

# Monitoring Reactive Systems with Dynamic Channels

Dante Zanarini Mauro Jaskelioff

CIFASIS - CONICET

Universidad Nacional de Rosario

{zanarini, jaskelioff}@cifasis-conicet.gov.ar

## Abstract

Given the increasingly sensitive data that web applications deal with, a lot of attention has been put into their security. Dynamic methods for ensuring confidentiality of secret data, such as monitors, are usually preferred due to their permissiveness and ability to adapt to dynamic features of web languages. One dynamic approach to confidentiality is through secure multi-execution, a technique which transforms programs into secure ones. A recent refinement of this technique led to a monitor for reactive systems such as web applications which is precise, in the sense that it raises an alarm exactly when a security condition is violated, and transparent, in the sense that the semantics of secure programs is preserved. A limitation of this and other approaches based on secure multi-execution is that there is a fixed set of channels with a fixed security level. However, most web applications create channels dynamically, even by doing something as trivial as adding a button to a page. Moreover, the security level of such new channel would be chosen dynamically. In this work, we overcome the limitation of assuming a fixed set of channels and introduce a model of reactive systems with dynamic channels and present a precise and transparent monitor for it.

**Categories and Subject Descriptors** D.3.3 [Programming Languages]: Language Constructs and Features; D.4.6 [Operating Systems]: Security and Protection—Information flow controls

**General Terms** Languages, Security

**Keywords** Information Flow, Reactive Systems, Dynamic Channels

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

PLAS 2014, July 29th, 2014, Uppsala, Sweden.  
Copyright © 2014 ACM 978-1-xxxx-xxxx-n/yy/mm...\$15.00.  
<http://dx.doi.org/10.1145/nnnnnnn.nnnnnn>

## 1. Introduction

Web applications are pervasive and deal with data of increasing significance. Therefore, many researchers have been working towards the goal of securing them. In particular, there have been many efforts to apply information flow techniques to web applications as a general mechanism to enforce confidentiality [7, 18, 20, 24, 35]. Given that several of the OWASP top-ten web vulnerabilities [37] can be rephrased in terms of information-flow problems, a general information-flow technique would be a substantial improvement over ad-hoc and purpose-specific approaches such as architectures to contain advertisement scripts [23], and browser extensions to control cache-based leaks [19], to name just a few.

However, it is difficult to obtain a general solution for controlling information flow in web applications. There are millions of deployed web pages. A good solution should raise a minimum of false alarms when applied to them, and hence it must be permissive and allow as many secure web pages as possible. Moreover, the advanced language features of web applications make them difficult to analyse. Therefore, researchers tend to work with simplified models of real web applications in order for the problem to be tractable.

Static Denning-style [14, 30, 36] enforcements of information flow, can reject secure programs, as they over-approximate the problem. Moreover, dynamic features of web applications can be difficult to work with in this style. Hence, in order to increase permissiveness and be able to cope with dynamic features of web applications, many approaches tend to be dynamic, usually in the form an execution monitor [1, 3, 4, 31], despite their inability to enforce sound and precise standard notions of information-flow policies, such as non-interference [25, 32].

Recently, Zanarini, Jaskelioff, and Russo [39] introduced a monitor for general reactive systems (such as web applications) which can precisely detect when a program leaks information, where the information leak is defined with respect to an input without secrets. The monitor is based on the idea of secure multi-execution (SME) [15], a technique which by executing the same program several times, once for each security level, transforms an arbitrary program into a secure

one. By improving SME so that the order of events is preserved across security levels, one can compare the original program with the transformed one and see if they differ, and construct in this way a security monitor which is both precise (only raises alarms when there is an information leak) and transparent (it does not affect the execution of secure programs.)

In reactive systems one has input channels on which events arrive, and output channels on which events are outputted. For example, in a web application, every element of the DOM tree has events associated to it (such as `onmouseover`), and therefore we can think that every element has an input channel associated to it. Having all events arrive through a single channel is not an option since we expect different nodes to have different security levels. An output event can be, for instance, a modification of an attribute of an existing node of the DOM, or the addition of a new one.

One limitation of the approach in [39] and of other approaches based on SME [5, 6, 15, 21, 28] is that there is a fixed set of channels, each with a fixed security level. However, web applications may modify the DOM during execution [8], and therefore the set of possible input channels may be modified during execution. That is, the set of channels is *dynamic*. For example, when adding a button to a web page, a new set of events will be possible (such as an `onClick` event). Such events must come through a new channel, which will have an application and context-specific security level. That is, depending on the context, the events from a newly created channel might be secret or public, and what’s more, the security level of an existing channel might change during execution. Therefore, in order to model more realistic reactive systems, a way of dealing with dynamic channels is needed. Note that, one could improvise a way of dealing with dynamic channels by assuming a countably infinite set of channels and assigning a “fresh” channel name to a new channel, but this approach has clear disadvantages: first and foremost, the security level for each channel is fixed and cannot be changed dynamically; second, the interpretation of a source language into a model with a fixed set of channels is more complex, and hence more error-prone.

Adding dynamic channels to a model has the disadvantage of making it more complex. The very question of what is a secure program in this context becomes non-trivial, as the security level of a channel might change during execution. However, if we are to tackle realistic web applications, it is a complexity that needs to be dealt with.

In this article, we extend the monitor for reactive systems introduced in [39] to account for dynamic channels. In order to do this, we need to extend the model of reactive systems and adapt the notion of non-interference. Notably, our extended monitor offers the same security and transparency guarantees as the one for static channels.

In Section 2, we introduce a model of reactive systems with dynamic channels, and show how a JavaScript-like language can be interpreted into such a model. In Section 3 we introduce notions of security for programs and for runs. In Section 4, we introduce our monitor for reactive systems with dynamic channels. In Section 5, we state the precision and transparency guarantees of the monitor. In Section 6 we discuss related work. Finally, in Section 7, we conclude. Proofs for all the results can be found in an online extended version [38].

## 2. Modelling Reactive Systems

We want to provide an abstract model for reactive systems. The model should be as abstract as possible in order to simplify away irrelevant features, but powerful enough so as to model all the pertinent features. The main example of a reactive system we have in mind is web applications. Therefore we will use this example to drive intuitions.

A reactive system consists of input events and of reactions to those events, which may be output events or simply an invisible change of the internal state of the system. Each input or output event occurs in a given channel. When thinking about web applications, a typical input event is a mouse click on an element of a web page, and a typical output event is putting some text on the screen. More precisely, a web page is represented by a DOM tree, input events are associated to elements of the DOM, and output events modify the DOM. An obvious way of adding security level information to a DOM, is adding an attribute to indicate the security level of the element, which may be modified during execution.

We model a reactive program as an interaction structure [17]. Our model is based on the model in [39], but extended to account for the dynamic opening and closure of channels. That is, we model reactive systems using an interaction structure indexed by a mapping of open channels to their corresponding security levels.

Throughout this article, we organise security levels in a lattice  $(\mathcal{L}, \sqsubseteq)$ , with the intention to express that data at level  $\ell_1$  can securely flow into data at level  $\ell_2$  when  $\ell_1 \sqsubseteq \ell_2$ . Let  $\text{Chan}$  be the set of channel names, and let  $\text{CS} = \text{Chan} \rightarrow \mathcal{L}$  be the type of channel systems, i.e. a channel system  $cs : \text{CS}$  is a partial function from channels to security levels. If  $c \in \text{dom}(cs)$ , we say that  $c$  is open in  $cs$ .

An event is a piece of data paired with the name of the communication channel associated to it. Given a channel system  $cs$ , the domain of  $cs$  indicates which channels are open, and therefore which events are compatible with  $cs$ . Take  $A$  to be a set; the set of  $A$ -events compatible with  $cs$  is:

$$E_A cs = \{(c, a) \in \text{Chan} \times A \mid c \in \text{dom}(cs)\}.$$

That is,  $E_A cs$  denotes the set of events on open channels of the channel system  $cs$ . Equivalently, the set of compatible events can be defined as

$$E_A cs = \{(c, a, \ell) \in \text{Chan} \times A \times \mathcal{L} \mid cs(c) = \ell\}.$$

It is not difficult to see that the two definitions are equivalent. However, as discussed later in the article, the latter is more convenient for analysing confidentiality in a dynamic-channel context. We denote with  $I$  the set of input values, and with  $O$  the set of output values.

**Definition 2.1** (Reactive Behaviour). A reactive behaviour  $\text{React} : \text{CS} \rightarrow \text{Set}$  is the interaction structure given by the following coinductive datatype indexed by a channel system.

$$\begin{aligned} \text{React} : \text{CS} \rightarrow \text{Set} = \\ \text{Read} : \forall cs. (E_I cs \rightarrow \text{React } cs) \rightarrow \text{React } cs \\ \text{Write} : \forall cs. E_O cs \rightarrow \text{React } cs \rightarrow \text{React } cs \\ \text{Step} : \forall cs. \text{React } cs \rightarrow \text{React } cs \\ \text{Stop} : \forall cs. \text{React } cs \\ \text{New} : \forall cs. ((c, \ell) : \text{Chan} \times \mathcal{L}) \rightarrow \text{React } (cs \oplus (c, \ell)) \\ \quad \rightarrow \text{React } cs \\ \text{Close} : \forall cs. (c : \text{Chan}) \rightarrow \text{React } cs \\ \quad \rightarrow \forall \ell. \text{React } (cs \oplus (c, \ell)) \end{aligned}$$

The idea is that  $\text{React } cs$  describes the execution behaviour of a reactive system whose open channels are described by a channel system  $cs$ . Because the behaviour datatype is defined coinductively, it is perfectly suited to model the behaviour of programs that run indefinitely.

Intuitively, constructor  $\text{Read}$  denotes a behaviour  $\text{React } cs$  that receives an input from an open channel determined by the environment ( $E_I cs$ ) and, based on that, decides how to continue ( $\text{React } cs$ ). Constructor  $\text{Write}$  represents a behaviour  $\text{React } cs$  which writes an output in a chosen channel ( $E_O cs$ ) and continues with another computation ( $\text{React } cs$ ). Constructor  $\text{Step}$  corresponds to a silent step, that is, a computation which does not affect the environment. Silent steps allows us to model divergence. Constructor  $\text{Stop}$  models termination. Reactive systems usually run forever, but we use this constructor to model abnormal termination as triggered by an unrecoverable error condition. Constructor  $\text{New}$  denotes a behaviour  $\text{React } cs$  which opens a new channel  $c$  of security level  $\ell$ , and continues with a behaviour on a channel system which extends  $cs$  with  $(c, \ell)$ . Note that in a behaviour  $r' = \text{New } (c, \ell) r : \text{React } cs$  the new channel  $c$  is open only in the continuation  $r$  and not on  $r'$ . Finally, constructor  $\text{Close}$  represents a behaviour on a channel system with an open channel  $c$  and continues with a behaviour with that channel closed. That is, the behaviour  $r' = (\text{Close } c r)$  has channel  $c$  open, but its continuation  $r$  must have it closed.

The operation  $\oplus$  adds a channel to a channel system. More precisely, given a channel  $c$  and a channel system  $cs$  such that  $c \notin \text{dom}(cs)$ ,  $cs \oplus (c, \ell) = cs \cup \{(c, \ell)\}$ . As a consequence of this definition, in a behaviour  $r' = \text{New } (c, \ell) r : \text{React } cs$ , we know that  $c$  is not open in  $cs$ , and  $r'$  can interact with  $c$  only in the continuation  $r$ . A similar observation applies to constructor  $\text{Close}$ . If  $r' = \text{Close } c r : \text{React } (cs \oplus (c, \ell))$ , then  $c$  is closed in  $cs$ , and the continuation  $r$  cannot interact with  $c$ .

By modelling reactive systems with an interaction structure we obtain a black-box model of the system. The constructors represent only the features which are relevant to our analysis of security and therefore we can abstract away all the rest of the complex features of the system. In this manner, we gain modularity since new programming features which are not related to I/O do not affect our formal results. Indexing by the channel system means that behaviours cannot do nonsense such as writing to a closed channel or opening an open channel. Therefore, working with behaviours is simplified since we do not need to consider absurd cases. Of course, a JavaScript program might be erroneous and try to perform an illegal operation, but this can be handled as an error, perhaps behaving as a  $\text{Stop}$ .

## 2.1 Semantics

A reactive behaviour is essentially a static description of the possible interactions that might occur during execution. In order to know what happens in a given run, we need to give reactive behaviours some semantics. We start by defining possibly infinite sequences.

**Definition 2.2** (Colists). Let  $A$  be a set. Consider the type of possibly infinite sequences of  $A$  to be coinductively defined as follows.

$$\text{Colist}_A = [] \mid A :: \text{Colist}_A$$

We write  $[a, b, c]$  to denote a finite colist  $a :: b :: c :: []$ .

The structure of input and output events of the system is given by colists. The set of all events with data from a set  $A$  is defined as  $\mathcal{E}_A = \text{Chan} \times A \times \mathcal{L}$ . It is easy to see that for any channel system  $cs$ , the set  $E_A cs$  of compatible events is included in  $\mathcal{E}_A$ .

**Definition 2.3** (Colists of input and output events). We define the set of colists of inputs  $\mathcal{I}$  and the set of colists of outputs  $\mathcal{O}$  for reactive systems as follows.

$$\mathcal{I} = \text{Colist}_{\mathcal{E}_I} \quad \mathcal{O} = \text{Colist}_{\mathcal{E}_O \cup \{\bullet, \circ\}}$$

The elements of an output colist can be events in  $\mathcal{E}_O$ , an invisible output ( $\bullet$ ), or an abnormal termination event ( $\circ$ ).

The rules in Figure 1 coinductively define an evaluation relation for reactive behaviours. This is really a *family* of relations indexed by a channel system. More precisely, we define a family of relations  $\Rightarrow_{cs} \subseteq (\text{React } cs \times \mathcal{I} \times \mathcal{O})$ , and write  $(t, i) \Rightarrow_{cs} o$  for  $(t, i, o) \in \Rightarrow_{cs}$ . Intuitively, feeding a colist of inputs  $i$  to a reactive behaviour  $t$  yields the output colist  $o \in \mathcal{O}$  iff  $(t, i) \Rightarrow_{cs} o$ .

Rule  $R_1$  produces no outputs ( $[]$ ) when no input events are present. Rule  $R_2$  consumes the first available input event ( $e$ ) if the event  $e$  is compatible with  $cs$ , and based on that, produces the output  $o$  ( $(f e, i) \Rightarrow_{cs} o$ ). Rule  $W$  outputs an event  $e$  ( $e :: o$ ), followed by the outputs triggered by program  $t$  ( $(t, i) \Rightarrow_{cs} o$ ). Rule  $N$  continues execution in the channel system obtained after opening a new channel ( $cs \oplus (c, \ell)$ ).

$$\begin{array}{c}
\text{R}_1 \frac{}{(\text{Read } f, \square) \Rightarrow_{cs} \square} \\
\text{R}_2 \frac{e \in E_I cs \quad (f \ e, i) \Rightarrow_{cs} o}{(\text{Read } f, e :: i) \Rightarrow_{cs} \bullet :: o} \\
\text{W} \frac{(t, i) \Rightarrow_{cs} o}{(\text{Write } e \ t, i) \Rightarrow_{cs} e :: o} \\
\text{N} \frac{(t, i) \Rightarrow_{cs \oplus (c, \ell)} o}{(\text{New } (c, \ell) \ t, i) \Rightarrow_{cs} \bullet :: o} \\
\text{C} \frac{cs = cs' \oplus (c, \ell) \quad (t, i) \Rightarrow_{cs'} o}{(\text{Close } c \ t, i) \Rightarrow_{cs} \bullet :: o} \\
\text{S} \frac{(t, i) \Rightarrow_{cs} o}{(\text{Step } t, i) \Rightarrow_{cs} \bullet :: o} \quad \text{E} \frac{}{(\text{Stop}, i) \Rightarrow_{cs} [o]}
\end{array}$$

**Figure 1.** Evaluation relation for reactive behaviours

Rule C continues execution in a channel system obtained after closing channel  $c$  (that is,  $cs'$ ). Rule S simply outputs  $\bullet$  when a silent computation step is performed. Finally, rule E signals an abnormal termination event and ends. Here we see the role of the termination event  $o$ : to indicate a termination different from  $\square$ .

It might seem strange that input events carry a security level. It would be more natural to consider input events as pairs  $(c, d)$  of a channel and a datum, and assign the security level of the event according to the mapping of  $c$  in the channel system at the time the event arrives. As a matter of fact, this is exactly the way we expect to *implement* the semantics of behaviours. Nevertheless, this “natural” semantics would be equivalent to the one in Fig. 1 as the predicate  $e \in E_I cs$  in  $R_2$  filters all non-compatible events. Of the two equivalent semantics, we have chosen the less intuitive semantics of Fig. 1 as it is better suited to formal analysis. In fact, considering input events equipped with a security level is precisely what is needed in order to define a notion of confidentiality for systems with dynamic channels (see Section 3).

## 2.2 Reactive Behaviour of a JavaScript-like Language

Reactive behaviours are independent of the features of the programming languages that generate them, which may be, for example, imperative or functional, call-by-name or call-by-value, etc.

In this section, we show how to construct reactive behaviours from programs in a JavaScript-like language. Figure 2 presents the syntax of a language, which is inspired by JavaScript models found in previous works [7, 27, 39].

Programs are defined as a list of event handlers. A handler has the form  $(ch(x)\{c\})$ , and associates a channel  $ch$  to a command  $c$  to be executed when an input message (stored

$$\begin{array}{l}
p ::= \cdot \mid h; p \\
h ::= ch(x) \{c\} \\
c ::= \text{skip} \\
\quad \mid c; c \\
\quad \mid x := e \\
\quad \mid \text{if } e \{c\} \{c\} \\
\quad \mid \text{while } e \{c\} \\
\quad \mid \text{out}(ch, e) \\
\quad \mid \text{new } h \\
\quad \mid \text{open}(ch, \ell) \\
\quad \mid \text{close}(ch)
\end{array}$$

**Figure 2.** A JavaScript-like language.

in variable  $x$ ) is received on channel  $ch$ . A command  $c$  is an imperative program augmented with output messages (out), handler creation (new), and dynamic creation (open) and deletion (close) of communication channels.

Command  $\text{open}(ch, \ell)$  opens a new communication channel  $ch$  at security level  $\ell$ . Command  $\text{close}(ch)$  closes channel  $ch$ . Command  $\text{new } h$  associates to an open channel a new event handler (possibly replacing a previous handler.) Opening an open channel, closing a closed channel, or associating a handler to a closed channel, results in a run-time error.

To simplify the construction of the reactive behaviour, we assume that every handler uses variable  $x$  to store the message received, and we extend the set of commands  $\mathbb{C}$  with a special symbol  $\phi$ .

A program state is a tuple  $(cs, hs, \mu, c)$ , where:

- ▶  $cs : \text{CS} \rightarrow \mathcal{L}$  is a partial mapping from channels to security levels: the set of channels being used by the program to interact with the environment;
- ▶  $hs : \text{dom}(cs) \rightarrow \mathbb{C} \cup \{\phi\}$  is the set of event handlers associated to input channels in  $cs$ ; if an input channel  $ch$  does not have an event handler associated to it, then a default handler  $hs(ch) = \phi$  is assumed;
- ▶  $\mu$  is a mapping from variables to values, the store; and
- ▶  $c$  is the command being executed in response to an input event; or  $\phi$  if the system is in a consumer state (the next interaction will be an input event).

We write  $(\text{State } cs)$  for the set of states whose first component is  $cs$ . Given a program state  $s = (cs, hs, \mu, c)$ , we can compute the next interaction of  $s$  with the environment. The set of possible interactions is defined as

$$\text{Interaction} = \{\downarrow, \uparrow_v^{ch}, \odot, \uparrow_\ell^{ch}, \downarrow^{ch}, \otimes\}.$$

Interaction  $\downarrow$  is raised in a consumer state, i.e. when  $c = \phi$ . Interaction  $\uparrow_v^{ch}$  is raised when  $s$  will produce an output message to an open channel  $ch$  with value  $v$ . Interactions  $\uparrow_\ell^c$

```

step (cs, hs, μ, c) = case c of
  ϕ                → (↓, (cs, hs, μ, ϕ))
  out(ch, e)       → if ch ∈ dom(cs)
                    then (↑[[e]]μch, (cs, hs, μ, ϕ))
                    else (⊙, (cs, hs, μ, ϕ))
  open(ch, ℓ)      → if ch ∉ dom(cs)
                    then (↑ℓch, (cs[ch ↦ ℓ], hs, μ, ϕ))
                    else (⊙, (cs, hs, μ, ϕ))
  close(ch)        → if ch ∈ dom(cs)
                    then (↓chch, (del(cs, ch), del(hs, ch), μ, ϕ))
                    else (⊙, (cs, hs, μ, ϕ))
  new ch(x) {c'}   → if ch ∈ dom(cs)
                    then (⊙, (cs, hs[ch ↦ c'], μ, ϕ))
                    else (⊙, (cs, hs, μ, ϕ))
  skip             → (⊙, (cs, hs, μ, ϕ))
  x := e           → (⊙, (cs, hs, μ[x ↦ [[e]]μ], ϕ))
  if e {c0} {c1} → if [[e]]μ = 0
                    then (⊙, (cs, hs, μ, c1))
                    else (⊙, (cs, hs, μ, c0))
  while e {c'}    → if [[e]]μ = 0
                    then (⊙, (cs, hs, μ, ϕ))
                    else (⊙, (cs, hs, μ, c'; while e {c'}))
  c0; c1        → let (i, (cs', hs', μ', c')) = step (cs, hs, μ, c')
                    in if c' = ϕ then (i, (cs', hs', μ', c1))
                    else (i, (cs', hs', μ, c'; c1))

```

**Figure 3.** Next interaction for a program state

and  $\downarrow^c$  are used to signal opening and closure of channels. A silent interaction is represented by  $\odot$ , and interaction  $\ominus$  represents termination. Our language has no instruction to signal that execution should terminate. However, interaction  $\ominus$  is used to denote an abnormal termination, as it may happen, for example, if a program tries to output a message on a closed channel. Figure 3 defines a function

$$\text{step} : \forall cs, cs'. \text{State } cs \rightarrow \text{Interaction} \times \text{State } cs',$$

which computes the next interaction with the environment of a state  $s : \text{State } cs$ , together with the new state  $s' : \text{State } cs'$  that results from such interaction.

The definition of step makes use of an evaluation function for expressions;  $[[e]]_\mu = v$  iff expression  $e$  evaluates to value  $v$  in memory  $\mu$ . We assume that  $[[\cdot]]$  is a side-effect free function. We write  $f[x \mapsto y]$  for the mapping that behaves like  $f$  except on  $x$ , where it maps to  $y$ , and  $\text{del}(f, a)$  for  $f \setminus \{(a, f a)\}$ .

If there is no command to execute, the system is ready to process the next input (interaction  $\downarrow$ ). Assignments, conditionals, loops and skip generate internal activities, raising interaction  $\odot$ . A sequence  $c_0; c_1$  of instructions generates the interaction of  $c_0$ . If  $ch$  is open in  $cs$ , command  $\text{out}(ch, e)$  generates an output message interaction  $\uparrow_{[[e]]_\mu}^{ch}$ ; otherwise an error condition  $\ominus$  is raised, since the program is trying to send a message to a closed channel. Command  $\text{new}(ch)\{c\}$  interacts silently  $\odot$ , and associates command  $c$  to channel  $ch$  in the handler set, provided that  $ch$  is open in the actual state. Otherwise,  $\ominus$  is raised. If  $ch$  is closed in  $cs$ , command  $\text{open}(ch, \ell)$  adds  $(ch, \ell)$  to the channel system and generates interaction  $\uparrow_\ell^{ch}$ . Command  $\text{close}(ch)$  raises inter-

```

[[−]] : ∀cs.State cs → React cs
[[s]] = let (i, s') = step(s) in
  case i of
    ⊙ → Step ([[s']])
    ⊙ → Stop
    ↑vch → Write (ch, v, cs(ch)) [[s']]
    ↑ℓch → New (ch, ℓ) [[s']]
    ↓chch → Close ch [[s']]
    ↓ → Read (λe. [[update(s', e)])]

```

update :  $\forall cs.State cs \times E_I cs \rightarrow \text{State } cs$

update  $((cs, hs, \mu, c), (ch, v, \ell)) = (cs, hs, \mu[x \mapsto v], hs(ch))$

**Figure 4.** Reactive behaviour for program states

action  $\downarrow^{ch}$ , provided that  $ch$  is open in the current state. If  $ch$  is open (close) in the current state,  $\text{open}(ch, \ell)$  ( $\text{close}(ch)$ ) raises  $\odot$ .

Iterating over step, one can generate a reactive behaviour  $t : \text{React } cs$  from a state  $s : \text{State } cs$ , as it is shown in Fig. 4.

**Definition 2.4** (Reactive behaviour for a JavaScript-like program). Given a channel system  $cs$  and a memory  $\mu$ , the interpretation of  $p$  in  $(cs, \mu)$  is defined as

$$[[p]]_{cs, \mu} = [[(cs, \text{handlers}(cs, p), \mu, \phi)],$$

where

```

handlers                : CS → C ∪ {ϕ}
handlers(cs, ·)         = ∅
handlers(cs, ch(x){c}; p) = let hs = handlers(cs, p)
                          in if ch ∈ dom(cs)
                              then {(ch, c)} ∪ hs
                              else hs

```

is the set of handlers in  $p$  with open channels in  $cs$ .

### 3. Security Conditions for Reactive Systems

Confidentiality for a reactive system can be defined in terms of a notion of similarity of colists of input and output events. In a system with dynamic channels, the security level of an event depends on the state of the system at the time of the arrival of the event. For example, an `onclick` event can be classified as public at first and as secret later. In terms of reactive behaviours, the security level of a channel might be changed by closing the channel and opening it at a different security level. In this context, the analysis of similarity of two colists for an observer at a certain security level is simplified when the security level is part of the event as we do not need to depend on the system state. This is the main reason for considering events on  $A$  as triples  $\text{Chan} \times A \times \mathcal{L}$ . Once the security level is included in the events, the definitions are analogous to those of systems with static channels. We recap the concepts about security of reactive systems presented in previous work [39], slightly modified to work in our setting.

### 3.1 Reactive Noninterference

The security level of an event is defined as  $\text{lvl}(c, v, \ell) = \ell$ . The predicate  $\text{visible}_\ell$  on events determines when an event is observable for an observer at level  $\ell$ :

$$\frac{\text{lvl}(e) \sqsubseteq \ell}{\text{visible}_\ell(e)} \quad \frac{}{\text{visible}_\ell(\circ)}$$

Termination ( $\circ$ ) is visible at all levels. Silent steps ( $\bullet$ ), on the other hand, are not visible at any security level.

A reactive behaviour is non-interferent when similar inputs produce similar outputs for observers at any security level.

**Definition 3.1** (Security for Reactive Behaviours). Given a family of similarity relations  $\sim_\ell$  on colists, we say that a reactive behaviour  $t : \text{React } cs$  is secure iff, for all  $\ell$  and input colists  $i, i'$  such that  $i \sim_\ell i'$ , if  $(t, i) \Rightarrow_{cs} o$  and  $(t, i') \Rightarrow_{cs} o'$ , then  $o \sim_\ell o'$ .

Note that the definition depends on a notion of similarity. Depending on the power of the attacker one wants to model, one chooses a notion of similarity that makes more or less distinctions and, as a result, one obtains stronger or weaker notions of security. Bohannon et al. [7] identify two notions of security with practical interest: ID- and CP-security.

**Definition 3.2** (ID-similarity). ID-similarity between colists for an observer at level  $\ell$  is formalised coinductively by the following rules.

$$\frac{}{\llbracket \sim_\ell^{\text{ID}} \rrbracket} \quad \frac{\neg \text{visible}_\ell(e) \quad s \sim_\ell^{\text{ID}} s'}{e :: s \sim_\ell^{\text{ID}} s'} \quad \frac{\neg \text{visible}_\ell(e) \quad s \sim_\ell^{\text{ID}} s'}{s \sim_\ell^{\text{ID}} e :: s'} \quad \frac{\text{visible}_\ell(e) \quad s \sim_\ell^{\text{ID}} s'}{e :: s \sim_\ell^{\text{ID}} e :: s'}$$

Intuitively, two colists are ID-similar at level  $\ell$  if there is no evidence that they produce different events for an observer at level  $\ell$ . In particular, an infinitely silent colist –such as one produced by a divergent computation– is ID-similar to any other colist because we never find evidence against it. A feature of ID-similarity is that it is symmetric and reflexive, but not transitive. If it were, we could not have an infinitely silent colist ID-similar to any other as transitivity would imply that every colist is similar to each other.

In order to take progress into account, and distinguish a productive colist from a silent one, we consider a stronger notion of non-interference for reactive systems called CP-security [7].

We coinductively define when a colist of events is not visible (silent) for an observer at level  $\ell$  as follows.

$$\frac{\neg \text{visible}_\ell(e) \quad \text{silent}_\ell(s)}{\text{silent}_\ell(e :: s)} \quad \frac{}{\text{silent}_\ell(\llbracket \rrbracket)}$$

We define a relation that identifies the next event that is visible to an observer at security level  $\ell$  (if it exists).

Intuitively, we say that  $s \triangleright_\ell e :: s'$  when  $e$  is the next event in  $s$  visible at level  $\ell$ . The following rules inductively define the relation  $\triangleright_\ell$ .

$$\frac{\text{visible}_\ell(e)}{e :: s \triangleright_\ell e :: s} \quad \frac{\neg \text{visible}_\ell(e) \quad s \triangleright_\ell e' :: s'}{e :: s \triangleright_\ell e' :: s'}$$

Note that the relation is inductively defined, which means that when  $s \triangleright_\ell e :: s'$ , the next  $\ell$ -visible event  $e$  of the colist  $s$  must come after a finite sequence of  $\ell$ -invisible events.

**Definition 3.3** (CP-similarity). CP-similarity between colists is defined coinductively by the following rules.

$$\frac{\text{silent}_\ell(s) \quad \text{silent}_\ell(s')}{s \sim_\ell^{\text{CP}} s'} \quad \frac{s \triangleright_\ell e :: s_1 \quad s' \triangleright_\ell e :: s'_1 \quad s_1 \sim_\ell^{\text{CP}} s'_1}{s \sim_\ell^{\text{CP}} s'}$$

As opposed to ID-similarity, the notion of CP-similarity requires proof that if one colist produces a visible event  $e$ , the next observable event in the other will be  $e$  and moreover, it will be produced in a finite number of steps. Also, unlike ID-similarity, CP similarity is an equivalence relation.

By instantiating Definition 3.1 to ID- and CP-similarity we obtain the notions of ID- and CP-security for reactive behaviours.

### 3.2 Run-based Security

It is well known that non-interference is not a safety property and cannot be precisely enforced by execution monitors [16, 32]. Therefore, we introduce a security condition which is defined on runs, and characterises the set of *secure inputs*, i.e. inputs for which a program does not leak information. One can think that if a program  $p$  receives an input  $i$  with no secrets, then executing  $p$  by feeding it  $i$  is secure, since there are no secrets to leak<sup>1</sup>. Using this intuition, we classify an input  $i$  as secure if it reveals the same information as the input where secrets have been erased.

We coinductively define the relation  $\blacktriangleright_\ell$  responsible for removing all the events unobservable at level  $\ell$ .

$$\frac{\text{silent}_\ell(s)}{s \blacktriangleright_\ell \llbracket \rrbracket} \quad \frac{s \triangleright_\ell e :: s' \quad s' \blacktriangleright_\ell s''}{s \blacktriangleright_\ell e :: s''}$$

Observe that given a colist  $s$ , there is a unique colist  $s'$  such that  $s \blacktriangleright_\ell s'$ . We will write  $s_{\blacktriangleright_\ell}$  for this unique colist, and refer to it as the *restriction* of  $s$  at level  $\ell$ .

Let us assume a level-indexed similarity relation  $\sim_\ell$  between colists. Two inputs for a program reveal the same secrets at a given security level  $\ell$  if, for an observer at level  $\ell$ , they are similar and induce similar outputs.

**Definition 3.4** ( $\approx_{\ell,t}$ ). Let  $t \in \text{React}$ ,  $\ell \in \mathcal{L}$ , and  $i, i'$  input colists such that  $i \sim_\ell i'$ ,  $(t, i) \Rightarrow o$  and  $(t, i') \Rightarrow o'$ . We say

<sup>1</sup> This observation is valid if we ignore covert channels, in particular, the termination channel.

that the program  $t$  reveals the same  $\ell$ -secrets when given the inputs  $i, i'$ , noted  $i \approx_{\ell, t} i'$ , iff  $o \sim_{\ell} o'$ .

Similarly to [6], we consider an input to be *secure* for a program  $t$  if it reveals the same information about the secrets as the input where secrets have been erased.

**Definition 3.5** (Secure input). Let  $t$  be a reactive behaviour. An input colist  $i$  is *secure for  $t$*  iff  $\forall \ell. i \approx_{\ell, t} i \blacktriangleright_{\ell}$ . We say that the input  $i$  is ID-secure (CP-secure) for  $t$  when  $\sim_{\ell}$  is instantiated to  $\sim_{\ell}^{\text{ID}}$  ( $\sim_{\ell}^{\text{CP}}$ ) in Definition 3.4.

**Example 3.6.** Consider the following program

$$\begin{aligned} p = & \text{c?}_0(\mathbf{x}) \{ \mathbf{r} := \mathbf{x} \}; \\ & \text{c?}_1(\mathbf{x}) \{ \text{if } \mathbf{r} \geq 1 \vee \mathbf{x} = 0 \\ & \quad \{ \text{new c?}_2(\mathbf{x}) \{ \text{out}(\text{c!}_0, \mathbf{r}) \} \} \\ & \quad \{ \text{while } 1 \{ \text{skip} \} \} \}; \\ & \text{c?}_2(\mathbf{x}) \{ \text{skip} \} \end{aligned}$$

Let us assume an initial memory ( $\mu_0 = \lambda x.0$ ) and an initial channel system

$$cs = \{ (\text{c?}_0, \mathbf{H}), (\text{c?}_1, \mathbf{L}), (\text{c?}_2, \mathbf{L}), (\text{c!}_0, \mathbf{L}) \},$$

where  $\mathbf{L}$  and  $\mathbf{H}$  are security levels such that  $\mathbf{L} \sqsubseteq \mathbf{H}$ . Let  $t : \text{React } cs$  be the interpretation of program  $p$  in  $cs$ . That is,  $t = \llbracket p \rrbracket_{cs, \mu_0}$ . Consider an input  $i$  with three messages as follows:

$$i = [(\text{c?}_0, v, \mathbf{H}), (\text{c?}_1, u, \mathbf{L}), (\text{c?}_2, 42, \mathbf{L})]$$

If  $v = 1$  and  $u = 0$ , then  $i$  is neither ID- nor CP-secure for  $t$ . However, if  $v = u = 1$ , we have that  $i$  is ID-secure, but not CP-secure for  $t$ . Taking  $v = u = 0$ , input  $i$  becomes ID- and CP-secure for  $t$ .

The evaluation relation defined in Figure 1 is only defined for compatible inputs, which means that for incompatible inputs evaluation will get stuck. However, as shown by the following example, erasing the secrets can turn a compatible input into an incompatible one.

**Example 3.7.** Consider the following program:

$$\begin{aligned} p = & \text{c?}_0(\mathbf{x}) \{ \mathbf{r} := \mathbf{x} \}; \\ & \text{c?}_1(\mathbf{x}) \{ \text{if } \mathbf{r} \geq 1 \\ & \quad \{ \text{open}(\text{c?}_2, \mathbf{L}); \\ & \quad \quad \text{new c?}_2(\mathbf{x}) \{ \text{out}(\text{c!}_0, \mathbf{r}) \} \} \\ & \quad \{ \text{skip} \} \}; \end{aligned}$$

Let  $cs = \{ (\text{c?}_0, \mathbf{H}), (\text{c?}_1, \mathbf{L}), (\text{c!}_0, \mathbf{L}) \}$ , be the initial channel system, let  $t : \text{React } cs$  be such that  $t = \llbracket p \rrbracket_{cs, \mu_0}$ . The insecure input  $i = [(\text{c?}_0, 1, \mathbf{H}), (\text{c?}_1, 1, \mathbf{L}), (\text{c?}_2, 1, \mathbf{L})]$  is compatible with  $t$ :

$$(t, i) \Rightarrow_{cs} [\bullet, \bullet, \bullet, \bullet, \bullet, (\text{c!}_0, 1, \mathbf{L})]$$

However, its restriction  $i \blacktriangleright_{\mathbf{L}} = [(\text{c?}_1, 1, \mathbf{L}), (\text{c?}_2, 1, \mathbf{L})]$  is not, as channel  $\text{c?}_2$  is never opened.

Hence, the filtering of a compatible input may be incompatible. Moreover, the fact that a program may get stuck with a filtered input does not imply that the (unfiltered) input is insecure for that program.

Therefore, in order to be able to use the notion of secure-run we need to account for incompatible input events in the evaluation relation. We extend the evaluation relation with the following rule:

$$R_3 \frac{e \notin E_I \text{ cs} \quad (\text{Read } f, i) \Rightarrow_{cs} o}{(\text{Read } f, e :: i) \Rightarrow_{cs} \bullet :: o}$$

Intuitively, Rule  $R_3$  simply discards the next input event if it is not compatible with the actual channel system. With this additional rule, the evaluation relation is defined for every input and it becomes a function.

We can now prove the following relation between ID-secure and CP-secure inputs:

**Lemma 3.8.** *Let  $t : \text{React } cs$  and  $i$  an input colist. If  $i$  is CP-secure for  $t$ , then  $i$  is ID-secure for  $t$ .*

### 3.3 From Run-based Security to Noninterference and Back

We have the notions of ID-security and CP-security which characterise the security of programs, and the notions of ID- and CP-secure inputs which characterise the security of runs. As it is shown in [39], the security for programs and runs is closely related for CP-similarity.

**Lemma 3.9** (Secure inputs and CP-security). *A reactive behaviour  $t \in \text{React } cs$  is CP-secure iff  $\forall i \in \mathcal{I}. i$  is CP-secure for  $t$ .*

It is easy to see that all inputs are CP-secure for a CP-secure program. In order to see why a program such that every input is CP-secure is a CP-secure program, observe that whenever two inputs  $i, i'$  are similar at some level  $\ell$  they have exactly the same restriction at that level.

$$i \sim_{\ell}^{\text{CP}} i' \implies i \blacktriangleright_{\ell} = i' \blacktriangleright_{\ell}$$

Using the transitivity of  $\sim_{\ell}^{\text{CP}}$  and the definition of secure input, one obtains that the output colists produced by  $t$  when executed on  $i$  and  $i'$  must be CP-similar, provided that  $i$  and  $i'$  are CP-secure.

The strong correspondence between security of programs and security of runs does not hold when we consider ID-security, and we only have the property that all inputs are secure for secure programs.

**Lemma 3.10** (Secure inputs and ID-security). *If a reactive behaviour  $t \in \text{React } cs$  is ID-secure, then every input  $i \in \mathcal{I}$  is ID-secure for  $t$ .*

Even though all inputs may be ID-secure for a program, the program might be interferent, as shown by the following example.

**Example 3.11.** Consider the following program

$$p = c?_0(x) \{ r := x \};$$

$$c?_1(x) \{ \text{if } r \geq 1$$

$$\quad \{ \text{out}(c!_0, r) \}$$

$$\quad \{ \text{while } 1 \{ \text{skip} \} \} \};$$

where we assume an initial channel system

$$cs = \{(c?_0, H), (c?_1, L), (c!_0, L)\}.$$

Every input is ID-secure for  $p$  but  $p$  is not an ID-secure program, as the following two ID-similar inputs show.

$$i = [(c?_0, 1, H), (c?_1, 0, L)]$$

$$\quad \sim_L^{\text{ID}}$$

$$i' = [(c?_0, 2, H), (c?_1, 0, L)]$$

Let  $t = \llbracket p \rrbracket_{cs, \mu_0}$ , the reactive behaviour obtained from  $p$  and  $\mu_0$  in  $cs$ . Then, we see that  $t$  is not ID-secure, since  $i$  and  $i'$  are ID-similar at level  $L$ , but their outputs are not ID-similar at level  $L$ .

$$(t, i) \Rightarrow [\bullet, \bullet, \bullet, \bullet, (c!_0, 1, L)]$$

$$\quad \sim_L^{\text{ID}}$$

$$(t, i') \Rightarrow [\bullet, \bullet, \bullet, \bullet, (c!_0, 2, L)]$$

Nevertheless, all inputs  $i$  are ID-secure for  $t$  (Definition 3.5). The key observation here is that  $t$  diverges for  $i \blacktriangleright_L$  (which coincides with  $i' \blacktriangleright_L$ ), since  $r$  was initially zero. In other words, the input colist without events on channel  $c?_0$  produces an output which is silent and infinite, and hence ID-similar to every other output.

## 4. Monitoring Reactive Behaviours

In this section, we present the main contribution of the paper: a monitor for reactive behaviours. The monitor is based on a multi-execution of a reactive behaviour: for each security level in  $\ell \in \mathcal{L}$ , a modified reactive behaviour is created, named the *producer* at level  $\ell$ . A producer at level  $\ell$  is fed only input events visible at security level  $\ell$ , and it produces outputs only at level  $\ell$ . Therefore, it is easy to see that each producer cannot leak secrets, since it never has any secrets available.

In order to detect if a run is secure or not, the following procedure is followed: if the original program produces an output event at security level  $\ell$ , then:

- ▶ only the producer at level  $\ell$  can produce an output at such security level;
- ▶ if such producer outputs the same event, it is safe, since producers do not leak secrets;
- ▶ if such producer outputs a different event, then the run is unsafe.

Hence by the coordinated execution of the original program and the producers, one can replicate the behaviour of the original program for as long as the program is secure, and detect when a run becomes insecure.

$$\begin{aligned} \text{prd} & : \forall \ell, cs. \text{React } cs \rightarrow \text{React } cs \downarrow \ell \\ \text{prd}_{\ell, cs}(\text{Read } f) & = \text{Read } (\lambda e. \text{prd}_{\ell, cs}(f e)) \\ \text{prd}_{\ell, cs}(\text{Write } e t) & = \text{if } \ell = \text{lvl}(e) \\ & \quad \text{then Write } e (\text{prd}_{\ell, cs} t) \\ & \quad \text{else Step } (\text{prd}_{\ell, cs} t) \\ \text{prd}_{\ell, cs}(\text{Step } t) & = \text{Step } (\text{prd}_{\ell, cs} t) \\ \text{prd}_{\ell, cs}(\text{New } (c, \ell') t) & = \text{if } \ell' \sqsubseteq \ell \\ & \quad \text{then New } (c, \ell') (\text{prd}_{\ell, cs \oplus (c, \ell')} t) \\ & \quad \text{else Step } (\text{prd}_{\ell, cs \oplus (c, \ell')} t) \\ \text{prd}_{\ell, cs}(\text{Close } c t) & = \text{if } cs \sqsubseteq \ell \\ & \quad \text{then Close } c (\text{prd}_{\ell, cs \setminus \{(c, \ell)\}} t) \\ & \quad \text{else Step } (\text{prd}_{\ell, cs \setminus \{(c, \ell)\}} t) \\ \text{prd}_{\ell, cs}(\text{Stop}) & = \text{Stop} \end{aligned}$$

**Figure 5.** Producer creation at level  $\ell$

### 4.1 Creation of Producers

A producer at level  $\ell$  is a modification of a reactive behaviour in such a way that it only outputs events at level  $\ell$ . Moreover, the channel system of the producer at level  $\ell$  is restricted to channels of at most level  $\ell$ . More precisely, the *restriction of a channel system*  $cs$  to level  $\ell$  is given by:

$$cs \downarrow \ell = \{(c, \ell') \in cs \mid \ell' \sqsubseteq \ell\}$$

The creation of producers is defined in Figure 5. When the original behaviour does a Read, the producer reads the event and continues execution with the producer corresponding to the reaction to that event. Because the channel system of a producer is restricted to  $\ell$ , the security level of this input event is necessarily less or equal to  $\ell$ . In the case of a Write, the producer will perform it only when the level of the channel is  $\ell$ , and do a silent step otherwise. Silent steps (Step) and stops (Stop) are simply copied. The creation (New) or closure (Close) of a channel will be replicated, as long as it is compatible with the restricted channel system (i.e. the security level  $\ell'$  of the channel being created or closed is such that  $\ell' \sqsubseteq \ell$ .) On the contrary, if the security level of the channel being created or closed is not compatible with the restricted channel system, a silent step is generated and the modification to the channel system is ignored.

Note how the type of the transformation ensures that initially the reactive behaviour does not has access to secrets since only channels with public events are open, and how this invariant is preserved by the function  $\text{prd}$  which creates producers. Also, it is easy to see that a producer at security level  $\ell$  will only perform Writes at level  $\ell$  by simple inspection of the definition of  $\text{prd}$ .

### 4.2 Small-step Semantics of Reactive Behaviours

The monitor works by analysing the next step that would be performed by a reactive behaviour and comparing it with a producer. In order to be able to refer easily to the next step of a behaviour we introduce a small-step execution



$$\begin{array}{c}
\text{sR}_1 \frac{e \in E_I(cs)}{(cs, \text{Read } f, e :: i) \xrightarrow{\bullet} (cs, f e, i)} \\
\text{sR}_2 \frac{e \notin E_I(cs)}{(cs, \text{Read } f, e :: i) \xrightarrow{\bullet} (cs, \text{Read } f, i)} \\
\text{sW} \frac{}{(cs, \text{Write } e t, i) \xrightarrow{e} (cs, t, i)} \\
\text{sS} \frac{}{(cs, \text{Step } t, i) \xrightarrow{\bullet} (cs, t, i)} \\
\text{sN} \frac{cs' = cs \oplus (c, \ell)}{(cs, \text{New } (c, \ell) t, i) \xrightarrow{\bullet} (cs', t, i)} \\
\text{sC} \frac{cs(c) = \ell \quad cs' = cs \setminus \{(c, \ell)\}}{(cs, \text{Close } c t, i) \xrightarrow{\bullet} (cs', t, i)}
\end{array}$$

**Figure 6.** Small step semantics for React

relation for React. The rules in Figure 6 define a relation  $\longrightarrow \subseteq \text{Conf} \times \mathcal{E}_{O \cup \{\bullet\}} \times \text{Conf}$ , where a configuration is a (dependent) tuple  $\text{Conf} = (cs : \text{CS}) \times \text{React } cs \times \mathcal{I}$  consisting of a channel system, a reactive behaviour for that channel system, and a colist of input events. As usual, we write  $\sigma \xrightarrow{e} \sigma'$  for  $(\sigma, e, \sigma') \in \longrightarrow$ .

Intuitively, a Read is only actually performed when the next input event is compatible (sR<sub>1</sub>). If the event is incompatible it is discarded without reading (sR<sub>2</sub>). A Write  $e t$  outputs an event  $e$  and continues executing  $t$  (sW). A Step  $t$  performs a silent step and continues with  $t$  (sS). A New  $(c, \ell) t$  performs a silent step and continues with  $t$  on a channel system extended with  $c$  open at level  $\ell$  (sN). A Close  $c t$  performs a silent step and continues with  $t$  on a channel system with channel  $c$  closed.

There are two ways in which a configuration may be final: when doing a Read on an empty input colist, and on a Stop. The predicates on configurations  $\text{end}_\bullet$  and  $\text{end}_\circ$  capture these two situations.

$$\frac{}{\text{end}_\bullet(cs, \text{Read } f, [])} \quad \frac{}{\text{end}_\circ(cs, \text{Stop}, i)}$$

### 4.3 A Precise Monitor

We define the monitor for reactive behaviours by a family of evaluation relations  $\Downarrow_{cs}$  indexed by a channel system. Like the evaluation relation in Figure 1, the monitor evaluation relation tell us which outputs a reactive behaviour might generate given a certain input. However, the monitor may also rise an alarm if it detects an insecure run. More precisely the monitor is given by a relation

$$\Downarrow_{cs} \subseteq \text{React } cs \times \mathcal{I} \times (\mathcal{L} \rightarrow \text{Conf}) \times \mathcal{O}_\epsilon$$

where  $\mathcal{O}_\epsilon$  is similar to the colist of output events  $\mathcal{O}$ , but it may also end with an alarm  $\epsilon$ . Formally, we coinductively

define  $\mathcal{O}_\epsilon$  as follows.

$$\mathcal{O}_\epsilon = [] \mid \epsilon \mid (\mathcal{E}_O \cup \{\bullet, \circ\}) :: \mathcal{O}_\epsilon$$

We write  $(t, i, \bar{\sigma}) \Downarrow_{cs} o$  for  $(t, i, \bar{\sigma}, o) \in \Downarrow_{cs}$  indicating that the reactive behaviour  $t$ , with input  $i$ , and producers  $\bar{\sigma}$  (we have one configuration for each producer at each security level, hence the vector notation) produce an output  $o$ .

The monitor evaluation relation is given by the rules in Figure 7. Rule SIL indicates it is always safe for a reactive behaviour to perform a silent step. Note that in doing the silent step the channel system may change, as it happens when opening or closing channels (see Fig. 6.) The other rules concern the case where a reactive behaviour wants to perform an operation which might produce a visible event, in which case the monitor needs to check that the corresponding producer outputs the same event. The most obvious case of visible event is when a reactive behaviour wants to output an event by doing a Write (rules VIS<sub>1-4</sub>.) Rule VIS<sub>1</sub> applies when the event to be written coincides with the output of the producer at the corresponding security level. In this case, the output is safe and the event is outputted. It might happen that the producer needs to do some silent steps before it outputs an event. Rule VIS<sub>2</sub> takes care of this case by advancing the computation on the producer. If the producer outputs an event, but it does not coincide with the one the reactive behaviour wants to output, we are in presence of an insecure run and an alarm is raised (rule VIS<sub>3</sub>.) If the producer configuration at the event level is final, then it cannot produce such an event and an alarm is raised (rule VIS<sub>4</sub>.) Finally, the last two rules (END<sub>1</sub> and END<sub>2</sub>) concerns the case where the reactive behaviour would finish executing. Before concluding that finishing would not leak information, the monitor needs to make sure that every producer finishes silently. Since we distinguish between finishing by emptying the input colist and finishing by a Stop (see Fig. 1) we need to treat the two cases differently. Therefore we need two rules with two different functions (done and stop) that will execute producers until they finish in the appropriate manner, or until they produce some visible output, raising an alarm. Note that even if the reactive behaviour terminates, some producer may diverge, and hence, the monitor may diverge. As we will see in the next section, this may only happen if the run is insecure.

The following example shows why, even if the reactive behaviour was going to end, rules END<sub>1</sub> and END<sub>2</sub> are needed.

**Example 4.1.** Consider the following program:

$$\begin{aligned}
p = & \text{c?}_0(x) \{ r := x \}; \\
& \text{c?}_1(x) \{ \text{if } r = 0 \\
& \quad \{ \text{out}(c!_0, 1) \} \\
& \quad \{ \text{skip} \} \};
\end{aligned}$$

$$\begin{array}{c}
\text{SIL} \frac{(cs, t, i) \xrightarrow{\bullet} (cs', t', i') \quad (t', i', \bar{\sigma}) \Downarrow_{cs'} o}{(t, i, \bar{\sigma}) \Downarrow_{cs} \bullet :: o} \\
\\
\text{VIS}_1 \frac{\text{lvl}(e) = \ell \quad \bar{\sigma}(\ell) \xrightarrow{e} \sigma' \quad (t, i, \bar{\sigma}[\ell \mapsto \sigma']) \Downarrow_{cs} o}{(\text{Write } e \ t, i, \bar{\sigma}) \Downarrow_{cs} e :: o} \\
\\
\text{VIS}_2 \frac{\text{lvl}(e) = \ell \quad \bar{\sigma}(\ell) \xrightarrow{\bullet} \sigma' \quad (\text{Write } e \ t, i, \bar{\sigma}[\ell \mapsto \sigma']) \Downarrow_{cs} o}{(\text{Write } e \ t, i, \bar{\sigma}) \Downarrow_{cs} \bullet :: o} \\
\\
\text{VIS}_3 \frac{\text{lvl}(e) = \ell \quad \bar{\sigma}(\ell) \xrightarrow{e'} \sigma' \quad e' \neq \bullet \quad e' \neq e}{(\text{Write } e \ t, i, \bar{\sigma}) \Downarrow_{cs} \epsilon} \\
\\
\text{VIS}_4 \frac{\text{lvl}(e) = \ell \quad (\text{end}_{\bullet}(\bar{\sigma}(\ell)) \vee \text{end}_o(\bar{\sigma}(\ell)))}{(\text{Write } e \ t, i, \bar{\sigma}) \Downarrow_{cs} \epsilon} \\
\\
\text{END}_1 \frac{}{(\text{Read } f, [], \bar{\sigma}) \Downarrow_{cs} \text{done}(\bar{\sigma})} \\
\\
\text{END}_2 \frac{}{(\text{Stop}, i, \bar{\sigma}) \Downarrow_{cs} \text{stop}(\bar{\sigma})}
\end{array}$$

**Figure 7.** Monitor for Reactive Behaviours

Let  $t$  be the reactive behaviour obtained from it for the initial memory  $\mu_0$  and channel system

$$cs = \{(c?_0, H), (c?_1, L), (c!_0, L)\}$$

and let  $i = [(c?_0, 1, H), (c?_1, 0, L)]$ . The evaluation of  $t$  with input  $i$  results in a finite sequence of invisible events.

$$(t, i) \Rightarrow_{cs} [\bullet, \bullet, \bullet, \bullet, \bullet]$$

However, the evaluation of  $t$  when fed with the restriction of  $i$  at level  $L$ , i.e.  $i_{\blacktriangleright L} = [(c?_1, 0, L)]$ , has an observable event.

$$(t, i_{\blacktriangleright L}) \Rightarrow_{cs} [\bullet, \bullet, (c!_0, 1, L)]$$

Therefore, the outputs are distinguishable at level  $L$ , and we conclude that  $i$  is not ID-secure for  $t$ .

The above example illustrates the need to run all producers and check that they terminate in the same manner as the reactive behaviour. In the monitor, this check is performed by the following functions.

$$\text{done}(\bar{\sigma}) = \begin{cases} [] & \text{if } \forall \ell. \text{end}_{\bullet}(\bar{\sigma}(\ell)) \\ \epsilon & \text{if } \exists \ell. \text{isWrite}(\bar{\sigma}(\ell)) \\ & \vee \text{end}_o(\bar{\sigma}(\ell)) \\ \bullet :: \text{done}(\text{next}(\bar{\sigma})) & \text{otherwise} \end{cases}$$

$$\text{stop}(\bar{\sigma}) = \begin{cases} \epsilon & \text{if } \exists \ell. \text{isWrite}(\bar{\sigma}(\ell)) \\ & \vee \text{end}_{\bullet}(\bar{\sigma}(\ell)) \\ \bullet :: \text{stop}(\text{next}(\bar{\sigma})) & \text{otherwise} \end{cases}$$

The predicate on configurations `isWrite` tell us if a configuration is a Write:

$$\text{isWrite}(cs, \text{Write } e \ t, i)$$

Function `next` advances each producer a silent step (when possible):

$$\begin{array}{l}
\text{next} \quad : \quad (\mathcal{L} \rightarrow \text{Conf}) \rightarrow (\mathcal{L} \rightarrow \text{Conf}) \\
\text{next}(\bar{\sigma})(\ell) = \begin{cases} \sigma' & \text{if } \bar{\sigma}(\ell) \xrightarrow{\bullet} \sigma' \\ \bar{\sigma}(\ell) & \text{otherwise.} \end{cases}
\end{array}$$

Given a channel system  $cs$ , a reactive behaviour  $t$ , and an input colist  $i$ , the monitor is initialised in the following manner

$$\text{monitor}(cs, t, i) = (t, i, \lambda \ell. (cs \downarrow_{\ell}, \text{prd}_{\ell, cs}(t), i))$$

That is, producers are created with function `prd` in a restricted channel system. Since inputs on closed channels are ignored there is no need to filter the input colist of each producer.

## 5. Properties of the Monitor

Ideally, a monitor should be precise and transparent. Precision means that the monitor should raise an alarm only for insecure runs. Transparency, on the other hand, means that secure runs should be indistinguishable from an ordinary execution (i.e. with the monitor not present.)

In order to state these properties formally, we define the predicate  $\text{ok} \subseteq \mathcal{O}_e$  which identifies output colists where the monitor has not raised an alarm. The predicate is defined coinductively by the following rules.

$$\frac{}{\text{ok}([])} \quad \frac{\text{ok}(s)}{\text{ok}(e :: s)}$$

The monitor for reactive behaviours of the previous section is precise with respect to ID-security.

**Theorem 5.1** (Precision for ID-secure runs). *Let  $cs$  be a channel system,  $t : \text{React } cs$  a reactive behaviour and  $i$  an input colist such that  $\text{monitor}(cs, t, i) \Downarrow_{cs} o$ . Then,*

$$i \text{ is ID-secure for } t \iff \text{ok}(o)$$

Moreover, the monitor is transparent with respect to both ID- and CP-security.

**Theorem 5.2** (Transparency for secure runs). *Let  $cs : \text{CS}$  be a channel system,  $t : \text{React } cs$  a reactive behaviour on  $cs$ , and  $i$  an input colist, such that  $(t, i) \Rightarrow_{cs} o$  and  $\text{monitor}(cs, t, i) \Downarrow_{cs} o'$ . Then,*

$$i \text{ is ID-secure for } t \implies \text{ok}(o') \wedge \forall \ell. o \sim_{\ell}^{\text{D}} o', \quad (1)$$

and

$$i \text{ is CP-secure for } t \implies \text{ok}(o') \wedge \forall \ell. o \sim_{\ell}^{\text{CP}} o'. \quad (2)$$

However, the monitor is not precise with respect to CP-security: it could happen that an insecure run causes the monitor to diverge instead of raising an alarm.

As a corollary of Theorem 5.2, Lemma 3.10, and Lemma 3.9, we have the following result on behaviours.

**Corollary 5.3** (Transparency for secure behaviours). *Let  $cs$  be a channel system,  $t$  an (ID/CP)-non-interferent reactive behaviour on  $cs$ , and  $i$  an input colist  $i$ . Then*

$$\left. \begin{array}{l} (t, i) \Rightarrow_{cs} o \\ \text{monitor}(cs, t, i) \Downarrow_{cs} o' \end{array} \right\} \Longrightarrow o \sim^{\text{ID/CP}} o'.$$

Therefore, the monitor will always preserve the semantics of secure programs.

## 6. Related Work

There are many security policies enforceable by execution monitoring [16, 22, 32], but non-interference is not one of them [25, 32]. Therefore, monitors often enforce properties stronger than non-interference [1, 3, 4, 29, 31] losing precision, and hence raising false alarms.

Our approach to monitoring is an extension of the monitor by Zanarini, Jaskelioff, and Russo [39]. The idea of transforming programs according to their security level is inspired by secure multi-execution as introduced by Devriese and Piessens [15] for interactive systems. However, the idea of having different execution threads according to their security level had been used before [9, 11]. Barthe et al. [5] show that secure multi-execution can also be achieved by code transformation. Rafnsson and Sabelfeld [28] make several improvements to secure multi-execution of interactive systems, such as lifting the totality assumptions on input channels and distinguishing between presence and content of messages. Differently from the above mentioned works which focused on interactive systems, Bielova et al. [6] adapt secure multi-execution for reactive systems and, similarly to this work, consider security in terms of runs.

Zheng and Myers [40] consider dynamic security labels for a statically typed language. Nevertheless, none of the works based on secure multi-execution described above [5, 6, 15, 28, 39] consider dynamic channels.

De Groef et al. [12] implement FlowFox, a web browser that supports an information flow control mechanism based on SME. The actual implementation of FlowFox allows for policies that assign labels to API method calls based on a state maintained by the security policy. Hence, the policy can record in its state which channels are open. However, this support for dynamic channels is not present in the formal model of FlowFox [13], and hence no formal guarantees are provided when using this feature.

The notions for reactive systems of ID- and CP-non-interference were introduced by Bohannon et al. [7]. As opposed to ID-security, the notion of CP-security is progress-sensitive. However, it is more difficult to enforce. Rafnsson and Sabelfeld [27] show that ID- and CP-security are

termination-insensitive, and therefore susceptible to brute force attacks that exploit the termination channel. Our proposal, similarly to other modern information-flow tools [10, 26, 33] cannot avoid this kind of leaks. However, for deterministic systems, the bandwidth of leaking information by exploiting outputs in combination with termination is logarithmic in the size of the secret [2] and can be reduced by applying buffering techniques [27].

Our model of reactive behaviours is an extension of the description of interactions in [21] which in turn is inspired by [34]. In order to cope with dynamic channels it has been extended to use the expressive power of interaction structures [17].

## 7. Conclusions

We propose a monitor for reactive systems in which channels can be opened and closed dynamically. Our approach extends the multi-execution monitor of Zanarini, Jaskelioff, and Russo [39] which could only handle a fixed set of channels. Due to the generality of the approach in [39] and a careful choice of how to extend it, we were able to adapt many notions to our more complex setting of dynamic channels with only slight modifications. The main issue to consider was how to model non-interference. Once we decided that the right way was to pair each event with a security level, the rest of the design decisions followed up naturally. Interestingly, we obtained similar results as the previous monitor: our monitor is precise with respect to ID-security and transparent with respect to both ID and CP-security.

The addition of dynamic channels means that a channel may be closed and then opened at a different security level, and therefore we need to treat security levels dynamically. This opens the door for an approach to declassification: by adding a declassification primitive, we may allow the temporary change of the security level of a channel and still be able to provide formal security guarantees.

Finally, we would like to implement the monitor in a real browser and develop some applications for it. This would provide a reality check for our model of reactive systems and suggest areas where improvement is needed.

## Acknowledgments

We thank the anonymous reviewers for their constructive criticism, and Exequiel Rivas and Alejandro Russo for their insightful comments and generous help. This work was funded by the *Agencia Nacional de Promoción Científica y Tecnológica* (PICT 2009-15) and *Consejo Nacional de Investigaciones Científicas y Técnicas* (CONICET).

## References

- [1] A. Askarov and A. Sabelfeld. Tight Enforcement of Information-Release Policies for Dynamic Languages. In *Proceedings of the 22nd IEEE Computer Security Foundations*

- Symposium*, Washington, DC, USA, 2009. IEEE Computer Society.
- [2] A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands. Termination-Insensitive Noninterference Leaks More Than Just a Bit. In *Proceedings of the 13th European Symposium on Research in Computer Security: Computer Security*, ESORICS '08. Springer-Verlag, 2008.
- [3] T. H. Austin and C. Flanagan. Efficient Purely-Dynamic Information Flow Analysis. In *Proc. ACM Workshop on Programming Languages and Analysis for Security (PLAS)*, June 2009.
- [4] T. H. Austin and C. Flanagan. Permissive dynamic information flow analysis. In *Proceedings of the 5th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, PLAS '10. ACM, 2010.
- [5] G. Barthe, J. M. Crespo, D. Devriese, F. Piessens, and E. Rivas. Secure multi-execution through static program transformation. In *Formal Techniques for Distributed Systems (FMODS/FORTE 2012)*, June 2012.
- [6] N. Bielova, D. Devriese, F. Massacci, and F. Piessens. Reactive non-interference for a browser model. In *Proceedings of the 5th International Conference on Network and System Security (NSS 2011)*, Sept. 2011.
- [7] A. Bohannon, B. C. Pierce, V. Sjöberg, S. Weirich, and S. Zdancewic. Reactive noninterference. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, CCS '09. ACM, 2009.
- [8] S. B. Byrne, R. S. Sutor, J. Robie, G. Nicol, M. Champion, S. Isaacson, I. Jacobs, L. Wood, C. Wilson, and A. L. Hors. Document Object Model (DOM) Level 1. W3C recommendation, W3C, Oct. 1998. <http://www.w3.org/TR/1998/REC-DOM-Level-1-19981001>.
- [9] R. Capizzi, A. Longo, V. N. Venkatakrishnan, and A. P. Sistla. Preventing Information Leaks through Shadow Executions. In *Proceedings of the 2008 Annual Computer Security Applications Conference*, ACSAC '08. IEEE Computer Society, 2008.
- [10] R. Chapman and A. Hilton. Enforcing security and safety models with an information flow analysis tool. In *Proc. of the 2004 annual ACM SIGAda international conference on Ada*. ACM, 2004.
- [11] M. Cristiá and P. Mata. Runtime Enforcement of Noninterference by Duplicating Processes and their Memories. In *Workshop de Seguridad Informática WSEGI 2009, Argentina*, 38 JAIIO, 2009.
- [12] W. De Groef, D. Devriese, N. Nikiforakis, and F. Piessens. FlowFox: a web browser with flexible and precise information flow control. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12. ACM, 2012.
- [13] W. De Groef, D. Devriese, N. Nikiforakis, and F. Piessens. Secure multi-execution of web scripts: Theory and practice. *Journal of Computer Security*, 22(4):469–509, 2014. ISSN 0926-227X. . URL <https://lirias.kuleuven.be/handle/123456789/442492>.
- [14] D. E. Denning and P. J. Denning. Certification of Programs for Secure Information Flow. *Commun. ACM*, 20(7):504–513, July 1977.
- [15] D. Devriese and F. Piessens. Noninterference through Secure Multi-execution. In *Proc. of the 2010 IEEE Symposium on Security and Privacy*, SP '10. IEEE Computer Society, 2010.
- [16] K. W. Hamlen, G. Morrisett, and F. B. Schneider. Computability classes for enforcement mechanisms. *ACM Trans. Program. Lang. Syst.*, Jan. 2006.
- [17] P. Hancock and A. Setzer. Interactive programs in dependent type theory. In P. Clote and H. Schwichtenberg, editors, *Computer Science Logic. 14th international workshop, CSL 2000*, Springer Lecture Notes in Computer Science, Vol. 1862, pages 317 – 331, 2000.
- [18] D. Hedin and A. Sabelfeld. Information-Flow Security for a Core of JavaScript. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society, 2012.
- [19] C. Jackson, A. Bortz, D. Boneh, and J. C. Mitchell. Protecting browser state from web privacy attacks. In *Proceedings of the 15th international conference on World Wide Web*, WWW '06. ACM, 2006.
- [20] D. Jang, R. Jhala, S. Lerner, and H. Shacham. An empirical study of privacy-violating information flows in JavaScript web applications. In *Proceedings of the 17th ACM conference on Computer and communications security*, CCS '10. ACM, 2010.
- [21] M. Jaskelioff and A. Russo. Secure multi-execution in Haskell. In *Proc. Andrei Ershov International Conference on Perspectives of System Informatics*, LNCS. Springer-Verlag, June 2011.
- [22] J. Ligatti, L. Bauer, and D. Walker. Edit Automata: Enforcement Mechanisms for Run-time Security Policies. *International Journal of Information Security*, Feb. 2005.
- [23] M. T. Louw, K. T. Ganesh, and V. N. Venkatakrishnan. AdJail: practical enforcement of confidentiality and integrity policies on web advertisements. In *Proceedings of the 19th USENIX conference on Security*, USENIX Security'10, Berkeley, CA, USA, 2010. USENIX Association.
- [24] J. Magazinius, A. Askarov, and A. Sabelfeld. A lattice-based approach to mashup security. In *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security*, ASIACCS '10. ACM, 2010.
- [25] J. McLean. A General Theory of Composition for Trace Sets Closed under Selective Interleaving Functions. In *Proceedings of the 1994 IEEE Symposium on Security and Privacy*, SP '94. IEEE Computer Society, 1994.
- [26] A. C. Myers, L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. Jif: Java Information Flow. Software release. Located at <http://www.cs.cornell.edu/jif>, July 2001.
- [27] W. Rafnsson and A. Sabelfeld. Limiting information leakage in event-based communication. In *Proceedings of the ACM SIGPLAN 6th Workshop on Programming Languages and Analysis for Security*, PLAS '11. ACM, 2011. ISBN 978-1-4503-0830-4.
- [28] W. Rafnsson and A. Sabelfeld. Secure Multi-execution: Fine-Grained, Declassification-Aware, and Transparent. In *Computer Security Foundations Symposium (CSF), 2013 IEEE 26th*, pages 33–48, June 2013.

- [29] A. Russo and A. Sabelfeld. Securing Timeout Instructions in Web Applications. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society, July 2009.
- [30] A. Sabelfeld and A. C. Myers. Language-Based Information-Flow Security. *IEEE J. Selected Areas in Communications*, 21(1):5–19, Jan. 2003.
- [31] A. Sabelfeld and A. Russo. From dynamic to static and back: Riding the roller coaster of information-flow control research. In *Proc. Andrei Ershov International Conference on Perspectives of System Informatics*, volume 0 of *LNCS*. Springer-Verlag, June 2009.
- [32] F. B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*, 2000.
- [33] D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières. Flexible Dynamic Information Flow Control in Haskell. In *Haskell Symposium*. ACM SIGPLAN, September 2011.
- [34] W. Swierstra and T. Altenkirch. Beauty in the Beast: A Functional Semantics of the Awkward Squad. In *Haskell '07: Proceedings of the ACM SIGPLAN Workshop on Haskell*, pages 25–36, 2007.
- [35] P. Vogt, F. Nentwich, N. Jovanovic, E. Kirda, C. Kruegel, and G. Vigna. Cross-Site Scripting Prevention with Dynamic Data Tainting and Static Analysis. In *Proc. Network and Distributed System Security Symposium*, Feb. 2007.
- [36] D. Volpano, G. Smith, and C. Irvine. A Sound Type System for Secure Flow Analysis. *J. Computer Security*, 4(3):167–187, 1996.
- [37] Williams and D. Wichers. OWASP Top 10 2010. <http://www.owasp.org/index.php/Top\10\2010>, 2010.
- [38] D. Zanarini and M. Jaskelioff. Monitoring Reactive Systems with Dynamic Channels: Extended Version, 2014. Available at <http://www.fceia.unr.edu.ar/~dante>.
- [39] D. Zanarini, M. Jaskelioff, and A. Russo. Precise Enforcement of Confidentiality for Reactive Systems. In *Computer Security Foundations Symposium (CSF), 2013 IEEE 26th*, pages 18–32, June 2013.
- [40] L. Zheng and A. C. Myers. Dynamic security labels and noninterference (extended abstract). In T. Dimitrakos and F. Martinelli, editors, *Formal Aspects in Security and Trust*, volume 173 of *IFIP International Federation for Information Processing*, pages 27–40. Springer US, 2005. ISBN 978-0-387-24050-3.