Denotational Fixed-Point Semantics for Constructive Scheduling of Synchronous Concurrency

Joaquín Aguado · **Michael Mendler** ·
Reinhard von Hanxleden · **Insa Fuhrmann**

Abstract The synchronous model of concurrent computation (SMoCC) is well established for programming languages in the domain of safety-critical reactive and embedded systems. Translated into mainstream C/Java programming, the SMoCC corresponds to a cyclic execution model in which concurrent threads are synchronised on a *logical clock* that cuts system computation into a sequence of *macro-steps*. A *causality analysis* verifies the existence of a schedule on memory accesses to ensure each macro-step is deadlock-free and determinate.

We introduce an abstract semantic domain $I(\mathbb{D}, \mathbb{P})$ and an associated denotational fixed point semantics for reasoning about concurrent and sequential variable accesses within a synchronous cycle-based model of computation. We use this domain for a new and extended behavioural definition of Berry's causality analysis in terms of approximation intervals. The domain $I(\mathbb{D}, \mathbb{P})$ extends the domain $I(\mathbb{D})$ from our previous work and fixes a mistake in the treatment of initialisations.

Based on this fixed point semantics we propose the notion of *Input Berry-constructiveness* (IBC) for synchronous programs. This new IBC class lies properly between *strong* (SBC) and *normal Berry-constructiveness* (BC) defined in previous work. SBC and BC are two ways to interpret the standard constructive semantics of synchronous programming, as exemplified by imperative SMoCC languages such as Esterel or Quartz. SBC is often too restrictive as it requires all variables to be initialised by the program. BC is too permissive because it initialises all variables to a fixed value, by default. Where the initialisation happens through the memory, *e.g.*, when carrying values from one synchronous tick to the next, then IBC is more appropriate.

IBC links two levels of execution, the macro-step level and the micro-step level. We prove that the denotational fixed point analysis for IBC, and hence Berry's causality analysis, is sound with respect to operational micro-level scheduling. The denotational model can thus be viewed as a compositional presentation of a synchronous scheduling strategy that ensures reactiveness and determinacy for imperative concurrent programming.

Joaquín Aguado · Michael Mendler
Faculty of Information Systems and Applied Computer Sciences, Bamberg University, Germany
E-mail: {joaquin.aguado,michael.mendler}@uni-bamberg.de

Reinhard von Hanxleden · Insa Fuhrmann
Department of Computer Science, Kiel University, Germany
E-mail: {rvh, ima}@informatik.uni-kiel.de

DRAFT 20.4.2015

1 Introduction

1.1 Motivation

Arguably the mathematically most satisfactory way to define a compositional programming
language semantics is the denotational approach which defines the semantics of a program
through a system of structurally-recursive equations involving continuous functions on
abstract semantic domains. Compositionality is built into a denotational model at the outset,
in the sense that the functional defining of the fixed-point semantics of a composite program
entirely depends on the abstract functional denotation of its components rather than their
syntax. As a consequence, algebraic axiomatisations for program verification and program
transformations can be derived from the properties of these functions in the abstract domains.

Unfortunately, denotational fixed point models for computationally rich programming
languages are notoriously hard to come by. A famous case in point is the long search for a
fully-abstract denotational model of the functional language PCF [13,1,45]. It is the tight
interaction of program components, in particular for non-deterministic concurrent systems,
that makes it hard to decouple a composite program into a system of continuous functions
in a simple way. It is often easier to understand the interaction behaviour of a concurrent
program operationally in terms of inductive relations rather than recursive functions. Hence,
many concurrent programming models or process algebras, for that matter, are based on
Plotkin-style structural operational semantics. Such models are then turned into an algebra
through notions of behavioural congruences and pre-congruences. Thereby abstracting from
behaviourally unobservable information carried by the operational rule system one achieves
the desired algebraic compositionality, see, *e.g.*, [66,10]. However, denotational semantics
generated in this fashion are essentially syntactic. Recursion is not explained by denotational
fixed points but by syntactic unfolding.

One can do better if the inductive operational rules satisfy certain structural constraints,
such as the GSOS or tyft/tyxt format. In these cases, general techniques are known to
derive independent denotational semantics based on the approximation of a process by finite
synchronization trees, see *e.g.* [2,46,29] for full-abstraction results for bisimulation-style
semantics. Still, these approximation-based denotational models have their own problems.
They are algebraically rather involved and depend on infinitary proof rules which fall outside
the scope of normal (Horn-style) equational reasoning. One classical instance of this problem
is the observation that, *e.g.*, bisimulation equivalence for process algebras with the empty
process 0, non-deterministic choice $p + q$, action prefix $a. p$ and recursion $\mu x. p(x)$, does not
admit a finitary denotational semantics based on complete partial orderings. Specifically,
the Park induction principle, $p(y) = y \Rightarrow \mu x. p(x) \leq y$, expressing that $\mu x. p(x)$ is the least
fixed point is inconsistent with monotonicity of the choice operator $+$. It is unclear if
bisimulation-style semantics can be finitely axiomatised in equational Horn logic, see *e.g.* [64].
Denotational fixed point semantics with Park induction seem to exist only in special cases,
such as acceptance testing, trace equivalence or simulation preorder [41,50,28].

While it is now clear that complex algebraic machinery is needed to reconcile genuinely
independent denotational and operational semantics for general non-deterministic process

DRAFT 20.4.2015

calculi, attention should turn once again to more special concurrent programming models. An early and successful starting point is the data-flow semantics of Kahn networks [47] which is fully-abstract for coroutine-style operational execution [48]. Kahn process nodes are sequential and deterministic and thus fairly restricted in modelling distributed systems. Yet, as discovered by Kok [49], non-determinism can be added to the Kahn model without losing full-abstraction using *(local) clocks* for the synchronisation of streams. This remarkable result brings into view an important special class of concurrent programming languages where denotational and operational approaches may go well together, known as the synchronous programming paradigm [33,9].

The Synchronous Model of Concurrent Computation (SMoCC) started in the 1980s with languages such as Statecharts [39,70], Lustre [19,34], Signal [32], Esterel [14,11] and Argos [62,63]. Developing concurrently with the emerging theory of process algebras, the SMoCC, from its beginning, has taken a practical programming perspective and targeted embedded and safety-critical systems in the automotive and avionics industries. The SMoCC languages have been very successful in these highly-demanding and complex domains. Part of this is due to their solid mathematical underpinning which inherits its robust logic from the design of digital synchronous circuits. Over the years, the quality-software assurance of the SMoCC paradigm has received attention in a wider range of applications. These include simulation Stateflow [35], web-orchestration [15], and music accompaniment [7] to mention a few. The SMoCC approach also has spread into functional programming [61] and mainstream imperative languages like C [16,52,36] or Java [67].

The SMoCC paradigm is based on a globally synchronous, locally asynchronous model of concurrent computation[1], which employs logical clocks to force asynchronous processes into a globally deterministic sequence of execution steps, called *macro steps* or *logical instants*. The SMoCC computations relate to classical automata in the sense that macro-steps correspond to automata transitions and configurations are discrete time points (automata states) on which system and environment can communicate (synchronise) with each other. At this level of modelling—under the Synchrony Hypothesis [9]—macro-steps appear as deterministic and functional input/output interactions. If this were all, synchronous programs could be analysed by the standard compositional techniques of the theory of synchronous automata, which fits both the denotational and the operational viewpoint equally well. Not much concurrency theory is needed for that.

However, there is a catch: The soundness of the automata model depends on the compiler verifying that the Synchrony Hypothesis is valid. Yet, the Synchrony Hypothesis is not compositional. The difficulty is that the SMoCC programs exhibit Mealy as opposed to Moore-style interaction. Since Mealy outputs depend instantaneously on the inputs and (in a typical SMoCC language) are also broadcast, the atomicity assumption creates a tangled causality cycle when the SMoCC automata are composed. Since each program acts as the environment of the other, the Synchrony Hypothesis expects each system to react faster than the other, and hence faster than itself! This is aggravated by the fact that in some SMoCC languages, such as Esterel or some version of Statecharts, the reaction of one component can depend on the absence of a reaction from another component. To resolve the paradoxes, *i.e.*, to prevent deadlock and non-determinism, the synchronous interaction must satisfy stringent causality requirements. Consequently, causality analyses have been a key component in the SMoCC compilers. Typically, these analyses correspond to the derivation of clock schedules

---

[1] This is sometimes referred to as the 'LAGS' model and not to be confused with the well-known but orthogonal 'GALS' model which features globally asynchronous and locally synchronous computations.

("clock calculus") for the activation of program statements [20, 8, 22, 80] or 3-valued circuit simulation ("ternary analysis") [12, 26, 76]. Edwards [25] and Potop-Butucaru et al. [71] provide good overviews of compilation challenges and approaches for concurrent languages, including synchronous languages. In this article we focus on causality in control-flow oriented SMoCCs such as Esterel or Quartz rather than data-flow oriented SMoCCs such as Lustre or Signal.

The techniques for causality analysis range from checking simple static criteria on control-dependencies to full-fledged data-dependent control-flow analysis. Proving the soundness of causality analyses necessarily requires maintaining some form of refinement ("constructive-ness" or "dependency") information about a lower-level asynchronous *micro-step* semantics. The first to observe this were Huizing, Gerth and de Roever [44] who showed that combining compositionality, causality and the Synchrony Hypothesis cannot be done within a single-levelled semantics (see also [73]). In other words, causality analysis establishes consistency of a *synchronous* macro-step with respect to an *asynchronous* micro-step execution model. This makes causality analyses and their soundness properties interesting from a concurrency theoretical point of view.

Different distributed execution platforms and memory models induce different degrees of uncontrollable non-determinism. They give thus rise to different notions of causality. A conservative, and thus robust, notion of causality among all the SMoCCs is the so-called *constructive* semantics of Esterel [14, 11] introduced by Berry in [12]. This is a pure macro-step semantics combining a structural operational semantics for macro-state transitions with a denotational fixed-point construction, also known as the *"must-cannot"* analysis, for computing causal reactions from every state. However, there do not seem to be soundness proofs for the causality analyses of Esterel relative to a micro-level scheduling in normal, i.e., unsynchronised memory.[2] The only available result on the lower-level operational soundness of the fixed-point construction is indirect, has never formally been proven and applies to the hardware translation given in [12]. At the hardware level it is known that constructiveness implies delay-insensitivity under non-inertial delays [59, 78, 65]. While this highlights the universal nature of the constructive semantics, it does not provide insights into the nature of constructiveness for software implementations of SMoCC languages. This question is now starting to be addressed in the literature. An interesting example is the more recent SMoCC language Quartz [74]. It has been given both macro-step operational semantics and a fixed-point semantics implementing an Esterel-style causality analysis [74, 31]. Talpin et. al. in [79] consider a combination of Signal and Quartz and prove that the constructive fixed-point semantics is sound for an operational micro-step semantics. In this article we proceed along similar ideas as [79] for Esterel-style imperative languages.

1.2 Contributions

In this article we prove the soundness of the denotational fixed point semantics for impera-tive SMoCC programs, commonly termed "constructive", with respect to their micro-step operational behaviour when compiled into multi-threaded shared memory code. To the best of our knowledge this is the first result of its kind for Esterel-style imperative programming.

---

[2] There is an informal sketch of a micro-step semantics in [12][Sec. 4.3] which is not developed further or formally related with the fixed point semantics for macro-steps.

DRAFT 20.4.2015

The recent constructive semantics that integrates Quartz and Signal [79] is based on a similar approach than the one proposed here in the sense that both are developed around similar mathematical structures, *i.e.*, fixed point on a lattice for representing signal statuses and Boolean values. The semantics framework of [79] unifies the behaviour of polychronous multi-clocked Signal networks and synchronous Quartz modules where synchronous Boolean variables are always present. In contrast, our approach significantly extends the standard 3-valued *"must-cannot"* semantics [12, 26, 76] with the effect that (i) it is able to handle *explicit initialisation* of signals, and (ii) it operates in a more structured domain of *information intervals* rather than flat ternary Kleene algebra. In the enriched domain we prove soundness of the fixed-point with respect to the micro-step operational execution. By "micro-step operational semantics" we mean a small-step semantics in which the reaction of a parallel composition for a single clock tick (rather than sequences of clocks) is (1) implemented by thread interleaving and (2) the execution does not use the must/cannot enriched statuses. E.g., the SOS reaction rules for Quartz in [74, 31] do not satisfy criterion (1). They give big-step semantics for full reaction instants. On the other hand, the operational semantics sketched by Berry in [12][Chap. 4.3] or by Talpin et.al. in [79] do not satisfy criterion (2).

Our main Theorem 1 strengthens the results presented in [4] where a similar fixed point semantics was introduced to prove that the sequentially constructive model of synchronous concurrent computation [37, 38] conservatively extends Berry's notion of constructiveness for Esterel. Specifically, we extend the work of [4] in three ways: Firstly, we correct a mistake preventing the denotational semantics of [4] from detecting deadlocks that can arise from concurrent initialisations (see our Ex. 19). Secondly, the results presented here imply that the fixed point analysis is not only sound for sequential constructiveness targeted in [4] but also for Esterel's more restrictive operational model of causality, characterised by *B-reactiveness* (Def. 4) and *SC-read-determinacy* (Def. 5). The combination of these two properties is a proper strengthening of the notion of Δ_*-constructiveness in the sense of [4], which corresponds to the notion of sequential constructiveness introduced in [37]. Thirdly, we introduce a new definition of constructiveness, called *IB-constructiveness* (Def. 9), to permit implicit initialisations through memory. It is more generous than the notions of constructiveness considered in [4] where all variables must be reinitialised, by the program or the environment, at every macro step. In other words, compared to [4] our semantics guarantees a *stronger* form of operational robustness for a *wider* class of programs.

1.3 Overview

Section 2 gives an abstract account of the SMoCC principle for imperative programs based on the consolidated language model pSCL and the operational notion of free scheduling. It also offers the definitions of important terms that will be used in the following sections, particularly of B-Admissibility and SC-Admissibility, which are both scheduling protocols restricting the free scheduling with different degrees of strength. The related terms of B-Reactiveness and SC-Reactiveness are also defined as well as the notion of X-Determinacy, parametric in X-Admissibility and its special case X-read-determinacy.

Section 3 is dedicated to the definition of the abstract domains and environments on which our denotational fixed point semantics is based. This includes the definition of the domain $\mathbb{D}$, whose four elements represent possible signal statuses and comprise representations needed for the handling of explicit initialization of signals. The semantics operates on closed intervals over $\mathbb{D}$ which represent predictions of variable statuses combined with a domain $\mathbb{P}$ that is

DRAFT 20.4.2015

capturing initialization statuses, yielding $I(\mathbb{D}, \mathbb{P})$ as working domain. Finally, the section introduces a domain $I(\mathbb{C})$ of program completion statuses.

Section 4 is the core of the article, where we put the introduced technical apparatus to form our denotational fixed point semantics for pSCL. The semantics induces three notions of constructiveness increasing in strength, *Berry-constructive* (BC), *Input Berry-constructive* (IBC) and *Strong Berry constructive* (SBC). This section finally contains our main Soundness Theorem 1 that states that IBC programs are B-reactive and SC-read-determinate.

Section 5 concludes this article and positions our work in reference to related work.

2 Operational Semantics of Synchronous Programs

2.1 Language Model

For our elaborations, we employ a language that focuses on the micro-step computations of a system. This language, referred to as pSCL[3] contains the necessary control structures for capturing multiple variable accesses as they occur inside macro-steps. pSCL abstracts syntactic and control particularities of existing synchronous languages not directly related to our analysis. This not only provides generality to the results but also avoids over-complicating our formal treatment. pSCL is *pure* in the sense that it manipulates Boolean *variables* from a finite set V, which carry information over time by changing value in $\mathbb{B} = \{0, 1\}$. A variable $s \in V$ with value $\gamma \in \mathbb{B}$ is denoted by s^γ. Here, $0, 1$ are used to code, respectively, the logical statuses *False* (absent, initialised) and *True* (present, updated) of a synchronous signal. The syntax of pSCL is given by the following BNF of operators:

$$
\begin{array}{llll}
P := & \varepsilon & \text{nothing} & \\
\mid & \pi & \text{pause} & \\
\mid & {\text{¡}}s & \text{s} = \text{false} & \text{(implicit unemit s in Esterel)} \\
\mid & !s & \text{s} = \text{true} & \text{(emit s in Esterel)} \\
\mid & s\ ?\ P : P & \text{if } s \text{ then } P \text{ else } P & \text{(present s then P else P in Esterel)} \\
\mid & P \| P & \text{fork } P \text{ par } P \text{ join} & \\
\mid & P\ ;\ P & & \\
\mid & rec\ p.\ P & p : P & \text{declare program label (implicit in Esterel loop)} \\
\mid & p & \text{goto } p & \text{jump to label (generalises Esterel iteration)}
\end{array}
$$

Since our syntax is abstract in the style of process algebras we also indicate the more concrete syntax as used in control-flow languages SCL [37] and Esterel on the right of each operator.

Intuitively, the *empty* statement ε indicates that a given program has been completed. That is, ε corresponds to the termination situation in which there are no further tasks to be performed in this or any subsequent macro-step. The *pause* control π forces a program to yield and wait for a global tick. This means that the execution cannot not proceed any further during the current macro-step but it will be resumed in the next instant. The *reset* (init) ¡s and *set* (update) $!s$ constructs modify the value of $s \in V$ to s^0 or s^1, respectively. The *conditional* control $s\ ?\ P : Q$ has the usual interpretation in the sense that depending on the status 1 or 0 of the guard variable s either P or Q are executed accordingly. *Parallel*

---

[3] This stands here for "pure Synchronous Constructive Language" indicating not only that signal variables in pSCL carry Boolean status but also that pSCL is a minimalistic version of control-flow synchronous languages in an abstract algebraic syntax.

DRAFT 20.4.2015

composition $P \parallel Q$ forks P and Q, so the statements of both are executed concurrently. This composition terminates (joins) when both components terminate, *i.e.*, both are completed in the sense of ε, not waiting in a pause π. When just one of the two components in $P \parallel Q$ terminates while the other pauses, then $P \parallel Q$ pauses. Otherwise, if one component terminates and the other does not pause or terminate then the computation continues from the statements of the other component until it terminates, too, or pauses. In the *sequential* composition $P \,;\, Q$, the statements of P are first completely executed. Then, the control is transferred to Q which, in turn, determines the behaviour of the composition thereafter. The operator $rec\,p.\,P$ introduces a recursion *label* or *process name* p that can be used in its body P to reiterate the process using p as a jump label. The semantics is so that $rec\,p.\,P$ is equivalent to its unfolding $P\{rec\,p.\,P/p\}$, where $P\{Q/p\}$ denotes syntactic substitution. As done in process algebras we can use rec to fold up recursive equation systems modelling arbitrary forward and backward jumps in control-flow graphs.

By default, a conditional binds tighter than sequential composition, which in turn binds tighter than parallel composition; the loop prefix $rec\,p$ has weakest binding power. As usual, brackets can be used for grouping statements to override the default associations. For instance, in the expression $rec\,p.\,x \;?\; \varepsilon : p;\,!y$ the scope of the loop extends to the end of the expression as in $rec\,p.\,((x \;?\; \varepsilon : p);\,!y)$ whereas $(rec\,p.\,x \;?\; \varepsilon : p);\,!y$ limits the scope and leave $!y$ outside the loop. Similarly, brackets are needed, as in $rec\,p.\,x \;?\; \varepsilon : (p;\,!y)$, to include $!y$ into the else branch of the conditional.

Recursion without restrictions is too powerful for our purposes. We impose the following three *well-formedness* conditions on pSCL expressions which suffices to model the static structure of many standard synchronous programming languages:

– No jumps out of an enclosing parallel composition. This does not limit the power of the language, as for example aborts, traps and general gotos as proposed for/provided by Esterel or SHIM [81, 82] can still be implemented by "chaining" jumps up the thread hierarchy, but has the advantage of a simple parallel/sequential control flow structure. Formally, in every loop $rec\,p.\,P$ the label p must not lie within the scope of a parallel operator $\parallel$. For instance, $rec\,q.\,P \parallel (!x;\,q)$ is not permitted while $P \parallel (rec\,q.\,!x;\,q)$ is accepted. This makes sure that the static control structure of a program is a *series-parallel graph* (see [24]) and the number of concurrently running threads is statically bounded by this graph. In particular any given static thread cannot be concurrently instantiated more than once; A fresh thread instance only runs sequentially after all previous instances of the same static thread.
– Every loop $rec\,p.\,P$ is *clock guarded*, *i.e.*, every free occurrence of label p in P lies within the sequential scope of a pause π. For instance, $rec\,p.\,\pi \,;\, p \,;\, {\textrm{¡}}s$ is clock guarded whereas $rec\,p.\,!s \,;\, p \,;\, {\textrm{¡}}s$ is not. Clock guarded processes are guaranteed to generate finite, terminating macro-steps. This corresponds to the standard requirement in Esterel to not have instantaneous loops.
– No loop label occurs both *free* and *bound* in an expression, where the notion of a free and bound label is as usual. This a standard restriction in process calculi, see e.g., [10]. For instance, $rec\,p.\,{\textrm{¡}}s \,;\, (rec\,q.\,p \,;\, q) \,;\, q)$ is not allowed, whereas $rec\,p.\,{\textrm{¡}}s \,;\, (rec\,r.\,p \,;\, r) \,;\, q)$ is accepted. This restriction avoids capturing of any free variable of $rec\,p.\,P$ by a loop recursion in P in the syntactic unfolding $P\{rec\,p.\,P/p\}$.

Henceforth, *programs* are assumed to be expressions satisfying these conditions. Programs without the rec construct will be called *finite* programs, or *fprogs* for short.

DRAFT 20.4.2015

The imperative statements of a pSCL program describe discrete changes of state at the level of micro-steps. An execution instance of a micro-step is called an *action*. The computation of a concurrent program gets described by a collection of *threads* (concurrent program fragments), each one performing actions independently and interacting with each other through shared memory. Generally, a computation depends on a distinction of actions happening sequentially after each other or concurrently. The sequential order is instantiated from sequential composition $P; Q$. Parallel composition $P \parallel Q$ is the construct that provides the thread topology for achieving concurrency. The resulting tree-like structure of the parallel construct determines statically which actions belong to which individual thread. At run-time, these static threads get instantiated and executed. Every one of such instantiations must have its own local *control-state* and, therefore, is considered a *process*. From this perspective, the *configuration* capturing the global state of a concurrent program at any given moment is determined by the local state of all its processes together with a shared *global memory*.

As in synchronous programming, a *micro-step* can take place when at least one process is *active*, *i.e.*, when it is able to execute an action realising a statement other than π. In this manner, a micro-step produces a change in the configuration resulting from a process executing an action that modifies its own local state and possibly the global memory. Thus, in the course of an instant, active processes induce micro-steps until every process either terminates or reaches a pause, thereby completing a *macro-step*. Then, from the resulting configuration, the environment can provide a fresh stimulus for continuing the computation with a new macro-step occurring in the next instant.

The interaction between processes at the micro-step level must be controlled according to some pre-established rules of *admissible* scheduling in order to enforce the Synchrony Hypothesis abstraction. For instance, suppose in $P \parallel Q$, program P performs a write to a variable x and Q concurrently reads x. Then, under the Synchrony Hypothesis the producer P (system) is faster than the consumer (environment)Q, or, equivalently, Q waits for P. A canonical notion of admissibility that enforces such causalities is the *"init;update;read"* protocol [37], which is referred to as the *"iur"* protocol in the following. It decrees that all initialisations ¡s must take place before any update !s which in turn must both be scheduled before any read, *i.e.*, any conditional test $s \; ? \; P : Q$ on s.

In the next Sec. 2.2 we define the notion of a free unconstrained execution for pSCL programs and then in Sec. 2.3 introduce the restriction imposed by the iur protocol. This defines the operational semantics of the class of *causal* pSCL programs for which we shall later, in Secs. 3 and 4, provide a suitable notion of *constructive* macro-step responses in terms of a denotational fixed-point analysis.

2.2 Micro-step Free Scheduling.

In our operational model, a *process T* is defined by its own current *control-state*, or *state* for short, which contains: (i) information about the precise position of T in the tree structure of forked processes and (ii) control-flow references to specific parts of the code. Formally, T is given by a triplet $\langle id, prog, next \rangle$ where we write $T.id$, $T.prog$ or $T.next$ for referring to the individual elements of T which are called, respectively, *identifier*, *current-program* and *next-control*. Concretely,

– $T.id$ is a non-empty sequence containing an alternation of natural numbers and the symbols l, r that always starts and ends with a number. For instance, $0.l.5$ and $1.r.3.l.7$

DRAFT 20.4.2015
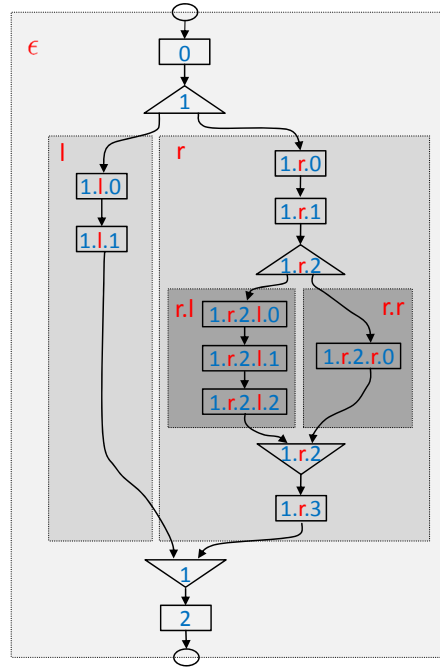
Fig. 1 A sequential-parallel program structure of thread identifiers.

are identifiers but $0.r$ and $r.1.r.2$ are not. Meta-variables to range over identifiers are ι, κ, possibly with indices.

– $T.prog$ is the pSCL expression that is currently scheduled to generate T's actions. Since these are pSCL expressions we use the meta-variables P, Q, *etc.*, to range over these.

– $T.next$ is a list of future program fragments that can be converted into actions sequentially after $T.prog$ has terminated. This list is extended when a sequential composition is executed in $T.prog$. We use the meta-variable Ks to range over next-controls.

The identifier $T.id$ localises process T uniquely in the sequential-parallel control flow of the program context which has generated T. The intuition is that the numbers in the identifier are counting the sequential steps taken by the program context. The symbols (l for *left* and r for *right*) recall the path of previous parallel forks from which the process has emerged. Where we are only interested in the depth of a process in the thread hierarchy, we use a *thread projection* function $th(\iota) \in \{l, r\}^*$ which drops from ι all sequencing numbers. The sequence $th(T.id)$ can be interpreted as the static thread identifier of process T.

Example 1 The serial-parallel graph in Fig. 1 gives an example of the thread identifiers generated by the fprog $P_\varepsilon = a_0 \,;\, (P_l \parallel P_r) \,;\, a_2$ with

$$P_l = a_{1.l.0} \,;\, a_{1.l.1} \qquad\qquad P_r = a_{1.r.0} \,;\, a_{1.r.1} \,;\, (P_{r.l} \parallel P_{r.r}) \,;\, a_{1.r.3}$$
$$P_{r.l} = a_{1.r.2.l.0} \,;\, a_{1.r.2.l.1} \,;\, a_{1.r.2.l.2} \qquad P_{r.r} = a_{1.r.2.r.0},$$

where all a_ι are primitive statements $\{\varepsilon, !s, \text{¡}s\}$. The subscripts ι indicate the thread identifier associated with the action a_ι when it is executed. In Fig. 1 these primitive statements are

shown as rectangular boxes with their identifier written inside it. Notice how the letters l and r (displayed in red colour) identify the static thread in which the action is executed. For instance the action $a_{0.r.2.l.1}$ is executed in the static thread $r.l$, which is the left child of the right child of the main thread. This is the projection $th(0.r.2.l.1) = r.l$. The first top level action a_0 is in the root thread ε, $i.e.$, $th(0) = \varepsilon$, where ε denotes the empty sequence. The hierarchical thread structure is visualised by the dotted gray background boxes.   □

Definition 1 To compare the sequential depth of processes, we use the *(partial) lexicographic order* $\prec$ on path identifiers. The natural numbers are ordered in the usual way, $i.e.$, $0 < 1 < 2\ldots$ while the symbols l, r are considered incomparable. Thus, for identifiers $\iota = d_1\ldots d_n$ and $\iota' = d_1'\ldots d_m'$ we have that $\iota \prec \iota'$ iff ι is a *proper prefix* of ι' or ι is *lexically below* ι', formally:

  – $n < m$ and $\forall 1 \leq j \leq n.\ d_j = d_j'$, **or**
  – there is $0 \leq i < n$ such that $d_{i+1} < d_{i+1}'$ and $\forall 1 \leq j \leq i$ we have $d_j = d_j'$.

We write $\preceq$ for the reflexive closure of $\prec$, $i.e.$, $\iota \preceq \iota'$ iff $\iota \prec \iota'$ or $\iota = \iota'$.   □

The order $\preceq$ contains both the thread hierarchy and sequencing. If $\iota \preceq \iota'$ then ι' is a sequential successor of ι in program order. If $\iota \npreceq \iota'$ and also $\iota' \npreceq \iota$ then both ι and ι' are concurrent. Note that there is no relationship between $\iota \prec \iota'$ and the prefix order on $th(\iota)$ and $th(\iota')$. The sequential successor ι', in general, can be both a descendant or an ancestor of ι in the thread hierarchy.

Example 2 For instance, in our example of Fig. 1, we have $1.r.2 \prec 1.r.2.l.1 \prec 1.r.3$ following the sequential program order but $1.l.0 \nprec 1.r.2.l.1$ and $1.r.2.l.1 \nprec 1.l.0$, because the labels l and r are incomparable. The actions with thread identifiers $1.l.0$ and $1.r.2.l.1$ are not sequentially ordered. They are executed in the concurrent threads $th(1.r.2.l.1) = r.l$ and $th(1.l.0) = l$. Observe that $1.r.2.l.1 \preceq 1.r.3$ but $th(1.r.2.l.1) = r.l$ is not a prefix of $r = th(1.r.3)$. In the other direction, the fork node $1.r.2$ is a sequential predecessor of $1.r.2.l.1$ and $r = th(1.r.2)$ is an ancestor of $r.l = th(1.r.2.l.1)$.   □

Formally, the *global memory* is a Boolean valuation function $\rho : V \to \mathbb{B}$ which indicates the current value for each variable $s \in V$. Any action of a process T (relative to a given memory ρ) produces a new memory ρ' and a set of *successor processes* T'. Thus, any action is completely specified by the *memory* function $\rho' := mem(T, \rho)$ and the *succession* function $T' := nxt(T, \rho)$. For any $x \in V$, the memory function is defined by:

$$mem(T, \rho)(x) := \begin{cases} 0 & \text{if } T.prog = \text{¡}s \text{ and } x = s \\ 1 & \text{if } T.prog = !s \text{ and } x = s \\ \rho(x) & \text{otherwise.} \end{cases}$$

This says that for a given variable $s \in V$, if T performs a reset $\text{¡}s$ then s is changed to 0, if T performs a set $!s$ then s is changed to 1, otherwise, s keeps its value from the previous memory. We define the succession $nxt(T, \rho)$ by case analysis on $T.prog$, where the sequential enumeration for identifier ι is computed by an *increment* function $inc(\iota)$ which increases by

DRAFT 20.4.2015

1 the last number of the identifier ι, *e.g.*, $inc(1.r.6) = 1.r.7$:

$$nxt(\langle \iota, P, [\,] \rangle, \rho) := \emptyset \quad \text{if } P \equiv \varepsilon, P \equiv \textit{¡s} \text{ or } P \equiv \textit{!s} \tag{1}$$

$$nxt(\langle \iota, P, Q{::}Ks \rangle, \rho) := \{ \langle inc(\iota), Q, Ks \rangle \} \quad \text{if } P \equiv \varepsilon, P \equiv \textit{¡s} \text{ or } P \equiv \textit{!s} \tag{2}$$

$$nxt(\langle \iota, P \,;\, Q, Ks \rangle, \rho) := \{ \langle \iota, P, Q{::}Ks \rangle \} \tag{3}$$

$$nxt(\langle \iota, rec\, p. P, Ks \rangle, \rho) := \{ \langle \iota, P\{rec\, p. P / p\}, Ks \rangle \} \tag{4}$$

$$nxt(\langle \iota, s \,?\, P : Q, Ks \rangle, \rho) := \begin{cases} \{ \langle inc(\iota), P, Ks \rangle \} & \text{if } \rho(s) = 1 \\ \{ \langle inc(\iota), Q, Ks \rangle \} & \text{otherwise} \end{cases} \tag{5}$$

$$nxt(\langle \iota, P \,\|\, Q, Ks \rangle, \rho) := \{ \langle \iota, \varepsilon, Ks \rangle, \langle \iota.l.0, P, [\,] \rangle, \langle \iota.r.0, Q, [\,] \rangle \}. \tag{6}$$

Let us explain the different cases one by one:

(1) If the program $T.prog$ is one of the basic statements (*i.e.*, empty ε, set $!s$ or reset $\textit{¡s}$) and the list of continuation processes in the next-control $T.next$ is empty $[\,]$, then the process (after execution) is terminated and disappears from the configuration. This is achieved by setting the succession to be the empty set.

(2) If $T.prog$ is one of the basic statements and the list of continuation processes in $T.next$ is a non-empty list $Q{::}Ks$, then we start Q in a new process with next-control Ks and a sequentially incremented index $inc(\iota)$.

(3) If $T.prog$ is a sequential composition $P \,;\, Q$ then we start P in a new process with the same identifier and add Q to the front of the next-control list. The identifier does not increment since we do not consider the new process $\langle \iota, P, Q{::}Ks \rangle$ a sequential successor but only a structural replacement.

(4) A loop $T.prog = rec\, p. P$ behaves like its unfolding $P\{rec\, p. P / p\}$, without modification to the identifier and next-controls.

(5) Next consider a process with conditional program $T.prog = s \,?\, P : Q$ in memory ρ. Depending on whether the memory value for the variable s is 1 or 0 we install the P or the Q branch, respectively, with an incremented identifier and the same next-control. The identifier is incremented because the branches are considered as being executed strictly after the conditional test, in sequential program order.

(6) Finally, executing a parallel program $T.prog = P \,\|\, Q$ instantiates the two sub-threads P and Q in their own process $\langle \iota.l.0, P, [\,] \rangle$ and $\langle \iota.r.0, Q, [\,] \rangle$, respectively, with a fresh and empty next-control but extended identifiers. The process P is the *initial* sequential statement of the *left* child of the parent process $\langle \iota, P \,\|\, Q, Ks \rangle$. Therefore, we add the suffix $l.0$ to the parent's identifier, and analogously $r.0$ for the right child Q. At the same time that the parent process forks its two children it transforms itself into a join process $\langle \iota, \varepsilon, Ks \rangle$. Since $\iota \prec \iota.l.0$ and $\iota \prec \iota.r.0$ both children have strictly larger identifiers. Since only processes with maximal identifiers are executable (details below), the join process must wait for the children to terminate before it can release the next-controls Ks, or terminate itself in case $Ks = [\,]$.

Note that there is no clause for the succession of a pausing process or a process label, *i.e.*, $nxt(\langle \iota, \pi, Ks \rangle, \rho)$ and $nxt(\langle \iota, p, Ks \rangle, \rho)$ are undefined. This is no problem since (i) program π is never executed in a micro-step action but only by the next global clock tick (see below), and (ii) we are only interested in the behaviour of *closed* pSCL expressions which do not have any free process labels.

Example 3 Consider the process $T_0 = \langle 0, \textit{¡x}\,;\, y \,?\, \pi : \textit{!x}, [\,] \rangle$. This process resets variable x and then either pauses or sets variable x depending on the value of variable y. Let us derive its behaviour in the formal semantics.

Starting from some initial memory ρ_0, executing T_0 yields a new memory $\rho_1 = mem(T_0, \rho_0)$ and a set of successors $S_1 = nxt(T_0, \rho_0)$. This first action breaks up the sequential composition operator ; according to rule (3). This results in $S_1 = \{T_1\}$ where $T_1 = \langle 0, \text{¡}x, [y \ ? \ \pi : !x] \rangle$. The action does not modify the memory, *i.e.*, $\rho_1 = \rho_0$. Proceeding with T_1 from ρ_1, we come to execute the reset ¡x following rule (2), obtaining $\rho_2 = mem(T_1, \rho_1)$ and successors $S_2 = nxt(T_1, \rho_1)$. Memory ρ_2 now assigns 0 to variable x, while y retains its initial value from ρ_0. The succession is $S_2 = \{T_2\}$ with $T_2 = \langle 1, y \ ? \ \pi : !x, [] \rangle$. Notice the increment of the identifier $T_2.id = 1 = inc(0) = inc(T_1.id)$ which reflects the fact that execution has passed a sequential composition operator. The conditional T_2 now reads the value of y in memory ρ_2 and passes control to the 'then' or 'else' branch:

- If $\rho_2(y) = \rho_0(y) = 1$ then the conditional executes the 'then' branch. We get $\rho_3 = mem(T_2, \rho_2) = \rho_2$ and $S_3 = nxt(T_2, \rho_2) = \{T_3\}$ with $T_3 = \langle 2, \pi, [] \rangle$ by rule (5). There are no micro-step rules for π which is forced to pause during the current synchronous instant. T_3 makes progress only at the next global clock tick where it transforms into $T_3' = \langle 0, \varepsilon, [] \rangle$ as described later.
- If $\rho_2(y) = \rho_0(y) = 0$ then $\rho_3 = mem(T_2, \rho_2) = \rho_2$ and $S_3 = nxt(T_2, \rho_2) = \{T_3\}$ with $T_3 = \langle 2, !x, [] \rangle$ by rule (5). From here, the execution of !x sets variable x and yields the new memory $\rho_4 = mem(T_3, \rho_3)$ with $\rho_4(x) = 1$ and $\rho_4(y) = \rho_2(y)$. Since $S_4 = nxt(T_3, \rho_3) = \emptyset$ by rule (1), there are no more processes from which we can continue. The execution of T_0 has terminated instantaneously in the current synchronous instant. $\qquad\square$

Let us combine the memory and succession functions for a single process to define the micro-steps of an arbitrary set of processes running concurrently. This requires the notion of a *configuration*, defined next:

Definition 2 A *configuration* is given by a pair (Σ, ρ), where ρ is the global memory and Σ, called the *process pool*, is a finite set of (closed) processes such that

- all identifiers are distinct, *i.e.*, for all $T_1, T_2 \in \Sigma$, if $T_1.id = T_2.id$ then $T_1 = T_2$;
- the sequential ordering of identifiers coincides with the thread hierarchy, *i.e.*, for all $T_1, T_2 \in \Sigma$, we have $T_1.id \preceq T_2.id$ iff $th(T_1.id)$ is a (not necessarily proper) prefix of $th(T_2.id)$;
- the identifiers form a full thread tree, *i.e.*, for each $T \in \Sigma$ and every proper prefix (ancestor) $t \in \{r, l\}^*$ of $th(T.id)$, there is a process $T' \in \Sigma$ of T with $th(T'.id) = t$.

A configuration (Σ, ρ) is *empty* if $\Sigma = \emptyset$. We call a process $T \in \Sigma$

- *pausing* when $T.prog = \pi$;
- *active* if it is not pausing and $T.id$ is $\preceq$-maximal (identifier order) in Σ;
- *waiting* if it is neither pausing nor active.

A configuration with memory ρ in which all the processes in Σ are waiting or pausing, is called *quiescent*. $\qquad\square$

Micro-sequences. From a given configuration (Σ, ρ) and a selection $T \in \Sigma$ of an active process, we can let T execute an action to produce a *micro-step*

$$(\Sigma, \rho) \xrightarrow{T} (\Sigma', \rho'), \tag{7}$$

where in the free scheduling there is no constraint on the selection of T other than it being active. The resulting memory $\rho' = mem(T, \rho)$ is computed directly from the *mem* function.

DRAFT 20.4.2015

The new process pool Σ' is obtained by removing T from Σ and replacing it by the set of successors generated by nxt, i.e., $\Sigma' = \Sigma \setminus \{T\} \cup nxt(T, \rho)$. Note that in the free schedule both the next process pool Σ' and the new memory ρ' only depend on the active process T that is executed and the current memory ρ. They do not depend on the other processes in Σ. Since the successor configuration is uniquely determined by (Σ, ρ) and T, we may write $(\Sigma', \rho') = T(\Sigma, \rho)$. In a *micro-sequence* the scheduler runs through a succession

$$(\Sigma_0, \rho_0) \xrightarrow{T_1} (\Sigma_1, \rho_1) \xrightarrow{T_2} \cdots \xrightarrow{T_k} (\Sigma_k, \rho_k) \tag{8}$$

of micro-steps obtained from the interleaving of process executions. We let $\twoheadrightarrow$ be the reflexive and transitive closure of $\rightarrow$. That is, we write

$$R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_k, \rho_k)$$

to express that there exists a micro-sequence R, not necessarily maximal, from configuration (Σ_0, ρ_0) to (Σ_k, ρ_k). The sequence R is a function mapping each index $1 \leq j \leq k$ to the process $R(j) = T_j$ executed at micro-step j and $len(R) = k$ is the length of the micro-sequence, i.e., the number of actions executed so far. We call any pair $(i, R(i))$ consisting of a micro-step index $1 \leq i \leq len(R)$ together with the process $R(i)$ executed at position i, a *process instance* of R. Further, it will be necessary to restrict a micro-sequence $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ to its *prefixes* $R@i : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_i, \rho_i)$ for $i \leq n = len(R)$.

Macro-steps. A *macro-step*, also called a *synchronous instant*, or *instant* for short, abbreviated

$$R : (\Sigma_0, \rho_0) \Longrightarrow (\Sigma_k, \rho_k) \tag{9}$$

is a maximal micro-sequence R that reaches a final quiescent configuration. Note that for any memory ρ, a configuration $(\emptyset, \rho)$ is trivially quiescent. For the sake of simplicity, sometimes we drop the mapping M from our relations $\twoheadrightarrow$ and $\Longrightarrow$. When (Σ_k, ρ_k) is quiescent but non-empty then no further micro-step is possible (which explains the term 'quiescent') since all processes are waiting for the clock to tick. Such a clock tick

$$(\Sigma_k, \rho_k) \Longrightarrow^{tick} (\Sigma', \rho') \tag{10}$$

consists of eliminating every pausing process with empty continuation $\langle \iota\, d, \pi, [\,] \rangle \in \Sigma_k$ and replacing every pausing process $\langle \iota\, d, \pi, Q::Ks \rangle \in \Sigma_k$ with a non-empty continuation by a new process $\langle \iota\, 0, Q, Ks \rangle \in \Sigma'$ preserving the sequential identifier of all ancestors but restarting the current thread at sequence number 0. The new memory ρ' preserves all internal and output variables but permits the environment to change any input variables for the next macro-step. For the investigations in this article, however, we are only interested in single macro-steps generated by the behaviour of pSCL expressions. Therefore, we will not be concerned with the modelling of successions of clock ticks.

Example 4 Let (Σ_1, ρ_0) be a configuration where ρ_0 gives value 0 to every variable and the process pool $\Sigma_1 = \{T_1\}$ consists of the root process: $T_1 = \langle 0, (!s\,;\,!t \parallel \mathbin{\textrm{¡}}s\,;\,t\,?\,\varepsilon : \pi), [\,] \rangle$. The complete computation graph for the free scheduling from (Σ_1, ρ_0) is depicted in Fig. 2. The processes are abbreviated as follows:

$$
\begin{aligned}
T_1 &= \langle 0, !s\,;\,!t \parallel \mathbin{\textrm{¡}}s\,;\,t\,?\,\varepsilon : \pi, [\,] \rangle & T_{31} &= \langle 0.l.0, !s, [!t] \rangle \\
T_{20} &= \langle 0, \varepsilon, [\,] \rangle & T_{32} &= \langle 0.r.0, \mathbin{\textrm{¡}}s, [t\,?\,\varepsilon : \pi] \rangle \\
T_{21} &= \langle 0.l.0, !s\,;\,!t, [\,] \rangle & T_{41} &= \langle 0.l.1, !t, [\,] \rangle \\
T_{22} &= \langle 0.r.0, \mathbin{\textrm{¡}}s\,;\,t\,?\,\varepsilon : \pi, [\,] \rangle & T_{42} &= \langle 0.r.1, t\,?\,\varepsilon : \pi, [\,] \rangle \\
T_{521} &= \langle 0.r.2, \varepsilon, [\,] \rangle & T_{522} &= \langle 0.r.2, \pi, [\,] \rangle
\end{aligned}
$$

DRAFT 20.4.2015
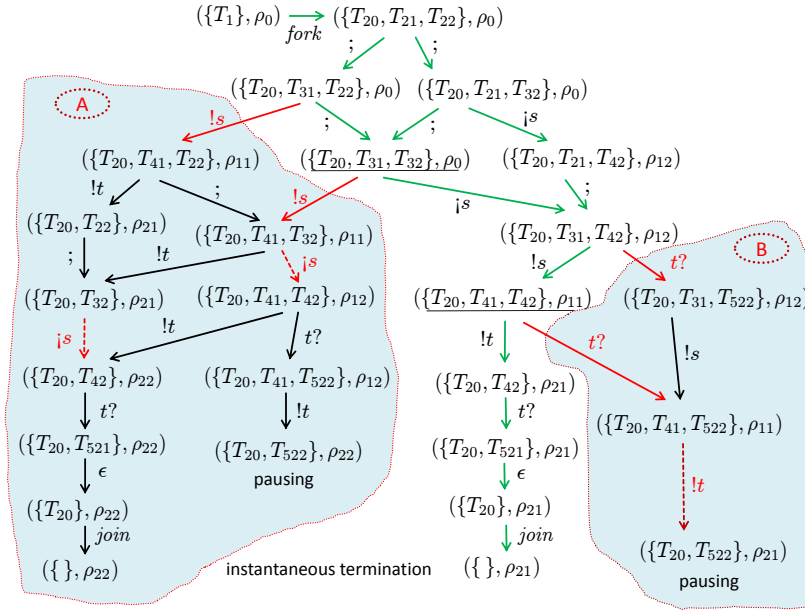
Fig. 2 The free scheduling graph of process T_1 of Ex. 4

Each edge in Fig. 2 is a single micro-step. For ease of explanation we do not use the selected process T_i as the label like in (7) but instead the primitive operator executed in the action, *i.e.*, a sequential composition ; (rule (3)), atomic set, reset or the empty statements $!s$, $!t$, $¡s$, ε (rules (1) and (2)) or a parallel composition $\parallel$ (rule (6)). The shaded regions named A and B will be explained later.

Since T_1 is active it can induce the micro-step $(\Sigma_1, \rho_0) \rightarrow (\Sigma_2, \rho_0)$ with the a succession $\Sigma_2 = \{T_{20}, T_{21}, T_{22}\}$ of three processes as a result of executing the parallel fork, the parent T_{20} and its two children T_{21} and T_{22}. Observe that in Σ_2 the two children are active but the parent with identifier 0 is waiting, because $0 \prec 0.l.0$ and $0 \prec 0.r.0$. The parent T_{20} now plays the role of a 'join' in the sense that it cannot execute any action until the two children terminate and its own identifier becomes maximal again. Let us suppose that T_{21} and T_{22} are scheduled in that order to get $(\Sigma_2, \rho_0) \twoheadrightarrow (\Sigma_4, \rho_0)$ with $\Sigma_4 = \{T_{20}, T_{31}, T_{32}\}$, where T_{31} and T_{32} are both active. The configuration (Σ_4, ρ_0) is underlined in Fig. 2. Notice that we reach exactly the same configuration if we first schedule T_{22} and then T_{21}. The concurrent execution of the sequential compositions in T_{21} and T_{22} is confluent, because there are no read or write accesses to variables. However, in (Σ_4, ρ_0) things become interesting since the chosen scheduling order will result in different configurations. For if $(\Sigma_4, \rho_0) \twoheadrightarrow (\Sigma_6, \rho_{11})$, with $\Sigma_6 = \{T_{20}, T_{41}, T_{42}\}$, results from scheduling T_{32} followed by T_{31}, then first the reset $¡s$ is performed and thereafter the set $!s$, so that $\rho_{11}(s) = 1$. On the other hand, if first T_{31} is picked and then T_{32} does its initial action, then $(\Sigma_4, \rho_0) \twoheadrightarrow (\Sigma_6, \rho_{12})$ with $\rho_{12}(s) = 0$. Although the resulting process pool Σ_6 is the same in both configurations, the global memory is not.

Continuing the schedule from configuration (Σ_6, ρ_{11}), also underlined in Fig. 2, we see that there is a race between the reading of variable t by T_{42} the writing to t by T_{41}. If we first execute T_{41}, then the conditional T_{42} will activate its 'then'-branch ε. Therefore, we eventually reach the configuration (Σ_9, ρ_{21}) with $\Sigma_9 = \{T_{20}\}$ where the memory satisfies $\rho_{21}(s) =$

DRAFT 20.4.2015

$\rho_{21}(t) = 1$. Now 'join' process T_{20} becomes active which instantaneously terminates reaching the quiescent configuration $(\{\}, \rho_{21})$. On the other hand, if in (Σ_6, ρ_{11}) the process T_{42} first gets to test the value of t, which is 0, before T_{41} sets it to 1, then the 'else'-branch is selected and we end up in the configuration (Σ_8, ρ_{21}) where $\Sigma_8 = \{T_{20}, T_{522}\}$. This configuration is also quiescent as it contains no active processes. Here, the 'join' process T_{20} is still waiting since it has a strictly smaller sequence number than process T_{522} which is pausing. No progress can be made until the next clock tick makes T_{522} disappear from the configuration, thereby activating T_{20} which then terminates instantaneously. Note that the conflict between T_{41} and T_{42} in (Σ_6, ρ_{11}) results in a non-determinism of control, *viz.* between terminating in the same instant or the next. □

Clearly, as demonstrated in Ex. 4 the selection strategy applied in the free scheduling of a program determines the final memory content and termination behaviour of a program in a macro-step. If we would consider pSCL as a just another clocked process algebra such as [42, 21, 56] or a model of general statecharts, *e.g.*, [70, 84, 40] (or in fact Java threads, for that matter) the non-determinism would not worry us. It is a natural consequence of the asynchrony of parallel execution. We can leave it in the responsibility of the versed programmer to harness her or his programs by explicit synchronisation through shared memory mutual exclusion algorithms (see [58]) in order to get rid of non-determinacy. Yet, this is not the right approach for synchronous programming where every program, by compilation, is required to code a deterministic Mealy machine. In synchronous programming it is the compiler which has to achieve determinate tick responses under pessimistic assumption on the varying degree of perturbations arising from the non-determinism of target run-time system.

In synchronous programming the programmer is supported by static schedulability and causality analyses. Often non-determinism can be eliminated by restricting the free scheduling to so-called *admissible* schedules that are natural for, or intended by, the programmer and at the same time reliably implemented on the chosen run-time platform by a trusted compiler.

Example 5 Consider Example 4 in which the non-determinacy of the tick response is due to races between the setting and resetting of variable s and the reading and writing of variable t. Suppose we compile the root process T_0 as a data-flow network in which the non-determinism maps to the concurrent execution of function blocks. Then it is easy to ensure that the data-flow always executes reset of a variable (value initialisation) before any set (value update) and all write accesses before the reads. This natural ordering prohibits the execution of the transitions shown in Fig. 2 as dashed arrows. It eliminates the path in region B in which the set action $!t$ happens after the reads $t?$ and the paths in region A with the resets $¡s$ occurring after the sets $!s$. The remaining admissible scheduling paths then all lead, deterministically, to instantaneous termination in configuration $(\{\}, \rho_{21})$. □

A canonical notion of admissibility to avoid causality locks is the "init;update;read" (iur) protocol, which forces the accesses of every variable to undergo strict cycles of first initialisations $(¡s)$, then updates $(!s)$ and finally reads $(s?)$. Moreover, the iur protocol can be refined by limiting the number of initialisations that are permitted during a single macro-step on any variable. Liberal notions of sequential constructiveness permitting more than one init;update;read cycle have recently been proposed [37, 38]. In the traditional model of synchronous programming—paradigmatically represented by Esterel—only one iur cycle is permitted. This leads to a more conservative notion of constructiveness which is the subject of this article and formalised in the next section.

Since well-formed pSCL programs are clock-guarded, we can unfold all loops and extract finite *rec*-free expressions that fully describe the program's macro step reactions. Therefore,

DRAFT 20.4.2015

as the main results in this article concern the scheduling of actions inside a single finite macro-step, it suffices to consider only finite, recursion-free pSCL programs, *i.e.*, fprogs.

2.3 Reactiveness and Determinacy.

All non-determinism of concurrent execution arises from two types of data races: write-write conflicts and write-read conflicts. To remove these races the iur scheduling protocol enforces precedence of resets over sets and of writes over reads. The strict ordering can be broken only if the variable accesses are confluent. A suitable notion of confluence has been introduced in [37,38].

Definition 3 (Confluence of Processes) Let T_1 and T_2 be two arbitrary processes and (Σ, ρ) a configuration. Then,

1. T_1, T_2 are called *conflicting* in (Σ, ρ) if both T_1 and T_2 are active in Σ and $T_1(T_2(\Sigma, \rho)) \neq T_2(T_1(\Sigma, \rho))$;
2. T_1, T_2 are *confluent* in (Σ, ρ), written $T_1 \sim_{(\Sigma, \rho)} T_2$, if there is no micro-sequence $(\Sigma, \rho) \twoheadrightarrow (\Sigma', \rho')$ such that T_1 and T_2 are conflicting in (Σ', ρ'). $\qquad\square$

Example 6 As an illustration consider once more Example 4. Processes T_{31} and T_{32} are conflicting in configuration $(\Sigma_4, \rho_0) = (\{T_{20}, T_{31}, T_{32}\}, \rho_0)$ because, as we have seen, both are active in this configuration and, moreover, different execution orders lead to different results. Since the action of T_{31} is !s (update) and the action of T_{32} is the reset ¡s (init), the scheduling protocol gives precedence to T_{32}. Similarly, T_{41} and T_{42} are in conflict in configuration (Σ_6, ρ_{12}) with $\Sigma_6 = \{T_{20}, T_{41}, T_{42}\}$ as can be seen from Fig. 2. For their part, processes T_{21} and T_{22} are independent or confluent in (Σ_2, ρ_0) with $\Sigma_2 = \{T_{20}, T_{21}, T_{22}\}$. This is so because in every micro-sequence $(\Sigma_2, \rho_0) \twoheadrightarrow (\Sigma', \rho')$ the only configuration in which both T_{21} and T_{22} are active is precisely (Σ_2, ρ_0). Furthermore, as can be seen from Fig. 2, the order of execution is unimportant in this case, namely $T_{21}(T_{22}(\Sigma_2, \rho_0)) = T_{21}(T_{22}(\Sigma_2, \rho_0)) = (\Sigma_4, \rho_0)$, where $\Sigma_4 = \{T_{20}, T_{31}, T_{32}\}$. Note that since the initial action of both T_{21} and T_{22} is the breaking up of the sequential composition, and thus not variable accesses, their ordering is unconstrained by the "init;update;read" scheduling protocol. $\qquad\square$

In this article we introduce a fairly stringent interpretation of the iur protocol derived from conservative SMoCCs such as Esterel or Quartz, which we term *Berry admissibility* (Def. 4 below). It uses confluence to permit "ineffective" sets after reads but is stronger than SC-admissibility [37], as it enforces the iur protocol on *all* accesses not just concurrent ones as in [37]. Whatever synchronisation protocol X we use —there may be many other interesting ones still to be discovered— the restriction to X-admissible executions not only reduces non-determinacy. Such synchronisation constraints may lead to deadlock, i.e., configurations in which no micro-step is possible without violating X-admissiblity. Thus we must care about X-reactiveness, *i.e.*, the property that a program does not get stuck when executed in an X-admissible fashion.

2.3.1 B-Admissibility and B-Reactiveness

The tighter the underlying notion of X-admissibility the more information we have from knowing that a program is X-reactive. If all X-admissible schedules are also Y-admissible then a program without deadlocks under X is also deadlock-free under Y. Here we introduce

DRAFT 20.4.2015

a suitable notion of admissibility that captures the essence of Esterel which is tighter than *SC-admissibility* introduced in [37, 38].

Definition 4 (Berry Admissibility and Reactiveness) A micro-sequence $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ is *Berry admissible* (*B-admissible*) iff

(1) *R* does not reset any variable that has been set before.

Formally, if $R(i)$ for $0 < i \leq n$ executes a set $!s$ for some $s \in V$, then no $R(j)$ for $i < j \leq n$ executes a reset $¡s$.

(2) *R* does not write any variable which has been read before, unless this late write is ineffective in the sense that the write is confluent with the read and the very same value has been written already before the read.

Formally, if $R(j)$ for $0 < j \leq n$ executes a conditional test $s ? P : Q$ for some P, Q, and $R(k)$ for $j < k \leq n$ performs a set $!s$ (reset $¡s$), then there exists an index $i < j$ before the read where $R(i)$ already executed a set $!s$ (reset $¡s$) and $R(j) \sim_{(\Sigma_j, \rho_j)} R(k)$.

An fprog *P* is called *Berry reactive* (*B-reactive*) if from every initial configuration $(\{\langle 0, P, [] \rangle\}, \rho_0)$ there is at least one B-admissible instant. □

Example 7 If a reset happens sequentially after a set (violating Def. 4(1)), as in $P_1 := !s ; ¡s$, then this violates the monotonicity of signal stabilisation. In the conservative delay-insensitive model of Esterel this is a hazard, since a concurrent environment could read either the first output value $s = 1$ (which is interpreted as an emit) or the second $s = 0$ (which is an initialization). This creates a write-write race, thus jeopardising determinism. P_1 does not have a B-admissible execution. The opposite ordering $P_2 := ¡s ; !s$ of a reset followed by a set is B-admissible, since it adheres to the monotonic stabilisation protocol.

A read-write race (violating Def. 4(2)) occurs in the sequential programs $P_3 := s ? !s : \varepsilon$ and $P_4 := s ? ¡s : \varepsilon$. The write accesses $!s$ and $¡s$, respectively, may effectively overwrite the externally controlled value of *s* which is tested in the conditionals. If we consider P_3 and P_4 to be environments of themselves then we run into a causality loop: the test $s?$ must wait until the program has set or reset its value, which however can only happen after the test has been executed. If *R* is the micro-sequence generated from P_3 with initial memory $\rho_0(s) = 1$ then it executes the set $!s$ after the read $s?$ without any set having happened before the read. Similarly, we get a reset $¡s$ after the read $s?$ in P_4 but this reset value has not been established before the read. Therefore, neither P_3 nor P_4 are B-reactive. In the hardware translation of Esterel, P_3 would be a delay loop $s = s$ which has two stable solutions $s = 0$ and $s = 1$, while P_4 generates essentially the feed-back system $s = \bar{s} \cdot s$ which may produces glitches before it settles at $s = 0$, if it stabilises at all.

P_3 and P_4 were acceptable if the status of *s* was already decided before the test $s?$. For instance, in $!s ; P_3$ the second $!s$ in P_3 is ineffective from the point of view of the read access because the status 1 on *s* is determined by the first $!s$ which occurs sequentially before the read. Thus, executing $!s ; P_3$ is B-admissible. P_4 can be executed admissibly in the form $¡s ; P_4$ which then bypasses the reset $¡s$ in P_4. On the other hand, $!s ; P_4$ would not be B-reactive because it generates a reset $¡s$ after a set $!s$. We note that all programs $P_1 – P_4$ are sequentially admissible [37] (called Δ_*-admissibility in [4]) because under sequential admissibility glitches can only be generated from concurrent accesses, not sequential ones as in P_1 and P_2. □

Example 8 Although each fprog $P := x ? !y : \varepsilon$ and $Q := y ? !x : \varepsilon$ is B-reactive, their concurrent composition fprog $P \| Q$ is not. There is only one initial memory ρ_0 from which

this has any B-admissible instants, viz. $\rho_0(x) = \rho_0(y) = 0$. Whenever initially $\rho_0(x) = 1$ (or $\rho_0(y) = 1$), then the write statement $!y$ in P (or $!x$ in Q) is executed *after* x (or y) has been read by P. This violates Def. 4(2) because there are no other writes before the read which would make the "late" write ineffective.   □

Example 9 All the scheduling sequences $R : (\{T_1\}, \rho_0) \twoheadrightarrow (\{\,\}, \rho_{21})$ of Ex. 4, following the transitions colored green in Fig. 2 are B-admissible. None of the scheduling sequences going through a red transition, entering region A or B, is B-admissible. The sequences entering region A are violating Def. 4(1) by resetting variable s (dashed red arrows labelled $\textrm{¡}s$) after s has been set (solid red arrows labelled $!s$). The sequences entering region B are breaking the constraint Def. 4(2) because variable t is set (dashed red arrow labelled $!t$) after it has been read (solid red arrows labelled $t?$), without any setting of variable t before the read. However, since at least one B-admissible scheduling sequence leads to completion, the program $!s \,;\, !t \,\|\, \textrm{¡}s \,;\, t \,?\, \varepsilon : \pi$ of Ex. 4 is *B-reactive*.   □

2.3.2 SC-Admissibility and SC-Read-Determinacy

When it comes to the question of determinacy then we want the underlying notion of *X*-admissibility to be as weak as possible. If a program analysis detects determinacy under all *X*-admissible executions, then the implied level of robustness depends on how much non-determinism is still permitted by *X*-admissible executions. For instance, if *X*-admissibility limits execution to a single micro-sequence, *e.g.* through a global linear priority ordering on all statements, then determinacy is trivial. On the other hand, knowing that a program is determinate under all free schedules, is a very strong (and rare) property for a program to have. To get more headroom for our main result we use SC-admissibility. In contrast to B-admissibility this admits writes-after-reads and resets-after-sets, if these are sequential successors in program order or confluent. The following definition is rephrased from [37, 38].

Definition 5 (SC-Admissibility and Reactiveness) A micro-sequence $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_k, \rho_k)$ is *SC-admissible* if for every two processes $R(i)$, $R(j)$ such that $0 < i < j \le n$ and either

(i) $R(j)$ performs a reset $\textrm{¡}s$ or set $!s$ on a variable s that was read (tested) by $R(i)$, or
(ii) $R(j)$ performs a reset $\textrm{¡}s$ on a variable s on which $R(i)$ performed a set $!s$,

the first $R(i)$ is sequentially before $R(j)$ in program order or both are confluent, *i.e.*, we have $R(i).id \preceq R(j).id$ or $R(i) \sim_{(\Sigma_i, \rho_i)} R(j)$.

An fprog P is *SC-reactive*, if from every initial configuration $(\{\langle 0, P, [\,] \rangle\}, \rho_0)$ there is at least one SC-admissible instant for P.   □

One can show that B-admissibility is more restrictive than SC-admissibility.

Proposition 1 *Every B-admissible micro-sequence is also SC-admissible.*

Proof Let $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ be a B-admissible micro-sequence, with processes instances $R(i)$ and $R(j)$ such that $0 < i < j \le n$. First note that by condition (1) of B-admissibility Def. 4, the situation (ii) of Def. 5 cannot occur. We only need to care about the situation (i), where $R(i)$ is a read and $R(j)$ a write of the same variable s. But then condition (2) of B-admissibility implies both are confluent, *i.e.*, $R(i) \sim_{(\Sigma_i, \rho_i)} R(j)$. This was to be shown.   □

DRAFT 20.4.2015

Example 10 The Ex. 4 is B-reactive and thus also SC-reactive. However, it does not have any SC-admissible scheduling sequences which are not B-admissible at the same time. None of the scheduling sequences entering regions A or B in Fig. 2 are SC-admissible. Let us look at what happens in region A. For instance, take the scheduling

$$R = T_1, T_{21}, T_{31}, T_{41}, T_{22}, T_{32}, T_{42}, T_{521}, T_{20} : (\{T_1\}, \rho_0) \rightarrow (\{\}, \rho_{22})$$

in which $R(3) = T_{31}$ performs a set $!s$ and later $R(6) = T_{32}$ performs a reset $¡s$. This violation of resets-before-sets is permitted under SC-admissibility only if the actions are sequentially ordered or confluent. The former is not the case, $T_{31}.id = 0.l.0 \not\preceq 0.r.0 = T_{32}.id$, because both processes are from concurrent threads. The latter is not the case either, because $T_{31} \not\sim_{(\{T_{20}, T_{31}, T_{22}\}, \rho_0)} T_{32}$. In fact, there is the (free) schedule $T_{22} : (\{T_{20}, T_{31}, T_{22}\}, \rho_0) \rightarrow (\{T_{20}, T_{31}, T_{32}\}, \rho_0)$ (underlined in Fig. 2) which reaches the configuration $(\{T_{20}, T_{31}, T_{32}\}, \rho_0)$ in which both T_{31} and T_{32} are active and conflicting (see Def. 3). Executing T_{31}, T_{32} from here leads to $(\{T_{20}, T_{41}, T_{42}\}, \rho_{12})$ while the swapped ordering T_{32}, T_{31} ends up in $(\{T_{20}, T_{41}, T_{42}\}, \rho_{11})$ which have different memories.

Similarly, one can show that the two concurrent processes T_{42} and T_{41} which read and set variable t are in conflict on every schedule that runs through region B. The critical configuration for region B is $(\{T_{20}, T_{41}, T_{42}\}, \rho_{11})$ (underlined in Fig. 2) in which processes T_{42} and T_{41} are in conflict with each other. □

Clearly, by Prop. 1, every B-reactive program is also SC-reactive. An X-reactive program is guaranteed not to deadlock under X-admissible execution. However, it may be non-determinate, *i.e.*, generate different final memory states. In defining determinacy precisely we meet another degree of freedom, depending on whether or not we permit the outcome at the end of an instant to be functionally dependent on the memory configuration at the beginning of the instant. For instance, we might distinguish, as done in Esterel V7, between *temporary* and *registered* variables. The value of a temporary variable is ephemeral and must be recomputed by the program at every instant. The value of a registered variable is provided by the environment in memory at the beginning of each instant. Hence, the final response may depend on the initial value of registered variables but not on the initial value of the temporary variables. This gives rise to the following definition, parametric in X-admissibility, where the notations $\twoheadrightarrow_X$ and $\Longrightarrow_X$ are used to indicate that the corresponding micro-sequence complies with a particular notion X of admissibility. E.g. $\twoheadrightarrow_B$ refers to a B-admissible micro-sequence and $\Longrightarrow_{SC}$ indicates a SC-admissible instant.

Definition 6 (X-Determinacy) For a given set of *temporary* variables $W \subseteq V$, an fprog P is *X-determinate for W* (X_W-*determinate*) iff the following two conditions hold:

1. For every fixed initial memory, P computes the same final memory in *all* X-admissible instants.

   Formally, if $(\langle \iota, P \rangle, \rho_0) \Longrightarrow_X (\Sigma_0, \gamma_0)$ and $(\langle \iota, P \rangle, \rho_0) \Longrightarrow_X (\Sigma_1, \gamma_1)$ then $\gamma_0 = \gamma_1$.

2. For every temporary variable in W, P either (*i*) computes the very same final value in *all* X-admissible instants, or (*ii*) it does not modify the initial memory value of this variable in *any* X-admissible instant. In other words, if P changes the value of a variable $x \in W$ in *any* X-admissible instant then this must be the final value for x in *all* X-admissible instants.

   Formally, for all $x \in W$ and γ_0, ρ_0, Σ_0: if $(\langle \iota, P \rangle, \rho_0) \Longrightarrow_X (\Sigma_0, \gamma_0)$ and $\gamma_0(x) \neq \rho_0(x)$, then for all γ_1, ρ_1, Σ_1 such that $(\langle \iota, P \rangle, \rho_1) \Longrightarrow_X (\Sigma_1, \gamma_1)$, we have $\gamma_1(x) = \gamma_0(x)$. □

DRAFT 20.4.2015

In this article we will treat two special cases: When W is the empty set, $W = \emptyset$, then X_W-determinacy is simply called *X-determinacy*. When $W = rd(P)$ is the set of read variables of P, defined by

$$rd(P) := \begin{cases} rd(P_1) \,\cup\, rd(P_2) & \text{if } P = P_1 \,\|\, P_2 \text{ or } P = P_1 \,;\, P_2 \\ \{s\} \,\cup\, rd(P_1) \,\cup\, rd(P_2) & \text{if } P = s \,?\, P_1 : P_2 \\ \emptyset & \text{otherwise} \end{cases}$$

then X_W-determinacy is referred to as *X-read-determinacy*. The following proposition is obvious, with Prop. 1:

Proposition 2 *Every X-read-determinate fprog is also X-determinate and every SC-determinate program is B-determinate.*

Note that purely *sequential* programs, *i.e.*, those without the concurrency operator $\|$, are trivially deterministic and hence SC-read-determinate. Sequential programs are also always SC-reactive. They can however fail B-reactiveness, *i.e.*, if their execution is not B-admissible because it generates a causal hazard in the access to a variable (see Ex. 7). This models the stronger interpretation of reactiveness in the more conservative SMoCCs like Esterel and Quartz which we deal with, here. Also, SC-read-determinacy is trivial for pure input variables which are never written by a program because their final value will always be the same as the initial value. Hence, all programs, including those containing the $\|$ operator, with *disjoint* input and output variables, are SC-read-determinate but possibly not B-reactive.

The following Ex. 11 brings home the problems causality poses for compositionality.

Example 11 All P_1, P_3, P_4 from Ex. 7 are purely sequential programs which are not B-reactive but SC-read-determinate. An fprog which is B-reactive but not SC-determinate is the parallel composition $P \,\|\, Q$, where $P := x \,?\, \varepsilon : !y$ and $Q := y \,?\, \varepsilon : !x$. The left component P sets y to 1 if x is 0 and the right sub-expression Q sets x to 1 if y is 0. Indeed, if both variables $x, y \in rd(P \,\|\, Q)$ are initially $\rho_0(x) = \rho_0(y) = 0$, the response of $P \,\|\, Q$ is non-determinate (under B-admissible scheduling). If P is first executed to termination and then Q, we get the final memory $\gamma_0(x) = 0, \gamma_0(y) = 1$; otherwise, if we first execute Q and then P, the result will be $\gamma_0(x) = 1, \gamma_0(y) = 0$. This is an internal non-determinism observable from a single fixed initial memory. $P \,\|\, Q$ is B-reactive but not B-determinate and thus neither SC-determinate.

That a program is non-determinate does not mean all its sub-programs must be non-determinate, too. *E.g.*, both fprogs P and Q in this example are SC-read-determinate. *E.g.*, the only read variable $x \in rd(P)$ is not touched by P and thus left to be controlled by the environment. This satisfies condition (2) of Def. 6. Note that the value of y is changed in the SC-admissible execution of P starting from $\rho_0(x) = 0$ and $\rho_0(y) = 0$ and its final value $\gamma_0(y) = 1$ is *not* the final value for all SC-admissible instants, *e.g.*, if $\rho_0(x) = 1$ and $\rho_0(y) = 0$ then we get $\gamma_0(y) = 0$. However, this is not a violation of Def. 6(2) because $y \notin rd(P)$.

Finally note that non-determinate programs can become determinate in context. *E.g.*, the SC-admissibility rules make sure that in $P \,\|\, Q \,\|\, !x$ the set $!x$ is executed before the test $x?$ in P, which means that P does not write $!y$ which prevents Q from writing $!x$, thereby avoiding an admissibility hazard with any earlier read of y by Q. Moreover, since the set $!x$ is executed before the read $x?$ by P, the set $!x$ by Q is confluent with the read. As a consequence, for any given initial memory ρ_0, all SC-admissible executions of $P \,\|\, Q \,\|\, !x$ produce the same determinate response γ_0 with $\gamma_0(x) = 1$ and $\gamma_0(y) = \rho_0(y)$ is the initial value. Thus, the fprog $P \,\|\, Q \,\|\, !x$ is SC-read-determinate.   □

In Sec. 4 below we shall give a sound denotational fixed point analysis to check a program is B-reactive and SC-read-determinate. Our fixed point characterisation defines the class of *input Berry constructive* (IBC) programs which includes more programs than the *strong Berry constructive* (SBC) programs introduced in [4]. The result established in [4], that every SBC program is SC-reactive and SC-determinate, is a corollary of the main Thm. 1 in this article which says that every IBC program is B-reactive and SC-read-determinate.

We first need to introduce the appropriate abstract semantical domains. This is done in the following Sec. 3.

3 Abstract Domains and Environments.

The constructiveness analysis on finite pSCL programs (fprogs) takes place in an abstract domain of information values which describe the sequential and concurrent interaction of signals. It accounts for data dependencies and can deal with the difference of a variable retaining its original initial value from the initial memory (pristine), being initialised to 0 and then either remaining 0 (signal absence) or being set to 1 (signal presence). This includes monotonic value changes from 0 to 1 and, essentially, corresponds to Berry's notion of constructiveness in Esterel [12], yet is able to deal with explicit initialisations which requires the ability to cope with prescriptive sequencing. This section introduces this abstract domain and its natural extension to environments, namely discrete structures able to maintain the information of a number of signal variables.

3.1 Value Domain $I(\mathbb{D})$ of Value Status.

Instead of distinguishing just two signal statuses "absent" and "present" as in traditional SMoCC, we consider the sequential behaviour of a variable (during each instant) as taking place in a linearly ordered 4-valued domain $\mathbb{D} = \{\bot \leq 0 \leq 1 \leq \top\}$. This requires to consider two additional *logical memory values*, namely $\bot$ and $\top$. The former indicates that the corresponding variable contains its initial memory value, *i.e.*, a pristine 0 or 1. The latter tells us that the variable value has passed from 1 to 0 at some point, independently of what the final memory result is. The linear ordering $\leq$ captures a trajectory through a *single* instance of the iur protocol. Observe the difference between the variable values $\mathbb{B} = \{0, 1\}$, which appear at "run-time" as defined in the operational semantics, and the signal statuses $\mathbb{D}$, which are the basis of constructiveness analysis. The latter lifts our description to a higher level in which the semantics of variables is enriched to reflect the fact that they are controlled by an implicit synchronisation protocol. Observe that the ordering $\leq$ in $\mathbb{D}$ is transitive which permits monotonic status changes from $\bot$ directly to 1, without first passing through 0. This means a program can set a variable (emit a signal in Esterel) which has not been explicitly reset. This matches the iur protocol, from which the notions of B/SC-admissibility are derived, which does not require an update to be preceded by an init operation. However, our fixed point semantics can be easily modified, without changing the domain $\mathbb{D}$, for the stronger requirement if needed.

We now go one step further in the abstraction. In the analysis we operate on *predictions* of variable statuses. Possible statuses of variables are approximated by closed *intervals* $I(\mathbb{D}) = \{[a, b] \mid a, b \in \mathbb{D}, a \leq b\}$ over $\mathbb{D}$. An interval $[a, b] \in I(\mathbb{D})$ in this 10-valued domain corresponds to the set of statuses $set([a, b]) = \{x \mid a \leq x \leq b\} \subseteq \mathbb{D}$. Intervals $[a, b]$ such that

$a < b$ denote *uncertain* information, *i.e.*, a potential non-deterministic response. Such a general interval represents an approximation to the final (stable) state of a variable from its two ends, the *lower bound a* and the *upper bound b*. An interval $[a,b]$ associated with a variable $x \in V$ can thus be read as follows: *"the executions of the statements so far ensure that x has currently status a, yet it cannot be excluded that some statements might be executed which could change (increase) the status of x up to b"*. In this vein, the intervals $[a,a]$ correspond to *decided*, or *crisp*, statuses which are naturally identified with the values $\bot = [\bot, \bot]$, $0 = [0,0]$, $1 = [1,1]$ and $\top = [\top, \top]$ of $\mathbb{D}$, respectively, *i.e.*, $\mathbb{D} \subset I(\mathbb{D})$. A variable $s \in V$ with status $\gamma \in I(\mathbb{D})$ is denoted by s^{γ}.

Example 12 When computing the reaction of fprog $\mathbf{i}s \,; x \,? \,!s : \varepsilon$, the interval for s will be $[0,1]$, assuming the status of x is not decided yet, say, $x^{[\bot, \top]}$. The status $s^{[0,1]}$ for variable s indicates that a reset $\mathbf{i}s$ *must* definitively be executed, but there is at least one set $!s$ that *can* potentially be executed, which is why the status of s ranges between 0 and 1.   □

On the domain $I(\mathbb{D})$ we can define two natural orderings:

  – The *point-wise* ordering $[a_1, b_1] \preceq [a_2, b_2]$ iff $a_1 \leq a_2$ and $b_1 \leq b_2$, and
  – the *(inverse) inclusion* ordering $[a_1, b_1] \sqsubseteq [a_2, b_2]$ iff $set([a_2, b_2]) \subseteq set([a_1, b_1])$,

which endow $I(\mathbb{D})$ with a full lattice structure for $\preceq$ and a lower semi-lattice structure for $\sqsubseteq$. The *point-wise* lattice $\langle I(\mathbb{D}), \preceq \rangle$ has minimum element $[\bot, \bot]$ and the minimum for the *inclusion* semi-lattice $\langle I(\mathbb{D}), \sqsubseteq \rangle$ is $[\bot, \top]$. The element $[\top, \top]$ is a maximal element for both orderings but it is the maximum only for $\preceq$. For $\sqsubseteq$ all singleton intervals $[a,a]$ are maximal. Join $\vee$ and meet $\wedge$ for the $\preceq$-lattice are obtained in the point-wise manner:

$$[a_1, b_1] \vee [a_2, b_2] = [max(a_1, a_2), max(b_1, b_2)]$$
$$[a_1, b_1] \wedge [a_2, b_2] = [min(a_1, a_2), min(b_1, b_2)].$$

In the inclusion $\sqsubseteq$-lattice the meet $\sqcap$ is

$$[a_1, b_1] \sqcap [a_2, b_2] = [min(a_1, a_2), max(b_1, b_2)].$$

The semi-lattice $\langle I(\mathbb{D}), \sqsubseteq \rangle$ does not possess joins, but it is *consistent complete*, *i.e.*, whenever in a nonempty subset $\emptyset \neq X \subseteq \mathbb{D}$ any two elements $x_1, x_2 \in X$ have an upper bound $y \in \mathbb{D}$, *i.e.*, $x_1 \sqsubseteq y$ and $x_2 \sqsubseteq y$, then there exists the least upper bound $\sqcup X = \sqcap \{y \mid \forall x \in X. \, x \sqsubseteq y\}$. This will give us least fixed points.

Fig. 3 illustrates the two-dimensional lattice structure of $I(\mathbb{D})$. The vertical direction (upwards, green arrows) corresponds to $\preceq$ and captures the sequential dimension of the statuses. The horizontal direction (right-to-left, blue arrows) is the inclusion ordering $\sqsubseteq$ and expresses the degree of precision of the approximation. The most precise status description is given by the crisp values on the left side, which are $\sqsubseteq$-maximal and are order-isomorphic to the embedded domain $\mathbb{D}$. The least precise information value is the interval $[\bot, \top]$ on the right. The following Ex. 13 illustrates how we can use the domain $I(\mathbb{D})$ in the fixed point analysis to navigate in both dimensions $\preceq$ and $\sqsubseteq$ for determining the instantaneous response of a program.

Example 13 Consider the fprog $P := (x \,? \,\varepsilon : (!y \parallel !z)) \parallel (y \,? \,\varepsilon : !x)$. Suppose that we execute P in a sequential thread in which all three variables are initially pristine, *i.e.*, with status $x^{\bot}$, $y^{\bot}$ and $z^{\bot}$. What are the final values of of the variables when P is completed? Since we do

DRAFT 20.4.2015

Fig. 3 Domain $I(\mathbb{D})$ for Approximating Signal Variable Statuses.

not know what the memory values of x and y are, we do not know how the branches are decided, *i.e.*, whether the first concurrent thread $x\,?\,\varepsilon : (!y \parallel !z)$ will execute ε or set both variables y and z in $!y \parallel !z$. Similarly, we cannot decide if the second thread $y\,?\,\varepsilon : !x$ sets x or not. Yet, what we do know is that the variables x, y or z *may* be set but *cannot* crash because there is no reset on any of them. So, the best approximation for the response of P, in terms of intervals $I(\mathbb{D})$, is the final status $x^{[\perp,1]}$, $y^{[\perp,1]}$, and $z^{[\perp,1]}$.

Now put P in parallel with the program $Q := {}_{\mathrm{i}}x \parallel !y$. Since Q certainly executes the reset ${}_{\mathrm{i}}x$ and no other write accesses to x, this produces the response x^0. Combining this with the status obtained from P gives the joint response $x^{0\vee[\perp,1]} = x^{[0,1]}$. This tells us that x must certainly be reset (viz. by Q) and then might be set (viz. by P). Notice how the interval $[\perp,1]$ has shrunk to $[0,1]$, which provides tighter information. What about variable y? It is set by Q and never reset, which means its status, after executing Q, is at least 1. By the iur protocol the set $!y$ must wait for any potential reset on y to have happened in the environment. In this case, P does not have a reset on y, so the set $!y$ of Q must go ahead, giving y^1 for the response of Q. This merges with the information from P to the joint response $y^{1\vee[\perp,1]} = y^1$.

But now we have narrowed down the status of y to a crisp 1, which implies that the conditional test $y\,?\,\varepsilon : !x$ in the second thread of P is decided. So we conclude that the set $!x$ must definitely be executed. Therefore, the status of x from P in our first approximation can now be tightened from $x^{[\perp,\top]}$ to x^1. Once we have that, the conditional $x\,?\,\varepsilon : (!y \parallel !z)$ in the first thread of P is decided, too, implying that the set $!z$ must be executed by P implying 1 as a lower bound for the status of z and an increase of information from $z^{[\perp,1]}$ to z^1. Since all three variables are now fixed to have crisp statuses x^1, y^1, z^1, the program $P \parallel Q$ is called strongly Berry constructive. From [4] this implies that $P \parallel Q$ is sequentially constructive, *i.e.*, SC-reactive and SC-determinate. From the results in this article it will follow that it is also B-reactive and SC-read-determinate. $\qquad\square$

Observe that the well-known ternary domain (Kleene) for the fixed-point analysis of Pure Esterel [12] or constructive Boolean circuits [65] is captured, as indicated in Fig. 3, by the inner part with values $[0,0]$ ("absent"), $[1,1]$ ("present") and $[0,1]$ ("undefined"). In ternary analysis all signal variables are implicitly assumed initialised, hence no need for $\perp$. Moreover, since there is no reset operator and thus programs cannot fail the monotonic single-change requirement, there is no need for $\top$ either, in languages such as Esterel, as long as initialisation of signals is implemented by the run-time rather than the program. This ternary fragment of $I(\mathbb{D})$ corresponds to three-valued Kleene logic with $\vee$ disjunction and

DRAFT 20.4.2015

Fig. 4 The domain $\mathbb{P}$ coding the initialisation status.

$\wedge$ logical conjunction. Fig. 3 visualises clearly how the 10-valued domain $I(\mathbb{D})$ offers an extended playground to represent the logic of explicit initialisation.

Interestingly, another recent approach to enrich the standard ternary domain is the constructive semantics by Talpin et.al. [79] for a multi-clocked synchronous (polysynchronous) data-flow language which integrates Quartz and Signal. This extension is based on a lattice $\mathscr{D}$ which extends $\{[0,0],[1,1],[0,1]\}$ by elements ? for representing unknown and $\frac{1}{2}$ for inconsistent signal statuses similar to our $\bot$ and $\top$. It also contain Boolean values for "true" and "false" (embedded as refinements of the present status) which our domain $I(\mathbb{D})$ does not model. On the other hand, the partial order $\mathscr{D}$ of [79] does not have an interval structure like $I(\mathbb{D})$, which is the key to modelling Esterel-style reaction to absence. This is not needed in the data-flow semantics of [79].

3.2 Semantic Domain $I(\mathbb{D}, \mathbb{P})$ of Signal Status.

There is one logical refinement to the domain $I(\mathbb{D})$ that we need to make in order to keep properly track of the completion of the initialisation phase on each variable. According to the synchronous protocol a set !s contained in a program can only go ahead if it is guaranteed that no reset ¡s on this variable is possibly outstanding. There is no information in the intervals of $I(\mathbb{D})$ to express that no reset is outstanding. For instance, the status $s^{[0,1]}$ specifies that the initialisation of s has been started and that there is a waiting update access on s, but it does not tell if there are any other resets ¡s still pending. However, this is important in the constructive scheduling, because only if the initialisation phase has been completed, the waiting update !s is permitted to proceed changing the status to $s^{[1,1]}$.

To capture the termination of the initialisation phase of the "init;update;read" protocol, we enrich the interval domain by an additional token $r \in \mathbb{P} = \{0, 1, 2\}$, called the *init status*. The status 2 expresses that the "init" phase is ongoing and a reset is still *predicted*. The status 1 means that no more resets are outstanding, *i.e.*, the init phase is completed, but the protocol is still *running*. Finally, if the "init;update" is *finished*, and thus the value of the variable determined, the init status 0 is obtained.

As for $I(\mathbb{D})$ there are natural sequential and information-theoretic orderings on $\mathbb{P}$ as seen in Fig. 4. The sequential ordering $\preceq$ is given by $0 \preceq 1 \preceq 2$ which reflects the fact that in sequential order a finished computation (0) must first become blocked at a set or a conditional test (1) to start a running protocol, before it reaches a predicted reset ¡s which witnesses an incomplete initialisation (2) for the reset variable s. In contrast, the information ordering on $\mathbb{P}$ is the opposite, $2 \sqsubseteq 1 \sqsubseteq 0$, which models the narrowing of behaviour that occurs when the status of variables becomes more and more decided. The init status 2 is least informative. It says that the protocol is contingent and that there may still be potential resets outstanding.

DRAFT 20.4.2015

Fig. 5 The extended interval domain $I(\mathbb{D}, \mathbb{P})$ including the init status $\mathbb{P} = \{0, 1, 2\}$.

With the value 1 the computation is still contingent but it guarantees that no resets are possible. Finally, 0 is the tightest status for it says that the protocol is finished actual and that no resets are possible.

The domain $(\mathbb{P}, \preceq, \sqsubseteq)$ is a lattice for both $\preceq$ and $\sqsubseteq$ in which only the semi-lattice structure will be relevant induced by the join operations $r_1 \vee r_2 = r_1 \sqcap r_2 = max(r_1, r_2)$. Our definition of constructive behaviours will be based on a fixed point analysis in the product domain

$$I(\mathbb{D}, \mathbb{P}) = \{([l,u], r) \mid [l,u] \in \mathbb{D}, r \in \mathbb{P}\} = I(\mathbb{D}) \times \mathbb{P}.$$

We will write a typical element $([l,u], r) \in I(\mathbb{D}, \mathbb{P})$ more compactly as $[l,u]{:}r$ and refer to the interval $[l,u]$ as the *value status* to separate it from the init status r. If $r = 0$ we simply write $[l,u]$ instead of $[l,u]{:}0$ or even a instead of $[a,a]{:}0$. In this fashion we naturally consider $\mathbb{D}$ as a subset of $I(\mathbb{D}, \mathbb{P})$. Generally, as before, when an interval is a singleton we write it as an element in $\mathbb{D}$, even if its init status is not 0. For instance, 0:1 is the same as $[0,0]{:}1$ or $\perp{:}2$ stands for $[\perp, \perp]{:}2$. These singleton intervals are contained within the dotted regions in Fig. 5.

The orderings $\sqsubseteq$ and $\preceq$ on $I(\mathbb{D}, \mathbb{P})$ are inherited component-wise from the corresponding orderings in the domains $I(\mathbb{D})$ and $\mathbb{P}$, respectively. The init status is logically part of the upper bound and so we define the upper projection on $I(\mathbb{D}, \mathbb{P})$ by stipulating $upp([l,u]{:}r) = [\perp, u]{:}r$, and for the lower projection we set $low([l,u]{:}r) = [l, \top]{:}2$. The same is obtained if we define the upper projection separately on $\mathbb{P}$ as the identity, *i.e.*, $upp(r) = r$ for all $r \in \mathbb{P}$ and the lower projection as the constant function $low(r) = 2$ for all $r \in \mathbb{P}$. Then, upp and low on $I(\mathbb{D}, \mathbb{P})$ are obtained component-wise from upp and low on $I(\mathbb{D})$ and $\mathbb{P}$, respectively.

Note that $I(\mathbb{D}, \mathbb{P})$ is essentially a tripling of $I(\mathbb{D})$, extending the domain $I(\mathbb{D})$ by the information contained in $\mathbb{P}$.[4] This is illustrated Fig. 5.

Example 14 Consider the cprog $P := {\text{¡}}s \, ; x \, ? \, !s : {\text{¡}}s$. Suppose we do not know anything about the status of x in the current environment. This is captured by the status $x^{[\perp, \top]{:}2}$ which is the $\sqsubseteq$-minimal element in $I(\mathbb{D}, \mathbb{P})$. It not only leaves open the full range $[\perp, \top]$ for the value status of x. The init status 2 models an unfinished "init" and a possible outstanding reset on x. Now, if the status of x is so maximally undetermined, the conditional $x \, ? \, !s : {\text{¡}}s$ is undecided.

---

[4] This extra bit for indicating predicted resets has been missing in our publication [4] where this fixed point analysis was introduced for the first time.

We cannot say if the initial reset ¡s in P is followed by the set !s or the reset ¡s. Consequently, the response of P for s will be $[0,1]{:}2$. The init status 2 indicates that the protocol execution of P on s is speculative and that there is a possible reset on s which may become active. The response of P on variable x, on the other hand, yields $x^{\perp:1}$ because the value status is guaranteed to remain pristine but that the computation is nevertheless speculative (because of the blocked conditional test on x).

When the state of x becomes decided with a crisp $x^0 = x^{[0,0]:0}$, then the conditional is switched through into the left branch containing the reset ¡s and the response of P for s refines into $0 = [0,0]{:}0$, too. When x is decided present x^1 then the conditional is unblocked and the set !s is executed. Hence, the response for s becomes $1 = [1,1]{:}0$. Both responses for s have init status 0 stating that the "init;update;read" protocol on s is completed.   □

Example 15 Consider a reset followed by a set, *i.e.*, the cprog $P := {¡x}\,;\,{!x}$. Let us schedule the actions of P starting from the sequential status $S_0 = x^\perp$, or equivalently, $S_0 = x^{[\perp,\perp]:0}$. This represents a fully determined initial memory of unknown value. The reset ¡x is the first action of P to be scheduled, raising the status of x to $S_1 = x^0$. The init status is still 0 because the reset terminates instantaneously. Thus, we reach the set $P' := {!x}$ as the continuation program. To be scheduled the set must wait for the completion of the init phase which depends on the concurrent environment. In the environment $C_0 := x^{[\perp,\top]:2}$ our sequential thread is blocked at the set. However, what we can conclude about the sequential response of P is that x undergoes a reset and then possibly a set, yielding the final status $S_2 := x^{[0,1]:1}$. We cannot put the lower bound to 1 because we have no guarantee that the set is actually executed. Also, the init status 1 informs the environment that the "init;update" in P is blocked but P does not produce any further resets, if it ever were to be continued. Assuming that P is running alone by itself we can strengthen the initial approximation C_0 of the environment by $C_1 := S_2$ and reanalyse P, again from the sequential status S_0. Now as we reach the set !x, the refined environment C_1 with init status 1 unblocks the set !x and we obtain the final sequential status $S_3 := x^1$.   □

The status of variables and their evolutions over time are kept in discrete structures, called *environments* $E : V \to I(\mathbb{D}, \mathbb{P})$ mapping each variable $x \in V$ to a status $E(x) \in I(\mathbb{D}, \mathbb{P})$. The orderings and (semi-)lattice operations are lifted to environments by stipulating

$$E_1 \trianglelefteq E_2 \text{ iff } E_1(x) \trianglelefteq E_2(x) \qquad \text{for } \trianglelefteq \in \{\preceq, \sqsubseteq\} \text{ and}$$
$$(E_1 \odot E_2)(x) = E_1(x) \odot E_2(x) \qquad \text{for } \odot \in \{\vee, \wedge, \sqcap\}.$$

If $E(x) = [a,b]{:}r$ then we will also write $x^{[a,b]:r} \in E$ and further $x^a \in E$ when $E(x) = [a,a]{:}0$. Using this notation we can view environments as sets of variable statuses $E = \{x^{[a,b]:r} \mid E(x) = [a,b]{:}r\}$ with the property that if $x^{[a,b]:r} \in E$ and $x^{[a',b']:r'} \in E$, then $a = a'$ and $b = b'$ and $r = r'$.

It is natural to identify the values $[a,b]{:}r \in I(\mathbb{D})$ with *constant* environments such that $([a,b]{:}r)(x) = [a,b]{:}r$ for all $x \in V$. An environment E is called *decided* if for all variables $x \in V$ there exists $b \in \mathbb{D}$ with $b{:}1 \sqsubseteq E(x)$; *crisp* if for all variables $x \in V$ there exists $b \in \mathbb{D}$ such that $b \sqsubseteq E(x)$; *ternary* if $E(x) \in \{0, 1, [0,1]\}$ for all variables $x \in V$; *crash-free* if $E(x) \preceq 1{:}2$ for all $x \in V$. An environment E in which all entries are one-sided lower (upper) intervals, *i.e.*, in which $x^{[a,b]:r} \in E$ implies $b = \top$ and $r = 2$ ($a = \perp$) is called a *lower* (*upper*) environment. Every environment can be separated into its lower and upper projections

$$low(E) := \{x^{[a,\top]:2} \mid x^{[a,b]:r} \in E\} \qquad\qquad upp(E) := \{x^{[\perp,b]:r} \mid x^{[a,b]:r} \in E\}$$

DRAFT 20.4.2015

so that

$$E = low(E) \sqcup upp(E) = \sqcap\{X \mid low(E) \sqsubseteq X \text{ and } upp(E) \sqsubseteq X\},$$

where the join exists since $low(E) \sqsubseteq E$ and $upp(E) \sqsubseteq E$, *i.e.*, $low(E)$ and $upp(E)$ are always consistent. Observe further that $low(E) = E \vee [\bot, \top]{:}2 = E \sqcap \top{:}2$ and $upp(E) = E \sqcap \bot{:}0$.

We use the set-like notation $\{\langle x_1^{\gamma_1}, x_2^{\gamma_2}, \ldots, x_n^{\gamma_n} \rangle\}$ to specify a finite environment that explicitly sets the status for the listed variables x_i and implicitly defines the status $\bot$ for all other variables $z \in V \setminus \{x_1, x_2, \ldots, x_n\}$. Then, the empty environment $\{\langle\rangle\} = \bot = [\bot, \bot]{:}0$ is the neutral element for $\vee$ which acts as the operator for set union.

Example 16 Let $S_1 = \{\langle x^0, y^{[0,\top]{:}2} \rangle\}$ and $S_2 = \{\langle x^{[\bot,1]{:}1}, z^{[0,1]} \rangle\}$. Then, $S_1 = \{\langle x^0 \rangle\} \vee \{\langle y^{[0,\top]{:}2} \rangle\}$, $S_2 = \{\langle x^{[\bot,1]{:}1} \rangle\} \vee \{\langle z^{[0,1]} \rangle\}$ and

$$S_1 \vee S_2 = \{\langle x^{0 \vee [\bot,1]{:}1}, y^{[0,\top]{:}2 \vee \bot}, z^{\bot \vee [0,1]} \rangle\} = \{\langle x^{[0,1]{:}1}, y^{[0,\top]{:}2}, z^{[0,1]} \rangle\}$$

$$S_1 \sqcap S_2 = \{\langle x^{0 \sqcap [\bot,1]{:}1}, y^{[0,\top]{:}2 \sqcap \bot}, z^{\bot \sqcap [0,1]} \rangle\} = \{\langle x^{[\bot,1]{:}1}, y^{[\bot,\top]{:}2}, z^{[\bot,1]} \rangle\}.$$

$\square$

3.3 Domain $I(\mathbb{C})$ of Completion Status

The completion status for an fprog P in concurrent environment C is given by a set of completion codes $cmpl\langle P, C \rangle \subseteq \mathbb{C} := \{\bot, 0, 1\}$ which model the uncertainty about the termination behaviour of P, analogous to the status intervals for signal variables. The code 0 stands for *instantaneous (normal) termination*, 1 for *pausing* and $\bot$ for *blocking* to model a situation when a program's control flow is stuck at a conditional test which cannot be decided. These completion codes $\mathbb{C}$ must not to be confused with the signal statuses in $\mathbb{D}$.

What is the information content of a subset $cmpl\langle P, C \rangle \subseteq \mathbb{C}$ of completion codes? When $c \in cmpl\langle P, C \rangle$ then c is a *possible* completion of P, but it is not guaranteed unless $cmpl\langle P, C \rangle = \{c\}$ is a singleton, in which case c *must* be the completion type of P. Otherwise, if $c' \neq c$ with $c' \in cmpl\langle P, C \rangle$, then c' is another type of completion that *can* happen for P in environment C. Complementarily, if $c \notin cmpl\langle P, C \rangle$ then c *cannot* occur. The "must" and "cannot" information —which is the basis for defining the completion semantics of programs in Esterel— is completely captured by the five subsets

$$I(\mathbb{C}) := \{\{\bot, 0\}, \{\bot, 1\}, \{\bot, 0, 1\}, \{0\}, \{1\}\}.$$

The sets $\{\}$, $\{0, 1\}$ and $\{\bot\}$ are missing because every program must at least possibly terminate instantaneously or possibly pause, and if a program possesses both possible codes 0 and 1 then this is so because some conditional test cannot be decided, which means it is blocked. So, $\bot$ must be a possible code for this program, too.[5]

---

[5] In other words, the free set-theoretic "collection semantics", which defines the completion code of a program as the set of all it possible completions (under a given choice of environments), would produce exactly the sets in $I(\mathbb{C})$. We could have defined $I(\mathbb{C})$ more generously as the set of subsets of completion codes $\mathscr{P}\{\bot, 0, 1\}$. However, our explicit description reveals more of the algebraic properties of $I(\mathbb{C})$ than $\mathscr{P}\{\bot, 0, 1\}$. For instance, it makes clear that the internal logic of $I(\mathbb{C})$ is not a Boolean algebra.

DRAFT 20.4.2015

The precise relation between $I(\mathbb{C})$ to the completion codes of Esterel [12] is given by defining the sets

$$must_k(P,C) := \{0 \mid cmpl\langle P,C\rangle = \{0\}\} \cup \{1 \mid cmpl\langle P,C\rangle = \{1\}\}$$
$$cannot_k(P,C) := \{0 \mid 0 \notin cmpl\langle P,C\rangle\} \cup \{1 \mid 1 \notin cmpl\langle P,C\rangle\}$$
$$can_k(P,C) := \{0,1\} \setminus cannot_k(P,C) \ = \ cmpl\langle P,C\rangle \setminus \{\bot\}$$

of codes that *must* and *cannot/can* be obtained by program P in environment C, respectively. We observe that $must_k(P,C) \cap cannot_k(P,C) = \emptyset$ and that both $must_k(P,C) \neq \{0,1\}$ and $cannot_k(P,C) \neq \{0,1\}$. This makes sense since must and cannot completions are contradictory and there is no program which must terminate and must pause at the same time, or cannot terminate and cannot pause at the same time. Since we do not consider completion codes for traps, every program can at least potentially terminate or pause. More specifically, $must_k(P,C)$ and $cannot_k(P,C)$ are either empty $\emptyset$ or a singleton set $\{0\}$ or $\{1\}$. Also, directly from the definition we find that if $must_k(P,C)$ is a singleton, then $cannot_k(P,C)$ is the complementary singleton set, *i.e.*, $must_k(P,C) = \{0\}$ implies $cannot_k(P,C) = \{1\}$ and $must_k(P,C) = \{1\}$ implies $cannot_k(P,C) = \{0\}$. Finally, $must_k(P,C) = \emptyset$ iff $\bot \in cmpl\langle P,C\rangle$ and $cannot_k(P,C) = \emptyset$ iff $cmpl\langle P,C\rangle = \{\bot,0,1\}$.

Note that (i) every P has at least one possible completion status, *i.e.*, $0 \in cmpl\langle P,C\rangle$ or $1 \in cmpl\langle P,C\rangle$ and (ii) if we cannot decide whether P terminates instantaneously or pauses then this is because we cannot decide if P completes at all, *i.e.*, if $\{0,1\} \subseteq cmpl\langle P,C\rangle$ then $\bot \in cmpl\langle P,C\rangle$. This explains why not all of the eight possible subsets of $\mathbb{C}$ can occur as the completion status of a program.

4 Denotational Semantics of Synchronous Programs.

Now that the technical apparatus of status intervals and environments is in place it is time to put it to use. What we will do in this section is to introduce an extended version of the causality analysis for Esterel, which includes initialisation. This analysis defines the class of constructive programs. This analysis performs an abstract program simulation using the interval environments $I(\mathbb{D},\mathbb{P})$ introduced above. To keep matters simple we consider only finite pSCL programs (fprogs), *i.e.*, programs without *rec*. This is without loss of generality. Since well-formed pSCL programs are clock-guarded, we can unfold all loops and extract finite *rec*-free expressions that fully describe the program's macro step reactions. We first describe the computation of completion codes in Sec. 4.1 and then the computation of program responses in Sec. 4.2.

4.1 Computing Completion Codes.

How are completion codes computed for a program P and environment C? As for the response semantics $\langle\langle P\rangle\rangle$ this is done by structural recursion on P. However, while the computation of the sets $must_k(P,C)$ and $cannot_k(P,C)$ in [12] is performed separately through a combinatorial construction, we here give a uniform and algebraic definition of the same information for $cmpl\langle P,C\rangle$. Specifically, we exploit that $I(\mathbb{C})$, like $I(\mathbb{D})$, forms a meet semi-lattice under the (inverse) inclusion ordering $\sqsubseteq$, *i.e.*, $\gamma_1 \sqsubseteq \gamma_2$ iff $\gamma_2 \subseteq \gamma_1$. The completion set $\{\bot,0,1\}$ is the minimal element in $I(\mathbb{C})$ and the meet $\sqcap$ is $\gamma_1 \sqcap \gamma_2 = \gamma_2 \sqcap \gamma_1 = \gamma_1$ if $\gamma_1 \sqsubseteq \gamma_2$ and $\gamma_1 \sqcap \gamma_2 = \{\bot,0,1\}$ if γ_1 and γ_2 are $\sqsubseteq$-incomparable. Let $\oplus$ be the strict lifting of Boolean

DRAFT 20.4.2015

summation to $\mathbb{C}$, *i.e.*, $0 \oplus 1 = 1 = 1 \oplus 0 = 1 \oplus 1$ and $0 \oplus 0 = 0$, while $x \oplus y = \bot$ iff $x = \bot$ or $y = \bot$. This can then further be lifted to completion sets, $\gamma_1 \oplus \gamma_2 := \{x \oplus y \mid x \in \gamma_1, y \in \gamma_2\}$. Notice that if we consider the completion codes 0 and 1 as numbers, then $\oplus$ is the same as *max*. Indeed, $\oplus$ on $I(\mathbb{C})$ is analogous to $\vee$ on $I(\mathbb{D})$. The *upper projection* is given by $upp(\gamma) := \gamma \cup \{\bot\}$. One shows that $\oplus$ and *upp* are well-defined on $I(\mathbb{C})$ and monotonic with respect to $\sqsubseteq$.

$$
\begin{aligned}
cmpl\langle P, C \rangle &:= 0 &&\text{if } P \text{ is one of } \varepsilon \text{ or } \mathrm{i}s \\
cmpl\langle !s, C \rangle &:= \begin{cases} \{0\} & \text{if } [\bot, \top]{:}1 \sqsubseteq C(s) \\ \{\bot, 0\} & \text{otherwise} \end{cases} \\
cmpl\langle \pi, C \rangle &:= 1 \\
cmpl\langle P \,\|\, Q, C \rangle &:= cmpl\langle P, C \rangle \oplus cmpl\langle Q, C \rangle \\[4pt]
cmpl\langle P \,;\, Q, C \rangle &:= \begin{cases} cmpl\langle P, C \rangle & \text{if } 0 \notin cmpl\langle P, C \rangle \\ cmpl\langle P, C \rangle \oplus cmpl\langle Q, C \rangle & \text{otherwise} \end{cases} \\[4pt]
cmpl\langle s \,?\, P : Q, C \rangle &:= \begin{cases} cmpl\langle P, C \rangle & \text{if } 1{:}1 \sqsubseteq C(s) \\ cmpl\langle Q, C \rangle & \text{if } 0{:}1 \sqsubseteq C(s) \\ upp(cmpl\langle P, C \rangle) \sqcap upp(cmpl\langle Q, C \rangle) & \text{otherwise} \end{cases}
\end{aligned}
$$

Fig. 6 Denotational analysis of completion codes for fprogs.

The function $cmpl\langle P, C \rangle \in I(\mathbb{C})$ is as described in Fig. 6. One shows by induction on P that if P is purely combinational, *i.e.*, it does not contain the π operator, then $cmpl\langle P, C \rangle = \{0\}$ or $cmpl\langle P, C \rangle = \{\bot, 0\}$. Furthermore, it is easy to see that the only way in which the status $\bot$ can enter the completion set is through the 'otherwise' case of a set or a conditional. More strictly, we have $\bot \in cmpl\langle P, C \rangle$ iff (i) the control flow reaches some set $!s$ in P which is blocked on the condition $[\bot, \top]{:}1 \not\sqsubseteq C(s)$, or (ii) there is some conditional $s \,?\, P' : Q'$ executed in P for which the guard variable s is undecided, *i.e.*, $1{:}1 \not\sqsubseteq C(s)$ and $0{:}1 \not\sqsubseteq C(s)$. The condition $[\bot, \top]{:}1 \sqsubseteq C(s)$ in the definition of $cmpl\langle !s, C \rangle$ requires that the init status of $C(s)$ is at most 1, *i.e.*, that initialisations $\mathrm{i}s$ are no longer possible. However, this does not constrain the value status. If we wanted to make a set $!s$ wait for at least one initialisation $\mathrm{i}s$ to take place, we could strengthen the condition $[\bot, \top]{:}1 \sqsubseteq C(s)$ to $[0, \top]{:}1 \sqsubseteq C(s)$.

Example 17 The completion intervals $\{0\}$ and $\{1\}$ are obtained from the pSCL expressions ε and π, respectively. The intervals $\{\bot, 0\}$ and $\{\bot, 1\}$ are the completion codes for expressions $x \,?\, \varepsilon : \varepsilon$ and $x \,?\, \pi : \pi$ in every concurrent environment C with $0{:}1 \not\sqsubseteq C$ and $1{:}1 \not\sqsubseteq C$. Finally, if x is undecided, we get $cmpl\langle x \,?\, \varepsilon : \pi, C \rangle = \{\bot, 0, 1\}$. The completion statuses $\{\bot, 0\}$ and $\{\bot, 1\}$ may also be obtained from programs $!x \,;\, \varepsilon$ and $!x \,;\, \pi$, respectively, in an environment C where $\bot{:}2 \preceq C(x)$. $\qquad\square$

4.2 Computing Program Responses.

The denotational semantics of a fprog P is given by a function $\langle\langle P \rangle\rangle_C^S$ that determines constructive information on the instantaneous response of P to an external stimulus consisting of a *sequential* environment S and a *concurrent* environment C. The sequential context S can be thought of as an initialisation under which P is activated. It represents knowledge about

the status of variables sequentially before P is started. In contrast, the parallel environment C contains the external stimulus which is concurrent with P. The lower bound $low \langle\!\langle P \rangle\!\rangle_C^S$ of the response tells us what P *must* write to the variables and the upper bound $upp \langle\!\langle P \rangle\!\rangle_C^S$ is the level that the variables *may* reach upon execution of P.

$$\langle\!\langle \varepsilon \rangle\!\rangle_C^S := S \qquad\qquad\qquad\qquad \langle\!\langle \pi \rangle\!\rangle_C^S := S$$

$$\langle\!\langle ¡s \rangle\!\rangle_C^S := \begin{cases} S \vee \{\!\{s^\top\}\!\} & \text{if } 1 \preceq S(s) \preceq \top \\ S \vee \{\!\{s^{\top:2}\}\!\} & 1{:}1 \preceq S(s) \\ S \vee \{\!\{s^0\}\!\} & \text{if } S(s) \preceq 0 \\ S \vee \{\!\{s^{0:2}\}\!\} & \text{if } \bot{:}1 \preceq S(s) \preceq 0{:}2 \\ S \vee \{\!\{s^{[0,\top]:2}\}\!\} & \text{otherwise} \end{cases} \qquad \langle\!\langle !s \rangle\!\rangle_C^S := \begin{cases} S \vee \{\!\{s^1\}\!\} & \text{if } [\bot,\top]{:}1 \sqsubseteq C(s) \\ S \vee \{\!\{s^{[\bot,1]}\}\!\} \vee \bot{:}1 & \text{otherwise} \end{cases}$$

$$\langle\!\langle P \,\|\, Q \rangle\!\rangle_C^S := \langle\!\langle P \rangle\!\rangle_C^S \vee \langle\!\langle Q \rangle\!\rangle_C^S$$

$$\langle\!\langle s \,?\, P : Q \rangle\!\rangle_C^S := \begin{cases} \langle\!\langle P \rangle\!\rangle_C^S & \text{if } 1{:}1 \sqsubseteq C(s) \\ \langle\!\langle Q \rangle\!\rangle_C^S & \text{if } 0{:}1 \sqsubseteq C(s) \\ S \vee upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot:1} \vee upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot:1} & \text{otherwise} \end{cases}$$

$$\langle\!\langle P \,;\, Q \rangle\!\rangle_C^S := \begin{cases} \langle\!\langle P \rangle\!\rangle_C^S & \text{if } 0 \notin cmpl\langle P, C \rangle \\ \langle\!\langle Q \rangle\!\rangle_C^{\langle\!\langle P \rangle\!\rangle_C^S} & \text{if } cmpl\langle P, C \rangle = \{0\} \\ \langle\!\langle P \rangle\!\rangle_C^S \vee upp\left(\langle\!\langle Q \rangle\!\rangle_C^{\langle\!\langle P \rangle\!\rangle_C^S} \right) & \text{otherwise} \end{cases}$$

Fig. 7 Denotational response analysis for fprogs (the function $cmpl\langle P,C \rangle$ is explained in Fig. 6.

The function $\langle\!\langle P \rangle\!\rangle_C^S$ is defined by recursion on the structure of the fprog P as seen in Fig. 7.

– The empty fprog $\langle\!\langle \varepsilon \rangle\!\rangle_C^S$ passes out its sequential stimulus S and does not add anything to it. The same applies to the pausing program π.

– The result of resetting a variable $\langle\!\langle ¡s \rangle\!\rangle_C^S$ depends on whether the sequential stimulus S already contains a status 1 for s or not and on the init status for s:
  – If $1 \preceq S(s) \preceq \top$, then the sequential status is $S(s) = [a,b]{:}r$ where the value status $[a,b]$ is one of $\{1, [1,\top], \top\}$ and the init status is $r = 0$. This indicates that s *must* have been set sequentially before the execution of the reset $¡s$. Hence, we must crash s since a change from 1 to 0 falls outside of the model. Also, $r = 0$ means that the scheduling control flow has reached the reset $¡s$ and since it terminates instantaneously the down-stream computation continues with the init status 0. All other variables $x \neq s$ retain their status from S. This is what $S \vee \{\!\{s^\top\}\!\}$ achieves, viz. $(S \vee \{\!\{s^\top\}\!\})(s) = S(s) \vee \{\!\{s^\top\}\!\}(s) = S(s) \vee \top = \top$ and $(S \vee \{\!\{s^\top\}\!\})(x) = S(x) \vee \{\!\{s^\top\}\!\}(x) = S(x) \vee \bot = S(x)$.
  – If $1{:}1 \preceq S(s)$ then $S(s) = [l,u]{:}r$ with a value status $[l,u]$ in the set $\{1, [1,\top], \top\}$ as above, but now the init status is $r \succeq 1$. Hence the up-stream computation must have set the variable but is still contingent, so that the $¡s$ is speculative. In this case we crash the value status and raise the init status to 2 since the reset is executed only speculatively. We must consider it as a possibly outstanding reset. The response, therefore is $S \vee \{\!\{s^{\top:2}\}\!\}$.

DRAFT 20.4.2015

- If $S(s) \preceq 0$ then the sequential status of s is one of $S(s) \in \{\perp, [0, \perp], 0\}$ again with init status 0. This says that the upstream computation has finished and s *cannot* have been set before. So we can execute the reset by returning $(S \vee \{\langle s^0 \rangle\})(s) = 0$. The init status stays 0 because the schedule passes the reset ¡s which terminates instantaneously.

- If $\perp{:}1 \preceq S(s) \preceq 0{:}2$ then $S(s) = [l, u]{:}r$ with $u \leq 0$ and $1 \preceq r$. The constraint $u \leq 0$ again guarantees that s is not set before while $1 \preceq r$ tells us that the up-stream schedule is contingent. Consequently, we must put the init status to 2 to record that the ¡s is only speculative. This gives the response $(S \vee \{\langle s^{0:2} \rangle\})(s) = 0{:}2$.

- Finally, the remaining cases are $S(s) = [a, b] : r$, where $a < 1$, $b \geq 1$ and $1 \preceq r$. These cases are subsumed by the constraint $[\perp, 1]{:}1 \preceq S(s) \preceq [0, \top]{:}2$. These statuses say that s *may* have been set before. We can neither be sure that a set on s must have happened earlier, nor that it cannot have happened. So, the execution of ¡s *may* crash the model, whence the result $S \vee \{\langle s^{[0,\top]:2} \rangle\}$ forces the value status of s to be $[0, \top]$. The init status must be 2 because the speculative control flow passes a reset.

- Setting a variable $\langle\langle !s \rangle\rangle_C^S$ updates the sequential environment S with the status s^1 for variable s. However, the "init;update;read" protocol permits a set $!s$ to be executed only if and when the init phase on s has been completed. This is checked by the condition $[\perp, \top]{:}1 \sqsubseteq C(s)$ on the environment which is the same as $C(s) \preceq \top{:}1$. If $C(s) \preceq \top{:}1$ then $C(s) = [l, u]{:}r$ with $r \preceq 1$. Thus, there cannot be any contingent reset still outstanding and we can execute the set $!s$ which terminates instantaneously. This gives the response $(S \vee \{\langle s^1 \rangle\})(s) = S(s) \vee 1$. On the other hand, if $C(s) \npreceq \top{:}1$, then the update $!s$ is blocked and only executed speculatively. In this case, the set $!s$ only forces the status of s to be in the interval $[\perp, 1]$. This leaves open if the set is actually executed or not. Also, the init status for all variables must be set to 1 in order to inform any sequential successor that its execution is only speculative rather than factual. Hence our definition of the response as $S \vee \{\langle s^{[\perp, 1]} \rangle\} \vee \perp{:}1$.

- The response of a parallel $\langle\langle P \| Q \rangle\rangle_C^S$ is obtained by letting each of the children P, Q react to the S and C environments, independently, and then combine their responses using $\vee$. This implements a logical disjunction on Boolean values and implements the idea that in B-admissible executions resets happen before any concurrent sets of a variable. If one of $\langle\langle P \rangle\rangle_C^S$ or $\langle\langle Q \rangle\rangle_C^S$ generates a crash, then the composition $\langle\langle P \| Q \rangle\rangle_C^S$ does so, too. Also the init status of combined with the join $\vee$ operator: The schedule of the "init;update" phases on a variable s in the parallel composition is completed, $\langle\langle P \| Q \rangle\rangle_C^S(s) \preceq \top{:}0$ if and only if the scheduling of both threads is completed, *i.e.*, if both $\langle\langle P \rangle\rangle_C^S(s) \preceq \top{:}0$ and $\langle\langle Q \rangle\rangle_C^S(s) \preceq \top{:}0$ Further, the schedule of $P \| Q$ is blocked and has a speculative reset, $\langle\langle P \| Q \rangle\rangle_C^S(s) \succeq \perp{:}2$ iff in one of the threads a reset is pending, *i.e.*, if $\langle\langle P \rangle\rangle_C^S(s) \succeq \perp{:}2$ or $\langle\langle Q \rangle\rangle_C^S(s) \succeq \perp{:}2$.

- In order to derive information about the variables' status under arbitrary SC-admissible scheduling, conditionals need to be evaluated cautiously. The result of a branching test $s ? P : Q$ can only be predicted if and when the value of s has been firmly established as a decided 0 or 1 under all possible SC-admissible schedules. The decision value for s is taken from the concurrent environment C. Accordingly, if $1{:}1 \sqsubseteq C(s)$ then $\langle\langle s ? P : Q \rangle\rangle_C^S$ behaves like $\langle\langle P \rangle\rangle_C^S$ and if $0{:}1 \sqsubseteq C(s)$ the result of the evaluation is $\langle\langle Q \rangle\rangle_C^S$. As long as the value of s is still undecided, *i.e.*, if $1{:}1 \not\sqsubseteq C(s)$ and $0{:}1 \not\sqsubseteq C(s)$, we cannot know if branch P or Q will be executed. However, at least the write accesses already recorded in the sequential environment S *must* become effective. This gives the condition

DRAFT 20.4.2015

$low \langle\!\langle s \ ? \ P : Q \rangle\!\rangle_C^S = low(S)$ for the lower bound. A write access *may* be produced by $s \ ? \ P : Q$ if it *may* be generated by S or by one of the branches P or Q. So, we speculatively compute the response of P and Q in the sequential environment $S \vee \bot{:}1$. This sets the init status of all variables to 1 (at least) in order to mark all write accesses in P and Q as speculative. This implies $upp \langle\!\langle s \ ? \ P : Q \rangle\!\rangle_C^S = upp(S) \vee upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1} \vee upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1}$ for the upper bound. Both can be expressed by the single equation $\langle\!\langle s \ ? \ P : Q \rangle\!\rangle_C^S = S \vee upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1} \vee upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1}$ which is seen as follows:

$$low(S \vee upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1} \vee upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1})$$
$$= low(S) \vee low(upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1}) \vee low \, upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1}$$
$$= low(S) \vee [\bot, \top] : 2 \vee [\bot, \top] : 2$$
$$= low(S) \vee [\bot, \top] : 2$$
$$= S \vee [\bot, \top] : 2 \vee [\bot, \top] : 2 \ = \ S \vee [\bot, \top] : 2 \ = \ low(S)$$

by the properties of $\vee$ and the projections and similarly

$$upp(S \vee upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1} \vee upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1})$$
$$= upp(S) \vee upp \, upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1}) \vee upp \, upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1}$$
$$= upp(S) \vee upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1} \vee upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1}.$$

Notice that $upp \langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1} \vee upp \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1}$ is the same as $upp(\langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1} \sqcap \langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1})$, the upper projection of the best over-approximation of both $\langle\!\langle P \rangle\!\rangle_C^{S \vee \bot{:}1}$ and $\langle\!\langle Q \rangle\!\rangle_C^{S \vee \bot{:}1}$. It is here that the meet operator $\sqcap$ is hidden in the semantics.

– The response of a sequential composition $P \, ; \, Q$ depends on a set of possible completion codes $cmpl\langle P, C \rangle \subseteq \{\bot, 0, 1\}$ from which we can tell whether P is known to terminate or pause or neither. The code 0 stands for instantaneous termination, 1 for pausing and $\bot$ for "unknown" or "blocked", to model the situation when P's control flow is stuck at a conditional test which cannot be decided. If $0 \notin cmpl\langle P, C \rangle$ then P cannot terminate instantaneously. In this case, Q will never be executed in the current instant, so that $\langle\!\langle P \, ; \, Q \rangle\!\rangle_C^S = \langle\!\langle P \rangle\!\rangle_C^S$. However, if $cmpl\langle P, C \rangle = \{0\}$, then P is guaranteed to terminate instantaneously. Thus, the overall response $\langle\!\langle P \, ; \, Q \rangle\!\rangle_C^S$ is that of Q reacting to the concurrent stimulus C and using the response $\langle\!\langle P \rangle\!\rangle_C^S$ as the sequential stimulus. Otherwise if $0 \in cmpl\langle P, C \rangle$ and $cmpl\langle P, C \rangle \neq \{0\}$, then this means that some conditional test on the execution path in P cannot be decided in C. Thus, it is not known yet how P will complete and, as a consequence, if Q will be executed. Therefore, we can only say a variable *must* be written by $P \, ; \, Q$ if it *must* be written by P in the present environments S and C. This leads to $low \langle\!\langle P \, ; \, Q \rangle\!\rangle_C^S = low \langle\!\langle P \rangle\!\rangle_C^S$. As regards upper bounds, a variable *may* be written if it *may* be written by Q with the response of P as its sequential stimulus: $upp \langle\!\langle P \, ; \, Q \rangle\!\rangle_C^S = upp \langle\!\langle Q \rangle\!\rangle_C^{\langle\!\langle P \rangle\!\rangle_C^S}$. One can show, as above in the case of conditionals, that both lower and upper bound equations can be combined into $\langle\!\langle P \, ; \, Q \rangle\!\rangle_C^S = \langle\!\langle P \rangle\!\rangle_C^S \vee upp \langle\!\langle Q \rangle\!\rangle_C^{\langle\!\langle P \rangle\!\rangle_C^S}$, or equivalently $\langle\!\langle P \, ; \, Q \rangle\!\rangle_C^S = \langle\!\langle P \rangle\!\rangle_C^S \sqcap \langle\!\langle Q \rangle\!\rangle_C^{\langle\!\langle P \rangle\!\rangle_C^S}$.

Example 18 Consider the fprog $P := (x \ ? \ \varepsilon : (!y \parallel !z)) \parallel (y \ ? \ \varepsilon : !x)$ with the environments $S = \{\!\langle \ \rangle\!\} = \bot$ and $C_0 = \{[\,]\} = [\bot, \top]{:}2$. The response $\langle\!\langle P \rangle\!\rangle_{C_0}^S$ is the information to be got from a single pass through P without letting P communicate with itself. In doing that the sequential environment S sums up the variable status that has been established by the upstream control

DRAFT 20.4.2015

flow as the execution reaches P. The environment C_0 accumulates our information about the global status of all variables, including the concurrent environment in which P is running. Considering that neither x nor y is decided in C_0, both the conditionals block. Since the updates $!x$, $!y$, $!z$ may possibly be executed and there is no later reset, the variables' expected status is at least $\perp$ and at most 1, *i.e.*, $\langle\!\langle P \rangle\!\rangle^S_{C_0} = \perp{:}1 \vee \{\!\langle x^{[\perp,1]}, y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\}$. The init status $\perp{:}1$ is imposed to record that the computation for all variables is incomplete, yet there is no contingent reset for any of them. Indeed, this is what the calculation using Fig. 7 obtains: The response of the first thread is

$$
\begin{aligned}
\langle\!\langle x \ ? \ \varepsilon : (!y \parallel !z) \rangle\!\rangle^S_{C_0} &= S \vee upp \, \langle\!\langle \varepsilon \rangle\!\rangle^{S \vee \perp:1}_{C_0} \vee upp \, \langle\!\langle !y \parallel !z \rangle\!\rangle^{S \vee \perp:1}_{C_0} \\
&= S \vee upp(S \vee \perp{:}1) \vee upp(\langle\!\langle !y \rangle\!\rangle^{S \vee \perp:1}_{C_0} \vee \langle\!\langle !z \rangle\!\rangle^{S \vee \perp:1}_{C_0}) \\
&= S \vee upp(S \vee \perp{:}1) \vee upp(S \vee \perp{:}1 \vee \{\!\langle y^1 \rangle\!\} \vee S \vee \perp{:}1 \vee \{\!\langle z^1 \rangle\!\}) \\
&= \perp \vee upp(\perp{:}1) \vee upp(\perp{:}1 \vee \{\!\langle y^1 \rangle\!\} \vee \perp{:}1 \vee \{\!\langle z^1 \rangle\!\}) \\
&= \perp{:}1 \vee upp\,\{\!\langle y^1, z^1 \rangle\!\} \ = \ \perp{:}1 \vee \{\!\langle y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\}.
\end{aligned}
$$

Similarly, we obtain $\langle\!\langle y \ ? \ \varepsilon : !x \rangle\!\rangle^S_{C_0} = \perp{:}1 \vee \{\!\langle x^{[\perp,1]} \rangle\!\}$ for the second thread. Joined together, the parallel composition then is

$$
\langle\!\langle P \rangle\!\rangle^S_{C_0} = \perp{:}1 \vee \{\!\langle y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\} \vee \{\!\langle x^{[\perp,1]} \rangle\!\} = \perp{:}1 \vee \{\!\langle x^{[\perp,1]}, y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\}
$$

as claimed.

Without further assumptions on the environment this is the end of the story, none of the variables' value status can be decided beyond $[\perp, 1]$. One shows that $cmpl\langle P, C_0 \rangle = \{\perp, 0\}$, *i.e.*, P does not terminate. Now put P in parallel with fprog $Q := \text{¡}x \parallel !y$, to continue the discussion begun in Ex. 11. Running Q from S and C_0 gives $\langle\!\langle Q \rangle\!\rangle^S_{C_0} = \perp{:}1 \vee \{\!\langle x^0, y^{[\perp,1]} \rangle\!\}$. The response is contingent because the set $!y$ cannot proceed in C_0 which does not exclude further resets on y. Therefore,

$$
C_1 = \langle\!\langle P \parallel Q \rangle\!\rangle^S_{C_0} = \perp{:}1 \vee \{\!\langle x^{[\perp,1]}, y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\} \vee \{\!\langle x^0, y^1 \rangle\!\} = \perp{:}1 \vee \{\!\langle x^{[0,1]}, y^1, z^{[\perp,1]} \rangle\!\}.
$$

This says that x must be reset but may be set later (stabilising without crash), y and z may remain pristine or stabilise at 1. In addition, the init status of all variables is 1, excluding any further possible resets arising from $P \parallel Q$. Notice that C_1 is a more precise description of the response compared to C_0, *i.e.*, $C_0 \sqsubset C_1$.

The remaining uncertainty arises because the single application of $\langle\!\langle P \parallel Q \rangle\!\rangle^S_{C_0}$ blocks the setting of y in the write access in Q. For this, $P \parallel Q$ needs to communicate with itself to find out that the set $!y$ can proceed. This is achieved by running a second pass, now feeding the concurrent environment C_1 instead of C_0. Since C_1 indicates a completed "init" phase for y the set $!y$ in Q is unblocked. We find $\langle\!\langle Q \rangle\!\rangle^S_{C_1} = \{\!\langle x^0, y^1 \rangle\!\}$. Since variable y is now a decided 1 the conditional in the second thread of P is turned off which makes the set $!x$ non-executable, so variable x cannot be set. The calculation for the second thread now is $\langle\!\langle y \ ? \ \varepsilon : !x \rangle\!\rangle^S_{C_1} = \langle\!\langle \varepsilon \rangle\!\rangle^S_{C_1} = S = \perp$. It terminates, *i.e.*, $cmpl\langle y \ ? \ \varepsilon : !x, C_1 \rangle = \{0\}$, as one shows without difficulty from the definition in Fig. 6. The first thread still does not terminate because x is still undecided in C_1 and we have $\langle\!\langle x \ ? \ \varepsilon : (!y \parallel !z) \rangle\!\rangle^S_{C_1} = \perp{:}1 \vee \{\!\langle y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\}$ as before. This means $\langle\!\langle P \rangle\!\rangle^S_{C_1} = \perp{:}1 \vee \{\!\langle y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\} \vee \perp = \{\!\langle y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\}$.

Thus, overall, this gives the refined response

$$
C_2 := \langle\!\langle P \parallel Q \rangle\!\rangle^S_{C_1} = \perp{:}1 \vee \{\!\langle y^{[\perp,1]}, z^{[\perp,1]} \rangle\!\} \vee \{\!\langle x^0, y^1 \rangle\!\} = \perp{:}1 \vee \{\!\langle x^0, y^1, z^{[\perp,1]} \rangle\!\}
$$

DRAFT 20.4.2015

which is a more precise status description, *i.e.*, $C_1 \sqsubset C_2$, since C_2 now also endows variable x with a decided value 0. As a result, the conditional in the first thread of P must execute $!z$ which finally resolves the status of z: $\langle\!\langle x ? \varepsilon : (!y \parallel !z)\rangle\!\rangle^S_{C_2} = \langle\!\langle !y \parallel !z)\rangle\!\rangle^S_{C_2} = \{\langle y^1, z^1\rangle\}$ which means

$$
\begin{aligned}
C_3 = \langle\!\langle P \parallel Q\rangle\!\rangle^S_{C_2} &= \langle\!\langle P\rangle\!\rangle^S_{C_2} \vee \langle\!\langle Q\rangle\!\rangle^S_{C_2} \\
&= \langle\!\langle x ? \varepsilon : (!y \parallel !z)\rangle\!\rangle^S_{C_2} \vee \langle\!\langle y ? \varepsilon : !x\rangle\!\rangle^S_{C_2} \vee \langle\!\langle Q\rangle\!\rangle^S_{C_2} \\
&= \{\langle y^1, z^1\rangle\} \vee \bot \vee \{\langle x^0, y^1\rangle\} = \{\langle x^0, y^1, z^1\rangle\}.
\end{aligned}
$$

The environment C_3, which satisfies $C_2 \sqsubset C_3$, is a crisp fixed point, $\langle\!\langle P \parallel Q\rangle\!\rangle^S_{C_3} = C_3$, in which the parallel composition $P \parallel Q$ terminates, *i.e.*, $cmpl\langle P \parallel Q, C_3\rangle = \{0\}$. $\qquad\square$

Ex. 18 is what we shall call a strongly Berry-constructive program (cf. Def. 7) which generates a crisp fixed point response. This implies (cf. Thm. 1) that the program is B-reactive and SC-read-determinate. There are however programs which cannot be scheduled because they contain a causal cycle which makes the schedule lock up. These deadlocks arise from the "init;update;read" protocol constraint that makes read accesses wait for the prior completion of all possible write accesses and sets wait for the completion of any possible resets. The following examples illustrates the two typical cases of deadlocks.

Example 19 The program $P_1 := !x \,;\, {\textrm{¡}}y \parallel !y \,;\, {\textrm{¡}}x$ is not constructive. Indeed it does not admit of any SC-admissible (and hence neither any B-admissible) schedule because in all its free schedules a reset action happens after a concurrent set action to the same variable, yet they are not confluent with each other. Hence, each schedule violates SC-admissibility. Also, the final memory is non-deterministic depending on the schedule. If we chose the sequence $!x \,;\, !y \,;\, {\textrm{¡}}x \,;\, {\textrm{¡}}y$ the final memory has $y = 0$, whereas if we schedule $!x \,;\, {\textrm{¡}}y \,;\, !y \,;\, {\textrm{¡}}x$ the we get $y = 1$. If we run the fixed point analysis the problem becomes visible as a deadlock: From $S := \bot$ and $C_0 := [\bot, \top]{:}2$ the two concurrent sets $!x$ and $!y$ both block so that $\langle\!\langle !x\rangle\!\rangle^S_{C_0} = \bot{:}1 \vee \{\langle x^{[\bot,1]}\rangle\}$ and $\langle\!\langle !y\rangle\!\rangle^S_{C_0} = \bot{:}1 \vee \{\langle y^{[\bot,1]}\rangle\}$. Then, because the sets guard the resets ¡y and ¡x, respectively, their init status is set to 2:

$$
\begin{aligned}
\langle\!\langle P_1\rangle\!\rangle^S_{C_0} &= \langle\!\langle !x \,;\, {\textrm{¡}}y \parallel !y \,;\, {\textrm{¡}}x\rangle\!\rangle^S_{C_0} \\
&= \langle\!\langle !x \,;\, {\textrm{¡}}y\rangle\!\rangle^S_{C_0} \vee \langle\!\langle !y \,;\, {\textrm{¡}}x\rangle\!\rangle^S_{C_0} \\
&= \langle\!\langle !x\rangle\!\rangle^S_{C_0} \vee upp\,\langle\!\langle {\textrm{¡}}y\rangle\!\rangle^{\langle\!\langle !x\rangle\!\rangle^S_{C_0}}_{C_0} \vee \langle\!\langle !y\rangle\!\rangle^S_{C_0} \vee upp\,\langle\!\langle {\textrm{¡}}x\rangle\!\rangle^{\langle\!\langle !y\rangle\!\rangle^S_{C_0}}_{C_0} \\
&= \bot{:}1 \vee \{\langle x^{[\bot,1]}\rangle\} \vee upp\,\langle\!\langle {\textrm{¡}}y\rangle\!\rangle^{\bot{:}1 \vee \{\langle x^{[\bot,1]}\rangle\}}_{C_0} \vee \bot{:}1 \vee \{\langle y^{[\bot,1]}\rangle\} \vee upp\,\langle\!\langle {\textrm{¡}}x\rangle\!\rangle^{\bot{:}1 \vee \{\langle y^{[\bot,1]}\rangle\}}_{C_0} \\
&= \bot{:}1 \vee \{\langle x^{[\bot,1]}\rangle\} \vee \{\langle y^{[\bot,1]}\rangle\} \\
&\qquad \vee upp(\bot{:}1 \vee \{\langle x^{[\bot,1]}\rangle\} \vee \{\langle y^{0:2}\rangle\}) \vee upp(\bot{:}1 \vee \{\langle y^{[\bot,1]}\rangle\} \vee \{\langle x^{0:2}\rangle\}) \\
&= \bot{:}1 \vee \{\langle x^{[\bot,1]}\rangle\} \vee \{\langle x^{[\bot,0]:2}\rangle\} \vee \{\langle y^{[\bot,1]}\rangle\} \vee \{\langle y^{[\bot,0]:2}\rangle\} \\
&= \bot{:}1 \vee \{\langle x^{[\bot,1]:2}\rangle\} \vee \{\langle y^{[\bot,1]:2}\rangle\}.
\end{aligned}
$$

In this updated environment $C_1 := \langle\!\langle P_1\rangle\!\rangle^S_{C_0}$ both variables still indicate contingent resets. As a consequence, in the next iteration the sets $!x$ and $!y$ again block, whence $\langle\!\langle P_1\rangle\!\rangle^S_{C_1} = C_1$. This fixed-point C_1 is not crisp (not even decided) and constitutes a scheduling deadlock. Observe that the deadlock is detected with the help of the init status not reducing from 2 to 1. In the fixed point semantics of [4] where the init status is missing P_1 would wrongly be classified as SC-constructive. This is a mistake that our extended semantics now fixes. $\qquad\square$

Example 20 Another unschedulable program is the "arbiter" $P_2 := x \; ? \; \varepsilon : !y \parallel y \; ? \; \varepsilon : !x$. It is not constructive because it fails to have any admissible schedules. Every execution order forces a set to happen concurrently after a read and both are not guaranteed to be confluent (depends on the initial memory). As one can verify, our domain-theoretic analysis of P_2 obtains $C_1 := \langle\langle P_1 \rangle\rangle_{C_0}^S = \bot{:}1 \vee \{\langle x^{[\bot,1]}, y^{[\bot,1]} \rangle\}$ and then $\langle\langle P_1 \rangle\rangle_{C_1}^S = C_1$, again choosing $S := \bot$ and $C_0 := [\bot, \top]{:}2$. The fixed point C_1 is undecided and therefore P_1 not (strongly) Berry-constructive (Def. 7). □

The completion codes $cmpl\langle P, C \rangle$ control the analysis of sequential composition. As long as P does not terminate or pause, a sequential successor Q only enters the calculation for $P \; ; \; Q$ to reduce the "may" (upper bound) information on signal statuses, never the "must" (lower bound) information. This is similar to the treatment of conditionals $s \; ? \; P : Q$ in which we block the "must" reaction of P and Q until variable s becomes decided. Until this happens the conditional does not terminate. One can show that termination and crisp reaction environments are closely related. For this we call an environment E *synchronized* when (i) $E(x) = [l, u]{:}0$ implies $l = u$, and (ii) $\bot{:}1 \preceq E(x)$ implies $\forall y. \; \bot{:}1 \preceq E(y)$, for all variables $x \in V$. As we shall see, all our environments will be *synchronized*. Hence the difference between a completed schedule marked by 0 and a contingent schedule marked by one of $\{1, 2\}$ is a feature of the whole environment rather than an individual variable.

Proposition 3 *Let S be synchronized then*

1. $\langle\langle P \rangle\rangle_C^S$ *is synchronized.*
2. *If S is a crisp sequential environment,* i.e., $S(x) \in \mathbb{D}$ *for all $x \in V$, then the response of a terminating or pausing fprog starting from S is crisp, too: If $cmpl\langle P, C \rangle = \{0\}$ or $cmpl\langle P, C \rangle = \{1\}$ then $\langle\langle P \rangle\rangle_C^S(x) \in \mathbb{D}$ for all $x \in V$. The converse also holds,* i.e., *if $\langle\langle P \rangle\rangle_C^S$ is crisp, then $\bot \notin cmpl\langle P, C \rangle$.*

Prop. 3 does not hold for decidedness: Although a program does not terminate it may be possible to constructively prove that its response is decided. *E.g.*, the fprog $s \; ? \; \varepsilon : \varepsilon$ does not complete in the concurrent environment $C(s) = [\bot, \top]{:}2$ but still has the decided response $\langle\langle s \; ? \; \varepsilon : \varepsilon \rangle\rangle_C^\bot(s) = \bot{:}1$, implying that the s remains pristine and environment-controlled.

Proposition 4 *For every reset-free fprog P, the sets $must_k(P, C)$ and $cannot_k(P, C)$ extracted from $cmpl\langle P, C \rangle$ as defined in Sec. 3.3 are identical to the completion semantics of Esterel [12].*

4.3 The Fixed Point Semantics and Constructivity.

While $\langle\langle P \rangle\rangle_C^S$ describes the instantaneous behaviour of P in a compositional fashion, the *constructive response* of P running by itself is obtained by the least fixed point

$$\mu C. \langle\langle P \rangle\rangle_C^S = \bigsqcup_{i \geq 0} C_i, \tag{11}$$

where $C_0 := [\bot, \top]{:}2$ and $C_{i+1} := \langle\langle P \rangle\rangle_{C_i}^S$. Note that the sequential environment S is not updated in the iteration. This reflects the fact that the fixed point approximates the reaction always from the beginning of and concurrent with P. In contrast, the environment S is an initialisation which captures the sequential history of the thread P which remains fixed each

DRAFT 20.4.2015

time the iteration takes place. The fixed point $\mu C . \langle\!\langle P \rangle\!\rangle_C^S$ closes P off against its concurrent environment C. It lets P communicate with itself by treating P as its own *concurrent* context.

For the fixed point to exist the termination function $cmpl\langle P, C \rangle$ and functional $\langle\!\langle P \rangle\!\rangle_C^S$ must be well-behaved. This is the content of the following Prop. 5. We do not use more than elementary fixed point theory over finite domains, here. For a detailed exposition of the technical background the reader is referred to [23].

Proposition 5 *Let P be an arbitrary fprog, S, E environments. Then,*

1. *The functional $cmpl\langle P, E \rangle$ is monotonic with respect to $\sqsubseteq$ in E.*
2. *The functional $\langle\!\langle P \rangle\!\rangle_C^S$ is inflationary in the sequential environment S with respect to $\preceq$.*
3. *The functional $\langle\!\langle P \rangle\!\rangle_E^S$ is monotonic with respect to $\sqsubseteq$ in both the concurrent environment E and the sequential environment S and monotonic for $\preceq$ in S.*

Note that $\langle\!\langle P \rangle\!\rangle_E^S$ is not in general inflationary in S wrt $\sqsubseteq$. For instance, if $[\bot, \top]{:}1 \sqsubseteq C(x)$ $\langle\!\langle !x \rangle\!\rangle_E^\bot(x) = 1$, but $\bot \not\sqsubseteq 1$. Also, because of the reaction to absence $\langle\!\langle P \rangle\!\rangle_E^S$ is not in general monotonic for $\preceq$ in the concurrent environment E.

Monotonicity (Prop. 5) together with finiteness of $I(\mathbb{D}, \mathbb{P})$ implies that the least fixed point $\mu C . \langle\!\langle P \rangle\!\rangle_C^S$ given by (11) is well-defined, for any sequential environment S, if we start from an initial concurrent environment C_0 that is a post-fixed point of $\langle\!\langle P \rangle\!\rangle_{\_}^S$, *i.e.*, if $C_0 \sqsubseteq \langle\!\langle P \rangle\!\rangle_{C_0}^S$. The trivial concurrent environment satisfying this is $C_0 = [\bot, \top]{:}2$ for all $x \in V$. This is the least element wrt $\sqsubseteq$ which codes null-information about the concurrent environment. With this choice of C_0, the sequential environment S is in fact completely arbitrary. We then have $C_i \sqsubseteq C_{i+1}$ and (11) is the stationary limit of this monotonically increasing sequence, which must exist because of the finiteness of $I(\mathbb{D}, \mathbb{P})$.

On the modelling side, the fixed point semantics, discussed so far, is able to accommodate different levels of synchronous constructiveness within it, as we will see next. Different notions of constructiveness are specified by means of certain properties in the fixed point response. First, the connection with Esterel can be made through the two versions of constructiveness introduced in [4]. Then, the denotational companion for the operational notion of IB-causality, namely IB-constructiveness (IBC), is identified and a soundness result is presented. The relationship between the various notions of constructiveness is also discussed.

The class of *strongly Berry-constructive* programs corresponds to the notion of constructiveness in Esterel, yet is able to manage explicit initialisations. This, as expected, can deal with a variable being reset to 0 and then either remaining 0 (signal absence) or being set to 1 (signal presence). Besides, it verifies proper initialisations as part of the constructiveness analysis. It holds the programmer responsible for proper initialisation, not the compiler or the run-time system. Thus, it is important to distinguish whether a variable retains its original value $\bot$ from the initial memory or not.

Definition 7 (Strong Berry-Constructiveness [4] SBC) An fprog P is *strongly Berry-constructive*, or *SBC*, iff for all variables $x \in V$ we have $(\mu C . \langle\!\langle P \rangle\!\rangle_C^\bot)(x) \in \{\bot, 0, 1\}$. $\qquad\square$

It is worth observing that in a SBC program the status $\bot$ for a variable corresponds to a witness for checking initialisations. It indicates that the variable is neither set nor reset by the program. If such a variable is read and thus used in a branching decision the program would be rejected, except for trivial cases. In other words, the resulting status $\bot$ from the fixed point indicates that the variable is indeed never accessed (set, reset or read) by the program.

DRAFT 20.4.2015

Example 21 Fprog $P := x \; ? \; \varepsilon : !y$ is not SBC since variable x (with status $\perp$) is not properly initialised in the code and thus it cannot be decided if variable y is set or not. The fixed point satisfies $\mu C. \langle\langle P \rangle\rangle^{\perp}_C(y) = [\perp, 1]{:}1$. In contrast, for the properly initialised fprog $!x \; ; \; P$ the fixed point will give us $\mu C. \langle\langle !x \; ; \; P \rangle\rangle^{\perp}_C = \{\langle x^1, y^{\perp} \rangle\}$ which is SBC.   □

On the other hand, the actual Esterel's semantics resets all signals to 0 by default, at the beginning of every instant. Thus, in this case, we need to look at *ternary* behaviours, *i.e.*, those which remain inside environments with $E(x) \in \{0, 1, [0, 1]\}$ for all $x \in V$. In order to keep the status of variables in the ternary domain, we could initialise with the reset construct and avoid sequentially forced resets from happening after sets. However, in the semantics $\langle\langle \_ \rangle\rangle^S$ one can emulate initialisation directly by running the fixed point in the sequential environment $S = 0$ instead of $S = \perp$. This give us the class of *Berry-constructive* programs:

Definition 8 (Berry-Constructiveness [4] BC) An fprog P is *Berry-constructive*, or *BC*, iff for all variables $x \in V$ we have $(\mu C. \langle\langle P \rangle\rangle^{0}_C)(x) \in \{0, 1\}$.   □

Example 22 The fprog P from Ex. 21 is BC because now $\mu C. \langle\langle P \rangle\rangle^{0}_C(y) = \{\langle x^0, y^1 \rangle\}$. In Esterel's hardware translation [12], the corresponding Boolean equations are $x = 0$ and $y = \bar{x} + 0$ which stabilise to $x = 0$ and $y = 1$. This depends on the initialisation of x to 0, however. On the other hand, $Q := x \; ? \; \varepsilon : !x$, which emits signal x if x is absent and does not emit it if x is present, is not BC: $\mu C. \langle\langle Q \rangle\rangle^{0}_C(x) = [0, 1]{:}1$. Its hardware translation would be an inverter loop, or combinational equation $x := \bar{x} + 0$, which may exhibit oscillations. Q is not SBC either since $\mu C. \langle\langle Q \rangle\rangle^{\perp}_C(x) = [\perp, 1]{:}1$.   □

Examples 21 and 22 show that SBC is properly more restrictive than BC. The difference between the two forms of Berry-constructiveness is whether we run the simulation with the sequential stimulus $\perp$ or 0, respectively.

The above are not sufficient for capturing SC-read-determinacy as given in Def. 6 which induces an open-world version of constructiveness that takes into consideration external inputs to the program. SC-read-determinacy is constructed from any arbitrary initial memory state that is not controlled by the code such as registered variables. The conditions imposed by this notion can, therefore, be read as follows. For all external inputs, there is always a schedule that does not lead to $\top$ (*reset safe*) and for all read variables and all such schedules either: the final memory value for the variable (temporary variable) is controlled by the program and always the same (0 or 1) *or* by the environment (registered variable) in which case it is not changed at all during the computation (*read safe*). In short, this specifies that every variable used for branching of control is either causally justified or never modified in the code, independently of the initial external input. This leads us to the following definition:

Definition 9 (IB-Constructiveness IBC) Fprog P is *Input Berry-constructive* (*IB-constructive* or *IBC*), iff its fixed point $C_* = \mu C. \langle\langle P \rangle\rangle^{\perp}_C$ is *safe* for P, that is:

- *reset-safe*: $\forall x \in V. \; C_*(x) \preceq 1{:}1$, and
- *read-safe*: $\forall x \in rd(P). \; C_*(x) \in \{\perp, 0, 1\}$.   □

One can show that the class of IBC programs lies between the SBC and the BC programs and that these inclusions are proper.

Example 23 The BC fprog $x \; ? \; \varepsilon : !y$ from Ex. 21, which is not SBC, is also IBC. The fixed point result is $\mu C. \langle\langle x \; ? \; \varepsilon : !y \rangle\rangle^{\perp}_C = \{\langle x^{\perp}, y^{[\perp, 1]{:}1} \rangle\}$. Since x is not properly initialised the status

DRAFT 20.4.2015

of y cannot be decided. This does not matter for IBC as y is not a read variable. Now take the program $Q := x \ ? \ (y \ ? \ \varepsilon : !y) : \varepsilon$. If we initialise with 0, we get $\mu C.\langle\!\langle Q \rangle\!\rangle_C^0 = \{\langle x^0, y^0 \rangle\}$, so Q is BC. Yet, it is not IBC because $\mu C.\langle\!\langle Q \rangle\!\rangle_C^\perp(y) = [\perp, 1]{:}1$ and $y \in rd(Q)$ is a read variable. The problem is that not every initial memory for Q admits of an IB-causal micro-step execution. Specifically, if $\rho_0(x) = 1$, then the sub-program $y \ ? \ \varepsilon : !y$ is scheduled which creates a read-write hazard. It reads the initial (environment-controlled) value of y and then, sequentially afterwards, may change it itself.   □

The result in [4] establishes that for every fprog P, if P is SBC then P is SC-reactive and SC-determinate. Here, we show the following stronger result:

Theorem 1 *For every fprog P, if P is IBC then it is B-reactive and SC-read-determinate.*

Thm. 1 gives a stronger soundness result for the application of the theory compared to Thm. 1 of [4] because it permits us to prove strictly stronger forms of reactiveness and determinacy for a strictly wider class of programs, considering that there are more IBC programs than SBC programs.

4.4 Soundness of the Denotational Fixed Point Semantics.

In the following we summarise the main elements of the proof for Thm. 1 stating that every IB-constructive fprog is B-reactive and SC-read-determinate, and a fortiori also sequentially constructive as introduced in [37, 38]. More details can be found in a technical report [5].

The key element in the soundness proof is to relate the abstract values in $\mathbb{D}$ and $\mathbb{P}$ used in the fixed point analysis with the operational behavior of process executions. These status values are interpreted as abstractions of the write accesses in a finite sequence of micro steps generating what we call the *sequential yield* of each thread. More precisely, a *sequential yield* is a function μ which assigns each possible thread identifier $\iota \in TI$ to a *sequential environment* $\mu(\iota) : V \to \mathbb{D} \times \mathbb{P}$ subject to the condition that $\iota \preceq \iota'$ implies $\mu(\iota') \preceq \mu(\iota)$. The idea is that $\mu(\iota)$ codes the local view of a thread instance ι about the sequential status of the variable values. So, if $\iota \prec \iota'$ then ι' is a (sequential) descendant of thread ι all of whose memory write accesses are visible to the waiting ancestor thread ι. The fact that the view of the ancestor ι is wider, also encompassing other threads (*e.g.*, siblings of ι and their descendants) running concurrently with ι, is captured by the constraint $\mu(\iota') \preceq \mu(\iota)$. The descendant ι' is behind the parent since the parent ι sees all variable accesses of all its active children while ι' only knows about its own.

With the following definition of the sequential yield we are interpreting the actions of a micro-sequence as an incremental update of a sequential state. The pairs in $\mathbb{D} \times \mathbb{P}$ are treated naturally as elements of $I(\mathbb{D}, \mathbb{P})$, *viz.* $(a, r) \in \mathbb{D} \times \mathbb{P}$ is the same as $[a, a]{:}r \in I(\mathbb{D}, \mathbb{P})$ and therefore written $a{:}r$. In this way, all operations on environments over $I(\mathbb{D}, \mathbb{P})$ can be used for the sequential environments, too.

Definition 10 (Sequential Yield) Let R be a finite sequence of micro-steps $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ and C an environment. We define the *sequential yield* $|R|_C : TI \to V \to \mathbb{D} \times \mathbb{P}$ of R by iteration through R, as follows: If $R = \varepsilon$, then $|R|_C(\iota)(x) := \perp = \perp{:}0$ for all $\iota \in TI$ and $x \in V$. Otherwise, suppose $R = R', T_n$ consists of a sequence $R' : (\Sigma_1, \rho_1) \twoheadrightarrow (\Sigma_{n-1}, \rho_{n-1})$ followed by a final action $T_n : (\Sigma_{n-1}, \rho_{n-1}) \to (\Sigma_n, \rho_n)$. Then, $|R|_C$ is computed from $|R'|_C$ by case analysis on the action T_n.

DRAFT 20.4.2015

Generally, the yield does not change for all threads concurrent to $T_n.id$, *i.e.*, for all $\kappa \in TI$ such that $\kappa \not\preceq T_n.id$ and $T_n.id \not\preceq \kappa$ we have $|R|_C(\kappa) := |R'|_C(\kappa)$. Also, if the next control is a non-empty list $T_n.next = Q::Ks'$ and the program $T_n.prog \in \{\varepsilon, !s, ¡s\}$ instantaneously terminates, then the execution of T_n installs the process $\langle inc(\iota), Q, Ks'\rangle$. This incremented thread inherits the sequential state from ι. In this case we put $|R|_C(inc(\iota)) := |R|_C(\iota)$. Otherwise, if $T_n.prog \in \{\varepsilon, !s, ¡s\}$ and $Ks = []$ is empty, then $|R|_C(inc(\iota)) := |R'|_C(\iota)$.

In all other cases, for ancestor and descendant threads κ, the new yield $|R|_C(\kappa)$ is determined according to the following clauses:

1. Executing a sequential composition or the empty statement does not change the yield. Formally, if $T_n.prog \in \{P\,;\,Q, \varepsilon\}$, then $|R|_C(\kappa) := |R'|_C(\kappa)$;

2. Executing a conditional test which is undecided in environment C raises the init status of the thread and its ancestors to 1; otherwise, if the test is decided in C the yield is preserved. Formally, if $T_n = \langle \iota, s\ ?\ P : Q, Ks\rangle$ and for all $b \in \{\bot, 0, 1\}$, $b{:}1 \not\sqsubseteq C(s)$, then we put $|R|_C(\kappa) := |R'|_C(\kappa) \vee \bot{:}1$ for all $\kappa \preceq inc(\iota)$; Otherwise, for $\kappa \not\preceq inc(\iota)$ or if $b{:}1 \sqsubseteq C(s)$ for some $b \in \{\bot, 0, 1\}$, then we define $|R|_C(\kappa) := |R'|_C(\kappa)$;

3. Upon forking a parallel process we copy the sequential status of the parent thread to its two children. Formally, if $T_n = \langle \iota, P\,\|\,Q, Ks\rangle$, then $|R|_C(\iota.l.0) = |R|_C(\iota.r.0) := |R'|_C(\iota)$ and for all $\kappa \neq \iota.r.0$ and $\kappa \neq \iota.l.0$ we have $|R|_C(\kappa) := |R'|_C(\kappa)$;

4. A set $!s$ increases the sequential yield of s in the executing thread and its ancestors and also the speculation status (for all variables) if the set is blocked by C due to a potentially pending reset. Formally, suppose $T_n = \langle \iota, !s, Ks\rangle$. Then, for all $inc(\iota) \prec \kappa$, $|R|_C(\kappa) := |R'|_C(\kappa)$ and for all $\kappa \preceq \iota$,
   - if $[\bot, \top]{:}1 \sqsubseteq C(s)$ then $|R|_C(\kappa)(s) := |R'|_C(\kappa)(s) \vee 1$ and $|R|_C(\kappa)(x) := |R'|_C(\kappa)(x)$ for all variables $x \neq s$. More compactly, $|R|_C(\kappa) := |R'|_C(\kappa) \vee \{\!\{s^1\}\!\}$;
   - if $[\bot, \top]{:}1 \not\sqsubseteq C(s)$ then $|R|_C(\kappa)(s) := |R'|_C(\kappa)(s) \vee 1{:}1$ and $|R|_C(\kappa)(x) := |R'|_C(\kappa)(x) \vee \bot{:}1$ for $x \neq s$. More compactly, $|R|_C(\kappa) := |R'|_C(\kappa) \vee \{\!\{s^1\}\!\} \vee \bot{:}1$.

5. A reset $¡s$ increases the sequential yield for s to 0 if the status is still smaller than 0, or to $\top$ if the status of s in the thread is already at or above 1. At the same time, if the thread has entered the speculative mode, then the reset $¡s$ raises the speculation status to 2. Formally, if $T_n = \langle \iota, ¡s, Ks\rangle$, then $|R|_C(\kappa)(x) := |R'|_C(\kappa)(x)$ for all $inc(\iota) \prec \kappa$ or $x \neq s$; Otherwise, for all $\kappa \preceq \iota$ we put
   - $|R|_C(\kappa)(s) := |R'|_C(\kappa)(s) \vee \top$ if $1 \preceq |R'|_C(\iota)(s) \preceq \top$;
   - $|R|_C(\kappa)(s) := |R'|_C(\kappa)(s) \vee \top{:}2$ if $1{:}1 \preceq |R'|_C(\iota)(s)$;
   - $|R|_C(\kappa)(s) := |R'|_C(\kappa)(s) \vee 0$ if $|R'|_C(\iota)(s) \preceq 0$;
   - $|R|_C(\kappa)(s) := |R'|_C(\kappa)(s) \vee 0{:}2$ if $\bot{:}1 \preceq |R'|_C(\iota)(s) \preceq 0{:}2$; $\square$

A special case is the totally pristine sequential yield $\mu_\bot$ with $\mu_\bot(\iota) = \bot$ for all $\iota \in TI$. This is the yield $|\varepsilon|_C$ of the empty micro sequence. Moreover, the yield operation is monotonic, *i.e.*, if R is a prefix of R' then $|R|_C(\iota) \preceq |R'|_C(\iota)$. Observe that if R does not have any write accesses to a variable x then the value status of x in the sequential yield remains $\bot$, the init status may raise to 1 but not to 2, *i.e.*, $|R|_C(\iota)(x) \preceq \bot{:}1$. For SC-admissible micro-sequences the yield gives a sound approximation of the final memory state:

Lemma 1 *Let $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ be an SC-admissible micro-step sequence and C an environment. Then, $|R|_C$ is consistent for the final memory ρ_n in the following sense:*

(i) If $|R|_C(Root.id)(x) \preceq \bot{:}2$ then $\rho_0(x) = \rho_n(x)$;
(ii) If $|R|_C(Root.id)(x) = b{:}r$ with $b \in \{0, 1\} \subset \mathbb{D}$ then $\rho_n(x) = b$;

*(iii) If $|R|_C(\mathit{Root}.id)(x) \succeq 1$ then there exists a micro step $1 \leq i \leq n$ such that $T_i.prog = !x$
and for all $T \in \Sigma_n$ with $T_i.id \preceq T.id$ we have $1 \preceq |R|_C(T.id)(x)$.* ☐

*(iv) Whenever in R a thread ι reads a variable x, the value status of x in any other concurrent
thread remains constant from this point onwards. In other words, no thread changes
the value status of x after it is has been read by another thread concurrent to it. For-
mally, suppose $R(i) = \langle \iota, x ? P : Q, Ks \rangle$ and for some $i \leq j \leq n$ and $\iota' \in TI$ we have
$|R@i|_C(\iota')(x) \wedge \top \neq |R@j|_C(\iota')(x) \wedge \top$. Then, $\iota' \preceq \iota$ or $\iota \preceq \iota'$.*

The strategy for proving Thm. 1, stating that every IB-constructive program is B-reactive
and SC-determinate, is to show that the fixed point $\mu C.\langle\!\langle P \rangle\!\rangle_C^{\perp} \in I(\mathbb{D}, \mathbb{P})$ computes sound
information about the sequential yield of every SC-admissible micro-step sequence R of P.
More specifically, we show that $\mu C.\langle\!\langle P \rangle\!\rangle_C^{\perp}$ is an abstract predictor for the SC-admissible
behavior of P in the sense that (i) the yield of every SC-admissible micro-sequence lies
within the window specified by $\mu C.\langle\!\langle P \rangle\!\rangle_C^{\perp}$ and (ii) there exists a B-admissible instant. This
is done by induction on the structure of P. However, since the fixed point of a composite
expression cannot be obtained from the fixed points of its sub-expressions, induction on P
for the full fixed point $\mu C.\langle\!\langle P \rangle\!\rangle_C^{\perp}$ does not work. Instead, we need to break up the fixed point
and do an outer induction along the iteration that obtains the fixed point in the limit. The
idea is to extract the logical meaning of a single iteration step $C_{i+1} = \langle\!\langle P \rangle\!\rangle_{C_i}^{S}$ as a conditional
specification of the SC-admissible behavior of P assuming a sequential environment S and
concurrent environment C_i. This can then be proven by induction on P.

The main observation is that a single application of the response functional $\langle\!\langle P \rangle\!\rangle_{C_i}^{S}$ covers
the behavior of an *initial slice* of any micro-sequence R generated from P, consisting of
an atomic "read;update" burst of P. This burst consists of all those statements of R that
can be executed solely based on the concurrent environment C_i to decide which branch to
take in a conditional and whether a set can go ahead or is blocked because of a pending
reset. At such a point, or if a conditional is undecided the slice stops. We have reached
the *stopping index* of the slice in R. In the slice, control branching is decided entirely in
terms of the variables whose values are decided in C_i and not on variables whose value may
be changing as a result of executing P. In particular, the execution in R covered by a slice
decided from C_i does not involve any communication between concurrent processes inside P.
Since effect of executing the slice is described by the response environment $C_{i+1} = \langle\!\langle P \rangle\!\rangle_{C_i}^{S}$,
the communication between threads is then handled by feeding back the result C_{i+1} as the
new concurrent environment in the next iteration $C_{i+2} = \langle\!\langle P \rangle\!\rangle_{C_{i+1}}^{S}$ of the response functional.

Definition 11 (Stopping Index) Let $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ be a micro-sequence and C an
environment. A process $T_i \in \Sigma_i$ for $0 \leq i < n$ is called *C-blocked* if T_i is active in Σ_i and either

- $T_i.prog$ is a branching $x ? Q : R$ and $\perp\!:\!1 \not\sqsubseteq C(x)$, $0\!:\!1 \not\sqsubseteq C(x)$ and $1\!:\!1 \not\sqsubseteq C(x)$, or
- $T_i.prog$ is a set $!x$ and $[\perp, \top]\!:\!1 \not\sqsubseteq C(x)$.

In all other cases, T_i is called *C-enabled*. Let $\langle \iota_P, P, Ks \rangle \in \Sigma_i$ be active in Σ_i. The *C-stopping
index* of P in R is the earliest step index $i \leq t \leq n$ such that one of the following holds:

- P pauses
- P has terminated instantaneously and handed over to Q in the next control $Ks = Q :: Ks'$
- all remaining active descendants $\langle \iota', P', Ks \rangle \in \Sigma_t$ with $\iota \preceq \iota'$ are *C-blocked*. ☐

Note that the *C*-stopping index of a program in a micro-sequence R may not exist if R is
not long enough so that R still has an active process from P in its last configuration and this

DRAFT 20.4.2015

process is not C-blocked. Also, if C is safe, *i.e.*, reset-safe and read-safe, then at its C-stop in R the program P must either pause or terminate instantaneously.

The C-stopping index of P in a micro-sequence R marks the maximal horizon until which we will follow the execution of P when using C as a predictor environment. However, this does not mean that the memory is actually consistent with the prediction C. The notion of *C-consistency* in the following Def. 12 fills this gap. A micro-sequence R is C-consistent if every read of every variable x for which C predicts a decided value, i.e., for which $b{:}1 \sqsubseteq C(x)$, for $b \in \{\bot, 0, 1\}$, the actual value of x in the memory coincides with this prediction b. So, if R happens to be C-consistent then the prediction will be sound up to the C-stopping index.

Definition 12 (*C-Consistency*) Let $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ be a micro sequence and C an environment. For any $0 \le i < n$, abbreviate by $\rho_i(x) \doteq b$ the condition that $\rho_i(x) = b$ if $b \in \{0, 1\}$ and $\rho_i(x) = \rho_0(x)$ if $b = \bot$. We say a read action $R(i).prog = x \ ? \ P : Q$ with $0 < i \le n$ is C-consistent in R if $b{:}1 \sqsubseteq C(x)$ for $b \in \{\bot, 0, 1\}$ implies $\rho_{i-1}(x) \doteq b$. R is called *C-consistent for a thread ι* if all read actions performed by all descendants of ι in R are C-consistent. □

The following two Props. 6 and 7 are instrumental. They relate the environment obtained from a single computation of the denotational function $\langle\langle P \rangle\rangle_C^S$ as upper and lower bounds of the sequential yield of a C-consistent execution of P. Also, the completion information $cmpl\langle P, C \rangle$ is related to the completion behaviour of P.

Proposition 6 (**Soundness of the Lower/Must Prediction**)
Let $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ be a micro sequence with an active process $\langle \iota_P, P, Ks \rangle$ in Σ_s, $0 < s \le n$, and C an environment such that R is C-consistent for ι_P and n the C-stopping index of P in R.

(i) If $cmpl\langle P, C \rangle = \{0\}$ then P instantaneously terminates at step n by executing an action of the form ε, $¡s$, $!s$; If $cmpl\langle P, C \rangle = \{1\}$ then P pauses at step n where the last of its descendants has reached the action π.

(ii) Suppose $S \preceq |R@s|_C(\iota_P) \preceq \top$ for some sequential environment S. Then, for each variable $x \in V$ there exists an index $s \le i \le n$ and a descendant thread $\iota \succeq \iota_P$ such that $\langle\langle P \rangle\rangle_C^S(x) \wedge \top \preceq |R@i|_C(\iota)(x) \vee [\bot, \top] \preceq \top$. Moreover, if $\bot \notin cmpl\langle P, C \rangle$ then $i = n$ and $\iota = \iota_P$.

(iii) If $S \preceq low \ |R@s|_C(\iota_P)$ then $\langle\langle P \rangle\rangle_C^S \preceq low \ |R@n|_C(\iota_P)$.

Proof Both parts (i) and (ii) are shown by induction on P. Regarding part (iii) we observe that under the assumptions of (ii) it follows that $low(\langle\langle P \rangle\rangle_C^S(x) \wedge \top) \preceq low(|R@i|_C(\iota)(x) \vee [\bot, \top]) = low \ |R@i|_C(\iota)(x) \preceq low \ |R@n|_C(\iota_P)(x)$ and therefore $\langle\langle P \rangle\rangle_C^S \preceq (low \ \langle\langle P \rangle\rangle_C^S) \wedge \top{:}2 = low \ \langle\langle P \rangle\rangle_C^S \wedge low(\top) = low(\langle\langle P \rangle\rangle_C^S \wedge \top) \preceq low \ |R@n|_C(\iota_P)$. □

Proposition 7 (**Soundness of Upper/Cannot Prediction**) *Let $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$ be a finite micro sequence with an active process $\langle \iota_P, P, Ks \rangle \in \Sigma_s$, $0 \le s \le n$, and C an environment such that R is C-consistent for ι_P. Suppose that all actions executed between s and n are from processes concurrent to ι_P or from descendants of P. In particular, there are no actions from the continuation list Ks. Then,*

(i) If $0 \notin cmpl\langle P, C \rangle$ then at least one descendant of P is active or pausing in Σ_n and if $1 \notin cmpl\langle P, C \rangle$ then not all descendants of P in Σ_n, if there are any, are pausing.

(ii) $upp \ |R@s|_C(\iota_P) \preceq S$ implies $upp \ |R@n|_C(\iota_P) \preceq \langle\langle P \rangle\rangle_C^S$.

DRAFT 20.4.2015

Proof We proceed by induction on the structure of the program and the length of the continuation list *Ks*. Note that the statements (i) and (ii) of the Prop. 7 hold trivially, if program P does not perform any actions between s and n. In this case, $upp |R@n|_C(\iota_P) = upp |R@s|_C(\iota_P) \preceq S \preceq \langle\langle P \rangle\rangle_C^S$ by the inflationary nature of the prediction (Prop. 5(2)). □

Proposition 8 *For every fprog P, if P is IBC then P is B-reactive.*

Proof Let P be an IBC program, *i.e.*, $C_* = \mu C.\langle\langle P \rangle\rangle_C^\perp$ is safe: for all $x \in V$, $C_*(x) \preceq 1{:}1$ and for all $x \in rd(P)$, $C_*(x) \in \{\perp, 0, 1\}$. Further, let (Σ_0, ρ_0) be an initial configuration in which program P appears as the sole active process in the pool, *i.e.*, $\Sigma_0 = \{\text{Root}\}$, where $\text{Root} = \langle \iota_P, P, [] \rangle$ and $\iota_P = \text{Root}.id = 0$.

We show that there must exist at least one B-admissible execution for P from any memory state. This proof demonstrates how the fixed point iteration can be used as a predictive B-admissible scheduler. We are going to build iteratively a contiguous sequence of B-admissible micro-sequences

$$(\Sigma_{n_0}, \rho_{n_0}) \overset{R_0}{\twoheadrightarrow} (\Sigma_{n_1}, \rho_{n_1}) \overset{R_1}{\twoheadrightarrow} (\Sigma_{n_2}, \rho_{n_2}) \overset{R_2}{\twoheadrightarrow} (\Sigma_{n_3}, \rho_{n_3}) \cdots \overset{R_{i-1}}{\twoheadrightarrow} (\Sigma_{n_i}, \rho_{n_i})$$

with $n_0 = 0$ and $n_{i-1} \leq n_i$, where in each scheduling round R_{i-1} we are pushing the execution as far as possible while staying C_{i-1}-enabled, where C_{i-1} is the sequence of concurrent environments generated by the fixed point iteration. Since the initial pool is $\Sigma_0 = \{\langle \iota_P, P, [] \rangle\}$, all threads in any of the process pools Σ_k reached during $R_0, R_1, \ldots, R_{i-1}$ are descendants of P. By construction, each descendant thread remaining active in round R_{i-1} is C_{i-1}-stopped in the final configuration Σ_{n_i}. For the fixed point C_*, which is safe, this means that in the corresponding end configuration $(\Sigma_{n_*}, \rho_{n_*})$ all threads descending from ι_P are either instantaneously terminated or pausing. Recall that in the final configuration $(\Sigma_{n_*}, \rho_{n_*})$ no set $!x$ can be C_*-blocked since C_* is reset-safe and no read x ? P' : Q' can be blocked because C_* is read-safe. Hence, at the fixed point, we have constructed a maximal micro sequence and thus reached the end of the macro step (instant). Here are the key invariants of the construction:

(I1) The yield of each partial schedule is in the range predicted by the fixed point approximation, *i.e.*, $C_i \sqsubseteq |R_0, R_1, \ldots, R_{i-1}|_{C_{i-1}}(\iota_P)$.
(I2) Each partial schedule $R_0, R_1, \ldots, R_{i-1}$ is B-admissible.
(I3) For every free schedule R' from $(\Sigma_{n_i}, \rho_{n_i})$, the extended schedule $R_0, R_1, \ldots, R_{i-1}, R'$ is C_i-consistent. Further, if $C_i(x) \preceq \top{:}1$ then R' does not contain a reset $¡x$.

The invariants (I1)–(I3) tell us that the full sequence $R = R_0, R_1, \ldots, R_*$ up to the fixed point, obtained as the result of our scheduling strategy, is C_*-consistent and that every conditional test performed in the full schedule R reads exactly the memory value predicted by the read-safe fixed point environment.

Base Case. Observe that the empty schedule ε is trivially B-admissible and its sequential yield $|\varepsilon|(\iota_P) = \perp$ lies in the environment $C_0 = [\perp, \top]{:}2$, *i.e.*, $C_0 \sqsubseteq \perp$. So, both (I1) and (I2) hold for the empty sequences. Regarding (I3) note that every free schedule R' starting in the configuration $(\Sigma_{n_0}, \rho_{n_0})$ is trivially C_0-consistent since no variable is decided in C_0. Since $C_0 \npreceq \top{:}1$ the schedule R' is not constrained regarding resets.

Step Case. By way of induction hypothesis (I1)–(I3), suppose we have constructed a B-admissible schedule $R_0, R_1, \ldots, R_{i-1}$ (I2) such that for every $j \leq i$ (using course-of-values induction) the yield of $R_0, R_1, \ldots, R_{j-1}$ with respect to C_{j-1} lies in the range predicted by

C_j (I1) and for every free schedule R' from $(\Sigma_{n_j}, \rho_{n_j})$ the extension $R_0, R_1, \ldots, R_{j-1}, R'$ is C_j-consistent (I3). Moreover, from (I3) we may assume that if $C_j(x) \preceq \top{:}1$ then R' is reset-free for x. From $(\Sigma_{n_i}, \rho_{n_i})$ we now continue to schedule all and only those actions that are active and C_i-enabled. We do this until ι_P stops under C_i, *i.e.*, until it completes or all remaining active threads are C_i-blocked. This procedure builds a round schedule R_i and leads to a configuration $(\Sigma_{n_{i+1}}, \rho_{n_{i+1}})$. Then, n_{i+1} is the C_i-stopping index of P in $R_0, R_1, \ldots, R_{i-1}, R_i$. If it happens that there is no active process in Σ_{n_i} which is C_i-enabled, then $\Sigma_{n_{i+1}} = \Sigma_{n_i}$ and $\rho_{n_{i+1}} = \rho_{n_i}$. In this case, we just move on to the next iteration round C_{i+1} of the fixed point without progressing the schedule. $\qquad\square$

Proposition 9 *For every fprog P, if P is IBC then P is SC-read-determinate.*

Proof Again let (Σ_0, ρ_0) be an initial configuration with $\Sigma_0 = \{\langle \iota_P, P, [\,] \rangle\}$ and $\iota_P = \text{Root}.id = 0$. Fix an SC-admissible instant $R : (\Sigma_0, \rho_0) \twoheadrightarrow (\Sigma_n, \rho_n)$, where $n = len(R)$. Observe that all processes in every pool Σ_i are descendants of ι_P. We cover the micro-sequence R incrementally with the results from the fixed point iteration, showing that R can only ever execute variable read accesses within the corridor predicted by the fixed point responses C_i, where $C_0 = [\bot, \top]{:}2$ and $C_{i+1} = \langle\!\langle P \rangle\!\rangle_{C_i}^{\bot}$. This exploits the soundness of lower and upper predictions, Props. 6 and 7.

More precisely, let i_{k+1} be the C_k-stopping index of P in R for $k \geq 0$. These must exist because R is an instant and thus a maximal micro-sequence. One shows by induction on k that R is C_k-consistent and

$$C_{k+1} \sqsubseteq |R@j|_{C_k}(\iota_P) \text{ for all } i_{k+1} \leq j \leq n$$

which says that the environment C_{k+1} is a sound approximation of the yield from i_{k+1} onwards. At the fixed point $C_* = \mu C. \langle\!\langle P \rangle\!\rangle_C^{\bot} \in \{\bot, 0, 1\}$, this implies that that R is C_*-consistent for ι_P and $C_* \sqsubseteq |R@n|_{C_*}(\iota_P)$. In view of Lem. 1 this shows that all SC-admissible instants R of P have the same deterministic final memory value and this memory value is the one computed by the IBC fixed point analysis. $\qquad\square$

5 Related Work

The usefulness of synchronisation primitives is well-established in main-stream concurrent programming. E.g., C++ and Java are based on a multi-threaded shared-memory execution model which provides synchronisation of methods to isolate threads and to ensure safety properties such as mutual exclusion. The clock synchronisation (pause) and associated scheduling constraints of our SMoCC approach may also be seen as a synchronisation pattern to ensure memory safety. It provides global snapshot barriers and pruning of thread interleaving with the aim of ensuring reactiveness and memory determinacy. The programmer must decide which synchronisation model is the right one for a given application context. In reactive and embedded systems the SMoCC has turned out to be a natural choice.

In terms of programming languages, the work presented here is at the interface between synchronous concurrent languages and C-like sequential languages, and is strongly influenced by both worlds. Edwards [25] and Potop-Butucaru et al. [71] provide good overviews of compilation challenges and approaches for concurrent languages, including synchronous languages. They discuss efficient mappings from Esterel to C, thus their work is related to ours in the sense that we present a means to express Esterel-style signal behavior and deterministic concurrency directly with variables in a C-like language. However, a key difference is that

DRAFT 20.4.2015

we do not "compile away" the concurrency as part of our signal-to-variable mapping, but fully preserve the original, concurrent semantics with shared variables.

Introducing the constructive causality classes SBC, IBC, BC we redress the synchronous model of computation, well-known in the embedded systems domain, for main-stream programming. There are already many proposals that extend C or Java with synchronous concurrency constructs. Reactive C [16] is an extension of C that employs the concepts of ticks and preemptions, but does not provide concurrency. FairThreads [17] are an extension introducing concurrency via native threads. PRET-C [6] and Synchronous C, a.k.a. SyncCharts in C [36], provide macros for defining synchronous concurrent threads. SC also permits dynamic thread scheduling, and thus would be a suitable implementation target for the pSCL language discussed here. SHIM [81], another C-like language, provides concurrent Kahn process networks with CCS-like rendezvous communication [43] and exception handling. SHIM has also been inspired by synchronous languages, but it does not use the synchronous programming model, instead relying on communication channels for synchronisation. None of these language proposals claims and proves to embed the concept of Esterel-style constructiveness into shared variables as we do here. As far as these language proposals include signals, they come as "closed packages" that do not, for example, allow to separate initialisations from updates.

As traditional sequential, single-core execution platforms are being replaced by multi-core/processing architectures, determinism is no longer a trade secret of synchronous programming but has become an important issue in shared memory concurrent programming. Powerful techniques have recently been developed to verify program determinism statically. For Java with structured parallelism, the tool DICE by Vechev *et al.* [83] performs static analysis to check that concurrent tasks do not interfere on shared array accesses. Leung *et al.* [54] present a *test amplification* technique based on a combination of instrumented test execution and static data-flow analysis to verify that the memory accesses of cyclic, barrier-synchronised, CUDA C++ threads do not overlap during a clock cycle (barrier interval). For polyhedral X10 programs with finish/async parallelism and affine loops over array-based data structures, Yuki *et al.* [85] describe an exact algorithm for static race detection that ensures deterministic execution.

These recently published analyses [83, 54, 85] are targeted at data-intensive, array/pointer/based code building on powerful arithmetical models and decision procedures for memory separation. Yet, they address determinism in more limited models of communication. SMoCC constructiveness concerns the determinism and reactivity of "control-parallel" rather than "data-parallel" synchronous programs and permits instantaneous communication between threads during a single tick. The challenge is to deal with feedbacks and reaction to absence, as in circuit design, which is difficult. The causality of the SMoCC memory accesses cannot necessarily be captured in terms of regular affine arithmetics as done in the polyhedral model of [83, 85] or reduced to a "small core of configuration inputs" as in [54]. Further, analyses such as [83, 54, 85] verify race-freedom for maximally strong data conflicts: Within the barrier no write must ever compete with a concurrent read or another conflicting write. Soundness of the analysis is straightforward under such full isolation. Full thread isolation is fine for Moore-style communication but does not hold in the SMoCCs whose hallmark is the Mealy model. Threads do in fact share variables during a clock phase and multi-emissions are permitted. Analysing SMoCC determinism, therefore, is tricky and arguing soundness of the constructivity analysis in the SMoCCs (our Thm. 1) is non-trivial. This is particularly true if reaction to absence is permitted, as in our work, which introduces non-monotonic system behaviour on which the standard (naive) fixed-point techniques fail.

DRAFT 20.4.2015

For functional programming languages, traditionally abstracting from the impurity of low-level scheduling, determinism on concurrent platforms also has become an issue. For instance, Kuper *et al.* [51] extend the IVar/LVar approach in Haskell to provide deterministic shared data-structures permitting multiple concurrent reads and writes. This extension, dubbed *LVish*, adds asynchronous event handlers and explicit value freezing to implement negative data queries. Since the negative information is transient, run-time exceptions are possible due to the race between freezing and writing. However, all error-free executions produce the same result which is called *quasi-determinism*. Because of the instantaneous communication and the negative information carried by the value status of shared data, the quasi-deterministic model of [51] is similar in spirit to our approach. However, there are at least two differences: First, our programming model deals with first-order imperative programs on Boolean data, while [51] considers higher-order λ-functions on more general "*atomistic*" data structures. Second, our $\langle\!\langle \_ \rangle\!\rangle$ constructivity includes *reactivity*, which is a liveness property, whereas [51] only address the safety property of non-interference. Our *two-dimensional* lattice $I(\mathbb{D})$ seems richer than the lifted domain *Freeze*$(\mathbb{D})$ of [51] which only distinguishes between the "unfrozen" statuses $[\bot, \top]$, $[0, \top]$, $[1, \top]$, $[\top, \top]$ (lower information) and the "frozen" statuses $[\bot, \bot]$, $[0, 0]$, $[1, 1]$ (crisp information). There do not seem to be genuine upper bound approximations expressible in *Freeze*$(\mathbb{D})$. It will be interesting to study the exact relationship between the two models on a common language fragment.

There is also a large body of related work investigating different notions of constructiveness. Causal Esterel programs on pure signals satisfy a strong scheduling invariant: they can be translated into constructive circuits which are *delay-insensitive* [18] under the non-inertial delay model, which can be fully decided using ternary Kleene algebra [65]. This makes Malik's work [59] on causality analysis of cyclic circuits applicable to constructiveness analysis of (instantaneous) Esterel program. This has been extended by Shiple *et al.* [78] to state-based systems, as induced by Esterel's pause operator, thus handling non-instantaneous programs as well. The algebraic transformations proposed by Schneider *et al.* [77] increase the class of programs considered constructive by permitting different levels of partial evaluation. Pnueli and Shalev's non-deterministic model of Statecharts [70] has been studied using an axiomatic semantics in intuitionistic logic [55], which subsequently has been extended to Esterel [57]. In [3] a game-theoretic approach is used to define a hierarchy of levels constructiveness using maximal post fixed points. However, none of these approaches considers imperative programming, separates initialisations and updates, or permits sequential writes within a tick as we do here.

Recently, Mandel et.al.'s clock domains [60] and Gemünde et.al.'s clock refinement [31] provide sequences of micro-level computations within an outer clock tick. This also increases sequential expressiveness albeit in an upside-down fashion compared to our approach. Our work on SC aims to reconstruct the scope of a synchronous instant on top of the primitive notion of sequential composition. Different classes of constructiveness are distinguished by how generous they are in bundling sequences of variable accesses from concurrent threads within a single clock tick. In the clock domains and clock refinement approach, clocks are the only sequencing mechanism, so micro-level sequencing is implemented in terms of lower-level clocks. It should be possible to combine our approach with that of [60, 31] by considering the sequential composition operator as a local micro-level clock nested inside an outer, and global, macro-level clock. This might generate a useful theory of causal clock abstractions.

Our work focuses on imperative, *i.e.*, control-dominated synchronous programs rather than data flow semantics. Recently, Talpin et.al. within the "Polycore" project have started

DRAFT 20.4.2015

important work on semantically integrating the control-flow synchronous language Quartz with the data flow language Signal. In [79] they present the first micro-step (or "small-step") operational fixed point semantics that is capable of executing both Signal code and the guarded actions of Quartz. The operational semantics models the behaviour of each variable in a 6-valued lattice domain $\mathscr{D}$ coding the signal statuses unknown (?), absent ($\bot$), present ($\top$), present-and-false (0), present-and-true (1) as well as inconsistency ($\sharp$). Based on the operational execution model they define the notion of constructive programs and prove a soundness theorem stating that each constructive program is deterministic.

One difference compared to our work is that the domain $I(\mathbb{D}, \mathbb{P})$ supports reaction to absence[6] which is a hall-mark of Esterel-style SMoCCs and motivates its richer interval structure. On the other hand, the polychronous language of [79] is richer than pSCL in that is has preemption and boolean data values which we do not consider. Finally, note that our definition of constructive programs (*e.g.*, Def. 9) is based on a genuine *denotational* semantics $\langle\langle \_ \rangle\rangle$, not an operational one as in [79]. E.g., it follows from our results that if two program P and Q generate the same response function $\langle\langle P \rangle\rangle_C^S = \langle\langle Q \rangle\rangle_C^S$ in S and C, then P and Q are behaviourally equivalent in all program contexts. Also, our operational semantics (*e.g.*. Sec. 2) uses free multi-threaded scheduling in a memory that is ignorant of the signal statuses. In particular, it does not perform any implicit enabledness, synchronisation or deadlock checks like the operational semantics of [79] does, in which execution maintains scheduling information through variable values in $\mathscr{D}$. Hence, our Soundness Theorem 1 which guarantees B-reactiveness and SC-determinacy makes a stronger soundness statement which is considerably more difficult to prove. Our result can be applied directly to standard imperative C/Java code which is not normally executed under a $\mathscr{D}$-instrumented run-time scheduler. Yet, given the limited language constructs of pSCL compared to [79], it would be very interesting to combine both approaches.

An acknowledged strength of synchronous languages is their formal foundation [9], which facilitates formal verification, timing analyses, and inclusion results of the type presented in this work. This formal foundation has been developed in several ways in the past; *e.g.*, Berry [12] presents several Plotkin-style structural operational semantics [69], as well as a definition in terms of circuits for Esterel. Our functional/algebraic approach based on $I(\mathbb{D}, \mathbb{P})$ generalizes the "must-cannot" analysis for constructiveness [12] and the ternary analysis for synchronous control flow [75] and circuits [59, 78]. The extension lies in the ability to deal with non-initialization ($\bot$) and re-initialization ($\top$) in sequential control flow, which the analyses [12, 75, 59, 78, 79] cannot handle. Due to the two-sided nature of intervals our semantics permits the modeling of instantaneous reaction to absence, a definitive feature of Esterel-style synchrony for control-flow languages. In contrast, the *balance equations* (see, e.g., [53]) or the *clock calculus* (see, e.g., [19, 68, 30, 79]) of synchronous reactive data flow do not handle reaction to absence. These analyses are concerned with inter-tick causality (i.e., in which ticks a signal is present) rather than intra-tick causality (i.e., presence or absence in a given tick) which we focus on here. Reflected into $I(\mathbb{D})$, Lustre clocks collapse the signal status (within a tick) to either $\bot$ (value not initialized or computed) or $[0, \top]$ (value computed). However, since each program abstracts to a continuous function on $I(\mathbb{D}, \mathbb{P})$-valued environments our model fits naturally into the Kahn-style fixed-points semantics and scheduling analysis for synchronous block diagrams [27, 72].

---

[6]   The oversampling feature of Signal may be seen as an implicit form of reaction to absence in the asynchronous data-flow part of [79]. Synchronous reaction to absence would map $\bot$ to $\top$ and $\top$ to $\bot$ in the domain $\mathscr{D}$ which does not seem to be expressible by the control-flow operators considered in [79].

DRAFT 20.4.2015

References

1. Abramsky, S., Jagadeesan, R., Malacaria, P.: Full abstraction for PCF. Information and Computation **163**(2), 409–470 (2000)
2. Aceto, L., Ingolfsdottir, A.: CPO models for compact GSOS languages. Information and Computation **129**(2), 107 – 141 (1996)
3. Aguado, J., Mendler, M.: Constructive semantics for instantaneous reactions. Theoretical Computer Science **241**, 931–961 (2011)
4. Aguado, J., Mendler, M., von Hanxleden, R., Fuhrmann, I.: Grounding synchronous deterministic concurrency in sequential programming. In: Proceedings of the 23rd European Symposium on Programming (ESOP'14), LNCS 8410, pp. 229–248. Springer, Grenoble, France (2014)
5. Aguado, J., Mendler, M., von Hanxleden, R., Fuhrmann, I.: Denotational fixed-point semantics for constructive scheduling of synchronous concurrency. Technical Report 96, University of Bamberg, Faculty of Information Systems and Applied Computer Sciences (2015). ISSN 0937-3349
6. Andalam, S., Roop, P.S., Girault, A.: Deterministic, predictable and light-weight multithreading using PRET-C. In: Proceedings of the Conference on Design, Automation and Test in Europe (DATE'10), pp. 1653–1656. Dresden, Germany (2010)
7. Baudart, G., Mandel, L., Pouzet, M.: Programming mixed music in ReactiveML. In: Proceedings of the First ACM SIGPLAN Workshop on Functional Art, Music, Modeling & Design, FARM '13, pp. 11–22. ACM, New York, NY, USA (2013)
8. Benveniste, A., Caillaud, B., Guernic, P.L.: Compositionality in dataflow synchronous languages: Specification and distributed code generation 1,2,3. Information and Computation **163**(1), 125–171 (2000)
9. Benveniste, A., Caspi, P., Edwards, S.A., Halbwachs, N., Guernic, P.L., de Simone, R.: The Synchronous Languages Twelve Years Later. In: Proc. IEEE, Special Issue on Embedded Systems, vol. 91, pp. 64–83. IEEE, Piscataway, NJ, USA (2003)
10. Bergstra, J., Ponse, A., Smolka, S. (eds.): Handbook of Process Algebra. Elsevier (2001)
11. Berry, G.: The foundations of Esterel. In: G. Plotkin, C. Stirling, M. Tofte (eds.) Proof, Language, and Interaction: Essays in Honour of Robin Milner, pp. 425–454. MIT Press, Cambridge, MA, USA (2000)
12. Berry, G.: The Constructive Semantics of Pure Esterel. Draft Book, Version 3.0, Centre de Mathématiques Appliqées, Ecole des Mines de Paris and INRIA, 2004 route des Lucioles, 06902 Sophia-Antipolis CDX, France (2002)
13. Berry, G., Curien, P.L., Lévy, J.J.: Full abstraction for sequential languages: the state of the art. In: M. Nivat, J.C. Reynolds (eds.) Algebraic Semantics, pp. 89–132. Cambridge University Press (1985)
14. Berry, G., Gonthier, G.: The Esterel synchronous programming language: Design, semantics, implementation. Science of Computer Programming **19**(2), 87–152 (1992)
15. Berry, G., Nicolas, C., Serrano, M.: Hiphop: A synchronous reactive extension for Hop. In: Proceedings of the 1st ACM SIGPLAN International Workshop on Programming Language and Systems Technologies for Internet Clients, PLASTIC '11, pp. 49–56. ACM, New York, NY, USA (2011)
16. Boussinot, F.: Reactive C: An extension of C to program reactive systems. Software Practice and Experience **21**(4), 401–428 (1991)
17. Boussinot, F.: Fairthreads: mixing cooperative and preemptive threads in C. Concurrency and Computation: Practice and Experience **18**(5), 445–469 (2006)
18. Brzozowski, J.A., Seger, C.J.H.: Asynchronous Circuits. Springer-Verlag, New York (1995)
19. Caspi, P., Pilaud, D., Halbwachs, N., Plaice, J.A.: Lustre: a declarative language for programming synchronous systems. In: Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL'87), pp. 178–188. ACM, Munich, Germany (1987)
20. Caspi, P., Pouzet, M.: A co-iterative characterization of synchronous stream functions. Electronic Notes in Theoretical Computer Science **11**(0), 1–21 (1998). CMCS'98, First Workshop on Coalgebraic Methods in Computer Science
21. Cleaveland, R., Lüttgen, G., Mendler, M.: An algebraic theory of multiple clocks. In: CONCUR '97, *LNCS*, vol. 1243, pp. 166–180. Springer (1997)
22. Cohen, A., Duranton, M., Eisenbeis, C., Pagetti, C., Plateau, F., Pouzet, M.: N-synchronous Kahn networks: A relaxed model of synchrony for real-time systems. In: Conference Record of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '06, pp. 180–193. ACM, New York, NY, USA (2006)

23. Davey, B.A., Priestley, H.A.: Introduction to Lattices and Order. Cambridge University Press (2002)
24. Duffin, R.: Topology of series-parallel networks. Journal of Mathematical Analysis and Applications **10**(2), 303 – 318 (1965)
25. Edwards, S.A.: Tutorial: Compiling concurrent languages for sequential processors. ACM Transactions on Design Automation of Electronic Systems **8**(2), 141–187 (2003)
26. Edwards, S.A., Lee, E.A.: The semantics and execution of a synchronous block-diagram language. Science of Computer Programming **48**(1), 21–42 (2003)
27. Edwards, S.A., Lee, E.A.: The Semantics and Execution of a Synchronous Block-Diagram Language. In: Science of Computer Programming, vol. 48. Elsevier (2003)
28. Ésik, Z.: Axiomatizing the least fixed point operation and binary supremum. In: P. Clote, H. Schwichtenberg (eds.) Computer Science Logic (CSL'00), LNCS 1862, pp. 302–316. Springer (2000)
29. Fiore, M., Moggi, E., Sangiorgi, D.: A fully abstract model for the π-calculus. Information and Computation **179**(1), 76–117 (2002)
30. Gamatié, A., Gonnord, L.: Static analysis of synchronous programs in Signal for efficient design of multi-clocked embedded systems. ACM Sigplan Notices **46**(5), 71–80 (2011)
31. Gemünde, M., Brandt, J., Schneider, K.: Clock refinement in imperative synchronous languages. EURASIP J. Emb. Sys. **2013**, 3 (2013)
32. Guernic, P.L., Goutier, T., Borgne, M.L., Maire, C.L.: Programming real time applications with SIGNAL. Proceedings of the IEEE **79**(9), 1321–1336 (1991)
33. Halbswachs, N.: Synchronous Programming of Reactive Systems. Kluwer Academic Publishers (1993)
34. Halbwachs, N., Caspi, P., Raymond, P., Pilaud, D.: The synchronous data-flow programming language LUSTRE. Proceedings of the IEEE **79**(9), 1305–1320 (1991)
35. Hamon, G.: A denotational semantics for Stateflow. In: EMSOFT'05: Proceedings of the 5th ACM International Conference on Embedded Software, pp. 164–172. ACM Press, New York, NY, USA (2005)
36. von Hanxleden, R.: SyncCharts in C—A Proposal for Light-Weight, Deterministic Concurrency. In: Proc. Int'l Conference on Embedded Software (EMSOFT'09), pp. 225–234. ACM, Grenoble, France (2009)
37. von Hanxleden, R., Mendler, M., Aguado, J., Duderstadt, B., Fuhrmann, I., Motika, C., Mercer, S., O'Brien, O.: Sequentially Constructive Concurrency—A conservative extension of the synchronous model of computation. In: Proc. Design, Automation and Test in Europe Conference (DATE'13), pp. 581–586. IEEE, Grenoble, France (2013)
38. von Hanxleden, R., Mendler, M., Aguado, J., Duderstadt, B., Fuhrmann, I., Motika, C., Mercer, S., O'Brien, O., Roop, P.: Sequentially Constructive Concurrency—A conservative extension of the synchronous model of computation. ACM Transactions on Embedded Computing Systems, Special Issue on Applications of Concurrency to System Design **13**(4s), 144:1–144:26 (2014)
39. Harel, D.: Statecharts: A visual formalism for complex systems. Science of Computer Programming **8**(3), 231–274 (1987)
40. Harel, D., Naamad, A.: The STATEMATE semantics of Statecharts. ACM Transactions on Software Engineering **5**(4), 293–333 (1996)
41. Hennessy, M.: Acceptance trees. J. ACM **32**(4), 896–928 (1985)
42. Hennessy, M., Regan, T.: A process algebra for timed systems. Information and Computation **117**, 221–239 (1995)
43. Hoare, C.A.R.: Communicating Sequential Processes. Prentice Hall, Upper Saddle River, NJ (1985)
44. Huizing, C., Gerth, R., de Roever, W.: Modeling Statecharts behavior in a fully abstract way. In: M. Dauchet, M. Nivat (eds.) 13th CAAP (CAAP '88), *Lecture Notes in Computer Science*, vol. 299, pp. 271–294. Springer, Nancy, France (1988)
45. Hyland, M., Ong, L.: On full abstraction for PCF: I, II and III. Information and Computation **163**(2), 285–408 (2000)
46. Ingólfsdóttir, A., Schalk, A.: A fully abstract denotational model for observational precongruence. Theoretical Computer Science **254**(1–2), 35–61 (2001)
47. Kahn, G.: The semantics of a simple language for parallel programming. In: J.L. Rosenfeld (ed.) Information Processing 74: Proceedings of the IFIP Congress 74, pp. 471–475. IFIP, North-Holland Publishing Co. (1974)
48. Kahn, G., MacQueen, D.B.: Coroutines and networks of parallel processes. In: IFIP Congress, pp. 993–998 (1977)
49. Kok, J.N.: Denotational semantics of nets with nondeterminism. In: B. Robinet, R. Wilhelm (eds.) ESOP, *LNCS*, vol. 213, pp. 237–249. Springer (1986)
50. Kozen, D.: A completeness theorem for Kleene algebras and the algebra of regular events. Information and Computation **110**(2), 366–390 (1994)
51. Kuper, L., Turon, A., Krishnaswami, N.R., Newton, R.R.: Freeze after writing: Quasi-deterministic parallel programming with LVars. In: Principles of Programming Languages (POPL'14), pp. 257–270. ACM, New York, USA (2014)

DRAFT 20.4.2015

52. Lavagno, L., Sentovich, E.: ECL: a specification environment for system-level design. In: Proc. 36th ACM/IEEE Conf. on Design Automation (DAC'99), pp. 511–516. ACM Press, New York, NY, USA (1999)
53. Lee, E.A., Messerschmitt, D.G.: Synchronous data flow. In: Proceedings of the IEEE, vol. 75, pp. 1235–1245. IEEE Computer Society Press (1987)
54. Leung, A., Gupta, M., Agarwal, Y., Gupta, R., Jhala, R., Lerner, S.: Verifying GPU kernels by test amplification. In: Programming Language Design and Implementation PLDI 2012, pp. 383–394. ACM, New York, USA (2012)
55. Luettgen, G., Mendler, M.: The intuitionism behind Statecharts steps. ACM Transactions on Computational Logic **3**(1), 1–41 (2002)
56. Lüttgen, G., von der Beeck, M., Cleaveland, R.: Statecharts via process algebra. In: Proc. 10th International Conference on Concurrency Theory CONCUR'99, pp. 399–414 (1999)
57. Lüttgen, G., Mendler, M.: Towards a model-theory for Esterel. In: F. Maraninchi, A. Girault, E. Rutten (eds.) SLAP 2002, *ENTCS*, vol. 65,5. Elsevier Science (2002)
58. Lynch, N.: Distributed Algorithms. Morgan Kaufmann Publishers (1996)
59. Malik, S.: Analysis of cyclic combinational circuits. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **13**(7), 950–956 (1994)
60. Mandel, L., Pasteur, C., Pouzet, M.: Time refinement in a functional synchronous language. In: ACM SIGPLAN Int. Symp. on Principles and Practice of Declarative Programming (PPDP'13), pp. 169–180. ACM, New York, NY, USA (2013)
61. Mandel, L., Pouzet, M.: ReactiveML: a reactive extension to ML. In: Proc. 7th ACM SIGPLAN Int'l Conf. on Principles and Practice of Declarative Programming, pp. 82–93 (2005)
62. Maraninchi, F.: The Argos language: Graphical representation of automata and description of reactive systems. In: IEEE Workshop on Visual Languages (1991)
63. Maraninchi, F., Rémond, Y.: Argos: An automaton-based synchronous language. Computer Languages **27**(27), 61–92 (2001)
64. Mendler, M., Lüttgen, G.: Is observational congruence axiomatisable in equational Horn logic? Information and Computation **208**(6), 634–651 (2010)
65. Mendler, M., Shiple, T.R., Berry, G.: Constructive boolean circuits and the exactness of timed ternary simulation. Formal Methods in System Design **40**(3), 283–329 (2012)
66. Milner, R.: Communication and Concurrency. Prentice Hall (1989)
67. Motika, C., von Hanxleden, R., Heinold, M.: Programming deterministice reactive systems with Synchronous Java (invited paper). In: Proceedings of the 9th Workshop on Software Technologies for Future Embedded and Ubiquitous Systems (SEUS 2013), IEEE Proceedings. Paderborn, Germany (2013)
68. Ngo, V.C., Talpin, J.P., Gautier, T.: Precise deadlock detection for polychronous data-flow specifications. In: Electronic System Level Synthesis Conference (ESLsyn), Proceedings of the 2014, pp. 1–6. IEEE (2014)
69. Plotkin, G.D.: A Structural Approach to Operational Semantics. Technical Report DAIMI FN-19, University of Aarhus, Denmark (1981). `http://homepages.inf.ed.ac.uk/gdp/publications/SOS.ps`
70. Pnueli, A., Shalev, M.: What is in a step: On the semantics of Statecharts. In: Proc. Int. Conf. on Theoretical Aspects of Computer Software (TACS'91), pp. 244–264. Springer, London, UK (1991)
71. Potop-Butucaru, D., Edwards, S.A., Berry, G.: Compiling Esterel. Springer (2007)
72. Pouzet, M., Raymond, P.: Modular static scheduling of synchronous data-flow networks: an efficient symbolic representation. In: EMSOFT, pp. 215–224 (2009)
73. de Roever, W.P., Lüttgen, G., Mendler, M.: What is in a step: New perspectives on a classical question. In: Z. Manna, D.A. Peled (eds.) Time for Verification, pp. 370–399. Springer LNCS 6200 (2010)
74. Schneider, K.: The synchronous programming language Quartz. Internal Report 375, Department of Computer Science, University of Kaiserslautern, Kaiserslautern, Germany (2009)
75. Schneider, K., Brandt, J., Schuele, T.: Causality analysis of synchronous programs with delayed actions. In: Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES'04), pp. 179–189. ACM, Washington D.C., USA (2004)
76. Schneider, K., Brandt, J., Schuele, T., Tuerk, T.: Maximal causality analysis. In: Conference on Application of Concurrency to System Design (ACSD'05), pp. 106–115. IEEE Computer Society, St. Malo, France (2005)
77. Schneider, K., Brandt, J., Schüle, T., Türk, T.: Improving constructiveness in code generators. In: F. Maraninchi, M. Pouzet, V. Roy (eds.) Int'l Workshop on Synchronous Languages, Applications, and Programming (SLAP'05), pp. 1–19. ENTCS, Edinburgh, Scotland, UK (2005)
78. Shiple, T.R., Berry, G., Touati, H.: Constructive Analysis of Cyclic Circuits. In: Proc. European Design and Test Conference (ED&TC'96), Paris, France, pp. 328–333. IEEE Computer Society Press, Los Alamitos, California, USA (1996)

DRAFT 20.4.2015

79. Talpin, J.P., Brandt, J., Gemünde, M., Schneider, K., Shukla, S.: Constructive polychronous systems. Science of Computer Programming **96**(3), 377–394 (2014)
80. Talpin, J.P., Ouy, J., Gautier, T., Besnard, L., Guernic, P.L.: Compositional design of isochronous systems. Science of Computer Programming **77**(2), 113–128 (2012)
81. Tardieu, O., Edwards, S.A.: Scheduling-independent threads and exceptions in SHIM. In: Proceedings of the International Conference on Embedded Software (EMSOFT'06), pp. 142–151. ACM, Seoul, South Korea (2006)
82. Tardieu, O., Edwards, S.A.: Instanteneous transitions in Esterel. In: Proceedings of Model Driven High-Level Programming of Embedded Systems (SLA++P'07). Braga, Portugal (2007)
83. Vechev, M., Yahav, E., Raman, R., Sarkar, V.: Automatic verification of determinism for structured parallel programs. In: R. Cousot, M. Martel (eds.) Static Analysis (SAS 2010), *LNCS*, vol. 6337, pp. 455–471. Springer (2010)
84. von der Beeck, M.: A comparison of Statecharts variants. In: H. Langmaack, W. de Roever, J. Vytopil (eds.) 3rd International School and Symposium on Formal Techniques in Real-time and Fault-tolerant Systems (FTRTFT '94), *Lecture Notes in Computer Science*, vol. 863, pp. 128–148. Springer, Lübeck, Germany (1994)
85. Yuki, T., Feautrier, P., Rajopadye, S., Saraswat, V.: Array dataflow analysis for polyhedral X10 programs. In: Principles and Practice of Parallel Programming (PPoPP 2013), pp. 23–34. ACM, New York, USA (2013)

DRAFT 20.4.2015