HAL
open science

Fine-grained and coarse-grained reactive noninterference

Pejman Attar, Ilaria Castellani

To cite this version:

Pejman Attar, Ilaria Castellani. Fine-grained and coarse-grained reactive noninterference. Trustworthy Global Computing 2013 - 8th International Symposium, Revised Selected Papers, Martín Abadi; Alberto Lluch-Lafuente, Aug 2013, Buenos Aires, Argentina. pp.21, 10.1007/978-3-319-05119-2_10 . hal-00915241

HAL Id: hal-00915241

https://inria.hal.science/hal-00915241

Submitted on 6 Dec 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fine-grained and coarse-grained reactive noninterference

Pejman Attar and Ilaria Castellani*

INRIA**

pejman.attar@inria.fr, ilaria.castellani@inria.fr

Abstract. We study bisimilarity and the security property of *noninterference* in a core *synchronous reactive language* that we name CRL.
In the synchronous reactive paradigm, programs communicate by means of broadcast events, and their parallel execution is regulated by the notion of *instant*. Within each instant, programs may emit events and get suspended while waiting for events emitted by other programs. They may also explicitly return the control to the scheduler, thereby suspending themselves until the end of the instant. An instant is thus a period of time during which all programs compute until termination or suspension.
In CRL there is no memory, and the focus is on the control structure of programs. An *asymmetric parallel operator* is used to implement a deterministic scheduling. This scheduling is fair – in the sense that it gives its turn to each parallel component – if all components are *cooperative*, namely if they always return the control after a finite number of steps.
We first prove that CRL programs are indeed cooperative. This result is based on two features of the language: the semantics of loops, which requires them to yield the control at each iteration of their body; and a delayed reaction to the absence of events, which ensures the monotonicity of computations (viewed as I/O functions on event sets) during instants.
Cooperativeness is crucial as it entails the *reactivity* of a program to its context, namely its capacity to input events from the context at the start of instants, and to output events to the context at the end of instants.
We define two bisimulation equivalences on programs, formalising respectively a *fine-grained observation* of programs (the observer is viewed as a program) and a *coarse-grained observation* (the observer is viewed as part of the context). As expected, the latter equivalence is more abstract than the former, as it only compares the I/O behaviours of programs at each instant, while the former also compares their intermediate results.
Based on these bisimulations, two properties of *reactive noninterference* (RNI) are proposed. Both properties are time-insensitive and termination-insensitive. Coarse-grained RNI is more abstract than fine-grained RNI, because it views the parallel operator as commutative and abstracts away from repeated emissions of the same event during an instant.
Finally, a type system guaranteeing both security properties is presented. Thanks partly to a design choice of CRL, which offers two separate constructs for loops and iteration, this type system allows for a precise treatment of termination leaks, which are an issue in parallel languages.

* Work partially supported by the french ANR 08-EMER-010 grant PARTOUT.

** Sophia Antipolis Center, 2004 route des Lucioles, 06902 Sophia Antipolis, France.

1 Introduction

Many systems of widespread use, such as web browsers and web applications, may be modelled as *reactive programs*, that is programs that listen and react to their environment in a continuous way, by means of events. Since the environment may include mutually distrusting parties, such as a local user and a remote web server, reactive programs should be able to protect the confidentiality of the data they manipulate, by ensuring a *secure information flow* from the inputs they receive from one party to the outputs they release to another party.

Secure information flow is often formalised via the notion of *noninterference* (NI), expressing the absence of dependency between secret inputs and public outputs (or more generally, between inputs of some confidentiality level to outputs of lower or incomparable level). Originally introduced in [12], NI has been studied for a variety of languages, ranging from standard imperative and functional languages [18,16] to process calculi based on CCS or the pi-calculus [11]. On the other hand, little attention has been paid to noninterference for reactive programs, with the notable exception of [13], [2] and [7].

We shall focus here on a particular brand of reactive programming, namely the *synchronous* one, which was first embodied in the synchronous language *SL* [9], an offspring of Esterel [6], and later incorporated into various programming environments, such as C, Java, Caml and Scheme. In the synchronous paradigm, the parallel execution of programs is regulated by a notion of *instant*. The model of *SL* departs from that of Esterel in that it assumes the reaction to the absence of an event to be postponed until the end of the instant. This assumption helps disambiguating programs and simplifying the implementation of the language. It is also essential to ensure the monotonicity of programs and their reactivity to the environment.

In this work, we will not explicitly model the interaction of a reactive program with the environment (this could be easily done but it would not bring any further insight). Instead, we concentrate on the interaction *within* a reactive program, making sure it regularly converges to a stable state (end of instant), in which the program is ready to interact with the environment. We call this property *cooperativeness* [1] or *internal reactivity*. In the sequel, we shall abandon the distinction between internal reactivity (among the components of a program) and *external reactivity* (towards the environment), to focus on the former.

This paper attempts to explore "secure reactive programming in a nutshell". To this end, we concentrate on a minimal reactive language without memory, consisting of standard sequential operators, an asymmetric parallel operator $\upharpoonleft$ (formalising a kind of *coroutine* parallelism under a deterministic scheduling), together with four typical reactive constructs, which we briefly describe next.

In our *Core Reactive Language CRL*, programs are made of parallel components s, s' – also called "threads" for simplicity in the following – combined with the operator $s \upharpoonleft s'$ and communicating by means of broadcast events. Threads may emit events, via a `generate` ev instruction, and get suspended while waiting for events to be emitted by other threads, through an `await` ev instruction. They may also explicitly yield the control to the scheduler, via a `cooperate`

instruction, thereby suspending themselves until the end of the current instant. An instant is therefore a period of time during which all threads compute until termination or suspension. Clearly, this is a logical rather than a physical notion of instant, since the termination of instants is determined by the collective behaviour of threads rather than by some physical clock. At the end of an instant, all threads are inactive and share the same view of emitted events. At instant change, a preemption construct `do` s `watching` ev allows some suspended parts of threads to be pruned off, thus implementing a time-out mechanism. Interaction with the environment is limited to the start and the end of instants: the environment injects events at the start of instants and collects them at the end.

The starting point of our work is the paper [2], which laid the basis for the study of noninterference in a synchronous reactive language. The present work improves on [2] in several respects, which we summarise below.

The language examined in [2] is similar to CRL but strictly more expressive, including imperative constructs, local declarations and a memory. Indeed, our asymmetric parallel operator $\upharpoonleft$, which gives priority to its left component, is inspired by that of [2]. Here, however, we adopt a slightly different semantics for $s \upharpoonleft s'$, which preserves the position of threads within a program, while the semantics of [2] swapped the positions of s and s' in $s \upharpoonleft s'$ in case s was suspended, reducing it to $s' \upharpoonleft s$. This simple change forces the scheduler in CRL to serve the same thread at the start of each instant, thus avoiding the so-called *scheduling leaks* of [2], and allowing for a more relaxed typing rule for $\upharpoonleft$, which is just the standard rule for symmetric parallel composition.

Moreover, reactivity was not a concern in [2]: as soon as they contained *while loops*, programs were not guaranteed to terminate or suspend within an instant. Hence, it only made sense to consider a fine-grained notion of noninterference. By contrast, in CRL all programs are reactive, thanks to a clear separation between the loop construct `loop` s and the iteration construct `repeat` exp `do` s, and to our semantics for loops, which requires them to yield the control at each iteration of their body. This makes it possible to define a notion of coarse-grained *reactive noninterference* (RNI), which accounts only for the I/O behaviour of programs within each instant. The coarse-grained RNI property has an advantage over the fine-grained one: it exploits in a more direct way the structure of reactive computations, and it recovers the flavour of big-step semantics within each instant, offering a more abstract NI notion for reactive programs.

Finally, our type system is more permissive than that of [2], thanks to the relaxed typing rule for parallel composition and to refined typing rules for the conditional. Both improvements are made possible by design choices of CRL.

The main contributions of this paper are: 1) the reactivity result, 2) the definition of two bisimulation equivalences for synchronous reactive programs, of different granularity. To our knowledge, semantic equivalences for reactive programs have only been studied previously by Amadio [4]; 3) the proposal of two properties of reactive noninterference, based on the above bisimulations, and 4) the presentation of a type system ensuring both noninterference properties.

The rest of the paper is organised as follows. Sections 2 and 3 present the syntax and the semantics of the language CRL. Section 4 is devoted to proving reactivity of CRL programs. Section 5 introduces the two bisimulation equivalences and gives some properties of them. In Section 6 we define our two NI properties. Section 7 presents the security type system and the proof of its soundness. Finally, future and related work are briefly discussed in Section 8.

The proofs of the results are mostly omitted and may be found in [5].

2 Syntax

In this section we introduce the syntax of CRL. Let Val be a set of values, ranged over by v, v', Var a set of variables, ranged over by x, y, z, and $Events$ a set of events, ranged over by ev, ev'. A fixed valuation function $V : Var \to Val$ for open terms is assumed, which however will be left implicit until Section 6.

Expressions. An expression $exp \in Exp$ may be a basic value, a variable, or the value returned by a function. Letting $\overrightarrow{exp}$ denote a tuple of expressions $exp_1, \ldots, exp_n$, the syntax of expressions is:

$$exp \in Exp ::= v \mid x \mid f(\overrightarrow{exp})$$

The evaluation of a function call $f(\overrightarrow{exp})$ is assumed to be instantaneous, and therefore so is the evaluation of an expression exp, denoted by $exp \rightsquigarrow v$, which is formally defined by the three rules:

$$\frac{}{v \rightsquigarrow v} \qquad \frac{V(x) = v}{x \rightsquigarrow v} \qquad \frac{\forall i \in \{1, \ldots, n\} \,.\, exp_i \rightsquigarrow v_i \qquad f(v_1, \ldots, v_n) = v}{f(\overrightarrow{exp}) \rightsquigarrow v}$$

Programs. We now present the syntax of CRL programs. Alongside with typical sequential operators, CRL includes four operators that are commonly found in reactive languages, $\texttt{cooperate}, \texttt{generate}\, ev, \texttt{await}\, ev$ and $\texttt{do}\, s\, \texttt{watching}\, ev$, as well as a binary *asymmetric parallel operator*, denoted by $\upharpoonleft$, which performs a deterministic scheduling on its components. This operator is very close to that used in [2] and, earlier on, in the implementation of *SugarCubes* [10]. However, while in [2] and [10] each parallel component was executing as long as possible, our operator $\upharpoonleft$ implements a form of *prioritised scheduling*, where the first component yields the control only when terminating or suspending (*late cooperation*), while the second yields it as soon as it generates an event that unblocks the first component (*early cooperation*). The syntax of programs is given by:

$$\begin{aligned} s \in Programs ::= {} & \texttt{nothing} \mid s; s \mid (s \upharpoonleft s) \mid \\ & \texttt{cooperate} \mid \texttt{generate}\, ev \mid \texttt{await}\, ev \mid \texttt{do}\, s\, \texttt{watching}\, ev \mid \\ & (\texttt{loop}\, s) \mid (\texttt{repeat}\, exp\, \texttt{do}\, s) \mid (\texttt{if}\, exp\, \texttt{then}\, s\, \texttt{else}\, s) \end{aligned}$$

Note that our language includes two different constructs for loops and iteration, in replacement of the standard *while loop* operator. This allows for a clear separation between nonterminating behaviours and iterative behaviours.

3 Semantics

This section presents the operational semantics of CRL. Programs proceed through a succession of instants, transforming sets of events. There are two transition relations, both defined on *configurations* of the form $\langle s, E\rangle$, where s is a program and $E \subseteq Events$ is an *event environment*, i.e. a set of present events.

Let us first give the general idea of these two transition relations:

1. The *small-step transition relation* describes the step-by-step execution of a configuration within a an instant. The general format of a transition is:
$$\langle s, E\rangle \rightarrow \langle s', E'\rangle$$
where:
   - s is the program to execute and s' is the residual program;
   - E is the starting event environment and E' is the resulting event environment: E' coincides with E if the transition does not generate any new event; otherwise $E' = E \cup \{ev\}$, where ev is the new generated event.
2. The *tick transition relation* describes the passage from one instant to the next, and applies only to suspended configurations. A transition of this kind has always the form:
$$\langle s, E\rangle \hookrightarrow \langle [s]_E, \emptyset\rangle$$
where the resulting event environment is empty and $[s]_E$ is a "reconditioning" of program s for the next instant, possibly allowing it to resume execution at the next instant even without the help of new events from the environment.

Before formally defining $\rightarrow$ and $\hookrightarrow$, we introduce the *suspension predicate* $\langle s, E\rangle\ddagger$, which holds when s is suspended in the event environment E, namely when all threads in s are waiting for events not contained in E, or have deliberately yielded the control for the current instant by means of a `cooperate` instruction.

The rules defining the predicate $\ddagger$ and the relations $\rightarrow$ and $\hookrightarrow$ are given in Fig. 3. The *reconditioning function* $[s]_E$ prepares s for the next instant: it erases all guarding `cooperate` instructions, as well as all guarding `do` s' `watching` ev instructions whose time-out event ev belongs to E (i.e. has been generated).

We assume programs are well-typed with respect to a standard type system that ensures that in the commands `if` exp `then` s_1 `else` s_2 and `repeat` exp `do` s the expression exp evaluates respectively to a boolean and to an integer $n \geq 1$.

Let us comment on the most interesting transition rules. The execution of a parallel program always starts with its left branch (Rule (par_1)). Once the left branch is over, the program reduces to its right branch (Rule (par_2)). If the left branch is suspended, then the right branch executes (Rule (par_3)) until unblocking the left branch. Thus *early cooperation* is required in the right branch. To avoid nondeterminism, a terminated right branch can only be eliminated if the left branch is suspended (Rule (par_4)). A `loop` s program executes its body cyclically: a `cooperate` instruction is systematically added in parallel to its body to avoid *instantaneous loops*, i.e. divergence within an instant[1] (Rule $(loop)$). A `do` s `watching` ev program executes its body until termination or suspension (Rule $(watch_1)$), reducing to `nothing` when its body terminates (Rule $(watch_2)$).

[1] In general, we shall call "instantaneous" any property that holds within an instant.

$$\langle cooperate, E\rangle\ddagger \quad (coop) \qquad \frac{ev \notin E}{\langle \texttt{await}\ ev, E\rangle\ddagger}\ (wait_s) \qquad \frac{\langle s, E\rangle\ddagger}{\langle \texttt{do}\ s\ \texttt{watching}\ ev, E\rangle\ddagger}\ (watch_s)$$

$$\frac{\langle s_1, E\rangle\ddagger}{\langle s_1; s_2, E\rangle\ddagger}\ (seq_s) \qquad \frac{\langle s_1, E\rangle\ddagger \qquad \langle s_2, E\rangle\ddagger}{\langle s_1 \upharpoonleft s_2, E\rangle\ddagger}\ (par_s) \qquad \frac{\langle s, E\rangle\ddagger}{\langle s, E\rangle \hookrightarrow \langle [s]_E, \emptyset\rangle}\ (tick)$$

Suspension Predicate and Tick Transition Rule

$$[\texttt{cooperate}]_E = \texttt{nothing} \qquad [\texttt{do}\ s\ \texttt{watching}\ ev]_E = \begin{cases} \texttt{nothing} & if \quad ev \in E \\ \texttt{do}\ [s]_E\ \texttt{watching}\ ev & otherwise \end{cases}$$

$$[\texttt{await}\ ev]_E = \texttt{await}\ ev \qquad [s_1; s_2]_E = [s_1]_E\ ;\ s_2 \qquad [s_1 \upharpoonleft s_2]_E = [s_1]_E \upharpoonleft [s_2]_E$$

Reconditioning Function

$$\frac{\langle s_1, E\rangle \to \langle s_1', E'\rangle}{\langle s_1; s_2, E\rangle \to \langle s_1'; s_2, E'\rangle}\ (seq_1) \qquad \langle \texttt{nothing}\ ; s, E\rangle \to \langle s, E\rangle \quad (seq_2)$$

$$\frac{\langle s_1, E\rangle \to \langle s_1', E'\rangle}{\langle s_1 \upharpoonleft s_2, E\rangle \to \langle s_1' \upharpoonleft s_2, E'\rangle}\ (par_1) \qquad \langle \texttt{nothing} \upharpoonleft s, E\rangle \to \langle s, E\rangle \quad (par_2)$$

$$\frac{\langle s_1, E\rangle\ddagger \quad \langle s_2, E\rangle \to \langle s_2', E'\rangle}{\langle s_1 \upharpoonleft s_2, E\rangle \to \langle s_1 \upharpoonleft s_2', E'\rangle}\ (par_3) \qquad \frac{\langle s, E\rangle\ddagger}{\langle s \upharpoonleft \texttt{nothing}, E\rangle \to \langle s, E\rangle}\ (par_4)$$

$$\langle \texttt{generate}\ ev, E\rangle \to \langle \texttt{nothing}, E \cup \{ev\}\rangle\ (gen) \qquad \frac{ev \in E}{\langle \texttt{await}\ ev, E\rangle \to \langle \texttt{nothing}, E\rangle}\ (wait)$$

$$\frac{\langle s, E\rangle \to \langle s', E'\rangle}{\langle \texttt{do}\ s\ \texttt{watching}\ ev, E\rangle \to \langle \texttt{do}\ s'\ \texttt{watching}\ ev, E'\rangle}\ (watch_1)$$

$$\langle \texttt{do}\ \texttt{nothing}\ \texttt{watching}\ ev, E\rangle \to \langle \texttt{nothing}, E\rangle \quad (watch_2)$$

$$\langle \texttt{loop}\ s, E\rangle \to \langle (s \upharpoonleft \texttt{cooperate}); \texttt{loop}\ s, E\rangle \quad (loop)$$

$$\frac{exp \rightsquigarrow n}{\langle \texttt{repeat}\ exp\ \texttt{do}\ s, E\rangle \to \langle \underbrace{s; \ldots ; s}_{n\ times}, E\rangle}\ (repeat)$$

$$\frac{exp \rightsquigarrow tt}{\langle \texttt{if}\ exp\ \texttt{then}\ s_1\ \texttt{else}\ s_2, E\rangle \to \langle s_1, E\rangle}\ (if_1) \qquad \frac{exp \rightsquigarrow ff}{\langle \texttt{if}\ exp\ \texttt{then}\ s_1\ \texttt{else}\ s_2, E\rangle \to \langle s_2, E\rangle}\ (if_2)$$

Small-step Transition Rules

Fig. 1. Operational Semantics of CRL

The small-step transition relation satisfies two simple properties.

Proposition 1. (Determinism)

Let $s \in$ Programs and $E \subseteq$ Events. Then:

$$s \neq \texttt{nothing} \quad \Rightarrow \quad \textit{either } \langle s, E\rangle \ddagger \ \textit{ or } \ \exists!\, s', E' \,.\, \langle s, E\rangle \rightarrow \langle s', E'\rangle$$

Proof. By inspecting the suspension and transition rules, it is immediate to see that at most one transition rule applies to each configuration $\langle s, E\rangle$.

Proposition 2. (Event persistence)

Let $s \in$ Programs and $E \subseteq$ Events. Then: $\quad \langle s, E\rangle \rightarrow \langle s', E'\rangle \ \Rightarrow \ E \subseteq E'$

Proof. Straightforward, since the only transition rule that changes the event environment E is the rule for `generate` ev, which adds the event ev to E.

We define now the notion of *instantaneous convergence*, which is at the basis of the reactivity property of CRL programs. Let us first introduce some notation.

The *timed multi-step transition relation* $\langle s, E\rangle \Rightarrow_n \langle s', E'\rangle$ is defined by:

$$\langle s, E\rangle \Rightarrow_0 \langle s, E\rangle$$

$$\langle s, E\rangle \rightarrow \langle s', E'\rangle \ \wedge \ \langle s', E'\rangle \Rightarrow_n \langle s'', E''\rangle \quad \Rightarrow \quad \langle s, E\rangle \Rightarrow_{n+1} \langle s'', E''\rangle$$

Then the *multi-step transition relation* $\langle s, E\rangle \Rightarrow \langle s', E'\rangle$ is given by:

$$\langle s, E\rangle \Rightarrow \langle s', E'\rangle \ \Leftrightarrow \ \exists\, n \,.\, \langle s, E\rangle \Rightarrow_n \langle s', E'\rangle$$

Note that the relation $\Rightarrow$ could also be defined as $\rightarrow^\star$.

The *immediate convergence* predicate is defined by:

$$\langle s, E\rangle \overset{\ddagger}{\curlyvee} \ \Leftrightarrow \ \langle s, E\rangle \ddagger \ \vee \ s = \texttt{nothing}$$

We may now define the relations and predicates of *instantaneous convergence* and *instantaneous termination*:

Definition 1. (Instantaneous convergence)

$$\begin{array}{lll} \langle s, E\rangle \Downarrow \langle s', E'\rangle & \text{if} & \langle s, E\rangle \Rightarrow \langle s', E'\rangle \ \wedge \ \langle s', E'\rangle \overset{\ddagger}{\curlyvee} \\ \langle s, E\rangle \Downarrow & \text{if} & \exists s', E' \,.\, \langle s, E\rangle \Downarrow \langle s', E'\rangle \end{array}$$

Definition 2. (Instantaneous termination)

$$\begin{array}{lll} \langle s, E\rangle \ \Downarrow\!\!\!\!\Downarrow \ E' & \text{if} & \langle s, E\rangle \Downarrow \langle \texttt{nothing}, E'\rangle \\ \langle s, E\rangle \ \Downarrow\!\!\!\!\Downarrow & \text{if} & \exists\, E' \,.\, \langle s, E\rangle \ \Downarrow\!\!\!\!\Downarrow \ E' \end{array}$$

The *timed* versions $\Downarrow_n$ and $\Downarrow\!\!\!\!\Downarrow_n$ of $\Downarrow$ and $\Downarrow\!\!\!\!\Downarrow$ are defined in the expected way.

The relation $\langle s, E\rangle \ \Downarrow \ \langle s', E'\rangle$ defines the overall effect of the program s within an instant, starting with the set of events E. Indeed, $\Downarrow$ may be viewed as defining the *big-step semantics* of programs within an instant[2]. As an immediate corollary of Proposition 1, the relation $\Downarrow$ is a function.

In the next section we prove that every configuration $\langle s, E\rangle$ instantaneously converges. This property is called *reactivity*.

[2] A direct definition of the big-step arrow $\Downarrow$ by a set of structural rules would be slightly more involved, as it would require calculating the output set E' as a fixpoint.

4 Reactivity

In this section we present our first main result, the reactivity of CRL programs. In fact, we shall prove a stronger property than reactivity, namely that every configuration $\langle s, E\rangle$ instantaneously converges in a number of steps which is bounded by the *instantaneous size* of s, denoted by $size(s)$. The intuition for $size(s)$ is that the portion of s that sequentially follows a `cooperate` instruction should not be taken into account, as it will not be executed in the current instant. Moreover, if s is a loop, $size(s)$ should cover a single iteration of its body.

To formally define the function $size(s)$, we first introduce an auxiliary function $dsize(s)$ (where"d" stands for "decorated") that assigns to each program an element of $(\mathbf{Nat} \times \mathbf{Bool})$. Then $size(s)$ will be the first projection of $dsize(s)$. Intuitively, if $dsize(s) = (n, b)$, then n is an upper bound for the number of steps that s can execute within an instant; and b is tt or ff depending on whether or not a `cooperate` instruction is reached within the instant. For conciseness, we let $n^\wedge$ stand for (n, tt), n stand for (n, ff), and n° range over $\{n^\wedge, n\}$.

The difference between $n^\wedge$ and n will essentially show when computing the size of a sequential composition: if the decorated size of the first component has the form $n^\wedge$, then a `cooperate` has been met and the counting will stop; if it has the form n, then n will be added to the decorated size of the second component.

Definition 3. (Instantaneous size)

The function size : Programs $\rightarrow$ **Nat** *is defined by:*

$$size(s) = n \quad \textit{if} \quad (dsize(s) = n \ \vee \ dsize(s) = n^\wedge).$$

where the function dsize : Programs $\rightarrow$ $(\mathbf{Nat} \times \mathbf{Bool})$ *is given inductively by:*

$$dsize(\texttt{nothing}) = 0 \qquad dsize(\texttt{cooperate}) = 0^\wedge$$

$$dsize(\texttt{generate}\ ev) = dsize(\texttt{await}\ ev) = 1$$

$$dsize(s_1; s_2) = \begin{cases} {n_1}^\wedge & \textit{if}\ dsize(s_1) = {n_1}^\wedge \\ (1 + n_1 + n_2)^\circ & \textit{if}\ dsize(s_1) = n_1 \ \wedge \ dsize(s_2) = {n_2}^\circ \end{cases}$$

$$dsize(s_1 \nmid s_2) = \begin{cases} (1 + n_1 + n_2)^\wedge & \textit{if}\ dsize(s_1) = {n_1}^\wedge \ \wedge \ dsize(s_2) = n_2 \\ (1 + n_1 + n_2)^\wedge & \textit{if}\ dsize(s_1) = n_1 \ \wedge \ dsize(s_2) = {n_2}^\wedge \\ (1 + n_1 + n_2)^\circ & \textit{if}\ dsize(s_1) = {n_1}^\circ \ \wedge \ dsize(s_2) = {n_2}^\circ \end{cases}$$

$$dsize(\texttt{repeat}\ exp\ \texttt{do}\ s) = (m + (m \times n))^\circ \qquad \textit{if}\ dsize(s) = n^\circ \ \wedge \ exp \rightsquigarrow m$$

$$dsize(\texttt{loop}\ s) = (2 + n)^\wedge \qquad \textit{if}\ dsize(s) = n^\circ$$

$$dsize(\texttt{do}\ s\ \texttt{watching}\ ev) = (1 + n)^\circ \quad \textit{if}\ dsize(s) = n^\circ$$

$$dsize(\texttt{if}\ exp\ \texttt{then}\ s_1\ \texttt{else}\ s_2) = \begin{cases} (1 + max\{n_1, n_2\})^\wedge, \textit{if}\ dsize(s_i) = {n_i}^\wedge, \\ (1 + max\{n_1, n_2\}), \textit{if}\ \ \textit{for}\ \ i \neq j \\ \qquad\qquad dsize(s_i) = n_i \ \wedge \ dsize(s_j) = {n_j}^\circ \end{cases}$$

The following lemma establishes that $size(s)$ decreases at each step of a small-step execution:

Lemma 1. (Size reduction within an instant)

$$\forall s \, \forall E \quad (\langle s, E \rangle \rightarrow \langle s', E' \rangle \;\Rightarrow\; size(s') < size(s))$$

The proof of this result is not entirely straightforward because of the use of the decorated size *dsize* in the definition of $size(s)$. The proof may be found in [5].

We are now ready to prove our main result, namely that every program s instantaneously converges in a number of steps that is bounded by $size(s)$.

Theorem 1. (Script reactivity) $\forall s, \forall E \quad (\exists \, n \leq size(s) \;\; \langle s, E \rangle \Downarrow_n)$

The proof proceeds by simultaneous induction on the structure and on the size of s. Induction on the size is needed for the case $s = s_1 \nmid s_2$. The detailed proof may be found in [5].

5 Fine-grained and coarse-grained bisimilarity

We now introduce two bisimulation equivalences (aka *bisimilarities*) on programs, which differ for the granularity of the underlying notion of observation. The first bisimulation formalises a *fine-grained observation* of programs: the observer is viewed as a program, which is able to interact with the observed program at any point of its execution. The second reflects a *coarse-grained observation* of programs: here the observer is viewed as part of the environment, which interacts with the observed program only at the start and the end of instants.

Let us start with an informal description of the two bisimilarities:

1. *Fine-grained bisimilarity* $\approx^{fg}$. In the bisimulation game, each small step must be simulated by a (possibly empty) sequence of small steps, and each instant change must be simulated either by an instant change, in case the continuation is observable (in the sense that it affects the event environment), or by an unobservable behaviour otherwise.
2. *Coarse-grained bisimilarity* $\approx^{cg}$. Here, each converging sequence of steps must be simulated by a converging sequence of steps, at each instant. For instant changes, the requirement is the same as for fine-grained bisimulation.

As may be expected, the latter equivalence is more abstract than the former, as it only compares the I/O behaviours of programs (as functions on sets of events) at each instant, while the former also compares their intermediate results.

Let us move now to the formal definitions of the equivalences $\approx^{fg}$ and $\approx^{cg}$.

We first extend the reconditioning function to the program `nothing` as follows:

Notation. $\llcorner s \lrcorner_E \stackrel{\text{def}}{=} \begin{cases} [s]_E & \text{if } \langle s, E \rangle \ddagger \\ s & \text{if } s = \texttt{nothing} \end{cases}$

Definition 4 (Fine-grained bisimulation).
A symmetric relation $\mathcal{R}$ on programs is a fg-bisimulation *if $s_1 \,\mathcal{R}\, s_2$ implies, for any $E \subseteq Events$:*

$$1)\ \langle s_1, E\rangle \rightarrow \langle s_1', E'\rangle \ \ \Rightarrow\ \ \exists\, s_2' \,.\ (\,\langle s_2, E\rangle \Rightarrow \langle s_2', E'\rangle \ \wedge\ s_1' \,\mathcal{R}\, s_2'\,)$$
$$2)\ \langle s_1, E\rangle\ddagger \ \ \Rightarrow\ \ \exists\, s_2' \,.\ (\,\langle s_2, E\rangle \Downarrow \langle s_2', E\rangle \ \wedge\ \llcorner s_1 \lrcorner_E \,\mathcal{R}\, \llcorner s_2' \lrcorner_E\,)$$

Then s_1, s_2 are fg-bisimilar*, $s_1 \approx^{fg} s_2$, if $s_1 \,\mathcal{R}\, s_2$ for some fg-bisimulation $\mathcal{R}$.*

The bisimilarity $\approx^{fg}$ is *weak*, in the terminology of process calculi, since it allows a single small step to be simulated by a (possibly empty) sequence of small steps. In the terminology of language-based security, an equivalence that abstracts away from the number of steps, thus allowing internal moves to be ignored, is called *time-insensitive*. Typically we have:

$$\texttt{nothing}\,;\,\texttt{generate}\,ev \ \approx^{fg}\ \texttt{generate}\,ev$$

$$\texttt{if}\ tt\ \texttt{then}\ s_1\ \texttt{else}\ s_2 \ \approx^{fg}\ s_1$$

The equivalence $\approx^{fg}$ is also *termination-insensitive*, as it cannot distinguish proper termination from suspension nor from internal divergence (recall that no divergence is possible within an instant and thus the execution of a diverging program always spans over an infinity of instants). For instance we have:

$$\texttt{nothing} \ \approx^{fg}\ \texttt{cooperate} \ \approx^{fg}\ \texttt{loop}\,\texttt{nothing}$$

Indeed, for any E the suspended behaviour $\langle \texttt{cooperate}, E\rangle\ddagger$ of the middle program can be simulated by the empty computation $\langle \texttt{nothing}, E\rangle \Downarrow \langle \texttt{nothing}, E\rangle$ of the left-hand program and by the two-step computation $\langle \texttt{loop}\,\texttt{nothing}, E\rangle \rightarrow \langle(\texttt{nothing} \upharpoonleft \texttt{cooperate})\,;\,\texttt{loop}\,\texttt{nothing}, E\rangle \rightarrow \langle \texttt{cooperate}\,;\,\texttt{loop}\,\texttt{nothing}, E\rangle\ddagger$ of the right-hand program, since $\llcorner \texttt{cooperate} \lrcorner_E = \texttt{nothing} = \llcorner \texttt{nothing} \lrcorner_E$, and $\llcorner \texttt{cooperate}\,;\,\texttt{loop}\,\texttt{nothing} \lrcorner_E = \texttt{loop}\,\texttt{nothing}$.

The last example shows that, while it weakly preserves small-step transitions, $\approx^{fg}$ does *not* preserve tick transitions. On the other hand, it detects the instant in which events are generated. In other words, it is sensitive to the *clock-stamp* of events. For instance, we have:

$$\texttt{nothing}\,;\,\texttt{generate}\,ev \ \not\approx^{fg}\ \texttt{cooperate}\,;\,\texttt{generate}\,ev$$

because in the left-hand program ev is generated in the first instant, while in the right-hand program it is generated in the second instant. Incidentally, this example shows that $\approx^{fg}$ is not preserved by sequential composition (as was to be expected given that $\approx^{fg}$ is termination-insensitive).

On the other hand, we conjecture that $\approx^{fg}$ is *compositional*, that is, preserved by parallel composition, because in the bisimulation game the quantification on the event environment is renewed at each step, thus mimicking the generation of events by a parallel component.

Finally, $\approx^{fg}$ is sensitive to the order of generation of events and to repeated emissions of the same event ("stuttering"). Typical examples are:

$$(\texttt{generate}\ ev_1 \upharpoonleft \texttt{generate}\ ev_2) \quad \not\approx^{fg} \quad (\texttt{generate}\ ev_2 \upharpoonleft \texttt{generate}\ ev_1)$$

$$\texttt{generate}\ ev \quad \not\approx^{fg} \quad (\texttt{generate}\ ev\,;\ \texttt{generate}\ ev)$$

In the last example, note that after generating the first event ev the right-hand program may be launched again in the event environment $E = \emptyset$, producing once more $E' = \{ev\}$. This cannot be mimicked by the left-hand program.

Definition 5 (Coarse-grained bisimulation).
A symmetric relation $\mathcal{R}$ on programs is a cg-bisimulation *if $s_1 \,\mathcal{R}\, s_2$ implies, for any $E \subseteq Events$:*

$$\langle s_1, E\rangle \Downarrow \langle s_1', E'\rangle \quad \Rightarrow \quad \exists\, s_2' \,.\ (\,\langle s_2, E\rangle \Downarrow \langle s_2', E'\rangle \ \wedge\ \llcorner s_1' \lrcorner_{E'} \ \mathcal{R}\ \llcorner s_2' \lrcorner_{E'}\,)$$

Then s_1, s_2 are cg-bisimilar*, $s_1 \approx^{cg} s_2$, if $s_1 \,\mathcal{R}\, s_2$ for some cg-bisimulation $\mathcal{R}$.*

The bisimilarity $\approx^{cg}$ compares the overall effect of two programs at every instant. Therefore, one may argue that $\approx^{cg}$ makes full sense when coupled with reactivity. Indeed, if $\approx^{cg}$ were applied to programs that diverge within the first instant (or to programs that are bisimilar for the first k instants and diverge in the following instant), it would trivially equate all of them. In the absence of reactivity, it would seem preferable to focus on a fine-grained bisimilarity such as $\approx^{fg}$, which is able to detect intermediate results of instantaneously diverging computations.

Like $\approx^{fg}$, the bisimilarity $\approx^{cg}$ is both time-insensitive and termination-insensitive. Indeed, as will be established by Theorem 2, $\approx^{fg}$ implies $\approx^{cg}$. Moreover, $\approx^{cg}$ is *generation-order-insensitive* and *stuttering-insensitive*. Typically:

$$(\texttt{generate}\ ev_1 \upharpoonleft \texttt{generate}\ ev_2) \quad \approx^{cg} \quad (\texttt{generate}\ ev_2 \upharpoonleft \texttt{generate}\ ev_1)$$

$$\texttt{generate}\ ev \quad \approx^{cg} \quad (\texttt{generate}\ ev\,;\ \texttt{generate}\ ev)$$

More generally, we can show that the equivalence $\approx^{cg}$ views the left-parallel composition $\upharpoonleft$ as a commutative operator:

Proposition 3. (Commutativity of $\upharpoonleft$ up to $\approx^{cg}$)

$$\forall s_1, s_2 \,.\quad s_1 \upharpoonleft s_2 \ \approx^{cg}\ s_2 \upharpoonleft s_1$$

On the other hand, $\upharpoonleft$ is associative modulo both equivalences $\approx^{fg}$ and $\approx^{cg}$:

Proposition 4. (Associativity of $\upharpoonleft$ up to $\approx^{fg}$ and $\approx^{cg}$)

$$\forall s_1, s_2, s_3 \,.\quad s_1 \upharpoonleft (s_2 \upharpoonleft s_3) \ \begin{matrix}\approx^{fg}\\ \approx^{cg}\end{matrix}\ (s_1 \upharpoonleft s_2) \upharpoonleft s_3$$

Let us recall that the asymmetric parallel operator $\upharpoonleft$ of [2] was not associative up to fine-grained semantics (a simple example was given in [2]).

We show now that $\approx^{fg}$ is strictly included in $\approx^{cg}$ (the strictness of the inclusion being witnessed by the examples given above):

Theorem 2. (Relation between the bisimilarities)

$$\approx^{fg} \subset \approx^{cg}$$

Proof. To prove $\approx^{fg} \subseteq \approx^{cg}$, it is enough to show that $\approx^{fg}$ is a cg-bisimulation. Let $s_1 \approx^{fg} s_2$. Suppose that $\langle s_1, E\rangle \Downarrow \langle s_1', E'\rangle$. This means that there exists $n \geq 0$ such that:

$$\langle s_1, E\rangle = \langle s_1^0, E^0\rangle \rightarrow \langle s_1^1, E^1\rangle \rightarrow \cdots \rightarrow \langle s_1^n, E^n\rangle = \langle s_1', E'\rangle \overset{\ddagger}{\underset{\Upsilon}{}}$$

Since $s_1 \approx^{fg} s_2$, by Clauses 1 and 2 of Definition 4 we have correspondingly:

$$\langle s_2, E\rangle = \langle s_2^0, E^0\rangle \Rightarrow \langle s_2^1, E^1\rangle \Rightarrow \cdots \Rightarrow \langle s_2^n, E^n\rangle \Downarrow \langle s_2', E'\rangle \qquad (*)$$

where $s_1^i \approx^{fg} s_2^i$ for every $i < n$ and $\llcorner s_1' \lrcorner_{E'} \approx^{fg} \llcorner s_2' \lrcorner_{E'}$. Then we may conclude since $(*)$ can be rewritten as $\langle s_2, E\rangle \Downarrow \langle s_2', E'\rangle$.

Coarse-grained bisimilarity is very close to the semantic equivalence proposed by Amadio in [4] for a slightly different reactive language, equipped with a classical nondeterministic parallel operator. By contrast, the noninterference notion of [2] was based on a fine-grained bisimilarity (although bisimilarity was not explicitly introduced in [2], it was *de facto* used to define noninterference) which, however, was stronger than $\approx^{fg}$, since it acted as a strong bisimulation on programs with an observable behaviour (i.e. affecting the event environment).

As argued previously, coarse-grained bisimilarity is a natural equivalence to adopt when reactivity is guaranteed. It allows one to recover the flavour of big-step semantics within instants. On the other hand, fine-grained bisimilarity seems a better choice when reactivity is not granted. Note that reactivity was not a concern in either [2] or [4]. Nevertheless, it had been thoroughly studied in previous work by Amadio et al. [3].

Finally, it should be noted that, since our left-parallel composition operator $\upharpoonleft$ is deterministic, we could as well have used trace-based equivalences rather than bisimulation-based ones. However, defining traces is not entirely obvious for computations proceeding through instants, as it requires annotating with clock-stamps the events or event sets that compose a trace (depending on whether the trace is fine-grained or coarse-grained). Moreover, bisimulation provides a convenient means for defining noninterference in our concurrent setting, allowing the notion of clock-stamp to remain implicit. Lastly, as we aim to extend our study to a fully-fledged distributed reactive language, including a notion of site and asynchronous parallelism between sites, for which determinism would not hold anymore, we chose to adopt bisimulation-based equivalences from the start.

This concludes our discussion on semantic equivalences. We turn now to the definition of noninterference, which is grounded on that of bisimulation.

6 Security property

In this section we define two noninterference properties for programs, which are based on the two bisimilarities introduced in Section 5. As usual when dealing with secure information flow, we assume a finite lattice $(\mathcal{S}, \leq)$ of *security levels*, ranged over by τ, σ, ϑ. We denote by $\sqcup$ and $\sqcap$ the join and meet operations on the lattice, and by $\bot$ and $\top$ its minimal and maximal elements.

In CRL, the objects that are assigned a security level are events and variables. An *observer* is identified with a downward-closed set of security levels (for short, a dc-set), i.e. a set $\mathcal{L} \subseteq \mathcal{S}$ satisfying the property: $(\tau \in \mathcal{L} \ \wedge \ \tau' \leq \tau) \ \Rightarrow \ \tau' \in \mathcal{L}$.

A type environment Γ is a mapping from variables and events to their types, which are just security levels τ, σ. Given a dc-set $\mathcal{L}$, a type environment Γ and an event environment E, the subset of E to which Γ assigns security levels in $\mathcal{L}$ is called the *$\mathcal{L}$-part* of E under Γ. Similarly, if $V : Var \to Val$ is a valuation, the subset of V whose domain is given levels in $\mathcal{L}$ by Γ is the *$\mathcal{L}$-part* of V under Γ.

Two event environments E_1, E_2 or two valuations V_1, V_2 are $=^{\Gamma}_{\mathcal{L}}$-equal, or indistinguishable by a $\mathcal{L}$-observer, if their $\mathcal{L}$-parts under Γ coincide:

Definition 6 ($\Gamma\mathcal{L}$-equality of event environments and valuations).
Let $\mathcal{L} \subseteq \mathcal{S}$ be a dc-set, Γ a type environment and V a valuation. Define:

$$E_1 =^{\Gamma}_{\mathcal{L}} E_2 \quad \textit{if} \quad \forall\, ev \in Events \ \ (\,\Gamma(ev) \in \mathcal{L} \ \Rightarrow \ (\,ev \in E_1 \Leftrightarrow ev \in E_2\,)\,)$$
$$V_1 =^{\Gamma}_{\mathcal{L}} V_2 \quad \textit{if} \quad \forall\, x \in Var \ \ (\,\Gamma(x) \in \mathcal{L} \ \Rightarrow \ V_1(x) = V_2(x)\,)$$

Let $\rightarrow_V, \Rightarrow_V, \Downarrow_V$ denote our various semantic arrows under the valuation V. Then we may define the indistinguishability of two programs by a fine-grained or coarse-grained $\mathcal{L}$-observer, for a given Γ, by means of the following two notions of $\Gamma\mathcal{L}$-bisimilarity:

Definition 7 (Fine-grained $\Gamma\mathcal{L}$-bisimilarity).
A relation $\mathcal{R}$ on programs is a fg-$\Gamma\mathcal{L}$-V_1V_2-bisimulation *if $s_1 \,\mathcal{R}\, s_2$ implies, for any E_1, E_2 such that $E_1 =^{\Gamma}_{\mathcal{L}} E_2$:*

1) $\langle s_1, E_1\rangle \rightarrow_{V_1} \langle s_1', E_1'\rangle \ \Rightarrow \ \exists\, s_2', E_2' \,.\, (\,\langle s_2, E_2\rangle \Rightarrow_{V_2} \langle s_2', E_2'\rangle \ \wedge \ E_1' =^{\Gamma}_{\mathcal{L}} E_2' \ \wedge \ s_1' \,\mathcal{R}\, s_2'\,)$

2) $\langle s_1, E_1\rangle\ddagger \ \ \Rightarrow \exists\, s_2', E_2' \,.\, (\,\langle s_2, E_2\rangle \Downarrow_{V_2} \langle s_2', E_2'\rangle \ \wedge \ E_1 =^{\Gamma}_{\mathcal{L}} E_2' \ \wedge \ \llcorner s_1 \lrcorner_{E_1} \ \mathcal{R} \ \llcorner s_2' \lrcorner_{E_2'}\,)$

3) and 4) : *Symmetric to* 1) *and* 2) *for $\langle s_2, E_2\rangle$ under valuation V_2.*

Then programs s_1, s_2 are fg-$\Gamma\mathcal{L}$-bisimilar*, $s_1 \approx^{fg}_{\Gamma\mathcal{L}} s_2$, if for any V_1, V_2 such that $V_1 =^{\Gamma}_{\mathcal{L}} V_2$, $s_1 \,\mathcal{R}\, s_2$ for some* fg-$\Gamma\mathcal{L}$-V_1V_2-bisimulation $\mathcal{R}$.

The *fg*-$\Gamma\mathcal{L}$-bisimilarity weakly preserves small-step transitions and convergence, while maintaining the $\Gamma\mathcal{L}$-equality on event environments. Note that, while in the definition of *fg*-bisimilarity it was possible to leave the valuation implicit, we need to make it explicit in the definition of *fg*-$\Gamma\mathcal{L}$-bisimilarity, as variables have security levels and are allowed to have different values if their level is not in $\mathcal{L}$.

The reason why a *fg*-$\Gamma\mathcal{L}$-V_1V_2-bisimulation is parameterised on two valuations V_1 and V_2, and the quantification on valuations in $\approx^{fg}_{\Gamma\mathcal{L}}$ is only performed at the beginning of the bisimulation game, rather than at each step as for event environments, is that programs have no means to change the valuation. In a more expressive language where the valuation could change, it would be necessary to include the valuation in the environment that is quantified at each step.

Definition 8 (Coarse-grained $\Gamma\mathcal{L}$-bisimilarity).
A relation $\mathcal{R}$ on programs is a cg-$\Gamma\mathcal{L}$-V_1V_2-bisimulation *if $s_1 \,\mathcal{R}\, s_2$ implies, for any E_1, E_2 such that $E_1 =^{\Gamma}_{\mathcal{L}} E_2$:*

1) $\langle s_1, E_1\rangle \Downarrow_{V_1} \langle s'_1, E'_1\rangle \;\Rightarrow\; \exists\, s'_2, E'_2 \,.\, (\; \langle s_2, E_2\rangle \Downarrow_{V_2} \langle s'_2, E'_2\rangle \;\wedge\; E'_1 =^{\Gamma}_{\mathcal{L}} E'_2 \;\wedge\; \llcorner s'_1 \lrcorner_{E'_1} \;\mathcal{R}\; \llcorner s'_2 \lrcorner_{E'_2} \;)$

2) *Symmetric to* 1) *for $\langle s_2, E_2\rangle$ under valuation V_2.*

Two programs s_1, s_2 are cg-$\Gamma\mathcal{L}$-bisimilar*, $s_1 \approx^{cg}_{\Gamma\mathcal{L}} s_2$, if for any V_1, V_2 such that $V_1 =^{\Gamma}_{\mathcal{L}} V_2$, $s_1 \,\mathcal{R}\, s_2$ for some* cg-$\Gamma\mathcal{L}$-V_1V_2-bisimulation $\mathcal{R}$.

Our *reactive noninterference* (RNI) properties are now defined as follows:

Definition 9 (Fine-grained and Coarse-grained RNI).
A program s is fg-secure in Γ if $s \approx^{fg}_{\Gamma\mathcal{L}} s$ for every dc-set $\mathcal{L}$.
A program s is cg-secure in Γ if $s \approx^{cg}_{\Gamma\mathcal{L}} s$ for every dc-set $\mathcal{L}$.

The following example, where the superscripts indicate the security levels of variables and events, illustrates the difference between *fg*-security and *cg*-security.

Example 1. The following program is *cg*-secure but not *fg*-secure:

$$s = \texttt{if}\ x^{\top} = 0\ \texttt{then}\ \texttt{generate}\ ev_1^{\bot} \nmid \texttt{generate}\ ev_2^{\bot}$$
$$\texttt{else}\ \texttt{generate}\ ev_2^{\bot} \nmid \texttt{generate}\ ev_1^{\bot}$$

If we replace the second branch of s by $\texttt{generate}\ ev_1^{\bot}\ ;\ \texttt{generate}\ ev_2^{\bot}$, then we obtain a program s' that is both *fg*-secure and *cg*-secure.

In general, from all the equivalences/inequivalences in page 11 we may obtain secure/insecure programs for the corresponding RNI property by plugging the two equivalent/inequivalent programs in the branches of a high conditional.

As expected, fine-grained security is stronger than coarse-grained security:

Theorem 3. (Relation between the RNI properties)

Let $s \in$ Programs. If s is fg-secure then s is cg-secure.

Proof. The proof consists in showing that for any Γ and $\mathcal{L}$, we have $\approx^{fg}_{\Gamma\mathcal{L}} \subseteq \approx^{cg}_{\Gamma\mathcal{L}}$. To this end, it is enough to show that for any pair of valuations V_1 and V_2, any fg-$\Gamma\mathcal{L}$-V_1V_2-bisimulation $\mathcal{R}$ is also a cg-$\Gamma\mathcal{L}$-V_1V_2-bisimulation. The reasoning closely follows that of Theorem 2 and is therefore omitted.

We conclude this section with an informal discussion about *scheduling leaks*. We speak of scheduling leak when the position of the scheduler at the start of an instant may depend on secrets tested in previous instants. We have mentioned already that, unlike the "swapping" operator $\upharpoonleft$ of [2], our operator $\nmid$ preserves the spatial structure of programs. As a consequence, the same parallel component is scheduled at the beginning of each instant, and the position of the scheduler is independent of any previous test. Thus the scheduling leaks arising with the operator $\upharpoonleft$, which implied a severe constraint in the type system of [2] (the addition of the condition $\sigma_i \leq \tau_j$ in Rule (PAR)), cannot occur anymore with $\nmid$. In particular, it may be shown that the scheduling leak example given in [2] does not arise if we replace $\upharpoonleft$ by $\nmid$. This point is explained in detail in [5].

7 Type system

We present now our security type system for CRL, which is based on those introduced in [8] and [17] for a parallel while language and already adapted to a reactive language in [2]. The originality of these type systems is that they associate pairs (τ, σ) of security levels with programs, where τ is a lower bound on the level of "writes" and σ is an upper bound on the level of "reads". This allows the level of reads to be recorded, and then to be used to constrain the level of writes in the remainder of the program. In this way, it is possible to obtain a more precise treatment of *termination leaks*[3] than in standard type systems.

Recall that a type environment Γ is a mapping from variables and events to security levels τ, σ. Moreover, Γ associates a type of the form $\overrightarrow{\tau} \rightarrow \tau$ to functions, where $\overrightarrow{\tau}$ is a tuple of types $\tau_1, \ldots, \tau_n$. Type judgements for expressions and programs have the form $\Gamma \vdash exp : \tau$ and $\Gamma \vdash s : (\tau, \sigma)$ respectively.

The intuition for $\Gamma \vdash exp : \tau$ is that τ is an *upper bound* on the levels of variables occurring in *exp*. According to this intuition, subtyping for expressions is *covariant*. The intuition for $\Gamma \vdash s : (\tau, \sigma)$ is that τ is a *lower bound* on the levels of events generated in s (the "writes"of s), and σ is an *upper bound* on the levels of events awaited or watched in s and of variables tested in s (the "reads" or *guards* of s, formally defined in Definition 11). Accordingly, subtyping for programs is *contravariant* in its first component, and *covariant* in the second.

The typing rules for expressions and programs are presented in Figure 7. The three rules that increase the guard type are (WATCHING), (REPEAT) and (COND1), and those that check it against the write type of the continuation are (SEQ), (REPEAT) and (LOOP). Note that there are two more rules for the conditional, factoring out the cases where either both branches terminate in one instant or both branches are infinite: indeed, in these cases no termination leaks can arise and thus it is not necessary to increase the guard level. In Rule (COND2), *FIN* denotes the set of programs *terminating in one instant*, namely

[3] Leaks due to different termination behaviours in the branches of a conditional. In classical parallel while languages, termination leaks may also arise in while loops. This is not possible in CRL, given the simple form of the `loop` construct. On the other hand, new termination leaks may originate from the possibility of suspension.

those built without using the constructs `await` ev, `cooperate` and `loop`. In Rule (Cond3), INF denotes the set of infinite or *nonterminating* programs, defined inductively as follows[4]:

- `loop` $s \in INF$;
- $s \in INF \Rightarrow$ `repeat` exp `do` $s \in INF$;
- $s_1 \in INF \ \vee \ s_2 \in INF \Rightarrow s_1; s_2 \in INF \ \wedge \ s_1 \nmid s_2 \in INF$
- $s_1 \in INF \ \wedge \ s_2 \in INF \Rightarrow$ `if` exp `then` s_1 `else` $s_2 \in INF$

Note that $FIN \cup INF \subset Programs$. Examples of programs that are neither in FIN nor in INF are: `await` ev, `if` exp `then` `nothing` `else` (`loop` s), and `do` (`loop` s) `watching` ev.

Definition 10 (Safe conditionals).

A conditional `if` exp `then` s_1 `else` s_2 *is* safe *if* $s_1, s_2 \in FIN$ *or* $s_1, s_2 \in INF$.

The reason for calling such conditionals "safe" is that they cannot introduce termination leaks, since their two branches have the same termination behaviour.

Note that the two sets FIN and INF only capture two specific *termination behaviours* of CRL programs, namely termination in one instant and nontermination. We could have refined further this classification of termination behaviours. Indeed, while only two termination behaviours are possible within each instant, due to reactivity (namely, proper termination and suspension), across instants there is an infinity of possible termination behaviours for programs: termination in a finite number k of instants, for any possible k, and nontermination. In other words, we could have defined a set FIN_k for each k and replaced Rule (Cond2) by a Rule Schema (Condk). The idea would remain the same: conditionals with uniform termination behaviours need not raise their guard level. For simplicity, we chose to focus on FIN and INF, leaving the finer analysis for future work.

We now prove that typability implies security via the classical steps:

Lemma 2 (Subject Reduction).

Let $\Gamma \vdash s : (\tau, \sigma)$. *Then* $\langle s, E\rangle \rightarrow \langle s', E'\rangle$ *implies* $\Gamma \vdash s' : (\tau, \sigma)$, *and* $\langle s, E\rangle\ddagger$ *implies* $\Gamma \vdash [s]_E : (\tau, \sigma)$.

Definition 11. (Guards and Generated Events)

1) For any s, *Guards*(s) *is the union of the set of events* ev *such that* s *contains an* `await` ev *or a* `do` s' `watching` ev *instruction (for some* s'*), together with the set of variables* x *that occur in* s *as argument of a function or in the control expression* exp *of an instruction* `repeat` exp `do` s' *or of an* unsafe *conditional* `if` exp `then` s_1 `else` s_2 *in* s.

2) For any s, *Gen*(s) *is the set of events* ev *such that* `generate` ev *occurs in* s.

The following Lemma establishes that if $\Gamma \vdash s : (\tau, \sigma)$, then τ is a *lower bound* on the levels of events in *Gen*(s) and σ is an *upper bound* on the levels of events and variables in *Guards*(s).

[4] Recall that in a `repeat` exp `do` s program, exp is supposed to evaluate to some $n \geq 1$.

$$(\textsc{Val})\quad \Gamma \vdash v : \bot \qquad (\textsc{Var})\quad \frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau} \qquad (\textsc{SubExp})\quad \frac{\Gamma \vdash exp : \sigma, \quad \sigma \leq \sigma'}{\Gamma \vdash exp : \sigma'}$$

$$(\textsc{Fun})\quad \frac{\Gamma \vdash \overrightarrow{exp} : \overrightarrow{\tau}, \quad \Gamma(f) = \overrightarrow{\tau} \to \tau, \quad \forall i \,.\, \tau_i \leq \tau}{\Gamma \vdash f(\overrightarrow{exp}) : \tau}$$

Typing rules for expressions

$$(\textsc{Nothing})\quad \Gamma \vdash \texttt{nothing} : (\top, \bot) \qquad (\textsc{Cooperate})\quad \Gamma \vdash \texttt{cooperate} : (\top, \bot)$$

$$(\textsc{Seq})\quad \frac{\Gamma \vdash s_1 : (\tau_1, \sigma_1), \quad \Gamma \vdash s_2 : (\tau_2, \sigma_2), \quad \sigma_1 \leq \tau_2}{\Gamma \vdash s_1 \,;\, s_2 : (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2)}$$

$$(\textsc{Par})\quad \frac{\Gamma \vdash s_1 : (\tau_1, \sigma_1), \quad \Gamma \vdash s_2 : (\tau_2, \sigma_2)}{\Gamma \vdash s_1 \nmid s_2 : (\tau_1 \sqcap \tau_2, \sigma_1 \sqcup \sigma_2)}$$

$$(\textsc{Generate})\quad \frac{\Gamma(ev) = \tau}{\Gamma \vdash \texttt{generate}\, ev : (\tau, \bot)} \qquad (\textsc{Await})\quad \frac{\Gamma(ev) = \sigma}{\Gamma \vdash \texttt{await}\, ev : (\top, \sigma)}$$

$$(\textsc{Watching})\quad \frac{\Gamma(ev) = \vartheta, \quad \Gamma \vdash s : (\tau, \sigma), \quad \vartheta \leq \tau}{\Gamma \vdash \texttt{do}\, s\, \texttt{watching}\, ev : (\tau, \vartheta \sqcup \sigma)}$$

$$(\textsc{Loop})\quad \frac{\Gamma \vdash s : (\tau, \sigma), \quad \sigma \leq \tau}{\Gamma \vdash \texttt{loop}\, s : (\tau, \sigma)} \qquad (\textsc{Repeat})\quad \frac{\Gamma \vdash exp : \vartheta \quad \Gamma \vdash s : (\tau, \sigma), \quad \vartheta \sqcup \sigma \leq \tau}{\Gamma \vdash \texttt{repeat}\, exp\, \texttt{do}\, s : (\tau, \vartheta \sqcup \sigma)}$$

$$(\textsc{Cond1})\quad \frac{\Gamma \vdash exp : \vartheta, \quad \Gamma \vdash s_i : (\tau, \sigma), \qquad i = 1, 2, \quad \vartheta \leq \tau}{\Gamma \vdash \texttt{if}\, exp\, \texttt{then}\, s_1\, \texttt{else}\, s_2 : (\tau, \vartheta \sqcup \sigma)}$$

$$(\textsc{Cond2})\quad \frac{\Gamma \vdash exp : \vartheta, \quad (\Gamma \vdash s_i : (\tau, \sigma) \quad \wedge \quad s_i \in FIN, \quad i = 1, 2\,), \quad \vartheta \leq \tau}{\Gamma \vdash \texttt{if}\, exp\, \texttt{then}\, s_1\, \texttt{else}\, s_2 : (\tau, \sigma)}$$

$$(\textsc{Cond3})\quad \frac{\Gamma \vdash exp : \vartheta, \quad (\Gamma \vdash s_i : (\tau, \sigma) \quad \wedge \quad s_i \in INF, \quad i = 1, 2\,), \quad \vartheta \leq \tau}{\Gamma \vdash \texttt{if}\, exp\, \texttt{then}\, s_1\, \texttt{else}\, s_2 : (\tau, \sigma)}$$

$$(\textsc{SubProg})\quad \frac{\Gamma \vdash s : (\tau, \sigma), \quad \tau' \leq \tau, \quad \sigma \leq \sigma'}{\Gamma \vdash s : (\tau', \sigma')}$$

Typing rules for programs

Fig. 2. Security type system

Lemma 3 (Guard Safety and Confinement).

1. *If* $\Gamma \vdash s : (\tau, \sigma)$ *then* $\Gamma(g) \leq \sigma$ *for every* $g \in Guards(s)$*;*
2. *If* $\Gamma \vdash s : (\tau, \sigma)$ *then* $\tau \leq \Gamma(ev)$ *for every* $ev \in Gen(s)$*.*

We now state the main result of this section, namely the soundness of the type system for fine-grained reactive noninterference (and thus, by Theorem 3, also for coarse-grained reactive noninterference). The proof involves some additional definitions and preliminary results, which are not given here but reported in [5].

Theorem 4 (Typability ⇒ Fine-grained RNI).

Let $s \in Programs$*. If* s *is typable in* Γ *then* s *is fg-secure in* Γ*.*

Note that programs s, s' of Example 1 are not typable (although *cg*-secure).

Example 2. The following programs are not typable and not secure, for any of the two security properties:

$$\texttt{await}\ ev_1^{\top}\ ;\ \texttt{generate}\ ev_2^{\bot}$$
$$\texttt{loop}\ (\texttt{generate}\ ev_2^{\bot}\ ;\ \texttt{await}\ ev_1^{\top})$$
$$\texttt{repeat}\ x^{\top}\ \texttt{do}\ \texttt{generate}\ ev^{\bot}$$
$$(\texttt{repeat}\ x^{\top}\ \texttt{do}\ \texttt{cooperate})\ ;\ \texttt{generate}\ ev^{\bot}$$
$$\texttt{do}\ (\texttt{cooperate}\ ;\ \texttt{generate}\ ev_2^{\bot})\ \texttt{watching}\ ev_1^{\top}$$

The insecure flows in the first two programs are *termination leaks*, due to the possibility of suspension. The second program illustrates the need for the condition $\sigma \leq \tau$ in Rule (Loop) (to produce a similar example with `repeat` we need at least three security levels). The fourth program shows why the guard level of `repeat` should be raised in Rule (Repeat).

We conclude with some examples illustrating the use of the conditional rules.

Example 3. The following programs s_i and s are all typable, with the given type:

$s_1 = \texttt{if}\ (x^{\top} = 0)\ \texttt{then}\ \texttt{await}\ ev_1^{\top}\ \texttt{else}\ \texttt{cooperate}$	type $(\top, \top)$
$s_2 = \texttt{if}\ (x^{\top} = 0)\ \texttt{then}\ \texttt{nothing}\ \texttt{else}\ \texttt{generate}\ ev^{\top}$	type $(\top, \bot)$
$s_3 = \texttt{if}\ (x^{\top} = 0)\ \texttt{then}\ \texttt{nothing}\ \texttt{else}\ (\texttt{loop}\ \texttt{nothing})$	type $(\top, \top)$
$s_4 = \texttt{if}\ (x^{\top} = 0)\ \texttt{then}\ (\texttt{loop}\ \texttt{nothing})\ \texttt{else}\ (\texttt{loop}\ \texttt{cooperate})$	type $(\top, \bot)$
$s = \texttt{generate}\ ev_2^{\bot}$	type $(\bot, \bot)$

Indeed, for all programs s_i the first component of the type (the write type) must be $\top$ because each of the Rules (Cond1) (Cond2) and (Cond3) prevents a "level drop" from the tested expression to the branches of the conditional, as in classical security type systems. On the other hand, the second component of the type (the guard type) will be $\bot$ for the safe conditionals s_2 and s_4, typed respectively using Rules (Cond2) and (Cond3), and $\top$ for the unsafe conditionals s_1 and s_3, typed using Rule (Cond1).

Then $s_2; s$ and $s_4; s$ are typable but not $s_1; s$ nor $s_3; s$.

8 Conclusion and related work

We have studied a core reactive language CRL and established a reactivity result for it, similar to those of [10,3] but based on different design choices. We also provided a syntactic bound for the length of the converging sequences.

We then proposed two RNI properties for the language, together with a security type system ensuring them. Our RNI properties rely on two bisimulation equivalences of different granularity. One of them, coarse-grained bisimilarity, is reminiscent of the semantic equivalence studied by Amadio in [4], which however was based on trace semantics. Our RNI properties also bear some analogy with the notions of *reactive noninterference* proposed in [7], and particularly with the termination-insensitive notion of ID-security (see also [15]), although the underlying assumptions of the model are quite different.

The model of cooperative threads of [1] is close in spirit to the model of CRL, but it is not concerned with synchronous parallelism. We should stress here that, to be appropriate for the study of a global computing setting, our synchronous model is intended to be part of a more general GALS model (Globally Asynchronous, Locally Synchronous), where various "synchronous areas" coexist and evolve in parallel, interacting with each other in an asynchronous way.

The idea of "slowing down" loops by forcing them to yield the control at each iteration, which is crucial for our reactivity result, was already used in [10] for a similar purpose. A similar instrumentation of loops was proposed in [14]. However, while in our work and in [10] a `cooperate` instruction is added *in parallel* with each iteration of the body of the loop, in [14] it is added *after* each iteration of the body. In a language that allows a parallel program to be followed in sequence by another program (which is not the case in [14]), our solution is more efficient in that it avoids introducing an additional suspension in case the body of the loop already contains one.

As regards future work, we expect some of our results - determinism, reactivity - to carry over smoothly to CRL extended with memory. However, some other properties like the commutativity of $\upharpoonright$ will not hold anymore in such setting, at least if the memory is freely shared among threads. Nevertheless, our bisimilarities and security properties would continue to make sense in such extended language. In the longer run, we plan to extend our study to a fully-fledged distributed reactive language, where programs are executed on different sites and may migrate from one site to the other. In this setting, execution would still be synchronous and reactive within each site (each site would be a "synchronous area" within a GALS model), but it would be asynchronous among different sites. A migrating thread would be integrated in the destination site only when this would become ready to react to its environment (whence the importance of local reactivity in each site). In a more expressive language with I/O blocking operations or other forms of abnormal termination, the enforcement of the reactivity property as well as the treatment of termination channels is likely to become more complex (although the time-out mechanism provided by the watching statement could be of some help here).

Acknowledgments

We thank Frédéric Boussinot for insightful discussions and feedback, and Bernard Serpette for useful comments on a previous version of this paper. We also thank the anonymous referees for helpful remarks and suggestions.

References

1. M. Abadi and G. Plotkin. A model of cooperative threads. In *Proceedings POPL 2009*, pages 29–40. ACM Press, 2009.
2. A. Almeida Matos, G. Boudol, and I. Castellani. Typing noninterference for reactive programs. *Journal of Logic and Algebraic Programming*, 72(2):124–156, 2007.
3. R. M Amadio and F. Dabrowski. Feasible reactivity for synchronous cooperative threads. *Electronic Notes in Theoretical Computer Science*, 154(3):33–43, 2006.
4. Roberto M Amadio. The SL synchronous language, revisited. *The Journal of Logic and Algebraic Programming*, 70(2):121–150, 2007.
5. P. Attar and I. Castellani. Fine-grained and coarse-grained reactive noninterference. INRIA Research Report, 2013.
6. G. Berry and G. Gonthier. The ESTEREL synchronous programming language: design, semantics, implementation. *Sci. of Comput. Programming*, 19, 1992.
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*, pages 79–90. ACM, 2009.
8. G. Boudol and I. Castellani. Noninterference for concurrent programs and thread systems. *Theoretical Computer Science*, 281(1):109–130, 2002.
9. F. Boussinot and R. de Simone. The SL synchronous language. *Software Engineering*, 22(4):256–266, 1996.
10. F. Boussinot and J-F. Susini. The SugarCubes Tool Box: A Reactive Java Framework. *Software: Practice and Experience*, 28(14):1531–1550, 1998.
11. R. Focardi and R. Gorrieri. Classification of security properties (part i: Information flow). In R. Focardi and R. Gorrieri, editors, *Foundations of Security Analysis and Design - Tutorial Lectures*, number 2171 in LNCS, 2001.
12. J. A. Goguen and J. Meseguer. Security policies and security models. In *Proceedings 1982 IEEE Symposium on Security and Privacy*, pages 11–20, 1982.
13. J. A. Goguen and J. Meseguer. Unwinding and inference control. In *Proceedings 1984 IEEE Symposium on Security and Privacy*, 1984.
14. A. Russo and A. Sabelfeld. Security for multithreaded programs under cooperative scheduling. In *Ershov Memorial Conference*, volume 4378 of *Lecture Notes in Computer Science*, pages 474–480. Springer, 2007.
15. A. Russo, D. Zanarini, and M. Jaskelioff. Precise enforcement of confidentiality for reactive systems. In *Proceedings of the 26th IEEE Computer Security Foundations Symposium*. IEEE, 2013.
16. A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.
17. G. Smith. A new type system for secure information flow. In *Proceedings of the 14th IEEE Computer Security Foundations Workshop*. IEEE, 2001.
18. D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):167–187, 1996.