteration similar to that of the Java bytecode verification.

Keywords Abstract interpretation · Information flow · Language based security

This work was partially supported by the Italian COFIN 2004 project “AIDA: Abstract Interpretation Design and Application”.

N. De Francesco · L. Martini (✉)
Dipartimento di Ingegneria dell’Informazione,
Università di Pisa, Via Diotisalvi, 2, 52126 Pisa, Italy
e-mail: luca.martini@iet.unipi.it

N. De Francesco
e-mail: nico@iet.unipi.it

1 Introduction

Protecting confidential information has ever been an issue on computer systems. One of the most widespread methods to ensure confidentiality is using some Discretionary Access Control (DAC) mechanisms, both to prevent unauthorized access and to permit authorized access to information, for a given policy of authorization. As reported by [35], a DAC is defined as follows:

“A means of restricting access to objects based on the identity of subjects and/or groups to which they belong. The controls are ‘discretionary in the sense that a subject with a certain access permission is capable of passing that permission (perhaps indirectly) on to any other subject.”

Unfortunately, the DAC mechanism can be in some cases inadequate. In fact, although it checks information release, it does not control its propagation. The *secure information flow* within programs in multilevel secure systems requires that information at a given security level does not flow to lower levels [18, 19]. Analyzing secure information flow allows a finer inspection of confidentiality than that obtained by using DAC mechanisms. In fact, checking information flows makes it possible to control, once given an access right, whether the accessed information is properly used, according to some confidentiality policy. Assume that x and y are variables. Examples of information flows are the instructions $x:=y$; and `if (y=0) then x:=0; else x:=1;`. In the first case there is an *explicit* information flow from y to x , while, in the second case, there is an *implicit* information flow: in both cases, checking the final value of x reveals information on the value of y .

We consider sequential programs communicating with the external environment by means of input and output channels. The program defines also a security policy by assigning a security level to each channel. A program has secure information flow if the observation of a channel having some security level does not reveal any information about the values input from channels associated with higher security levels. The language includes dynamic structures and pointers.

We give a framework for analyzing secure information flow based on Abstract Interpretation (AI). AI [13–15] is a method for analyzing programs in order to collect approximate information about their run-time behavior. It is based on a non-standard semantics, that is a semantic definition in which a simpler (abstract) domain replaces the standard (concrete) one, and the operations are interpreted on the new domain. Using this approach different analyses can be systematically defined. Moreover, the proof of the correctness of the analysis can be done in a standard way.

In the paper we define a concrete instrumented operational semantics which handles, in addition to execution aspects, the level of the information flows of the program. The basis of the approach is that each value is annotated by a security level, representing the lub of the levels of the implicit and explicit flows on which the value depends. Also each channel is associated with a security level, representing the lub of the levels of the data present in the channel. The level of the input data is assumed to be that specified for the corresponding channel by the security policy. The level of data flowing through the variables and structures of the program is calculated dynamically taking into account the information flows. We define a collecting semantics which associates with each program point (instruction) the set of concrete states in which the machine can be when the point is reached.

We remark that the secure information flow property is defined in terms of independence of values produced for channels with a given security level from values taken from channels with higher security levels. A main result of the paper is the soundness of the concrete semantics. We prove that the program is secure if in all states of the collecting semantics the level of each channel is lower than or equal to that specified by the policy defined by the program. This theorem justifies the concrete semantics by relating it with the noninterference property. It shows that the concrete semantics correctly manipulates security annotations of data, that is data are associated with the security level on which they actually depend.

The abstract domains are obtained from the concrete ones by keeping the security levels and forgetting the actual values. A main point is the abstract domain of

references. A reference in the concrete domain holds the creation point (that is the new instruction) of the structure it refers to. This enables the analysis to distinguish between structures created at different instructions. In the abstract semantics, each new instruction represents a different abstract object, and, to perform aliasing control, each abstract reference holds the set of the all possible abstract objects it can refer to.

Starting from the concrete semantics, different abstractions can be performed to statically analyze information flow. Higher is the abstraction, higher is the efficiency of the analysis, but lower is the level of accuracy. The most popular approach to static secure information flow is based on *security typing*: each variables has a security annotation which can be seen as a part of its type and secure information flow is checked by means of type-inference: see, for example [1, 8, 22, 28, 33, 40, 41], while a recent survey is [39]. The advantage of these approaches is efficiency, since algorithms based on type inference can be used. Security typing can be modeled in our framework as the highest abstraction of the concrete semantics. In the paper we show a different abstraction, which is more precise than standard security typing, since it produces a more refined information. In fact, it is able to certify a larger class of programs with respect to typing approaches to security. A state of the abstract semantics is a table having a row for each instruction. Each row is the abstraction of all concrete states in which the machine can be when executing the corresponding instruction. The program is safe if in all rows of the abstract semantics the level of each channel is lower than or equal to that specified by the policy defined by the program. A main difference between this approach and security typing is that we do not assign a security level to each variable, but to each pair variable-instruction. The security type computed by the abstract semantics for a variable x and instruction t is the least upper bound of the secrecy levels of the information flows on which x depends when t is executed in any possible computation. Hence, we allow a register to hold data with a different secrecy level for different instructions of the program, while in security typing a same security type is computed for each variable and must hold for the whole program. We can call our approach *instruction-level security typing*, in that it infers for each variable a security type for each instruction of the program. Instruction-level security typing is inspired by bytecode verification [31], applied to Java bytecode, which is the intermediate code produced by Java compilers. The bytecode verifier, in order to accomplish safety checks on the code, performs an instruction-level typing algorithm: for each instruction it infers a type for every register and stack element. The instruction-level security typing can be performed

by an efficient fixpoint iteration algorithm, similar to that used by bytecode verification. In the paper we also present a tool implementing the method. The instruction-level security typing was firstly sketched in [3,16].

Section 2 describes some scenarios related to information flow checking, Sect. 3 introduces the languages and the definition of secure information flow. Section 4 defines the concrete semantics. Section 5 describes the abstract domains, while Sect. 6 defines the abstract semantics. In Sect. 7 a prototype tool that implements the proposed analysis is shown, together with some examples. Finally, Sect. 8 discusses related work and concludes.

2 Motivations

In this section we describe some scenarios to better clarify how tracing information flows can be essential to preserve confidentiality.

Example 1 (Trojan horse) Most commercial operating systems implement DAC mechanisms to ensure confidentiality. Users are arranged in groups, processes runs on behalf of the users and can access the resources depending on a set of permissions. There can be different ways of accessing resources (for instance *read*, *write*, *execute*). In UNIX-like systems, the permissions for the files can be viewed through the `ls` command. Suppose for instance that in a directory there are two files, one called `secreing.gpg` and owned by the administrator (`root` user) and the other called `myfile`, owned by the unprivileged user `john`. The former file can be read only by its owner and contains confidential information, while the latter can be read and written by every user.

```
root@myhost[mydir]# ls -l
total 8192
-rw-rw-rw-  2 john users      4096 2005-09-14 11:43 myfile
-rw-----  2 root wheel     4096 2005-11-25 12:01 secreing.gpg
```

Now, the administrator can redirect the content of the secret file into the public file, for instance using the well-known command `cat`, thus allowing unprivileged users to access the confidential information:¹

```
root@myhost[mydir]# cat secreing.gpg >>
myfile
```

One may argue that this is a too naive example, because: (i) the administrator should be able to downgrade the privacy level of data (ii) the administrator executes the `cat` command of his own will and explicitly. Although there is sense in these objections, the problem

¹ The `cat` process is able to read the `secreing.gpg` file because runs on behalf of the user `root`.




Fig. 1 A Java Card that hosts three applets. The applet B can disclose some A's data to applet C, without violating the firewall policy

is that the administrator could disclose the information without knowing it, because this action is “buried” into a program that he considers useful and harmless. This program could be either erroneous or malicious (in this latter case the term *trojan horse* is used). The administrator should be conscious that a program can downgrade information before executing it, to properly decide if this downgrading is acceptable or not.

Example 2 (Java Card Firewall) A Java Card [12] is a smart card running a Java Virtual Machine, the Java Card Virtual Machine (JCVM), and it is becoming a secure token in various fields, such as banking and public administration. The Java Card system was designed to speed up the development of applications (applets) and to increase portability. The JCVM is single-threaded, but more than one applet can coexist on the same card. Applets are normally isolated through the Java Card Firewall mechanism. This firewall allows an applet to access external objects only through an object sharing mechanism, called shareable interface. The firewall is based on an access control policy and therefore does not control information propagation.

Consider now a card that hosts three applets, say A, B, and C, each issued by a different commercial entity (see Fig. 1). Suppose that some partnership exists such that A must share some data with B and likewise B with C in order to be able to cooperate. Therefore the applets must be programmed to inform the firewall of this cooperation. Even if the applet A does not want that its confidential data would be propagated to C, it cannot rely on the firewall to control the behavior of B. Instead, before initiating cooperation with B, it must trust B's code.

Example 3 (A tax calculation) This example is similar to the previous, but belongs to a different context. Imagine that an user would write out his/her tax return electronically, using a computer application. This application is




Fig. 2 An application that calculates electronically tax returns. How the user can be sure that his confidential data are not disclosed to third parties?

a server that holds users data, and, upon the requests of the user, provides them the tax rate they have to apply to their income. The application must contact the Internet to know the current income thresholds and the current tax rates, as illustrated in Fig. 2.

The question is: how can the user trust that the program do not leak any information about his/her income to the party that provides the thresholds and the rates? An access control mechanism could only either block the entire communication or allow the data to flow without any checking. Thus, a more fine-grained control is necessary.

3 The model

We consider the simple language illustrated in Fig. 3. Besides basic data, the language handles dynamic structures. We indicate with k a literal value and with s, f, x, a , respectively, generic structure, field, variable and channel name. E represents the expressions and C the commands. Each instruction is labeled by a label $t \in \mathcal{B} = \{0, 1, \dots, n-1\}$, where n is the number of instructions in the program. We denote by New the subset of the new instructions in \mathcal{B} .

Every program P can retrieve data from a set of input channels and can send data to a set of output channels. If a is an input channel, the input command $a?x$ takes an item from a and assign it to variable x . The output command $a!e$ sends the value of expression e over the output channel a , provided that e is an expression returning a basic type (int). Given an input channel a , execution of an input command on a produces an *input action* (k, a)

```

P ::= {D; C}
D ::= T x | in σ a | out σ a | D ; D
T ::= int | S
S ::= struct s {D}
C ::= t : x = E | t : x.f = E | t : x = new s | t : a?x | t : a!E | t : skip
      | t : if(E) C else C; | t : while(E) C; | C; C
E ::= k | E Op E | x | x.f
Op ::= + | - | && | == | >

```

Fig. 3 Language grammar

where k is the value taken from a . Analogously, the execution of an output command on an output channel a produces an *output action* (k, a) where k is the value inserted into a . In the following, we denote as Names_I (respectively, Names_O) the set of input (output) channels used by a program; moreover $\text{Names} = \text{Names}_I \cup \text{Names}_O$ and $\text{Names}_I \cap \text{Names}_O = \emptyset$.

We assume that programs are type correct, that is:

- in any expression $x.f$, the variable x is a reference to structure with a field of name f ,
- in any expression $E_1 \text{ Op } E_2$, E_1 and E_2 are of the same type,
- in any assignment $x = E$, the expression E is of the same type of the variable x ,
- in any assignment $x.f = E$, the variable x is a reference to structure with a field of name f , and the expression E is of the same type of the field f ,
- in any output statement $a!E$, the expression E is of type int ,
- in any input statement $a?x$, the variable x is of type int .

These constraints can be easily checked by standard type checking algorithms.

The input and output channels represent the external environment in which the program is executed, that is all the interactions of the program occur by means of the channels. An external server is not able to inspect the internal state of the program, but can only send/receive messages to/from the program by means of the input/output channels. A *security policy* assigns to each input and output channel a security level, representing a fixed degree of secrecy. The security policy is expressed by the declaration of the channels. A channel a is declared by using the keyword in (out) to indicate that is an input (output) channel and by indicating also its security level. Security levels are defined as a finite lattice $(\mathcal{L}, \sqsubseteq_{\mathcal{L}})$, ranged over by σ, τ, \dots and partially ordered by $\sqsubseteq_{\mathcal{L}}$. In the following we indicate by $S : \text{Names} \rightarrow \mathcal{L}$ the security policy specified by the channels declarations.

Definition 1 (*secure information flow*) Let P be a program and S a security policy for P . Given $\sigma \in \mathcal{L}$, let us denote by $\text{Names}^{\sqsubseteq \sigma} = \{a \in \text{Names} \mid S(a) \sqsubseteq \sigma\}$ the set of channels with security level lower than or equal to σ . Given a computation C of the program, let $\text{actions}(C, \sigma)$ denote the sequence of input and output actions involving the channels in $\text{Names}^{\sqsubseteq \sigma}$ performed by C . We say that P has σ -secure information flow (is σ -secure) under S if, assuming that the external environment sends to the program the same sequence of values on each input channel in $\text{Names}^{\sqsubseteq \sigma}$, for any

<p>P1</p> <pre>1: a?x; 2: b!x;</pre>	<p>P2</p> <pre>1: y=1; 2: a?x; 3: if (x==0) 4: y=0; else 5: skip; 6: b!y;</pre>	<p>P3</p> <pre>1: a?x; 2: if (x==0) 3: d?y; else 4: skip;</pre>	<p>P4</p> <pre>1: a?x; 2: while (x>0){ 3: b!1; 4: x=x-1; }</pre>
<p>P5</p> <pre>1: d?x; 2: while (x>0){ 3: b!1; 4: a?x; }</pre>	<p>P6</p> <pre>1: a?x; 2: while (x>0){ 3: x=x-1; } 4: b!1;</pre>	<p>P7</p> <pre>1: a?y; 2: y=0; 3: b!y;</pre>	<p>P8</p> <pre>1: s1=new S; 2: s2=new S; 3: a?x; 4: if (x) 5: s3=s1; else 6: s3=s2; 7: s3.f=1;</pre>

Fig. 4 Some examples

two computations C_1 and C_2 of the program, it holds $actions(C_1, \sigma) = actions(C_2, \sigma)$.

We say that P has secure information flow (is secure) if it is σ -secure for each $\sigma \in \mathcal{L}$.

If a program is secure, an external attacker having secrecy level σ cannot infer information that is more secret than σ from a σ -secure program since he can inspect only input and output channels with levels that are lower than or equal to σ .

Similar properties are also referred in literature as *noninterference* properties.

The above definition considers computations that may or may not terminate, and therefore the property is *termination-sensitive* [40]. A weaker termination-insensitive property could consider only finite computations.

Note that, in general an attacker could gather secure information by observing the so-called *covert channels* [29], e.g. the time occurring between the input/output operations, or the power consumption. We do not consider this kind of noninterference.

Let us show some examples of programs. Consider the programs in Fig. 4 and suppose that a and d are input channels and b is an output channel. Moreover assume $S(a) = H, S(b) = S(d) = L$, with $L \sqsubset H$. Since in this example there are only two security levels, we can say that channels b and d are public, while channel a is private.

Program P1 shows an explicit insecure information flow, since the value output on channel b depends on the value input from a : private information is made available to a public observer. Program P2 is insecure because it is possible to know if the private input is zero by observing the value present on the public output channel. In program P3 the private value affects the contents of input channel d , from which an item is taken only if

the input is zero. Note that we consider observable both the input and the output channels. In program P4 the number of the values output on channel b depends on the input value. In program P5 the first iteration of the `while` is driven by a low value, while the following iterations depend on high level information. Also program P6 may have an illicit information flow, even though the value output on channel b is always the same: it is possible that, due to an infinite loop, no value is output on channel b . Program P7 is secure, since the output value, which is constant, does not depend on the input: even though y is written with a high value, afterward it is assigned a constant value, and this one is given as an output. Consider program P8 and suppose that S is an user-defined structure with two int fields f and g , and that $s1, s2, s3$ are references of type S . Please notice that, depending on the value taken from the high level input channel a , instruction 7 updates field f of two different objects (created at the first two instructions). Now consider the two cases in which instruction 7 is followed by: (i) $8: b!s1.g$; (ii) $8: b!s1.f$; . In case (i) the program is secure because field g of object created at instruction 1 is the same in any computation. On the contrary, in case (ii), the value of the field $s1.f$ depends on the input: by aliasing, the assignment in instruction 7 could have modified it.

4 Concrete semantics

In this section we define the concrete semantics of the language. To take into account the security level of data, we annotate each value v flowing through the variables and the structure fields with a security level, representing the least upper bound of the security levels of the explicit and implicit information flows on which v depends. A value is a pair (v^e, σ) , where v^e is an execution value and σ a security level. The domains of the concrete semantics are shown in Fig. 5. An execution value may be an integer $k \in \mathbb{Z}$ or a reference to an user-defined structure. A reference is in turn a pair (ℓ, t) , where $\ell \in \mathcal{A}^e$ is a heap address and $t \in \text{New}$ is the label of the instruction which created the corresponding structure. This tag will be useful in the abstraction to coalesce into a same abstract structure all structures created at the same instruction. The memory is represented by means of two functions: one denoted by μ , that

$$\begin{aligned}
 v \in \mathcal{V} &= (\mathbb{Z} \cup \mathcal{A}) \times \mathcal{L} & \mathcal{A} &= \mathcal{A}^e \times \text{New} \\
 \mu \in \mathcal{M} &= \text{Var} \rightarrow \mathcal{V} & \xi \in \Xi &= \mathcal{A} \rightarrow \mathcal{M}_{\text{struct}} \\
 c \in \mathcal{C} &= \text{Names} \rightarrow (\mathbb{Z}^* \times \mathcal{L}) & q \in \mathcal{Q} &= \mathcal{B} \times \text{Env} \times \mathcal{M} \times \Xi \times \mathcal{C}
 \end{aligned}$$

Fig. 5 Domains of the concrete semantics

$$\text{Const} \frac{}{\langle k, \mu, \xi \rangle \xrightarrow{E} \langle k, \perp_{\mathcal{L}} \rangle} \quad \text{Op} \frac{\langle E_1, \mu, \xi \rangle \xrightarrow{E} \langle v_1^e, \sigma_1 \rangle, \langle E_2, \mu, \xi \rangle \xrightarrow{E} \langle v_2^e, \sigma_2 \rangle}{\langle E_1 \text{ op } E_2, \mu, \xi \rangle \xrightarrow{E} \langle v_1^e \text{ op } v_2^e, \sigma_1 \sqcup_{\mathcal{L}} \sigma_2 \rangle}$$

$$\text{Value}_x \frac{}{\langle x, \mu, \xi \rangle \xrightarrow{E} \mu(x)} \quad \text{Value}_{x.f} \frac{\mu(x) = ((\ell, t), \sigma_1), \xi(\ell, t)(f) = (v^e, \sigma_2)}{\langle x.f, \mu, \xi \rangle \xrightarrow{E} \langle v^e, \sigma_1 \sqcup_{\mathcal{L}} \sigma_2 \rangle}$$

Fig. 6 Concrete semantics of expressions

associates every variable with its value, and the other, denoted by ξ , that associates the addresses (references) with the respective structure instances. Every structure in the heap can be represented by a memory whose variables are the fields. Valid fields names are in the domain \mathcal{F} . We denote by \mathcal{M}_S the domain of memories having the fields of structure S as variables, and by $\mathcal{M}_{\text{struct}}$ the set: $\mathcal{M}_{\text{struct}} = \bigcup \{ \mathcal{M}_S \mid S \text{ used in } P \}$. The state of input and output channels $c \in \mathcal{C}$ is a mapping from the names of the channels to pairs (s, σ) , where $s \in \mathbb{Z}^*$ is a finite sequence of values and σ a security level. We use \cdot to denote sequence concatenation, and λ to denote the empty sequence. Initially, the security level of each input channel a is set to $S(a)$, that is the security level defined for a by the security specification. As a consequence, each value taken from an input channel a is annotated with $S(a)$. The security level of the output channels is initially set to the minimum level $\perp_{\mathcal{L}}$. The security level of the channels can be modified by the computation, when the channel is accessed.

The concrete semantics is defined by means of a set of rules: the rules for expressions are shown in Fig. 6 and the rules for instructions in Fig. 7. Let us consider the rules for expressions, defining a relation $\xrightarrow{E} \subseteq (\text{expr} \times \mathcal{M} \times \Xi) \times \mathcal{V}$. Rule **Const** assigns the bottom security level to any constant value. Rule **Op** calculates the security level of the result of an operation as the least upper bound (from now on, lub) of the security levels of the operands. Rule **Value_x** returns the value of the variable in the memory. Rule **Value_{x.f}** annotates the resulting value with the lub of the security levels of the reference and of the value stored in the field.

The rules for instructions (Fig. 7) define a relation $\longrightarrow \subseteq \mathcal{Q} \times \mathcal{Q}$ between the states of the computation. The set of concrete states is $\mathcal{Q} = \mathcal{B} \times \text{Env} \times \mathcal{M} \times \Xi \times \mathcal{C}$, where $\text{Env} = \mathcal{B} \rightarrow \mathcal{L}$. Each state $q \in \mathcal{Q}$ is a tuple $\langle t, \rho, \mu, \xi, c \rangle$ describing the configuration of the machine when executing the command t : μ and ξ define the values of variables and structures fields, while c represent the status of the channels. We also keep in each state a *security environment* $\rho \in \text{Env}$. The environment ρ assigns to every program point a security level representing the level of the implicit flow under which the corresponding command is executed. In the following, given an instruction label t and a set \mathcal{Q} of states, we use the notation $\mathcal{Q}(t)$ to denote the set of states in \mathcal{Q} corresponding to instruction t . A value (v^e, τ) evaluated, assigned or tested while the execution is under a security environment σ , changes its security level into $\sigma \sqcup \tau$. The

Fig. 7 Concrete semantics of commands

$$\text{Skip} \frac{}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}(t), \rho, \mu, \xi, c \rangle}$$

$$\text{Assign}_{t:x=E} \frac{\langle E, \mu, \xi \rangle \xrightarrow{E} \langle v^e, \sigma \rangle}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}(t), \rho, \mu[x \leftarrow (v^e, \rho(t)) \sqcup_{\mathcal{L}} \sigma], \xi, c \rangle}$$

$$\text{Assign}_{t:x.f=E} \frac{\langle E, \mu, \xi \rangle \xrightarrow{E} \langle v^e, \sigma_1 \rangle, \mu(x) = ((\ell, t_1), \sigma_2), \sigma_3 = \sigma_1 \sqcup_{\mathcal{L}} \sigma_2 \sqcup_{\mathcal{L}} \rho(t)}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}(t), \rho, \mu, \xi[(\ell, t_1).f \leftarrow (v^e, \sigma_3)], c \rangle}$$

$$\text{New}_{t:x=new S} \frac{\text{fresh}(\xi) = \ell}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}(t), \rho, \mu[x \leftarrow ((\ell, t), \rho(t))], \xi[(\ell, t) \leftarrow \mu_{S\perp}], c \rangle}$$

$$\text{Input}_{t:a?x} \frac{c(a) = (k \cdot s, \sigma), a \in \text{Names}_I}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}(t), \rho, \mu[x \leftarrow (k, \rho(t)) \sqcup_{\mathcal{L}} \sigma], \xi, c[a \leftarrow (s, \rho(t)) \sqcup_{\mathcal{L}} \sigma] \rangle}$$

$$\text{Output}_{t:b!E} \frac{\langle E, \mu, \xi \rangle \xrightarrow{E} \langle k, \sigma_1 \rangle, c(b) = (s, \sigma_2), b \in \text{Names}_O}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}(t), \rho, \mu, \xi, c[b \leftarrow (s \cdot k, \rho(t)) \sqcup_{\mathcal{L}} \sigma_1 \sqcup_{\mathcal{L}} \sigma_2] \rangle}$$

$$\text{If}_{t:\text{if } (E) \text{ C} \text{ else } C, (\text{true})} \frac{\langle E, \mu, \xi \rangle \xrightarrow{E} \langle \text{true}, \sigma \rangle}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}_{\text{true}}(t), \rho[t' \leftarrow \rho(t')] \sqcup_{\mathcal{L}} \sigma \rangle_{\forall t' \in \text{scope}(t)}, \mu, \xi, c \rangle}$$

$$\text{While}_{t:\text{while } (E) \text{ C} \text{ (true)}} \frac{\langle E, \mu, \xi \rangle \xrightarrow{E} \langle \text{true}, \sigma \rangle}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}_{\text{true}}(t), \rho[t' \leftarrow \rho(t')] \sqcup_{\mathcal{L}} \sigma \rangle_{\forall t' \in \text{scope}(t)}, \mu, \xi, c \rangle}$$

environment, initially set to $\perp_{\mathcal{L}}$ for all commands, can be updated by the conditional and repetitive commands. With $succ(t)$ we indicate the successive instruction to be executed. All commands have only one successor, except the conditional and repetitive commands that have two successors, depending on the value of the guard; they are denoted by $succ_{true}(t)$ and $succ_{false}(t)$. We assume that the first instruction of the program has label t_0 and that for the last instruction is $succ(t) = \text{end}$.

$$scope(t) = \begin{cases} \bigcup \{scope(t') \mid t' \in C_1 \cup C_2\} & \text{if } t: \text{if}(E) \{C_1\} \text{ else } \{C_2\} \\ \bigcup \{scope(t') \mid t' \in C_1 \cup C_2\} & \text{if } t: \text{while}(E) \{C_1\} C_2 \\ \{t\} & \text{otherwise} \end{cases}$$

Rule **Assign** _{$t,x=e$} annotates the security level of the value to be assigned with the lub of the security level resulted by the evaluation of the expression and the environment of the instruction t . The notation $\mu[x \leftarrow (v^e, \sigma)]$ stands for the memory obtained by μ by updating the contents for the variable x with the value (v^e, σ) . Rule **Assign** _{$t,x,f=e$} annotates the value to be assigned with the lub of (1) the security level resulted by the evaluation of the expression, (2) the security level of the reference, and (3) the environment of t . In the rule, the notation $\xi[(\ell, t).f \leftarrow v]$ indicates the heap ξ' obtained from ξ by updating the field f of the structure located at address ℓ (and created at instruction t) with the value v .

Rule **New** contains the notation $\xi[(\ell, t) \leftarrow \mu_{S\perp}]$, meaning that, during the execution of instruction t , in the heap ξ a new structure of type S is created at address ℓ , its fields containing the default value. We assume the default value is the pair $(0, \perp_{\mathcal{L}})$. In the premise of the rule the function $\text{fresh} : \Xi \rightarrow \mathcal{A}^e$ is used to find a free location in the heap to store the new structure.

Rule **Input** takes a value from the specified input channel and assigns it to the destination variable, annotated with the lub of the level σ of the channel and the environment of t . Also the level of the channel is updated in the same way. As a consequence, if $\rho(t)$ is higher than σ , the level of the channel is upgraded, to record the fact that the manipulation of the channel depends on an information flow with level $\sigma \sqsubseteq_{\mathcal{L}} \rho(t)$. Analogously, in the **Output** rule, the level of the specified output channel can be upgraded taking into account the level of the value and that of the environment of the instruction.

The **If** and **While** rules, whatever branch is chosen, affect the environment of all the instructions belonging to the scope of the command, taking into account the level of the condition. The set $scope(t)$ contains all the instructions that can be executed or not depending on the condition. In the **If** case, $scope(t)$ includes all

the instructions belonging to only one branch starting from the **if**. For the **while** command, $scope(t)$ includes the instructions belonging to the loop (the *true* part of the **while**) and also all instructions after the loop until the end of the program (the *false* part). The inclusion of these instructions takes into account the possibility of an infinite loop: in this case, the commands following the loop will never be executed.

The function $scope: \mathcal{B} \rightarrow \wp(\mathcal{B})$ is defined as follows:

where C_1 and C_2 are sequences of instructions. Note that we adopt a termination-sensitive noninterference property. A weaker termination-insensitive property, which is defined only for terminating executions, can be analyzed by using a different definition of $scope(t)$, obtained by modifying the second clause of the above definition. The clause becomes $\bigcup \{scope(t') \mid t' \in C\}$ if $t: \text{while}(E) \{C\}$. In this way we consider as scope of a **while** command only the body of the command and not the continuation of the program. Hence, the level of data manipulated after the **while** is not affected by the level of the condition.

Updating the environment is necessary to trace implicit flow: the value of the condition (with its security level) drives the execution of the instructions in $scope(t)$. The table shows only the rule to be applied when the condition is true. The rule to be applied when the condition is false (not shown) is equal except that has $succ_{false}$ instead of $succ_{true}$.

Definition 2 (initial state) Given an initial configuration $i_0 : \text{Names}_I \rightarrow \mathbb{Z}^*$ of the input channels, the initial state is defined as $q(i_0) = \langle t_0, \rho_{\perp}, \mu_0, \xi_{\lambda}, c_0 \rangle$, where $\forall t \in \mathcal{B}, \rho_{\perp}(t) = \perp_{\mathcal{L}}$, and $\forall x \in X, \mu_0(x) = (0, \perp_{\mathcal{L}})$, and ξ_{λ} is the heap with empty domain (i.e. the everywhere undefined function). The state c_0 is such that for all $a \in \text{Names}_I, c_0(a) = (i_0(a), \mathcal{S}(a))$ and for all $a \in \text{Names}_O, c_0(a) = (\lambda, \perp_{\mathcal{L}})$.

We now define a collecting semantics, which associates with each instruction the set of states in which the instruction can be executed in any computation.

First we define an alignment operation $\text{align}(Q)$ which, given a set of states Q , aligns all the states corresponding to the same instruction. $\text{align}(Q)$ increments Q with some extra states for each instruction t : for each state $q \in Q(t)$, a state q' is added to Q having the same execution values occurring in q , but where the security levels of the environment, memory variables, fields of

Fig. 8 Auxiliary functions for merging

$$\begin{array}{llll}
\max_E : \mathcal{Q} \times \mathcal{B} \rightarrow \mathcal{L} & \max_E(Q, t) & = \bigsqcup_{\mathcal{L}} \{\rho(t) \mid \langle t', \rho, \mu, \xi, c \rangle \in Q\} \\
\max_{\mathcal{M}} : \mathcal{Q} \times \text{Var} \rightarrow \mathcal{L} & \max_{\mathcal{M}}(Q, x) & = \bigsqcup_{\mathcal{L}} \{\sigma \mid \langle t, \rho, \mu, \xi, c \rangle \in Q, \mu(x) = (v^e, \sigma)\} \\
\max_{\Xi} : \mathcal{Q} \times \mathcal{A} \times \mathcal{F} \rightarrow \mathcal{L} & \max_{\Xi}(Q, (\ell, t), f) & = \bigsqcup_{\mathcal{L}} \{\sigma \mid \langle t, \rho, \mu, \xi, c \rangle \in Q, \xi(\ell, t)(f) = (v^e, \sigma)\} \\
\max_C : \mathcal{Q} \times \text{Names} \rightarrow \mathcal{L} & \max_C(Q, a) & = \bigsqcup_{\mathcal{L}} \{\sigma \mid \langle t, \rho, \mu, \xi, c \rangle \in Q, c(a) = (s, \sigma)\}
\end{array}$$

Fig. 9 The align function

$$\begin{array}{l}
\text{align}(\langle t, \rho, \mu, \xi, c \rangle, Q) = \langle t, \rho', \mu', \xi', c' \rangle \\
\text{with:} \\
\forall t' \in \mathcal{B} : \rho'(t') = \max_E(Q(t), t') \\
\forall x \in \text{Var} : \mu'(x) = \text{up}(\mu(x), \max_{\mathcal{M}}(Q(t), x)) \\
\forall (\ell, t') \in \text{dom}(\xi), f \in \text{dom}(\xi(\ell, t')) : \xi'(\ell, t')(f) = \text{up}(\xi(\ell, t')(f), \max_{\Xi}(Q(t), \ell, t', f)) \\
\forall a \in \text{Names} : c'(a) = \text{up}(c(a), \max_C(Q, a)) \\
\text{align}(Q) = (\bigcup_{q \in Q} \text{align}(q, Q)) \cup Q
\end{array}$$

structures and channels are upgraded to the lub in \mathcal{L} of the levels occurring in the states in $Q(t)$ for the same items.

In Fig. 8 some auxiliary functions used in the alignment process are shown. Let Q be a set of states: then $\max_{\mathcal{M}}(Q, x)$ is the lub of the security levels of x in the memories occurring in the states of Q . For each $t \in \mathcal{B}$, $\max_E(Q, t)$ is the lub of the values of $\rho(t)$ in the environment occurring in the states of Q . For each field f of each structure created at instruction $(\ell, t) \in \mathcal{A}$, $\max_{\Xi}(Q, \ell, t, f)$ is the lub of the values held by the field f in the heap occurring in the states of Q . For each channel $a \in \text{Names}$, $\max_C(Q, a)$ is the lub of the security levels held by the channel a in the states of Q .

The definition of **align** is shown in Fig. 9. Given a value $v = (v^e, \tau)$, with $v^e \in (\mathbb{Z} \cup \mathcal{A})$, $\text{up}(v, \sigma) = (v^e, \tau \sqcup_{\mathcal{L}} \sigma)$ is the value obtained by maintaining the execution part of the value and upgrading the annotation of v .

For example, consider to have a memory μ with only one variable x and only a channel a and consider the states:

$$\begin{array}{l}
q_1 = \langle 2, \rho, \mu(x) = (2, L), \xi, c(a) = (1 \cdot 2, H) \rangle \\
q_2 = \langle 2, \rho, \mu(x) = (3, H), \xi, c(a) = (2 \cdot 3, L) \rangle \\
q_3 = \langle 3, \rho, \mu(x) = (4, L), \xi, c(a) = (3, L) \rangle
\end{array}$$

We have that:

$$\text{align}(\{q_1, q_2, q_3\}) = \{q_1, q_2, q_3, q_4, q_5\}$$

where:

$$\begin{array}{l}
q_4 = \langle 2, \rho, \mu(x) = (2, H), \xi, c(a) = (1 \cdot 2, H) \rangle \text{ and} \\
q_5 = \langle 2, \rho, \mu(x) = (3, H), \xi, c(a) = (2 \cdot 3, H) \rangle
\end{array}$$

State q_4 is equal to q_1 with respect to numeric values, but variable x contains an high value, which is the maximum value of x in the set of states with program counter 2. Analogously, state q_5 has the same values of q_2 but aligns the level of channel a to H . State q_3 is the only one with program counter 3 and thus no new state needs to be added.

Definition 3 (*concrete next operator next*) Given a set of concrete states $Q \subseteq \mathcal{Q}$, the application of the **next** operator yields the aligned set of states that are either in Q , or reached in one step of computation starting from a state in Q .

$$\text{next}(Q) = \text{align}(Q \cup \{q' \mid \exists q \in Q : q \longrightarrow q'\})$$

Proposition 1 (*monotonicity of next*) *next is monotone in* $(\wp(Q), \subseteq)$.

The concrete collecting semantics $\text{sem} \in \wp(\mathcal{Q})$ is the set of all aligned concrete states belonging to all executions.

Definition 4 (*collecting semantics*) The concrete collecting semantics $\text{sem} \in \wp(\mathcal{Q})$ is the lub of the following increasing chain, defined for all $n \in \mathbb{N}$:

$$\begin{array}{l}
\text{sem}_0 = \{q(i_0) \mid \forall i_0 \in (\text{Names}_I \rightarrow \mathbb{Z}^*)\} \\
\text{sem}_{n+1} = \text{next}(\text{sem}_n)
\end{array}$$

Performing **align** at each step of sem_n aligns the security annotations of the states corresponding to the join point of different branches of a conditional instruction, in order to properly manage implicit flows. This function is necessary for the soundness of the concrete semantics, proved by Theorem 1. Consider, for example, program P2 of Fig. 4. If we consider an execution in which the input value is 0, the branch *true* of the if command is executed, and at instruction 6 the state is $q = \langle 6, \rho, \mu, \xi_\lambda, c \rangle$, with $\rho(4) = \rho(5) = H$, $\mu(y) = (0, H)$ and where the annotation H of 0 records the implicit flow of level H under which the assignment to y has been performed. On the other hand, if the input value is different from 0, variable x is not affected in the conditional command and the state $q' = \langle 6, \rho, \mu', \xi_\lambda, c' \rangle$ is reached, where $\mu'(y) = (1, \perp_{\mathcal{L}})$. This state does not represent the implicit flow, since the level of the value held by y is low. Instead, the contents of y has been affected also in this case by the implicit flow of level H . Since the alignment operation is applied to the chain sem_n , there exists at least one j such that sem_j contains a state $\langle 6, \rho, \mu'', \xi_\lambda, c' \rangle$ where $\mu''(y) = (1, H)$. This state derives from the alignment of

q to the levels occurring in q' and represents the effect of the implicit flow on y in the case in which the *false* branch has been chosen.

The security annotations used by the concrete semantics can be tied to the secure information flow property via a soundness theorem. To prove such theorem we need to define an equivalence relation between states.

Definition 5 (σ -equivalence) Let $\sigma \in \mathcal{L}$.

- Given two heaps $\xi_1, \xi_2 \in \Xi$, we define a relation $\equiv_{\sigma}^{\xi_1, \xi_2} \subseteq \mathcal{V} \times \mathcal{V}$ as the greatest relation satisfying:
 $v_1 \equiv_{\sigma}^{\xi_1, \xi_2} v_2$ if and only if one of the following cases holds:
 - $\sigma_1 \not\sqsubseteq_{\mathcal{L}} \sigma$ and $\sigma_2 \not\sqsubseteq_{\mathcal{L}} \sigma$
 - $\sigma_1 \sqsubseteq_{\mathcal{L}} \sigma, \sigma_2 \sqsubseteq_{\mathcal{L}} \sigma, v_1 = (k, \sigma_1), v_2 = (k, \sigma_2)$
 - $\sigma_1 \sqsubseteq_{\mathcal{L}} \sigma, \sigma_2 \sqsubseteq_{\mathcal{L}} \sigma, v_1 = (l_1, t_1, \sigma_1), v_2 = (l_2, t_2, \sigma_2)$
and $\forall f \in \mathcal{F} : \xi_1(l_1).f \equiv_{\sigma}^{\xi_1, \xi_2} \xi_2(l_2).f$
- Two environments ρ_1 and ρ_2 are σ -equivalent ($\rho_1 \equiv_{\sigma}^E \rho_2$) iff
 $\forall t \in \mathcal{B} : \text{either } \rho_1(t) \sqsubseteq_{\mathcal{L}} \sigma \text{ and } \rho_2(t) \sqsubseteq_{\mathcal{L}} \sigma \text{ or } \rho_1(t) \not\sqsubseteq_{\mathcal{L}} \sigma \text{ and } \rho_2(t) \not\sqsubseteq_{\mathcal{L}} \sigma$.
- Two concrete pairs memory-heap (μ_1, ξ_1) and (μ_2, ξ_2) are σ -equivalent ($(\mu_1, \xi_1) \equiv_{\sigma}^{\mathcal{M}, \Xi} (\mu_2, \xi_2)$) iff $\forall x : \mu_1(x) \equiv_{\sigma}^{\xi_1, \xi_2} \mu_2(x)$.
- Two concrete channels $a_1 = (\delta_1, \sigma_1)$ and $a_2 = (\delta_2, \sigma_2)$ are σ -equivalent ($a_1 \equiv_{\sigma}^C a_2$) if one of the following cases hold:
 - $\sigma_1 \sqsubseteq_{\mathcal{L}} \sigma, \sigma_2 \sqsubseteq_{\mathcal{L}} \sigma$ and $\delta_1 = \delta_2$
 - $\sigma_1 \not\sqsubseteq_{\mathcal{L}} \sigma$ and $\sigma_2 \not\sqsubseteq_{\mathcal{L}} \sigma$
- Two concrete states $q_1 = \langle t_1, \rho_1, \mu_1, \xi_1, c_1 \rangle$ and $q_2 = \langle t_2, \rho_2, \mu_2, \xi_2, c_2 \rangle$ are σ -equivalent ($q_1 \equiv_{\sigma}^Q q_2$) iff

$$t_1 = t_2 \quad \rho_1 \equiv_{\sigma}^E \rho_2 \quad (\mu_1, \xi_1) \equiv_{\sigma}^{\mathcal{M}, \Xi} (\mu_2, \xi_2) \\ \forall a \in \text{Names} : c_1(a) \equiv_{\sigma}^C c_2(a)$$

The relation $\equiv_{\sigma}^Q \in \mathcal{Q} \times \mathcal{Q}$ is an equivalence relation.

Lemma 1 Let be $\mu_1, \mu_2 \in \mathcal{M}$ and $\xi_1, \xi_2 \in \Xi$, such that $(\mu_1, \xi_1) \equiv_{\sigma}^{\mathcal{M}, \Xi} (\mu_2, \xi_2)$ and an expression $E \in \text{Expr}$. If

$$\langle E, \mu_1, \xi_1 \rangle \xrightarrow{E} v_1, \text{ and} \\ \langle E, \mu_2, \xi_2 \rangle \xrightarrow{E} v_2,$$

then,

$$v_1 \equiv_{\sigma}^{\xi_1, \xi_2} v_2$$

Proof By induction on the structure of expressions. \square

The following lemma states that a) the execution of a non-conditional instruction preserves σ -equivalence; and b) σ -equivalence is also preserved with a conditional instruction, provided that the tested condition is less secret than σ .

Lemma 2 Let be $q_1 = \langle t, \rho_1, \mu_1, \xi_1, c_1 \rangle \equiv_{\sigma}^Q q_2 = \langle t, \rho_2, \mu_2, \xi_2, c_2 \rangle$. If $q_1 \longrightarrow q'_1 = \langle t'_1, \rho'_1, \mu'_1, \xi'_1, c'_1 \rangle$ and $q_2 \longrightarrow q'_2 = \langle t'_2, \rho'_2, \mu'_2, \xi'_2, c'_2 \rangle$, then

$$\rho'_1 \equiv_{\sigma}^E \rho'_2 \quad (\mu_1, \xi_1) \equiv_{\sigma}^{\mathcal{M}, \Xi} (\mu_2, \xi_2) \quad \forall a \in \text{Names} : \\ c'_1(a) \equiv_{\sigma}^C c'_2(a)$$

Moreover, if

- t is neither an `if` nor a `while` instruction; or
- t is either an `if` or a `while` instruction and the guard E is such that

$$\langle E, \mu_1, \xi_1 \rangle \xrightarrow{E} (v_1^e, \sigma'_1) \text{ with } \sigma'_1 \sqsubseteq_{\mathcal{L}} \sigma$$

then it is also true that $t'_1 = t'_2$ (and therefore $q'_1 \equiv_{\sigma}^Q q'_2$)

Proof By examining all possible kinds of instruction.

(Assign_{t;x=e}) The rule affects only the memory. By Lemma 1, we have that the expression E , once evaluated, returns two σ -equivalent values, let say $v_1 = (v_1^e, \sigma_1)$ and $v_2 = (v_2^e, \sigma_2)$ (we suppose for simplicity that v_1 and v_2 are simple values and not references). Moreover, suppose $\rho_1(t) = \tau_1$ and $\rho_2(t) = \tau_2$. The memories μ'_1 and μ'_2 differ from μ_1 and μ_2 only for the values $\mu'_1(x)$ and $\mu'_2(x)$ that are updated with $v'_1 = (v_1^e, \sigma_1 \sqcup_{\mathcal{L}} \tau_1)$ and $v'_2 = (v_2^e, \sigma_2 \sqcup_{\mathcal{L}} \tau_2)$, respectively. Now, we have two cases:

- $\sigma_1 \sqsubseteq_{\mathcal{L}} \sigma$ and $\sigma_2 \sqsubseteq_{\mathcal{L}} \sigma, \tau_1 \sqsubseteq_{\mathcal{L}} \sigma$ and $\tau_2 \sqsubseteq_{\mathcal{L}} \sigma$
Since $v_1 \equiv_{\sigma}^{\xi_1, \xi_2} v_2$ and $(\sigma_1 \sqcup_{\mathcal{L}} \tau_1) \sqsubseteq_{\mathcal{L}} \sigma$ and $(\sigma_2 \sqcup_{\mathcal{L}} \tau_2) \sqsubseteq_{\mathcal{L}} \sigma$. Therefore, also $v'_1 \equiv_{\sigma}^{\xi_1, \xi_2} v'_2$.
- $(\sigma_1$ and $\sigma_2 \not\sqsubseteq_{\mathcal{L}} \sigma)$ or $(\tau_1$ and $\tau_2 \not\sqsubseteq_{\mathcal{L}} \sigma)$ Then $(\sigma_1 \sqcup_{\mathcal{L}} \tau_1) \not\sqsubseteq_{\mathcal{L}} \sigma$ and $(\sigma_2 \sqcup_{\mathcal{L}} \tau_2) \not\sqsubseteq_{\mathcal{L}} \sigma$. Therefore, $v'_1 \equiv_{\sigma}^{\xi_1, \xi_2} v'_2$.

Then, it holds that $(\mu'_1, \xi_1) \equiv_{\sigma}^{\mathcal{M}, \Xi} (\mu'_2, \xi_2)$.

(Assign_{t;x,f=e}) This rule affects only the heap. Again, since by Lemma 1 the expression E returns two σ -equivalent values and $\rho_1 \equiv_{\sigma}^E \rho_2$, we can make a reasoning similar to the previous case.

(New_{t;x=new s}) Here, both the memory and the heap change. Suppose that $\text{fresh}(\xi_1) = \ell_1$ and $\text{fresh}(\xi_2) = \ell_2$. It holds that $\forall y \in \text{Var}, y \neq x, \mu'_1(y) = \mu_1(y) \equiv_{\sigma}^{\xi_1, \xi_2} \mu_2(y) = \mu'_2(y)$. Then $\mu'_1(y) \equiv_{\sigma}^{\xi_1, \xi_2} \mu'_2(y)$. Since the function `fresh` returns two clean addresses, it holds also that $\mu'_1(y) \equiv_{\sigma}^{\xi'_1, \xi'_2} \mu'_2(y)$. Moreover $\mu'_1(x) = (\ell_1, t, \rho_1(t)) \equiv_{\sigma}^{\xi_1, \xi_2} \mu'_2(y) = (\ell_2, t, \rho_2(t))$ since the two new instances are equal and $\rho_1 \equiv_{\sigma}^E \rho_2$.

(Input_{t;a?x}) This rule modifies the channel a and the variable x . However, the thesis trivially derives from the σ -equivalence of memories, environments and channels before the input operation.

(Output_{t;b?E}) When applying this rule, only the channel b is modified. Again, the thesis holds from the σ -equivalence of the environments and of the two results of the expression E (by Lemma 1)

(If, While) Here, only the program counter and the environment can change. Let $\langle E, \mu_1, \xi_1 \rangle \xrightarrow{E} (k_1, \sigma'_1)$ and $\langle E, \mu_2, \xi_2 \rangle \xrightarrow{E} (k_2, \sigma'_2)$ be the two transitions that happen when evaluating the guards. By Lemma 1, only two cases may apply:

$\sigma'_1 \sqsubseteq_{\mathcal{L}} \sigma$ **and** $\sigma'_2 \sqsubseteq_{\mathcal{L}} \sigma$ Then, it holds that the execution values of the guards are the same ($k_1 = k_2$) and therefore $t'_1 = t'_2$. Moreover the value of the environment for instructions in the scope may be updated, in both cases, with a level $\sqsubseteq_{\mathcal{L}} \sigma$, thus preserving σ -equivalence.

$\sigma'_1 \not\sqsubseteq_{\mathcal{L}} \sigma$ **and** $\sigma'_2 \not\sqsubseteq_{\mathcal{L}} \sigma$ The value of the environment for instructions in the scope is upgraded, in both cases, with a level $\not\sqsubseteq_{\mathcal{L}} \sigma$, thus preserving σ -equivalence. \square

The following theorem states the soundness of the collecting semantics. It proves that the management of data performed by the concrete semantics correctly represent the level of the information flows of the program.

Theorem 1 (soundness of the concrete semantics) *A program P has secure information flow under a security policy \mathcal{S} if for each concrete state $\langle t, \rho, \mu, \xi, c \rangle \in \text{sem}$, for each channel a , if $c(a) = (\delta, \sigma)$, $\delta \in \mathbb{Z}^*$, then $\sigma \sqsubseteq S(a)$.*

Proof (sketch) The proof is made by proving σ -security for a generic σ . Consider two concrete executions starting from the same configurations of the input channels belonging to $\text{Names}^{\sqsubseteq \sigma}$. Until a conditional or iterative instruction has not been reached with a high conditional expression (i.e. with level $\not\sqsubseteq \sigma$), by Lemma 2 the two executions perform the same instructions and reach at each step σ -equivalent states. If an input or output instruction is executed on a channel a such that $S(a) \sqsubseteq_{\mathcal{L}} \sigma$, by the hypothesis the level of the channel is not upgraded to a level $\not\sqsubseteq_{\mathcal{L}} \sigma$ and thus the output values are σ -equivalent and have equal numeric part. If no **if** or **while** with high guard is reached, the property is satisfied. If, instead, an **if** or **while** command is reached, say t , with a high guard, then the two executions can be made up of different sequences of instructions. Let $q_1 = \langle t, \rho_1, \mu_1, \xi_1, c_1 \rangle$ and $q_2 = \langle t, \rho_2, \mu_2, \xi_2, c_2 \rangle$ the states reached by the two computations respectively before the execution of t and let $q_1^1 = \langle t_1^1, \rho_1^1, \mu_1, \xi_1, c_1 \rangle$ and $q_2^1 = \langle t_2^1, \rho_2^1, \mu_2, \xi_2, c_2 \rangle$, respectively, the states after the execution of the rule **If** or **While** with label t . By Lemma 2 it holds that $\rho_1^1 \stackrel{E}{=} \rho_2^1$. As a consequence, by the **If** and

While rules, for each instruction $t' \in \text{scope}(t)$, it holds $\rho_1^1(t') \not\sqsubseteq_{\mathcal{L}} \sigma$ and $\rho_2^1(t') \not\sqsubseteq_{\mathcal{L}} \sigma$. While instructions that belong to $\text{scope}(t)$ are executed, since the annotation of values is upgraded to the level of the environment of the instructions, it holds: i) if a variable or a field of a structure is updated, then the stored value has annotation $\not\sqsubseteq_{\mathcal{L}} \sigma$ in both executions; and ii) no input or output channel a with $S(a) \sqsubseteq_{\mathcal{L}} \sigma$ is affected, otherwise sem would not respect the condition imposed by the theorem (since the environment of the instruction is $\not\sqsubseteq_{\mathcal{L}} \sigma$, if a channel is updated, the annotation of the channel would rise to a level $\not\sqsubseteq_{\mathcal{L}} \sigma$). Two cases are possible:

Case 1 At least one subsequent instruction of the program does not belong to $\text{scope}(t)$ (there is not a loop). Let \bar{t} be the first instruction not belonging to $\text{scope}(t)$, that is \bar{t} is the join point of the two branches. Both computations reach \bar{t} . Let $\bar{q}_1 = \langle \bar{t}, \bar{\rho}_1, \bar{\mu}_1, \bar{\xi}_1, \bar{c}_1 \rangle$ and $\bar{q}_2 = \langle \bar{t}, \bar{\rho}_2, \bar{\mu}_2, \bar{\xi}_2, \bar{c}_2 \rangle$, respectively, the corresponding states. Note that, since the definition of rule **If** and **While** upgrades the environment of all instructions in $\text{scope}(t)$, it holds $\bar{\rho}_1 \stackrel{E}{=} \rho_1 1 \stackrel{E}{=} \bar{\rho}_2 \stackrel{E}{=} \rho_2 1$. Let i and j be the minimum indexes of the chain sem_n such that $\bar{q}_1 \in \text{sem}_i$ and $\bar{q}_2 \in \text{sem}_j$. Due to the **align** operation applied by **next**, there are two states in $\text{sem}_{\max(i,j)}(\bar{t})$, $\hat{q}_1 = \langle \bar{t}, \hat{\rho}_1, \hat{\mu}_1, \hat{\xi}_1, \hat{c}_1 \rangle$ and $\hat{q}_2 = \langle \bar{t}, \hat{\rho}_2, \hat{\mu}_2, \hat{\xi}_2, \hat{c}_2 \rangle$, corresponding resp. to the alignment of \bar{q}_1 and \bar{q}_2 , i.e. with the same execution values resp. of \bar{q}_1 and \bar{q}_2 , but with the security levels upgraded to the lub of all the states in $\text{sem}_{\max(i,j)}(\bar{t})$. We have that $(\hat{\mu}_1, \hat{\xi}_1) \stackrel{\mathcal{M}, \Xi}{=} (\hat{\mu}_2, \hat{\xi}_2)$ since $(\mu_1, \xi_1) \stackrel{\mathcal{M}, \Xi}{=} (\mu_2, \xi_2)$ and a variable or a structure field in \hat{q}_1, \hat{q}_2 has level $\sqsubseteq_{\mathcal{L}} \sigma$ after **align** only if has not been affected by either of the two computations. A similar reasoning can be made for channels. Thus it holds that $\hat{q}_1 \stackrel{\mathcal{Q}}{=} \hat{q}_2$. The above reasoning can be iterated following the two computations from resp. \hat{q}_1 and \hat{q}_2 .

Case 2 All subsequent instructions of the program belong to $\text{scope}(t)$ (there is a loop). By hypothesis, channels with security level $\sqsubseteq_{\mathcal{L}} \sigma$ cannot be affected from this point on by either of the two computations and thus the property is satisfied. \square

5 Abstract domains

In this section we define the abstract domains of the abstract interpretation and we prove that they are connected to the concrete domains by means of a correctness condition. Let us recall the definition of Galois Insertion.

Definition 6 (Galois Insertion) Let (C, \sqsubseteq) and (A, \sqsubseteq) be two complete lattices. Two functions $\alpha: C \mapsto A$ and

Fig. 10 Lattice of abstract values

$$\begin{aligned}
 \mathcal{A}^{\natural} &= \wp(\mathbf{New}), \mathcal{V}^{\natural} = (\{\bullet\} \cup \mathcal{A}^{\natural}) \times \mathcal{L} \text{ ranged over by } v_1^{\natural}, v_2^{\natural}, \dots \\
 \sqsubseteq_{\mathcal{V}} &: v_1^{\natural} \sqsubseteq_{\mathcal{V}} v_2^{\natural} \text{ iff } v_1^{\natural} = (T_1, \sigma_1) \wedge v_2^{\natural} = (T_2, \sigma_2) \wedge T_1 \subseteq T_2 \wedge \sigma_1 \sqsubseteq_{\mathcal{L}} \sigma_2 \vee \\
 &\quad \vee v_1^{\natural} = (\bullet, \sigma_1) \wedge v_2^{\natural} = (\bullet, \sigma_2) \wedge \sigma_1 \sqsubseteq_{\mathcal{L}} \sigma_2 \\
 \sqcup_{\mathcal{V}} &: v_1^{\natural} \sqcup_{\mathcal{V}} v_2^{\natural} = \begin{cases} (T_1 \cup T_2, \sigma_1 \sqcup_{\mathcal{L}} \sigma_2) & \text{if } v_1^{\natural} = (T_1, \sigma_1) \wedge v_2^{\natural} = (T_2, \sigma_2) \\ (\bullet, \sigma_1 \sqcup_{\mathcal{L}} \sigma_2) & \text{if } v_1^{\natural} = (\bullet, \sigma_1) \wedge v_2^{\natural} = (\bullet, \sigma_2) \\ \top_{\mathcal{V}} & \text{otherwise} \end{cases} \\
 \sqcap_{\mathcal{V}} &: v_1^{\natural} \sqcap_{\mathcal{V}} v_2^{\natural} = \begin{cases} (T_1 \cap T_2, \sigma_1 \sqcap_{\mathcal{L}} \sigma_2) & \text{if } v_1^{\natural} = (T_1, \sigma_1) \wedge v_2^{\natural} = (T_2, \sigma_2) \\ (\bullet, \sigma_1 \sqcap_{\mathcal{L}} \sigma_2) & \text{if } v_1^{\natural} = (\bullet, \sigma_1) \wedge v_2^{\natural} = (\bullet, \sigma_2) \\ \perp_{\mathcal{V}} & \text{otherwise} \end{cases} \\
 \alpha_{\mathcal{V}}^1(v) &= \begin{cases} (\bullet, \sigma) & v = (k, \sigma), k \in \mathbb{Z} \\ (\{t\}, \sigma) & v = ((\ell, t), \sigma), (\ell, t) \in \mathcal{A} \end{cases} \quad y \in \wp(\mathcal{V}), \alpha_{\mathcal{V}}(y) = \bigsqcup_{v_i \in y} \alpha_{\mathcal{V}}^1(v_i) \\
 \gamma_{\mathcal{V}}(v^{\natural}) &= \begin{cases} \{(k, \sigma') \mid k \in \mathbb{Z}, \sigma' \sqsubseteq_{\mathcal{L}} \sigma\} & v^{\natural} = (\bullet, \sigma) \\ \{((\ell, t), \sigma') \mid t \in T, (\ell, t) \in \mathcal{A}, \sigma' \sqsubseteq_{\mathcal{L}} \sigma\} & v^{\natural} = (T, \sigma) \\ \mathcal{V} & v^{\natural} = \top_{\mathcal{V}} \\ \emptyset & v^{\natural} = \perp_{\mathcal{V}} \end{cases}
 \end{aligned}$$

$\gamma: A \mapsto C$ form a *Galois insertion* between (C, \sqsubseteq) and (A, \sqsubseteq) , iff all the following conditions hold:

- **α -Monotonicity:** $\forall y, y' \in C. y \subseteq y' \Rightarrow \alpha(y) \sqsubseteq \alpha(y')$
- **γ -Monotonicity:** $\forall a, a' \in A. a \sqsubseteq a' \Rightarrow \gamma(a) \subseteq \gamma(a')$
- **Galois:** $\forall y \in C. y \subseteq \gamma(\alpha(y))$
- **Insertion:** $\forall a \in A. \alpha(\gamma(a)) = a$

Showing that the concrete and the abstract domains are connected by a Galois Insertion will be necessary for proving that the abstract flow equations converge to a fixpoint [25]. The abstract domains are obtained by eliminating from the concrete values both the execution values and execution addresses. Every value maintains instead its security annotation. Simple values (`int`) are no longer held and are represented with a \bullet symbol. In order to make the heap finite we abstract onto the same element different structures created at the same label. Moreover, an abstract address ℓ^{\natural} is composed of a set of labels in `New`. In this way ℓ^{\natural} records all the possible creation points of the structures pointed to by it during the computation. The operations defined on the lattice of abstract values ($\mathcal{V}^{\natural}, \sqsubseteq_{\mathcal{V}}^{\natural}, \sqcup_{\mathcal{V}}^{\natural}, \sqcap_{\mathcal{V}}^{\natural}, \perp_{\mathcal{V}}^{\natural}, \top_{\mathcal{V}}^{\natural}$) are reported in Fig. 10. The abstraction of a set of simple values is the lub of their security levels. We assume that $\alpha_{\mathcal{V}}$ returns the bottom element of \mathcal{V}^{\natural} if applied to the the empty set. Dually for the concretization function. The same for the other abstraction functions. The abstraction of a set of concrete references is an abstract reference that contains both the lub of their security levels and a set T of instruction points. T contains all the instructions at which the structure referenced is created. For example, if $v_1 = ((\ell_1, t_1), \sigma_1)$ and $v_2 = ((\ell_2, t_2), \sigma_2)$, then $v^{\natural} = \alpha_{\mathcal{V}}(\{v_1, v_2\}) = (\{t_1, t_2\}, \sigma_1 \sqcup_{\mathcal{L}} \sigma_2)$.

An abstract memory $\mu^{\natural} \in \mathcal{M}_X^{\natural} = X \rightarrow \mathcal{V}^{\natural}$ maps every variable in the set X to an abstract value (see Fig. 11). Two abstract memories can be compared only

$$\begin{aligned}
 \mathcal{M}_X^{\natural} &= X \rightarrow (\mathcal{L} \cup \mathcal{A}^{\natural}) \text{ ranged over by } \mu_1, \mu_2, \dots \\
 \sqsubseteq_{\mathcal{M}} &: \mu_1 \sqsubseteq_{\mathcal{M}} \mu_2 \text{ iff } \forall x \in X. \mu_1(x) \sqsubseteq_{\mathcal{A}} \mu_2(x) \\
 \sqcup_{\mathcal{M}} &: (\mu_1 \sqcup_{\mathcal{M}} \mu_2)(x) = \mu_1(x) \sqcup_{\mathcal{V}} \mu_2(x) \\
 \sqcap_{\mathcal{M}} &: (\mu_1 \sqcap_{\mathcal{M}} \mu_2)(x) = \mu_1(x) \sqcap_{\mathcal{V}} \mu_2(x) \\
 \perp_{\mathcal{M}} &: \forall x \in X. \perp_{\mathcal{M}}(x) = \perp_{\mathcal{V}} \\
 \top_{\mathcal{M}} &: \forall x \in X. \top_{\mathcal{M}}(x) = \top_{\mathcal{V}} \\
 \alpha_{\mathcal{M}}(y)(x) &= \alpha_{\mathcal{V}}(\{\mu(x) \mid \mu \in y\}) \\
 \gamma_{\mathcal{M}}(\mu^{\natural}) &= \{\mu \in \mathcal{M} \mid \forall x \in X. \mu(x) \in \gamma_{\mathcal{V}}(\mu^{\natural}(x))\}
 \end{aligned}$$

Fig. 11 Lattice of abstract memories

$$\begin{aligned}
 \Xi^{\natural} &= \mathbf{New} \rightarrow \mathcal{M}_{\text{struct}}^{\natural} \text{ ranged over by } \xi_1, \xi_2, \dots \\
 \sqsubseteq_{\Xi} &: \xi_1 \sqsubseteq_{\Xi} \xi_2 \text{ iff } \forall t \in \mathbf{New}. \xi_1(t) \sqsubseteq_{\mathcal{M}} \xi_2(t) \\
 \sqcup_{\Xi} &: (\xi_1 \sqcup_{\Xi} \xi_2)(t) = \xi_1(t) \sqcup_{\mathcal{M}} \xi_2(t) \\
 \sqcap_{\Xi} &: (\xi_1 \sqcap_{\Xi} \xi_2)(t) = \xi_1(t) \sqcap_{\mathcal{M}} \xi_2(t) \\
 \perp_{\Xi} &: \forall t \in \mathbf{New}. \perp_{\Xi}(t) = \perp_{\mathcal{M}} \\
 \top_{\Xi} &: \forall t \in \mathbf{New}. \top_{\Xi}(t) = \top_{\mathcal{M}} \\
 \alpha_{\Xi}(y)(t) &= \alpha_{\mathcal{M}}(\{\xi(\ell, t) \mid \xi \in y, \ell \in \mathcal{A}^e\}) \\
 \gamma_{\Xi}(\xi^{\natural}) &= \{\xi \in \Xi \mid \forall (\ell, t) \in \mathcal{A}. \xi(\ell, t) \in \gamma_{\mathcal{M}}(\xi^{\natural}(t))\}
 \end{aligned}$$

Fig. 12 Lattice of abstract heaps

if their domains are the same. When $X = \mathbf{Var}$ we omit the subscript and indicate the domain with \mathcal{M}^{\natural} .

An abstract heap $\xi^{\natural} \in \Xi^{\natural} = \mathbf{New} \rightarrow \mathcal{M}_{\text{struct}}^{\natural}$ is a map from structure creation points to abstract memories representing fields contents. Two heaps $\xi_1^{\natural}, \xi_2^{\natural}$ can be compared only if each abstract address points to structures of the same type, i.e. $\forall t \in \mathbf{New}. \xi_1^{\natural}(t)$ and $\xi_2^{\natural}(t)$ are comparable memories (see Fig. 12).

Input and output channels are represented in the abstract domain $\mathcal{C}^{\natural} = \mathbf{Names} \rightarrow \mathcal{L}$ with tuples of security levels, one for each channel (see Fig. 13).

Proposition 2 $\alpha_{\mathcal{V}}$ and $\gamma_{\mathcal{V}}$ form a Galois Insertion.

Proof (α -Monotonicity) Let $y, y' \in \wp(\mathcal{V})$. Since $\alpha_{\mathcal{V}}(y) = \bigsqcup_{v_i \in y} \alpha_{\mathcal{V}}^1(v_i)$ and $y \subseteq y'$, we can write

$$\begin{aligned}
\mathcal{C}^{\natural} &= \text{Names}_I \rightarrow \mathcal{L} \text{ ranged over by } c_1, c_2 \dots \\
\sqsubseteq_c &: c_1 \sqsubseteq_c c_2 \text{ iff } \forall a \in \text{Names}. c_1(a) \sqsubseteq_{\mathcal{L}} c_2(a) \\
\sqcup_c &: (c_1 \sqcup_c c_2)(a) = c_1(a) \sqcup_{\mathcal{L}} c_2(a) \\
\sqcap_c &: (c_1 \sqcap_c c_2)(a) = c_1(a) \sqcap_{\mathcal{L}} c_2(a) \\
\perp_c &: \forall a \in \text{Names}. \perp_c(a) = \perp_{\mathcal{L}} \\
\top_c &: \forall a \in \text{Names}. \top_c(a) = \top_{\mathcal{L}} \\
\alpha_{\mathcal{C}}(y)(a) &= \alpha_{\mathcal{L}}(\{c(a) \mid c \in y\}) \\
\gamma_{\mathcal{C}}(c^{\natural}) &= \{c \in \mathcal{C} \mid \forall a \in \text{Names}. c(a) = (s, \sigma), s \in \mathbb{Z}^+, \sigma \sqsubseteq_{\mathcal{L}} c^{\natural}(a)\}
\end{aligned}$$

Fig. 13 Lattice of abstract channels

$\alpha_{\mathcal{V}}(y') = \alpha_{\mathcal{V}}(y) \sqcup_{\mathcal{V}} \bigsqcup_{v_i \in y \setminus y'} \alpha_{\mathcal{V}}^1(v_i)$. The properties of the lub $\sqcup_{\mathcal{V}}$ assure therefore that $\alpha_{\mathcal{V}}(y) \sqsubseteq_{\mathcal{A}} \alpha_{\mathcal{V}}(y')$.

(γ -Monotonicity) Let $v_1^{\natural}, v_2^{\natural} \in \mathcal{V}^{\natural}$. If $v_1^{\natural} \sqsubseteq_{\mathcal{V}} v_2^{\natural}$ then either both v_1^{\natural} and v_2^{\natural} are addresses or both are simple values. Let us consider the former case, in which $v_1^{\natural} = (T_1, \sigma_1)$ and $v_2^{\natural} = (T_2, \sigma_2)$. We will prove $\gamma_{\mathcal{V}}(v_1^{\natural}) \subseteq \gamma_{\mathcal{V}}(v_2^{\natural})$ showing that every $v = (\ell, t, \sigma) \in \gamma_{\mathcal{V}}(v_1^{\natural})$ belongs also to $\gamma_{\mathcal{V}}(v_2^{\natural})$. $t \in T \wedge \sigma \sqsubseteq_{\mathcal{L}} \sigma_1$ since $v \in \gamma_{\mathcal{V}}(v_1)$. On the other hand, $T_1 \subseteq T_2 \wedge \sigma_1 \sqsubseteq_{\mathcal{L}} \sigma_2$, since $v_1^{\natural} \sqsubseteq_{\mathcal{V}} v_2^{\natural}$. Then, $t \in T_1 \subseteq T_2 \Rightarrow t \in T_2$ and $\sigma \sqsubseteq_{\mathcal{L}} \sigma_1 \sqsubseteq_{\mathcal{L}} \sigma_2 \Rightarrow \sigma \sqsubseteq_{\mathcal{L}} \sigma_2$, thus proving that $v \in \gamma_{\mathcal{V}}(v_2^{\natural})$. If v_1^{\natural} and v_2^{\natural} are simple values, the proof is similar.

(Galois) Let $y \in \wp(\mathcal{V})$ and $v \in y$. We will show that $v \in \gamma_{\mathcal{V}}(\alpha_{\mathcal{V}}(y))$ to prove the thesis.

$$\begin{aligned}
v \in y &\Rightarrow \{v\} \subseteq y \Rightarrow \text{(by monotonicity of } \alpha_{\mathcal{V}}) \\
&\Rightarrow \alpha_{\mathcal{V}}(\{v\}) \sqsubseteq_{\mathcal{V}} \alpha_{\mathcal{V}}(y) \\
&\Rightarrow \text{(by monotonicity of } \gamma_{\mathcal{V}}) \\
&\Rightarrow \gamma_{\mathcal{V}}(\alpha_{\mathcal{V}}(\{v\})) \subseteq \gamma_{\mathcal{V}}(\alpha_{\mathcal{V}}(y)).
\end{aligned}$$

Thus, it suffices to prove that $v \in \gamma_{\mathcal{V}}(\alpha_{\mathcal{V}}(\{v\})) = \gamma_{\mathcal{V}}(\alpha_{\mathcal{V}}^1(v))$. Suppose that y contains only addresses: then $\alpha_{\mathcal{V}}^1(v) = \alpha_{\mathcal{V}}^1((\ell, t, \sigma)) = v^{\natural} = (t, \sigma)$. It is straightforward, by definition of $\gamma_{\mathcal{V}}$, that $v \in \gamma_{\mathcal{V}}(v^{\natural})$. Analogously for the case in which y contains only simple values. If y contains both addresses and simple values, $\alpha_{\mathcal{V}}(y) = \top_{\mathcal{V}}$ and $\gamma_{\mathcal{V}}(\top_{\mathcal{V}}) = \wp(\mathcal{V})$ that is the top element of the lattice (\mathcal{V}, \subseteq) .

(Insertion) Let $v \in \mathcal{V}^{\natural}, v = (T, \sigma), y = \gamma_{\mathcal{V}}(v)$. Then, by definition of abstraction and concretization functions,

$$\begin{aligned}
\alpha_{\mathcal{V}}(y) &= \bigsqcup_{v_i \in y} \alpha_{\mathcal{V}}(v_i) = \bigsqcup_{t \in T \sigma' \sqsubseteq_{\mathcal{L}} \sigma} \alpha_{\mathcal{V}}((\ell, t), \sigma') \\
&= \bigsqcup_{t \in T \sigma' \sqsubseteq_{\mathcal{L}} \sigma} (\{t\}, \sigma') = (T, \sigma)
\end{aligned}$$

Similarly for a $v = (\bullet, \sigma)$. \square

Proposition 3 $\alpha_{\mathcal{M}}$ and $\gamma_{\mathcal{M}}$ form a Galois Insertion.

Proof (**α -Monotonicity**) Let $y, y' \in \wp(\mathcal{M}_X)$ and $x \in X$.

$$\begin{aligned}
y \subseteq y' &\Rightarrow \{\mu(x) \mid \mu \in y\} \subseteq \{\mu(x) \mid \mu \in y'\} \\
&\Rightarrow \text{(by monotonicity of } \alpha_{\mathcal{V}}) \\
&\Rightarrow \alpha_{\mathcal{V}}(\{\mu(x) \mid \mu \in y\}) \sqsubseteq_{\mathcal{V}} \alpha_{\mathcal{V}}(\{\mu(x) \mid \mu \in y'\}) \\
&\Rightarrow \text{(by definition of } \alpha_{\mathcal{M}}) \\
&\Rightarrow \alpha_{\mathcal{M}}(y)(x) \sqsubseteq_{\mathcal{V}} \alpha_{\mathcal{M}}(y')(x)
\end{aligned}$$

(γ -Monotonicity) Let $\mu, \mu' \in \mathcal{M}^{\natural}$.

$$\begin{aligned}
\mu \sqsubseteq_{\mathcal{M}} \mu' &\Leftrightarrow \forall x \in X. \mu(x) \sqsubseteq_{\mathcal{V}} \mu'(x) \\
&\Rightarrow \text{(by monotonicity of } \gamma_{\mathcal{V}}) \\
&\Rightarrow \forall x \in X. \gamma_{\mathcal{V}}(\mu(x)) \subseteq \gamma_{\mathcal{V}}(\mu'(x)) \\
&\Rightarrow \text{(by definition of } \gamma_{\mathcal{M}}) \\
&\Rightarrow \gamma_{\mathcal{M}}(\mu) \subseteq \gamma_{\mathcal{M}}(\mu').
\end{aligned}$$

(Galois) Let $y \in \wp(\mathcal{M}_X)$ and $\mu \in y$.

$$\begin{aligned}
\{\mu\} &\subseteq y \\
&\Rightarrow \text{(by monotonicity of } \alpha_{\mathcal{M}}) \\
&\Rightarrow \alpha_{\mathcal{M}}(\{\mu\}) \sqsubseteq_{\mathcal{M}} \alpha_{\mathcal{M}}(y) \\
&\Rightarrow \text{(by monotonicity of } \gamma_{\mathcal{M}}) \\
&\Rightarrow \gamma_{\mathcal{M}}(\alpha_{\mathcal{M}}(\{\mu\})) \subseteq \gamma_{\mathcal{M}}(\alpha_{\mathcal{M}}(y)).
\end{aligned}$$

Thus, it suffices to show that $\mu \in \gamma_{\mathcal{M}}(\alpha_{\mathcal{M}}(\{\mu\}))$ to prove the thesis.

$$\begin{aligned}
&\text{(by definition of } \alpha_{\mathcal{M}}) \forall x \in X \alpha_{\mathcal{M}}(\{\mu\})(x) = \alpha_{\mathcal{V}}(\{\mu(x)\}) \\
&\text{(by Proposition 2)} \quad \forall x \in X \{\mu(x)\} \subseteq \gamma_{\mathcal{V}}(\alpha_{\mathcal{V}}(\{\mu(x)\})) \Rightarrow \\
&\text{(by definition of } \gamma_{\mathcal{M}}) \quad \Rightarrow \mu \in \gamma_{\mathcal{M}}(\alpha_{\mathcal{M}}(\{\mu\}))
\end{aligned}$$

(Insertion) Let $\mu^{\natural} \in \mathcal{M}_X^{\natural}$ and let $x \in X$.

$$\begin{aligned}
\alpha_{\mathcal{M}}(\gamma_{\mathcal{M}}(\mu^{\natural}))(x) &= \text{(by definition of } \alpha_{\mathcal{M}}) \\
&= \alpha_{\mathcal{V}}(\{\mu(x) \mid \mu \in \gamma_{\mathcal{M}}(\mu^{\natural})\}) \\
&= \text{(by definition of } \gamma_{\mathcal{M}}) \\
&= \alpha_{\mathcal{V}}(\{\mu(x) \mid \mu \in \{\mu^{c'} \mid \forall x' \in X. \mu^{c'}(x') \in \gamma_{\mathcal{V}}(\mu^{\natural}(x'))\}\}) \\
&= \text{(by set definition)} \\
&= \alpha_{\mathcal{V}}(\{\mu(x) \mid \mu(x) \in \gamma_{\mathcal{V}}(\mu^{\natural}(x))\}) \\
&= \text{(by set definition)} \\
&= \alpha_{\mathcal{V}}(\gamma_{\mathcal{V}}(\mu^{\natural}(x))) \\
&= \text{(by Proposition 2)} \\
&= \mu^{\natural}(x)
\end{aligned}$$

\square

Proposition 4 α_{Ξ} and γ_{Ξ} form a Galois Insertion.

Proposition 5 $\alpha_{\mathcal{C}}$ and $\gamma_{\mathcal{C}}$ form a Galois Insertion.

We omit proofs for Propositions 4 and 5 since they are similar to the previous.

The abstract domain of states is $Q^\sharp = \mathcal{B} \rightarrow (\mathcal{L} \times \mathcal{M}^\sharp \times \Xi^\sharp \times \mathcal{C}^\sharp)$. It contains all functions associating the instruction labels \mathcal{B} with elements in $(\mathcal{L} \times \mathcal{M}^\sharp \times \Xi^\sharp \times \mathcal{C}^\sharp)$. Given an abstract state $q^\sharp \in Q^\sharp$, and an instruction label $t \in \mathcal{B}$, $q^\sharp(t) = \langle \sigma, \mu^\sharp, \xi^\sharp, c^\sharp \rangle$ is a tuple composed of a security level representing the security environment of t , an abstract memory, heap and channels. We use $q^\sharp(t).env$ to denote σ . We denote by

$$dom(q^\sharp) = \left\{ t \mid q^\sharp(t) = \langle \sigma, \mu^\sharp, \xi^\sharp, c^\sharp \rangle \wedge \mu^\sharp \neq \perp_{\mathcal{M}} \wedge \xi^\sharp \neq \perp_{\Xi} \wedge c^\sharp \neq \perp_{\mathcal{C}} \right\}$$

the instruction addresses to which q^\sharp assigns a defined value for memory, heap and channels. We have that $(Q^\sharp, \sqsubseteq_Q)$ is a lattice, where, the operation \sqsubseteq_Q is defined as the pointwise application of the corresponding operation on the fields of the abstract states. Let us now consider the abstraction and concretization functions between the concrete and abstract domains of the states.

$\alpha_Q : \wp(Q) \rightarrow Q^\sharp$ is defined as follows. Let Q be a set of concrete states in $Q = \mathcal{B} \times Env \times \mathcal{M} \times \Xi \times \mathcal{C}$. For each $t \in \mathcal{B}$, it is $\alpha_Q(Q)(t) = \langle \sigma, \mu^\sharp, \xi^\sharp, c^\sharp \rangle$ where

$$\sigma = \bigsqcup_{\mathcal{L}} \{ \rho(t) \mid \langle t', \rho, \mu, \xi, c \rangle \in Q \},$$

$$\mu^\sharp = \alpha_{\mathcal{M}}(\{ \mu \mid \langle t, \rho, \mu, \xi, c \rangle \in Q \}),$$

$$\xi^\sharp = \alpha_{\Xi}(\{ \xi \mid \langle t, \rho, \mu, \xi, c \rangle \in Q \}),$$

$$c^\sharp = \alpha_{\mathcal{C}}(\{ c \mid \langle t, \rho, \mu, \xi, c \rangle \in Q \}).$$

If an instruction t does not occur in Q , then the abstraction functions $\alpha_{\mathcal{M}}$, α_{Ξ} and $\alpha_{\mathcal{C}}$ will produce bottom values, excluding t from $dom(\alpha_Q(Q))$. Note that the security environment of an instruction t (whether t is in $dom(\alpha_Q(Q))$ or not) in the abstract state is the lub of the security environments assigned to t by all states in Q , while the abstract memory, heap and channels associated with t are the lub of the abstractions of the concrete memories, heaps and channels, respectively, occurring the states of Q corresponding to the execution of the instruction with label t . For the concretization function $\gamma_Q : Q^\sharp \rightarrow \wp(Q)$ we have:

$$\begin{aligned} \gamma_Q(q^\sharp) &= \{ \langle t, \rho, \mu, \xi, c \rangle \mid t \in dom(q^\sharp), \\ &\quad \forall t' \in \mathcal{B}, \rho(t') \sqsubseteq q^\sharp(t').env, q^\sharp(t) = \langle \sigma, \mu^\sharp, \xi^\sharp, c^\sharp \rangle, \\ &\quad \mu = \gamma_{\mathcal{M}}(\mu^\sharp), \xi = \gamma_{\Xi}(\xi^\sharp), c = \gamma_{\mathcal{C}}(c^\sharp) \} \end{aligned}$$

Theorem 2 α_Q and γ_Q form a Galois Insertion.

Proof (α -**Monotonicity**) We want to prove that, $\forall Q, Q' \in \wp(Q)$, such that $Q \subseteq Q'$, $\alpha_Q(Q) \sqsubseteq_Q \alpha_Q(Q')$. From the definition of α_Q it is straightforward to prove $dom(\alpha_Q(Q)) \subseteq dom(\alpha_Q(Q'))$. It remains to prove that, given $t \in \mathcal{B}$ with $\alpha_Q(Q)(t) = \langle \sigma, \mu^\sharp, \xi^\sharp, c^\sharp \rangle$, $\alpha_Q(Q')(t) =$

$\langle \sigma', \mu'^\sharp, \xi'^\sharp, c'^\sharp \rangle$, it is: (1) $\sigma \sqsubseteq_{\mathcal{L}} \sigma'$, (2) $\mu^\sharp \sqsubseteq_{\mathcal{M}} \mu'^\sharp$, (3) $\xi^\sharp \sqsubseteq_{\Xi} \xi'^\sharp$ and (4) $c^\sharp \sqsubseteq_{\mathcal{C}} c'^\sharp$. The first inequality holds by definition of α_Q and the set inclusion between Q and Q' . The second, the third and the fourth are proven by the fact $Q \subseteq Q' \Rightarrow \{ \langle t, \rho, \mu, \xi, c \rangle \in Q \} \subseteq \{ \langle t, \rho, \mu, \xi, c \rangle \in Q' \}$.

(γ -**Monotonicity**) To prove this property, we must show that $\forall q_1^\sharp, q_2^\sharp \in Q^\sharp, q_1^\sharp \sqsubseteq_Q q_2^\sharp, \gamma_Q(q_1^\sharp) \subseteq \gamma_Q(q_2^\sharp)$. Firstly, we note that $q_1^\sharp \sqsubseteq_Q q_2^\sharp$ implies $dom(q_1^\sharp) \subseteq dom(q_2^\sharp)$. Other consequences of the hypothesis are that, if $q_1^\sharp(t) = \langle \sigma_1, \mu_1^\sharp, \xi_1^\sharp, c_1^\sharp \rangle$ and $q_2^\sharp(t) = \langle \sigma_2, \mu_2^\sharp, \xi_2^\sharp, c_2^\sharp \rangle$, then $\sigma_1 \sqsubseteq_{\mathcal{L}} \sigma_2$, $\mu_1^\sharp \sqsubseteq_{\mathcal{M}} \mu_2^\sharp$, $\xi_1^\sharp \sqsubseteq_{\Xi} \xi_2^\sharp$ and $c_1^\sharp \sqsubseteq_{\mathcal{C}} c_2^\sharp$. By γ -monotonicity for \mathcal{M} , Ξ and \mathcal{C} , we have that $\gamma_{\mathcal{M}}(\mu_1^\sharp) \subseteq \gamma_{\mathcal{M}}(\mu_2^\sharp)$, $\gamma_{\Xi}(\xi_1^\sharp) \subseteq \gamma_{\Xi}(\xi_2^\sharp)$ and $\gamma_{\mathcal{C}}(c_1^\sharp) \subseteq \gamma_{\mathcal{C}}(c_2^\sharp)$. All these facts suffice to state that, by construction of $\gamma_Q, \gamma_Q(q_1^\sharp) \subseteq \gamma_Q(q_2^\sharp)$.

(**Galois**) Let $Q \in \wp(Q)$ and $q \in Q, q = \langle t, \rho, \mu, \xi, c \rangle$. We will show that $q \in \gamma_Q(\alpha_Q(Q))$ to prove that $Q \subseteq \gamma_Q(\alpha_Q(Q))$.

$$\begin{aligned} q \in Q &\Rightarrow \{q\} \subseteq Q \Rightarrow (\text{by monotonicity of } \alpha_Q) \\ &\Rightarrow \alpha_Q(\{q\}) \sqsubseteq_Q \alpha_Q(Q) \Rightarrow (\text{by monotonicity of } \gamma_Q) \\ &\Rightarrow \gamma_Q(\alpha_Q(\{q\})) \subseteq \gamma_Q(\alpha_Q(Q)). \end{aligned}$$

Thus, it suffices to prove that $q \in \gamma_Q(\alpha_Q(\{q\}))$. We have that $\alpha_Q(\{q\}) = q^\sharp$, with $dom(q^\sharp) = \{t\}$ and $\alpha_Q(\{q\})(t) = \langle \rho(t), \alpha_{\mathcal{M}}(\mu), \alpha_{\Xi}(\xi), \alpha_{\mathcal{C}}(c) \rangle$. We have that $q \in \gamma_Q(\alpha_Q(\{q\})) = \{ \langle t', \rho', \mu', \xi', c' \rangle \mid \forall t' \in \mathcal{B}, \rho'(t) \sqsubseteq_{\mathcal{L}} \rho(t), \mu' \in \gamma_{\mathcal{M}}(\alpha_{\mathcal{M}}(\mu)), \xi' \in \gamma_{\Xi}(\alpha_{\Xi}(\xi)), c' \in \gamma_{\mathcal{C}}(\alpha_{\mathcal{C}}(c)) \}$ that, by Propositions 3, 4 and 5 clearly contains q .

(**Insertion**) We have to prove that $\forall q^\sharp \in Q^\sharp, \alpha_Q(\gamma_Q(q^\sharp)(t).env) = q^\sharp$. We will show that the thesis hold analyzing each field of a generic abstract state for the instruction t . For the environment we have:

$$\begin{aligned} \alpha_Q(\gamma_Q(q^\sharp)(t).env) &= (\text{by definition of } \alpha_Q) \\ &= \bigsqcup_{\mathcal{L}} \{ \rho(t) \mid \langle t', \rho(t), \mu, c \rangle \in \gamma_Q(q^\sharp) \} \\ &= (\text{by definition of } \gamma_Q) \\ &= \bigsqcup_{\mathcal{L}} \{ \rho(t) \sqsubseteq_{\mathcal{L}} q^\sharp.env \} = q^\sharp.env \end{aligned}$$

and for the memory:

$$\begin{aligned} \alpha_Q(\gamma_Q(q^\sharp)(t).mem) &= (\text{by definition of } \alpha_Q) \\ &= \alpha_{\mathcal{M}}(\{ \mu \mid \langle t, \rho(t), \mu, c \rangle \in \gamma_Q(q^\sharp) \}) \\ &= (\text{by definition of } \gamma_Q) \\ &= \alpha_{\mathcal{M}}(\gamma_Q(q^\sharp).mem) \\ &= (\text{by Proposition 3}) \\ &= q^\sharp.mem \end{aligned}$$

where, given $q^{\sharp} \in \mathcal{Q}^{\sharp}$ and $t \in \mathcal{B}$ with $q^{\sharp}(t) = \langle \rho, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle$, $q^{\sharp}(t).mem$ indicates μ^{\sharp} . The heap and channels parts are omitted because similar to the memory case. \square

6 Abstract semantics and correctness

In this section we give an abstract semantics of the language that allows us to finitely execute the program in the abstract domain.

Figure 14 describes the abstract semantics of expressions. The rules of the abstract semantics for instructions are shown in Fig. 15. They define a relation $\xrightarrow{C_{\sharp}^{\sharp}}$ between the abstract states: if the premise of the rule is true, the rule transforms the state q^{\sharp} in the way described by the rule. There is only one rule for `if` and `while`: in both cases, besides propagating the state unchanged to the successors, the field `env` of all the instructions in `scope(t)` are updated. Rules **Value**_{*x,f*} and **Assign**_{*x,f=e*} need some explanations. In the abstract semantics, the structure addresses are lost and the references, besides the security level, contain the set T of possible creation points. Then, in order to obtain the abstract value `x.f` needed by Rule **Value**_{*x,f*}, it is necessary to compute the

lub of $\xi^{\sharp}(t_i)(f)$ for all the t_i in the set T . Similarly, to execute Rule **Assign**_{*x,f=e*}, an assignment must be performed for each abstract structure that x might refer to.

Definition 7 (*next^{sharp} operator*) Given an abstract state q^{\sharp} , the application of the `nextsharp` operator yields the state reached in one step of computation from each instruction:

$$\text{next}^{\sharp}(q^{\sharp}) = \bigsqcup_{\mathcal{Q}} \left\{ \bar{q}^{\sharp} \mid q^{\sharp} \xrightarrow{C_{\sharp}^{\sharp}} \bar{q}^{\sharp} \right\}$$

Proposition 6 (*monotonicity of next^{sharp}*) `nextsharp` is monotone in $(\mathcal{Q}^{\sharp}, \sqsubseteq^{\sharp})$.

Definition 8 (*initial abstract state q_0^{\sharp}*) For the initial state q_0^{\sharp} we have $\text{dom}(q_0^{\sharp}) = \{t_0\}$ and $q_0^{\sharp}(t_0) = \langle \perp_{\mathcal{L}}, \perp_{\mathcal{M}}, \perp_{\Xi}, c_0^{\sharp} \rangle$, where for all $a \in \text{Names}_I$, $c_0^{\sharp}(a) = \mathcal{S}(a)$ and for all $a \in \text{Names}_O$, $c_0^{\sharp}(a) = \perp_{\mathcal{L}}$.

Definition 9 (*abstract semantics*) The abstract semantics $\text{sem}^{\sharp} \in \mathcal{Q}^{\sharp}$ is the least upper bound in $(\mathcal{Q}^{\sharp}, \sqsubseteq^{\sharp})$ of the following increasing chain, defined for all $n \in \mathbb{N}$:

$$\begin{aligned} \text{sem}_0^{\sharp} &= q_0^{\sharp} \\ \text{sem}_{n+1}^{\sharp} &= \text{next}^{\sharp}(\text{sem}_n^{\sharp}) \end{aligned}$$

Fig. 14 Abstract expressions semantics

$$\begin{array}{l} \text{Const} \frac{}{\langle k, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (\bullet, \perp_{\mathcal{L}})} \quad \text{Op} \frac{\langle E_1, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (\bullet, \tau_1), \langle E_2, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (\bullet, \tau_2)}{\langle E_1 \text{ op } E_2, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (\bullet, \tau_1 \sqcup_{\mathcal{L}} \tau_2)} \\ \\ \text{Value} \frac{}{\langle x, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} \mu^{\sharp}(x)} \quad \text{Value} \frac{\mu^{\sharp}(x) = (T, \sigma), \bigsqcup_{v, t \in T} \xi^{\sharp}(t)(f) = (w, \tau)}{\langle x.f, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (w, \tau \sqcup_{\mathcal{L}} \sigma)} \end{array}$$

Fig. 15 Abstract semantics of commands

$$\begin{array}{l} \text{Skip} \frac{q^{\sharp}(t) = \langle \sigma, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle, q^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp}, \xi'^{\sharp}, c'^{\sharp} \rangle}{\bar{q}^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp} \sqcup_{\mathcal{M}} \mu^{\sharp}, \xi'^{\sharp} \sqcup_{\Xi} \xi^{\sharp}, c'^{\sharp} \sqcup_{\mathcal{C}} c^{\sharp} \rangle} \\ \\ \text{Assign} \frac{q^{\sharp}(t) = \langle \sigma, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle, q^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp}, \xi'^{\sharp}, c'^{\sharp} \rangle, \langle E, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (w, \tau)}{\bar{q}^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp} \sqcup_{\mathcal{M}} \mu^{\sharp} [x \leftarrow (w, \sigma \sqcup_{\mathcal{L}} \tau)], \xi'^{\sharp} \sqcup_{\Xi} \xi^{\sharp}, c'^{\sharp} \sqcup_{\mathcal{C}} c^{\sharp} \rangle} \\ \\ \text{Assign} \frac{q^{\sharp}(t) = \langle \sigma, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle, q^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp}, \xi'^{\sharp}, c'^{\sharp} \rangle, \langle E, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (w, \tau_1), \mu^{\sharp}(x) = (T, \tau_2), \tau_3 = \sigma \sqcup_{\mathcal{L}} \tau_1 \sqcup_{\mathcal{L}} \tau_2}{\bar{q}^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp} \sqcup_{\mathcal{M}} \mu^{\sharp}, \xi'^{\sharp} \sqcup_{\Xi} \xi^{\sharp} [t_j.f \leftarrow (w, \tau_3)]_{\forall t_j \in T}, c'^{\sharp} \sqcup_{\mathcal{C}} c^{\sharp} \rangle} \\ \\ \text{New} \frac{q^{\sharp}(t) = \langle \sigma, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle, q^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp}, \xi'^{\sharp}, c'^{\sharp} \rangle}{\bar{q}^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp} \sqcup_{\mathcal{M}} \mu^{\sharp} [x \leftarrow (t, \sigma)], \xi'^{\sharp} \sqcup_{\Xi} \xi^{\sharp}, c'^{\sharp} \sqcup_{\mathcal{C}} c^{\sharp} \rangle} \\ \\ \text{Input} \frac{q^{\sharp}(t) = \langle \sigma, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle, q^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp}, \xi'^{\sharp}, c'^{\sharp} \rangle, \tau = c^{\sharp}(a) \sqcup_{\mathcal{L}} \sigma}{\bar{q}^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp} \sqcup_{\mathcal{M}} \mu^{\sharp} [x \leftarrow \tau], \xi'^{\sharp} \sqcup_{\Xi} \xi^{\sharp}, c'^{\sharp} \sqcup_{\mathcal{C}} c^{\sharp} [a \leftarrow \tau] \rangle} \\ \\ \text{Output} \frac{q^{\sharp}(t) = \langle \sigma, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle, q^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp}, \xi'^{\sharp}, c'^{\sharp} \rangle, \langle E, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (\bullet, \tau)}{\bar{q}^{\sharp}(\text{succ}(t)) = \langle \sigma', \mu'^{\sharp} \sqcup_{\mathcal{M}} \mu^{\sharp}, \xi'^{\sharp} \sqcup_{\Xi} \xi^{\sharp}, c'^{\sharp} \sqcup_{\mathcal{C}} c^{\sharp} [b \leftarrow \sigma \sqcup_{\mathcal{L}} \tau] \rangle} \\ \\ \text{If, while} \frac{q^{\sharp}(t) = \langle \sigma, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle, \langle E, \mu^{\sharp}, \xi^{\sharp} \rangle \xrightarrow{E_{\sharp}^{\sharp}} (\bullet, \tau), q^{\sharp}(\text{succ}_{\text{true}}(t)) = \langle \sigma', \mu'^{\sharp}, \xi'^{\sharp}, c'^{\sharp} \rangle, q^{\sharp}(\text{succ}_{\text{false}}(t)) = \langle \sigma'', \mu''^{\sharp}, \xi''^{\sharp}, c''^{\sharp} \rangle}{\bar{q}^{\sharp}(\text{succ}_{\text{true}}(t)) = \langle \sigma' \sqcup_{\mathcal{L}} \tau, \mu'^{\sharp} \sqcup_{\mathcal{M}} \mu^{\sharp}, \xi'^{\sharp} \sqcup_{\Xi} \xi^{\sharp}, c'^{\sharp} \sqcup_{\mathcal{C}} c^{\sharp} \rangle, \bar{q}^{\sharp}(\text{succ}_{\text{false}}(t)) = \langle \sigma'' \sqcup_{\mathcal{L}} \tau, \mu''^{\sharp} \sqcup_{\mathcal{M}} \mu^{\sharp}, \xi''^{\sharp} \sqcup_{\Xi} \xi^{\sharp}, c''^{\sharp} \sqcup_{\mathcal{C}} c^{\sharp} \rangle, \forall t' \in \text{scope}(t) : \bar{q}^{\sharp}(t').env = \tau \sqcup_{\mathcal{L}} q^{\sharp}(t').env} \end{array}$$

Proposition 7 (Expression semantics approximation)

$$\forall \langle E, \mu, \xi \rangle : \langle E, \mu, \xi \rangle \xrightarrow{E} v \Rightarrow \langle E, \alpha_{\mathcal{M}}(\mu), \alpha_{\Xi}(\xi) \rangle \xrightarrow{E^{\sharp}} v^{\sharp}$$

with $\alpha_{\mathcal{V}}^1(v) \sqsubseteq_{\mathcal{V}} v^{\sharp}$

Proof We prove the correspondences between each pair of rules in concrete and abstract semantics for expressions.

- (Const)** $v^{\sharp} = \alpha_{\mathcal{V}}^1(v)$ is true since $\alpha_{\mathcal{V}}^1((k, \perp_{\mathcal{L}})) = \perp_{\mathcal{L}}$.
- (Value_x)** $\mu^{\sharp} = \alpha_{\mathcal{M}}(\mu) \Rightarrow \mu^{\sharp}(x) = \alpha_{\mathcal{V}}^1(\mu(x)) \forall x \in Var$ by definition of $\alpha_{\mathcal{M}}$.
- (Value_{x,f})** We want to prove $\alpha_{\mathcal{V}}^1(v^e, \sigma_1 \sqcup_{\mathcal{L}} \sigma_2) \sqsubseteq_{\mathcal{V}} (w, \sigma \sqcup_{\mathcal{L}} \tau)$. First of all, since $\mu^{\sharp} = \alpha_{\mathcal{M}}(\mu)$ then $T = \{t\}$ and $\sigma = \sigma_1$. Moreover, $v = (w, \tau) = \xi^{\sharp}(t)(f) = \alpha_{\Xi}(\xi)(t)(f) = \bigsqcup_{\mathcal{V}} \xi(\ell_i, t)(f) \sqsubseteq_{\mathcal{V}} \xi(\ell)(f) = (v^e, \sigma_2)$, by which $\xi(\ell_i, t)(f) \sqsubseteq_{\mathcal{V}} \xi(\ell)(f)$, the thesis.
- (Op)** The thesis is $\alpha_{\mathcal{V}}^1(v_1^e \text{ op } v_2^e, \sigma_1 \sqcup_{\mathcal{L}} \sigma_2) = (\bullet, \sigma_1 \sqcup_{\mathcal{L}} \sigma_2) \sqsubseteq_{\mathcal{V}} (\bullet, \tau_1 \sqcup_{\mathcal{L}} \tau_2)$. The proof can be conducted by structural induction. Supposing that $\alpha_{\mathcal{V}}^1(v_i^e, \sigma_i) \sqsubseteq_{\mathcal{V}} (\bullet, \tau_i), i = 1, 2$, then $\sigma_i \sqsubseteq_{\mathcal{L}} \tau_i$, leading to $(\sigma_1 \sqcup_{\mathcal{L}} \sigma_2) \sqsubseteq_{\mathcal{L}} (\tau_1 \sqcup_{\mathcal{L}} \tau_2)$. \square

Proposition 8 (Command semantics correspondence)

$$\forall q \in \mathcal{Q} : q \longrightarrow q' \Rightarrow \alpha_{\mathcal{Q}}(\{q\}) \xrightarrow{C^{\sharp}} q^{\sharp'} \quad \text{such that}$$

$$\alpha_{\mathcal{Q}}(\{q'\}) \sqsubseteq_{\mathcal{Q}} q^{\sharp'}.$$

Proof As in the previous proposition, we will prove that each abstract rule correctly approximates its correspondent concrete version. From the definition of abstraction function we have that $\forall i \neq t, \alpha_{\mathcal{Q}}(\{q\})(i) = (\rho(i), \perp_{\mathcal{M}^{\sharp}}, \perp_{\mathcal{C}^{\sharp}})$. Please note that the application of an abstract semantics rule upgrades only the part of the abstract state referring to the successor(s) of t and that, in every case, the environment is only upgraded. Let us suppose $q = \langle t, \rho, \mu, \xi, c \rangle, q' = \langle t', \rho', \mu', \xi', c' \rangle, q^{\sharp'}(t') = \langle \sigma', \mu^{\sharp'}, \xi^{\sharp'}, c^{\sharp'} \rangle$. We proceed by cases depending on the type of instruction with label t .

- (Assign_{t;x=e})** Only the memory status is affected by this rule. Therefore we have to prove that if $\mu' = \mu [x \leftarrow (v^e, \rho(t) \sqcup \sigma)] \wedge \mu^{\sharp} = \alpha_{\mathcal{M}}(\mu)$, then $\alpha_{\mathcal{M}}(\mu') \sqsubseteq_{\mathcal{M}} \mu^{\sharp} [x \leftarrow (w, \rho(t) \sqcup \tau)]$. This is true, since by the premises of the rules and Prop. 7, we can suppose $\alpha_{\mathcal{V}}^1((v^e, \sigma)) \sqsubseteq_{\mathcal{V}} (w, \tau)$.
- (Assign_{t;x.f=e})** In this case, only the heap state is affected, therefore we have to show that $\alpha_{\Xi}(\xi') \sqsubseteq_{\Xi} \xi^{\sharp'}$. Since $\mu^{\sharp'} = \alpha_{\mathcal{M}}(\mu)$, then $T = t_1$ and $\tau_2 = \sigma_2$. Moreover, by Prop. 7 applied to the premises of the rules, we have $\alpha_{\mathcal{V}}(v^e, \sigma_1) \sqsubseteq_{\mathcal{V}} (w, \tau_1)$. Then we can conclude

$$\alpha_{\mathcal{V}}((v^e, \sigma_3)) \sqsubseteq_{\mathcal{V}} (w, \tau_3) \quad (*)$$

$\alpha_{\Xi}(\xi')(t) = \alpha_{\Xi}(\xi [\ell, f \leftarrow (v^e, \sigma_3)])(t) = \xi^{\sharp}(t), \forall t \in \text{New}, t \neq t_1$, since only the memory for an object created at label t_1 is changed by application of the concrete rule. On the other hand,

$$\alpha_{\Xi}(\xi')(t_1) = \alpha_{\mathcal{M}}(\{\xi'(\ell_j, t_1) | \ell_j \neq \ell\}) \sqcup_{\mathcal{M}} \alpha_{\mathcal{M}}(\xi'(\ell, t_1)) = \xi^{\sharp}(t_1) \sqcup_{\mathcal{M}} \alpha_{\mathcal{M}}(\xi'(\ell, t_1))$$

Therefore, since $\xi^{\sharp'}(t_1) = \xi^{\sharp}(t_1) \sqcup_{\mathcal{M}} \xi^{\sharp}(t_1) [f \leftarrow (w, \tau_3)]$, it remains to prove that:

$$\alpha_{\mathcal{M}}(\xi'(\ell, t_1)) \sqsubseteq_{\mathcal{M}} \xi^{\sharp}(t_1) [f \leftarrow (w, \tau_3)] \quad (**)$$

However, we can note that only the field f is changed by the application of the rules, then, using (*), it is easy to show that also (**) and the thesis hold.

(New_{t;x=new s}) While in the concrete rule both the memory and the heap change, in the abstract rule the heap state remains the same. That happens because, in the concrete rule, we create a new object in a fresh location, filling its fields with default (bottom) values. Since in the abstraction $\xi^{\sharp}(t)$ is the lub between all instances created at label t , the new object does not give any contribution. It remains to prove that:

$$\alpha_{\mathcal{M}}(\mu [x \leftarrow (\ell, t, \rho(t))]) \sqsubseteq_{\mathcal{M}} \alpha_{\mathcal{M}}(\mu) [x \leftarrow (t, \rho(t))]$$

which is straightforward from: $\alpha_{\mathcal{V}}((\ell, t, \rho(t))) = (t, \rho(t))$.

(Input_{t;a?x}) In this case the memory and the input states change. For the memory it suffices to note that $\sigma = \tau \wedge \alpha_{\mathcal{V}}(k, \sigma \sqcup_{\mathcal{L}} \rho(t)) = (\bullet, \tau \sqcup_{\mathcal{L}} \rho(t))$. For the input we must show that: $\alpha_{\mathcal{C}}(c) \sqsubseteq_{\mathcal{C}} c^{\sharp'}$, that is true since only the a channel is changed and $c'(a) = (s, \rho(t) \sqcup_{\mathcal{L}} \sigma) \wedge c^{\sharp'}(a) = \rho(t) \sqcup_{\mathcal{L}} \tau$.

(Output_{t;b?E}) From the premises and Proposition 7 we have that:

$$\alpha_{\mathcal{V}}(k, \sigma_1) \sqsubseteq_{\mathcal{V}} (\bullet, \tau) \Rightarrow \sigma_1 \sqsubseteq_{\mathcal{L}} \tau_1$$

From $c^{\sharp} = \alpha_{\mathcal{C}}(c)$ we have $\sigma_2 = \tau_2$. Then $\alpha_{\mathcal{C}}(c') \sqsubseteq_{\mathcal{C}} c^{\sharp'}$ is true since only the b channel is changed and $c'(b) = (s, \rho(t) \sqcup_{\mathcal{L}} \sigma_1 \sqcup_{\mathcal{L}} \sigma_2) \wedge c^{\sharp'}(b) = \rho(t) \sqcup_{\mathcal{L}} \tau_1 \sqcup_{\mathcal{L}} \tau_2$.

(If, While) We have $q' = \langle t', \rho', \mu, \xi, c \rangle$ with $\rho'(\bar{t}) = \rho(\bar{t})$ for $\bar{t} \notin \text{scope}(t), \rho'(\bar{t}) = \rho(\bar{t}) \sqcup_{\mathcal{L}} \tau$ if $\bar{t} \in \text{scope}(t)$ and the tested condition evaluates to (true, τ) or (false, τ). Now, since by Prop. 7 the tested condition in the abstract case must evaluate to τ' with $\tau \sqsubseteq_{\mathcal{L}} \tau'$ we have that $\alpha_{\mathcal{Q}}(q')(\bar{t}).\text{env} \sqsubseteq_{\mathcal{L}} q^{\sharp'}(\bar{t}).\text{env}$. The abstract state $\alpha_{\mathcal{Q}}(q')$ will have bottom value for

memory, heap and channels for the instruction labels different from t' . If $\bar{t} = t'$ then the memory, heap and channels are unchanged, and $\alpha_Q(q') \sqsubseteq_Q q^{\sharp'}$. \square

Proposition 9 *The application of align preserves abstraction. That is, $\forall Q \in \wp(Q)$,*

$$\alpha_Q(\text{align}(Q)) = \alpha_Q(Q)$$

Proof (sketch) Since $\text{align}(Q)$ contains all the states in Q , plus some states for which the security levels are lifted to the lub of states in Q , these latter states do not contribute to the lub due to the application of α_Q . \square

Proposition 10 (Local correctness) next^{\sharp} is a safe approximation of next:

$$\forall Q \in \wp(Q) : \text{next}(Q) \subseteq \gamma(\text{next}^{\sharp}(\alpha_Q(Q)))$$

Proof It suffices to prove that:

$$\alpha_Q(\text{next}(Q)) \sqsubseteq_Q \text{next}^{\sharp}(\alpha_Q(Q)) \quad (\dagger)$$

Indeed, applying γ -monotonicity to (\dagger) , we can conclude:

$$\text{next}(Q) \subseteq \gamma(\alpha_Q(\text{next}(Q))) \subseteq \gamma(\text{next}^{\sharp}(\alpha_Q(Q)))$$

where the first subset operation is given by the Galois property (Theorem 2). On the other hand, equation (\dagger) can be derived directly from Proposition 8 and from the definitions of next and next^{\sharp} operators.

Using the definition of next^{\sharp} , we can rewrite the right-hand member of (\dagger) as:

$$\begin{aligned} \text{next}^{\sharp}(\alpha_Q(Q)) &= \text{next}^{\sharp} \left(\bigsqcup_{q'' \in Q} \alpha_Q(\{q''\}) \right) \\ &= \bigsqcup_{q^{\sharp} \rightarrow q^{\sharp'}, q^{\sharp} \in \alpha_Q(Q)} q^{\sharp'} \end{aligned}$$

On the other hand, we can use the definition of next to obtain:

$$\begin{aligned} \alpha_Q(\text{next}(Q)) &= \text{(by Definition 3)} \\ &= \alpha_Q(\text{align}(Q \cup \{q' | \exists q \in Q : q \longrightarrow q'\})) \\ &= \text{(by Proposition 9)} \\ &= \alpha_Q(Q \cup \{q' | \exists q \in Q : q \longrightarrow q'\}) \\ &\sqsubseteq_Q \text{(by monotonicity of } \alpha_Q) \\ &\sqsubseteq_Q \bigsqcup_{q \rightarrow q', q \in Q} \alpha(\{q'\}) \end{aligned}$$

From these two results we can rewrite (\dagger) as:

$$\bigsqcup_{q \rightarrow q', q \in Q} \alpha(\{q'\}) \sqsubseteq_Q \bigsqcup_{q^{\sharp} \rightarrow q^{\sharp'}, q^{\sharp} \in \alpha(Q)} q^{\sharp'} \quad (\dagger\dagger)$$

Since Proposition 8 holds, each state contributing to the left-hand lub is surely less or equal of a state contributing to the right-hand lub. The lattice properties ensure that, given a, b, c, d in a lattice (A, \sqsubseteq) , with $a \sqsubseteq b \wedge c \sqsubseteq d$, then $(a \sqcup c) \sqsubseteq (b \sqcup d)$. By applying this property, we can conclude $(\dagger\dagger)$ and therefore the thesis. \square

Theorem 3 (Global correctness) $\alpha_Q(\text{sem}) \sqsubseteq_Q \text{sem}^{\sharp}$.

Proof (sketch) The theorem can be proved by induction on the length of the chains sem_i and sem_i^{\sharp} , observing that $\alpha_Q(\text{sem}_0) = \text{sem}_0^{\sharp}$ (base step) and applying Proposition 10 (induction step). \square

A consequence of the above theorem is the following corollary. Its meaning is that we can use the abstract as a means to check secure information flow.

Corollary 1 *If, for any $t \in \mathcal{B}$, with $\text{sem}^{\sharp}(t) = \langle t, \mu^{\sharp}, \xi^{\sharp}, c^{\sharp} \rangle$, it holds that $\forall a \in \text{Names}, c^{\sharp}(a) \sqsubseteq_{\mathcal{L}} S(a)$, then the considered program has secure information flow.*

Proof (sketch) Combining Theorem 3 with Theorem 1. \square

7 A prototype tool

A prototype tool (lflow)² that, given a program, constructs its abstract semantics sem^{\sharp} , has been developed. lflow accepts programs written in the language described in Sect. 3. The lattice \mathcal{L} has been defined as the simplest two-level chain $\{L, H\}$, with $L \sqsubseteq_{\mathcal{L}} H$, but the tool can be easily extended to manage generic lattices. lflow has been written in C++, using Flex [36] and Bison [20] as scanner and parser generators. After having parsed the input file, lflow builds the initial abstract state q_0^{\sharp} . Then, starting from q_0^{\sharp} , it performs a least fixpoint computation using the Kildall working list algorithm [27]. Finally, it dumps sem^{\sharp} . Giving lflow a “verbose” switch, it is possible to dump also each step of the fixpoint calculation.

As an example, consider the application of the algorithm to programs P5 and P8 in Fig. 4. In Fig. 16, we summarize the abstract execution, showing the result of the algorithm (sem^{\sharp}) in the two cases. Let us briefly explain how the state in Fig. 16a is computed for P5. Initially, the entry point of the program is inserted in the working list and abstractly executed. Every instruction brings its successor into the working list, and until instruction 4 is executed, the states are unchanged from their default value. Execution of instruction 4 aligns the value of x to

² lflow is freely available at the URL: <http://www.ing.unipi.it/~o1103499>.

Fig. 16 Abstract semantics of the programs **a** P5 and **b** P8, calculated using *lflow*

Instruction	env	x	a	b	d
1:d?x	L	L	H	L	L
2:While (x > 0)	L	H	H	H	L
3:b!1	H	H	H	H	L
4:a?x	H	H	H	H	L
end:	H	H	H	H	L

(a)

Instruction	env	x	s1	s2	s3	1.f	1.g	2.f	2.g
1:s1=new S	L	L	∅, L	∅, L	∅, L	L	L	L	L
2:s2=new S	L	L	1, L	∅, L	∅, L	L	L	L	L
3:a?x	L	L	1, L	2, L	∅, L	L	L	L	L
4:if(x) 5: else 6:	L	H	1, L	2, L	∅, L	L	L	L	L
5:s3=s1	H	H	1, L	2, L	∅, L	L	L	L	L
6:s3=s2	H	H	1, L	2, L	∅, L	L	L	L	L
7:s3.f=1	H	H	1, L	2, L	{1, 2}, H	L	L	L	L
8:???	H	H	1, L	2, L	{1, 2}, H	H	L	H	L

(b)

H. Then, when the while instruction is newly executed, the environment of all the instructions in its scope (3,4) is upgraded. The new execution of the loop lifts the security level of channel *b* to *H* (because of the environment, see Rule **Output**), thus making the program insecure. In Fig. 16b we show the abstract semantics for program P8. We can notice that, before executing instruction 7, *s3* may refer either to the object created at 1 or to the object created at 2. After the abstract execution of instruction 7, the field *f* of both the two abstract objects is upgraded.

7.1 The tax example

Reconsider the tax example described in Sect. 2, and suppose a program that implements the tax calculation. User data are stored in a list. Each element of the list is an instance of *User*, a structure (see Fig. 17) that contains three fields: *id* (the user ID), *incomeAmount* (the yearly income of the current user) and *next* (the reference to the next element of the list).

The program communicates, through I/O channels, with the users and with the net, that provides information about the current tax rates. In particular:

- the input channel *user* is used to receive the user *id*,
- the input channel *request* is used to receive the type of the user request,

```

struct User {
    int id;
    int incomeAmount;
    struct User next;
};

struct User user_list ,u,u1;
int n,x,flag ,code ,rate ,tmp;

in user L;
in request L;
in income H;
in from_net L;
out to_net L;
out tax H;
    
```

Fig. 17 Declarations (tax example)

- the input channel *income* is used to receive the user *income*,
- the input channel *from_net* is used to receive the tax rates and the threshold, from the net,
- the output channel *to_net* is used to request to the net:
 - the income threshold that discriminates between the application of the low and the high tax rate (0),
 - the low rate (1),
 - the high rate (2).
- the output channel *tax* is used to communicate to the user the tax rate he/she has to apply.

We set the security level of all the channels to **L** (low, public) except for the *tax* and the *income* channel, whose level is **H** (high, private). The *tax* channel holds confidential data since the applied rate depends on the user income. Recall that, according to our method, there is no need to assign a security level to the variables used in the program. The program is an infinite loop that waits for and serves requests from the users (see Fig. 18). The variable *tmp* holds the last value read from the channel *request* and contains the request code. Requests can be of four types:

- registration to the service (code 0),
- modification of the user income (code 1),
- request of the tax rate to be applied to the user income (code 2),
- un-registration to the service (code 3).

In the first case (Fig. 19) the user is inserted into the list, with a fictitious income of 0. In the second (see Fig. 20), the program, after having found the correct user instance in the list, sets the user income to the value taken from the confidential channel *income*. Otherwise, if the user was not registered, before setting the income, the program registers the user by adding a new element to the list. With the third type of request, an user asks the server which tax rate he/she must apply to his income (see Fig. 21). To answer, the program must contact the net that provides the current threshold and the two rates.

```

user_list = 0;
while (1) {
  user?x;
  request?tmp;          // (*)
  if (tmp == 0) {
    // Registration of a new user (see Fig. 19)
  } else {
    if (tmp == 1) {
      // Set the income for the user u.
      // Register the user if not registered yet.
      // (see Fig. 20)
    } else {
      if (tmp == 2) {
        // Send the user the tax rate due (see Fig. 21)
      } else {
        if (tmp == 3 && user_list != 0) {
          // delete an user (see Fig. 22)
        } else {
          skip;
        };
      };
    };
  };
}; // end of the main loop

```

Fig. 18 Program structure (tax example)

```

if (tmp == 0) {
  u = new User;
  u.id = x;
  u.incomeAmount = 0;
  u.next = user_list;
  user_list = u;
} else { /* ... */ };

```

Fig. 19 The user registration (tax example)

```

if (tmp == 1) {
  u = user_list; // search the user
  while (u != 0 && u.id != x) { u = u.next; };
  if (u == 0){
    u = new User; // create the user
    u.id = x;
    u.next = user_list;
    user_list = u;
  } else { skip; };
  income?tmp; // get the income (**)
  u.incomeAmount = tmp; // set the income
} else { /* ... */ };

```

Fig. 20 The modification of the income (tax example)

Finally, with the fourth type of request (see Fig. 22), an user can ask the program to delete his/her data from the list.

Analyzing the program with *lflow*, we can observe that, in the abstract semantics, all the channel but `to_net` respect their specification. In Fig. 23, the channels status, produced by the tool for the main loop instruction, is shown. The two security levels shown for each channel are the level specified by the policy and the level calculated by the abstract semantics, respectively. As the extract shows, the security level of `to_net` has

```

if (tmp == 2) {
  u = user_list; // find the user
  while (u != 0 && u.id != x) { u=u.next; };
  to_net!0; // request the threshold
  from_net?n; // get the threshold
  if (u != 0) {
    if (u.incomeAmount < n) {
      code = 1; // request the low rate
    } else {
      code = 2; // request the high rate
    };
    to_net!code; // request the due tax rate
    from_net?rate; // get the tax rate
    tax!rate; // send the tax rate due
  } else { skip; };
} else { /* ... */ };

```

Fig. 21 The request of the tax rate (tax example)

```

if (tmp == 3 && user list != 0) {
  if (user_list.id == x) { // delete the head of the list
    user_list = user_list.next;
  } else {
    u = user_list; // find the user
    u1 = user_list.next;
    while (u1 != 0 && u1.id != x) {
      u = u1;
      u1 = u1.next;
    };
    if (u1 != 0) {
      u.next = u1.next; // delete the user
    } else { skip; };
  } else { skip; };
}

```

Fig. 22 The deletion of an user (tax example)

```

State for instruction 106:While (Const){ 29 30 105 }
Var code          H
Var high_rate     L
Var low_rate      L
Var n             L
Var rate          L
Var tmp           H
Var u             L { 34 49 }
Var u1           L { 34 49 }
Var user_list     L { 34 49 }
Var x            L
34.incomeAmount  H
34.label         L
34.next          L {34 49 }
49.incomeAmount  H
49.label         L
49.next          L {34 49 }
Channel from_net L L
Channel income   H H
Channel tax      H H
Channel to_net   L H
Channel user     L L
LFP iterations: 183

```

Fig. 23 Extract of the output of the *lflow* tool used with the tax example. The final channels state for the main while instruction is shown. The *last line* shows the number of iterations needed to reach the fixpoint

increased, in the abstract semantics, to the **H** security level. This is due to the part of the program that handles the request with code 2. Since the program requests

```

if (tmp == 2) {
  u = user_list; // find the user
  while (u != 0 && u.id != x ) { u=u.next; };
  if (u != 0) {
    to_net!0;           // request the threshold
    from_net?n;        // get the threshold
    to_net!1;          // request the low rate
    from_net?low_rate; // get the low rate
    to_net!2;          // request the high rate
    from_net?high_rate; // get the high rate
    if (u.incomeAmount < n){ // select the due rate
      rate = low_rate;
    } else {
      rate = high_rate;
    };
    tax!rate;          // send the rate
  } else { skip; };
} else { /* ... */ };

```

Fig. 24 The request of the tax rate - safe version (tax example)

the net either the high or the low rate through the public channel `to_net`, an external observer that has read access to this channel can infer if the user income amount is below the threshold or not, thus violating the policy. To overcome this problem, the code in Fig. 21 can be substituted by the code in Fig. 24. In this case, since the program requests both the low and the high rate, there is no way for the attacker to know which rate is communicated to the user, since the channel `tax` is confidential. Analyzing the new version of the program with `lflow`, the abstract semantics produced by the tool respects the policy. Please note that the variable `tmp` is used both to temporarily store the value read from the private channel `income` and to store the value read from the public channel `request`. Therefore, the variable `tmp` has a different security type after the instruction marked with (\star) (see Fig. 19) and after instruction $(\star\star)$ (see Fig. 20). This would not have been possible with a Volpano-like typing [41], that would have assigned a high security type to `tmp` throughout the whole program, thus leading to deeming the program as insecure. To make this program typable in a Volpano-like typing, two different variables could be used in place of `tmp`: one to store private data, and another to store public data. In general, it may be not easy for the programmer to understand which variables need duplication and this may require a specific prior analysis. Moreover, if many variables needs duplication, the overhead may become significant, especially in embedded systems (e.g. Java Cards [12]) where the memory space is a concern.

7.2 Complexity

Let us now give a short account of the complexity of our analysis: for space complexity, it is $O(N \cdot \log(M) \cdot n)$ where $N = \sharp(\text{Var}) + \sharp(\text{New})$ if the maximum number of

fields of each structure is constant, M is the number of elements in \mathcal{L} and n is the number of program points. The time complexity is theoretically $O(N^2 \cdot M \cdot n)$: every application of an abstract rule has a linear complexity in N due to the lub operation on the abstract memory and heap, and, in the worst case, the abstract state of every instruction can assume up to $O(N \cdot M)$ different values during the verification process. However, in practice, the number of abstract executions is much smaller. As suggested in [30] the dataflow analysis can be conducted at the level of the basic blocks instead of single instructions, saving only the state for the beginning of each basic block and calculating the others on the fly. This helps to reduce the space complexity to $O(N \cdot \log(M) \cdot B)$, and the time complexity to $O(N^2 \cdot M \cdot B)$, where B is the number of basic blocks.

8 Related work and conclusions

A recent survey of works on secure information flow is [39]. The problem has been coped with mainly by means of typing. In type-based approaches, each variable is assigned a security level, which is part of the type of the variable and secure information flow is checked by means of a type system; see, for example, [1, 8, 33, 41, 43]. The work [4] handles secure information flow in object oriented languages. They extend the syntax of a Java-like language with security annotations and build a type system that enforce noninterference. Their language is richer than ours since it includes classes and methods. However, the pointer analysis is only syntactical, while our analysis can distinguish between object instances created at different instructions. In [37, 38] references, exceptions and `let`-polymorphism are treated for a call-by-value λ -calculus.

As compared with typing, AI can give a finer inspection of information flows, since it executes the program, even though in an abstract way. In order to check input/output non-interference, it is not necessary to associate security levels to variables: a variable, during its life, can hold data with different security levels without affecting non-interference, provided that the output channels contain data with levels that are lower than or equal to the channel's level. Security typing can be obtained in our framework by collapsing the rows of the abstract semantics (corresponding to the instructions) into a unique row containing the pointwise lub of the rows. We think that also explicit declassification (as proposed, for instance, in [34]) can be suitably handled by abstract interpretation. The language could be extended with a new command $a!^\sigma \text{ E}$ that sends the result of the expression `E` to the channel `a` forcing the

security annotation of the expression to be equal to σ . The concrete semantics rule could be:

$$\text{Declassification} \frac{\langle E, \mu, \xi \rangle \xrightarrow{E} (k, \tau_1), c(b) = (s, \tau_2), b \in \text{Names}_O}{\langle t, \rho, \mu, \xi, c \rangle \longrightarrow \langle \text{succ}(t), \rho, \mu, \xi, c [b \leftarrow (s \cdot k, \sigma \sqcup_{\mathcal{L}} \tau_2)] \rangle}$$

and the non-interference property of Definition 1 has to be changed in order to exclude the output performed using this new command.

Other papers based on AI [42,21] takes as abstract domain the lattice of levels and perform an AI with almost the same power of typing (in terms of class of certified programs). Thus they do not exploit all the power of abstract interpretation. For example, they can not certify program P7 above. On the other hand, the focus of [21] is the definition of a framework based on AI able to represent a parametrized notion of non-interference. Approaches that are able to cope with “temporary breaking of security”, similar to that presented by program P7, are based on theorem proving [24,26]. AI is also exploited in [2] to annotate programs with pre and postconditions defining variable dependencies.

An early work for analyzing secure information flow in high-level languages was presented by Mizuno and Schmidt [32]. They describe the execution of programs in a rich language (non-void procedures with local and global variables are considered) by means of a denotational semantics that handles execution values and security levels. The full semantics is then approximated, using an AI, by an abstract semantics that handles only security levels. Though they provide a full proof that this approximation is correct, they do not prove that the full semantics enforces the security properties. On the contrary, we justify the use of the concrete semantics by means of Theorem 1.

A recent paper by Hunt and Sands [23] introduces a security typing system that is flow-sensitive, i.e. it allows variables to hold different security levels in different instructions of the program, as our analysis does. Even though their analysis does not deal neither with channels nor with dynamic structures, and therefore a straightforward comparison is not possible, it seems that they achieve the same precision of our method. Moreover, the paper presents an algorithm that transforms a program that is typable in their flow-sensitive analysis into another program that can be also typed in a Volpano-like typing [41].

Some previous papers of the team to which the authors belong cope with the definition of abstractions suitable to check secure information flow, based on the annotation of data with security levels. The works [6,7,9,10] handle secure information flow in stack based

machine languages, while the papers [5,17] consider high level languages, including parallel ones. In these papers

abstract transition systems are used, possibly having a high number of states: the same instruction may belong to different states, characterized by different security environments and memories. This corresponds to perform a less approximate abstract interpretation than instruction level security typing. The number of states being high, the abstraction is not suitable to be directly used for a definition of a tool for checking secure information flow. In fact there is a need for other techniques to be combined with this abstraction method: in the above papers we used model checking to complete the verification process (a similar combination of abstraction and model checking is used in [11]). In the present paper, instead, the abstract semantics is a table composed of a row for each program point and is built by an efficient fixpoint algorithm using the abstract rules. Finally, the previous papers of the authors do not cope with pointers and dynamic structures, here handled by a suitable abstract domain.

References

1. Abadi, M., Banerjee, A., Heintze, N., Riecke, J.G.: A core calculus of dependency. In: 26th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages Proceedings, January 20–22, 1999, San Antonio, pp. 147–160. ACM, New York (1999)
2. Amtoft, T., Banerjee, A.: Information flow analysis in logical form. In: Giacobazzi, R. (ed.) SAS 2004 (11th Static Analysis Symposium), Verona, Italy, August 2004, Vol. 3148 of Lecture Notes in Computer Science, pp. 100–115. Springer, Berlin (2004)
3. Avvenuti, M., Bernardeschi, C., De Francesco, N.: Java bytecode verification for secure information flow. ACM SIGPLAN Notices **38**(12), 20–27 (2003)
4. Banerjee, A., Naumann, D.: Secure information flow and pointer confinement in a Java-like language. In: 15th IEEE Security Foundations Workshop (CSFW’02) Proceedings. IEEE, 2002
5. Barbuti, R., Bernardeschi, C., De Francesco, N.: Abstract interpretation of operational semantics for secure information flow. Inform. Process. Lett. **83**(2), 101–108 (2002)
6. Barbuti, R., Bernardeschi, C., De Francesco, N.: Checking security of Java bytecode by abstract interpretation. In: SAC ’02: Proceedings of the 2002 ACM symposium on Applied computing, March 10–14, 2002, Madrid, Spain, pp. 229–236. ACM New York (2002)
7. Barbuti, R., Bernardeschi, C., De Francesco, N.: Analyzing information flow properties in assembly code by abstract interpretation. Comput. J. **47**(1), 25–45 (2004)

8. Barthe, G., Rezk, T.: Non-interference for a JVM-like language. In: TLDI '05: Proceedings of the 2005 ACM SIGPLAN international workshop on types in languages design and implementation, pp. 103–112, New York, NY, USA, 2005. ACM, New York
9. Bernardeschi, C., De Francesco, N.: Combining abstract interpretation and model checking for analysing security properties of Java bytecode. In: Cortesi, A. (ed.) Third International Workshop on Verification, Model Checking and Abstract Interpretation Proceedings, VMCAI 2002, Venice, January 21–22, 2002, Proceedings, Vol. 2294 of Lecture Notes in Computer Science, pp. 1–15. Springer, Berlin (2002)
10. Bernardeschi, C., De Francesco, N., Lettieri, G.: An abstract semantics tool for secure information flow of stack-based assembly programs. *Microprocess. Microsyst.* **26**(8), 391–398 (2002)
11. Bieber, P., Cazin, J., Girard, P., Lanet, J.-L., Wiels, V., Zanon, G.: Checking secure interactions of smart card applets. In: ESORICS 2000 Proceedings (2000)
12. Chen, Z.: Java Card Technology for Smart Cards: Architecture and Programmer's Guide. Addison-Wesley Longman Publishing, (2000)
13. Cousot, P., Cousot, R.: Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: 4th annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages Proceedings, pp. 238–252, Los Angeles, (1977)
14. Cousot, P., Cousot, R.: Abstract interpretation frameworks. *J. Logic Comput.* **2**, 511–547 (1992)
15. Cousot, P., Cousot, R.: Inductive definitions, semantics and abstract interpretations. In: ACM (ed) Conference Record of the Nineteenth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '92), Albuquerque, New Mexico, January 1992, pp. 83–94. ACM, New York (1992)
16. De Francesco, N., Martini, L.: Abstract interpretation to check secure information flow in programs with input–output security annotations. In: Dimitrakos, T., Martinelli, F., Ryan, P. Y., Schneider, S., (eds) Formal Aspects in Security and Trust: Third International Workshop, FAST 2005, Newcastle upon Tyne, UK, July 18–19, 2005, Revised Selected Papers, Lecture Notes in Computer Science, pp. 63–80. Springer, Berlin (2006)
17. De Francesco, N., Santone, A., Tesei, L.: Abstract interpretation and model checking for checking secure information flow in concurrent systems. *Fundam. Inf.* **54**(2–3), 195–211 (2003)
18. Denning, D.E.: A lattice model of secure information flow. *Commun. ACM* **19**(5), 236–243 (1976)
19. Denning, D.E., Denning, P.J.: Certification of programs for secure information flow. *Commun. ACM* **20**(7), 504–513 (1977)
20. Donnely, C., Stallman, R.: Bison, the YACC-compatible parser generator. Free Software Foundation, November (1995)
21. Giacobazzi, R., Mastroeni, I.: Abstract non-interference: parameterizing non-interference by abstract interpretation. In: Jones, N. D., Leroy, X. (eds.) Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, (POPL '04), Venice, Italy, January 14–16, 2004, pp. 186–197, ACM, New York (2004)
22. Heintze, N., Riecke, J. G.: The SLam calculus: programming with secrecy and integrity. In: Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, January 19–21, 1998, San Diego, CA, USA, pp. 365–377. ACM, New York (1998)
23. Hunt, S., Sands, D.: On flow-sensitive security types. In: Morrisett, J. G., Jones, S. L. P. (eds.) Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2006, Charleston, South Carolina, USA, January 11–13, 2006, pp. 79–90. ACM, New York (2006)
24. Jacobs, B., Pieters, W., Warnier, M.: Statically checking confidentiality via dynamic labels. In: Workshop on Issues in the Theory of Security proceedings, Long Beach, CA, United States, January 20, 2005. ACM, New York (2005)
25. Jones, N. D., Nielson, F.: Abstract interpretation: a semantic based tool for program analysis. In: Abramsky, S., Gabbay, D.M., Maibaum, T.S.E. (eds.) Handbook of Logic in Computer Science, Vol. 4, 527–636. Oxford University Press, Oxford (1995)
26. Joshi, R., Leino, K.M.: A semantic approach to secure information flow. *Sci. Comput. Programm.* **37**(1–3), 113–138 (2000)
27. Kildall, G.A.: A unified approach to global program optimization. In: Proceedings of the 1st Annual ACM Symposium on Principles of Programming Languages, Boston, Massachusetts, October 1973, pp. 194–206 (1973)
28. Kobayashi, N., Shirane, K.: Type-based information flow analysis for low-level languages. In: Informal Proceedings of the 3rd Asian Workshop on Programming Languages and Systems (APLAS'02) (2002)
29. Lampson, B.W.: A note on the confinement problem. *Commun. ACM*, **16**(10), 613–615 (1973)
30. Leroy, X.: Java bytecode verification: algorithms and formalizations. *J. Automat. Reason.* **30**(3–4), 235–269 (2003)
31. Lindholm, T., Yellin, F.: Java Virtual Machine Specification, 2nd edn. Addison-Wesley Longman Publishing (1999)
32. Mizuno, M., Schmidt, D.A.: A security flow control algorithm and its denotational semantics correctness proof. *Formal Aspects Comput.* **4**, 727–754 (1992)
33. Myers, A.C.: Jflow: practical mostly-static information flow control. In: 26th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages Proceedings, January 20–22, 1999, San Antonio, TX, pp. 228–241. ACM, New York (1999)
34. Myers, A.C., Liskov, B.: A decentralized model for information flow control. In: SOSP '97: Proceedings of the sixteenth ACM symposium on Operating systems principles, pp. 129–142. ACM, New York (1997)
35. A guide to understanding Discretionary Access Control in trusted systems. Technical Report NCSC-TG-003 Version 1, National Computer Security Center, 1987
36. Paxson, V.: Flex, a fast scanner generator, version 2.5, March 1995
37. Pottier, F., Conchon, S.: Information flow inference for free. In: Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00), Montreal, Canada, September 18–21, 2000, SIGPLAN Notices **35**(9), pp. 46–57 (2000)
38. Pottier, F., Simonet, V.: Information flow inference for ML. In 29th ACM Symposium on Principles of Programming Languages Proceedings, Portland, January 16–18, 2002, pp. 319–330 (2002)

39. Sabelfeld, A., Myers, A.C.: Language-based information-flow security. *IEEE J. Select. Areas Commun.* **21**(1), 5–19 (2003)
40. Smith, G., Volpano, D.: Secure information flow in a multi-threaded imperative language. In: *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, January 19–21, 1998, San Diego, pp. 1–10. ACM, New York (1998)
41. Volpano, D., Smith, G., Irvine, C.: A sound type system for secure flow analysis. *J. Comput. Secur.* **4**(3), 167–187 (1996)
42. Zanotti, M.: Security typings by abstract interpretation. In: *Static Analysis, 9th International Symposium, SAS 2002, Madrid, Spain, September 17–20, 2002, Proceedings, Vol. 2477 of Lecture Notes in Computer Science*, pp. 360–375, Springer, Berlin (2002)
43. Zdancewic, S., Myers, A.C.: Secure information flow via linear continuations. *Higher Order Symbol. Comput.* **15**(2–3), 209–234, Kluwer, Dordrecht (2002)