

A Progress-Sensitive Flow-Sensitive Inlined Information-Flow Control Monitor (Extended Version)

Andrew Bedford^a, Stephen Chong^c, Josée Desharnais^a, Elisavet Kozyri^b, Nadia Tawbi^a

^aUniversité Laval, Québec, Canada

^bCornell University, New York, USA

^cHarvard University, Massachusetts, USA

Abstract

We present a novel progress-sensitive, flow-sensitive hybrid information-flow control monitor for an imperative interactive language. Progress-sensitive information-flow control is a strong information security guarantee which ensures that a program’s progress (or lack of) does not leak information. Flow-sensitivity means that this strong security guarantee is enforced fairly precisely: our monitor tracks information flow per variable and per program point. We illustrate our approach on an imperative interactive language.

Our hybrid monitor is inlined: source programs are translated, by a type-based analysis, into a target language that supports dynamic security levels. A key benefit of this is that the resulting monitored program is amenable to standard optimization techniques such as partial evaluation.

One of the distinguishing features of our hybrid monitor is that it uses sets of levels to track the different *possible* security types of variables. This feature allows us to distinguish outputs that *never* leak information from those that *may* leak information.

Keywords: information-flow control, non-interference, inline monitor, flow-sensitive, progress-sensitive

1. Introduction

We increasingly rely on computer systems to safeguard confidential information, and maintain the integrity of critical data and operations. But in our highly interconnected world, these trusted systems often need to communicate with untrusted parties. Trusted systems risk leaking confidential information to untrusted parties, or allowing input from untrusted parties to corrupt data or the operation of the trusted system.

Information-flow control is a promising approach to enable trusted systems to interact with untrusted parties, providing fine-grained application-specific control of confidential and untrusted information. Static mechanisms for information-flow control (such as security type systems [1, 2]) analyze a program before execution to determine whether the program’s execution will satisfy the appropriate information flow requirements. This has low runtime overhead, but can generate many false positives. Hybrid mechanisms ([3]) have been proposed to eliminate some of these false positives. They combine static analysis with runtime monitoring in order to increase the precision of the analysis at the expense of higher runtime overhead. Compared to hybrid mechanisms, purely dynamic mechanisms ([4]), which enforce information flow policies only by monitoring execution and without using any static analysis, are either unsound or less permissive [5].

In this work, we enforce confidentiality policies using a

novel hybrid information-flow control monitor for an imperative interactive language. It is an extension of the work presented at IFIP SEC 2016 [6]. The key features of our monitor are as follows.

Our monitor is *progress-sensitive* [7]: it prevents confidential information from being leaked via progress channels. Information leaks through a progress channel when a program’s progress (or lack of) depends on confidential information and is observable by an adversary. It is a generalization of termination-sensitive information security to interactive systems (i.e., systems that interact with an external environment at runtime).

Our monitor is *flow-sensitive*: the security level associated with program variables may change during the execution. Flow sensitivity increases the precision of the monitor, meaning that it is able to accept more programs.

Our language has *channel-valued variables*: communication channels are constants that can be assigned to program variables. This language feature allows realistic communication scenarios to be modelled in our language (e.g., where the same code may communicate with users of arbitrary security levels). Most previous work on language-based information-flow control require that the channel used for an input or output operation be statically known, and allow only a single communication channel per security level.

Our monitor is *inlined*: source programs are translated into a target language that supports dynamic security lev-

els [8]. The type-based translation inserts commands to track the security levels of program variables and contexts, and to control information flow. A key benefit of inlining the monitor is that the resulting monitored program in the target language is amenable to standard optimization techniques such as partial evaluation [9].

Our monitor is *hybrid*: it uses both dynamic and static enforcement techniques. The translation to the target language performs a static analysis. If the program is statically determined to be insecure, then the program is statically rejected. Otherwise, the translation of the program dynamically tracks the statically unknown security levels of variables, and ensures that no leak occurs at runtime.

Our main contributions are as follows.

- We present an extended version of the hybrid monitor first presented in our previous article [6]. The extension consists in generalizing the flow- and progress-sensitive enforcement to general lattices.
- We use sets of levels during static analysis to represent the possible levels that can be associated with a variable during execution. When general lattices are enforced, sets of levels are a more accurate representation than the *unknown level* U introduced in our previous article [6]. Thus, the distinction that our static analysis makes between outputs that *never* leak information and outputs that *may* leak information is more precise when using sets of levels. This more precise distinction leads to fewer inlined commands, and thus to less runtime overhead and increased permissiveness of the monitor.
- Previous work treats channels and regular variables differently, by associating distinct type structures during analysis. Our approach demonstrates that modeling the sensitivity of the information conveyed by these two entities leads to a striking similar type structure.
- We present two additional ideas to increase the precision of the static analysis and the permissiveness of the dynamic analysis: propagating constraints on the set of possible security levels and using conditional updates. We plan on integrating these ideas into our future mechanisms.
- We prove that our inlined monitor is sound and that the semantics of the original program is preserved, as long as the program is secure.

1.1. Examples

We present several examples of programs in our source language, to both provide background on information-flow control, and highlight some of the features of our hybrid monitor. For simplicity, we assume that the variables `lowValue`, `medValue`, `highValue`, `lowChannel`, `medChannel` and `highChannel` exist, have arbitrary values and have

the suggested security levels (L for low, M for medium and H for high).

Explicit and implicit flows. An *insecure explicit information flow* occurs when a confidential value is output to a public channel. An *insecure implicit information flow* [10] occurs when the decision to perform output on a public channel depends on confidential information. This violates security because an observer of the public channel will see whether the output occurred, and might thus learn confidential information. Techniques for tracking and controlling implicit and explicit information flow at the language level are well known [1, 2, 5], and are used in this work. The following program exhibits both insecure explicit and insecure implicit information flow. Our approach will reject this program statically.

```
(* insecure explicit flow *)
send highValue to lowChannel;
if highValue > 0 then
  (* insecure implicit flow *)
  send 42 to lowChannel
end
```

Unknown security levels. Our source language supports variables whose security level could be statically unknown. Consider the following program, where the output may or may not be secure, depending on the value of `lowValue`.

```
if lowValue > 0
  then c := highChannel
  else c := lowChannel
end;
send highValue to c
```

Listing 1: Statically uncertain channel level

Purely static mechanisms would reject this program entirely, and indeed, to the best of our knowledge, all previous work either cannot express this program or statically reject it. This is because it would be unsound to statically treat `c` as being a private channel, since that might incorrectly allow private values to be sent to public channel `lowChannel`. Similarly, it would be unsound to statically treat `c` as a public channel, since that might incorrectly allow values read from private channel `highChannel` to be treated as public values. By contrast, our hybrid approach recognizes that the security of this program depends on the runtime value of `c`, and instruments it to track whether `c` refers to a high-security channel or a low-security channel, in order to intervene only in the latter case. We use sets of levels in our security types in order to track the different possibilities during our static analysis (e.g., `c` is $\{L, H\}$ after the conditional).

To push the permissiveness a little more, we treat value variables similarly as channel variables. Take Listing 2 for example.

```

if lowValue > 0
  then x := lowValue  (*x is L*)
  else x := highValue (*x is H*)
end;
send x to lowChannel

```

Listing 2: Statically uncertain variable level

Traditionally, the level of variable `x` after the `if` command would be statically approximated as being high. Such an approximation would cause the program to be statically rejected at the `send` command. By tracking the possible levels of variables, we can detect that a safe executions can happen and choose to instrument the program instead.

Progress channels. The progress of a program, which is observable through its outputs, can also reveal information. For example, in the following program, whether or not the output on the low-security channel occurs depends on whether the preceding loop terminates, which in turn depends on confidential information.

```

while highValue > 0 do
  skip
end; (*loop may diverge*)
send 42 to lowChannel

```

Listing 3: Progress leak

Although this example leaks only 1 bit of information, progress channels can be used to leak a significant amount of information [7]. The most common way to prevent leaks through progress channels is to forbid loops whose execution depends on confidential information [11, 12], but it leads to the rejection of many secure programs, such as the following.

```

while highValue > 0 do
  highValue := highValue - 1
end; (*loop always terminates*)
send 42 to lowChannel

```

Listing 4: Loop that always terminates

To accept such programs, we follow Moore et al. [13] and use an oracle (conservative and assumed correct) to statically determine the termination behavior of loops. If the oracle determines that a loop always terminates (like the one in Listing 4), then we know that no following output could leak information through progress channels. On the other hand, if the oracle says that it may diverge (like the one in Listing 3), then we must take into account the fact that an output following the loop’s execution could leak information.

In our approach, the oracle is a parameter and is based on termination analysis methods brought from the literature such as the one described in Cook et al. [14].

The inlined monitor itself may introduce leaks of confidential values through progress channels. The decision that a monitor makes to allow or not an execution, may depend on confidential values. Thus, outputs following

this decision could leak these confidential values. Consider the following example:

```

if highValue > 0
  then skip
  else send highValue to c
end;
send 42 to lowChannel

```

Listing 5: Leaking through monitor decision

If during execution `c` is `lowChannel`, then the decision to allow this `send` leaks the value of the condition, `highValue > 0`, to `lowChannel`.

The danger of a monitor to leak confidential values through progress channels becomes greater when the monitor enforces confidentiality policies on a multilevel lattice. Consider the following example, taken from Kozyri et al. [15].

```

if medValue > 0
  then x := highValue  (*x is H *)
  else x := lowValue   (*x is L *)
end;
send x to c
send 1 to lowChannel

```

Listing 6: Variable level sensitivity

Suppose that the unknown channel `c` of the first `send` command happens to be of medium security level at runtime. If the monitor does not halt the execution of the first `send` command, then the second `send` command would be reached. Then, if the output to `lowChannel` occurred, it would leak information about `medValue`.

To prevent such a leak, we combine our previous work [6] with the dynamic mechanism of Kozyri et al. [15] and keep track of the level of information that could be leaked through progress channels. More specifically, for Listing 5, we detect that one of the branch of the `if` contains a `send` command that may be blocked during runtime and update the level of information that could be leaked through progress channels to H (i.e., the level of the condition), hence preventing future outputs to channels that are lower than H . Similarly, for Listing 6, we detect that the first `send` command may be blocked during execution and that the value of its parameters (`x` and `c`) depends of information of level M . For this reason, the level of information that could be leaked through progress channels is updated to M after that first output, hence preventing future outputs to channels that are lower than M .

1.2. Structure

In Section 2, we present the imperative language used to illustrate our approach. Section 3 defines the non-interference property. Section 4 describes our typed-based instrumentation mechanism, explains the type system, and presents the target language in which the instrumented programs are written; it is an extension of the source language with dynamic security levels. Section 5 proves that

the instrumented programs are non-interferent. Section 6 presents two ways to increase the precision and permissiveness of our monitor. Section 7 is a summary of related work. Finally, we conclude in Section 8.

2. Source language

Source programs are written in a simple imperative language with commands for receiving and sending information.

2.1. Syntax

Let \mathcal{V} be a set of identifiers for variables, and \mathcal{C} a set of predefined communication channels. The syntax is as follows.

(variables)	x	\in	$\mathcal{V} \cup \mathcal{C}$
(integer constants)	n	\in	\mathbb{Z}
(expressions)	e	$::=$	$x \mid n \mid e_1 \text{ op } e_2 \mid \text{read } x$
(commands)	cmd	$::=$	$\text{skip} \mid x := e \mid \text{if } e \text{ then } cmd_1 \text{ else } cmd_2 \text{ end} \mid$ $\text{while } e \text{ do } cmd \text{ end} \mid cmd_1; cmd_2 \mid \text{send } x_1 \text{ to } x_2$

Values are integers (we use zero for false and nonzero for true), or channel names. The symbol **op** stands for arithmetic or logic binary operators on integers.

We suppose that the interaction of a program with its environment (which can be a user or another program) is done through channels. Channels can be, for example, files, users, network channels, keyboards, computer screens, etc. Without loss of generality, we consider that each channel consists of one value. The expression **read** x returns the current value in channel x (without modifying the channel's content). Command **send** x_1 **to** x_2 sends the value of variable x_1 to channel x_2 and overwrites the current value in the channel. In other words, it outputs the value of x_1 to channel x_2 . The security levels of these channels are designated in advance by some security administrator.

Note that in order to keep the syntax simple, we chose to enforce certain constraints in the semantics and type system rather than in the syntax. For example, it is the semantics and type system that verify that x_1 is an integer variable and x_2 is a channel during the execution of a **send** command.

2.2. Semantics

A memory $m : \mathcal{V} \uplus \mathcal{C} \rightarrow \mathbb{Z} \uplus \mathcal{C}$ is a partial map from variables and channels to values, where the value of a channel is the last value sent to this channel. More precisely a memory is the disjoint union of two (partial) maps of the following form:

$$m_v : \mathcal{V} \rightarrow \mathbb{Z} \uplus \mathcal{C}, \quad m_c : \mathcal{C} \rightarrow \mathbb{Z},$$

where \uplus stands for the disjoint union operator. We omit the subscript whenever the context is clear. We write

$m(e) = r$ to indicate that the evaluation of expression e under memory m returns r .

The semantics of the source language is mostly standard and is illustrated in Figure 1. Program configurations are tuples $\langle cmd, m, o \rangle$ where cmd is the command to be evaluated, m is the current memory and o is the current output trace. A transition between two configurations is denoted by the \rightarrow symbol. We write \rightarrow^* for the reflexive transitive closure of the \rightarrow relation.

We write $v :: vs$ for sequences where v is the first element of the sequence, and vs is the rest of the sequence. We write ϵ for the empty sequence. An output trace is a sequence of output events: it is of the form $o = (v_0, ch_0) :: (v_1, ch_1) :: \dots$ where $v_k \in \mathbb{Z}$ is an integer value, and $ch_k \in \mathcal{C}$ is a channel, $k \in \mathbb{N}$. The rule for sending a value appends a new output event to the end of the trace. We abuse notation and write $o :: (v, ch)$ to indicate event (v, ch) appended to trace o .

We write $\langle cmd, m, \epsilon \rangle \downarrow o$ if execution of configuration $\langle cmd, m, \epsilon \rangle$ can produce trace o , where o may be finite or infinite. For finite o , $\langle cmd, m, \epsilon \rangle \downarrow o$ holds if there is a configuration $\langle cmd', m', o \rangle$ such that $\langle cmd, m, \epsilon \rangle \rightarrow^* \langle cmd', m', o \rangle$. For infinite o , $\langle cmd, m, \epsilon \rangle \downarrow o$ holds if for all traces o' such that o' is a finite prefix of o , we have $\langle cmd, m, \epsilon \rangle \downarrow o'$.

3. Security

For our purposes, we assume a finite lattice of security levels $(\mathcal{L}, \sqsubseteq)$ which contains at least two elements: L for the bottom of the lattice and H for the top of the lattice, i.e. $\forall \ell \in \mathcal{L}, L \sqsubseteq \ell \wedge \ell \sqsubseteq H$. We define an execution as ℓ -secure if the outputs to channels of level ℓ or lower do not reveal any information about the inputs of channels that are not lower than or equal to ℓ . This is a standard form of non-interference (e.g., [1, 2]) adapted for our particular language model.

Before formally defining non-interference, we first introduce some helpful technical concepts. The *projection of trace o to security level ℓ* , written $o \upharpoonright \ell$, is its restriction to output events whose channels' security levels are less than or equal to ℓ . Formally,

$$\begin{aligned} \epsilon \upharpoonright \ell &= \epsilon \\ ((v, ch) :: o) \upharpoonright \ell &= \begin{cases} (v, ch) :: (o \upharpoonright \ell) & \text{if } levelOfChan(ch) \sqsubseteq \ell \\ o \upharpoonright \ell & \text{otherwise} \end{cases} \end{aligned}$$

where $levelOfChan(ch)$ denotes the security level of channel ch (typically specified by the administrator).

We say that two memories m and m' are ℓ -equivalent if they agree on the content of variables (including channel variables) whose security levels are ℓ or lower.

$$\begin{array}{c}
\text{(SKIP)} \quad \langle \mathbf{skip}, m, o \rangle \longrightarrow \langle \mathbf{stop}, m, o \rangle \\
\\
\text{(ASSIGN-CH)} \quad \frac{e \in \mathcal{C}}{\langle x := e, m, o \rangle \longrightarrow \langle \mathbf{stop}, m[x \mapsto e], o \rangle} \\
\\
\text{(ASSIGN)} \quad \frac{e \notin \mathcal{C} \quad m(e) = r}{\langle x := e, m, o \rangle \longrightarrow \langle \mathbf{stop}, m[x \mapsto r], o \rangle} \\
\\
\text{(SEND)} \quad \frac{m(x_1) = v \in \mathbb{Z} \quad (ch = x_2 \in \mathcal{C} \vee ch = m(x_2) \in \mathcal{C})}{\langle \mathbf{send} \ x_1 \ \mathbf{to} \ x_2, m, o \rangle \longrightarrow \langle \mathbf{stop}, m[ch \mapsto v], o :: (v, ch) \rangle} \\
\\
\text{(SEQ1)} \quad \frac{\langle cmd_1, m, o \rangle \longrightarrow \langle \mathbf{stop}, m', o' \rangle}{\langle cmd_1; cmd_2, m, o \rangle \longrightarrow \langle cmd_2, m', o' \rangle} \\
\\
\text{(SEQ2)} \quad \frac{\langle cmd_1, m, o \rangle \longrightarrow \langle cmd'_1, m', o' \rangle \quad cmd'_1 \neq \mathbf{stop}}{\langle cmd_1; cmd_2, m, o \rangle \longrightarrow \langle cmd'_1; cmd_2, m', o' \rangle} \\
\\
\text{(IF)} \quad \frac{m(e) \neq 0 \implies i = 1 \quad m(e) = 0 \implies i = 2}{\langle \mathbf{if} \ e \ \mathbf{then} \ cmd_1 \ \mathbf{else} \ cmd_2 \ \mathbf{end}, m, o \rangle \longrightarrow \langle cmd_i, m, o \rangle} \\
\\
\text{(LOOP1)} \quad \frac{m(e) \neq 0}{\langle \mathbf{while} \ e \ \mathbf{do} \ cmd \ \mathbf{end}, m, o \rangle \longrightarrow \langle cmd; \mathbf{while} \ e \ \mathbf{do} \ cmd \ \mathbf{end}, m, o \rangle} \\
\\
\text{(LOOP2)} \quad \frac{m(e) = 0}{\langle \mathbf{while} \ e \ \mathbf{do} \ cmd \ \mathbf{end}, m, o \rangle \longrightarrow \langle \mathbf{stop}, m, o \rangle}
\end{array}$$

Figure 1: Semantics of the source language

Definition 1 (Progress-sensitive non-interference) *We say that a program p satisfies progress-sensitive non-interference if for any $\ell \in \mathcal{L}$, and for any two memories m and m' that are ℓ -equivalent, and for any trace o such that $\langle p, m, \epsilon \rangle \downarrow o$, then there is some trace o' , such that $\langle p, m', \epsilon \rangle \downarrow o'$ and $o \upharpoonright \ell = o' \upharpoonright \ell$.*

This definition of non-interference is progress-sensitive in that it assumes that an observer can distinguish an execution that will not produce any additional observable output (due to termination or divergence) from an execution that will make progress and produce additional observable output. Progress-insensitive definitions of non-interference typically weaken the requirement that $o \upharpoonright \ell = o' \upharpoonright \ell$ to instead require that $o \upharpoonright \ell$ is a prefix of $o' \upharpoonright \ell$, or vice versa.

Among the previously presented examples, only List-

ing 4 satisfies progress-sensitive non-interference. Nevertheless, we statically accept the programs of Listings 1, 2, 5 and 6, since we transform them into programs that satisfy progress-sensitive non-interference.

4. Type-based instrumentation

We enforce non-interference by translating a source program to a target program that tracks the security levels of its variables and checks the security of output commands. The translation performs a type-based static analysis of the source program, and rejects programs that contain outputs whose executions will always be blocked by our monitor (i.e., the translation fails).

In this section, we first present the security types for the source language followed by the description of the target language, which extends the source language with runtime representation of security levels. We then present the translation from the source language to the target language.

4.1. Source language types

Source language types are defined according to the following grammar.

$$\begin{array}{ll}
(\text{security labels, } Lab) & \ell ::= \mathcal{P}(\mathcal{L}) \setminus \{\emptyset\} \\
(\text{value types, } ValT) & \sigma ::= int_\ell \mid int_\ell \ chan \\
(\text{variable types, } VarT) & \tau ::= \sigma_\ell
\end{array}$$

Security labels are non-empty sets of security levels. They represent the possible security levels of a variable at runtime. If a security label contains more than one element, it means that its security level is statically unknown (see Listings 1 and 2 on pages 2 and 3 for examples).

Value types are the types of integers and channels. Type int_ℓ is the type of integers whose values are of security level ℓ , and type $int_\ell \ chan$ is the type of a channel whose values are of security level ℓ .

Variable types associate a security level with a value type. Intuitively, $\sigma_{\ell'}$ represents the type of a variable whose value type is σ , and whose variable type is ℓ' . The latter is an upper bound of the information level influencing the value of the variable. When a variable type ℓ' is associated with a value type ℓ for a channel, it means that the sensitivity of the content of the channel is ℓ , and the sensitivity of the channel itself is ℓ' . This is the same approach that was used by Bedford et al. [6].

For example, consider the program in Listing 1, page 2. Immediately following the conditional command, the type system gives variable c the type $(int_{\{L,H\}} \ chan)_{\{L\}}$. This type reflects that only low information determines which channel is assigned to c (i.e., variable `lowValue` determines c 's value), and whether c is a low channel or a high channel is statically unknown.

Similarly, when a variable type ℓ' is associated with a value type ℓ for an integer, it means that the sensitivity of the integer is ℓ , and the sensitivity of ℓ is ℓ' . This is

the same approach that was used by Kozyri et al. [15] to prevent leakage through guarded sends.

For example, consider the program in Listing 6, page 3. Immediately following the conditional command, the type system gives variable x the type $(int_{\{L,H\}})_{\{M\}}$. This type reflects that an information of medium level determines which value is assigned to x (i.e., variable `medValue` determines x 's value), and that the information contained within x is either of low security level or of high security level. It is necessary to keep track of the context level in which x has been assigned its value. This information is used to halt the execution of the second `send` and prevent observers from deducing that `medValue` is less or equal to 0.

4.2. Sets of levels

Note that we use sets of levels not only to increase the precision of the analysis, but also because we *have to* due to our use of channel variables. Indeed, one of the issues that we encountered is the fact that we cannot conservatively approximate the level of a channel variable, due to the fact that we both read and write on channels. Listing 7 illustrates why.

```

if lowValue > 0 then
  c := lowChannel
else
  c := highChannel
end
send highValue to c (*Pessimist: c is L*)
x := read c (*Pessimist: c is H*)
send x to lowChannel

```

Listing 7: We cannot be pessimistic about channel variables

After the conditionals, c has type $(int_{\{L,H\}}chan)_{\{L\}}$ because it contains either a low or high channel and its value is assigned in a context of level L . Our typing system accepts this program, but makes sure that a runtime checks are inserted. If the condition `lowValue > 0` happens to be true at runtime, then the execution of the program will be stopped at the first `send` command thanks to the inserted runtime checks. Similarly, if the condition is false, then the execution of the program will be stopped at the second `send`. The uncertainty is unavoidable in the presence of flow sensitivity and channel variables. Indeed, we point out that we cannot be pessimistic about the level of c in this program. The output command suggests that a safe approximation for c would be a low security level. Yet, the input command suggests that a safe approximation for c would be a *high* security level, which contradicts the previous observation.

Consequently, in order to accept the program in Listing 7, we chose to use sets of security levels. As a consequence, we will obtain fewer false positives as we do not consider the worst possible case in our analysis, we leave it to the execution to check whether the information flow turns out to be secure or not.

We derive two relations that allow us to compare the sets of possible levels.

Definition 2 *The relations \sqsubseteq_s , surely less than, and \sqsubseteq_m , maybe less than, are defined as follows*

$$\begin{aligned} \ell_1 \sqsubseteq_s \ell_2 &= (\forall e_1 \in \ell_1, e_2 \in \ell_2. e_1 \sqsubseteq e_2). \\ \ell_1 \sqsubseteq_m \ell_2 &= (\exists e_1 \in \ell_1, e_2 \in \ell_2. e_1 \sqsubseteq e_2). \end{aligned}$$

Intuitively, we have $\ell \sqsubseteq_s \ell'$ when we can be sure statically that $\ell \sqsubseteq \ell'$ will be true at runtime, and we have $\ell \sqsubseteq_m \ell'$ when it is possible that $\ell \sqsubseteq \ell'$ at runtime.

Definition 3 *The supremum on sets of levels is defined as follows*

$$\ell_1 \sqcup \ell_2 = \{e_1 \sqcup e_2 \mid e_1 \in \ell_1, e_2 \in \ell_2\}.$$

Recall that sets of levels represent the possible security levels that can occur. Hence a supremum between specific levels must be mimicked, between sets of levels, by the set of all possible results of such suprema. Consequently, we call it a supremum even if the corresponding preorder relation is not used here.

We instrument source programs to track at runtime the security levels that are statically unknown. In order to track these security levels, our target language allows their runtime representation.

4.3. Syntax and semantics of target language

Our target language is inspired by the work of Zheng and Myers [8], which introduced a language with first-class security levels, and a type system that soundly enforces non-interference in this language. The syntax of our target language is defined as follows:

(variables)	x	\in	$\mathcal{V} \cup \mathcal{C}$
(level variables)	\tilde{x}	\in	\mathcal{V}_{level}
(integer constants)	n	\in	\mathbb{Z}
(basic levels)	k	\in	\mathcal{L}
(level expressions)	l	$::=$	$k \mid \tilde{x} \mid l_1 \sqcup l_2$
(integer expressions)	exp	$::=$	$x \mid n \mid exp_1 \mathbf{op} exp_2 \mid$ $\mathbf{read} x$
(expressions)	e	$::=$	$exp \mid l$
(commands)	cmd	$::=$	$\mathbf{skip} \mid (x_1, \dots, x_n) := (e_1, \dots, e_n) \mid$ $\mathbf{if} e \mathbf{then} cmd_1 \mathbf{else} cmd_2 \mathbf{end} \mid cmd_1; cmd_2 \mid$ $\mathbf{while} e \mathbf{do} cmd \mathbf{end} \mid \mathbf{send} x_1 \mathbf{to} x_2 \mid$ $\mathbf{if} l_1 \sqsubseteq l_2 \mathbf{then} (\mathbf{send} x_1 \mathbf{to} x_2) \mathbf{else} \mathbf{fail} \mathbf{end}$

The main difference between our source language and target language is that it adds support for *level variables*, a runtime representation of security levels. These level variables will allow target programs to verify that certain conditions are met before sending the contents of a variable to a channel. This is the goal of the new `send` command, nested in a conditional – we call it a *guarded send*.

For simplicity, we assume that security levels can be stored only in a restricted set of variables $\mathcal{V}_{level} \subseteq \mathcal{V}$. Thus,

the variable part m_v of a memory m now has the following type

$$m_v : (\mathcal{V}_{level} \rightarrow \mathcal{L}) \uplus (\mathcal{V} \setminus \mathcal{V}_{level} \rightarrow \mathbb{Z} \uplus \mathcal{C})$$

Furthermore we assume that \mathcal{V}_{level} contains variables $_pc$ and $_hc$, and two level variables x_{val} and x_{ctx} associated with each variable $x \in \mathcal{V} \setminus \mathcal{V}_{level}$. Variables $_pc$ and $_hc$ hold the security levels of the context and halting context respectively, they represent the security level in which a command is executed. We will discuss in more details what they represent in Section 4.4. Variables x_{val} and x_{ctx} are used to track the security levels of variables at runtime. For example, if x is a channel variable of security type $(int_\ell chan)_{\ell'}$, then the values of these variables should be $x_{val} = \ell$ and $x_{ctx} = \ell'$ (this will be ensured by our instrumentation).

The simultaneous assignment $(x_1, \dots, x_n) := (e_1, \dots, e_n)$ is introduced for the sake of clarity. Any assignment implies an immediate update of the concerned level variables. For all common commands, the semantics of the target language is the same as in the source language.

4.4. Instrumentation as a type system

Our instrumentation algorithm is specified as a type system in Figure 2. Its goal is to inline monitor actions in the program under analysis, thereby generating a safe version of the program, or to reject the program when it contains obvious leaks of information. The inlined actions are essentially updates of level variables and checks on these variables in order to control the execution of potentially leaking **send** commands. After a check, a **send** command is either executed if it is safe or its execution is prevented and the program is aborted.

The typing rules of expressions have judgements of the form $\Gamma \vdash e : \sigma_\ell$, stating that σ_ℓ is the type of e under the typing environment $\Gamma : \mathcal{V} \uplus \mathcal{C} \rightarrow VarT$. The instrumentation judgements are of the form

$$\Gamma, pc, hc \vdash cmd : t, h, \Gamma', \llbracket cmd \rrbracket$$

where Γ, Γ' are typing environments (initially empty) and cmd is the command under analysis. The program context, $pc \in Lab$, is used to keep track of the security level in which a command is executed, in order to detect implicit flows. The halting context, $hc \in Lab$, is used to detect progress channels leaks. It represents the level of information that could have caused the program to *halt* (due to a failed guarded send command) or *diverge* (due to an infinite loop). In other words, it is the level of information that could be leaked through progress channels by an output. Variable $h \in Lab$ corresponds to the halting context after executing command cmd . The termination type $t \in \mathcal{T} = \{T, D, M_\ell\}$ of a command is used to keep the halting context up to date. We distinguish three termination types: T means that a command terminates for all memories, D means that a command diverges for all

memories, M_ℓ means that a command's termination is unknown statically; the subscript is used to indicate on which level(s) the termination depends. For example, the termination of the loop in Listing 3 is $M_{\{H\}}$ because it can either terminate or diverge at runtime, and this depends on information of level H . The loop in Listing 4 on the other hand is of termination type T because, no matter what the value of variable h is, it will always eventually terminate. Similarly, a loop whose condition is always true will have termination type D since it always diverges. Recall that, of course, the precision of this analysis depends on the precision of the oracle. Variable $\llbracket cmd \rrbracket$ is the generated instrumented version of command cmd and is often presented using a macro whose name starts with *gen*.

The instrumentation of a program p begins by inserting commands to initialize a few level variables: $_pc$, $_hc$ are initialized to L , as well as the level variables x_{ctx} and x_{val} for each variable $x \in \mathcal{V}$ appearing in p . Similarly, level variables c_{ctx} and c_{val} associated with each channel c used in p are also initialized, but the latter rather gets initialized to $levelOfChan(c)$, which is an input parameter for the analysis. After the initialization, the instrumentation is given by the rules of Figure 2. We now explain these rules.

Rules (S-CHAN) and (S-INT) specify the type of channel and integer constants respectively.

Rule (S-VAR) infers the type of a variable from the environment Γ .

Rule (S-OP) infers the type of an expression from the types of its operands and it reflects the fact that operations on channels are not allowed.

Rule (S-READ) returns the security type of the current value in channel c . This rule highlights the striking similarity between type $int_\ell chan_{\ell_c}$ of a channel and type $(int_\ell)_{\ell_c}$ of a variable that may subsequently store expression **read** c . Here, the security label ℓ of channel's value type becomes the security label of variable's value type. Similarly for the security label ℓ_c of channel's variable type. To the best of our knowledge, this is the first work that tries to use similar type structures for channels and ordinary variables in an enforcement mechanism.

Rule (S-ASSIGN) updates the type of x in the environment with the type of expression e . The supremum with pc is used to prevent implicit flows. Its instrumentation is given by the following macro, which represents a simultaneous assignment $x := e, x_{val} := e_{val}$, and $x_{ctx} := e_{ctx} \sqcup _pc$.

$$\begin{aligned} \text{genassign} = \\ (x, x_{val}, x_{ctx}) := (e, e_{val}, e_{ctx} \sqcup _pc) \end{aligned}$$

We write e_{ctx} to represent the expression made of the supremum of variables appearing in expression e . For example if $e = x + \text{read } c$, then $e_{ctx} = x_{ctx} \sqcup c_{ctx}$. If $e = x + y$ then $e_{ctx} = x_{ctx} \sqcup y_{ctx}$. The idea is the same for e_{val} .

Rule (S-SEND) checks whether a send command is statically safe by requiring $(pc \sqcup hc \sqcup \ell_x \sqcup \ell'_x \sqcup \ell'_c) \sqsubseteq_s \ell_c$ (i.e., all possible values of the left-hand side are always lower or equal to the right-hand side). The variables on the left-hand side correspond to the level of information that can

$$\begin{array}{c}
\text{(S-CHAN)} \\
\frac{\text{levelOfChan}(ch) = \ell}{\Gamma \vdash ch : (int_{\{\ell\}} chan)_{\{\ell\}}} \\
\\
\text{(S-INT)} \quad \Gamma \vdash n : (int_{\{L\}})_{\{L\}} \\
\\
\text{(S-VAR)} \quad \frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau} \\
\\
\text{(S-READ)} \quad \frac{\Gamma \vdash c : int_{\ell} chan_{\ell_c}}{\Gamma \vdash \mathbf{read} c : (int_{\ell})_{\ell_c}} \\
\\
\text{(S-OP)} \quad \frac{\Gamma \vdash e_1 : (int_{\ell_1})_{\ell'_1} \quad \Gamma \vdash e_2 : (int_{\ell_2})_{\ell'_2}}{\Gamma \vdash e_1 \mathbf{op} e_2 : (int_{\ell_1 \sqcup \ell_2})_{\ell'_1 \sqcup \ell'_2}} \\
\\
\text{(S-SKIP)} \quad \Gamma, pc, hc \vdash \mathbf{skip} : T, hc, \Gamma, \mathbf{skip} \\
\\
\text{(S-ASSIGN)} \quad \frac{\Gamma \vdash e : \sigma_{\ell'_e}}{\Gamma, pc, hc \vdash x := e : T, hc, \Gamma[x \mapsto \sigma_{pc \sqcup \ell'_e}], \mathbf{genassign}} \\
\\
\text{(S-SEND)} \quad \frac{\Gamma(x) = (int_{\ell_x})_{\ell'_x} \quad \Gamma(c) = (int_{\ell_c} chan)_{\ell'_c} \quad (pc \sqcup hc \sqcup \ell_x \sqcup \ell'_x \sqcup \ell'_c) \sqsubseteq_s \ell_c}{\Gamma, pc, hc \vdash \mathbf{send} x \mathbf{to} c : T, hc, \Gamma, \mathbf{send} x \mathbf{to} c} \\
\\
\text{(S-GSEND)} \quad \frac{\Gamma(c) = (int_{\ell_c} chan)_{\ell'_c} \quad \Gamma(x) = (int_{\ell_x})_{\ell'_x} \quad (pc \sqcup hc \sqcup \ell_x \sqcup \ell'_x \sqcup \ell'_c) \not\sqsubseteq_s \ell_c \quad (pc \sqcup hc \sqcup \ell_x \sqcup \ell'_x \sqcup \ell'_c) \sqsubseteq_m \ell_c}{\Gamma, pc, hc \vdash \mathbf{send} x \mathbf{to} c : T, pc \sqcup hc \sqcup \ell'_x \sqcup \ell'_c, \Gamma, \mathbf{gensend}} \\
\\
\text{(S-IF)} \quad \frac{\Gamma \vdash e : (int_{\ell_e})_{\ell'_e} \quad pc' = pc \sqcup \ell_e \sqcup \ell'_e \quad h_d = \mathit{hasGSend}(\{cmd_1, cmd_2\}, pc') \quad h = \bigcup_{j \in \{1,2\}} (h_j \sqcup h_d \sqcup \mathit{level}(t_1 \oplus_{pc'} t_2)) \quad \Gamma, pc', hc \vdash cmd_j : t_j, h_j, \Gamma_j, \llbracket cmd_j \rrbracket \quad j \in \{1, 2\}}{\Gamma, pc, hc \vdash \mathbf{if} e \mathbf{then} cmd_1 \mathbf{else} cmd_2 \mathbf{end} : t_1 \oplus_{pc'} t_2, h, \Gamma_1 \sqcup_{pc'} \Gamma_2, \mathbf{genif}} \\
\\
\text{(S-LOOP1)} \quad \frac{\Gamma = \Gamma \sqcup_{pc'} \Gamma' \quad \Gamma \vdash e : (int_{\ell_e})_{\ell'_e} \quad O(e, cmd, \Gamma) = t_o \quad pc' = pc \cup (pc \sqcup \ell_e \sqcup \ell'_e) \quad hc' = hc \cup (hc \sqcup \mathit{level}(t') \sqcup h') \quad h_d = \mathit{hasGSend}(\{cmd\}, pc') \quad \Gamma, pc', hc' \vdash cmd : t', h', \Gamma', \llbracket cmd \rrbracket}}{\Gamma, pc, hc \vdash \mathbf{while} e \mathbf{do} cmd \mathbf{end} : t_o, h_d \sqcup h' \sqcup \mathit{level}(t_o), \Gamma, \mathbf{genwhile}} \\
\\
\text{(S-LOOP2)} \quad \frac{\Gamma \neq \Gamma \sqcup_{pc'} \Gamma' \quad \Gamma \vdash e : (int_{\ell_e})_{\ell'_e} \quad \Gamma, pc \sqcup \ell_e \sqcup \ell'_e, hc \vdash cmd : t', h', \Gamma', \llbracket cmd \rrbracket \quad pc' = pc \cup (pc \sqcup \ell_e \sqcup \ell'_e) \quad hc' = hc \cup (hc \sqcup \mathit{level}(t') \sqcup h') \quad \Gamma \sqcup_{pc'} \Gamma', pc', hc' \vdash \mathbf{while} e \mathbf{do} cmd \mathbf{end} : t'', h'', \Gamma'', \mathbf{instCode}}{\Gamma, pc, hc \vdash \mathbf{while} e \mathbf{do} cmd \mathbf{end} : t'', h'', \Gamma'', \mathbf{instCode}} \\
\\
\text{(S-SEQ1)} \quad \frac{\Gamma, pc, hc \vdash cmd_1 : D, h, \Gamma_1, \llbracket cmd_1 \rrbracket}{\Gamma, pc, hc \vdash cmd_1; cmd_2 : D, h, \Gamma_1, \llbracket cmd_1 \rrbracket}} \\
\\
\text{(S-SEQ2)} \quad \frac{t_1 \neq D \quad \Gamma, pc, hc \vdash cmd_1 : t_1, h_1, \Gamma_1, \llbracket cmd_1 \rrbracket \quad \Gamma_1, pc, h_1 \vdash cmd_2 : t_2, h_2, \Gamma_2, \llbracket cmd_2 \rrbracket}}{\Gamma, pc, hc \vdash cmd_1; cmd_2 : t_1 \text{ ; } t_2, h_2, \Gamma_2, \llbracket cmd_1 \rrbracket; \llbracket cmd_2 \rrbracket}}
\end{array}$$

Figure 2: Instrumentation and typing rules for the source language

be revealed by the output to c . If so, the instrumentation inserts the send command as it is.

Rule (S-GSEND) checks whether a send command may be safe, by requiring $(pc \sqcup hc \sqcup \ell_x \sqcup \ell'_x \sqcup \ell'_c) \sqsubseteq_m \ell_c$. The instrumentation then transforms it into a guarded send, as follows

```
gensend =
  if _pc  $\sqcup$  _hc  $\sqcup$  x_val  $\sqcup$  x_ctx  $\sqcup$  c_ctx  $\sqsubseteq$  c_val then
    (send x to c)
  else
    fail
  end;
  _hc := _pc  $\sqcup$  _hc  $\sqcup$  x_ctx  $\sqcup$  c_ctx;
```

The halting context must record the possible failure of a guarded send at runtime, and hence, it is updated with the level of information that influences its success/failure. Particularly, the halting context is updated with the sensitivity of the context and of the two variables involved, the channel variable [6] and the regular variable [15]. For example, Listing 8 shows why c_{ctx} must be included in this update.

```
if x > 0 (*H at runtime *)
  then c := lowChannel
  else c := highChannel
end;
send highValue to c (*guarded send*)
send lowValue to lowChannel
```

Listing 8: Dangerous runtime halting

Assume that x is high and false at runtime. Then the first guarded send is accepted, but allowing an output on a low security channel subsequently would leak information about x . This is because, had x been true, then the first send would have failed. However, c has level $int_{\{L,H\}} chan_{\{H\}}$ after the conditional. Updating $_hc$ with c_{ctx} will affect the check of all subsequent guarded send and prevent such leaks.

If none of the send rules can be applied, then the program is statically rejected.

Before explaining the rules for composed commands, we first need to define a few functions and operators. For the conditional rules, we need a supremum of environments.

Definition 4 *The supremum of two environments is given as $dom(\Gamma_1 \sqcup_{pc} \Gamma_2) = dom(\Gamma_1) \cup dom(\Gamma_2)$, and*

$$(\Gamma_1 \sqcup_{pc} \Gamma_2)(x) = \begin{cases} \Gamma_i(x) & \text{if } x \in dom(\Gamma_i) \setminus dom(\Gamma_j), \\ & \{i, j\} = \{1, 2\} \vee \Gamma_1(x) = \Gamma_2(x) \\ (int_{\ell_1 \cup \ell_2} chan)_{(\ell'_1 \cup \ell'_2) \sqcup_{pc}} & \text{if } \Gamma_1(x) = (int_{\ell_1} chan)_{\ell'_1} \wedge \\ & \Gamma_2(x) = (int_{\ell_2} chan)_{\ell'_2} \wedge \\ & \Gamma_1(x) \neq \Gamma_2(x) \\ (int_{\ell_1 \cup \ell_2})_{(\ell'_1 \cup \ell'_2) \sqcup_{pc}} & \text{if } \Gamma_1(x) = (int_{\ell_1})_{\ell'_1} \wedge \\ & \Gamma_2(x) = (int_{\ell_2})_{\ell'_2} \wedge \\ & \Gamma_1(x) \neq \Gamma_2(x) \\ Error & \text{otherwise.} \end{cases}$$

When a typing inconsistency occurs, e.g., when a variable is used as an integer in one branch and as a channel in another, the analysis stops and an error is returned, causing the program to be statically rejected.

The function $level : \mathcal{T} \rightarrow \mathcal{P}(\mathcal{L}) \setminus \{\emptyset\}$ returns the *termination level* (i.e., the level that the termination depends on) and is defined as:

$$level(t) = \begin{cases} \{L\} & \text{if } t \in \{T, D\} \\ \ell & \text{if } t = M_\ell \end{cases}$$

Two operators are used to compose terminations types, \oplus , used in the typing of conditionals, and \circledast , used in the typing of sequences. They are defined as follows.

$$t_1 \oplus_{pc} t_2 = \begin{cases} t_1 & \text{if } t_1 = t_2 \in \{T, D\} \\ M_{pc \sqcup (\ell_1 \cup \ell_2)} & \text{otherwise,} \\ & \ell_1 = level(t_1), \ell_2 = level(t_2) \end{cases}$$

$$t_1 \circledast t_2 = \begin{cases} M_{\ell_1 \cup \ell_2} & \text{if } t_1 = M_{\ell_1} \text{ and } t_2 = M_{\ell_2} \\ t_i & \text{if } t_j = T, \{i, j\} = \{1, 2\} \\ D & \text{otherwise} \end{cases}$$

The following example illustrates the use of these operators:

```
if highValue then
  while 1 do skip end (* D *)
else
  skip (* T *)
end
```

The termination type of an **if** command is computed using the \oplus operator, from three parameters: the termination types of each of the two branches and the security level of the guard condition. Hence, in this example, we obtain $D \oplus_H T = M_{\{H\}}$.

One more function needs to be defined. Its motivation is given by the following example. Assume that in Listing 9 c turns out to be a low channel at runtime. If the last send command is reached and executed it would leak information about **highValue**. The same leak would happen if instead of the guarded send we had a diverging loop.

```
if highValue > 0 then
  if  $\ell \sqsubseteq \ell'$ 
    then send highValue to c
    else fail end;
  end;
  send lowValue to lowChannel
```

Listing 9: A guarded send can generate a progress leak

To prevent this kind of leak, we verify if the branches of a conditional contains guarded send commands using the function

$$hasGSend : \mathcal{P}(\text{Cmd}) \times \text{Lab} \rightarrow \text{Lab},$$

where Cmd is the set of commands and Lab the set of security labels. If one of the commands given in parameter contains a guarded send, then it returns the security label given in parameter, otherwise L is returned. This function is used to update the halting context to pc when there is a risk that one of the branch halts the execution. Without this update, hc could leak information about the condition in a subsequent send.

Rule (S-IF) specifies the typing of an **if** command. When typing an **if** command, we type the two branches under pc' , which is the supremum of the conditional's guard expression and current context. The resulting typing environments, Γ_1 and Γ_2 , then contain the security levels that variables may have after executing the first or second branch. The typing environment returned by the **if** is the join of those two, defined in Definition 4, so that it contains the possibilities of both branches. Similarly, variable h is used to calculate the possible values that the halting context may have after the conditional, hence the union.

Its instrumentation is given by the following macro:

```
genif =  _oldpcν := _pc;      _pc := _pc ⊔ e_val ⊔ e_ctx;
        if e then           [[cmd2]];
        _pc := _pc ⊔ e_val ⊔ e_ctx;  update(mv1)
        [[cmd1]];           end;
        update(mv2)       _pc := _oldpc
        else
```

where `update` is defined as follows

```
update(mv) =
  if mv = ∅ skip;
  else
    for each x ∈ mv  x_ctx := x_ctx ⊔ _pc;
    if _hc ∈ mv  _hc := _hc ⊔ _pc;
```

and where t_j is the termination type of cmd_j , mv_j is the set of modified variables in cmd_j (we include $_hc$ in this set if the termination of the two branches can differ i.e. if $\neg(t_1 = t_2 \in \{T, D\})$ or if at least one guarded send occurs in the other branch), and $e_{val} \sqcup e_{ctx}$ is the guard condition's level expression.

The instrumented code starts by saving the current context to $_oldpc^\nu$ (the symbol ν indicates that it is a fresh variable). The program context is updated with the security level of the guard condition. The **if** itself is then generated.

The function `update` generates the command `skip`; if the parameter set mv is empty otherwise it generates updates of the context level of each modified variable in the other branch as well as the update of $_hc$ if necessary. The underlying reason is to ensure that the value of these level variables is at least pc .

In a situation like the following listing, this function allows to update y 's level, to protect x .

```
y := 0;
if x > 0 then (*H at runtime*)
  y := 1
else skip end;
send x to lowChannel
```

Listing 10: Modified variables

Here, even if the else branch is taken at runtime, the level of variable y must be updated. Otherwise, information about x would be revealed by the send command (even with a guarded send).

Rules (S-LOOP1), (S-LOOP2) specify the **while** command typing. They involve computing a least fixed point to derive the right typing environment. This is necessary because of the flow sensitivity feature. Typing rule (S-LOOP2) is applied recursively until a least fixpoint is found, at which point (S-LOOP1) is applied and its result is returned. Note that, since the calculation of the fixpoint is defined in a constructive way, the rule will not accept any other fixpoints than a least fixpoint and it will always compute the same least fixpoint. The union operator is used to update the pc' and hc' variables so that we keep track of all their possible values. Due to our use of finite lattices, and the monotonicity of the union and supremum on levels, it is easy to show that this computation converges, the proof is given in Appendix B, Lemma 7. The typing relies on an oracle O that returns the termination type of the loop (t_o). It is worth noting that the call to the oracle is performed statically. Calling it dynamically would enhance precision, but increase significantly the overhead. If the loop contains guarded send commands, which could fail and reveal information about the condition of the loop, then we update the halting context to prevent this leak. The presence of at least one guarded send command is detected using the function `hasGSend()`.

```
genwhile = _oldpcν := _pc;      end;
          while e do           _pc := _pc ⊔ e_val ⊔ e_ctx;
          _pc := _pc ⊔ e_val ⊔ e_ctx;  update(mv);
          [[cmd]];              _pc := _oldpc
```

The inserted commands are similar to those of the **if** command. The level variables and halting context are updated after the loop in case an execution does not enter the loop. The context needs to be updated at the beginning of each iteration as the value, and hence level, of expression e may change.

Rule (S-SEQ1) is applied if cmd_1 always diverges; we then ignore cmd_2 , as it will never be executed. Otherwise, rule (S-SEQ2) is applied. The halting context returned is h_2 instead of $h_1 \sqcup h_2$ because h_2 already takes into account h_1 .

Examples of instrumented programs are available in Appendix A.

5. Soundness

In order to prove that the instrumented program generated by Figure 2 correctly enforces non-interference, we need to adapt the definitions of non-interference and ℓ -equivalent memories to our target language, because of level variables. Recall that a memory for the target language is the union of two maps of the following form:

$$\begin{aligned} m_v &: (\mathcal{V}_{level} \rightarrow \mathcal{L}) \uplus (\mathcal{V} \setminus \mathcal{V}_{level} \rightarrow \mathbb{Z} \uplus \mathcal{C}), \\ m_c &: \mathcal{C} \rightarrow \mathbb{Z}. \end{aligned}$$

We write $\text{dom}_l(m) := \text{dom}(m_v) \cap \mathcal{V}_{level} = m_v^{-1}(\mathcal{L})$. A memory m is called *complete* for a program p if

- $\{x_{\text{val}}, x_{\text{ctx}}\} \subseteq \text{dom}_l(m)$ for any $x \in \mathcal{C} \cup \text{dom}(m) \setminus \mathcal{V}_{level}$ that appears in p .
- $_pc$ and $_hc$ are in $\text{dom}_l(m)$
- if $c \in \mathcal{C}$ appears in p , then $m_v(c_{\text{val}}) = \text{levelOfChan}(c)$ and $m_v(c_{\text{ctx}}) = L$
- if $m(x) \in \mathcal{C}$ then $m_v(x_{\text{val}}) = \text{levelOfChan}(m(x))$.

The first two conditions ensure that level variables exist in the domain of the memory, whereas the last ones makes sure that it is compliant with the security policy for channels.

The definition of ℓ -equivalent memories, which is based on [15], must handle level variables. Whenever the level variable x_{ctx} corresponding to a variable x is such that $m(x_{\text{ctx}}) \sqsubseteq \ell$ in one memory, then it must have the same value in both memories, otherwise a leak can happen.

Definition 5 *ℓ -equivalent memories.* We say that two complete memories of the target language m_1 and m_2 are ℓ -equivalent, written $m_1 \equiv_\ell m_2$, iff they satisfy the following properties

1. if $m_i(_pc) \sqsubseteq \ell$ for some $i \in \{1, 2\}$, then $m_1(_pc) = m_2(_pc)$. The same property holds for $_hc$.
2. $x \in m_{i,v}^{-1}(\mathbb{Z} \cup \mathcal{C})$, $\wedge m_i(x_{\text{ctx}}) \sqsubseteq \ell$, for some $i = 1, 2$, then
 - $m_1(x_{\text{ctx}}) = m_2(x_{\text{ctx}})$
 - $m_1(x_{\text{val}}) = m_2(x_{\text{val}})$
 - if $m_1(x_{\text{val}}) \sqsubseteq \ell$ then $m_1(x) = m_2(x)$
3. $c \in \mathcal{C} \wedge \text{levelOfChan}(c) \sqsubseteq \ell \Rightarrow m_{1,c}(c) = m_{2,c}(c)$

It may seem surprising that the memories may differ on a level variable such as x_{val} ; this is because they may differ on the value of high variables. To see this, here is a variation of Listing 1.

```

if highValue > 0
  then c := highChannel
  else c := lowChannel
end;

```

Listing 11: The security type of c 's content is sensible

In this example, the content of c depends on a private condition, and hence its level variable c_{ctx} should be H ; moreover, variable c_{val} , containing the level of the content of c , may have different values in two ℓ -equivalent memories. Another example, for non-channel variables, is Listing 6, reproduced below, where `medValue` has security level M , with $L \sqsubseteq M \sqsubseteq H$.

```

if medValue > 0
  then x := highValue  (*x is H *)
  else x := lowValue   (*x is L *)
end;
send x to c
send 1 to lowChannel

```

In this example, variable x_{val} contains information of level M . If c turns out to be of level M at execution then reaching or not the last send command will depend on x_{val} . The program will be blocked or not, and this will reveal which branch was taken in the preceding conditional. Interestingly, this problem does not arise when the lattice is restricted to only two levels, $\{L, H\}$, as argued in Kozyri et al. [15].

Here is the definition of non-interference for the target language. The difference from Definition 1 is the requirement of the memories to be complete and the use of Definition 5.

Definition 6 *Progress-sensitive non-interference.* We say that a program p satisfies progress-sensitive non-interference if for any $\ell \in \mathcal{L}$, and for any two complete memories m and m' that are ℓ -equivalent, and for any trace o such that $\langle p, m, \epsilon \rangle \downarrow o$, then there is some trace o' , such that $\langle p, m', \epsilon \rangle \downarrow o'$ and $o \upharpoonright \ell = o' \upharpoonright \ell$.

Using these updated definitions, we prove that the instrumented programs are non-interferent.

Theorem 1 (Soundness of enforcement) *If a program p is well typed according to the type system of Figure 2, then the generated program $\llbracket p \rrbracket$ satisfies progress sensitive non-interference.*

We also show that the instrumentation preserves the semantics of the original program. That is, the instrumentation of a program p , written $\llbracket p \rrbracket$, produces exactly the same output as p as long as it is allowed to continue; it may be stopped at some point to prevent a leak. If m is a memory for the target language we write \hat{m} for the restriction of m to $\mathcal{V} \setminus \mathcal{V}_{level}$.

Theorem 2 (Semantics preservation) *Let p be a program, m a memory, and o, o' output traces. Then*

$$\begin{aligned} \langle \llbracket p \rrbracket, m, \epsilon \rangle \downarrow o &\Rightarrow \langle p, \hat{m}, \epsilon \rangle \downarrow o \\ \langle p, \hat{m}, \epsilon \rangle \downarrow o \wedge \langle \llbracket p \rrbracket, m, \epsilon \rangle \downarrow o' &\Rightarrow o \preceq o' \vee o' \preceq o \end{aligned}$$

where $o' \preceq o$ means that o' is a prefix of o .

The proofs are available in Appendix B.

6. Increasing precision and permissiveness

During the course of this work, we thought of two ways to improve the precision of our static analysis and permissiveness of our dynamic analysis. While we chose not to use them into this work (to keep things as simple as possible), we think they are worth pointing out.

6.1. Security type constraints

In Listing 13, only executions where c is a high channel will get past the first guarded send. For this reason, we can consider, for the rest of the analysis, that its type is $int_{\{H\}}chan_{\{L\}}$ instead of $int_{\{L,H\}}chan_{\{L\}}$.

```

if lowValue > 0 then
  c := lowChannel
else
  c := highChannel
end;
send highValue to c; (*will be transformed
into a guarded send*)
(*to reach here, c must be {H}*)
x := read c; (*so x is {H}*)
send x to lowChannel (*always leaks, so
statically rejected*)

```

Listing 13: Constraint on the security type of a channel variable

The same idea applies to integer variables. For example, we know that the instructions after the first **send** of Listing 14 will only be reached if variable x is low. For this reason, we can consider that x 's type after the **send** will be $(int_L)_{\{L\}}$ instead of $(int_{\{L,H\}})_{\{L\}}$.

```

if lowValue > 0 then
  c := lowChannel
else
  c := highChannel
end;
x := read c; (*x is {L,H}*)
send x to lowChannel; (*will be
transformed into a guarded send*)
(*to reach here, x must be {L}*)
send x to lowChannel (*no need to
transform into a guarded send*)

```

Listing 14: Constraint on the security type of an integer variable

Using these constraints in our static analysis would lead to the insertion of fewer guarded sends, and thus, less runtime overhead.

6.2. Conditional updates of the halting context

While we chose to always update the halting context $_hc$ after a guarded send with

$$_hc := _pc \sqcup _hc \sqcup x_{ctx} \sqcup c_{ctx},$$

there are cases where we can be more precise. One such case is illustrated in Listing 15.

```

if medValue then
  x := read highChannel
else
  x := read highChannel2
end;
send x to c; (*guarded send*)
send lowValue to lowChannel

```

Listing 15: Example where $_hc$ does not need to be updated with x_{ctx}

In this example, while x 's value may differ, its type is constant and equal to $(int_{\{H\}})_{\{M\}}$. Since the value of $medValue$ does not affect its type, it means that it has no influence on the guarded send's decision to block or allow the output. Hence, in this case, the update to $_hc$ variable after the guarded send does not need to include variable x_{ctx} .

Similarly, when a channel c whose type is constant is used in a guarded send command, the update to $_hc$ does not need to include variable c_{ctx} .

Hence, using conditional updates would allow the dynamic analysis to be more permissive as the updates to $_hc$ are less conservative.

7. Related work

There has been much research into language-based techniques for controlling information flow over the last two decades. In this section, we focus on hybrid techniques for information-flow control. Hybrid techniques are attractive as the combination of static and dynamic analyses offers potentially multiple advantages such as low runtime overhead, increased precision and flexibility.

Le Guernic et al. [3] present the first hybrid information-flow control monitor. The enforcement is based on a monitor that is able to perform static checks during the execution. The enforcement is not flow-sensitive. Le Guernic, in [16], extends this work to concurrent programs. Russo and Sabelfeld [5] generalize their work, presenting a series of hybrid monitors that differ on the action to perform in the event of a security violation. They also state that purely dynamic enforcements are more permissive than purely static enforcements but they cannot be used in case of flow-sensitivity. They propose a hybrid flow-sensitive enforcement based on calling static analysis during the execution. This enforcement is not progress sensitive.

Kozyri et al. [15] show that it is not trivial to design dynamic enforcement mechanisms that support general lattices, are flow-sensitive, and do not leak information through termination enforced by the monitor. Their mechanism supports all these three features by using *metalabels*, which are labels on labels. A metalabel represents the sensitivity of the corresponding label, in the same way that ℓ represents the sensitivity of σ in a variable type σ_ℓ .

Bedford et al. [17] generate instrumented code, enforcing information flow based on static analysis (i.e., an

information-flow monitor is inlined). The approach supports channel variables and is flow sensitive, but does not take into account leaks due to progress. Also, the inlined monitor does not use dynamic security levels, but employs a heavy-handed approach which is not as amenable to standard optimization techniques as the present one. The target language is not formally defined and no soundness proof of the instrumented code is provided.

Moore et al. [13] consider precise enforcement of flow-insensitive progress-sensitive security. Progress sensitivity is also based on an oracle’s analysis, but they call upon it dynamically while we do it statically. We have also introduced additional termination types to increase the permissiveness of the monitor.

Chudnov and Naumann [18] inline a flow-sensitive hybrid monitor (based on a monitor of Russo and Sabelfeld [5]) and extend it to Javascript [19]. They prove its soundness by showing that the execution of the inlined monitor is bisimilar to the execution of a non-inlined monitor. We inline a flow-sensitive progress-sensitive hybrid monitor and, as we did not have already have a non-inlined monitor, we proved its soundness by showing that the output traces produced by two ℓ -equivalent executions will always be the same.

Magazinius et al. [20] present on-the-fly inlining of a dynamic information security monitor. We speculate that we could extend their ideas to allow on-the-fly instrumentation.

Assaf et al. [21] enforce termination-insensitive non-interference on a simple imperative language that support pointers using hybrid monitor. To deal with pointer-aliasing, they use a set of memory locations, which is similar to our sets of levels. One difference is that we use sets to distinguish outputs that never leak information from those that may leak information. Since we only insert runtime checks for the outputs that may leak information, our monitor introduces less overhead at runtime.

Askarov and Sabelfeld [22] use hybrid monitors to enforce information security in dynamic languages. In this setting, dynamic evaluation of programs (e.g., eval statements in JavaScript) requires on-the-fly static analysis of programs. They provide a model to define non-interference that is suitable to progress-sensitivity and they quantify information leaks due to termination [7].

Hritcu et al. [23] introduces an error-handling mechanism that allows all errors (even those caused by an information-flow control violation) to be safely recoverable. They support dynamic levels. To help prevent leaks through covert channels, they provide a discretionary access control mechanism called clearance that allows them to put an upper bound on the *pc*. Contrarily to our approach, the detection (and prevention) of leaks through progress channels is not done automatically.

Askarov et al. [24] introduce a hybrid monitoring framework capable of handling concurrent programs. They illustrate their approach on a simple imperative language similar to ours, but it does not support channel variables.

In their framework, each thread is guarded by its own local monitor (progress- and flow-sensitive). There is also a single global monitor that synchronizes the threads. Like us, they make use of an oracle to approximate the termination behaviour of branches. This oracle is called upon at runtime (making it a kind of on-the-fly static analysis), whereas ours is called only statically. The main difference between their approach and ours, excluding the concurrency of course, is the fact that our monitor is inlined whereas theirs is not.

8. Conclusion

We have presented a hybrid information flow enforcement mechanism, which detects and prevents leaks that may occur through the data-flow or the progress of a program. It uses information inferred during a phase of static analysis to instrument the program; this helps to reduce the number of false positives during the execution. The instrumented program uses level variables, a simple yet powerful way, to perform its dynamic analysis. This instrumented code can then be partially evaluated in order to reduce the amount of added commands.

Our main contributions are the following.

(a) We present an extended version of the hybrid monitor first presented in our previous article [6]. It is capable of enforcing flow- and progress-sensitive information security on general lattices. It is more precise and introduces less overhead than currently available solutions (e.g., [11, 13]) for two reasons: it makes use of a static termination oracle and does not approximate the level of a variable at the join of a conditional. Since our monitor is inlined, it can be easily optimized using classical partial evaluation techniques, [9].

(b) We prove the soundness of our inlined monitor and that the semantics of the original program is preserved, as long as it is secure.

(c) We show that, thanks to the use of sets of levels, it is possible to distinguish outputs that *never* leak information from outputs that *may* leak information.

(d) We present two ideas to increase the precision of the static analysis and the permissiveness of the dynamic analysis: propagating constraints on the set of possible security levels and using conditional updates. We chose not to use them in this work in order to increase readability.

Future work. Future work includes extensions to concurrency, declassification and information leakage due to timing. We would like to scale up the approach to deal with real world languages and to test it on elaborate programs. The use of abstract interpretation, [25], to enhance the static analysis is also to be considered in future work.

Acknowledgments. We would like to thank Fred B. Schneider and the reviewers for their valuable comments. Andrew Bedford, Josée Desharnais, and Nadia Tawbi are supported by grants from Laval University and NSERC.

Elisavet Kozyri is supported by AFOSR grants F9550-16-0250 and grants from Microsoft.

References

- [1] A. Sabelfeld and A. C. Myers, "Language-based information-flow security," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 1, 2003.
- [2] D. Volpano, C. Irvine, and G. Smith, "A sound type system for secure flow analysis," *Journal of computer security*, vol. 4, no. 2, pp. 167–187, 1996.
- [3] G. Le Guernic, A. Banerjee, T. Jensen, and D. A. Schmidt, "Automata-based confidentiality monitoring," *Asian Computing Science Conference*, 2006.
- [4] T. H. Austin and C. Flanagan, "Efficient purely-dynamic information flow analysis," in *Proceedings of the Workshop on Programming Languages and Analysis for Security*, 2009.
- [5] A. Russo and A. Sabelfeld, "Dynamic vs. static flow-sensitive security analysis," in *CSF*. IEEE Computer Society, 2010, pp. 186–199.
- [6] A. Bedford, S. Chong, J. Desharnais, and N. Tawbi, "A progress-sensitive flow-sensitive inlined information-flow control monitor," in *ICT Systems Security and Privacy Protection - 31st IFIP TC 11 International Conference, SEC 2016, Ghent, Belgium, May 30 - June 1, 2016, Proceedings*, 2016, pp. 352–366.
- [7] A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands, "Termination-insensitive noninterference leaks more than just a bit," in *Proceedings of the European Symp. on Research in Computer Security: Computer Security*, 2008.
- [8] L. Zheng and A. C. Myers, "Dynamic security labels and noninterference," in *Formal Aspects in Security and Trust*. Springer, 2005, pp. 27–40.
- [9] N. D. Jones, C. K. Gomard, and P. Sestoft, *Partial evaluation and automatic program generation*. Prentice Hall, 1993.
- [10] D. E. Denning, "A lattice model of secure information flow," *Communications of the ACM*, vol. 19, pp. 236–243, May 1976.
- [11] K. R. O'Neill, M. R. Clarkson, and S. Chong, "Information-flow security for interactive programs," in *CSFW*. IEEE, 2006.
- [12] G. Smith and D. Volpano, "Secure information flow in a multi-threaded imperative language," in *POPL*, 1998.
- [13] S. Moore, A. Askarov, and S. Chong, "Precise enforcement of progress-sensitive security," in *CCS 2012*, 2012.
- [14] B. Cook, A. Podelski, and A. Rybalchenko, "Proving program termination," *Commun. ACM*, vol. 54, no. 5, pp. 88–98, May 2011.
- [15] E. Kozyri, J. Desharnais, and N. Tawbi, "Block-safe information flow control," Department of Computer Science, Cornell University, Tech. Rep., Aug. 2016.
- [16] G. L. Guernic, "Automaton-based confidentiality monitoring of concurrent programs," in *CSF*, 2007.
- [17] A. Bedford, J. Desharnais, T. G. Godonou, and N. Tawbi, "Enforcing information flow by combining static and dynamic analysis," in *Proceedings of the International Symposium on Foundations & Practice of Security*, 2013.
- [18] A. Chudnov and D. A. Naumann, "Information flow monitor inlining," in *Proceedings of the 23rd IEEE Security Foundations Symposium*, 2010.
- [19] —, "Inlined information flow monitoring for javascript," in *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '15. New York, NY, USA: ACM, 2015, pp. 629–643. [Online]. Available: <http://doi.acm.org/10.1145/2810103.2813684>
- [20] J. Magazinius, A. Russo, and A. Sabelfeld, "On-the-fly inlining of dynamic security monitors," *Computers & Security*, vol. 31, no. 7, pp. 827–843, 2012.
- [21] M. Assaf, J. Signoles, F. Tronel, and É. Total, "Program transformation for non-interference verification on programs with pointers," in *IFIP International Information Security Conference*. Springer, 2013, pp. 231–244.
- [22] A. Askarov and A. Sabelfeld, "Tight enforcement of information-release policies for dynamic languages," in *CSF*, 2009.
- [23] C. Hritcu, M. Greenberg, B. Karel, B. C. Pierce, and G. Morrisett, "All your ifexception are belong to us," in *Security and Privacy (SP), 2013 IEEE Symposium on*. IEEE, 2013, pp. 3–17.
- [24] A. Askarov, S. Chong, and H. Mantel, "Hybrid monitors for concurrent noninterference," in *Computer Security Foundations Symposium*, 2015.
- [25] P. Cousot and R. Cousot, "Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints," in *POPL*, R. M. Graham, M. A. Harrison, and R. Sethi, Eds. ACM, 1977, pp. 238–252.

Appendix A. Examples

In order to simplify the examples, we assume that the variables `lVal`, `mVal`, `hVal` and channels `lChan`, `mChan`, `hChan` already exist, have arbitrary values and have the suggested security levels. We also assume that all other level variables have been initialized to L . Boolean values are represented as integer where 0 means false and any other integer means true.

Guarded send

The following example illustrates a situation where the guarded send is used to prevent a possible leak of information.

```

if mVal > 0 then
  c := mChan
else
  c := hChan
end;
send hVal to c; (*may leak information*)

```

Here is its instrumentation :

```

(*if*)
_oldpc1 := _pc;
if lVal > 0 then
  _pc := _pc  $\sqcup$  lValval  $\sqcup$  lValctx  $\sqcup$  L;

  (*assign*)
  (c, cval, cctx) := (lChan, lChanval, lChanctx  $\sqcup$  _pc);

  (c, cctx) := (c, cctx  $\sqcup$  _pc); (*update*)
else
  _pc := _pc  $\sqcup$  lValval  $\sqcup$  lValctx  $\sqcup$  L;

  (*assign*)
  (c, cval, cctx) := (hChan, hChanval, hChanctx  $\sqcup$  _pc);

  (c, cctx) := (c, cctx  $\sqcup$  _pc); (*update*)
end
_pc := _oldpc1;

(*guarded send*)
if _hc  $\sqcup$  _pc  $\sqcup$  hValval  $\sqcup$  hValctx  $\sqcup$  cctx  $\sqsubseteq$  cval then
  (send hVal to c)
else fail end;
_hc := _pc  $\sqcup$  _hc  $\sqcup$  hValctx  $\sqcup$  cctx;

```

Divergence

Commands after a loop that always diverges are ignored. Hence, the following program is statically safe.

```

while 1 do
  skip
end;
send hVal to lChan

```

Even if it is statically safe, it is instrumented :

```

(*while*)
_oldpc1 := _pc;
while 1 do
  _pc := _pc  $\sqcup$  L  $\sqcup$  L;
  (*skip*)
  skip;
end;
_pc := _pc  $\sqcup$  L  $\sqcup$  L;
_pc := _oldpc1

```

After partial evaluation, it results in the following program:

```

while 1 do
  skip
end

```

Notice that because the code is statically safe, the partial evaluation is able to get rid of the instructions added by the instrumentation algorithm. This is because, if the code is statically safe, then the conditions of the guarded send commands are all true. If guarded send commands are not needed, then level variables are also not needed.

Appendix B. Proofs

We prove that the type system of Figure 2 generates non-interferent programs (i.e., we prove its soundness).

Theorem 1 (Soundness of enforcement) *If a program p is well typed according to the type system of Figure 2, then the generated program $\llbracket p \rrbracket$ satisfies progress sensitive non-interference.*

The theorem is a consequence of the following results.

Lemma 3 *If two memories are ℓ -equivalent, then they agree on every expressions involving level variables that were generated by the instrumentation and whose value is $\sqsubseteq \ell$; these expressions include $_pc$, $_hc$, e_{val} , e_{ctx} , etc.*

Sketch. The proof is by induction. The idea is to show that all such expressions either only use level variables of the form x_{ctx} which are $\sqsubseteq \ell$, on which ℓ -equivalent memories agree by definition, or that whenever these expressions use some x_{val} , then they *protect* this potentially dangerous variable by x_{ctx} . The program context, $_pc$, is modified in the conditional and the loop, where the supremum is taken from the level of the condition. \square

Notation If $\Gamma(x) = \sigma_{\ell_x}$, we write $\Gamma_{ctx}(x)$ to mean ℓ'_x . If $\Gamma(x) = (int_{\ell_x} chan)_{\ell'_x}$ or $(int_{\ell_x})_{\ell'_x}$, then we write $\Gamma_{val}(x)$ to mean ℓ_x .

The following lemma states that security types of variables calculated by the typing system of Figure 2 include all possible runtime values.

Lemma 4 For any execution starting with a complete memory, if the runtime memory is m when an instrumented command $\llbracket cmd \rrbracket$ is executed, and m' is the memory after the execution, assuming the typing generated by Figure 2 is

$$\Gamma, pc, hc \vdash cmd : t, hc', \Gamma', \llbracket cmd \rrbracket,$$

we have

$$(1) m(_pc) \in pc \text{ and } m(_pc) = m'(_pc)$$

and the following invariants in the program execution

$$(2) m(_hc) \in hc$$

$$(3) m(x_{val}) \in \Gamma_{val}(x)$$

$$(4) m(x_{ctx}) \in \Gamma_{ctx}(x)$$

Proof. The proof is by structural induction.

Base Cases:

Initially, we have $pc = hc = \{L\}$, $m(_pc) = m(_hc) = L$, similarly for x_{val} and x_{ctx} , so the conditions are true at the initial state.

Case $\llbracket skip \rrbracket$ is trivial.

Case $\llbracket x := e \rrbracket = (x, x_{val}, x_{ctx}) := (e, e_{val}, e_{ctx} \sqcup _pc)$

Since the execution of this code does not modify $_pc$ or $_hc$, we have (1) and (2).

For every variable t_{val} in expression e_{val} , we have that $m(t_{val}) \in \Gamma_{val}(t)$. Similarly, for every variable t_{ctx} in expression e_{ctx} , we have that $m(t_{ctx}) \in \Gamma_{ctx}(t)$. Hence, by the definition of the supremum on sets (Definition 3) and (S-OP), we have (3) (i.e., $m'(x_{val}) \in \Gamma'_{val}(x)$). We also have (4) due to the updated environment returned (S-ASSIGN).

Case $\llbracket send\ x\ to\ c \rrbracket$

Subcase $\llbracket send\ x\ to\ c \rrbracket = send\ x\ to\ c$

Trivial since no variable type is modified.

Subcase $\llbracket send\ x\ to\ c \rrbracket =$

```

if  $\_pc \sqcup \_hc \sqcup x_{val} \sqcup x_{ctx} \sqcup c_{ctx} \sqsubseteq c_{val}$ 
then  $(send\ x\ to\ c)$ 
else fail
end;
 $\_hc := \_pc \sqcup \_hc \sqcup x_{ctx} \sqcup c_{ctx};$ 

```

Initially we have that $m(_pc) \in pc, m(_hc) \in hc, m(x_{ctx}) \in \Gamma_{ctx}(x)$ and $m(c_{ctx}) \in \Gamma_{ctx}(c)$. Only $_hc$ is modified by this command, hence we have (1), (3) and (4). We also have (2) by the supremum on sets and the hc' returned by (S-GSEND).

Induction Cases:

Case $\llbracket cmd_1; cmd_2 \rrbracket$

We can use induction on $\llbracket cmd_1 \rrbracket$, with the following hypothesis:

$$\bullet \Gamma, pc, hc \vdash cmd_1 : t_1, hc_1, \Gamma_1, \llbracket cmd_1 \rrbracket.$$

Induction hypothesis gives us that m'' (the memory after executing $\llbracket cmd_1 \rrbracket$) and Γ_1 satisfy (1)-(4). If $t_1 = D$ then $\llbracket cmd_1; cmd_2 \rrbracket$ is simply $\llbracket cmd_1 \rrbracket$, and we are done since all executions diverge. If $t_1 \neq D$, then we have for $\llbracket cmd_2 \rrbracket$ that:

$$\bullet \Gamma_1, pc, h_1 \vdash cmd_2 : t_2, h_2, \Gamma_2, \llbracket cmd_2 \rrbracket$$

Since the same program context as for cmd_1 is fed to the typing rule, then by the induction hypothesis and the fact that m'' and Γ_1 satisfy (1)-(4), we have the result for m' and Γ_2 .

Case $\llbracket if\ e\ then\ cmd_1\ else\ cmd_2\ end \rrbracket =$

```

 $\_oldpc' := \_pc;$ 
if  $e$  then
 $\_pc := \_pc \sqcup e_{val} \sqcup e_{ctx};$ 
 $\llbracket cmd_1 \rrbracket;$ 
update( $mv_1$ )
end;
update( $mv_2$ )
 $\_pc := \_oldpc$ 
else
 $\llbracket cmd_2 \rrbracket;$ 
update( $mv_2$ )
end;

```

For (1) we observe that $_pc$ has the same value before and after the instrumented code. Moreover, before cmd_i , $i = 1, 2$, the value of $_pc$ is in pc' , the context used to type these commands. For (3), by induction, we have that the invariants are true after the execution of $\llbracket cmd_1 \rrbracket$ or $\llbracket cmd_2 \rrbracket$. Hence we have (3). The **update** function in each branch inserts commands that, for every modified variable x , updates its x_{ctx} with $_pc$. This corresponds to what is done in the supremum of environments (Definition 4), so we have (4). Finally, since the **update** function updates the $_hc$ with $_pc$ when the termination behavior of the two branches can differ, and that the $_hc$ is updated by **update** if one of the branch contains a guarded send, we have (2) as this corresponds to the union done in (S-IF).

Case $\llbracket while\ e\ do\ cmd\ end \rrbracket$

```

 $\_oldpc' := \_pc;$ 
while  $e$  do
 $\_pc := \_pc \sqcup e_{val} \sqcup e_{ctx};$ 
 $\llbracket cmd \rrbracket;$ 
update( $mv$ );
 $\_pc := \_oldpc$ 
end;

```

The case is similar to the conditional. Variable $_pc$ is updated at the beginning of each iteration, it belongs to $pc \sqcup \ell_e \sqcup \ell'_e$, as wanted for the induction step involving $\llbracket cmd \rrbracket$, and we get (1), as pc' contains all the possible values of $_pc$ when $\llbracket cmd \rrbracket$ is executed. By induction, we have that the invariants are true after executing $\llbracket cmd \rrbracket$. Hence we have (3). We also have (2) and (4) using the same arguments as in the conditional. \square

The following proposition shows that any step of two executions performed from ℓ -equivalent memories results in ℓ -equivalent outputs. Theorem 1 follows as a corollary from it and from the next lemma.

Proposition 5 Let $m_i, i = 1, 2$ be two ℓ -equivalent complete memories, o_i be a trace and $\llbracket cmd \rrbracket$ be a command generated by Figure 2, that is

$$\Gamma, pc, hc \vdash cmd : t, h, \Gamma', \llbracket cmd \rrbracket.$$

Then if both $\llbracket cmd \rrbracket, m_i$ terminate, that is, we have maximal executions $\langle \llbracket cmd \rrbracket, m_i, o_i \rangle \longrightarrow^* \langle \mathbf{stop}, m'_i, o'_i \rangle$, then the following statements are invariants:

- (a) the memories are ℓ -equivalent (Def 5)
- (b) ℓ -projections of observations are equal

Proof. The proof is by structural induction.

Base Cases:

Initially, we have $m_i, i = 1, 2$ two ℓ -equivalent memories where the level variables have been initialized, so (a) and (b) are straightforward.

Case $\llbracket \mathbf{skip} \rrbracket$ is trivial.

Case $\llbracket x := e \rrbracket = (x, x_{val}, x_{ctx}) := (e, e_{val}, e_{ctx} \sqcup _pc)$

We have to prove (a), as (b) is trivial since this command does not modify the output traces. For ℓ -equivalence, since $_pc, _hc$ and the contents of channels are not modified, we obtain conditions (1) and (3) of Definition 5. For condition (2): since m_1 and m_2 are ℓ -equivalent, if $m'_1(x, x_{val}, x_{ctx}) \neq m'_2(x, x_{val}, x_{ctx})$, then either $m_i(x_{val}) \not\sqsubseteq \ell$ (in which case m'_1 and m'_2 are still ℓ -equivalent) or $m'_i(x_{ctx}) = m_i(_pc \sqcup e_{ctx}) \not\sqsubseteq \ell$ (by Lemma 3). Meaning that there is at least one variable t_{ctx} in e_{ctx} such that $m_i(t_{ctx}) \not\sqsubseteq \ell$, or that $_pc \not\sqsubseteq \ell$. Hence, m'_1 and m'_2 are still ℓ -equivalent. If $m'_1(x, x_{val}, x_{ctx}) = m'_2(x, x_{val}, x_{ctx})$, then m'_1 and m'_2 are ℓ -equivalent.

Case $\llbracket \mathbf{send } x \text{ to } c \rrbracket$

Subcase $\llbracket \mathbf{send } x \text{ to } c \rrbracket = \mathbf{send } x \text{ to } c$

By the typing rule, we have

$$pc \sqcup hc \sqcup \ell_x \sqcup \ell'_x \sqcup \ell'_c \sqsubseteq_s \ell_c$$

using the property of supremum, combining with Lemma 4 and converting the notation, we obtain

$$\begin{aligned} m_i(_pc) &\in pc \sqsubseteq_s \Gamma_{val}(c) \\ m_i(_hc) &\in hc \sqsubseteq_s \Gamma_{val}(c) \\ m_i(x_{val}) &\in \Gamma_{val}(x) \sqsubseteq_s \Gamma_{val}(c) \\ m_i(x_{ctx}) &\in \Gamma_{ctx}(x) \sqsubseteq_s \Gamma_{val}(c) \\ m_i(c_{ctx}) &\in \Gamma_{ctx}(c) \sqsubseteq_s \Gamma_{val}(c) \end{aligned}$$

By the definition of \sqsubseteq_s , this implies $m_i(_pc \sqcup _hc \sqcup x_{val} \sqcup x_{ctx} \sqcup c_{ctx}) \sqsubseteq m_i(c_{val})$, which means that the sending of x on c is safe. Hence we obtain (a) and (b) by ℓ -equivalence of the memories.

Subcase $\llbracket \mathbf{send } x \text{ to } c \rrbracket =$

```

if  $\_pc \sqcup \_hc \sqcup x_{val} \sqcup x_{ctx} \sqcup c_{ctx} \sqsubseteq c_{val}$ 
  then  $(\mathbf{send } x \text{ to } c)$ 
  else fail
end;
 $\_hc := \_pc \sqcup \_hc \sqcup x_{ctx} \sqcup c_{ctx}$ ;

```

We have to prove (a) and (b). For (a), we only have to prove conditions (1) and (3) as only the value of $_hc$ and the content of the channel may change.

If $m_i(_pc \sqcup _hc \sqcup x_{ctx} \sqcup c_{ctx}) \sqsubseteq \ell$ for one memory $i \in \{1, 2\}$ then, by ℓ -equivalence, it is also the case for the other memory and we get (1). It also implies that the two memory agree on x_{val} and c_{val} as they are ℓ -equivalent. This means that they either both fail, in which case there is nothing more to prove, or the **send** command is executed in both. If it is executed, then $m_i(_pc \sqcup _hc \sqcup x_{val} \sqcup x_{ctx} \sqcup c_{ctx}) \sqsubseteq m_i(c_{val})$, for $i = 1, 2$, and by ℓ -equivalence again, $o_i \upharpoonright \ell = o'_i \upharpoonright \ell$ in both executions and (b) is true, as well as (3).

If $m'_i(_hc) = m_i(_pc \sqcup _hc \sqcup x_{ctx} \sqcup c_{ctx}) \not\sqsubseteq \ell$ for $i = 1, 2$, then we get (1) and we have that the **send** command could succeed in m_1 and fail in m_2 . But since $m_i(_pc \sqcup _hc \sqcup x_{ctx} \sqcup c_{ctx}) \not\sqsubseteq \ell$, we know that it cannot succeed on a channel whose level is $\sqsubseteq \ell$ and we get (b) and (3).

Induction Cases:

Case $\llbracket cmd_1; cmd_2 \rrbracket$

The hypothesis gives that m_1 and m_2 satisfy (a) and (b). Then we can use induction on $\llbracket cmd_1 \rrbracket$, with the following hypotheses:

- $\Gamma, pc, hc \vdash cmd_1 : t_1, h_1, \Gamma_1, \llbracket cmd_1 \rrbracket$.
- $\langle \llbracket cmd_1 \rrbracket, m_i, o_i \rangle \longrightarrow^* \langle \mathbf{stop}, m''_i, o''_i \rangle$

Induction hypothesis gives us that m''_1 and m''_2 satisfy (a) and (b). We have for $\llbracket cmd_2 \rrbracket$ that:

- $\Gamma_1, pc, h_1 \vdash cmd_2 : t_2, h_2, \Gamma_2, \llbracket cmd_2 \rrbracket$
- $\langle \llbracket cmd_2 \rrbracket, m''_i, o''_i \rangle \longrightarrow^* \langle \mathbf{stop}, m'_i, o'_i \rangle$

By the induction hypothesis and the fact that m''_1 and m''_2 satisfy (a) and (b), we obtain that m'_1 and m'_2 satisfy (a) and (b).

Case $\llbracket \mathbf{if } e \text{ then } cmd_1 \text{ else } cmd_2 \text{ end} \rrbracket =$

```

 $\_oldpc' := \_pc$ ;       $\_pc := \_pc \sqcup e_{val} \sqcup e_{ctx}$ ;
if  $e$  then           $\llbracket cmd_2 \rrbracket$ ;
 $\_pc := \_pc \sqcup e_{val} \sqcup e_{ctx}$ ;  $\mathbf{update}(mv_1)$ 
 $\llbracket cmd_1 \rrbracket$ ;          end;
 $\mathbf{update}(mv_2)$        $\_pc := \_oldpc'$ 
else

```

Assume that m_1 and m_2 satisfy (a) and that $o_1 \upharpoonright \ell = o_2 \upharpoonright \ell$. If $m_i(e_{val} \sqcup e_{ctx}) = \sqcup_{t \in Var(e)} m_i(t_{val} \sqcup t_{ctx}) \sqsubseteq \ell$, $i = 1, 2$ (symmetry is given by ℓ -equivalence), we have the result by induction since both memories take the same branch, say i .

Now assume that $m_i(e_{\text{val}} \sqcup e_{\text{ctx}}) \not\sqsubseteq \ell$, $i = 1, 2$. If $m_1(e) = m_2(e)$ we are in the same situation as above. So assume w.l.o.g. that $m_1(e)$ is true but $m_2(e)$ is false; hence under memory m_i , code $\llbracket cmd_i \rrbracket$ will be executed.

If both commands terminate, that is:

- $\langle \llbracket cmd_1 \rrbracket, m_1, o_1 \rangle \longrightarrow^* \langle \text{stop}, m'_1, o'_1 \rangle$
- $\langle \llbracket cmd_2 \rrbracket, m_2, o_2 \rangle \longrightarrow^* \langle \text{stop}, m'_2, o'_2 \rangle$

Then we know, by Lemma 6(c), that

$$o'_1 \upharpoonright \ell = o_1 \upharpoonright \ell = o_2 \upharpoonright \ell = o'_2 \upharpoonright \ell.$$

because before executing cmd_i , variable $_pc$ is updated with the value of $m_i(e_{\text{val}} \sqcup e_{\text{ctx}})$ which is $\not\sqsubseteq \ell$. Hence, the execution of $\llbracket cmd_i \rrbracket$ will produce no output on channels of level $\sqsubseteq \ell$. Hence, (b) is proven for the induction step.

This also proves (a) on channels, but we need to prove it on variables (level and non level).

Let $x \in \mathcal{V}$ and $mv = mv_1 \cup mv_2$, where mv_i the set of variables that may be modified in cmd_i , $i = 1, 2$. By definition, variable $_hc$ is also included in this set if the termination of the branches may differ or if one of them contains a guarded send. If $x \notin mv$ and the m_i 's agree on x , x_{ctx} and x_{val} , then the m_i 's also agree on x , x_{ctx} and x_{val} .

If $x \in mv_i$, we have that $x_{\text{ctx}} \sqsupseteq _pc \sqcup e_{\text{val}} \sqcup e_{\text{ctx}}$ due to the commands inserted by the update function. Thus we have (a).

Case $\llbracket \text{while } e \text{ do } cmd \text{ end} \rrbracket =$

```

_oldpc' := _pc;           end;
while e do                 _pc := _pc \sqcup e_{val} \sqcup e_{ctx};
_oldpc := _pc \sqcup e_{val} \sqcup e_{ctx}; update(mv);
\llbracket cmd \rrbracket;              _pc := _oldpc

```

As a loop is essentially a (possibly infinite) sequence of **if**, the case is similar to the conditional. Assume that m_1 and m_2 satisfy (a) and (b). If $m_i(e_{\text{val}} \sqcup e_{\text{ctx}}) = \sqcup_{t \in \text{Var}(e)} m_i(t_{\text{val}} \sqcup t_{\text{ctx}}) \not\sqsubseteq \ell$, then we know that the execution of $\llbracket cmd \rrbracket$ will produce no output on channels of level $\sqsubseteq \ell$ since variable $_pc$ is updated before entering the loop, and that every variable that is or could have been modified by $\llbracket cmd \rrbracket$ will have a t_{ctx} that is $\not\sqsubseteq \ell$ due to the update function. Hence, we have (a) and (b).

If instead we have that $m_i(e_{\text{val}} \sqcup e_{\text{ctx}}) \sqsubseteq \ell$, then we have the result by induction since both memories will always take the same branch. \square

Lemma 6 *With the premises of the previous theorem, if one step of $\llbracket cmd \rrbracket$ fails or diverges for one memory then it also fails or diverges for the other or no more output on channels of level ℓ or lower will be performed on that execution. More precisely, for $i = \{1, 2\}$*

- (a) $m_i(_hc) \sqsubseteq m'_i(_hc)$

(b) *if only one execution fails or diverges, say m_1 , then $o_2 \upharpoonright \ell = o'_2 \upharpoonright \ell$ and $m'_2(_hc) \not\sqsubseteq \ell$.*

(c) *if $m_i(_pc) \not\sqsubseteq \ell$ or $m_i(_hc) \not\sqsubseteq \ell$ then $o_i \upharpoonright \ell = o'_i \upharpoonright \ell$*

(d) *if both executions fail or both diverge, then the ℓ -projections of observations are equal*

Proof. (a) is straightforward since $_hc$ is always included in the right-hand side when updating $_hc$.

(b) is proven by induction, following the lines of Proposition 5. The interesting cases are the guarded send sub-case and the inductive cases.

Case $\llbracket \text{send } x \text{ to } c \rrbracket$

Let's assume that the send command is transformed into a guarded send. The only case where only one execution fails, say m_1 , is one where $m'_2(_hc) \not\sqsubseteq \ell$, and where no observation is made on a channel of level ℓ or lower, as wanted.

Case $\llbracket cmd_1; cmd_2 \rrbracket$

Let m''_i and o''_i be the memories and output traces after executing cmd_1 . Let m'_i and o'_i be the memories and output traces after executing $cmd_1; cmd_2$. If one execution fails or diverges, say m_1 on cmd_1 , then by induction, $m''_2(_hc) \not\sqsubseteq \ell$, and by (c), we obtain $o_2 \upharpoonright \ell = o''_2 \upharpoonright \ell$. By (a) and (c) and induction, we then get $o_2 \upharpoonright \ell = o'_2 \upharpoonright \ell$, as wanted. By (a), and (c) again, we also obtain $m'_2(_hc) \not\sqsubseteq \ell$. If m_1 fails or diverges on cmd_2 instead, the argument is similar.

Case $\llbracket \text{if } e \text{ then } cmd_1 \text{ else } cmd_2 \text{ end} \rrbracket$

There are two situations in which only one execution, say m_1 , can fail: (1) only one execution executes a guarded send or (2) they both execute the same guarded send but the result is different. In case (1), we have that $m_i(_pc) \not\sqsubseteq \ell$ (a consequence of Proposition 5 and Lemma 3), and that at least one of the **if**'s branch contains a guarded send, which is always followed by an update to variable $_hc$. Since one of the branch modifies variable $_hc$, it also means that the update function updates $_hc$ to $_pc$, which is $\not\sqsubseteq \ell$. Hence, we have (b) in this case. In case (2), the only way that the result of a guarded send can be different is if $m_i(_pc \sqcup _hc \sqcup x_{\text{ctx}} \sqcup c_{\text{ctx}}) \not\sqsubseteq \ell$, where x is the variable sent and c the channel on which the send occurs. In this case, the update to $_hc$ that immediately follows the guarded send will ensure that $m_2(_hc) \not\sqsubseteq \ell$. Hence we also have (b) in this case.

Similarly, for only one execution to diverge, two things must be true: (1) $m_i(_pc) \not\sqsubseteq \ell$, and (2) the termination type of the **if** command is M . Since the termination type is M , we have that the update function inserted after the conditional updates the $_hc$ to $_pc$, which is $\not\sqsubseteq \ell$. Hence, we also have (b) in this case.

Case `[[while e do cmd end]]`

The argument is similar to the `if` command.

For (c), there are two cases where there are observations, the send case and the guarded send case. For the latter, the condition is taken care of by the guard. For the former, if $m_i(_pc) \not\sqsubseteq \ell$ or $m_i(_hc) \not\sqsubseteq \ell$ then, by Lemma 4 (1) and (2), one of the sets pc or hc contains a security level $\ell' \not\sqsubseteq \ell$. By the typing rule, this implies that $\ell' \sqsubseteq_s \Gamma_{\text{val}}(c)$. By definition of \sqsubseteq_s , all elements of $\Gamma_{\text{val}}(c)$ are greater than or equal to ℓ' and hence, again by Lemma 4 (3), $\ell' \sqsubseteq_s m_i(c_{\text{val}})$ and $\ell \not\sqsubseteq_s m_i(c_{\text{val}})$, thus $o_i \upharpoonright \ell = o'_i \upharpoonright \ell$.

Finally, for (d), we have three cases: both executions fail, both executions diverge and $m_i(_pc \sqcup e_{\text{val}} \sqcup e_{\text{ctx}}) \sqsubseteq \ell$ or both executions diverge and $m_i(_pc \sqcup e_{\text{val}} \sqcup e_{\text{ctx}}) \not\sqsubseteq \ell$, where e is the guard expression of the loop that diverges. If both executions fail, then there will be no more outputs and so the ℓ -projections remain equivalent. If both executions diverge and $m_i(_pc \sqcup e_{\text{val}} \sqcup e_{\text{ctx}}) \sqsubseteq \ell$, then by Lemma 3 we have that the ℓ -projections of observations are equal. If both executions diverge and $m_i(_pc \sqcup e_{\text{val}} \sqcup e_{\text{ctx}}) \not\sqsubseteq \ell$, then the update to the $_pc$ inside the body of the loop ensures that $m_i(_pc) \not\sqsubseteq \ell$ before executing cmd , and so that there will be no more outputs on channels of level lower or equal to ℓ . \square

Lemma 7 (Fixed-point) *The fixed-point computed by typing rules (S-LOOP1) and (S-LOOP2) always converges to a value.*

Proof. Let Γ, Γ' be typing environments and pc, pc' be sets of levels. We say that $(\Gamma, pc) \sqsubseteq (\Gamma', pc')$ if:

$$\begin{aligned} \text{dom}(\Gamma) &\subseteq \text{dom}(\Gamma') \wedge \\ pc &\subseteq pc' \wedge \\ (\forall x \in \text{dom}(\Gamma)) & \\ &\Gamma_{\text{val}}(x) \subseteq \Gamma'_{\text{val}}(x) \wedge \Gamma_{\text{ctx}}(x) \subseteq \Gamma'_{\text{ctx}}(x) \end{aligned}$$

Since this relation is a partial order (i.e., reflexive, transitive, antisymmetric), to prove that the computation of the fixpoint in (S-LOOP2) always converges, we only need to prove that applying the rule iteratively builds a chain of environments that satisfy the relation. More precisely, we prove that the sequence $(\Gamma_0 \sqcup_{pc'_0} \Gamma'_0, pc'_0) \sqsubseteq (\Gamma_1 \sqcup_{pc'_1} \Gamma'_1, pc'_1) \sqsubseteq \dots$ converges, where Γ_0 is the initial environment, Γ'_0 is the environment obtained after the first analysis of the loop body cmd , pc'_0 is the context in which the first analysis of cmd is performed. The result of analyzing cmd , the loop body, under Γ_i is Γ'_i . Note that $\Gamma_i = (\Gamma_{i-1} \sqcup_{pc'_{i-1}} \Gamma'_{i-1})$, is the environment under which the loop is recursively analyzed at step i .

By Definition 4, we have

$$\text{dom}(\Gamma_{i-1} \sqcup_{pc'_{i-1}} \Gamma'_{i-1}) \subseteq \text{dom}(\Gamma_i \sqcup_{pc'_i} \Gamma'_i).$$

Since $pc'_i = pc'_{i-1} \cup (pc'_{i-1} \sqcup \ell_{e_{i-1}} \sqcup \ell'_{e_{i-1}})$, where $\ell_{e_{i-1}}$ and $\ell'_{e_{i-1}}$ are the levels associated to the conditional expression of the loop at step $i-1$, we have $pc'_{i-1} \subseteq pc'_i$. Let $\Gamma = (\Gamma_{i-1} \sqcup_{pc'_{i-1}} \Gamma'_{i-1})$ and $\Gamma' = (\Gamma_i \sqcup_{pc'_i} \Gamma'_i)$. Let $x \in \text{dom}(\Gamma)$.

Case x is not modified in the body of the loop :

We have that $\Gamma_{\text{val}}(x) = \Gamma'_{\text{val}}(x)$ and $\Gamma_{\text{ctx}}(x) = \Gamma'_{\text{ctx}}(x)$ by Definition 4.

Case x is modified in the body of the loop:

We have that $\Gamma_{\text{val}}(x) \subseteq \Gamma'_{\text{val}}(x)$ by Definition 4. Since Γ is calculated using pc'_{i-1} , we have that $\Gamma_{\text{ctx}}(x) \sqcup pc'_{i-1} = \Gamma'_{\text{ctx}}(x)$ and since Γ' is calculated using pc'_i , we have that $\Gamma'_{\text{ctx}}(x) \sqcup pc'_i = \Gamma'_{\text{ctx}}(x)$. Hence, we have that $\Gamma_{\text{ctx}}(x) \subseteq \Gamma'_{\text{ctx}}(x)$ by $pc'_{i-1} \subseteq pc'_i$ and by Definition 4. Hence, a fixpoint will always be reached since \mathcal{L} is finite. \square

Theorem 2 (Semantics preservation.) *Let p be a program, m a memory, and o an output trace. Then*

$$\begin{aligned} \langle \llbracket p \rrbracket, m, \epsilon \rangle \downarrow o &\Rightarrow \langle p, \hat{m}, \epsilon \rangle \downarrow o \\ \langle p, \hat{m}, \epsilon \rangle \downarrow o \wedge \langle \llbracket p \rrbracket, m, \epsilon \rangle \downarrow o' &\Rightarrow o \preceq o' \vee o' \preceq o \end{aligned}$$

where $o' \preceq o$ means o' is a prefix of o .

Note that \hat{m} is a memory for the source language, and hence it is of type $\mathcal{V} \uplus \mathcal{C} \rightarrow \mathbb{Z} \uplus \mathcal{C}$ whereas m can, in addition, map variables to levels.

Sketch. By structural induction. The program generated by our instrumentation contains the same commands as the original program, in the same order. The only difference being the additional assignments on level variables and checks. For this reason, the only non-trivial case is the send command, since it modifies the output trace, or halts the program. Hence assume that $cmd = \mathbf{send } x \mathbf{ to } c$. There are two cases: (1) $\llbracket cmd \rrbracket = cmd$ and (2) $\llbracket cmd \rrbracket = \mathbf{if } _pc \sqcup _hc \sqcup x_{\text{val}} \sqcup x_{\text{ctx}} \sqcup c_{\text{ctx}} \sqsubseteq_s c_{\text{val}} \mathbf{ then (send } x_1 \mathbf{ to } x_2) \mathbf{ else fail end; } _hc := _pc \sqcup _hc \sqcup x_{\text{ctx}} \sqcup c_{\text{ctx}}$.

In the first case, the claim is trivial. In the second case, the send is guarded by a condition. If this condition is true, the sending will happen and the output trace will be updated with $o:(m(x_1), m(x_2))$, as would have been done by cmd . Otherwise, the program $\llbracket cmd \rrbracket$ will be stopped, and hence, no more output will happen, although cmd could produce other outputs. \square