A Calculus for Esterel

If can, can. If no can, no can.

SPENCER P. FLORENCE, Northwestern University, U.S.A.

SHU-HUNG YOU, Northwestern University, U.S.A.

JESSE A. TOV, Northwestern University, U.S.A.

ROBERT BRUCE FINDLER, Northwestern University, U.S.A.

The language Esterel has found success in many safety-critical applications, such as fly-by-wire systems and nuclear power plant control software. Its imperative style is natural to programmers building such systems and its precise semantics makes it work well for reasoning about programs.

Existing semantics of Esterel generally fall into two categories: translation to Boolean circuits, or operational semantics that give a procedure for running a whole program. In contrast, equational theories enable reasoning about program behavior via equational rewrites at the source level. Such theories form the basis for proofs of transformations inside compilers or for program refactorings, and defining program evaluation syntactically.

This paper presents the first such equational calculus for Esterel. It also illustrates the calculus's usefulness with a series of example equivalences and discuss how it enabled us to find bugs in Esterel implementations.

CCS Concepts: • **Theory of computation** → *Operational semantics*;

Additional Key Words and Phrases: Esterel, Synchronous Reactive Programming, Semantics

ACM Reference Format:

Spencer P. Florence, Shu-Hung You, Jesse A. Tov, and Robert Bruce Findler. 2019. A Calculus for Esterel: If can, can. If no can, no can.. *Proc. ACM Program. Lang.* 3, POPL, Article 61 (January 2019), 29 pages. <https://doi.org/10.1145/3290374>

1 INTRODUCTION

The language Esterel has found success in many safety-critical applications. It has been used in the creation and verification of the maintenance and test computer, landing gear control computer, and virtual display systems of civilian and military aircraft at Dassault Aviation (Berry et al. 2000); the control software of the N4 nuclear power plants; the Airbus A320 fly-by-wire system; and the specification of part of Texas Instrument's digital signal processors (Benveniste et al. 2002).

This success with real time and embedded systems in domains that need strong guarantees can be partially attributed to its computational model. Esterel treats computation as a series of deterministic reactions to external stimuli. All parts of a reaction complete in a single, discrete time step, called an *instant*. Furthermore, in this synchronous reactive paradigm (Benveniste and Berry 1991; Benveniste et al. 2002), each instant is isolated from interference by the outside environment

Authors' addresses: Spencer P. Florence, Northwestern University, U.S.A., spencer.florence@eecs.northwestern.edu; Shu-Hung You, Northwestern University, U.S.A., shu-hung.you@eecs.northwestern.edu; Jesse A. Tov, Northwestern University, U.S.A., jesse@eecs.northwestern.edu; Robert Bruce Findler, Northwestern University, U.S.A., robby@eecs.northwestern.edu.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/1-ART61

<https://doi.org/10.1145/3290374>

once the reaction begins. In addition, instants exhibit deterministic concurrency; each reaction may contain concurrent threads without execution order affecting the result of the computation.

This combination of synchronous reactions with deterministic concurrency makes formulating the semantics a challenging problem. Existing semantics tend to take two forms. The first, and most widely used, are semantics that give meaning to programs through a translation to circuits. These semantics are excellent for compilation and optimization. However they are not ideal for programmers, who would rather reason in terms of the source program, not its compiled form.

The second form are operational semantics that eschew term rewriting in favor of decorating terms with various flavors of code pointers and state annotations to track execution. These semantics are easier for programmers to reason with but give meaning only to whole programs. They do not lend themselves to compositional reasoning about program fragments, which programmers need.

To obtain the best of both of these approaches, we build on Plotkin (1975) and Felleisen and Hieb (1992)'s work on equational theories of programming languages. These theories model languages with a set of axioms that specify when source-level terms are equivalent. As a result, they provide a single framework for both reasoning about how a program will run (e.g. reduce to an answer) using only the source text of the program, and for justifying program transformations in host of applications: compiler transformations, refactorings, program derivations, etc.

This paper reports on the first equational theory of Kernel Esterel (Berry 2002). Developing such a theory is tricky because of the highly non-local nature of evaluation in Esterel. To maintain determinism and synchrony, evaluation in one thread of execution may affect code arbitrarily far away away in the program, even if that evaluation does not directly modify shared state. For instance, the selection of a particular branch of execution in one thread may immediately unblock a different thread of execution. The selection of the other branch may render the entire program invalid. These non-local execution and correctness issues are at the heart of Esterel's notions of *Logical Correctness* and *Constructiveness*, and have informed the choice of techniques used for previous semantics. The circuit semantics match both notions well because they are intimately tied to whether or not a given cyclic circuit settles. The operational semantics handle these properties by performing full program passes on each evaluation step to both propagate execution information to the entire program, and determine which locations in the program are safe to evaluate. A calculus, however, cannot use either of those techniques. To this end our calculus borrows from Felleisen and Hieb (1992)'s equational theory of state and Potop-Butucaru (2002)'s Constructive Operational Semantics to give the first calculus for Esterel.

The remainder of this paper consists of seven sections. Section 2 provides an introduction to Esterel and to the specific syntax we use for Kernel Esterel. Section 3 explains the semantics and our central results, which have all been checked in Agda. With the semantics defined, the paper moves on to discuss implications of specific aspects of our semantics. Section 4 discusses constructiveness and how it interacts with our semantics. Section 5 gives some example equivalences that our calculus supports and discusses others that it does not. Our semantics is executable and section 6 discusses how we test that our semantics is faithful to preexisting semantics and implementations. In short, we designed and implemented an executable version of our semantics and used it to find bugs in Esterel implementations. We also automatically typeset the figures in the paper from the semantics and use it to test all of the examples in the paper. Section 7 discusses a standard reduction that we designed to aid in testing but have not proven, and we conclude with a discussion of related work in section 8.

$$\begin{aligned}
p, q ::= & (\text{signal } S \ p) \mid (\text{seq } p \ q) \mid (\text{emit } S) \mid (\text{present } S \ p \ q) \mid (\text{par } p \ q) \\
& \mid \text{nothing} \mid \text{pause} \mid (\text{loop } p) \mid (\text{suspend } p \ S) \mid (\text{trap } p) \mid (\text{exit } n) \\
& \mid (\text{shared } s := e \ p) \mid (+= s \ e) \mid (\text{var } x := e \ p) \mid (:= x \ e) \mid (\text{if } x \ p \ q) \\
S \in & \text{signal variables} & x \in & \text{sequential variables} \\
s \in & \text{shared variables} & e \in & \text{host expressions}
\end{aligned}$$

Fig. 1. Esterel Syntax

2 A SENSE OF ESTEREL

This section provides some background on Esterel both to introduce the language to those not familiar with it and to orient Esterel experts with the particular notation we have chosen for Kernel Esterel. Figure 1 shows the syntax we use for our Esterel calculus.

Evaluation of an Esterel program is unlike conventional programming languages in that it proceeds in a series of *instants*. Each instant happens in, essentially, no time and appears atomic from the outside. An instant is triggered by a change in the state of the outside world. The external state changes of the world are communicated to Esterel via *signals*. Within each instant, each signal can either be present (set), absent (not set), or in an indeterminate state, where it is not yet known if it will be present or absent. Once a signal's value becomes known in a specific instant, it cannot change. Accordingly, the outside world may, in between instants, set a signal to present or it may not, indicating that its value is as yet undetermined (as the program itself may set it).¹ Once the instant begins these signal values cannot be modified by outside world, preventing interference with the computation. Once an instant completes, the Esterel program will have decided the value of all of its signals. The outside world can then observe these values, and respond by setting different signals for the next instant. The value of signals does not carry over between instants.

Esterel is typically used as an embedded language, where the outside world is some other programming language, e.g., C for reactive, real-time systems (Potop-Butucaru et al. 2007), Bigloo (Serrano and Weis 1995) and JavaScript for GUIs (Berry et al. 2011), or Racket (Flatt and PLT 2010) for medical prescriptions (Florence et al. 2015). The external language controls when instants take place and sets up the signal environment for each instant. From the perspective of the host language, the atomicity of instants gives Esterel a notion of discrete, logical time. Each instant represents one tick of the clock, and the host language controls the “clock speed” by explicitly starting instants.

2.1 Conditioning on Signals: present

Esterel programs can also have local signals that they use to communicate internally. Let us consider a few example programs that use internal signals to get a sense of how Esterel programs evaluate. Figure 2 shows a first example.

This program has two external signals, S01 and S02,² through which it will communicate its output. The signal form is a binding form that introduces a local signal (here named SL) available in its body. Signals that are free in the entire program are the ones that support communication with the host language, external to Esterel. At the beginning of the instant the values of SL, S01, and S02 are not known.

```

(signal SL
  (seq (emit SL)
    (present SL
      (emit S01)
      (emit S02))))

```

Fig. 2: A First Example

¹For those familiar with Esterel: free signals in programs in our calculus correspond to input-output signals in Esterel.

²We prefix all signal names with an S.

The `seq` form is sequential composition, so this program first emits `SL`, which means the signal `SL` is known to be present for this entire instant. Next, the program evaluates a signal conditional, written using the `present` keyword in Esterel. When a signal is known to be present, a `present` form is equivalent to its first sub-expression, in this case `(emit S01)`. So this program emits `S01` and then terminates, ending the instant with `S01` present and with `S02` absent.

Esterel also supports a deterministic form of parallelism and indeed if we replace the sequential composition in figure 2 with parallel composition, as shown in figure 3, the program is guaranteed to behave identically. Specifically, the `present` form in the second arm of the `par` (conceptually) blocks until the signal `SL` is emitted or we learn it cannot be emitted in this instant. So the first arm of the `par` is the only part of the program that can progress, and once it performs the `(emit SL)`, that unblocks `present` form, enabling `(emit S01)` to happen.

In order for a `present` expression to become unblocked and evaluate the second sub-expression, the Esterel program must determine that given signal cannot be emitted in this instant. One way this can happen is that there are no occurrences of `(emit SL)`. So, if we remove the `(emit SL)` from our running example, as shown in figure 4, then the program will emit the signal `S02`.

The way that `present` works helps guarantee Esterel's form of deterministic concurrency. Until a particular signal's value becomes known, the program simply refuses to make a choice about which branch to run. This style of conditional raises many interesting

questions about how apparent cyclic references interact with each other, however. For example, what should the program in figure 5 do? (`nothing` is the Esterel equivalent of `unit` or `void` in other languages.) How such programs behave is well-studied in the Esterel community and touches on the notions of logical correctness and constructiveness, which we return to in section 2.4.

2.2 Running for Multiple Instants: pause

So far, all of the example programs have terminated in a single instant but, in general, an Esterel program might run to some intermediate state and then pause. When all of the parallel branches of some program have paused or terminated, then the instant terminates. During the next instant, however, evaluation picks up right where it left off, with whatever remains of the program.

The `pause` expression brings the host language's notion of logical time into Esterel. From the Esterel programmer's perspective, every instruction in the language is *instantaneous*—taking zero logical time—with the exception of `pause`, which takes one unit of time. This effectively stops a thread of execution when it reaches `pause`, until the host language starts the next time step. At the start of the next instant, one unit of time has passed, so the pauses will have had enough “time” to complete and the program will resume from that point.

```
(signal SL
  (par (emit SL)
        (present SL
          (emit S01)
          (emit S02))))
```

Fig. 3: This time with `par`

```
(signal SL
  (present SL
    (emit S01)
    (emit S02)))
```

Fig. 4: A signal never emitted

```
(signal SL1
  (signal SL2
    (par (present SL1 (emit SL2) nothing)
          (present SL2 (emit SL1) nothing))))
```

Fig. 5: Cyclic signal dependencies

```
(signal SL
  (par (seq pause
        (emit SL))
        (present SL
          (emit S01)
          (emit S02))))
```

Fig. 6: Multiple instants

As an example, consider the program in figure 6. As it starts, the second arm of the par blocks, as with the example in figure 3. The first arm of the par first evaluates pause, which means that that arm of the par has terminated for the instant, and cannot reach the (emit SL) until the next instant. Accordingly the present takes the else branch, safe in the knowledge that no (emit SL) can happen this instant. In the next instant, the program resumes from each pause it hit the previous instant. Therefore only SL is emitted in the second instant.

2.3 Determining That a Signal Cannot be Emitted: Can

Determining whether or not a signal can be emitted is not simply a matter of eliminating untaken branches in present expressions that have executed and then checking the remaining emit expressions. Sometimes, a present may be blocked on some as-yet indeterminate signal, but portions of its branches are already known to be unreachable because other signal values are known, enabling us to declare that some signal is absent.

For example, consider the program in figure 7. The par's first sub-expression is a present and its second is a seq. The present expression is blocked on SL1. Of course, the last expression in the seq expression emits SL1 but beware: it is preceded by another present expression that may or may not pause. If it does pause, then the (emit SL1) happens in a future instant (so we take the "else" branch of the present on SL1). If it does not pause, then the (emit SL1) happens in the current instant (and so we take the "then" branch of the present on SL1). Let's work through how Esterel evaluates this complex interplay of signals and branches of par expressions.

First SL2 is emitted. Once it is, it is clear which branch the inner present expressions will take, even though they cannot yet run because we do not yet know about SL1. In particular, neither one can take their "else" branch and thus none of the (emit SL3) expressions can be evaluated. Accordingly we can now reduce the (present SL3 pause nothing) to nothing. From there we can evaluate (emit SL1), which unblocks the present on SL1, which emits the output signal S01.

The most important step in this sequence was when Esterel decided that SL3 cannot be emitted. The decision procedure for determining when a signal cannot be emitted in the current instant is called Can. It follows the same reasoning we have here, but accounts for other details of the core language of Esterel. For example, if the first sub-expression of a seq cannot terminate in a given instant, Can will rule out any emissions in the second sub-expression.

The full definition is given in figures 15 and 16, and is explained in section 3.2.

2.4 Getting Stuck: Logical Correctness and Constructivity

The style of instantaneous decision making in Esterel, facilitated via Can, leads to programs with no meaning, even though a traditional programming language would give them meaning. Such programs are called *logically incorrect*.

```
(signal SL1
  (signal SL2
    (signal SL3
      (par (present SL1
            (present SL2
              (emit S01)
              (emit SL3))
            (present SL2
              (emit S02)
              (emit SL3)))
          (seq
            (emit SL2)
            (seq
              (present SL3 pause nothing)
              (emit SL1)))))))
```

Fig. 7: Can

Logical correctness can be thought of as a consequence of the instantaneous nature of decision making in Esterel: if non-pause expressions take no logical time, then decisions about the value of a signal are communicated instantly and that value cannot change. Therefore, the program should behave as if that value was determined at the start of the instant. Therefore, there should be exactly one value for each signal. Some programs, however, have zero or multiple possible assignments. Consider the program in figure 8. No matter the definition of Can, S1 cannot be set to either present or absent. If S1 were present, the program would take the first branch of the condition, and the program would terminate without having emitted S1. If S1 were set to absent, the program would chose the second branch and emitting the signal. Both executions lead to a contradiction, therefore there are no valid assignments of signals in this program. This program is logically incorrect.

The opposite is true of the signals in the program in figure 9. Here, if S1 is chosen to be present, the conditional will take the first branch and S1 will be emitted, justifying the choice of signal value. However, if the signal is chosen

```
(signal S1
  (present S1 nothing (emit S1)))
```

Fig. 8: No possible value for S1

to be absent, the signal will not be emitted and the choice of absence is also justified. Thus there are two possible assignments to the signals in this program, and this program is also logically incorrect.

A related notion, *constructiveness*, arises from an order of execution imposed by seq and present. All decisions in the first part of a seq must be made before decisions in the second part and the value of a signal being conditioned

```
(signal S1
  (present S1 (emit S1) nothing))
```

Fig. 9: Too many values for S1

on by present must be determined before decisions within either of its branches can be made. Decisions that may affect sibling branches in a par, however, may happen in any order.

To ensure these ordering constraints, Esterel imposes an order on information propagation: decisions about the presence of a signal can only be used by the portion of the program that is after (in the sense of the ordering imposed by seq and present) it is emitted. Thus, programs that are logically correct may still be rejected because there is no order in which to run the program that will arrive at the single, valid assignment. Such programs are called non-constructive.³ Not all logically correct programs are constructive, but the converse is true: all constructive programs are logically correct. Put another way, making a guess about the value of a signal and backtracking if the guess turns out to be wrong would admit logically correct, but non-constructive, programs.

Succinctly, a program is constructive if it is logically correct, and the values of signals can be determined without any speculation: a signal is present only after it has been emitted, and a signal is absent only after Can determines it cannot be emitted.

Example non-constructive programs are shown in figure 10. The first program has only one possible assignment for S1, as it is emitted by both branches of the conditional. Because present requires that the value of S1 be known before executing a sub-expression, however, there is no valid order in which to execute the code, and the program is rejected as non-constructive. A similar phenomena can be seen in the second program in figure 10, but with seq.

```
(signal S1
  (present S1 (emit S1) (emit S1)))
(signal S1
  (seq (present S1
        nothing
        nothing)
        (emit S1)))
```

Fig. 10: Constructiveness examples

The two ordering constraints can interact. In the example in figure 11, the (emit SL1) is in a seq that may or may not pause, which prevents us from determining if SL2 is emitted.

Non-constructive programs are handled two different ways by Esterel implementations. Some approximate constructiveness with a conservative static analysis and reject programs they cannot

³The use of the name “constructive” arises from connections to constructive logic (Mendler et al. 2012).

prove constructive on all inputs. This is the default behavior of Esterel v5 (Berry 2000). Others treat non-constructivity as a runtime error, raising an error if, during an instant, the program cannot determine a value for all signals. This is the behavior of Hiphop.js (Berry et al. 2011), and Esterel v5 when used with the `-I` flag.

In the circuit semantics for Esterel, a non-constructive program is one that, when compiled to a circuit, will cause the circuit to misbehave, never settling because of instantaneous cyclic dependencies between inputs and outputs of some of the gates. That is, a program is constructive if and only if its circuit stabilizes within some fixed delay (Berry 2002; Mendler et al. 2012).

```
(signal SL1
  (signal SL2
    (par (present SL1 (emit SL2) nothing)
      (seq (present SL2 pause nothing)
        (emit SL1))))))
```

Fig. 11: Getting stuck

Non-constructive programs usually get stuck in our calculus, but they do not always. The issues here are subtle and revisited in section 4.

2.5 Loops, suspend, Non-local Exits, Variables, and the Host Language

Our calculus also covers the rest of Kernel Esterel. The `(trap p)` and `(exit n)` forms allow non-local control. Roughly speaking, `(exit n)` will abort execution up to the $n+1^{\text{th}}$ enclosing `(trap p)`, reducing it to nothing. These can be used for exception handling, but also for non-exceptional control flow. For example, it may be simpler to express some repeating task on the assumption it never terminates and then, in parallel, use `exit` to abort the task when necessary. Kernel Esterel's trap is a simplified form of Esterel's trap where traps are named and exits refer to those names.

The loop form is an infinite loop, running its body, `p`, over and over, but with a constraint that the loop's body can be started at most once in any instant. This means that the body of a loop must either pause or exit at least once in every iteration, thereby ensuring that instants always terminate. One subtle ramification of this point is that two separate iterations of a loop may run within a single instant, but only in the situation where we finish an iteration that was started in a previous instant and then start a new one in the current instant (which must then pause or exit). We return to this point in section 3.3.

Loops that fail this condition are called *instantaneous* and programs with such loops are not constructive. In our calculus, we handle this by reducing a loop in such a way that the program gets stuck if the loop were to be instantaneous.

The suspend form has a subtle semantics. If we are starting a suspend for the first time, it simply runs the body. But, if we are picking up from a previous instant where we paused in the body of a suspend, then we test the signal. If it is present, the entire suspend is paused until the next instant. If it is not present, evaluation continues within the suspend, picking up at the pause.

The suspend form is used to implement many useful, high-level behaviors. One straight-forward use is to implement a form of multiplexing, where some portion of the input signals are used directly by several different sub-pieces of the computation at once, and another portion of the input determines which of those computation is the desired output. For example, an ALU might, in parallel, both add and multiply its inputs and store the output in the same place. The suspend form can be used to control whether the addition or multiplication computation happens.

Another use of suspend is in task management. As a workflow is progressing there may be a task that runs at an interval that varies over time. This repeating task is important, but there may be an occasional situation where some more important task takes precedence. suspend can be used to pause the subcomputation corresponding to the repeating task, and resume it later without losing its current state.

And finally, Esterel has two forms of variables: sequential variables ($\mathbf{x}$) and shared variables (lowercase $\mathbf{s}$). Both of these variables refer to values and expressions in a host language, into which Esterel is embedded. For example, in Esterel v5 (Berry 2000) the host language is a subset of C, whereas in Hiphop.js (Berry et al. 2011) the host language is JavaScript.

Sequential variables are conventional mutable variables. To ensure deterministic concurrency, they may be used only sequentially (any given variable may not appear free in both branches of any specific `par` expression).

Shared variables, on the other hand, may be modified or looked at in multiple branches of a `par` expression. However, restrictions apply to ensure that the order of modifications is not observable. In particular, the program's execution cannot be influenced by the value of the variable until after all modifications have been performed (in a manner reminiscent of, but simpler than, Kuper and Newton (2013)'s LVars).

Shared variables start each instant with their old values, carried over from the previous instant. Multiple writes to a shared variable within an instant are collected with an associative and commutative operation, which throws away the value from the previous instant. After all possible writes are collected, the shared variable's value is available. Tracking if a shared variable is writable uses the same mechanism as tracking whether or not a signal is set, and shared variables are subject to the same logical correctness and constructiveness constraints as signals.

For simplicity, our calculus restricts shared and sequential variables to be natural numbers. Shared variables use `+` as the only combining operation. Sequential variables also support a conventional conditional expression, `if`, that tests if the value is 0 or not.

For a fuller explanation of these features and how they behave, start with Potop-Butucaru et al. (2007)'s book *Compiling Esterel*, especially the first two chapters. The semantic rules in figure 14 also provide more details on how these constructs work.

3 THE ESTEREL CALCULUS

The core of our calculus is a reduction relation on program expressions that corresponds to a single-step of computation within one instant. This relation captures a notion of simplification, where each computational step brings us closer to a final answer. Thus, the reduction induces an evaluator for the language. Furthermore, the reflexive, symmetric, and transitive closure of the relation together with its closure over arbitrary contexts gives rise to an equivalence relation between programs terms, which is our calculus.

The remainder of this section explores the definitions that comprise the calculus, specifically the definitions shown in figure 12. Section 3.1 shows the basic notion of reduction that our calculus supports and section 3.2 describes our `Can` function. The judgment form $\vdash_{\text{CB}}$ captures how signals are to be used in Esterel programs, and is described in section 3.3. Finally section 3.4 gives the definitions of `Eval` and $\rightsquigarrow$, and the central result of this work, namely that `Eval` is a function.

Before diving into the rules, however, we need to extend the $\mathbf{p}$ non-terminal to track information about the term as it reduces. Figure 13 shows the two extensions. First, the $(\text{loop } \mathbf{p} \ \mathbf{q})$ expression form is similar to a $(\text{seq } \mathbf{p} \ (\text{loop } \mathbf{q}))$ and is used by the loop reduction rule (discussed in section 3.1).

The other extension is the $(\mathbf{q} \ \boldsymbol{\theta} \ . \ \mathbf{p})$ expression form. It pairs an environment ($\boldsymbol{\theta}$) with an Esterel expression. The environment records what we have learned about the signals and variables in this instant for the contained subexpression, and various rules either add information to the $\boldsymbol{\theta}$ or exploit information recorded as the program reduces. We keep the environments local to specific expressions in order to facilitate local reasoning.

$p \rightarrow q$	$p \equiv_e q$
Our notion of reduction; the primitive computational steps of our calculus	The reflexive, transitive, symmetric, context closure of $\rightarrow$
Can : $p \theta \rightarrow \{ S: (\text{Setof } S), K: (\text{Setof } \kappa), sh: (\text{Setof } s) \}$	Eval : $p \theta \rightarrow (\text{Setof } S)$
Determines the signals an expression can emit	Runs the program for a single instant and returns the emitted signals
$\vdash_{CB} p$	$\Downarrow : \text{complete} \rightarrow p$
A well-formedness condition on programs ensuring that signals and variables are well-behaved	Prepares a fully-reduced program for the next instant

Fig. 12. An overview of the main definitions

$p, q ::= \dots$	Empty Environment: $\{\}$
$(\varrho \theta. p)$	Singleton Environments: $\{ \langle \text{« var »} \mapsto \langle \text{« val »} \rangle \}$
$(\overline{\text{loop}} p q)$	$\{ S \mapsto \text{status} \}$
status ::= present	$\{ s \mapsto \langle n, \text{shared-status} \rangle \}$
absent	$\{ x \mapsto n \}$
unknown	
shared-status ::= ready old new	Environment Composition: $\theta \leftarrow \theta$
complete ::= done $(\varrho \theta^c. \text{done})$	$(\theta_1 \leftarrow \theta_2)(S) = \theta_2(S) \text{ if } S \in \text{dom}(\theta_2)$
done ::= stopped paused	$(\theta_1 \leftarrow \theta_2)(S) = \theta_1(S) \text{ if } S \notin \text{dom}(\theta_2)$
stopped ::= nothing (exit n)	... ditto for s and x
paused ::= pause	
$(\text{seq paused } q)$	Complete Environments: θ^c
$(\overline{\text{loop}} \text{ paused } q)$	A complete environment is one
$(\text{par } \text{ paused } \text{ paused})$	where no signals are unknown
$(\text{suspend } \text{ paused } S)$	and all shared variables are ready
$(\text{trap } \text{ paused})$	
E ::= (seq E q)	Resetting Environments: $\lfloor \theta^c \rfloor$
$(\overline{\text{loop}} E q)$	Resetting a complete environment
$(\text{par } E q)$	updates all signals to unknown
$(\text{par } p E)$	and all shared variables to old
$(\text{suspend } E S)$	
$(\text{trap } E)$	Restricting the Domain: $(\theta \setminus \{S\})$
$[\]$	Restricting the domain of an
	environment removes the
	binding for S
Metafunctions:	Embedded host language expressions
$\downarrow^p : \text{stopped} \rightarrow \text{stopped}$	e : host expressions
$\downarrow^p \text{ nothing} = \text{nothing}$	FV(e) : all x and s that appear free in e
$\downarrow^p (\text{exit } 0) = \text{nothing}$	$\mathcal{E}val[\![e, \theta]\!]$: evaluation; produces n
$\downarrow^p (\text{exit } n) = (\text{exit } n-1)$	

Fig. 13. Supplemental Structures

	$(\text{signal } \mathbf{S} \ \mathbf{p}) \rightarrow (\varrho \{ \mathbf{S} \mapsto \text{unknown} \}. \mathbf{p})$	[signal]
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(\text{emit } \mathbf{S})]) \rightarrow (\varrho (\boldsymbol{\theta} \leftarrow \{ \mathbf{S} \mapsto \text{present} \}). \text{E}[\text{nothing}])$	[emit]
signals	where $\boldsymbol{\theta}(\mathbf{S}) \in \{ \text{present}, \text{unknown} \}$	
	$(\varrho \ \boldsymbol{\theta}. \mathbf{p}) \rightarrow (\varrho (\boldsymbol{\theta} \leftarrow \{ \mathbf{S} \mapsto \text{absent} \}). \mathbf{p})$	[absence]
	where $\mathbf{S} \in \text{dom}(\boldsymbol{\theta}), \mathbf{S} \notin \text{Can}_\varrho \llbracket (\varrho \ \boldsymbol{\theta}. \mathbf{p}), \{\} \rrbracket.S, \boldsymbol{\theta}(\mathbf{S}) = \text{unknown}$	
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(\text{present } \mathbf{S} \ \mathbf{p} \ \mathbf{q})]) \rightarrow (\varrho \ \boldsymbol{\theta}. \text{E}[\mathbf{p}])$ where $\boldsymbol{\theta}(\mathbf{S}) = \text{present}$	[is-present]
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(\text{present } \mathbf{S} \ \mathbf{p} \ \mathbf{q})]) \rightarrow (\varrho \ \boldsymbol{\theta}. \text{E}[\mathbf{q}])$ where $\boldsymbol{\theta}(\mathbf{S}) = \text{absent}$	[is-absent]
shared variables	$(\varrho \ \boldsymbol{\theta}. \text{E}[(\text{shared } \mathbf{s} := \mathbf{e} \ \mathbf{p})]) \rightarrow (\varrho \ \boldsymbol{\theta}. \text{E}[(\varrho \{ \mathbf{s} \mapsto \langle \mathbf{n}, \text{old} \rangle \}. \mathbf{p})])$	[shared]
	where $\text{FV}(\mathbf{e}) \subset \text{dom}(\boldsymbol{\theta}), \forall \mathbf{s} \in \text{FV}(\mathbf{e}). \boldsymbol{\theta}(\mathbf{s}) = \langle _, \text{ready} \rangle, \mathbf{n} = \mathcal{E}val \llbracket \mathbf{e}, \boldsymbol{\theta} \rrbracket$	
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(+= \mathbf{s} \ \mathbf{e})]) \rightarrow (\varrho (\boldsymbol{\theta} \leftarrow \{ \mathbf{s} \mapsto \langle \mathcal{E}val \llbracket \mathbf{e}, \boldsymbol{\theta} \rrbracket, \text{new} \rangle \}). \text{E}[\text{nothing}])$	[set-old]
	where $\boldsymbol{\theta}(\mathbf{s}) = \langle _, \text{old} \rangle, \text{FV}(\mathbf{e}) \subset \text{dom}(\boldsymbol{\theta}), \forall \mathbf{s} \in \text{FV}(\mathbf{e}). \boldsymbol{\theta}(\mathbf{s}) = \langle _, \text{ready} \rangle$	
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(+= \mathbf{s} \ \mathbf{e})]) \rightarrow (\varrho (\boldsymbol{\theta} \leftarrow \{ \mathbf{s} \mapsto \langle \mathbf{n} + \mathcal{E}val \llbracket \mathbf{e}, \boldsymbol{\theta} \rrbracket, \text{new} \rangle \}). \text{E}[\text{nothing}])$	[set-new]
	where $\boldsymbol{\theta}(\mathbf{s}) = \langle \mathbf{n}, \text{new} \rangle, \text{FV}(\mathbf{e}) \subset \text{dom}(\boldsymbol{\theta}), \forall \mathbf{s} \in \text{FV}(\mathbf{e}). \boldsymbol{\theta}(\mathbf{s}) = \langle _, \text{ready} \rangle$	
	$(\varrho \ \boldsymbol{\theta}. \mathbf{p}) \rightarrow (\varrho (\boldsymbol{\theta} \leftarrow \{ \mathbf{s} \mapsto \langle \mathbf{n}, \text{ready} \rangle \}). \mathbf{p})$	[readyness]
sequential variables	where $\mathbf{s} \in \text{dom}(\boldsymbol{\theta}), \mathbf{s} \notin \text{Can}_\varrho \llbracket (\varrho \ \boldsymbol{\theta}. \mathbf{p}), \{\} \rrbracket.\text{sh}, \boldsymbol{\theta}(\mathbf{s}) = \langle \mathbf{n}, \text{shared-status} \rangle,$	
	$\text{shared-status} \in \{ \text{old}, \text{new} \}$	
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(\text{var } \mathbf{x} := \mathbf{e} \ \mathbf{p})]) \rightarrow (\varrho \ \boldsymbol{\theta}. \text{E}[(\varrho \{ \mathbf{x} \mapsto \mathcal{E}val \llbracket \mathbf{e}, \boldsymbol{\theta} \rrbracket \}. \mathbf{p})])$	[var]
	where $\text{FV}(\mathbf{e}) \subset \text{dom}(\boldsymbol{\theta}), \forall \mathbf{s} \in \text{FV}(\mathbf{e}). \boldsymbol{\theta}(\mathbf{s}) = \langle _, \text{ready} \rangle$	
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(:= \mathbf{x} \ \mathbf{e})]) \rightarrow (\varrho (\boldsymbol{\theta} \leftarrow \{ \mathbf{x} \mapsto \mathcal{E}val \llbracket \mathbf{e}, \boldsymbol{\theta} \rrbracket \}). \text{E}[\text{nothing}])$	[set-var]
	where $\mathbf{x} \in \text{dom}(\boldsymbol{\theta}), \text{FV}(\mathbf{e}) \subset \text{dom}(\boldsymbol{\theta}), \forall \mathbf{s} \in \text{FV}(\mathbf{e}). \boldsymbol{\theta}(\mathbf{s}) = \langle _, \text{ready} \rangle$	
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(\text{if } \mathbf{x} \ \mathbf{p} \ \mathbf{q})]) \rightarrow (\varrho \ \boldsymbol{\theta}. \text{E}[\mathbf{p}])$ where $\mathbf{x} \in \text{dom}(\boldsymbol{\theta}), \boldsymbol{\theta}(\mathbf{x}) \neq 0$	[if-true]
	$(\varrho \ \boldsymbol{\theta}. \text{E}[(\text{if } \mathbf{x} \ \mathbf{p} \ \mathbf{q})]) \rightarrow (\varrho \ \boldsymbol{\theta}. \text{E}[\mathbf{q}])$ where $\boldsymbol{\theta}(\mathbf{x}) = 0$	[if-false]
$\circ$	$(\varrho \ \boldsymbol{\theta}_1. \text{E}[(\varrho \ \boldsymbol{\theta}_2. \mathbf{p})]) \rightarrow (\varrho (\boldsymbol{\theta}_1 \leftarrow \boldsymbol{\theta}_2). \text{E}[\mathbf{p}])$	[merge]
seq	$(\text{seq nothing } \mathbf{q}) \rightarrow \mathbf{q}$	[seq-done]
	$(\text{seq (exit } \mathbf{n}) \ \mathbf{q}) \rightarrow (\text{exit } \mathbf{n})$	[seq-exit]
trap	$(\text{trap stopped}) \rightarrow \downarrow^p \text{stopped}$	[trap]
par	$(\text{par nothing done}) \rightarrow \text{done}$	[par-nothing]
	$(\text{par (exit } \mathbf{n}) \ \text{paused}) \rightarrow (\text{exit } \mathbf{n})$	[par-1exit]
	$(\text{par (exit } \mathbf{n}_1) (\text{exit } \mathbf{n}_2)) \rightarrow (\text{exit } \max(\mathbf{n}_1, \mathbf{n}_2))$	[par-2exit]
	$(\text{par } \mathbf{p} \ \mathbf{q}) \rightarrow (\text{par } \mathbf{q} \ \mathbf{p})$	[par-swap]
	$(\text{suspend stopped } \mathbf{S}) \rightarrow \text{stopped}$	[suspend]
loop	$(\text{loop } \mathbf{p}) \rightarrow (\text{loop } \mathbf{p} \ \mathbf{p})$	[loop]
	$(\text{loop (exit } \mathbf{n}) \ \mathbf{q}) \rightarrow (\text{exit } \mathbf{n})$	[loop^stop-exit]

Fig. 14. Reduction Rules

3.1 Reduction Rules

The rules given in figure 14 govern how computation takes place within a single instant.

The first rule, **[signal]**, reduces a **signal** expression to a **q** expression by introducing a singleton θ that binds the signal to unknown.

Once a signal has an entry in a relevant θ , the **[emit]** rule records that a signal is present (using the composition operator $\leftarrow$ from figure 13) and eliminates the **emit** expression. The side-condition ensures that the environment θ does not already indicate that the signal is absent.

In order for an **emit** to fire, it must be in the body of a **q** in only a specific set of positions, as captured by the **E** contexts, shown in figure 13. They include the first sub-expression of a **seq** expression, the first sub-expression of a $(\overline{\text{loop}} \text{ p } \text{q})$ expression, either branch of a **par**, the body of a **suspend** or a **trap**. Notably this rule does not allow an expression like $(\text{q } \theta. (\text{seq } \text{p } (\text{emit } \text{S})))$ to reduce its **emit** expression because the **p** could be **pause**, delaying the **(emit S)** to the next instant. More generally, the expressions captured by **E** are guaranteed to happen in the current instant.

As we saw in section 2.3, **Can** determines if a signal cannot be emitted. The rule **[absence]** uses Can_q (a variation of **Can** that is explained in section 3.2) to determine that a signal cannot be emitted and records that information in a θ expression, if that information is not yet recorded.

Once the status of a signal is recorded as either present or absent, the **[is-present]** and **[is-absent]** rules can reduce present expressions.

The rules **[shared]**, **[set-old]**, **[set-new]**, and **[readyness]** handle shared variables in a manner similar to how the previous set of rules handle signals. The **[shared]** rule introduces a new environment that binds the shared variable using the **e** in the shared expression to determine the default value of the variable using the host language's evaluation function. The rules **[set-old]** and **[set-new]** modify a shared variable depending on whether it has been modified in the current instant or not. If the status of a shared variable in the environment is **old**, it is being modified for the first time in the current instant and the rule **[set-old]** replaces the old value in the environment with the new value. If the status of a shared variable is **new**, it has already been modified in the current instant and the rule **[set-new]** adds the current value and the new value in the **+=** expression and stores the result in the environment. One subtlety of note here: the **[shared]** rule creates an environment which marks the shared variable as **old**, not **new**. This is because the value initially given to a shared variable represents its default value rather than its initial value, and so acts as if this value was set in the previous instant. Finally, the **[readyness]** rule makes a variable change from **writable** to **readable**. This occurs if Can_q 's result does not contain the shared variable **s**, which means it will not be modified in this instant and thus we can update the environment to mark the variable as **ready**. Furthermore, the side-conditions on the **[shared]**, **[set-new]**, and **[set-old]** rules (as well as the corresponding rules for sequential variables) ensure that these rules can fire only if, for every shared variable used in the host language expression, that variable safe to be read, e.g. is marked as **ready** in θ .

The rules **[var]**, **[set-var]**, **[if-true]**, and **[if-false]** cover sequential variables. Unlike the rules for signals or shared variables, these rules do not refer to **Can**. These variables are not allowed to be free in two different arms of any **par** expression, so they can be freely read and written.

The final rule that handles **q** expressions is **[merge]**. It combines two environments, lifting an inner environment out to an outer one and composing them into a single environment.

There are two rules for sequential composition. If the first sub-expression is **nothing**, then we replace the entire expression with the second branch. If the first sub-expression is an **exit** expression, however, then the entire sequence exits, discarding the second part of the **seq** expression.

The next rule handles `trap`. Once the body of a `trap` has finished evaluating, it will either be an `exit` expression or nothing, which the $\downarrow^p$ (figure 13) function handles.

The `par` rules are a little more interesting. The first three refer to the definitions of **stopped** and **done** in figure 13 and handle the situations when both branches are finished for the instant. If one side has reduced to nothing, the `[par-nothing]` rule reduces to the other one. If one side has exited and the other is **paused**, the `[par-1exit]` rule preempts the other branch of the `par` by bubbling the `exit` up. If both sides have exited the `[par-2exit]` rule reduces the expression to whichever `exit` will reach the farthest up `trap`. The `[par-swap]` rule switches the branches, allowing `[par-nothing]` and `[par-1exit]` to match regardless of which branch is `exit` or nothing.

The `[suspend]` rule reduces to its body when its body has either exited or reduced to nothing.

This leaves us with one last pair of rules: `[loop]` and `[loop^stop-exit]`. Intuitively, we would like an expression like `(loop p)` to reduce simply to just `(seq p (loop p))`, duplicating the body `p` into a `seq` expression which becomes the current iteration of the loop.

Such a rule could give rise to infinite loops within a single instant, however, which is forbidden in Esterel. We capture this constraint in our calculus with the $\overline{\text{loop}}$ expression form. It is introduced only by the reduction rule for `loop`, and is meant to capture a single unrolling of the loop; the first sub-expression is the part of the loop that runs in the current instant and the second sub-expression is the body of the loop, saved to be used in the next instant. There is no rule that eliminates a $\overline{\text{loop}}$ when the first sub-expression is nothing (unlike `seq`, which has the `[seq-done]` rule). As such, programs get stuck when they contain instantaneous loops.

One thing to note about these rules: with the exception of `[par-swap]`, they are strongly normalizing. The proof is given as *noetherian* in Agda code in the supplementary material.

3.2 The Can Function

This section describes the function `Can`, a conservative analysis of the state of an Esterel program that determines its behavior. Our definition is inspired by Berry (2002)’s definition, generalized to support `q` expressions and modified to handle a reduction semantics rather than one based on annotating the program with program counters.

This function computes a conservative approximation to the behavior of some given Esterel expression with respect to some knowledge about signals and shared variables that is encapsulated in an environment, θ . In particular, it computes a set of signals (S), a set of exit codes (κ), and a set of shared variables (s). Any S that is not in the result is guaranteed not to be emitted in the current instant (although if some S is in the result, it may or may not be emitted in the current instant). The same holds for any shared variable in the result: if an s is not the result, then it is guaranteed that s cannot be updated again in the current instant. If the s is in the result, then it may or may not be written to. The exit codes capture whether or not the given expression pauses, reduces to nothing, or exits. If the expression may reduce to nothing, then the code `nothin` will be in the result. If the expression may pause, then the code `paus` will be in the result. If the expression may exit with the code `n`, then the code `n` will be in the result. Thus, if any of those specific codes are *not* in the result, then we know the expression does not have the corresponding behavior.

The notation we use for the records in the definition of `Can` is similar to many record notations, but we use the precise one in Pierce (2002)’s book *Types and Programming Languages*. We write `Can`’s result as a record with three fields, where curly braces construct records, e.g., the `emit` case of `Can` returns a record with a singleton set of signals (containing S), a singleton set of exit codes (containing `nothin`) and the empty set of shared variables. Selecting a field from a record uses dot notation. For example, `Can[p,  $\theta$ ].S` selects the “S” field from a call to `Can`.

$\kappa ::= \text{nothin} \mid \text{paus} \mid n$

$\text{Can} : \mathbf{p} \ \boldsymbol{\theta} \rightarrow \{ S : (\text{Setof } \mathbf{S}), K : (\text{Setof } \kappa), \text{sh} : (\text{Setof } \mathbf{s}) \}$

$\text{Can}[(\varrho \ \boldsymbol{\theta}_1, \mathbf{p}), \boldsymbol{\theta}_2] = \{ S = \text{Can}_\varrho[(\varrho \ \boldsymbol{\theta}_1, \mathbf{p}), \boldsymbol{\theta}_2].S \setminus \text{dom}(\boldsymbol{\theta}_1),$
 $K = \text{Can}_\varrho[(\varrho \ \boldsymbol{\theta}_1, \mathbf{p}), \boldsymbol{\theta}_2].K,$
 $\text{sh} = \text{Can}_\varrho[(\varrho \ \boldsymbol{\theta}_1, \mathbf{p}), \boldsymbol{\theta}_2].\text{sh} \setminus \text{dom}(\boldsymbol{\theta}_1) \}$

$\text{Can}[\text{nothing}, \boldsymbol{\theta}] = \{ S = \emptyset, K = \{ \text{nothin} \}, \text{sh} = \emptyset \}$

$\text{Can}[\text{pause}, \boldsymbol{\theta}] = \{ S = \emptyset, K = \{ \text{paus} \}, \text{sh} = \emptyset \}$

$\text{Can}[(\text{exit } n), \boldsymbol{\theta}] = \{ S = \emptyset, K = \{ n \}, \text{sh} = \emptyset \}$

$\text{Can}[(\text{emit } \mathbf{S}), \boldsymbol{\theta}] = \{ S = \{ \mathbf{S} \}, K = \{ \text{nothin} \}, \text{sh} = \emptyset \}$

$\text{Can}[(\text{present } \mathbf{S} \ \mathbf{p} \ \mathbf{q}), \boldsymbol{\theta}] = \text{Can}[\mathbf{p}, \boldsymbol{\theta}]$

where $\boldsymbol{\theta}(\mathbf{S}) = \text{present}$

$\text{Can}[(\text{present } \mathbf{S} \ \mathbf{p} \ \mathbf{q}), \boldsymbol{\theta}] = \text{Can}[\mathbf{q}, \boldsymbol{\theta}]$

where $\boldsymbol{\theta}(\mathbf{S}) = \text{absent}$

$\text{Can}[(\text{present } \mathbf{S} \ \mathbf{p} \ \mathbf{q}), \boldsymbol{\theta}] = \{ S = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].S \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].S,$
 $K = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].K \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].K,$
 $\text{sh} = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].\text{sh} \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].\text{sh} \}$

$\text{Can}[(\text{suspend } \mathbf{p} \ \mathbf{S}), \boldsymbol{\theta}] = \text{Can}[\mathbf{p}, \boldsymbol{\theta}]$

$\text{Can}[(\text{seq } \mathbf{p} \ \mathbf{q}), \boldsymbol{\theta}] = \text{Can}[\mathbf{p}, \boldsymbol{\theta}]$

where $\text{nothin} \notin \text{Can}[\mathbf{p}, \boldsymbol{\theta}].K$

$\text{Can}[(\text{seq } \mathbf{p} \ \mathbf{q}), \boldsymbol{\theta}] = \{ S = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].S \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].S,$
 $K = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].K \setminus \{ \text{nothin} \} \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].K,$
 $\text{sh} = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].\text{sh} \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].\text{sh} \}$

$\text{Can}[(\text{loop } \mathbf{p}), \boldsymbol{\theta}] = \text{Can}[\mathbf{p}, \boldsymbol{\theta}]$

$\text{Can}[(\text{loop } \mathbf{p} \ \mathbf{q}), \boldsymbol{\theta}] = \text{Can}[\mathbf{p}, \boldsymbol{\theta}]$

$\text{Can}[(\text{par } \mathbf{p} \ \mathbf{q}), \boldsymbol{\theta}] = \{ S = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].S \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].S,$
 $K = \{ \max(\kappa_1, \kappa_2) \mid \kappa_1 \in \text{Can}[\mathbf{p}, \boldsymbol{\theta}].K, \kappa_2 \in \text{Can}[\mathbf{q}, \boldsymbol{\theta}].K \},$
 $\text{sh} = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].\text{sh} \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].\text{sh} \}$

$\text{Can}[(\text{trap } \mathbf{p}), \boldsymbol{\theta}] = \{ S = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].S, K = \{ \downarrow^\kappa x \mid x \in \text{Can}[\mathbf{p}, \boldsymbol{\theta}].K \}, \text{sh} = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].\text{sh} \}$

$\text{Can}[(\text{signal } \mathbf{S} \ \mathbf{p}), \boldsymbol{\theta}] = \{ S = \text{Can}[\mathbf{p}, \boldsymbol{\theta} \leftarrow \{ \mathbf{S} \mapsto \text{absent} \}].S \setminus \{ \mathbf{S} \},$
 $K = \text{Can}[\mathbf{p}, \boldsymbol{\theta} \leftarrow \{ \mathbf{S} \mapsto \text{absent} \}].K,$
 $\text{sh} = \text{Can}[\mathbf{p}, \boldsymbol{\theta} \leftarrow \{ \mathbf{S} \mapsto \text{absent} \}].\text{sh} \}$

where $\mathbf{S} \notin \text{Can}[\mathbf{p}, \boldsymbol{\theta} \leftarrow \{ \mathbf{S} \mapsto \text{unknown} \}].S$

$\text{Can}[(\text{signal } \mathbf{S} \ \mathbf{p}), \boldsymbol{\theta}] = \{ S = \text{Can}[\mathbf{p}, \boldsymbol{\theta}_2].S \setminus \{ \mathbf{S} \}, K = \text{Can}[\mathbf{p}, \boldsymbol{\theta}_2].K, \text{sh} = \text{Can}[\mathbf{p}, \boldsymbol{\theta}_2].\text{sh} \}$
 where $\boldsymbol{\theta}_2 = \boldsymbol{\theta} \leftarrow \{ \mathbf{S} \mapsto \text{unknown} \}$

$\text{Can}[(\text{shared } \mathbf{s} := \mathbf{e} \ \mathbf{p}), \boldsymbol{\theta}] = \{ S = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].S, K = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].K, \text{sh} = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].\text{sh} \setminus \{ \mathbf{s} \} \}$

$\text{Can}[(\text{+} = \mathbf{s} \ \mathbf{e}), \boldsymbol{\theta}] = \{ S = \emptyset, K = \{ \text{nothin} \}, \text{sh} = \{ \mathbf{s} \} \}$

$\text{Can}[(\text{var } \mathbf{x} := \mathbf{e} \ \mathbf{p}), \boldsymbol{\theta}] = \text{Can}[\mathbf{p}, \boldsymbol{\theta}]$

$\text{Can}[(\text{:= } \mathbf{x} \ \mathbf{e}), \boldsymbol{\theta}] = \{ S = \emptyset, K = \{ \text{nothin} \}, \text{sh} = \emptyset \}$

$\text{Can}[(\text{if } \mathbf{x} \ \mathbf{p} \ \mathbf{q}), \boldsymbol{\theta}] = \{ S = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].S \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].S,$
 $K = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].K \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].K,$
 $\text{sh} = \text{Can}[\mathbf{p}, \boldsymbol{\theta}].\text{sh} \cup \text{Can}[\mathbf{q}, \boldsymbol{\theta}].\text{sh} \}$

$\downarrow^\kappa : \kappa \rightarrow \kappa$
$\downarrow^\kappa \text{nothin} = \text{nothin}$
$\downarrow^\kappa \text{paus} = \text{paus}$
$\downarrow^\kappa 0 = \text{nothin}$
$\downarrow^\kappa n = n-1$
if $n > 0$

$\max : \kappa \ \kappa \rightarrow \kappa$ $\max(\kappa_1, \kappa_2)$ computes the maximum of κ_1 and κ_2 where we define $\text{nothin} < \text{paus} < 0 < 1 < \dots$
--

Fig. 15. Can Function

$$\begin{aligned}
\text{Can}_q : (q \ \theta. \ p) \ \theta &\rightarrow \{ S: (\text{Setof } S), K: (\text{Setof } \kappa), \text{sh}: (\text{Setof } s) \} \\
\text{Can}_q \llbracket (q \ \theta. \ p), \theta_2 \rrbracket &= \text{Can}_q \llbracket (q \ (\theta \setminus \{S\}). \ p), \theta_2 \leftarrow \{ S \mapsto \text{absent} \} \rrbracket \\
&\text{where } S \in \text{dom}(\theta), \\
&\quad \theta(S) = \text{unknown}, \\
&\quad S \notin \text{Can}_q \llbracket (q \ (\theta \setminus \{S\}). \ p), \theta_2 \leftarrow \{ S \mapsto \text{unknown} \} \rrbracket.S \\
\text{Can}_q \llbracket (q \ \theta. \ p), \theta_2 \rrbracket &= \text{Can}_q \llbracket (q \ (\theta \setminus \{S\}). \ p), \theta_2 \leftarrow \{ S \mapsto \theta(S) \} \rrbracket \\
&\text{where } S \in \text{dom}(\theta) \\
\text{Can}_q \llbracket (q \ \theta_i. \ p), \theta_2 \rrbracket &= \text{Can} \llbracket p, \theta_2 \rrbracket
\end{aligned}$$

Fig. 16. The Can Function for q Expressions (cases are checked in order)

The three results from Can interact with each other in order to determine the overall result. Consider the two seq cases. In the first one, the side-condition says that nothin is not in the $\mathbf{K}$ field for the first sub-expression of the seq , p . Accordingly, we know that p does not reduce to nothing, thus it must either exit or pause. Since it exits or pauses, we know that none of the behavior of q is relevant as it will not be evaluated in this instant and so the result of Can for the entire seq expression is just its result for the p expression. This means that

$$\text{Can} \llbracket (\text{seq pause (emit } S)), \theta \rrbracket.S$$

is the empty set, since the emit must happen in the next instant.

In the second seq case, we know that nothin is a possible result code, and thus p might reduce to nothing so we have to combine the result of the p and q recursive calls. Mostly this amounts to taking the union, but note that the $\mathbf{K}$ case removes nothin from the codes in the result of p before performing the union. This removal accounts for the fact that, even if p reduces to nothing, q must also reduce to nothing for the seq expression to reduce to nothing. For example,

$$\text{Can} \llbracket (\text{seq nothing pause}), \theta \rrbracket.K$$

correctly contains only the exit code paus .

The $\overline{\text{loop}}$ expression form, in contrast, always ignores the second sub-expression, because we know that the second sub-expression can affect only future instants.

Various other cases in the definition of Can reflect the semantics of the different constructs in similar ways. The cases handling present consult the given θ to see if the status of the signal is known and look only at the corresponding branch of the present expression if so. The rule for par takes into account the same behavior that the four par rules in the reduction relation do when computing the codes for the entire expression out of the codes of the subexpressions. The trap case uses the metafunction $\downarrow^*$ to adjust the exit codes in a manner that mimics how trap expressions reduce. Since the shared form introduces a new variable, its case in Can removes that variable from the results, as it is lexically scoped. In each of these cases, Can ignores the $\mathbf{e}$ expressions, as it does not reason about the behavior of the host language.

This leaves the signal and q cases. Consider how Can handles signal expressions. The second signal case is the more straightforward one. It says that the result for the entire signal form is the same as the result for the body of a signal form when it is analyzed with no knowledge about the signal. But there would be a problem with the Can function if that were the only case.

To motivate the first signal case in Can , consider this call:

$$\text{Can} \llbracket (\text{signal } S2 \ (\text{present } S2 \ (\text{emit } S1) \ \text{nothing})), \{\} \rrbracket.S$$

If we took the second signal case in `Can`, then this would return a set containing `S1`. It actually returns the empty set. The first signal case calls `Can` with `S2` set to unknown and checks to see if `S2` is not present in the “`S`” portion of the result. It is not (because there are no `(emit S2)` expressions), so `Can` then sets `S2` to absent and reprocesses its body. This time, because `S2` is known to be absent, `Can` considers only the last sub-expression of the present, thereby ignoring the `(emit S1)` and returning the empty set of signals.

In isolation, analyzing the body twice seems like overkill, especially because it triggers exponential behavior in the number of nested signal forms.⁴ But consider this call to `Can`:

```
Can[(signal S1
      (seq (present S1 pause nothing)
            (signal S2 (present S2
                        (emit S1)
                        nothing))))), {}].S
```

This example input is a bit complex, but first notice that the inner signal expression is the same as the previous example (and there are no other `(emit S1)` expressions), so we know that `S1` is not going to be emitted. If `Can` did not have that first signal case, then it could not learn that `S1` cannot be emitted and thus we would not be able to use the **[absence]** rule on this expression, and the program would remain stuck, unable to reduce the first present.

Finally, for the `q` case, the `Can` function dispatches to `Canq` (figure 16). The `Canq` function looks complex, but it is essentially the same as the two signal cases. It is broken out into its own function because `q` binds multiple signals at once; so `Canq` recurs through the structure of the environment, considering each of the signals that are bound. The first case of `Canq` corresponds to the first signal case in `Can`; the second case in `Canq` corresponds to the second signal case, and the last case in `Canq` corresponds to the situation where there are no more signals bound in Θ (and Θ , can be dropped as it contains only information about `s` and `x` variables, which `Can` does not need).

3.3 Reincarnation, Schizophrenia, and Correct Binding

The signal form seems to be something close to a variable binding form, familiar from conventional λ -calculus based programming languages. It is, however, not the same and a significant source of subtlety in Esterel. The Esterel community has explored these issues in great detail and in this section, we try to bring across the basic points and then explain how our calculus handles them.

The two central issues are the phenomena of schizophrenic and reincarnated signals. To understand them, first recall the central tenant of Esterel signals: every signal must have exactly one value in a given instant. Now, consider the example program in figure 17. During the first instant of execution the signal `SZ` will be absent, as the program pauses before emitting it. In the next instant we pick up where we left off. The first thing that happens is that we emit `SZ`. Then the loop body restarts. Because we have re-entered the loop body and encountered the signal expression afresh, the `SZ` should now be absent. But this means that the signal `SZ` has two different values in a single instant!

In the literature, signals which are duplicated by a loop body in within one instant are called *reincarnated*. If a reincarnated signal obtains different values in each of its incarnations, it is called *schizophrenic*. Schizophrenic signals, however, merely appear to violate the single-value-per-instant

```
(loop
  (signal SZ
    (seq pause
          (emit SZ))))
```

Fig. 17: A Schizophrenic Loop

⁴This exponential behavior affected the testing of our semantics against existing Esterel semantics and implementations; see section 6 and section 7 for more.

rule. Because instantaneous loops are banned, the number of times a loop body can be entered is bounded. This means that the number of reincarnations of any signal is also bounded. Therefore we consider each incarnation to, in fact, be a separate signal, removing the apparent violation.

This resolution shows up directly in Esterel compilers and circuit semantics. Naive treatment of schizophrenic signals can cause unstable loops in the corresponding circuit, breaking the guarantee that all constructive programs translate to stable circuits. Therefore, many Esterel compilers duplicate parts of loop bodies with schizophrenic signals to remove the apparent violation of the single-value-per-instant rule, avoiding cross-loop cycles (Berry 2002; Potop-Butucaru et al. 2007; Schneider and Wenz 2001). In short, each incarnation of a signal gets a separate wire.

Esterel semantics such as the Constructive Operational Semantics (Potop-Butucaru 2002) and the Constructive Behavioral Semantics (Berry 2002) take a different approach, handling such signals by carefully arranging to “forget” a schizophrenic signal’s first value when the second one is needed.

Our semantics takes an approach inspired by the circuit perspective, meaning we do not treat signals in a conventional way. More precisely, we do not assume the variable convention (Barendregt 1984), nor do we include an α rule. Indeed, we think of signals as if they name wires.

This perspective means that schizophrenic and reincarnated signals are, at first glance, handled very simply. We just duplicate the bodies of loops in the **[loop]** rule, so each signal will end up in a different q , potentially bound to a different value—akin to the strategy that circuit semantics employ. This approach, however, does raise a significant concern: what happens if the **[merge]** rule moves q expressions in such a way that the environment captures variables it did not bind before? Our calculus avoids this problem by working only with programs that have *correct binding*, as captured by the $\vdash_{CB}$ judgment form in figure 18. (The $\vdash_{CB}$ judgment also ensures that sequential variables are used in at most one branch of any par, which is not related to the concerns of schizophrenia, but does ensure determinism and is convenient to include here.)

To understand the correct binding judgment, first look at the seq rule. It says that the bound signals of the first sub-expression must be distinct from the free signals of the second. Since the **[merge]** rule moves binders based on the definition of **E** (in figure 13), it can move a q out from the first sub-expression only. Thus, in order to preserve the binding structure of the expression as we reduce, we need only make sure that a q that moves out of the first sub-expression of a seq does not capture a signal in the second sub-expression, which is precisely what the premise avoids.

The other rules all generally follow this reasoning process for their premises. The suspend rule’s premise follows exactly that reasoning, as binder may be lifted out past the **S**. The par rule’s second and third premises also follow exactly the same reasoning. The first premise of par is necessary to avoid the situation where the same signal is bound in both branches and then is lifted out from both. The fourth premise ensures that sequential variables are used properly.

The loop rule must ensure that the bound and free signals of its subexpression do not overlap, as it reduces by duplicating its first subexpression into a $l\overline{oop}$, which acts like a seq expression (so the intuition for seq applies, but with both subexpressions being the same one). Similarly, because $(l\overline{oop} \ p \ q)$ behaves like $(seq \ p \ (loop \ q))$, the premises of its rule are just the premises of the seq and loop rules, combined.

THEOREM 3.1.

$$\forall \ p, q, C. \quad \vdash_{CB} C[p] \Rightarrow p \multimap q \Rightarrow \vdash_{CB} C[q]$$

This theorem states that, no matter which context an expression reduces in (with C as given in figure 19), if the expression had correct binding before reduction, it does afterwards, too. The proof is given as $\multimap_1$ -preserve-CB in the Agda code in the supplementary material. From this we conclude that programs with correct binding cannot exhibit incorrect variable capture.

Fig. 18. Correct Binding

It should also be noted that any Esterel program that uses its sequential variables correctly either already has correct binding or can be renamed into one that has correct binding (introducing new wires, of course) before reducing the program. Thus, the restriction that our calculus handles only programs with correct binding is not severe, as any already correct program can be transformed into one which is well behaved in our calculus.

3.4 Evaluating Programs

Now that we have established the correct binding invariant and defined the primitive notions of reduction, we can turn to the definition of the evaluator. It is shown on the top-left of figure 19. It accepts a program and an initial environment (that captures what the host language sets the input signals to), and it returns the signals that were emitted at the end of the instant. The output of the evaluator ignores shared variables. However, values of shared variables can be indirectly returned by introducing new signals whose presence depends on the values of shared variables.

The $\equiv_e$ relation is the symmetric, transitive, reflexive closure of the $\rightarrow$ relation, which is the compatible closure of the $\rightarrow$ reduction relation. The symmetric case has an additional premise $\vdash_{CB} p$ to ensure that all of the intermediate terms used in $\equiv_e$ have correct binding.

$$\frac{(\varrho \ \theta. \mathbf{p}) \equiv_e (\varrho \ \theta^c. \mathbf{done})}{\text{Eval}(\mathbf{p}, \theta) = \{ \mathbf{S} \in \text{dom}(\theta^c) \mid \theta^c(\mathbf{S}) = \text{present} \}}$$

$$\frac{\mathbf{p} \rightarrow \mathbf{q}}{\mathbf{p} \equiv_e \mathbf{q}} \quad \frac{\mathbf{p} \equiv_e \mathbf{q} \quad \vdash_{\text{CB}} \mathbf{p}}{\mathbf{q} \equiv_e \mathbf{p}}$$

$$\frac{}{\mathbf{p} \equiv_e \mathbf{p}} \quad \frac{\mathbf{p}_1 \equiv_e \mathbf{p}_2 \quad \mathbf{p}_2 \equiv_e \mathbf{p}_3}{\mathbf{p}_1 \equiv_e \mathbf{p}_3}$$

$$\frac{\mathbf{p} \rightarrow \mathbf{q}}{C[\mathbf{p}] \rightarrow C[\mathbf{q}]}$$

$$\frac{\mathbf{p}_1 \rightarrow \mathbf{p}_2 \quad \mathbf{p}_2 \rightarrow^* \mathbf{p}_3}{\mathbf{p}_1 \rightarrow^* \mathbf{p}_3} \quad \frac{}{\mathbf{p} \rightarrow^* \mathbf{p}}$$

$$\begin{aligned}
C ::= & (\text{signal } S \ C) \\
& | (\text{seq } C \ \mathbf{q}) \\
& | (\text{seq } \mathbf{p} \ C) \\
& | (\text{loop } C \ \mathbf{q}) \\
& | (\text{loop } \mathbf{p} \ C) \\
& | (\text{present } S \ C \ \mathbf{q}) \\
& | (\text{present } S \ \mathbf{p} \ C) \\
& | (\text{par } C \ \mathbf{q}) \\
& | (\text{par } \mathbf{p} \ C) \\
& | (\text{loop } C) \\
& | (\text{suspend } C \ S) \\
& | (\text{trap } C) \\
& | (\text{shared } s := e \ C) \\
& | (\text{var } \mathbf{x} := e \ C) \\
& | (\text{if } \mathbf{x} \ C \ \mathbf{q}) \\
& | (\text{if } \mathbf{x} \ \mathbf{p} \ C) \\
& | (\varrho \ \theta. C) \\
& | []
\end{aligned}$$

Fig. 19. Eval

$\rightsquigarrow : \text{complete} \rightarrow \mathbf{p}$	
$\rightsquigarrow(\varrho \ \theta^c. \mathbf{p})$	$= (\varrho \ \lfloor \theta^c \rfloor. \rightsquigarrow \mathbf{p})$
$\rightsquigarrow \text{pause}$	$= \text{nothing}$
$\rightsquigarrow \text{nothing}$	$= \text{nothing}$
$\rightsquigarrow(\text{loop } \mathbf{p} \ q)$	$= (\text{seq } \rightsquigarrow \mathbf{p} \ (\text{loop } q))$
$\rightsquigarrow(\text{seq } \mathbf{p} \ q)$	$= (\text{seq } \rightsquigarrow \mathbf{p} \ q)$
$\rightsquigarrow(\text{par } \mathbf{p} \ q)$	$= (\text{par } \rightsquigarrow \mathbf{p} \ \rightsquigarrow q)$
$\rightsquigarrow(\text{suspend } \mathbf{p} \ S)$	$= (\text{suspend } (\text{seq } (\text{present } S \ \text{pause } \text{nothing}) \ \rightsquigarrow \mathbf{p}) \ S)$
$\rightsquigarrow(\text{trap } \mathbf{p})$	$= (\text{trap } \rightsquigarrow \mathbf{p})$
$\rightsquigarrow(\text{exit } n)$	$= (\text{exit } n)$

Fig. 20. Next Instant

The definition of Eval is written using a notation that assumes the central result of this paper, namely that Eval is a (partial) function:

THEOREM 3.2.

$$\begin{aligned} & \forall S_1, S_2, \theta, \mathbf{p}. \\ & \vdash_{\text{CB}} (\varrho \ \theta. \mathbf{p}) \Rightarrow \\ & \text{Eval}(\mathbf{p}, \theta) = S_1 \Rightarrow \\ & \text{Eval}(\mathbf{p}, \theta) = S_2 \Rightarrow \\ & S_1 = S_2 \end{aligned}$$

The above theorem states that if $\mathbb{S}_1$ and $\mathbb{S}_2$ are both sets of signals satisfying the Eval judgment in figure 19, then $\mathbb{S}_1$ and $\mathbb{S}_2$ must be equal. The proof is given as $\text{eval} \equiv_e$ -consistent in the Agda code in the supplementary material.

This theorem is a corollary of the consistency of $\equiv_e$, which states that if two expressions are $\equiv_e$, then there is an expression that both reduce to, under the transitive reflexive closure of the compatible closure of the reduction relation:

THEOREM 3.3.

$$\begin{aligned} \forall \mathbf{p}, \mathbf{q}. \vdash_{\text{CB}} \mathbf{p} \Rightarrow \mathbf{p} \equiv_e \mathbf{q} \Rightarrow \\ \exists \mathbf{r}. \mathbf{p} \rightarrow^* \mathbf{r} \wedge \mathbf{q} \rightarrow^* \mathbf{r} \end{aligned}$$

The proof is given as $\equiv_e$ -consistent, and it follows from the confluence of reduction.

Our semantics supports multiple instants via a transformation that prepares a complete expression for the next instant, $\Rightarrow$, shown in figure 20. It makes four modifications to the expression. First, it resets all signals to unknown and all shared variables to old via $\lfloor \theta^c \rfloor$ (defined in figure 13). Second, it replaces the pause expressions where the program stopped with nothing. Third, it replaces each loop expression with a loop and seq. Finally, it adds a present expression to suspend expressions that have paused. The present serves to conditionally pause the body of the suspend in the next instant. The result is an expression suitable for use with Eval in the next instant.

4 ON CONSTRUCTIVENESS

Logical correctness and constructiveness are key for any correct semantics of Esterel. For examples of these properties see section 2.4. We follow the definition of constructiveness given by the constructive operational semantics (COS) evaluator as referenced by Berry (2002) and described by Potop-Butucaru (2002): non-constructive programs reduce to stuck terms (that are not **complete**).

In our semantics, for many expressions, this is also the case. But, it is not the case for all of them because reductions that occur in arbitrary program contexts sometimes give Can more information than it “should” have (more precisely, more information that it would get by running the program directly). This extra information means that reductions in our calculus can transform some non-constructive programs into constructive ones that can still reduce.

For an example, consider the expression in figure 21. If we restrict our attention to the outside part of the term (the way that the COS semantics does), it reduces only by replacing the outer signal form with a q expression. At that point, the expression appears to be stuck because Can is unable to prove that S1 is not emitted (and thus the [absence] rule does not apply) and the present expression does not reduce (because S1 is unknown).

There are reductions that can occur, however, at the inner signal expression, revealing information to Can, and enabling it to determine that S1 is absent.

Specifically, the calculus can reduce in this context:

```
(q { S1 ↦ unknown }.
  (present S1
    []
    nothing))
```

```
(signal S1
  (present S1
    (signal S2
      (seq (emit S2)
        (present S2
          nothing
          (emit S1))))
    nothing))
```

Fig. 21: A Non-constructive Program

and thus it can turn the inner signal form into a q and perform the `emit`, resulting in the expression in figure 22. Being able to reduce in that context is effectively “peeking” ahead into the future non-constructively.

Once those reductions happen, Can is able to determine that S_1 cannot be emitted and now the `[absence]` rule can fire, eventually reducing the original expression to

$$(q \{ S_1 \mapsto \text{absent} \}. \text{nothing})$$

In sum, our calculus equates some non-constructive programs to constructive programs with the same logical behavior. Although we are not satisfied with this aspect of our calculus and believe that it deserves further study, such a relaxation of constructiveness is not unprecedented (Tardieu 2007).

```
(q { S1 ↦ unknown }.
  (present
   S1
   (q { S2 ↦ present }.
    (seq nothing
      (present S2
        nothing
        (emit S1))))))
  nothing))
```

Fig. 22: An expression equivalent to the expression in figure 21

5 WHAT THE CALCULUS CAN AND CANNOT PROVE

Our semantics lends itself to establishing equivalences between program fragments because any two expressions that are $\equiv_e$ to each other always produce the same result in the evaluator:

THEOREM 5.1.

$$\begin{aligned} &\forall \mathbf{p}, \boldsymbol{\theta}_1, \mathbb{S}_1, \mathbf{q}, \boldsymbol{\theta}_2, \mathbb{S}_2. \\ &\vdash_{\text{CB}} (q \boldsymbol{\theta}_1. \mathbf{p}) \Rightarrow \\ &(q \boldsymbol{\theta}_1. \mathbf{p}) \equiv_e (q \boldsymbol{\theta}_2. \mathbf{q}) \Rightarrow \\ &\text{Eval}(\mathbf{p}, \boldsymbol{\theta}_1) = \mathbb{S}_1 \Rightarrow \\ &\text{Eval}(\mathbf{q}, \boldsymbol{\theta}_2) = \mathbb{S}_2 \Rightarrow \\ &\mathbb{S}_1 = \mathbb{S}_2 \end{aligned}$$

This theorem is a straightforward consequence of $\equiv_e$ being consistent; the proof is given as $\equiv_e \Rightarrow \text{eval}$ in the Agda code in the supplementary material.

The remainder of this section explores various equivalences (shown in figure 23) as well as some limitations of the calculus. The proofs of the equivalences are all given in `agda/calculus-examples.agda` in the supplementary material.

The first example, theorem 5.2, shows that we can rearrange signal forms. This example works well in our calculus. It requires only that the body expression has correct binding, allowing us to rearrange adjacent signal forms arbitrarily.

Next, theorem 5.3 shows that if an `emit` is followed by a `present`, the `present` can always be replaced by the taken branch. This example exposes a first limitation of the calculus. Although it is still true, our calculus cannot prove this equivalence without the signal form being visible in an evaluation context surrounding the `seq` form.

In a dual to theorem 5.3, theorem 5.4 shows that if we know that neither branch of the `present` expression can emit S , we can replace the `present` form with its second subexpression.

Theorem 5.5 lets us lift a `seq` expression that starts with an `emit` out of a `par` branch. Intuitively, this equivalence is a consequence of Esterel’s deterministic parallelism. Because `emit` is instantaneous and does not depend on the status of any signal, we can do it in parallel to q or before q starts, whichever is more convenient.⁵

⁵This fact is crucial for many Esterel compilers, which attempt to generate static schedules for concurrent code (Potop-Butucaru et al. 2007).

THEOREM 5.2.

$$\begin{aligned} &\forall S_1, S_2, p. \\ &\vdash_{CB} p \Rightarrow \\ &(\text{signal } S_1 \\ &\quad (\text{signal } S_2 \\ &\quad \quad p)) \equiv_e \\ &(\text{signal } S_2 \\ &\quad (\text{signal } S_1 \\ &\quad \quad p)) \end{aligned}$$

THEOREM 5.3.

$$\begin{aligned} &\forall S, p, q. \\ &\vdash_{CB} p \Rightarrow \\ &(\text{signal } S \\ &\quad (\text{seq}(\text{emit } S) \\ &\quad \quad (\text{present } S \\ &\quad \quad \quad p \\ &\quad \quad \quad q))) \equiv_e \\ &(\text{signal } S \\ &\quad (\text{seq}(\text{emit } S) \\ &\quad \quad p)) \end{aligned}$$

THEOREM 5.4.

$$\begin{aligned} &\forall S, p, q. \vdash_{CB} q \Rightarrow \\ &(\forall \text{ status.} \\ &\quad S \notin (\text{Can } p \{ S \mapsto \text{status} \}).S) \Rightarrow \\ &(\forall \text{ status.} \\ &\quad S \notin (\text{Can } q \{ S \mapsto \text{status} \}).S) \Rightarrow \\ &(\text{signal } S \\ &\quad (\text{present } S \\ &\quad \quad p \\ &\quad \quad q)) \equiv_e \\ &(\text{signal } S \\ &\quad q) \end{aligned}$$

THEOREM 5.5.

$$\begin{aligned} &\forall S, p, q. \\ &\vdash_{CB} (\text{par } p \ q) \Rightarrow \\ &(\text{signal } S \\ &\quad (\text{par}(\text{seq} \\ &\quad \quad (\text{emit } S) \\ &\quad \quad \quad p) \\ &\quad \quad q)) \equiv_e \\ &(\text{signal } S \\ &\quad (\text{seq}(\text{emit } S) \\ &\quad \quad (\text{par } p \\ &\quad \quad \quad q)))) \end{aligned}$$

THEOREM 5.6.

$$\begin{aligned} &\forall n, p, q. \\ &\vdash_{CB} p \Rightarrow \\ &q \in \text{done} \Rightarrow \\ &p \equiv_e q \Rightarrow \\ &(\text{trap} \\ &\quad (\text{par}(\text{exit } n+1) \\ &\quad \quad p)) \equiv_e \\ &(\text{par}(\text{exit } n) \\ &\quad (\text{trap } p)) \end{aligned}$$

THEOREM 5.7.

$$\begin{aligned} &\forall p, q, S. \\ &\vdash_{CB} (\text{seq}(\text{signal } S \ p) \\ &\quad q) \Rightarrow \\ &(\varrho \{ \}. (\text{seq}(\text{signal } S \ p) \\ &\quad \quad q)) \equiv_e \\ &(\varrho \{ \}. (\text{signal } S \\ &\quad \quad (\text{seq } p \ q))) \end{aligned}$$

Fig. 23. Equivalences Provable in our Calculus

When a trap is outside a par, our calculus allows us to push the trap inside, in some situations. Theorem 5.6 is one such. This calculation requires p to be equivalent to some **done** expression q , but that is a weakness of our calculus. In fact, the two expressions are observably equivalent without any assumptions.

Theorem 5.7 further generalizes Theorem 5.2 to rearrange binding forms across other expressions. In this example, the signal form is pulled out of the seq expression. In general, these two expressions are observably equivalent even without the q expression outside. Our calculus cannot prove it, however, because the calculus needs an outer q expression in order to perform a [merge] in the middle of the proof.

We explored a calculus that includes a “lifting” rule that allows us to move a q term up and down in an evaluation context. This rule makes it difficult to establish confluence of the calculus, however, as the would-be lifting rule and the [merge] rule interact with each other in complex ways. In particular, our evaluation contexts do not have unique decomposition, due to par. Accordingly, a use of the lifting rule from one side of a par expression can block a use of the [merge] rule from the other side. We conjecture that a lifting rule would be confluent, but have not proven it. If we did have such a rule, then we believe we would be able to prove theorem 5.7 without the need for an enclosing empty q expression and even be able to relax one of the assumptions of theorem 5.6, assuming only that q is **complete**.

Our calculus also cannot reason effectively with expressions that span multiple instants. For example, the expression $(\text{seq } (\text{loop pause}) \ q)$ is equivalent to (loop pause) , but our calculus cannot prove it. Similarly, a common pattern is to emit a signal and pause in a loop, and also to run that loop in parallel with some code that looks at the signal. Our calculus would not be able to propagate the signal's presence because of the pause.

Another deep lack in our calculus is the ability to reason about input signals. In order for our calculus to work with a signal, it must be bound by signal so knowledge about it can be manipulated via the rules for q expressions. Input signals, in contrast, might or might not be set by the environment and our calculus cannot perform the required conditional reasoning.

6 TESTING

As we are developing a new semantics for Esterel and Esterel is a well-established language, a natural concern is whether our semantics captures Esterel or some other, subtly different language. In order to mitigate this concern, we tested our semantics against two Esterel implementations: Esterel v5 (Berry 2000) and Hiphop.js (Berry et al. 2011), as well as an executable version of Potop-Butucaru et al. (2007)'s COS semantics. Perhaps unsurprisingly, we also discovered bugs in both of the implementations during this process (as random testing can be extremely effective (Yang et al. 2011)). The remainder of this section describes the testing process and the bugs discovered.

6.1 Testing for Conformance

In order to test our model against the existing semantics and implementations, we had to build some software libraries:

- **Redex COS model** We built a model of the COS semantics in Redex (Felleisen et al. 2009). The semantics is a faithful model of the COS semantics because it is a rule-for-rule translation of the COS semantics; aside from a few syntactic differences (notably more parentheses), it mirrors Potop-Butucaru et al. (2007)'s model exactly, enabling us to simulate the behavior of the semantics on any example program.
- **Redex calculus model** Our calculus is implemented in Redex; the rules shown in all of the figures are generated automatically from the Redex source code, and the Redex model also enables us to explore the reduction of any example program.
- **Agda/Redex bridge** We built a library that can translate the reduction sequences generated by Redex into proofs in Agda, ensuring that the specific, concrete terms which reduce in Redex also reduce the same way in Agda. This process accepts a specific term and a reduction sequence. It produces a proof, which then is submitted to the Agda compiler for verification.
- **Redex/Hiphop.js bridge** We built a library that can translate Redex expressions into Hiphop.js programs and then evaluate them. We also built a translator for a subset of Hiphop.js programs that can translate them into Redex so they can be checked against the calculus and the COS model. This translator does not accept all Hiphop.js programs, because Hiphop.js programs embed JavaScript code and our model cannot evaluate the JavaScript.
- **Redex/Esterel v5 bridge** We also built a translator that produces Esterel v5 programs from Redex terms in the COS model and in our calculus.

Using these libraries we can test all four implementations of Esterel (the COS semantics, the Esterel v5 compiler, Hiphop.js, and our calculus) against each other.

There is one subtle point about testing our calculus. Because it is a calculus, we need an algorithm that can determine which of the many possible reductions we should take in order to find an effective path to a **complete** state (if one exists). To do this, we identified a subset of the possible reductions in a way that acts like a standard reduction, guaranteeing that we find a **complete** state if the

program is constructive, and that reduction gets stuck if it is not. This reduction relation is given in figure 27 with some supplementary definitions given in figure 28; it is explained briefly in section 7. We use this reduction relation to guide the calculus, verifying that each step in the standard reduction is also possible in the calculus. There are no proofs about our standard reduction but we use it only to test our calculus against other implementations, as described in this section.

These libraries give us the ability to, given an Esterel program, determine if it produces the same signals across multiple instants. But we also need a source of Esterel programs to test. For that purpose, we used two approaches.

First, we took the Hiphop.js test suite, which consists of 130 Hiphop.js programs. Of those, four use `pre`, a construct that is not in Kernel Esterel, and were excluded from our tests. An additional 84 use JavaScript in some non-trivial way, and therefore could also not be run in our model. Our calculus produces the same results on the remaining 42 program as Hiphop.js.

The translation of the Hiphop.js tests into our model produces programs that have a large number of signals, which causes problems for the process that finds reductions in the calculus. In short, the problem is that the exponential behavior in `Can` triggers significant performance problems in the calculus, enough so that running these tests appears not to be feasible. To mitigate this issue these tests are run only against the standard reduction, which updates signal values in bulk in a single step, and does not preform the exponential analysis on the top most environment, greatly reducing evaluation time.

Second, we used Redex's capability to generate random Esterel expressions and run them in all of the implementations to see if they agree. We have discovered (and fixed) errors in our calculus using this method, and we currently have no known bugs. We have run over 1,800,000 random tests and they still do periodically find counterexamples, but they find only known bugs in the implementations.

This random testing process proved invaluable in debugging the calculus, catching several errors that cannot be found via the proofs in Agda. For example, late in the development process, the random tester found that an old version of the `[shared]` rule was incorrect. The old version initialized the shared variables status to `new`, but the COS specifies that the initial status is `old`. This bug does not invalidate any of the theorems in Agda, but it does violate the property that our calculus and the other implementations agree. That is, the properties we can effectively check via random testing are stronger than those we can check via proof (in practice).

6.2 Bugs Discovered

During the process of validating our calculus, we discovered four bugs in Hiphop.js and one bug in the Esterel v5 compiler. All of the bugs have been confirmed by the authors of the systems. All but one of the Hiphop.js bugs have been fixed.

The bug in Esterel v5 is exhibited by a translation of the program in figure 24, where e_o evaluates to 0, e_i evaluates to 1, and e_{outer} refers to `s-outer`.

This program is non-constructive and it gets stuck in our both our calculus and the COS semantics; it cannot reduce the inner `shared` because the initialization of the signal depends on `s-outer`, but `s-outer` cannot be read because there is a write pending. The Esterel v5 compiler runs the program, incorrectly setting `s-inner` to 0.

This program also demonstrates one of the bugs we found in Hiphop.js. Of the other three bugs, one of them was an internal error, crashing Hiphop.js on the program (`trap (suspend (exit 0) S1))`).

```
(shared s-outer :=  $e_o$ 
  (seq (shared s-inner :=  $e_{outer}$ 
        nothing)
    (+= s-outer  $e_i$ )))
```

Fig. 24: A Bug Found in the Esterel v5 Compiler

$\text{set-absent} : \theta (\text{Setof } S) \rightarrow \theta$ $\text{set-absent}(\theta, S) =$ $\text{set-absent}(\theta \leftarrow \{s \mapsto \text{absent}\}, S \setminus \{s\})$ where $S \in S$ $\text{set-absent}(\theta, S) =$ θ	$\text{set-ready} : \theta (\text{Setof } s) \rightarrow \theta$ $\text{set-ready}(\theta, S) =$ $\text{set-ready}(\theta \leftarrow \{s \mapsto \langle n, \text{ready} \rangle\}, S \setminus \{s\})$ where $s \in S, \theta(s) = \langle n, \text{shared-status} \rangle$ $\text{set-ready}(\theta, S) =$ θ
$\Pi_{\parallel} : \text{done done} \rightarrow \text{done}$ $\text{nothing } \Pi_{\parallel} \text{ done} = \text{done}$ $\text{done } \Pi_{\parallel} \text{ nothing} = \text{done}$ $(\text{exit } n_1) \Pi_{\parallel} (\text{exit } n_2) = (\text{exit } \max(n_1, n_2))$ $(\text{exit } n) \Pi_{\parallel} \text{ paused} = (\text{exit } n)$ $\text{paused } \Pi_{\parallel} (\text{exit } n) = (\text{exit } n)$	

Fig. 26. Standard Reduction Auxiliary Metafunctions

The next bug was triggered by the expression (suspend nothing S1), and produced an error in terms of the undefined value from Hiphop.js's host language, JavaScript.

The final bug is exhibited by the program in figure 25. Both in our calculus and in the COS semantics, the Can function can determine that S-inner cannot be emitted, and that therefore S-outer cannot be emitted. Therefore the program is constructive, both signals are absent, and the program reduces to nothing. However this program appeared to be non-constructive to Hiphop.js.

```
(signal S-outer
  (signal S-inner
    (seq
      (present S-outer nothing nothing)
      (present S-inner
        (emit S-outer)
        nothing))))
```

Fig. 25: A Bug Found in Hiphop.js

7 STANDARD REDUCTION

Our standard reduction exists only in Redex (not in Agda, unlike the rest of the semantics). We use it to help with our testing process, as described in section 6.

Figure 27 shows the reduction rules. There are four differences between the rules of the calculus and the rules of the standard reduction. First, expressions reduce only if they have an outer q . Second, the [absence] and [readyness] rules set as many signals or variables as they can in a single step. Third, the [absence] and [readyness] rules use Can_q in the calculus and Can in the standard reduction. In the standard reduction, the extra analysis that Can_q performs is not necessary. Finally, the rules are oriented so that at most one applies at each step. There are two pieces to this orientation: restricting the context in which the rules may apply and restricting the [absence] and [readyness] rules so they apply only when no other rule applies.

To understand how the rules are oriented, consider the [absence] and [readyness] rules. They require the body to be either done or blocked, where blocked is given in figure 28. It captures when an expression cannot reduce because it needs the value of a signal or shared variable that is not known or ready.

The context restriction is captured by the $\theta \vdash E \text{ det}$ judgment. The judgment is designed to restrict the choice of sub-expression in par terms so only one side is considered for reduction.

signals	$(\varrho \ \theta. E[(\text{signal } S \ p)]) \rightarrow (\varrho \ \theta. E[(\varrho \ \{ S \mapsto \text{unknown} \}. p)])$ where $\theta \vdash E \text{ det}$ [signal]
	$(\varrho \ \theta. E[(\text{emit } S)]) \rightarrow (\varrho \ (\theta \leftarrow \{ S \mapsto \text{present} \}). E[\text{nothing}])$ [emit] where $\theta \vdash E \text{ det}$, $\theta(S) \in \{ \text{present}, \text{unknown} \}$
	$(\varrho \ \theta. p) \rightarrow (\varrho \ \text{set-absent}(\theta, S). p)$ [absence] where $(\theta \vdash p \text{ blocked or } p \in \text{done})$, $S = \{ S \in \text{dom}(\theta) \mid \theta^c(S) = \text{unknown} \} \setminus \text{Can}[\![p, \theta]\!].S, S \neq \emptyset$
	$(\varrho \ \theta. E[(\text{present } S \ p \ q)]) \rightarrow (\varrho \ \theta. E[p])$ where $\theta \vdash E \text{ det}$, $\theta(S) = \text{present}$ [is-present]
	$(\varrho \ \theta. E[(\text{present } S \ p \ q)]) \rightarrow (\varrho \ \theta. E[q])$ where $\theta \vdash E \text{ det}$, $\theta(S) = \text{absent}$ [is-absent]
shared variables	$(\varrho \ \theta. E[(\text{shared } s := e \ p)]) \rightarrow (\varrho \ \theta. E[(\varrho \ \{ s \mapsto \langle n, \text{old} \rangle \}. p)])$ [shared] where $\theta \vdash E \text{ det}$, $\text{FV}(e) \subset \text{dom}(\theta)$, $\forall s \in \text{FV}(e). \theta(s) = \langle _, \text{ready} \rangle$, $n = \mathcal{E}val[\![e, \theta]\!]$
	$(\varrho \ \theta. E[(+= s \ e)]) \rightarrow (\varrho \ (\theta \leftarrow \{ s \mapsto \langle n + \mathcal{E}val[\![e, \theta]\!], \text{new} \rangle \}). E[\text{nothing}])$ [set-new] where $\theta \vdash E \text{ det}$, $\theta(s) = \langle n, \text{new} \rangle$, $\text{FV}(e) \subset \text{dom}(\theta)$, $\forall s \in \text{FV}(e). \theta(s) = \langle _, \text{ready} \rangle$
	$(\varrho \ \theta. E[(+= s \ e)]) \rightarrow (\varrho \ (\theta \leftarrow \{ s \mapsto \langle \mathcal{E}val[\![e, \theta]\!], \text{new} \rangle \}). E[\text{nothing}])$ [set-old] where $\theta \vdash E \text{ det}$, $\theta(s) = \langle _, \text{old} \rangle$, $\text{FV}(e) \subset \text{dom}(\theta)$, $\forall s \in \text{FV}(e). \theta(s) = \langle _, \text{ready} \rangle$
	$(\varrho \ \theta. p) \rightarrow (\varrho \ \text{set-ready}(\theta, S_2). p)$ [readyness] where $(\theta \vdash p \text{ blocked or } p \in \text{done})$, $\{ S \in \text{dom}(\theta) \mid \theta^c(S) = \text{unknown} \} \setminus \text{Can}[\![p, \theta]\!].S = \emptyset$, $S_1 = \{ s \in \text{dom}(\theta) \mid \theta^c(s) = \langle \text{ev}, \text{shared-status} \rangle \}$, $\text{shared-status} \in \{ \text{new}, \text{old} \}$, $S_2 = S_1 \setminus \text{Can}[\![p, \theta]\!].\text{sh}, S_2 \neq \emptyset$
sequential variables	$(\varrho \ \theta. E[(\text{var } x := e \ p)]) \rightarrow (\varrho \ \theta. E[(\varrho \ \{ x \mapsto \mathcal{E}val[\![e, \theta]\!]. p)])$ [var] where $\theta \vdash E \text{ det}$, $\text{FV}(e) \subset \text{dom}(\theta)$, $\forall s \in \text{FV}(e). \theta(s) = \langle _, \text{ready} \rangle$
	$(\varrho \ \theta. E[(:= x \ e)]) \rightarrow (\varrho \ (\theta \leftarrow \{ x \mapsto \mathcal{E}val[\![e, \theta]\!]. p)])$ [set-var] where $\theta \vdash E \text{ det}$, $x \in \text{dom}(\theta)$, $\text{FV}(e) \subset \text{dom}(\theta)$, $\forall s \in \text{FV}(e). \theta(s) = \langle _, \text{ready} \rangle$
	$(\varrho \ \theta. E[(\text{if } x \ p \ q)]) \rightarrow (\varrho \ \theta. E[p])$ where $\theta \vdash E \text{ det}$, $x \in \text{dom}(\theta)$, $\theta(x) \neq 0$ [if-true]
	$(\varrho \ \theta. E[(\text{if } x \ p \ q)]) \rightarrow (\varrho \ \theta. E[q])$ where $\theta \vdash E \text{ det}$, $\theta(x) = 0$ [if-false]
seq	$\varrho \ (\varrho \ \theta_1. E[(\varrho \ \theta_2. p)]) \rightarrow (\varrho \ (\theta_1 \leftarrow \theta_2). E[p])$ where $\theta_1 \vdash E \text{ det}$ [merge]
	$(\varrho \ \theta. E[(\text{seq nothing } q)]) \rightarrow (\varrho \ \theta. E[q])$ where $\theta \vdash E \text{ det}$ [seq-done] $(\varrho \ \theta. E[(\text{seq (exit } n) \ q)]) \rightarrow (\varrho \ \theta. E[(\text{exit } n)])$ where $\theta \vdash E \text{ det}$ [seq-exit]
trap	$(\varrho \ \theta. E[(\text{trap stopped})]) \rightarrow (\varrho \ \theta. E[\downarrow^p \text{ stopped}])$ where $\theta \vdash E \text{ det}$ [trap]
par	$(\varrho \ \theta. E[(\text{par stopped done})]) \rightarrow (\varrho \ \theta. E[\text{stopped } \sqcap_{\parallel} \text{done}])$ where $\theta \vdash E \text{ det}$ [parr]
	$(\varrho \ \theta. E[(\text{par paused stopped})]) \rightarrow (\varrho \ \theta. E[\text{paused } \sqcap_{\parallel} \text{stopped}])$ where $\theta \vdash E \text{ det}$ [parl]
	$(\varrho \ \theta. E[(\text{suspend stopped } S)]) \rightarrow (\varrho \ \theta. E[\text{stopped}])$ where $\theta \vdash E \text{ det}$ [suspend]
	$(\varrho \ \theta. E[(\text{loop } p)]) \rightarrow (\varrho \ \theta. E[(\text{loop } p \ p)])$ where $\theta \vdash E \text{ det}$ [loop]
	$(\varrho \ \theta. E[(\text{loop } (\text{exit } n) \ q)]) \rightarrow (\varrho \ \theta. E[(\text{exit } n)])$ where $\theta \vdash E \text{ det}$ [loop ^ stop-exit]

Fig. 27. Standard Reduction Rules

$\frac{\theta \vdash e \text{ blocked}}{\theta \vdash (\text{var } x := e \ p) \text{ blocked}}$		$\frac{\theta \vdash e \text{ blocked}}{\theta \vdash (:= x \ e) \text{ blocked}}$	$\boxed{\theta \vdash p \text{ blocked}}$
$\frac{\theta(S) = \text{unknown}}{\theta \vdash (\text{present } S \ p \ q) \text{ blocked}}$	$\frac{\theta \vdash p \text{ blocked}}{\theta \vdash (\text{suspend } p \ S) \text{ blocked}}$	$\frac{\theta \vdash p \text{ blocked}}{\theta \vdash (\text{trap } p) \text{ blocked}}$	
$\frac{\theta \vdash p \text{ blocked} \quad \theta \vdash q \text{ blocked}}{\theta \vdash (\text{par } p \ q) \text{ blocked}}$	$\frac{\theta \vdash p \text{ blocked}}{\theta \vdash (\text{par } p \ \text{done}) \text{ blocked}}$	$\frac{\theta \vdash q \text{ blocked}}{\theta \vdash (\text{par } \text{done} \ q) \text{ blocked}}$	
$\frac{\theta \vdash p \text{ blocked}}{\theta \vdash (\text{seq } p \ q) \text{ blocked}}$		$\frac{\theta \vdash p \text{ blocked}}{\theta \vdash (\overline{\text{loop}} \ p \ q) \text{ blocked}}$	
$\frac{\theta \vdash e \text{ blocked}}{\theta \vdash (\text{shared } s := e \ p) \text{ blocked}}$		$\frac{\theta \vdash e \text{ blocked}}{\theta \vdash (+= s \ e) \text{ blocked}}$	
$\frac{s \in \text{FV}(e) \quad \theta(s) = \langle n, \text{old} \rangle}{\theta \vdash e \text{ blocked}}$	$\frac{s \in \text{FV}(e) \quad \theta(s) = \langle n, \text{new} \rangle}{\theta \vdash e \text{ blocked}}$	$\boxed{\theta \vdash e \text{ blocked}}$	
$\frac{\theta \vdash E \text{ det}}{\theta \vdash (\text{seq } E \ q) \text{ det}}$	$\frac{\theta \vdash E \text{ det}}{\theta \vdash (\overline{\text{loop}} \ E \ q) \text{ det}}$	$\frac{}{\theta \vdash [] \text{ det}}$	$\boxed{\theta \vdash E \text{ det}}$
$\frac{\theta \vdash E \text{ det}}{\theta \vdash (\text{suspend } E \ S) \text{ det}}$		$\frac{\theta \vdash E \text{ det}}{\theta \vdash (\text{trap } E) \text{ det}}$	
$\frac{\theta \vdash E \text{ det}}{\theta \vdash (\text{par } E \ q) \text{ det}}$	$\frac{\theta \vdash E \text{ det}}{\theta \vdash (\text{par } \text{done} \ E) \text{ det}}$	$\frac{\theta \vdash E \text{ det} \quad \theta \vdash p \text{ blocked}}{\theta \vdash (\text{par } p \ E) \text{ det}}$	

Fig. 28. Standard Reduction Auxiliary Judgments

There are three par rules: one that always allows reductions in the left-hand side, and two others that allow reduction on the right, but only when the left is either done or blocked.

Otherwise, the reduction rules in the standard reduction parallel those in calculus. Of course, we do not take the compatible closure; instead we just reduce in evaluation contexts.

There is one subtle point of this standard reduction: it is not the standard reduction relation for our calculus, but rather the one for a slightly smaller calculus. In particular, it does not bypass constructiveness, unlike the calculus (as discussed in section 4). It reaches a stuck state instead.

8 RELATED WORK

Three decades of work on Esterel have resulted in a diversity of semantic models. A fundamental difference of our semantics is that ours is a calculus rather than merely a reduction system—it has an equational theory. In general, the nature of our semantics makes it better for discussing transformations to Esterel programs and, more specifically, it lets us prove equivalences in arbitrary contexts, like those discussed in section 5. By the same token, our semantics is not particularly well-suited to implementing Esterel.

Prior semantics of Esterel can be broadly categorized as follows:

- Macrostep operational semantics compute the result of an instant in big-step style, where evaluation relates the state of a program at the beginning of an instant to the state at the end.
- Microstep operational semantics compute the result of an instant as a series of small-step style transitions until the instant is considered terminated. Our semantics is in this style.
- Circuit semantics gives meaning to Esterel programs by translation to Boolean circuits.

Semantics of Esterel are also classified as logical or constructive. Logical semantics are simpler, but give meaning to programs that are logically correct but non-constructive. Constructive semantics use constructive information propagation to enforce both (Berry 2002).

Finally, semantics of Esterel are distinguished by how much of the language they cover. Some cover all of Kernel Esterel, while others cover only Pure Esterel, which omits shared and sequential variables.

Berry and Gonthier (1992) give two operational semantics of Esterel, a macrostep logical semantics called the behavioral semantics and a microstep logical semantics called the execution semantics. They have proved these equivalent and, for the latter, proved a confluence theorem.

Berry (2002) gives an update to the logical behavioral (macrostep) semantics to make it constructive. The logical behavioral semantics requires existence and uniqueness of a behavior, without explaining how it could be computed, while the constructive behavioral semantics introduces *Can* to do it in an effective but restricted way. Berry (2002) also gives the state behavioral semantics, another macrostep semantics.

The constructive operational semantics (COS) first appears in Potop-Butucaru (2002)’s thesis. It is a microstep semantics that uses program decorations track control flow and avoid rewriting the program. The COS model, like the constructive behavioral semantics, avoids giving meaning to non-constructive programs, but unlike the behavioral semantics, it provides a guide toward efficient implementation.

Like some of those semantics, our semantics handles the larger language (Kernel Esterel) and accounts for constructiveness. Unlike all of those semantics, our semantics works by term rewriting, substituting equals for equals and simplifying programs, which makes it a good basis for proving program fragment equivalences.

Circuit semantics, such as those that appear in Berry (2002) and Potop-Butucaru et al. (2007), are the semantics generally used by Esterel implementations like Hiphop.js and Esterel v5. These implementations use circuits as an intermediate representation during compilation.

These semantics can be used for program optimization (as ours can too), but in a different way than ours. Because the translation to circuits is complex, it is difficult to connect transformations done at the circuit level back to the source level, so these semantics would not form a good basis for, say, refactoring tools, or other more human-centered tasks. Another contrast is that these semantics are much better suited to whole-program optimizations (e.g., using existing CAD tools to simplify the circuit) whereas our semantics is better suited to local transformations.

Additional semantics include Tardieu (2007)’s, who uses a different technique than constructiveness to eliminate logically incorrect programs. It handles only Pure Esterel.

Our semantics of Esterel closely follows the work of Felleisen and Hieb (1992) on semantics for programs with state. In particular, we borrow the idea of internalizing state into terms using a q term that binds a partial store embedded at any level in a term. However, unlike Felleisen and Hieb, because of Esterel’s par construct, we do not have unique decomposition into an evaluation context and redex. This means that the **[merge]** rule can merge from both sides of a par into the same position, which complicates our proofs. Our semantics also has to handle many Esterel-specific notions, which are not a concern in Felleisen and Hieb (1992)’s work.

ACKNOWLEDGMENTS

Thanks to Daniel Feltey, Christos Dimoulas, and Vincent St-Amour for their feedback on drafts of this paper. Thanks to Colin Vidal, Gérard Berry, and Manuel Serrano for their help with Hiphop.js, Esterel v5 and, more generally, understanding Esterel’s semantics and feedback on this work. Thanks to Matthias Felleisen for his help with how to structure the proofs and help understanding the semantics of state. A special thanks to Gérard Berry for suggesting that we use Newman’s lemma; without it we would probably never have finished the proof of confluence. This work was supported by NSF proposal CCF-1526109.

BIBLIOGRAPHY

- H. Barendregt. The Lambda Calculus: its Syntax and Semantics. North Holland, 1984.
- Albert Benveniste and Gérard Berry. The Synchronous Approach to Reactive and Real-Time Systems. *Proceedings of the IEEE* 79(9), 1991.
- Albert Benveniste, Paul Caspi, Stephen A. Edwards, Nicolas Halbwachs, Paul Le Guernic, and Robert de Simone. The Synchronous Languages 12 Years Later. *Proceedings of the IEEE* 91(1), 2002.
- Gérard Berry. The Esterel v5 Language Primer Version v5_91. École des Mines de Paris, CMA and INRIA, Sophia-Antipolis, France, 2000.
- Gérard Berry. The Constructive Semantics of Pure Esterel (Draft Version 3). 2002.
- Gérard Berry, Amar Bouali, Xavier Fornari, Emmanuel Ledinot, Eric Nassor, and Robert de Simone. ESTEREL: A Formal Method Applied to Avionic Software Development. *Science of Computer Programming* 36(1), pp. 5–25, 2000.
- Gérard Berry and Georges Gonthier. The Esterel Synchronous Programming Language: Design, Semantics, Implementation. *Science of Computer Programming* 19(2), pp. 87–152, 1992.
- Gérard Berry, Cyprien Nicolas, and Manuel Serrano. HipHop: A Synchronous Reactive Extension for Hop. In *Proc. PLASTIC*, 2011.
- Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. Semantics Engineering with PLT Redex. MIT Press, 2009.
- Matthias Felleisen and Robert Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science* 103(2), pp. 235–271, 1992.
- Matthew Flatt and PLT. Reference: Racket. PLT Inc., PLT-TR-2010-1, 2010. <http://racket-lang.org/tr1/>
- Spencer P. Florence, Burke Fetscher, Matthew Flatt, William H. Temps, Tina Kiguradze, Dennis P. West, Charlotte Niznik, Paul R. Yarnold, Robert Bruce Findler, and Steven M. Belknap. POP-PL: A Patient-Oriented Prescription Programming Language. In *Proc. GPCE*, 2015.
- Lindsey Kuper and Ryan R. Newton. LVars: lattice-based data structures for deterministic parallelism. In *Proc. Workshop on Functional High-performance Computing (FHPC)*, 2013.
- Michael Mendler, Thomas R. Shiple, and Gérard Berry. Constructive Boolean circuits and the exactness of timed ternary simulation. *Formal Methods in System Design* 40, pp. 283–329, 2012.
- Benjamin Pierce. Types and Programming Languages. MIT Press, 2002.

- Gordon Plotkin. Call-by-name, Call-by-value, and the λ -Calculus. *Theoretical Computer Science* 1(2), pp. 125–159, 1975.
- Dumitru Potop-Butucaru. Optimizations for Faster Simulation of Esterel Programs. PhD dissertation, Ecole des Mines de Paris, 2002.
- Dumitru Potop-Butucaru, Stephen A. Edwards, and Gérard Berry. Compiling Esterel. Springer, 2007.
- Klaus Schneider and Michael Wenz. A New Method for Compiling Schizophrenic Synchronous Programs. In *Proc. International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES)*, 2001.
- Manuel Serrano and Pierre Weis. Bigloo: a portable and optimizing compiler for strict functional languages. In *Proc. Static Analysis Symposium*, 1995.
- Olivier Tardieu. A Deterministic Logical Semantics for Pure Esterel. *ACM Transactions on Programming Languages and Systems* 29(2), 2007.
- Xuejun Yang, Yang Chen, Eric Eide, and John Regehr. Finding and Understanding Bugs in C Compilers. In *Proc. Programming Language Design and Implementation (PLDI)*, 2011.