

Precise Enforcement of Confidentiality for Reactive Systems

Dante Zanarini, Mauro Jaskelioff
CIFASIS – CONICET
Universidad Nacional de Rosario
{zanarini, jaskelioff}@cifasis-conicet.gov.ar

Alejandro Russo
Chalmers University of Technology
russo@chalmers.se

Abstract—In the past years, researchers have been focusing on applying information flow security to web applications. These mechanisms should raise a minimum of false alarms in order to be applicable to millions of existing web pages. A promising technique to achieve this is secure multi-execution (SME). If a program is already secure, its secure multi-execution produces the same output events; otherwise, this correspondence is intentionally broken in order to preserve security. Thus, there is no way to know if unexpected results are due to bugs or due to semantics changes produced by SME. Moreover, SME provides no guarantees on the relative ordering of output events from different security levels. We argue that these shortcomings limit the applicability of SME.

In this article, we propose a scheduler for secure multi-execution which makes it possible to preserve the order of output events. Using this scheduler, we introduce a novel combination between monitoring and SME, called *multi-execution monitor*, which raises alarms only for actions breaking the non-interference notion of ID-security for reactive systems. Additionally, we show that the monitor guarantees transparency even for CP-similarity, a progress-sensitive notion of observation.

I. INTRODUCTION

In recent years there has been an increasing interest in applying information flow control as a general mechanism to preserve confidentiality on web applications (e.g. [11, 19, 22, 27, 46]). The adoption of this technology promises to reduce the need for ad-hoc and purpose-specific counter-measures (e.g. architectures to contain advertisement scripts [26], browser extensions to control cache-based leaks [20], etc.) In fact, several of the OWASP top-ten web vulnerabilities [49] can be rephrased in terms of information-flow problems.

In a web scenario, where millions of web pages have been written and deployed, it is important to provide permissive information-flow mechanisms, i.e. mechanisms that raise as few false alarms as possible. Traditional Denning-style [16, 35, 48] information-flow enforcements perform over-approximations that could lead to reject secure programs. Driven by permissiveness and dynamic features of web scripting languages, researchers tend to adopt dynamic techniques in the form of execution monitors (e.g. [2, 4, 5, 37]). Despite efforts to push the limits of dynamic information-flow, execution monitors are still not capable of enforcing sound and precise information-flow policies [29,

39], and must therefore reject possibly secure and useful web pages.

Recently, Devriese and Piessens [17] devised an alternative dynamic approach, called secure multi-execution (SME), based on the idea of executing the same program several times, once for each security level. As opposed to previous enforcement mechanisms, this novel technique works as a black-box approach; it only requires applying specific actions when inputs and outputs (I/O) are produced. Secure multi-execution does not require either static analysis or execution monitoring. As claimed by their authors, secure multi-execution enforces a specific version of non-interference with significantly better precision than traditional static and/or other dynamic techniques. More precisely, they prove that termination-sensitive non-interferent programs, which terminate under normal execution, match the behaviour produced under secure multi-execution. In contrast, if the program is leaking information, secure multi-execution will change its semantics in order to enforce security. It is claimed that this approach is the first one to achieve both soundness and precision.

Although promising, secure multi-execution suffers from some drawbacks which may limit the applicability of this technique. More concretely, we identify the following weaknesses.

- *Precision* As described above, it is postulated that secure-multi execution is *precise*, and thus more permissive than static analysis and execution monitoring [7, 9, 17, 24]. We argue that such comparison might be somewhat unfair. While static analysis and execution monitors are capable of accepting or rejecting programs, secure multi-execution just runs them. Trading permissiveness by uncertainty, secure multi-execution makes it not possible to distinguish when the semantics of a program has been altered in order to preserve security. As a result, users might experience that their programs do not behave as expected without knowing if it is due to software errors or due to security reasons. We believe that a better terminology to refer to the precision of secure multi-execution is *transparency* [25], i.e. the ability of an enforcement mechanism to preserve the semantics of executions which already obey the security policy in question.

► *Order of events* Secure multi-execution is claimed to be transparent for terminating runs of termination-sensitive non-interferent programs. Under scrutiny, this claim only holds if the interleaving of events from different security levels is not relevant for the computation. In fact, output events might be arbitrarily interleaved when coming from different security levels. Not preserving the order of events might be problematic, for instance, in web pages with complex DOM-elements reacting to the same event (e.g. mouse click.)¹

► *Schedulers* In previous works on SME [7, 17], the soundness and transparency arguments are given for a particular scheduler, called `selectlowprio`. This scheduler prioritises the execution of the copy of the program associated to the lowest security level. Due to this choice, secure multi-execution rules out leaks through the *external timing covert channel*, i.e. revealing confidential data by precisely measuring the time in which external (and observable) events are triggered. However, one major drawback of this scheduler is that it requires a total extension of the lattice order. In web scenarios, however, web domains are often modelled as incomparable security levels [27], which prohibits the use of such scheduler. Authors in [24] discuss alternative scheduling strategies for arbitrary lattices which guarantee different security policies for different security levels.

The main contribution of this paper is *multi-execution monitoring*, a novel combination of monitoring and secure multi-execution. This technique respects the interleaving of events from different security levels, and thus provides better transparency results than secure multi-execution. More importantly, multi-execution monitoring allows to detect insecure commands with precision, i.e. the monitor only raises an alarm when a command breaks the non-interference notion of ID-security [11] for reactive systems.

Intuitively, the idea of multi-execution monitoring is simply to monitor a program by comparing it with its secure multi-execution. A multi-execution monitor runs a program simultaneously with its secure multi-execution version. The two programs will be in sync (perform exactly the same I/O operations) for as long as the execution is secure. If one version tries to do something different from the other, then the monitor reports that the program is insecure.

The most important contributions of this paper can be summarised as follows.

► Inspired by the coalgebraic theory of systems [21], we propose a novel semantics for programs based on *interaction trees*. This formulation treats programs as black-boxes, about which nothing is known except what is inferred from their I/O interactions with the environment. In this manner, we gain modularity since new program-

ming features related to internal operations do not affect our formal results.

► We define a scheduler that significantly improves the transparency guarantees for secure multi-execution. With this scheduler, we can guarantee that secure multi-execution preserves the order of events and progress of secure programs.

► We introduce a multi-execution monitor which can precisely detect when commands violate ID-security. This feature, not only allows us to notify users when programs are malicious, but also enables the debugging of insecure programs. Multi-execution monitoring gives good transparency guarantees for the non-interference notions of ID- and CP-security [11]. In fact, these transparency guarantees make it possible to report leaks as traces of the original program.

The paper is organised as follows. In Section II we introduce reactive interaction trees, our model of reactive systems, show how JavaScript-like programs might be interpreted on it, and define the notion of ID-secure programs. Secure multi-execution is presented in Section III. Section IV presents a scheduler for secure multi-execution which preserves the order of events and provides better transparency guarantees. In Section V, we define the security condition on executions that the multi-execution monitor will enforce. Multi-execution monitoring is introduced in Section VI. In Section VII we discuss related work. Finally, in Section VIII we conclude and discuss future work.

Complete proofs for all the results can be found in an online extended version [50].

II. REACTIVE SYSTEMS AND NON-INTERFERENCE

We model reactive programs as *interaction trees*, i.e. data structures in the form of trees describing every possible interaction with the environment. We assume a set of channels $Chan$, input values in a set I , and output values in a set O . An *event* is a piece of data paired with the name of the communication channel associated to it. Let $E_A = Ch \times A$ denote the set of events of type A . Interaction trees are then defined as the following inductive datatype.

Definition 1 (Reactive Interaction Trees).

$$\begin{aligned} React = & \text{Read } (E_I \rightarrow React) \\ & | \text{Write } (E_O \times React) \\ & | \text{Step } React \end{aligned}$$

Intuitively, constructor *Read* denotes a program that receives an input from a channel determined by the environment (E_I) and, based on that, decides how to continue (*React*). Constructor *Write* represents programs which write an output in a chosen channel (E_O) and continue with another computation (*React*). Finally, constructor *Step* corresponds to a silent step, that is, a computation which

¹The W3C consortium specifies the expected behaviour for such scenarios. <http://www.w3.org/TR/DOM-Level-2-Events/events.html>

does not affect the environment. Silent steps allows us to model divergence. We do not model termination (*React* trees are necessarily infinite) as reactive systems are usually meant to run forever. However, we could easily model termination by adding a new constructor *Stop*.

Interaction trees have no notion of state, and therefore are more abstract than concrete labelled transition systems and make the semantics and proof machinery simpler.

The idea of modelling the interactions of reactive programs through an infinite tree comes from a coalgebraic view of systems [21]. An interaction tree is the carrier of the final coalgebra of a functor, implementations of concrete systems are given by coalgebras for this functor, and the interpretation of a program into *React* arises from the universal property of the final coalgebra.

A. Semantics

An interaction tree is essentially a static description of the possible inputs and outputs that might occur during execution. To know exactly which interactions occur for a given run, we need to provide an evaluation relation. We start by defining possibly infinite sequences.²

Definition 2 (Colists). Let A be a set. Consider the type of possibly infinite sequences of A to be coinductively defined as follows.

$$\text{Colist}_A = \llbracket \mid A :: \text{Colist}_A$$

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

The structure of input and output events of the system is given by colists.

Definition 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.

$$\mathcal{I} = \text{Colist}_{E_I} \quad \mathcal{O} = \text{Colist}_{E_o \cup \{\bullet\}}$$

The elements of an output colist are either an event E_o , or an invisible output \bullet .

The evaluation relation $\Rightarrow \subseteq (\text{React} \times \mathcal{I}) \times \mathcal{O}$ is coinductively defined by the rules in Figure 1, where we write $(t, i) \Rightarrow o$ for $(t, i, o) \in \Rightarrow$. Intuitively, feeding an interaction tree t with a colist of input events $i \in \mathcal{I}$ yields the output colist $o \in \mathcal{O}$ iff $(t, i) \Rightarrow o$.

Rule (R_1) produces no outputs (\llbracket) when no input events are present. Rule (R_2) consumes the first available input event (e), and based on that, produces the output o ($(f(e), i) \Rightarrow o$). Rule (W) outputs an event e ($e :: o$), and then the outputs triggered by program t ($(t, i) \Rightarrow o$). Rule (S) simply outputs \bullet when a silent computation step is performed.

²We distinguish lists (finite sequences), colists (possibly infinite sequences), and streams (infinite sequences). However, we overload the notation for constructors.

$$\begin{array}{l} (R_1) \frac{}{(Read\ f, \llbracket) \Rightarrow \llbracket} \\ (R_2) \frac{(f(e), i) \Rightarrow o}{(Read\ f, e :: i) \Rightarrow \bullet :: o} \\ (W) \frac{(t, i) \Rightarrow o}{(Write(e, t), i) \Rightarrow e :: o} \\ (S) \frac{(t, i) \Rightarrow o}{(Step\ t, i) \Rightarrow \bullet :: o} \end{array}$$

Figure 1. Evaluation relation for interaction trees

Interaction trees are not concerned with the features of the programming language used to code a reactive system. This level of abstraction allows us to apply our technique and results to, for instance, different imperative or functional languages. For each language, it is enough to describe how interaction trees are generated from programs. To illustrate this point, we briefly describe interaction trees for reactive JavaScript-like programs.

B. Interaction trees for a JavaScript-like language

Despite its simplicity, *React* is able to model the interactions of complex languages. As an example of that, Figure 2 presents a language inspired by [32]. This language is a subset of JavaScript and describes many of its features. Expressions are side-effect free and denote strings, numbers, and boolean values. Event handlers can change the state of the system as well as define new ones. Input and output channels are disjoint since we focus on how programs react with the environment rather than themselves. Programs (p) are defined as a sequence of event handlers. Event handlers (h) indicate which commands (c) to execute when an input arrives to a channel ($ch(x) \{c\}$). Most of the commands are self-explanatory. However, some of them require further explanation. Command **out** (ch, e) outputs the value denoted by e into channel ch . Command **new** h declares a new event handler (or replaces an existing one). Command **eval** (e) dynamically evaluates the instructions denoted by a string expression e .

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

Figure 2. A JavaScript-like language. Symbols ch and e range over channels and expressions, respectively.

Throughout the examples of this article, we will make the following assumptions: events have integer values; there are two input channels $L?$, $H?$ and two output channels $L!$, $H!$; events received on $L?$ and events sent to $L!$

are considered public events; events received on $H?$ and events sent to $H!$ are considered private or secret events. Programs in this language may be seen as a loop which reads an event and executes the handler associated to it. Such a handler may produce some outputs, end silently or diverge. We can interpret every JavaScript-like program using interaction trees. More specifically, there exists a function $\llbracket - \rrbracket : Prg \rightarrow M \rightarrow React$, that given a program $p \in Prg$ and a memory $\mu \in M$, it gives us the resulting interaction tree $\llbracket p \rrbracket(\mu)$ (see Appendix A for details). This tree denotes the interactions that may happen when program p is run under memory μ . Here, M is the set of memories, i.e. mappings of variables to values.

The programs of this language have a special structure: no input event may be handled *inside* a handler. However, this structure plays no role once we abstract away programs by interpreting them into interaction trees. The same observation can be made about many features of this language, such as assignments, conditionals, loops, and dynamic code evaluation, as well as features not present in it, e.g., objects and DOM-trees.

Example 4. Consider the following JavaScript-like program.

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

$$L?(x) \{ \mathbf{if} \ r = 0$$

$$\quad \{ \mathbf{out} \ (L!, 0) \}$$

$$\quad \{ \mathbf{while} \ 1 \ \mathbf{do} \ \mathbf{skip} \};$$

The program diverges when the secret value stored in r is different from zero. The interpretation of p for a memory μ is the following interaction tree.

$$\llbracket p \rrbracket(\mu) = Read \ (\lambda(ch, v). \mathbf{case} \ ch \ \mathbf{of}$$

$$\quad H? \rightarrow Step \ (\llbracket p \rrbracket(\mu[r \mapsto v]))$$

$$\quad L? \rightarrow Step \ (\mathbf{if} \ \mu(r) = 0$$

$$\quad \quad \mathbf{then} \ Write \ ((L!, 0), \llbracket p \rrbracket(\mu))$$

$$\quad \quad \mathbf{else} \ \mathit{diverge})$$

$$\quad _ \rightarrow Step \ (\llbracket p \rrbracket(\mu)))$$

We write $f[x \mapsto y]$ for the function that behaves like f , except on x where it maps to y . We denote with *diverge* the infinite sequence of *Steps* (i.e. $\mathit{diverge} = Step \ \mathit{diverge}$.)

C. Reactive Non-Interference

We organise security levels in a lattice $(\mathcal{L}, \sqsubseteq)$, with the intention to express that data at level ℓ_1 can securely flow into data at level ℓ_2 when $\ell_1 \sqsubseteq \ell_2$, and we associate a security level to each channel. The security level of an event e (noted $lvl(e)$) is determined by the security level of its channel. We define the predicate $visible_\ell$ on events that determines when an event is observable for an observer at level ℓ .

$$\frac{lvl(e) \sqsubseteq \ell}{visible_\ell(e)}$$

Output colists may also have a silent event \bullet . Hence, for output colists we extend the predicate so that silent events are not visible at any security level ($\forall \ell. \neg visible_\ell(\bullet)$).

Definition 5 (ID-similarity). ID-similarity between colists for an observer at level ℓ is formalised coinductively by the following rules.

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

$$\frac{\neg visible_\ell(e) \quad s \sim_\ell^{ID} s'}{s \sim_\ell^{ID} e :: s'} \quad \frac{visible_\ell(e) \quad s \sim_\ell^{ID} s'}{e :: s \sim_\ell^{ID} e :: s'}$$

ID-similarity is reflexive and symmetric, but it is not transitive. In fact, the absent of transitivity turns this relationship meaningful for divergent computations. Note that two colists will be considered ID-similar if there is hope that they may produce the same observable events. In particular, a colist s will be similar to every other colist if it consist of nothing but infinitely many events not visible at level ℓ . Hence, transitivity would imply that every pair of colists is similar, but this is clearly not the case.

We define the non-interference notion of ID-security for reactive programs [9, 11] as follows.

Definition 6 (ID-security). An interaction tree t is ID-secure iff, for every level ℓ and input colists i, i' such that $(t, i) \Rightarrow o$ and $(t, i') \Rightarrow o', i \sim_\ell^{ID} i'$ implies $o \sim_\ell^{ID} o'$.

III. SECURE MULTI-EXECUTION

Secure multi-execution runs a program multiple times, once for each security level, but giving I/O operations a level-dependent semantics. Outputs to a channel at security level ℓ are only performed in the execution corresponding to level ℓ . Inputs from a channel at security level ℓ are only available to executions corresponding to a level ℓ_e , where $\ell \sqsubseteq \ell_e$. It is clear that the secure multi-execution of programs does not leak information. The execution at level ℓ produces outputs only at that level, while consuming data with less or equal confidentiality than ℓ .

We will define secure multi-execution directly on interaction trees as a transformation that takes an interaction tree and returns another one representing its secure multi-execution. The returned interaction tree is guaranteed to be non-interferent, i.e. it does not leak secrets (see Definition 6). The transformation proceeds in two steps:

- The original interaction tree *React* gives rise to a receiver interaction tree *Receive* and to one producer interaction tree *Produce* for each security level. These new types of interaction trees are generated by functions rcv and prd_ℓ , respectively (explained below). Following the modelling of reactive systems [9–11], the interaction tree *Receive* captures actions related to obtaining input data from the environment, while *Produce* captures actions related to producing outputs.

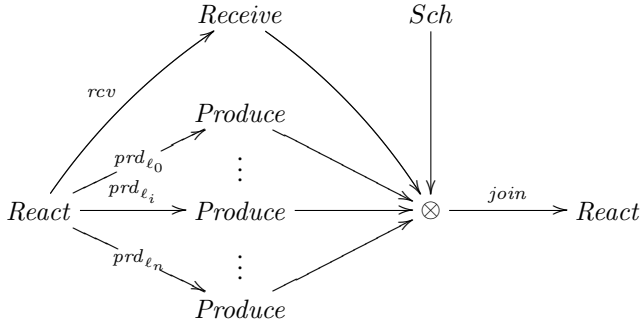


Figure 3. Transformation providing secure-multi execution

► The *Receive* and *Produce* interaction trees are re-interpreted back into an interaction tree *React* with the help of a scheduler *Sch*. This is done by the function *join*.

Figure 3 illustrates the transformation process.

A. Receiver and Producers

Interaction trees *Receive* are coinductively defined as follows.

$$\begin{aligned} \text{Receive} &= \text{RnQ} (E_I \rightarrow \text{Receive}) \\ &| \text{Step Receive} \end{aligned}$$

The transformation from the original interaction tree to a receiver simply transforms writes into silent steps, and reads into “read and queue” (constructor *RnQ*).

$$\begin{aligned} rcv &: \text{React} \rightarrow \text{Receive} \\ rcv (\text{Read } f) &= \text{RnQ} (\lambda e. rcv (f(e))) \\ rcv (\text{Write } (e, t)) &= \text{Step} (rcv(t)) \\ rcv (\text{Step } t) &= \text{Step} (rcv(t)) \end{aligned}$$

When an input is obtained at security level ℓ , secure multi-execution dictates that it should be observable for executions at level ℓ' such that $\ell \sqsubseteq \ell'$. In order to distribute data to the appropriate executions, we use queues of events. Constructor *RnQ* is then denoting the fact that every input data is obtained from the environment and placed in the appropriate queues (see function *join*).

Interaction trees *Produce* are coinductively defined as follows.

$$\begin{aligned} \text{Produce} &= \text{Reuse} (E_I \rightarrow \text{Produce}) \\ &| \text{Write} (E_O \times \text{Produce}) \\ &| \text{Step Produce} \end{aligned}$$

The producer function is parameterised by a security level (prd_ℓ). Reads are replaced by fetching data obtained by the receiver (constructor *Reuse*). How data is reused is explained in detail in the function *join*. Writes are only performed if the security level of the channel coincides with

the security level of the producer.

$$\begin{aligned} prd_\ell &: \text{React} \rightarrow \text{Produce} \\ prd_\ell (\text{Read } f) &= \text{Reuse} (\lambda e. prd_\ell (f(e))) \\ prd_\ell (\text{Write } (e, t)) &= \text{if } \ell = lvl(e) \\ &\quad \text{then Write } (e, prd_\ell(t)) \\ &\quad \text{else Step } (prd_\ell(t)) \\ prd_\ell (\text{Step } t) &= \text{Step } (prd_\ell(t)) \end{aligned}$$

Interaction trees of the kind *Produce* are isomorphic to *React*. Moreover, *Receive* can be embedded into *React*. In this light, functions *rcv* and prd_ℓ could be defined to simply produce *React* interaction trees instead. However, we have chosen to introduce new types of interaction trees in order to make the presentation more intuitive.

B. Obtaining ID-secure interaction trees

Once we have obtained the receiver and the producers, we proceed to join them into a single interaction tree *React*. We do this by choosing commands from different trees, as dictated by a scheduler. Schedulers are modelled as streams of elements in $\mathcal{L} \cup \{\star\}$, where the symbol \star accounts for the receiver, while each $\ell \in \mathcal{L}$ accounts for the producer at level ℓ . We denote with *Sch* the set of all schedulers.

In order to receive data from the receiver, we equip each producer with a queue of input events. We model this *system of queues* by a function from levels to queues of input events, i.e., $Q = \mathcal{L} \rightarrow \text{Queue} (E_I)$. As dictated by secure multi-execution, input events are distributed to producers only when the security level of the producer is equal or higher than the security level of the event. We define the operator \oplus responsible for the proper distribution of input events to producers via queues. More precisely, given $q \in Q$ and an event $e \in E_I$, we define

$$\begin{aligned} q \oplus e &= \lambda \ell. \text{if } lvl(e) \sqsubseteq \ell \\ &\quad \text{then enqueue } (q(\ell), e) \\ &\quad \text{else } q(\ell) \end{aligned}$$

where the function *enqueue* just adds an event to a queue. Observe that function \oplus modifies a system of queues q by appending a new event e to each queue corresponding to a level ℓ such that $lvl(e) \sqsubseteq \ell$.

Function *join*, the most interesting function of the transformation, takes a scheduler, a receiver, a producer for each security level, and a system of queues. As a result, it calculates an ID-secure interaction tree. More precisely, the type signature of *join* is as follows.

$$\text{join} : \text{Sch} \times \text{Receive} \times (\mathcal{L} \rightarrow \text{Produce}) \times Q \rightarrow \text{React}$$

The function is defined by pattern-matching on the scheduler. If the scheduler dictates that it is the turn of the receiver (\star), *join* reproduces each *Step* in the receiver’s interaction tree until it finds a *RnQ*. At this point *join* should perform a *Read*. However, in order to avoid leaking secrets through the termination channel, it must first check that all producers

have consumed their input queues and that they are waiting for more input. This situation is verified by the predicate $sync(p, q)$, which holds iff the next interaction in every producer is *Reuse*, and $q(\ell)$ is empty for all ℓ . Hence, when the producers are synched, a *Read* can be performed and the input event added to the corresponding queues. If producers are not synched (for example, because there is a producer whose first action is to do a *Step*) then *join* should execute producers in order to try to synchronise them. This execution of producers is performed by the function *next*.

$$\begin{aligned} join \ (\star :: s, r, p, q) &= \text{case } r \text{ of} \\ Step \ r' &\rightarrow Step \ (join \ (\star :: s, r', p, q)) \\ RnQ \ f &\rightarrow \text{if } sync(p, q) \text{ then} \\ &\quad Read \ (\lambda e. join \ (s, f(e), p, q \oplus e)) \\ &\quad \text{else let } (p', q') = next(p, q) \\ &\quad \text{in } Step \ (join \ (\star :: s, r, p', q')) \end{aligned}$$

Here, $next(p, q) = (\lambda \ell. (step(p, q, \ell))_1, \lambda \ell. (step(p, q, \ell))_2)$ is the function that tries to make a single step at every security level ℓ using the function *step*. The function *step*, in turn, tries to make a single step at a single level. The subindices in the definition of *next* denote pair projections. The function *step* is defined as follows.

$$step(p, q, \ell) = \begin{cases} (p_\ell, q(\ell)) & \text{if } p(\ell) = Step \ p_\ell \\ (f(e), q_\ell) & \text{if } p(\ell) = Reuse \ f \\ & \quad \wedge q(\ell) = e :: q_\ell \\ (p(\ell), q(\ell)) & \text{otherwise} \end{cases}$$

This function simply makes a computation step on a producer unless there is a *Write* or a *Reuse* with an empty queue.

Continuing with the definition of *join*, if the scheduler dictates that it is the turn of the producer at level ℓ , *join* will inspect the producer tree corresponding to level ℓ and execute it until it finds a write. If that producer would perform a write, a *Write* is added to the resulting tree. If the producer tries to reuse an event when there is none, it just yields the execution; if, on the other hand, there is an event, it gets consumed. If the producer makes a *Step*, *join* will replicate it.

$$\begin{aligned} join \ (\ell :: s, r, p, q) &= \text{case } p(\ell) \text{ of} \\ Write \ (o, p_\ell) &\rightarrow Write \ (o, join \ (s, r, p[\ell \mapsto p_\ell], q)) \\ Step \ p_\ell &\rightarrow Step \ (join \ (\ell :: s, r, p[\ell \mapsto p_\ell], q)) \\ Reuse \ f &\rightarrow \text{case } q(\ell) \text{ of} \\ e :: es &\rightarrow Step \ (join \ (\ell :: s, r, p[\ell \mapsto f(e)], q[\ell \mapsto es])) \\ [] &\rightarrow Step \ (join \ (\ell :: s, r, p, q)) \end{aligned}$$

Given a scheduler, secure multi-execution for interaction trees is defined by the following function.

$$\begin{aligned} sme &: Sch \times React \rightarrow React \\ sme \ (s, t) &= join \ (s, rcv(t), \lambda \ell. prd_\ell(t), \lambda \ell. []) \end{aligned}$$

The next proposition states the security of *sme*.

Proposition 7. *Let $t \in React$ and $s \in Sch$. Then, the*

interaction tree $sme \ (s, t)$ is ID-secure.

Similarly to [7, 9, 17, 24], we can prove that the transformation preserves the semantics of ID-secure programs when the interleaving of events at different security levels is not relevant (see the extended version [50] for the precise statement and proof.)

The security and transparency propositions [50] quantify over all schedulers. This might seem surprising, as clearly one can choose a bad scheduler. For instance, we could choose the scheduler that always chooses the receiver. Such a scheduler would never issue a *Write*, and therefore would always diverge. As discussed in Section II, a silent infinite colist is ID-similar to every other one. In particular, the output colist $\perp = \bullet :: \perp$ is infinite and silent at every level, and therefore ID-similar to every output colist. Therefore, the trivial transformation $bad(t) = diverge$ satisfies the security and transparency propositions. After all, *diverge* is ID-secure and produces an output (\perp) ID-similar to any other output. We believe that secure multi-execution can do better than the trivial diverging transformation, but in order to show it we need to state better formal guarantees.

In the next section, we present one of the main contributions of this paper. We show how the program under execution and its input induce a scheduler that significantly improves the transparency guarantees of secure multi-execution.

IV. ORDER-PRESERVING SCHEDULER

The definition of secure multi-execution in the previous section is parameterised by a scheduler. However, if we are interested in the order in which events from different levels are produced, the choice of scheduler is of paramount importance. The standard precision result for secure multi-execution [7, 9, 17, 24] ensures that the order of output events is preserved only when looking at a given security level in isolation. However, in certain scenarios (such as the monitor in Section VI) one needs to take into account the interleaving of events from different security levels. Therefore, a stronger guarantee is required. Example 8, although very simple, illustrates this point.

Example 8. We define program p with just one handler as follows:

$$\begin{aligned} p = \mathbb{L}?(x) \{ &\text{out} \ (\mathbb{L}!, x); \\ &\text{if } x > 10 \ \{ \text{out} \ (\mathbb{H}!, x) \} \\ &\quad \{ \text{skip} \}; \\ &\text{out} \ (\mathbb{L}!, 1) \}; \end{aligned}$$

This program is non-interferent. Since there is no handler for channel $\mathbb{H}?$, every secret input is ignored.

It is easy to see that, for all inputs, every event on channel $\mathbb{H}!$ is preceded by an event on channel $\mathbb{L}!$ with exactly the same value. However, the secure multi-execution of p with

a $\text{select}_{\text{lowprio}}$ scheduler has a different behaviour for a high observer: every event on channel $H!$ is preceded by an event with constant 1 on $L!$.

In order to preserve the order of events we will look at the order of events generated by the original program. That is, we will use the execution of the original program to guide secure multi-execution. If the original execution issues a *Read* command, the scheduler chooses the receiver, identified as \star , to run. Observe that this is the only interaction tree under secure multi-execution capable of issuing such command. Instead, if the original execution issues a *Write* command to a channel at level ℓ , then the producer $p(\ell)$ is run. Observe that $p(\ell)$ is the only interaction tree under secure multi-execution that is able to perform *Writes* into channels at level ℓ . However, if the original execution issues a *Step*, there is no information from which to decide what to schedule next. To account for this situation, we extend our definition of schedulers (*Sch*) in Section III-B with the element \circ . Finally, if read commands are issued by the execution of the original program under an empty colist of input events, it does not really matter which program under secure multi-execution gets scheduled. After all, the execution of the original program has stopped (see rule (R_1) in Figure 1). More precisely, the order-preserving scheduler, called *ops*, is defined as follows.

$$\begin{aligned} ops & : \text{React} \times \mathcal{I} \rightarrow \text{Sch} \\ ops (\text{Read } f, e :: i) & = \star :: (ops (f(e), i)) \\ ops (\text{Read } f, []) & = \circ :: (ops (\text{Read } f, [])) \\ ops (\text{Write } (e, t), i) & = lvl(e) :: ops (t, i) \\ ops (\text{Step } t, i) & = \circ :: ops (t, i) \end{aligned}$$

The scheduler takes the interaction tree of the program to be executed under secure multi-execution (*React*), the colist of inputs (\mathcal{I}), and returns the scheduling policy (*Sch*). Observe how reads in the presence of inputs are mapped to the receiver (\star), while writes are mapped into producers with the same security level as the channel ($lvl(e)$).

As a consequence of adding symbol \circ to the scheduler, we need to extend the definition of *join* with the following additional case.

$$join (\circ :: s, p, r, q) = Step (join (s, p, r, q))$$

When *join* finds the symbol \circ in the scheduling policy, it simply makes a *Step*.

A. Transparency guarantees for the scheduler *ops*

The order-preserving scheduler *ops* allows secure multi-execution to provide better transparency guarantees than the ones previously shown.

Theorem 9 (Transparency for ID-secure trees). *Let t be an ID-secure interaction tree, and i an input colist such that $(t, i) \Rightarrow o$. If $(sme(ops(t, i), t), i) \Rightarrow o'$ then $\forall \ell. o \sim_{\ell}^{ID} o'$.*

The theorem above states that the output of the original program and its secure multi-execution are ID-similar for ID-secure interaction trees. This means that for any observer at some level ℓ , the order of ℓ -visible events is preserved. This is an improvement over previous results since it considers the interleaving of events. Nevertheless, satisfying transparency for ID-secure programs does not guarantee that secure multi-execution performs any progress, e.g., *sme* might always diverge (see discussion at the end of Section III.) In order to guarantee progress, we consider a stronger notion of non-interference for reactive systems called CP-security [11].

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

$$\frac{\neg visible_{\ell}(e) \quad silent_{\ell}(s)}{silent_{\ell}(e :: s)} \quad \frac{}{silent_{\ell}([])}$$

We define a relation that identifies the next event that is visible to an observer at level ℓ (if exists). Intuitively, we say that $s \triangleright_{\ell} e :: s'$ when e is the next event in s visible at level ℓ . The following rules inductively define the relation \triangleright_{ℓ} .

$$\frac{visible_{\ell}(e)}{e :: s \triangleright_{\ell} e :: s} \quad \frac{\neg visible_{\ell}(e) \quad s \triangleright_{\ell} e' :: s'}{e :: s \triangleright_{\ell} e' :: s'}$$

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

Definition 10 (CP-similarity). CP-similarity between colists is defined coinductively by the following rules.

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

As opposed to ID-similarity, CP-similarity is an equivalence relation. Moreover, CP-similarity guarantees progress by asking one colist to be as productive as the other, i.e. the two colists either produce the same visible event in a finite number of steps, or both become silent. We now define when a program is CP-secure.

Definition 11 (CP-security). An interaction tree t is CP-secure iff, for every level ℓ and input colists i, i' such that $(t, i) \Rightarrow o$ and $(t, i') \Rightarrow o', i \sim_{\ell}^{CP} i'$ implies $o \sim_{\ell}^{CP} o'$.

CP-security is strictly stronger than ID-security: any CP-secure program is ID-secure [11]. We show that secure multi-execution is also transparent with respect to CP-secure programs when using the order preserving scheduler.

Theorem 12 (Transparency for CP-secure trees). *Let t be a CP-secure interaction tree, and i an input colist such that $(t, i) \Rightarrow o$. If $(sme(ops(t, i), t), i) \Rightarrow o'$ then $\forall \ell. o \sim_{\ell}^{CP} o'$.*

The transparency theorem for CP-secure trees is a significant improvement over previous results for secure multi-execution in reactive systems. Previous results did not show that secure multi-executions approaches fulfilling the security and transparency properties were any better than the transformation that always produces diverging runs. The above theorem, however, is able to guarantee progress for CP-secure programs as well as event-order preservation. Hence, if a CP-secure program produces a visible event, its secure multi-execution is forced to produce it too.

V. SECURE INPUTS

We want to precisely detect leaks of secret information in reactive systems. Non-interference, a property of programs, cannot be precisely enforced by an execution monitor [29, 39]. More importantly, it may not be a desirable property to enforce. For instance, many web applications deployed on the web might be harmless most of the time, but leak information only in certain situations. That is, they may leak information in certain runs, but not on others. In this light, it is not surprising that some information-flow monitors accept runs of interferent programs as long as they do not leak information. For instance, monitors in [2, 5, 19, 34, 37, 42] accept the runs of program `if public == 42 then public := secret else skip` when the public input is different from 42. With this in mind, we define a security condition on runs (rather than on programs), by characterising the inputs for which programs do not leak secret information (we ignore leaks due to covert channels.) In order to define this notion, we need to present an auxiliary relation.

We coinductively define the relation \blacktriangleright_ℓ responsible for removing all the events unobservable at level ℓ .

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

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 ℓ .

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

Definition 13 ($\approx_{\ell,t}$). 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 that the program t reveals the same ℓ -secrets when given the inputs i, i' , noted $i \approx_{\ell,t} i'$, iff $o \sim_\ell o'$.

Similarly to [9], 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 14 (Secure input). Let t be an interaction tree. An input colist i is *secure for* t iff $\forall \ell. i \approx_{\ell,t} i_{\blacktriangleright_\ell}$. We say

$$p = \text{H?}(x) \{ r := x \};$$

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

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

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

Figure 4. ID-insecure program with ID-secure inputs

that the input i is ID-secure (CP-secure) for t when \sim_ℓ is instantiated to \sim_ℓ^{ID} (\sim_ℓ^{CP}) in Definition 13.

It is desirable to establish a connection between programs and their secure inputs, so as to be able to transfer security properties from one notion to the other. Fortunately, there is a close relationship between CP-secure programs and CP-secure inputs.

Lemma 15 (Secure inputs and CP-security). *A reactive interaction tree $t \in \text{React}$ is CP-secure iff $\forall i \in \mathcal{I}. i$ is CP-secure for t .*

When it comes to ID-security, the relationship between secure programs and secure inputs is not that strong.

Lemma 16 (Secure inputs and ID-security). *If a reactive interaction tree $t \in \text{React}$ is ID-secure, then $\forall i \in \mathcal{I}. i$ is ID-secure for t .*

It is easy to prove that all inputs are secure for a secure program, in both the ID-security and the CP-security case. However, it might not be obvious that a program such that every input is CP-secure is a CP-secure program. Note that whenever two inputs i, i' are similar at some level ℓ they have exactly the same restriction at that level.

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

By definition of secure input and transitivity of \sim_ℓ^{CP} , we can conclude that the output streams produced by t with i and i' must be CP-similar, provided that i and i' are CP-secure.

Lemma 16 indicates that if a program is ID-secure, then every input for that program is ID-secure, i.e. running the program under those inputs leaks the same amount of information as if secrets had been erased. The converse, however, does not hold. We illustrate this point with the following example.

Example 17. Consider the program p in Figure 4. Every input is ID-secure for p but p is not an ID-secure program, as the following two ID-similar inputs show.

$$i = [(\text{H?}, 1), [(\text{L?}, 0)]] \sim_L^{\text{ID}} i' = [(\text{H?}, 2), (\text{L?}, 0)]$$

Let μ_0 be the initial memory where every variable is initialised to 0, and let $t = \llbracket p \rrbracket(\mu_0)$, the interaction tree obtained from p and μ_0 . 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 .

$$\begin{aligned} (t, i) &\Rightarrow [\bullet, \bullet, \bullet, \bullet, (\perp!, 1)] \\ &\quad \approx_L^{ID} \\ (t, i') &\Rightarrow [\bullet, \bullet, \bullet, \bullet, (\perp!, 2)] \end{aligned}$$

Nevertheless, all inputs i are ID-secure for t (Definition 14). 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 H ? produces an output which is silent and infinite, and hence ID-similar to every other output.

The following theorems state the transparency of secure inputs for secure multi-execution with our order-preserving scheduler. That is, the theorems show that outputs of a given program under secure multi-execution are not observably different when provided with a secure input.

Theorem 18 (Transparency for ID-secure inputs). *Let t be an interaction tree, and let i be an input colist ID-secure for t such that $(t, i) \Rightarrow o$. If $(sme(ops(t, i), t), i) \Rightarrow o'$ then $\forall \ell. o \sim_\ell^{ID} o'$.*

The theorem above uses \sim_ℓ^{ID} and therefore assures transparency under this notion of observation, i.e. differences in outputs due to divergence are not observable and therefore not captured. If one wants to distinguish productive outputs from divergence, then CP-similarity is the right notion of observation.

Theorem 19 (Transparency for CP-secure inputs). *Let t be an interaction tree, and let i be an input colist CP-secure for t such that $(t, i) \Rightarrow o$. If $(sme(ops(t, i), t), i) \Rightarrow o'$ then $\forall \ell. o \sim_\ell^{CP} o'$.*

VI. MULTI-EXECUTION MONITOR

An important problem of the secure multi-execution approach is that it makes programs non-interferent by modifying its semantics. Consequently, it is difficult to detect if, and when, programs behave maliciously. To remedy this situation, we present a monitor capable of precisely detecting when an input is insecure for a program.

Our monitor executes the original program and its secure multi-execution in parallel, checking at each step that both executions would produce the same output command. If outputs differ, we are in presence of an information leak, and thus execution is aborted. If, on the other hand, executions remain synchronised, the output command is safe to be executed. It is crucial for the monitor to work that secure multi-execution is run under a scheduler that preserves the order of output events such as the scheduler ops from Section IV.

At any point during execution, the monitor might need to signal an alarm. We define a new kind of datatype that can represent the outputs of the monitor. This datatype, written \mathcal{O}_ϵ , is similar to a colist of output events, but it may end

$$\begin{aligned} & \frac{lvl(e) = \ell \quad p(\ell) = Write(e, p_\ell) \quad (t, i, p[\ell \mapsto p_\ell], q) \Downarrow o}{(W_1) \quad (Write(e, t), i, p, q) \Downarrow e :: o} \\ & \frac{lvl(e) = \ell \quad p(\ell) = Write(e', p_\ell) \quad e \neq e' \quad (Write(e, t), i, p, q) \Downarrow \epsilon}{(W_2) \quad (Write(e, t), i, p, q) \Downarrow \epsilon} \\ & \frac{lvl(e) = \ell \quad p(\ell) = Reuse f_\ell \quad q(\ell) = e' :: q_\ell \quad (Write(e, t), i, p[\ell \mapsto f_\ell(e')], q[\ell \mapsto q_\ell]) \Downarrow o}{(W_3) \quad (Write(e, t), i, p, q) \Downarrow \bullet :: o} \\ & \frac{lvl(e) = \ell \quad p(\ell) = Reuse f_\ell \quad q(\ell) = []}{(W_4) \quad (Write(e, t), i, p, q) \Downarrow \epsilon} \\ & \frac{lvl(e) = \ell \quad p(\ell) = Step p_\ell \quad (Write(e, t), i, p[\ell \mapsto p_\ell], q) \Downarrow o}{(W_5) \quad (Write(e, t), i, p, q) \Downarrow \bullet :: o} \\ & \frac{(t, i, p, q) \Downarrow o}{(S) \quad (Step t, i, p, q) \Downarrow \bullet :: o} \\ & \frac{(f(e), i, p, q) \Downarrow o}{(R_1) \quad (Read f, e :: i, p, q) \Downarrow \bullet :: o} \\ & \frac{}{(R_2) \quad (Read f, [], p, q) \Downarrow go(p, q)} \end{aligned}$$

Figure 5. Semantics for the multi-execution monitor

with an alarm ϵ . More formally, we coinductively define \mathcal{O}_ϵ as follows.

$$\mathcal{O}_\epsilon = [] \mid \epsilon \mid O :: \mathcal{O}_\epsilon$$

At a first glance, it seems enough to simply run the interaction tree of the program under consideration in parallel with the one obtained from the sme transformation. However, in order to precisely detect violations of the security policy a new relation is needed. The monitor is expressed by the relation $\Downarrow \subseteq React \times \mathcal{I} \times (\mathcal{L} \rightarrow Produce) \times Q \times \mathcal{O}_\epsilon$ defined in Figure 5, where we write $(t, i, p, q) \Downarrow o$ whenever $(t, i, p, q, o) \in \Downarrow$. The intuition is that whenever $(t, i, p, q) \Downarrow o$, the evaluation of the interaction tree t , together with the input i , a producer for each security level p , and a system of queues q , results in an output o . Note that components t and i pertain to the execution of the original program, while p and q to its secure multi-execution.

Rules (W_1) to (W_5) concern the case where the monitored interaction tree wants to do a $Write$ of an event at security level ℓ . In rule (W_1) , the producer at level ℓ ($p(\ell)$) tries to write the same event, hence the event is performed ($e :: o$). In rule (W_2) , on the other hand, the producer at level ℓ tries to write a different event ($e \neq e'$), so the monitor raises an alarm (ϵ) and aborts execution. In rule (W_3) and (W_4) , the producer at level ℓ tries to reuse some input events. If there is an event available ($q(\ell) = e' :: q_\ell$), rule (W_3) provides

the information to the producer ($p[\ell \mapsto f_\ell(e')]$). In contrast, if there is no event available ($q(\ell) = []$), rule (W_4) aborts execution (ϵ). Finally, in rule (W_5), the producer at level ℓ makes a silent step, so a silent event is performed by the monitor (\bullet).

In rule (S), the interaction tree tries to make a silent step ($Step$), in which case the monitor outputs a silent event (\bullet). Rules (R_1) and (R_2) are related to consuming input data. If there is an input event ($e :: i$), rule (R_1) consumes it ($f(e)$). If, on the other hand, there is no event, rule (R_2) determines the result of the evaluation depending on the state of system of queues, as determined by the function go .

The function go tries to examine if the termination behaviour of the interaction tree matches its secure multi-execution version. There are three possible situations. In the first one, every producer has consumed its input queue and the next command to be executed is a *Reuse*, thus matching the *Read* event. That is, the producers are synchronised ($sync(p, q)$), and go can simply return the empty colist. In the second situation, if there is one *Write* event as the next command to be executed by a producer at level ℓ , noted $writer(p(\ell))$, then an alarm is raised, since executions are out of sync. Lastly, it may happen that none of the two conditions above apply, e.g., there are no writes but some of the producers are on *Step* commands. In this situation, go should make progress on the producers until some of the two first situations occur. This search, however, may lead to divergence. More specifically, we define go as follows.

$$go(p, q) = \begin{cases} [] & \text{if } sync(p, q) \\ \epsilon & \text{if } \exists \ell, writer(p(\ell)) \\ \bullet :: go(next(p, q)) & \text{otherwise} \end{cases}$$

where the predicate $sync$ and the function $next$ are those defined in Section III. It may seem odd to go through the complication of defining the function go when the program was going to end anyway, but the following example shows why this is needed:

Example 20. Consider the following reactive program:

$$p = \text{H?}(x) \{ r := x \}; \\ \text{L?}(x) \{ \mathbf{if} \ r = 0 \\ \quad \{ \mathbf{out}(L!, 1) \} \\ \quad \{ \mathbf{skip} \} \};$$

Let t be the interaction tree obtained from it for the initial memory μ_0 , and let $i = [(H, 1), (L, 0)]$. The evaluation of t with input i , will result in a finite sequence of invisible events.

$$(t, i) \Rightarrow [\bullet, \bullet, \bullet, \bullet, \bullet]$$

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

$$(t, i_{\blacktriangleright L}) \Rightarrow [\bullet, \bullet, (L, 1)]$$

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

In order to detect cases like this one, where the normal execution ends, but its execution with the secrets erased would produce an output, the monitor needs to make sure that the execution of the producers with the available inputs will not produce an output, as done by rule (R_2).

A. Properties of multi-execution monitoring

This section describes the most important contribution of this work. We establish the security policy enforced by our monitor as well as the transparency guarantees. To simplify notation, given an interaction tree t and input colist i , we define $monitor(t, i) = (t, i, \lambda \ell. prd_\ell(t), \lambda \ell. i_{\blacktriangleright \ell})$ as the initial configuration of our monitor for auditing program t under the input events i . We define the predicate $ok \subseteq \mathcal{O}_\epsilon$ identifying outputs colists where the monitor has not raised an alarm. The predicate is defined coinductively by the following rules.

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

It is easy to see that an output is not ok precisely when it ends with an alarm ϵ . Clearly, we can embed any output $o \in \mathcal{O}_\epsilon$ such that $ok(o)$ in \mathcal{O} . To avoid additional notation, we perform such embedding transparently.

The following theorem states that the multi-execution monitor is able to precisely detect ID-secure inputs.

Theorem 21 (Precision for ID-secure runs). *Let $t : React$ and let i be an input colist such that $monitor(t, i) \Downarrow o$. Then*

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

This theorem states that, if and only if the monitor raises an alarm ($\neg ok(o)$), the run i is not ID-secure. Consequently, the run has tried to leak more information than the one observed in a run where secrets are not present in the system (Definition 14). In this case, having detected such condition, the monitor assures that the program under surveillance is not ID-secure (contrapositive of Lemma 16). Differently from most of the dynamic monitoring techniques for confidentiality (e.g. [2, 5, 19, 28, 37, 40, 42]), our monitor does not raise false alarms due to some imprecision in the analysis of information flow inside a program.

When our monitor does not raise an alarm, we cannot infer the ID-security of the program, only the ID-security of the observed run. It could be the case that the program is ID-secure, in which case the monitor will never raise an alarm (Lemma 16), but also that the program is interferent and the input was ID-secure. This last case is common in dynamic monitors which accept non-leaking runs of interferent programs [2, 5, 19, 34, 37, 42].

Once an alarm is raised, the multi-execution monitor is capable to report a trace where the detected insecurity gets

revealed. More precisely, if we have that $monitor(t, i) \Downarrow o$ and $\neg ok(o)$, o is a trace that exposes the detected insecurity. The preservation in the order of events, i.e., the use of the scheduler ops , is essential to assert that o is indeed a trace of the monitored program.

The next theorem establishes that the monitor is transparent for ID-secure runs. In particular, the interleaving of events from different security levels is not altered (\sim_{ℓ}^{ID}).

Theorem 22 (Transparency for ID-secure runs). *Let $t : React$ and i an input colist such that $(t, i) \Rightarrow o$ and $monitor(t, i) \Downarrow o'$. Then,*

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

When considering CP-security, we can guarantee that the monitor will not raise an alarm and be transparent for CP-secure inputs. However, a CP-insecure input may cause the monitor to diverge without raising an alarm, as the monitor cannot predict if another visible event will be found.

Theorem 23 (CP-precise monitor). *Let $t : React$ and i an input colist such that $(t, i) \Rightarrow o$ and $monitor(t, i) \Downarrow o'$. Then,*

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

Raising an alarm clearly reveals information by the termination channel [1]. To mitigate this particular leak, our monitor can be adapted to include techniques for suppression or modification of outputs [25]. For instance, and considering ID-security, the monitor could silently diverge when an insecurity is detected. Alternatively, the monitor could change the outputs of the program to match those dictated by SME under any scheduler. In both cases, the execution would be the same as that of an ID-secure program.

VII. RELATED WORK

Precise dynamic enforcement of non-interference: A series of work characterises the security policies enforceable by execution monitoring [18, 25, 39]. As a result of that, it is known that non-interference is not a safety property (see [29, 39] for a proof), and therefore not enforceable by execution monitors. The main argument for that claim relies on the fact that non-interference relates a pair of execution traces, while safety properties refer to a single one. Despite being inherently imprecise, researchers propose execution monitors to enforce, in the shape of a safety property, a stronger version of non-interference (e.g. [2, 4, 5, 34, 37].) Although these monitors stop the execution of potentially dangerous programs, they still reject some non-interferent ones. Motivated by theorem proving techniques, Darvas, Hähnle, and Sands [15] show how to cast non-interference into a safety property by composing programs with a copy of themselves. This technique is known as self-composition (term coined in [8]) and it has been used to exploit known techniques for program verification [8, 45]. It is an open

question if some sort of self-composition could precisely, and dynamically, detect when programs violate confidentiality. This work shows that it is possible to precisely, and dynamically, detect when the notion of non-interference ID-security [11] gets violated.

Secure multi-execution: The closest related work to ours is secure multi-execution. From a systems perspective, Capizzi et al. [12] describe the idea of secure multi-execution for two security levels using the term *shadow executions*. Similarly to this work, the authors do not cover timing covert channels. However, they do not give any transparency guarantees for secure programs when using shadow executions. Similarly to Capizzi et al., but looking to obtain a secure Linux kernel, Cristiá and Mata [14] consider similar ideas as secure multi-execution using two security levels and ignoring timing channels. Differently from this work, their method is formalised for a specific programming language. Devriese and Piessens [17] introduce secure multi-execution. In that work, authors evaluate the practicality of their ideas in the Spidermonkey JavaScript engine and prove theoretical results for a specific scheduler. Barthe et al. [7] show how to achieve secure multi-execution by code transformation. While their transformation is defined over a specific programming language, ours in Section III is described for interaction trees and thus it is more general. Under the same scheduler and security condition as Devriese and Piessens, Barthe et al. prove the soundness and precision of the transformed code. Focusing only on implementation issues, Jaskelioff and Russo [23] provide secure multi-execution for Haskell programs via a library. In that work, the authors propose a pure description of the I/O operations of programs that influenced the adoption of interaction trees in this paper. In order to deal with timing leaks, authors in [7, 17] require a total order of the lattice and choose a specific scheduler. Instead, authors in [24] describe a range of schedulers capable of preventing timing leaks which depends on the comparability of the elements in the lattice. In this work, we ignore the external timing covert channel for the sake of simplicity and precision of our enforcement. However, due to the modularity of our approach, we could easily apply practical black-box techniques [3, 51] to mitigate timing leaks. Bielova et al. [9] adapt secure multi-execution for web browsers. Similarly to this work, they consider a notion of secure runs for which they can provide transparency guarantees, i.e., that the behaviour of those runs is not altered by secure multi-execution. An extended version of this work describes, as an implementation optimization, situations where it is not necessary to run a browser per security level in the lattice. None of the works described above [7, 9, 12, 14, 17, 24] can preserve the interleaving of events generated at different security levels as well as report when insecurities occur. Instead, at the price of not considering timing covert channels, our work describes a scheduling strategy capable of preserving such order and

detecting insecure actions violating ID-security. Focusing on extending secure multi-execution, and independent of this work, Rafnsson and Sabelfeld [33] propose a scheduler that is also able to preserve the order of events for interactive programs. Their work lifts the totality assumptions on input channels, i.e. that inputs are always available, and introduce means for declassification. Our monitor might be able to benefit from these results since it uses secure-multi execution underneath.

Faceted values: Focusing on gaining performance, Austin and Flanagan [6] proposed a semantics based on faceted values that simulates multiple executions in one run. Differently from this work, execution with faceted values requires a full description of the underlying programming language semantics. They provide no formal guarantees that the interleaving of output events at different security levels is preserved. Similarly to secure multi-execution, this approach is not capable of detecting when insecurities occur during the execution of programs.

Non-interference for reactive systems: Bohannon et al. [11] define several non-interference notions for reactive systems including ID- and CP-security. While CP-security provides stronger guarantees, it is more difficult to enforce by information-flow techniques. As noticed by Rafnsson and Sabelfeld [32], ID- and CP-security are termination-insensitive, making it possible to leak secret values by brute force attacks. Modern information-flow tools like Jif [30] (based on Java), SPARK Examiner [13] (based on Ada), and the sequential version of LIO [42] (based on Haskell) are not strong enough to avoid these leaks. For deterministic systems like the ones we consider, the bandwidth of leaking information by exploiting outputs in combination with termination is logarithmic in the size of the secret, i.e. it takes exponential time in the size of the secret to leak its whole value [1]. Nevertheless, the bandwidth can be reduced by applying buffering techniques [32]. We could easily adapt our multi-execution monitor to do that.

Interaction trees: The interaction trees used to model reactive systems in this work are based on the coalgebraic view of systems [21]. Swierstra and Altenkirch [43, 44] use interaction trees to provide a functional model of some of the features of Haskell’s I/O monad such as mutable state, interactive programming and concurrency. Differently from us, they do not consider reactive systems.

VIII. SUMMARY

We propose multi-execution monitoring, a novel technique combining execution monitoring and secure multi-execution. This technique precisely detects actions that reveal information under the notion of ID-security. Consequently, we keep alarms to the minimum. We also prove that the monitor provides good transparency guarantees for the progress-sensitive non-interference notion of CP-security. To achieve these results, we rely on a scheduling strategy for secure

multi-execution which allows us to preserve the interleaving of events, and the notions of ID-secure and CP-secure inputs.

Having the foundations for our multi-execution monitor, we can take our work into several future directions. Interaction trees can be easily adapted to model interactive programs [43, 44] (by adding one constructor *Stop* and modifying the *Read* constructor to let the program, instead of the environment, choose the channel.) However, interactive programs would require the modification of the secure multi-execution mechanism in order to consider default values for producers reading data from higher security levels. Modelling non-determinism is another interesting direction to explore. To model non-determinism, instead of considering one interaction with the environment at the time, one could consider finite sets of them [21]. This modification demands a change in the notion of similarity depending on attackers’ power [31, 41, 47]. Clearly, this extension requires to extend secure multi-execution to account for non-determinism while being permissive, which is an open challenge. Declassification [38], or intended release of information, is a desirable feature of any practical information-flow system. Taking the declassification policy of delimited release [36], Askarov and Sabelfeld [2] show techniques to dynamically enforce it. We believe that our monitor could enforce declassification permissively. For that, we would extend the interaction tree model with a special constructor indicating a declassification action. Every time that the original program reads a secret data that would eventually be declassified, our monitor forwards it to the public producers. Then, when reaching a declassification point, the original program and the public producers need to be sync: they should both release the same values; otherwise, the original program might be releasing more information than expected.

Acknowledgments: We thank David Sands for suggesting a more appropriate reference to his work and the anonymous reviewers for insightful comments and bringing several references to our attention. This work was funded by the Swedish research agencies VR and STINT and the *Agencia Nacional de Promoción Científica y Tecnológica* (PICT 2009-15)

REFERENCES

- [1] 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.
- [2] 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.
- [3] A. Askarov, D. Zhang, and A. C. Myers. Predictive black-box mitigation of timing channels. In *Proc. of*

- the 17th ACM conference on Computer and Communications Security. ACM, 2010.
- [4] 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.
- [5] 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.
- [6] T. H. Austin and C. Flanagan. Multiple facets for dynamic information flow. In *Proceedings of the 39th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages, POPL '12*. ACM, 2012.
- [7] 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 (FMOODS/FORTE 2012)*, June 2012.
- [8] G. Barthe, P. R. D'Argenio, and T. Rezk. Secure Information Flow by Self-Composition. In *Proceedings of the 17th IEEE workshop on Computer Security Foundations, CSFW '04*. IEEE Computer Society, 2004.
- [9] 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.
- [10] A. Bohannon and B. C. Pierce. Featherweight Firefox: formalizing the core of a web browser. In *Proceedings of the 2010 USENIX conference on Web application development, WebApps'10*. USENIX Association, 2010.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] A. Darvas, R. Hähnle, and D. Sands. A theorem proving approach to analysis of secure information flow (preliminary version). In R. Gorrieri, editor, *Workshop on Issues in the Theory of Security, WITS*, 2003.
- [16] D. E. Denning and P. J. Denning. Certification of Programs for Secure Information Flow. *Commun. ACM*, 20(7):504–513, July 1977.
- [17] 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.
- [18] K. W. Hamlen, G. Morrisett, and F. B. Schneider. Computability classes for enforcement mechanisms. *ACM Trans. Program. Lang. Syst.*, Jan. 2006.
- [19] D. Hedin and A. Sabelfeld. Information-Flow Security for a Core of JavaScript. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society, 2012.
- [20] 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.
- [21] B. Jacobs and J. Rutten. A Tutorial on (Co)Algebras and (Co)Induction. *Bulletin of the European Association for Theoretical Computer Science*, 62:222–259, 1997.
- [22] 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.
- [23] 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.
- [24] V. Kashyap, B. Wiedermann, and B. Hardekopf. Timing- and Termination-Sensitive Secure Information Flow: Exploring a New Approach. In *Proc. of IEEE Symposium on Sec. and Privacy*. IEEE, 2011.
- [25] J. Ligatti, L. Bauer, and D. Walker. Edit Automata: Enforcement Mechanisms for Run-time Security Policies. *International Journal of Information Security*, Feb. 2005.
- [26] 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.
- [27] 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.

- [28] J. Magazinius, A. Russo, and A. Sabelfeld. On-the-fly Inlining of Dynamic Security Monitors. In *Proceedings of the International Information Security Conference, SEC*, 2010.
- [29] 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.
- [30] 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.
- [31] W. Rafnsson, D. Hedin, and A. Sabelfeld. Securing Interactive Programs. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society, 2012.
- [32] 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.
- [33] W. Rafnsson and A. Sabelfeld. Secure multi-execution: fine-grained, declassification-aware, and transparent. Submitted, Feb. 2013.
- [34] A. Russo and A. Sabelfeld. Securing Timeout Instructions in Web Applications. In *Proc. IEEE Computer Sec. Foundations Symposium*. IEEE Computer Society, July 2009.
- [35] A. Sabelfeld and A. C. Myers. Language-Based Information-Flow Security. *IEEE J. Selected Areas in Communications*, 21(1):5–19, Jan. 2003.
- [36] A. Sabelfeld and A. C. Myers. A Model for Delimited Information Release. In *Proc. International Symp. on Software Security (ISSS'03)*, volume 3233 of LNCS, pages 174–191. Springer-Verlag, Oct. 2004.
- [37] 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, LNCS*. Springer-Verlag, June 2009.
- [38] A. Sabelfeld and D. Sands. Dimensions and Principles of Declassification. In *Proc. IEEE Computer Sec. Foundations Workshop*, pages 255–269, June 2005.
- [39] F. B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*, 2000.
- [40] P. Shroff, S. Smith, and M. Thober. Dynamic Dependency Monitoring to Secure Information Flow. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium, CSF '07*. IEEE Computer Society, 2007.
- [41] G. Smith and D. Volpano. Secure Information Flow in a Multi-threaded Imperative Language. In *Proc. ACM Symp. on Principles of Programming Languages*, pages 355–364, Jan. 1998.
- [42] D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières. Flexible Dynamic Information Flow Control in Haskell. In *Haskell Symposium*. ACM SIGPLAN, September 2011.
- [43] W. Swierstra. *A Functional Specification of Effects*. PhD thesis, University of Nottingham, November 2008.
- [44] 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.
- [45] T. Terauchi and A. Aiken. Secure information flow as a safety problem. In *Proceedings of the 12th international conference on Static Analysis, SAS'05*. Springer-Verlag, 2005.
- [46] 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.
- [47] D. Volpano and G. Smith. Probabilistic Noninterference in a Concurrent Language. *J. Computer Security*, 7(2–3), Nov. 1999.
- [48] D. Volpano, G. Smith, and C. Irvine. A Sound Type System for Secure Flow Analysis. *J. Computer Security*, 4(3):167–187, 1996.
- [49] Williams and D. Wichers. OWASP Top 10 2010. http://www.owasp.org/index.php/Top10_2010, 2010.
- [50] D. Zanarini, M. Jaskelioff, and A. Russo. Precise Enforcement of Confidentiality for Reactive Systems: Extended Version, 2013. Available at <http://www.fceia.unr.edu.ar/~dante>.
- [51] D. Zhang, A. Askarov, and A. C. Myers. Predictive mitigation of timing channels in interactive systems. In *Proc. of the 18th ACM conference on Computer and Communications Security*. ACM, 2011.

APPENDIX A.

INTERACTION TREES FOR JAVASCRIPT-LIKE PROGRAMS

We show how to obtain the corresponding interaction tree for programs in the JavaScript-like language of Figure 2. We assume that every handler uses variable x to store the message received, and we extend the set of commands \mathbb{C} with an special symbol \dagger . Making these assumptions, we can reinterpret p as mapping from channels to the set $\mathbb{C} \cup \{\dagger\}$ as follows.

$$\begin{aligned} (\cdot)_m &= \lambda ch. \dagger \\ (ch(x)\{c\}; p)_m &= \lambda ch'. (\text{if } ch = ch' \text{ then } c \text{ else } (p)_m) \end{aligned}$$

$$\begin{array}{c}
\frac{}{(p_m, \mu, \phi) \xrightarrow{\downarrow} (p_m, \mu, \phi)} \\
\frac{}{(p_m, \mu, \mathbf{skip}) \xrightarrow{\odot} (p_m, \mu, \phi)} \\
\frac{(\mu, e) \Downarrow v}{(p_m, \mu, x := e) \xrightarrow{\odot} (p_m, \mu[x \mapsto v], \phi)} \\
\frac{}{(p_m, \mu, \mathbf{new} \text{ } ch(x)\{c\}) \xrightarrow{\odot} (p_m[ch \mapsto c], \mu, \phi)} \\
\frac{(\mu, e) \Downarrow v}{(p_m, \mu, \mathbf{out} \text{ } (ch, e)) \xrightarrow{\uparrow_v^{ch}} (p_m, \mu, \phi)} \\
\frac{}{(p_m, \mu, \mathbf{eval}(s)) \xrightarrow{\odot} (p_m, \mu, \mathbf{parse}(s))} \\
\frac{(\mu, e) \Downarrow 0}{(p_m, \mu, \mathbf{if} \text{ } e \{c\}\{c'\}) \xrightarrow{\odot} (p_m, \mu, c')} \\
\frac{(\mu, e) \Downarrow v \quad v \neq 0}{(p_m, \mu, \mathbf{if} \text{ } e \{c\}\{c'\}) \xrightarrow{\odot} (p_m, \mu, c)} \\
\frac{(\mu, e) \Downarrow 0}{(p_m, \mu, \mathbf{while} \text{ } e \text{ } \mathbf{do} \text{ } c) \xrightarrow{\odot} (p_m, \mu, \phi)} \\
\frac{(\mu, e) \Downarrow v \quad v \neq 0}{(p_m, \mu, \mathbf{while} \text{ } e \text{ } \mathbf{do} \text{ } c) \xrightarrow{\odot} (p_m, \mu, c; \mathbf{while} \text{ } e \text{ } \mathbf{do} \text{ } c)} \\
\frac{(p_m, \mu, c) \xrightarrow{i} (p'_m, \mu', \phi)}{(p_m, \mu, c; c') \xrightarrow{i} (p'_m, \mu', c')} \\
\frac{(p_m, \mu, c) \xrightarrow{i} (p'_m, \mu', c_0) \quad c_0 \neq \phi}{(p_m, \mu, c; c') \xrightarrow{i} (p'_m, \mu', c_0; c')}
\end{array}$$

Figure 6. Next interaction for a reactive state

The definition above simplifies our transformation, since we can treat uniformly all input events. We note $(p)_m$ as p_m for simplicity.

A state s is a tuple (p_m, μ, c) , where

- ▶ p_m is a mapping from channels to $\mathbb{C} \cup \{\phi\}$ (the re-interpretation of a reactive program p as a mapping);
- ▶ μ is a mapping from variables to values, the memory;
- ▶ $c \in \mathbb{C} \cup \{\phi\}$ is the command being executed by the machine in response to an input event, or ϕ if the system is ready to receive an input.

We name *State* the set of states for our JavaScript-like reactive program. Given a state $s = (p_m, \mu, c)$, we can compute the next interaction of s with the environment. An interaction is an element in the set $Signal = \{\downarrow, \uparrow_v^{ch}, \odot\}$. Interaction \downarrow is raised if s is a consumer state, i.e. if $c = \phi$. Signal \uparrow_v^{ch} is raised when c will produce an output message (ch, v) . Otherwise, the next step of s is

$$\begin{array}{l}
\llbracket - \rrbracket_{st} : State \rightarrow React \\
\llbracket s \rrbracket_{st} = \mathbf{let} \ (i, (p_m, \mu, c)) = \mathit{step}(s) \ \mathbf{in} \\
\quad \mathbf{case} \ i \ \mathbf{of} \\
\quad \odot \ \rightarrow \ \mathit{Step} \ (\llbracket (p_m, \mu, c) \rrbracket_{st}) \\
\quad \uparrow_v^{ch} \ \rightarrow \ \mathit{Write} \ (ch, v, \llbracket (p_m, \mu, c) \rrbracket_{st}) \\
\quad \downarrow \ \rightarrow \ \mathit{Read} \ (\lambda(ch, v). \llbracket (p_m, \mu[x \mapsto v], p_m(ch)) \rrbracket_{st})
\end{array}$$

Figure 7. Generation of interaction trees for Javascript-like programs

silent, represented by signal \odot . Figure 6 defines a function $\mathit{step} : State \rightarrow Signal \times State$. We write $s \xrightarrow{i} s'$ for $\mathit{step}(s) = (i, s')$.

The definition of step makes use of some auxiliary functions. The first one is an evaluation function \Downarrow , such that $(\mu, e) \Downarrow v$ iff expression e evaluates to value v in memory μ . We assume that \Downarrow is a side-effect free function. Second, function parse takes a string s and returns the command denoted by it if successfully parsed. For simplicity, if there is an error when parsing s , we assume parse returns \mathbf{skip} . The rules are self-explanatory and therefore we do not discuss them.

Given a state $s = (p_m, \mu, c)$, and using function step , Figure 7 shows how to map interactions from s into interaction trees. The effects on the environment (i.e. \odot , \downarrow , and \uparrow_v^{ch}) are simply mapped into the constructors of interaction trees. The interesting case is the one related to input events. When processing input events, the interaction tree describes that, when an event arrives, the corresponding event handler is invoked $(p_m(ch))$, and variable x gets updated with the received value.

We now define function $\llbracket - \rrbracket$ from Section II-B as follows.

$$\begin{array}{l}
\llbracket - \rrbracket \quad : \ Prg \times M \rightarrow React \\
\llbracket p \rrbracket(\mu) = \llbracket (p_m, \mu, \phi) \rrbracket_{st}
\end{array}$$