

The CURIO Language Specification and its Machine-Verified Confluence Proof

Malcolm Dowse and Andrew Butterfield*
Trinity College Dublin

July 24, 2005

Contents

1	Introduction	2
2	API Models and I/O Contexts	3
2.1	Definition	3
2.2	Four Examples	6
2.3	Discussion	11
3	Curio – a Language for Reasoning About I/O	13
3.1	Language Primitives	14
3.2	Examples	15
3.3	Semantics	16
3.4	Convergence/Divergence and the Implementation	19
3.5	Convergence is Recursively Enumerable	21
4	The Machine-Verified Confluence Proof	24
4.1	Introduction	24
4.2	Preliminaries	25
4.3	Initial results	29
4.4	Analysing Failure	32
4.5	Failure implies Divergence	39
4.6	Confluence of Reduction	45
A	Sparkle Proof Sections	50

*Supported by Enterprise Ireland Basic Research Grant SC-2002-283.

1 Introduction

We present the specification and a simple implementation of the language CURIO, a small monadic concurrent functional language. This language was designed to give a semantics to concurrent I/O in pure functional languages by way of modelling the API directly. Central to CURIO is the fact that program execution is deterministic if an I/O model obeys a pre-condition. We give the details of the machine-verified proof of this confluence result.

All theorems and lemmas have been fully machine verified with Sparkle [dMvEP01], a semi-automated LCF-style [Pau87] proof assistant specifically designed for reasoning about lazy functional languages. Our metalanguage is Core-Clean, a stripped down version of the functional language Clean [PvE01] with no support for ad-hoc polymorphism (type classes) or I/O. We use Haskell [PHWH03] syntax in this document since it will be more familiar to most readers and, anyway, the differences are slight. The one main difference is the presence of strictness annotation, but this can be emulated in Haskell using (\$) or seq and we skim over any references to it in this document.

The proof style is denotational, so we quantify and induct over elements of domains, not syntax. It also means that proofs tend to become cluttered with uninteresting but unavoidable results concerning \perp . The main confluence proof requires about 100 separate results, many of which relate solely to how \perp propagates.

The semantics of a less sophisticated precursor to CURIO can be found in [DBvE05], published in the proceedings of IFL 2004.

2 API Models and I/O Contexts

We begin by introducing our way of modelling I/O and the notion of an I/O contexts which is required to ensure determinism. These will both be used to give the semantics of a real language in Section 3.

2.1 Definition

One of the goals of our research is to describe the interaction of programs with its environment by way of modelling the behaviour of each individual action on a “world state”.

The definition of an I/O model is as follows: in the metalanguage it is the following 4-tuple parameterised by five types:

$$\mathfrak{s} :: \text{IOModel } \nu \alpha \rho \omega \varsigma \triangleq \langle \text{af} :: \alpha \rightarrow \omega \rightarrow (\omega, \nu) \\ , \text{wa} :: \alpha \rightarrow \omega \rightarrow \text{Bool} \\ , \text{ap} :: \varsigma \rightarrow \alpha \rightarrow \text{Bool} \\ , \text{pf} :: \rho \rightarrow \varsigma \rightarrow (\varsigma, \varsigma) \rangle$$

The five types take the following roles, all of which will be explained in more detail later: ν : return values from actions; α : actions; ρ : “parameters” to the splitting of contexts; ω : world state; ς : I/O contexts. As we shall see, these four components are enough to describe state-based I/O which permits both a **fork**-like concurrency primitive and communication.

When giving general properties, we always assume the existence of some arbitrary, implicit I/O model called \mathfrak{s} which binds the five types ν , α , ρ , ω and ς , and the four functions **af**, **wa**, **ap** and **pf**. It is important to remember that the five types are domains with a \perp element, and the four functions can be any arbitrary computable function whose results may be undefined (\perp)¹.

The four functions are now explained in detail.

The API – **af** and **wa**

The function $\text{af} :: \alpha \rightarrow \omega \rightarrow (\omega, \nu)$ defines the state-transformer for each action. Elements of type α identify I/O actions which can be performed and for any action $a :: \alpha$, $\text{af } a :: \omega \rightarrow (\omega, \nu)$ is the state-transformer for that action. This describes how the action changes the global state of type ω , and what return value of type ν it yields. Type ν is typically a sum-type capable of storing **Ints**, **Bools**, **Chars** or any other value that an action might need to return.

¹This probably isn’t ideal, especially for the world model, but it’s not serious. Any set-based model can be turned into a domain-based one by (1) turning the sets into “flat” domains by adding a bottom element, (2) making the functions bottom-preserving, and (3) showing that the functions are indeed Turing-computable.

$$\begin{aligned}
|||_s, \text{ally}_s & : \alpha \rightarrow \alpha \rightarrow \mathbb{B} \\
a_l |||_s a_r & \triangleq \forall w \in \omega. \forall w_2 \in \omega. \mathbf{wa} a_l w = \text{FALSE} \wedge \mathbf{wa} a_r w = \text{FALSE} \implies \forall v_l \in \nu. \forall v_r \in \nu. \\
& \quad (\exists w_1 \in \omega. \mathbf{af} a_l w = (w_1, v_l) \wedge \mathbf{af} a_r w_1 = (w_2, v_r)) \\
& \quad \iff \\
& \quad (\exists w_1 \in \omega. \mathbf{af} a_r w = (w_1, v_r) \wedge \mathbf{af} a_l w_1 = (w_2, v_l)) \\
\text{ally}_s(a_l, a_r) & \triangleq \forall w \in \omega. \forall w_1 \in \omega. \forall v \in \nu. \mathbf{wa} a_l w = \text{FALSE} \wedge \mathbf{af} a_l w = (w_1, v) \wedge \\
& \quad \neg(\mathbf{wa} a_r w = \text{TRUE}) \implies \mathbf{wa} a_r w = \mathbf{wa} a_r w_1 \\
\sqsubseteq_s, \diamond_s & : \varsigma \rightarrow \varsigma \rightarrow \mathbb{B} \\
c_1 \sqsubseteq_s c_2 & \triangleq \forall a \in \alpha. \mathbf{ap} a c_1 \implies \mathbf{ap} a c_2 \\
c_l \diamond_s c_r & \triangleq \forall a_l \in \alpha. \forall a_r \in \alpha. \mathbf{ap} c_l a_l \wedge \mathbf{ap} c_r a_r \implies \\
& \quad a_l |||_s a_r \wedge \text{ally}_s(a_l, a_r) \wedge \text{ally}_s(a_r, a_l)
\end{aligned}$$

Figure 1: Relations on actions and contexts

If communication is to be permitted then there must be occasions in which an action is waiting for something to occur and cannot proceed. The function $\mathbf{wa} :: \alpha \rightarrow \omega \rightarrow \text{Bool}$ indicates exactly this. An action a can only be performed in world w when $\mathbf{wa} a w = \text{FALSE}$. We say an action a is **stalled** (in world state w) if $\mathbf{wa} a w = \text{TRUE}$.

Figure 1 defines two relations on actions and the API.

- $a_1 |||_s a_2$: for any two actions a_1 and a_2 , if neither are stalled then the order in which they are executed is irrelevant – both with regard to their effect on world state and their return values. $|||_s$ is symmetric (but not necessarily transitive or reflexive.)
- $\text{ally}_s(a_1, a_2)$: If action a_1 is not stalled then performing it cannot cause action a_2 , if also not stalled, to then become stalled. The word “ally” describes the fact that action a_1 will not obstruct or hinder action a_2 – they are in effect working with each other.

The $a_1 |||_s a_2$ condition allows for two possibilities: either both actions succeed for both orderings or for both orderings one action fails. For example, if $\mathbf{af} a_1 w = (w_1, v_1)$ and $\mathbf{af} a_2 w = (w_2, v_2)$ then it is still possible that $\mathbf{af} a_2 w_1 = \perp$ and $\mathbf{af} a_1 w_2 = \perp$.

I/O Contexts – ap

The functions \mathbf{af} and \mathbf{wa} define the API. What is now needed is some means of reining in the power of a concurrent (sub-)program by giving it only a

limited set of actions which it is allowed to perform.

We call these permission sets **I/O contexts**. I/O contexts are elements of the type ς and the function $\mathbf{ap} :: \varsigma \rightarrow \alpha \rightarrow \mathbf{Bool}$ defines the actions permitted by any context. A context c can be thought of as the set of actions a such that $\mathbf{ap} \ c \ a = \mathbf{TRUE}$. Each (sub-)program has a context associated with it at run-time. The context determines what actions the (sub-)program is allowed to perform, and if used in a controlled fashion it gives a mechanism for ensuring determinism.

Figure 1 defines two important relations on contexts.

- $c_1 \sqsubseteq_s c_2$: any action permitted by context c_1 is also permitted by c_2 .
- $c_1 \diamond_s c_2$: for all actions a_1 and a_2 , permitted by contexts c_1 and c_2 respectively, it is true that $a_1 \parallel_s a_2$, $\mathbf{ally}_s(a_1, a_2)$ and $\mathbf{ally}_s(a_2, a_1)$.

(\diamond_s is symmetric and \sqsubseteq_s is a pre-order – by definition.)

Enforcing Determinism – pf

Say a process running in context c forks into two processes running in contexts c_l and c_r . To guarantee deterministic behaviour:

- The order in which actions in the two child sub-programs are performed should be irrelevant. That is to say: if the run-time system can make a choice between doing an action a_l permitted by c_l or an action a_r permitted by c_r , then neither action can impede the other causing it to become stalled, and when both are finally executed their order shouldn't affect the resultant world state or the actions' return values: $c_l \diamond_s c_r$.
- No child process should be allowed to perform an action forbidden by the parent context: $c_l \sqsubseteq_s c$ and $c_r \sqsubseteq_s c$.

These properties are usually undecidable. This is a problem, since we want to give describe a language implementation. Therefore we assume the existence of a function which has been proved to obey these exact properties. The function $\mathbf{pf} :: \rho \rightarrow \varsigma \rightarrow (\varsigma, \varsigma)$ splits a context returning two new ones for the two concurrent left and right sub-programs. Elements of the type ρ give the programmer some flexibility with regard to how he or she wishes the current context to be split. \mathbf{PRE}_s , as defined in Figure 2, is the pre-condition that \mathbf{pf} must obey and we will show later that if it does, any I/O performing program on that I/O model remains deterministic.

Sometimes we don't need the full power of \mathbf{PRE}_s , just something a lot weaker. \mathbf{pre}_s only guarantees the second of the two properties mentioned above.

$$\begin{aligned}
\text{PRE}_s &\triangleq \forall p \in \rho. \forall c \in \varsigma. \forall c_l \in \varsigma. \forall c_r \in \varsigma. \mathbf{pf} \ p \ c = (c_l, c_r) \implies c_l \sqsubseteq_s c \wedge c_r \sqsubseteq_s c \wedge c_l \diamond_s c_r \\
\text{pre}_s &\triangleq \forall p \in \rho. \forall c \in \varsigma. \forall c_l \in \varsigma. \forall c_r \in \varsigma. \mathbf{pf} \ p \ c = (c_l, c_r) \implies c_l \sqsubseteq_s c \wedge c_r \sqsubseteq_s c
\end{aligned}$$

Figure 2: The pre-conditions PRE_s and pre_s

2.2 Four Examples

To give an idea of the flexibility of this approach, here are a few examples of complete I/O models all of which obey PRE_s . (All lemmas have been machine-verified.)

Communication buffers

I/O model bfft is a simple 1-1 communication buffer. World state is of type $[\text{Int}]$, and there are two actions: $\text{SEND } i$ places i at the end of the list (returning 0, a token value); RCVE removes the first element from the list and returns it. If the current context is SUBC TRUE one can only send, if it is SUBC FALSE one can only receive, and if it is TOPC both are permitted. One can only split the context if it is TOPC . This will give one sub-program the right to send, and the other the right to receive.

Lemma 2.1. PRE_{bfft} .

Proof. If $\mathbf{pf}_{\text{bfft}}$ splits TOPC into SUBC FALSE and SUBC TRUE :

- $(\text{SUBC FALSE}) \diamond_{\text{bfft}} (\text{SUBC TRUE})$: Only a send and a receive can be performed. $\text{RCVE} \parallel_{\text{bfft}} (\text{SEND } i)$ holds, because if the receive isn't stalled, the buffer must be non-empty, and therefore the actions must affect different parts of the buffer and be order independent. $\text{ally}_{\text{bfft}}(\text{RCVE}, \text{SEND } i)$ is trivial because a send is always non-stalled. $\text{ally}_{\text{bfft}}(\text{SEND } i, \text{RCVE})$ is true because adding an element to a buffer cannot cause a receive to become stalled.
- $(\text{SUBC FALSE}) \sqsubseteq_{\text{bfft}} \text{TOPC}$ and $(\text{SUBC TRUE}) \sqsubseteq_{\text{bfft}} \text{TOPC}$: trivial (TOPC doesn't forbid any actions.)

□

Many-to-Many Mutexes

I/O model lock allows many processes to be synchronised. World state is of type Bool indicating whether or not the lock is set. There are three actions: LOCK sets world state to TRUE ; UNLOCK sets world state to FALSE ; WAIT will stall until world state is FALSE , then proceed leaving it unchanged.

$\text{bffr} :: \text{IOModel } \nu_{\text{Bffr}} \alpha_{\text{Bffr}} \rho_{\text{Bffr}} \omega_{\text{Bffr}} \varsigma_{\text{Bffr}} \triangleq \langle \text{af}_{\text{Bffr}}, \text{wa}_{\text{Bffr}}, \text{ap}_{\text{Bffr}}, \text{pf}_{\text{Bffr}} \rangle$

$$\begin{array}{llll}
\nu_{\text{Bffr}} & \triangleq & \text{Int} & \alpha_{\text{Bffr}} & \triangleq & \text{SEND Int} \mid \text{RCVE} & \omega_{\text{Bffr}} & \triangleq & [\text{Int}] \\
\rho_{\text{Bffr}} & \triangleq & \text{Bool} & \varsigma_{\text{Bffr}} & \triangleq & \text{TOPC} \mid \text{SUBC Bool} & & & \\
\text{af}_{\text{Bffr}} & (\text{SEND } i) & is & \triangleq & (is ++ [i], 0) \\
\text{af}_{\text{Bffr}} & \text{RCVE} & (i : is) & \triangleq & (is, i) \\
\text{wa}_{\text{Bffr}} & (\text{SEND } i) & is & \triangleq & \text{FALSE} \\
\text{wa}_{\text{Bffr}} & \text{RCVE} & is & \triangleq & \text{null } is \\
\text{ap}_{\text{Bffr}} & \text{TOPC } a & & \triangleq & \text{TRUE} \\
\text{ap}_{\text{Bffr}} & (\text{SUBC } b) & (\text{SEND } i) & \triangleq & b \\
\text{ap}_{\text{Bffr}} & (\text{SUBC } b) & \text{RCVE} & \triangleq & \text{not } b \\
\text{pf}_{\text{Bffr}} & b & \text{TOPC} & \triangleq & (\text{SUBC } b, \text{SUBC } (\text{not } b))
\end{array}$$

Figure 3: bffr – a 1-to-1 communication buffer

Contexts are either “TRUE”, meaning all actions are permitted, or “FALSE”, meaning just “UNLOCK” and “WAIT” are allowed.

Contexts can be split as many times as one likes, but the child context will always be “FALSE”. This means that although a program’s context *can* be TRUE at the top-level, allowing “LOCK” to be performed, “LOCK” can never be performed concurrently with any other action. The world state doesn’t keep track of how many processes are waiting for the mutex to be released, so if it was released then locked again, a waiting process might miss this event.

There are no (useful) return values in this I/O model.

Lemma 2.2. PRE_{lock} .

Proof. The function pf_{Lock} always splits contexts into FALSE and FALSE:

- Proving $\text{FALSE} \diamond_{\text{lock}} \text{FALSE}$: The combinations of actions are (1) two unlocks, (2) two waits or (3) an unlock and a wait. $\text{UNLOCK} \parallel_{\text{lock}} \text{UNLOCK}$ and $\text{ally}_{\text{lock}}(\text{UNLOCK}, \text{UNLOCK})$ hold because UNLOCK changes state in the same way, and is never stalled. $\text{WAIT} \parallel_{\text{lock}} \text{WAIT}$ and $\text{ally}_{\text{lock}}(\text{WAIT}, \text{WAIT})$ is true since WAIT doesn’t affect world state, is order-independent and also cannot cause any action to become stalled. $\text{UNLOCK} \parallel_{\text{lock}} \text{WAIT}$ and $\text{ally}_{\text{lock}}(\text{UNLOCK}, \text{WAIT})$: If both unlocking and waiting are non-stalled, then their order is irrelevant since wait doesn’t change the world state – nothing can cause unlock to become stalled, and no amount of unlocking can cause a wait to become stalled.
- $\text{FALSE} \sqsubseteq_{\text{lock}} b$ for all b : True by reflexivity if $b = \text{FALSE}$, and if

$\text{lock} :: \text{IOModel } \nu_{\text{Lock}} \alpha_{\text{Lock}} \rho_{\text{Lock}} \omega_{\text{Lock}} \varsigma_{\text{Lock}} \triangleq \langle \text{af}_{\text{Lock}}, \text{wa}_{\text{Lock}}, \text{ap}_{\text{Lock}}, \text{pf}_{\text{Lock}} \rangle$

$$\begin{aligned}
\nu_{\text{Lock}} &\triangleq () & \omega_{\text{Lock}} &\triangleq \text{Bool} & \alpha_{\text{Lock}} &\triangleq \text{LOCK} \mid \text{UNLOCK} \mid \text{WAIT} \\
\rho_{\text{Lock}} &\triangleq () & \varsigma_{\text{Lock}} &\triangleq \text{Bool} \\
\text{af}_{\text{Lock}} \quad \text{LOCK} \quad b &\triangleq (\text{TRUE}, ()) \\
\text{af}_{\text{Lock}} \quad \text{UNLOCK} \quad b &\triangleq (\text{FALSE}, ()) \\
\text{af}_{\text{Lock}} \quad \text{WAIT} \quad b &\triangleq (b, ()) \\
\text{pf}_{\text{Lock}} \quad () \quad c &\triangleq (\text{FALSE}, \text{FALSE}) \\
\text{wa}_{\text{Lock}} \quad a \quad b &\triangleq \begin{cases} b, & a = \text{WAIT} \\ \text{FALSE}, & \text{otherwise} \end{cases} \\
\text{ap}_{\text{Lock}} \quad b \quad a &\triangleq \begin{cases} \text{FALSE}, & a = \text{LOCK}, b = \text{FALSE} \\ \text{TRUE}, & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 4: lock – a mutex

$b = \text{TRUE}$ then it is true because context “TRUE” doesn’t forbid any actions.

□

An Integer Variable

The I/O model ivar is an integer variable which can be written to and read using the actions READI and WRITEI . Neither of these can ever be stalled.

The context is either NONEC (no actions are permitted), READC (only reading is allowed) or WRITEC (reading and writing are both allowed.) When splitting a context, if the context is NONEC or READC then the parameter is ignored and the left and right context remain the same as that of the parent. When splitting context “WRITEC”, however, depending on which of “LEFTWR”, “RIGHTWR” or “BOTHRD” is given as parameter, either one side can be allowed to do everything leaving nothing for the other side, or both sides can be allowed to only read the integer.

Lemma 2.3. PRE_{ivar} .

Proof. Since no action can be stalled, the $\text{ally}_{\text{ivar}}$ property holds automatically for any pair of actions.

- Proving \diamond_{ivar} : the function pf_{IVar} can split contexts in three different ways. $\text{NONEC } \diamond_{\text{ivar}} \text{ NONEC}$ and $\text{WRITEC } \diamond_{\text{ivar}} \text{ NONEC}$ are trivially true, since no actions are permitted by NONEC . When proving $\text{READC } \diamond_{\text{ivar}} \text{ READC}$, only “READI” is allowed, and $\text{READI} \parallel_{\text{ivar}} \text{READI}$ is true since it doesn’t change the world state.

$$\text{ivar} :: \text{IOModel } \nu_{\text{IVar}} \alpha_{\text{IVar}} \rho_{\text{IVar}} \omega_{\text{IVar}} \varsigma_{\text{IVar}} \triangleq \langle \text{af}_{\text{IVar}}, \text{wa}_{\text{IVar}}, \text{ap}_{\text{IVar}}, \text{pf}_{\text{IVar}} \rangle$$

$\nu_{\text{IVar}} \triangleq$	Int	$\rho_{\text{IVar}} \triangleq$	$\text{LEFTWR} \mid \text{RIGHTWR} \mid \text{BOTHRD}$
$\omega_{\text{IVar}} \triangleq$	Int	$\varsigma_{\text{IVar}} \triangleq$	$\text{NONEC} \mid \text{READC} \mid \text{WRITEC}$
		$\alpha_{\text{IVar}} \triangleq$	$\text{READI} \mid \text{WRITEI Int}$
af_{IVar}	READI	$i \triangleq$	(i, i)
af_{IVar}	$(\text{WRITEI } i_1)$	$i \triangleq$	$(i_1, 0)$
wa_{IVar}	a	$i \triangleq$	FALSE
ap_{IVar}	NONEC	$a \triangleq$	FALSE
ap_{IVar}	READC	$\text{WRITEI} \triangleq$	FALSE
ap_{IVar}	READC	$\text{READI} \triangleq$	TRUE
ap_{IVar}	WRITEC	$a \triangleq$	TRUE
pf_{IVar}	p	$\text{NONEC} \triangleq$	$(\text{NONEC}, \text{NONEC})$
pf_{IVar}	p	$\text{READC} \triangleq$	$(\text{READC}, \text{READC})$
pf_{IVar}	LEFTWR	$\text{WRITEC} \triangleq$	$(\text{WRITEC}, \text{NONEC})$
pf_{IVar}	RIGHTWR	$\text{WRITEC} \triangleq$	$(\text{NONEC}, \text{WRITEC})$
pf_{IVar}	BOTHRD	$\text{WRITEC} \triangleq$	$(\text{READC}, \text{READC})$

Figure 5: `ivar` – a shared integer variable

- $c_1 \sqsubseteq_{\text{ivar}} c_2$ for all c_1, c_2 as split by pf_{IVar} : It is clear that $\text{NONEC} \sqsubseteq_{\text{ivar}} \text{READC}$ and $\text{READC} \sqsubseteq_{\text{ivar}} \text{WRITEC}$, so it can be seen directly from the definition of pf_{IVar} that when contexts are split this property is always obeyed.

□

Terminal I/O

As a small real-world I/O example, consider the model of terminal I/O given in Figure 6.

There are two actions: (`PUTC` c) writes character c , and `GETC` reads a character. The model is not perfect, since it does not capture any interesting temporal properties of `stdin/stdout`. There is no notion of absolute time – each outputted `Char` gives rise to 0 or more inputted characters instantaneously. It is adequate, nonetheless, and obeys the simple property that if characters are available for input then outputting characters cannot change that. The model also lets us exploit the fact that `stdin` and `stdout`

$$\begin{aligned}
\text{term} &:: \text{IOModel } \nu_{\text{Term}} \alpha_{\text{Term}} \rho_{\text{Term}} \omega_{\text{Term}} \varsigma_{\text{Term}} \triangleq \langle \text{af}_{\text{Term}}, \text{wa}_{\text{Term}}, \text{ap}_{\text{Term}}, \text{pf}_{\text{Term}} \rangle \\
\nu_{\text{Term}} &\triangleq \text{Char} & \alpha_{\text{Term}} &\triangleq \text{PUTC Char} \mid \text{GETC} \\
\omega_{\text{Term}} &\triangleq ([\text{Char}], \text{TermIO}) & \rho_{\text{Term}}, \varsigma_{\text{Term}} &\triangleq \text{IOCXT Bool Bool} \\
&& \text{TermIO} &\triangleq \text{TERMIO (Char} \rightarrow ([\text{Char}], \text{TermIO})) \\
\text{wa}_{\text{Term}} (\text{PUTC } c) (cs, t) && &\triangleq \text{FALSE} \\
\text{af}_{\text{Term}} (\text{PUTC } c) (cs, \text{TERMIO } f) && &\triangleq ((cs \# \text{fst } (f c), \text{snd } (f c)), ' ') \\
\text{wa}_{\text{Term}} \text{GETC } (cs, t) && &\triangleq \text{null } cs \\
\text{af}_{\text{Term}} \text{GETC } (c : cs), t && &\triangleq ((cs, t), c) \\
\text{ap}_{\text{Term}} (\text{IOCXT } b _) (\text{PUTC } c) && &\triangleq b \\
\text{ap}_{\text{Term}} (\text{IOCXT } _ b) \text{GETC} && &\triangleq b \\
\text{pf}_{\text{Term}} (\text{IOCXT } bp_p \ bg_p) (\text{IOCXT } bp \ bg) && &\triangleq \\
&& &(\text{IOCXT } (bp \&\& bp_p) (bg \&\& bg_p), \text{IOCXT } (bp \&\& \text{not } bp_p) (bg \&\& \text{not } bg_p))
\end{aligned}$$

Figure 6: `term` – a model for Terminal I/O

are usually two separate handles, and we may want a process to wait for input on one whilst another outputs data on the other. (This is probably only useful for filtering programs such as `grep`. Interactive programs would usually require a lock-step synchronisation of input and output).

The proof of PRE_{term} is not difficult. It permits the same degree of concurrency as the `bfft` model. Reading may take place in parallel with writing in this semantics since (1) if a `Char` is ready for input then writing a character will not change that and (2) if a character is not available for input then the read will stall until one is. The function pf_{Term} guarantees that if many processes are running concurrently then at most one can call `PUTC` and at most one can perform `GETC`.

The `term` model may be seen, in one sense, as just a generalisation of `bfft`. Consider the type `TermIO` once again. One possible instance of it is `loopBack`, defined as follows:

$$\begin{aligned}
\text{loopBack} &:: \text{TermIO} \\
\text{loopBack} &= \text{TermIO } (\backslash c \rightarrow ([c], \text{loopBack}))
\end{aligned}$$

This really just models the “semantics” of a user who re-inputs every character outputted to him/her. Therefore it is no different to a communication buffer between the sending and receiving processes.

2.3 Discussion

The small examples give a flavour of what is possible. One of the most surprising and unusual aspects of this system, we believe, is that we don't need to explicitly mention communication channels at all. By permitting actions to be temporarily stalled, one then has enough machinery to synchronise processes and safely transfer information via the world state.

The key to guaranteeing determinism is making sure that when a context c is split into c_l and c_r to allow concurrency, an action in one context cannot influence the behaviour of an action in the other context. This means that an action permitted by c_l cannot block an action permitted by c_r (or vice versa) and the order in which they are executed must be irrelevant. This rules out competition for a limited resource. Some typically non-deterministic constructs are multiple-writer streams, multiple-reader streams and shared, mutable variables.

Actions can become stalled, however, like “RCVE” in the buffer example. Roughly speaking, an action a_1 can only cause a (not necessarily different) action a_2 to become stalled if they can under no circumstances be executed concurrently. Also, if an action a_1 can successfully predict whether another action a_2 is going to stall then a_1 cannot be run concurrently with any action which changes whether or not a_2 is stalled. For example, in the `bfft` model neither of the sub-contexts, which allow just sending or just receiving, could be modified to permit an action which returns whether or not the buffer is empty. However, the receiving context could allow one to wait until the buffer is non-empty and the sending context could allow one to wait until the buffer is empty, and neither would admit non-determinism.

There are a few interesting sub-classes of actions.

- Observer actions, which never change world state (like “WAIT” and “READI”). If a_1 is an observer action, then for any a_0 , $\text{ally}_s(a_1, a_0)$ and for any two observer actions a_1 and a_2 , $a_1 \parallel_s a_2$ holds.
- Actions which always return the same value (like “LOCK”, “WRITEI 7”, or an action which, say, increments a counter without indicating its value.) If action a_1 is of this form then $a_1 \parallel_s a_1$.
- Actions which can never be stalled, like “SEND i ” and “WRITEI i ”. If action a_1 is never stalled then for any action a_0 , $\text{ally}_s(a_0, a_1)$.
- Commutative contexts: if it is true that $c \diamond_s c$ for some c , (which will be true if there is some p such that $\text{pf } p \ c = (c, c)$, as is the case with READC in the `ivar` example) then the actions permitted by c in a monadic setting form a commutative monad. In other words, there are absolutely no constraints on the ordering of actions.

It should also be noted that many “reasonable” properties of I/O models are not always present. These include: a context which permits all actions;

a context which permits no actions; a way of splitting a context such that all permissions are given to one side only; a way splitting contexts in a symmetric way (if a context can be split into (c_l, c_r) can it also be split into (c_r, c_l) ?). These weren't necessary for the confluence proof, so we didn't bother to include them, but their existence would probably make things somewhat tidier.

3 Curio – a Language for Reasoning About I/O

In this section we introduce the CURIO language. This is a small functional language with concurrency, interprocess communication and monadic constructs which together let the user write expressive programs which perform I/O. This expressivity is due to our ability to enforce determinism.

In reality, CURIO is less a full language specification than a rigorous semantics for a collection of powerful I/O primitives which can be wrapped around a pure functional language using solely its denotational semantics. This overall approach is not new. In fact, at the outset it bears many similarities to that taken by the Haskell community:

“Our semantics is stratified in two levels: an inner denotational semantics which describes the behaviour of pure terms, while an outer monadic semantics describes the behaviour of IO computations.” [Pey01]

The Concurrent Haskell language [PGF96] is probably the lazy functional language with the most fully-fledged semantics for I/O. What makes CURIO different to the semantics of Concurrent Haskell is the nature of this outer operational semantics.

Concurrent Haskell’s semantics for I/O is in effect a *co-inductive* one. I/O actions represent labelled transitions in a CCS-style process calculus. Each action is a distinct observable event, and the meaning of a program is solely determined by the (possibly non-terminating) order in which actions occur.

CURIO’s semantics is *inductive* rather than co-inductive. There exists a “world-state” which the program interacts with and modifies when doing I/O, and the observable effect of a program is its resultant world-state. Whereas with Concurrent Haskell all actions are observable, with CURIO it is only the cumulative effect of a finite number of actions over a program’s lifetime that is observable.

Naturally there are pros and cons to both approaches. The most immediate advantage to the CCS-style semantics is that by using it one can distinguish two infinite programs (for example, a program which loops continuously and, say, an internet server designed never to terminate). It is also highly elegant.

What makes the semantics of CURIO somewhat more expressive (if less elegant) is that it ascribes to each action a precise effect on world state, and therefore allows us to distinguish actions which do and do not interfere with one another. Insofar as a language semantics is a vehicle for compiler optimisations and formal proofs, this would appear to make it better – but we cannot tell at this early stage.

3.1 Language Primitives

A program in CURIO is any element of the type $\text{Prog}_{\nu\alpha\rho} \beta$ and the five I/O primitives are as follows: (unless otherwise specified it is assumed that the types ν , α and ρ are all bound by some I/O model).

$$\begin{aligned} (>>=) &:: \text{Prog}_{\nu\alpha\rho} \beta \rightarrow (\beta \rightarrow \text{Prog}_{\nu\alpha\rho} \gamma) \rightarrow \text{Prog}_{\nu\alpha\rho} \gamma \\ \text{return} &:: \beta \rightarrow \text{Prog}_{\nu\alpha\rho} \beta \\ \text{action} &:: \alpha \rightarrow \text{Prog}_{\nu\alpha\rho} \nu \\ \text{test} &:: \alpha \rightarrow \text{Prog}_{\nu\alpha\rho} \beta \rightarrow \text{Prog}_{\nu\alpha\rho} \beta \rightarrow \text{Prog}_{\nu\alpha\rho} \beta \\ \text{par} &:: \rho \rightarrow \text{Prog}_{\nu\alpha\rho} \beta \rightarrow \text{Prog}_{\nu\alpha\rho} \gamma \rightarrow (\beta \rightarrow \gamma \rightarrow \epsilon) \rightarrow \text{Prog}_{\nu\alpha\rho} \epsilon \end{aligned}$$

Central to these language primitives is the notion of an I/O context. Each (sub-)program is executed within a context $c \in \varsigma$, and that context dictates which actions that (sub-)program can perform, and therefore affects the programs outcome.

The full semantics is given further on, but very briefly: the first two are the familiar monadic $>>=$ and **return**; **action** performs a primitive action; **test** performs one of two programs depending on whether an action is allowed by the current context; **par** runs two programs concurrently.

A program is said to be a **value** (or **evaluated**) if it is of the form **return** v , for some v . A program is said to be an **action** if it is of the form **action** a , for some a . Programs of the form **par** $p \ m_l \ m_r$ (*) are abbreviated using an infix notation as $m_l \ || \ | \ m_r$.

Of the five original types, α and ρ end up being part of the language's syntax, ς is best understood as an internal run-time data-structure in the language implementation, and ω is part of the world model (i.e. part of the language's semantics).

CURIO is a non-strict language. None of the five primitives evaluate their arguments. Therefore programs can be partial or infinite, as would be expected in a real lazy language. Consider the following example program: if the current context allows a read then an infinite number of reads are performed concurrently, otherwise -1 is returned.

```
let reads = par BOTHRD (action READI) reads f
in test READI reads (return (-1))
```

The above example serves to demonstrate why our style of operational semantics will have to be a bit unusual. We don't have lists, or "let" constructs in our language yet we are using them freely when writing programs. Since programs like the above are permitted, our operational semantics can't just manipulate a program as normal syntax. Instead we must use the denotational semantics of the metalanguage and build an operational semantics which acts on and modifies elements of domains – we reason about, induct

over and manipulate the domain of type $\text{Prog}_{\nu\alpha\rho} \beta$, not the syntax of programs of type $\text{Prog}_{\nu\alpha\rho} \beta$. This has some immediate consequences. Since all our language constructs are lazy, and since there are terms of type $\text{Prog}_{\nu\alpha\rho} \beta$ which don't have an outermost constructor, we must also describe the program's behaviour when sub-programs are undefined, or \perp .

3.2 Examples

We now give example programs making use of Haskell's do-notation.

These examples are specific to certain I/O models, so the types ν , α and ρ will differ. Given some I/O model, say `bfft`, to avoid clutter we use $\text{Prog}_{\text{bfft}} \beta$ as a convenient shorthand for $\text{Prog}_{\nu_{\text{bfft}} \alpha_{\text{bfft}} \rho_{\text{bfft}}} \beta$.

Communication Buffer Examples

The following program attempts to send a list of integers along a communication buffer. The Boolean return value indicates whether the current context permitted the SEND action.

```
sendInts :: [Int] → Progbfft Bool
sendInts []      = return True
sendInts (i:is) =
  test (Send i)
    (action (Send i) >>= \_ -> sendInts is)
    (return False)
```

The program `rcveInts c` attempts to retrieve c integers from the communication buffer. If receiving is not permitted by the program's context then it returns `[]`, the empty list.

```
rcveInts :: Int → Progbfft [Int]
rcveInts c = test Rcv (rcveInts' c) (return [])
  where rcveInts' | c <= 0      = return []
                | otherwise = do
                    i <- action Rcv
                    is <- rcveInts (c-1)
                    return (i:is)
```

Mutex Examples

The program `lockPar` runs two `lock` programs in parallel combining the two resultant return values into a tuple. This is completely general since in the `lock` model the ρ type has just one element anyway, `()`.

```
lockPar :: Proglock β → Proglock γ → Proglock (β, γ)
lockPar pl pr = par () pl pr (\v1 vr -> (v1, vr))
```

Integer Variable Examples

The following program applies a function f to the integer variable returning TRUE if the context permitted reading and writing.

```
applyFn :: (Int → Int) → Progivar Bool
applyFn f = do
  b <- test (WriteI 0) (return True) (return False)
  if b then do
    i <- action ReadI
    action (WriteI (f i))
  return b
```

Terminal I/O Examples

The function `putStr` writes a string to `stdout` (or fails if that is not allowed).

```
putStr :: String → Progterm ()
putStr [] = return ()
putStr (c:cs) = action (PutC c) >>= \_ -> putStr cs
```

`stdPermissions` returns a string which indicates which actions the current context permits.

```
stdPermissions :: Progterm String
stdPermissions = do
  in <- test GetC (return "stdin") (return "no stdin")
  out <- test (PutC 'x') (return "stdout") (return "no stdout")
  return (in ++ ", " ++ out)
```

The program `getPutC c` outputs character `c` whilst concurrently requesting a character from input. The entire program returns the character that was read.

```
getPutC :: Char → Progterm Char
getPutC c = par
  (IOcxt True False) (action (PutC c)) (action GetC) (\_ c1 -> c1)
```

3.3 Semantics

The non-deterministic single-step semantics can be found in Figure 7. We use an SOS-style notation [Plo81].

All the reduction rules describe the behaviour of what we call a world/program pair, written $w \Vdash m$, where w is the world and m is the program. This allows one to describe how a program and world-state interact over time. The context $c \in \zeta$ in which a program is run also affects how the program behaves, so reduction rules are also annotated with the current context.

$$\frac{w \Vdash m \uparrow^c}{w \Vdash m \gg= f \uparrow^c} \qquad \frac{w \Vdash m \downarrow^c}{w \Vdash m \gg= f \downarrow^c} \quad (m \text{ not a value})$$

$$w \Vdash \text{return } v \gg= f \longrightarrow^c w \Vdash f v \quad \frac{w \Vdash m \longrightarrow^c w_1 \Vdash m_1}{w \Vdash m \gg= f \longrightarrow^c w_1 \Vdash m_1 \gg= f}$$

ap c a	wa a w	af a w	Behaviour of action a
\perp			$w \Vdash \text{action } a \uparrow^c$
FALSE			$w \Vdash \text{action } a \uparrow^c$
TRUE	\perp		$w \Vdash \text{action } a \uparrow^c$
TRUE	FALSE	\perp	$w \Vdash \text{action } a \uparrow^c$
TRUE	FALSE	(w_1, v)	$w \Vdash \text{action } a \longrightarrow^c w_1 \Vdash \text{return } v$
TRUE	TRUE		$w \Vdash \text{action } a \downarrow^c$

ap c a	Behaviour of test $a m_1 m_2$
\perp	$w \Vdash \text{test } a m_1 m_2 \longrightarrow^c w \Vdash \perp$
FALSE	$w \Vdash \text{test } a m_1 m_2 \longrightarrow^c w \Vdash m_2$
TRUE	$w \Vdash \text{test } a m_1 m_2 \longrightarrow^c w \Vdash m_1$

$$w \Vdash \text{return } v \downarrow^c$$

$$w \Vdash \perp \uparrow^c$$

$$\frac{\text{pf } p c = \perp}{w \Vdash m_l \parallel \parallel_*^p m_r \uparrow^c}$$

$$\text{pf } p c = (c_l, c_r) \left\{ \begin{array}{l} \frac{w \Vdash m_l \uparrow^{c_l}}{w \Vdash m_l \parallel \parallel_*^p m_r \uparrow^c} \quad \frac{w \Vdash m_r \uparrow^{c_r}}{w \Vdash m_l \parallel \parallel_*^p m_r \uparrow^c} \\ w \Vdash \text{return } v_l \parallel \parallel_*^p \text{return } v_r \longrightarrow^c w \Vdash \text{return } v_l * v_r \\ \frac{w \Vdash m_l \longrightarrow^{c_l} w' \Vdash m'_l}{w \Vdash m_l \parallel \parallel_*^p m_r \longrightarrow^c w' \Vdash m'_l \parallel \parallel_*^p m_r} \quad (m_r \neq \perp) \\ \frac{w \Vdash m_r \longrightarrow^{c_r} w' \Vdash m'_r}{w \Vdash m_l \parallel \parallel_*^p m_r \longrightarrow^c w' \Vdash m_l \parallel \parallel_*^p m'_r} \quad (m_l \neq \perp) \\ \frac{w \Vdash m_l \downarrow^{c_l} \quad w \Vdash m_r \downarrow^{c_r}}{w \Vdash m_l \parallel \parallel_*^p m_r \downarrow^c} \quad (m_l, m_r \text{ not both values}) \end{array} \right.$$

Figure 7: Non-deterministic Single-Step Semantics for CURIO

There are three reduction relations, \longrightarrow^c , \uparrow^c and \downarrow^c .

$$\begin{aligned} w \Vdash m \longrightarrow^c w' \Vdash m' &\triangleq \text{“}w \Vdash m \text{ can reduce to } w' \Vdash m' \text{ in context } c\text{”} \\ w \Vdash m \uparrow^c &\triangleq \text{“}w \Vdash m \text{ can fail in context } c\text{”} \\ w \Vdash m \downarrow^c &\triangleq \text{“}w \Vdash m \text{ is in normal form in context } c\text{”} \end{aligned}$$

We say a world/program pair *is* in normal form (rather than can be in normal form) because, as Lemma 4.13 in Section 4 shows, despite non-determinism, if a program has converged to normal form then it cannot either fail or reduce. Similarly, if a world/program pair can fail or reduce, then it cannot be in normal form. Failure may seem a slightly curious thing to include, but, as the implementation in the next subsection will show, it is necessary. A single reduction step in CURIO can correspond to many individual steps in the operational semantics of the metalanguage. Equivalently, a failed single-step reduction may denote a never-ending sequence of reduction steps in the metalanguage.

Values (programs of the form `return v`) on their own are in normal form. If a world/program pair is in normal form but not a value, then we say it is **stalled**. If this occurs, the convergence to normal form is caused by stalled actions within the program.

action a attempts to perform action a . If it isn't permitted by the context, **action** a always fails. If it is permitted, it may fail, be stalled or reduce to the action's return value all depending on the API. If it reduces, it will modify the world state. **test** $a m_t m_f$ allows a program to query the context in which it is being run. If action a is permitted in the current context then **test** $a m_t m_f$ executes m_t , otherwise it executes m_f . Usually m_f would be a sort of exception handler, returning a value indicating that certain actions were locked out by the current context. Programs of the form $m \gg= f$ behave in the normal monadic style: m is reduced continually until it is a value `return v` for some v , then $f v$ is reduced. If m at any stage fails or becomes stalled, the same will happen to $m \gg= f$. A program \perp always fails.

A **base program** refers to any program of the form `return v`, \perp , `return v >>= f`, **action** a , **test** $a m_t m_f$ or `return v_l |||_*^p return v_r`. The behaviour of these programs is always entirely deterministic for a given context and world. It is concurrency on its own which introduces non-determinism. $m_l |||_*^p m_r$, executed in context c , runs m_l and m_r in parallel in contexts c_l and c_r respectively, where $\mathbf{pf} p c = (c_l, c_r)$. If two concurrent sub-programs can either reduce or fail, then either side may be chosen arbitrarily. This continues until

- one side fails, causing both programs in parallel to fail.
- m_l and m_r become values `return v_l` and `return v_r` respectively,

in which case the parallel execution terminates, becoming the value `return $v_l * v_r$` .

- one side becomes stalled and the other side converges to normal form (that is: it is either a value or also stalled), causing the concurrent execution of both programs to be stalled.

This non-deterministic single-step form of semantics is the most “obviously correct” way of describing concurrency. The side-condition on the parallel reduction rules that $m_l \neq \perp$ and $m_r \neq \perp$ is undesirable but hard to avoid. If we check at the outset whether both the left and right sub-programs are values it makes the implementation *much* simpler. Unfortunately this means forcing both m_l and m_r to an outermost constructor.

3.4 Convergence/Divergence and the Implementation

Figure 8 contains information about how the language is implemented in the metalanguage.

Programs in CURIO are elements of a higher-order algebraic type. The use of an algebraic type is necessary since we need to be precise about how there are exactly five ways of constructing a program. This means that for our machine-verified proofs the types of `>>=` and `par` become monomorphic. This isn’t a serious problem. The property of having only five constructors can be informally guaranteed in a real language using class and module interfaces, and it is not necessary for a program to be able to query another program’s outermost constructor, which is something that algebraic types specifically allow.

The function `nexts` implements single-step reduction, and non-determinism is implemented by supplying it with an additional parameter of type `Guess`. This is, roughly, a stream of Boolean values and it guides the reduction algorithm’s search for a redex in the presence of concurrency. We existentially quantify over its values to show that reducing to a particular reduct is possible. If a world/program pair is in normal form in some context then `nexts` returns `CONVERGED`. Otherwise it returns `REDUCT (w_1, m_1)`, for some w_1 , m_1 , or fails (\perp). These three possibilities define the three single-step reduction relations, \downarrow , \longrightarrow and \uparrow .

To make the move from a single-step (or reduction) semantics to a big-step (or evaluation) semantics we must investigate the repeated single-step reduction of a world/program pair. The single-step semantics is non-deterministic, so non-determinism must also have some presence in a big-step semantics. First off, one must define a non-deterministic evaluation relation. $w \Vdash m \Downarrow^c w' \Vdash m'$ means “ $w \Vdash m$ can, after zero or more single-step reductions in context c yield $w' \Vdash m'$, which is in normal form.”

This is a rather weak property. It would be nice if we knew that $w \Vdash m$ would *always* reduce to a normal form $w' \Vdash m'$ in some context c . This

$$\begin{aligned}
\text{Prog } \nu \alpha \rho &\triangleq \text{ BIND } (\text{Prog } \nu \alpha \rho) (\nu \rightarrow \text{Prog } \nu \alpha \rho) \\
&| \text{ RET } \nu \\
&| \text{ ACTION } \alpha \\
&| \text{ TEST } (\text{Prog } \nu \alpha \rho) (\text{Prog } \nu \alpha \rho) \\
&| \text{ PAR } \rho (\text{Prog } \nu \alpha \rho) (\text{Prog } \nu \alpha \rho) (\nu \rightarrow \nu \rightarrow \nu)
\end{aligned}$$

$$\begin{aligned}
\text{Reduction } \beta &\triangleq \text{ REDUCT } \beta \mid \text{ CONVERGED} \\
\text{next}_s &:: \text{ Guess } \rightarrow \varsigma \rightarrow (\omega, \text{Prog } \nu \alpha \rho) \rightarrow \text{Reduction } (\omega, \text{Prog } \nu \alpha \rho) \\
\text{rdce}_s &:: \text{ Nat } \rightarrow [\text{Guess}] \rightarrow \varsigma \rightarrow (\omega, \text{Prog } \nu \alpha \rho) \rightarrow (\omega, \text{Prog } \nu \alpha \rho)
\end{aligned}$$

$$\begin{aligned}
&\text{rdce}_s (i + 1) [g:gs] c (w, m) = (w'', m'') \\
&\iff \\
&\left(\exists_{w' \in \omega} \cdot \exists_{m' \in \text{Prog } \nu \alpha \rho} \cdot \begin{array}{l} \text{next}_s g c (w, m) = \text{REDUCT } (w', m') \wedge \\ \text{rdce}_s i gs c (w', m') = (w'', m'') \end{array} \right)
\end{aligned}$$

$$\text{rdce}_s 0 gs c (w, m) = (w, m)$$

$$w \Vdash m \longrightarrow^c w' \Vdash m' \triangleq \exists_{g \in \text{Guess}} \cdot \text{next}_s g c (w, m) = \text{REDUCT } (w', m')$$

$$w \Vdash m \downarrow^c \triangleq \exists_{g \in \text{Guess}} \cdot \text{next}_s g c (w, m) = \text{CONVERGED}$$

$$w \Vdash m \uparrow^c \triangleq \exists_{g \in \text{Guess}} \cdot \text{next}_s g c (w, m) = \perp$$

$$w \Vdash m \overset{i}{\longrightarrow}^c w' \Vdash m' \triangleq \exists_{gs \in [\text{Guess}]} \cdot \text{rdce}_s i gs c (w, m) = (w', m')$$

$$w \Vdash m \longrightarrow^c w' \Vdash m' \triangleq \exists_{i \in \mathbb{N}} \cdot w \Vdash m \overset{i}{\longrightarrow}^c w' \Vdash m'$$

$$w \Vdash m \overset{i}{\downarrow}^c w' \Vdash m' \triangleq w \Vdash m \overset{i}{\longrightarrow}^c w' \Vdash m' \wedge w' \Vdash m' \downarrow^c$$

$$w \Vdash m \downarrow^c w' \Vdash m' \triangleq \exists_{i \in \mathbb{N}} \cdot w \Vdash m \overset{i}{\downarrow}^c w' \Vdash m'$$

$$\begin{aligned}
w \Vdash m \overset{i}{\downarrow}^c w' \Vdash m' &\triangleq \forall_{gs \in [\text{Guess}]} \cdot \text{rdce}_s i gs c (w, m) = (w', m') \\
&\wedge w' \Vdash m' \downarrow^c
\end{aligned}$$

$$w \Vdash m \downarrow^c w' \Vdash m' \triangleq \exists_{i \in \mathbb{N}} \cdot w \Vdash m \overset{i}{\downarrow}^c w' \Vdash m'$$

$$w \Vdash m \uparrow^c \triangleq \neg \exists_{w' \in \omega} \cdot \exists_{m' \in \text{Prog } \nu \alpha \rho} \cdot w \Vdash m \downarrow^c w' \Vdash m'$$

Figure 8: Implementation Details

is expressed as $w \Vdash m \Downarrow^c w' \Vdash m'$. Divergence in CURIO, expressed as $w \Vdash m \Uparrow^c$, means that a program can under no circumstances reduce to some normal form in context c . The relationship between \uparrow and \Uparrow is subtle. The former is failure in the denotational semantics of the metalanguage, the latter is failure in the operational semantics of our language. Theorem 4.2 in Section 4 states that if PRE_s , then $w \Vdash m \Uparrow^c$ implies $w \Vdash m \Uparrow$.

The relation \Downarrow is implemented with rdce_s and expresses the repeated single-step reduction of a world/program pair until it is in normal form. It requires a list, or stream of **Guesses**, one for each single-step reduction (we don't mention any boundary or definedness conditions, but they do exist.) Convergence, \Downarrow , differs to \Downarrow in that it universally quantifies over the guesses. It is trivially true that $w \Vdash m \Downarrow^c w' \Vdash m'$ implies $w \Vdash m \Downarrow^c w' \Vdash m'$ and the confluence proof shows that if PRE_s holds, $w \Vdash m \Downarrow^c w' \Vdash m'$ implies $w \Vdash m \Downarrow^c w' \Vdash m'$. Sometimes \Downarrow and \Downarrow are annotated with a number which denotes how many reduction steps took place.

Proposition 3.1. *If $w \Vdash m \Downarrow^c w_1 \Vdash m_1$ then it is not the case that $w \Vdash m \Uparrow^c$.*

Proof. Direct from the definition of \Uparrow . \square

Proposition 3.2. *Each of the following statements imply those below it:*

1. $w \Vdash m \Downarrow^c w_1 \Vdash m_1$
2. $w \Vdash m \Downarrow^c w_1 \Vdash m_1$
3. $w \Vdash m \twoheadrightarrow^c w_1 \Vdash m_1$

Proof. (1) implies (2) since if a property holds for all **Guesses** it must hold for some **Guess**. (2) implies (3) because the only extra condition on the former is that $w_1 \Vdash m_1 \Downarrow^c$ is in normal form. \square

It should also be noted that although it's an implementation in a real language, we have no interest at the moment in its efficiency.

3.5 Convergence is Recursively Enumerable

It is reassuring to note that although reduction etc. are defined in the form of an existential quantification over reduction steps, this is just for convenience.

We now prove separately that \Downarrow is recursively enumerable. This is a (machine-verified) proof that there exists a function

$$\text{run}_s :: [\text{Guess}] \rightarrow \varsigma \rightarrow (\omega, \text{Prog } \nu \alpha \rho) \rightarrow (\omega, \text{Prog } \nu \alpha \rho)$$

such that

$$\begin{aligned} \mathbf{run}_s \text{ } gsc (w, m) &= (w_1, m_1) \\ &\iff \\ \exists_{i \in \mathbb{N}}. \mathbf{rdce}_s \text{ } i \text{ } gsc (w, m) &= (w_1, m_1) \wedge w_1 \Vdash m_1 \downarrow^c \end{aligned}$$

Once this is proved, convergence and divergence can be expressed in a more intuitive manner as follows:

$$\begin{aligned} w \Vdash m \not\Downarrow^c w_1 \Vdash m_1 &\iff \exists_{gs \in [\text{Guess}]} . \mathbf{run}_s \text{ } gsc (w, m) = (w_1, m_1) \\ w \Vdash m \Downarrow^c w_1 \Vdash m_1 &\iff \forall_{gs \in [\text{Guess}]} . \mathbf{run}_s \text{ } gsc (w, m) = (w_1, m_1) \\ w \Vdash m \Uparrow^c &\iff \forall_{gs \in [\text{Guess}]} . \mathbf{run}_s \text{ } gsc (w, m) = \perp \end{aligned}$$

Showing that this holds is a non-trivial task. If one writes an implementation which just repeatedly applies \mathbf{next}_s , perhaps an infinite number of times, then there is no structure to induct over. We must build an implementation which internally constructs an intermediate list.

The proof relies on a more general (and quite powerful) lemma. Consider the following two functions:

$$\begin{aligned} \mathbf{iterate} &:: (\beta \rightarrow \beta) \rightarrow \beta \rightarrow [\beta] \\ \mathbf{iterate} \text{ } f \text{ } x &\stackrel{\Delta}{=} x : \mathbf{iterate} \text{ } f \text{ } (f \text{ } x) \\ \mathbf{downtil} &:: (\beta \rightarrow \beta) \rightarrow (\beta \rightarrow \text{Bool}) \rightarrow \beta \rightarrow [\beta] \\ \mathbf{downtil} \text{ } f \text{ } p \text{ } x &\stackrel{\Delta}{=} x : \mathbf{if} (p \text{ } x) \text{ then } [] \text{ else } (\mathbf{downtil} \text{ } f \text{ } p \text{ } (f \text{ } x)) \end{aligned}$$

$\mathbf{iterate}$ is a standard Haskell library function. The term $\mathbf{iterate} \text{ } f \text{ } x$ creates an infinite list, each successive element of which contains the next iteration of function f to an initial value x . $\mathbf{downtil}$ is somewhat similar except there is an extra computable predicate $p :: \beta \rightarrow \text{Bool}$ which indicates whether the iteration of f should stop. Therefore $\mathbf{downtil}$ may or may not return an infinite list.

The following lemma gives a useful relationship between the two. It shows that given side-conditions relating p and f , some i applications of f to x_0 yields an x_1 such that $p \text{ } x_1$ if and only if $\mathbf{last} (\mathbf{downtil} \text{ } f \text{ } p \text{ } x_0) = x_1$ and x_1 is defined.

Lemma 3.1.

$$\begin{aligned} \forall_{f \in \beta \rightarrow \beta}. \forall_{p \in \beta \rightarrow \text{Bool}}. p \perp = \perp \wedge f \perp = \perp \wedge (\forall_{x' \in \beta}. p \text{ } x' \neq \text{FALSE} \implies f \text{ } x' = \perp) \\ \implies \forall_{x_0 \in \beta}. \forall_{x_1 \in \beta}. \\ (\exists_{i \in \mathbb{N}}. x_1 = \mathbf{iterate} \text{ } f \text{ } x_0 \text{ } !! \text{ } i \wedge p \text{ } x_1) \\ \iff \\ (x_1 = \mathbf{last} (\mathbf{downtil} \text{ } f \text{ } p \text{ } x_0) \wedge x_1 \neq \perp) \end{aligned}$$

Proof. The proof needs quite a few extra lemmas, and on the whole requires an extremely careful treatment of non-termination. In particular, with the given side-conditions

- If $f x = \perp$ then `iterate f x` = $[x, \perp, \perp, \perp, \dots]$.
- If $p x = \perp$ then `dountil f p x = x : \perp` and `iterate f x` = $[x, \perp, \perp, \perp, \dots]$.
- If $p x = \text{TRUE}$ then `dountil f p x` = $[x]$ and `iterate f x` = $[x, \perp, \perp, \perp, \dots]$.
- If $p x = \text{FALSE}$ and $f x = \perp$ then `dountil f p x = x : (\perp : \perp)`.

To prove the \implies direction we induct over i , the number of iterations required. We must show that if an iteration results in \perp then it cannot revert back to a non- \perp term and also prove that if $p x = \text{FALSE}$ then `last (dountil f p x)` = `last (dountil f p (f x))`.

To prove the \impliedby direction we must prove first that

- if $l = \text{length (dountil f p x)}$ and $l \neq \perp$ then $p (\text{iterate f x !! } (l - 1))$.
- if for all $i, 0 \leq i < k, p (\text{iterate f x !! } i) = \text{FALSE}$ then `dountil f p x !! k = iterate f x !! k`.
- various other results to do with `last`, `iterate` and infinite lists.

We can show that if `(dountil f p x)` is infinite then `last (dountil f p x)` will be \perp (thus proving a contradiction) and if `(dountil f p x)` is a specific finite length then i will be that length minus one. \square

To prove the final result we construct the function `nextWraps` which treats the state of the programs evolution as a 4-tuple containing (1) the world state, (2) the program, (3) the current fresh list of `Guess` and (4) a Boolean value indicating whether the previous iteration resulted in a world/program pair in normal form.

`nextWraps` :: $\varsigma \rightarrow (\omega, \text{Prog } \nu \alpha \rho, [\text{Guess}], \text{Bool}) \rightarrow (\omega, \text{Prog } \nu \alpha \rho, [\text{Guess}], \text{Bool})$

Now, with a suitable wrapper, the function

`last (dountil (nextWraps c) fth4 (w, m, gs, FALSE))`

forms the implementation of `runs gs c (w, m)`, where `fth4` returns the fourth element from a 4-tuple.

Proposition 3.3. *If $w \Vdash m \Downarrow^c w_1 \Vdash m_1$ and $w \Vdash m \Downarrow^c w_2 \Vdash m_2$ then it is true that $m_1 = m_2$ and $w_1 = w_2$*

Proof. Obvious, since we have shown that reduction to some normal form defines a function, not just a relation. \square

4 The Machine-Verified Confluence Proof

4.1 Introduction

A non-deterministic reduction system is said to be confluent if for a given term all possible reduction sequences eventually yield the same normal form, or no normal form at all. In this section we prove that reduction in CURIO is confluent when PRE_5 holds – that \Downarrow is equivalent to \Downarrow . This powerful property means that although our definition of concurrency is a natural, non-deterministic one involving arbitrary choices, this arbitrariness is contained and has no effect on the overall outcome. For us, this means that, for all reduction orders, a program will either always terminate with the same resultant world/program pair or always diverge.

That the confluence proof has been machine-verified is also, on its own, a relatively notable result. Confluence proofs for the λ -calculus have been machine-verified before (in Coq [Hue94], and Isabelle/HOL [Nip96]) but we have yet to see one in an LCF style. Perhaps there is a good reason for this – confluence proofs usually wouldn’t require one to prove properties about a *program*. This, however, is the approach we took. We prove that a simple *implementation* of CURIO is confluent.

Instructions on the actual implementation and how the machine-readable form of the proofs can be obtained is available in Appendix A.

Terminology

Confluence is also known as the “Church-Rosser” property after the authors of the original proof in 1935 for the λ -calculus [CR36]. Formal definitions of confluence tend to differ slightly depending on which texts are read.

Barendregt, in the standard reference text on the λ -calculus [Bar84], defines a reduction system to be confluent if it obeys the “diamond property”. This means that if $A \twoheadrightarrow B$ and $A \twoheadrightarrow C$ then there is some D such that $B \twoheadrightarrow D$ and $C \twoheadrightarrow D$, where \twoheadrightarrow is the reflexive, transitive closure of the reduction relation ‘ \longrightarrow ’. In Term Rewriting Systems [Ter03], however, the diamond property (and confluence) is defined to be the property “ $A \longrightarrow B$ and $A \longrightarrow C$ implies there is some D such that $B \longrightarrow D$ and $C \longrightarrow D$ ”. It is the latter definition of confluence and the diamond property which is closest to what we use, but because our operational semantics also has a notion of failure this muddies the water somewhat.

There is also a certain amount of confusion concerning the differences between (finite) reduction *sequences* and reduction *strategies*. The pure λ -calculus is confluent, yet certain reduction strategies may, for a given term, not result in a normal form when others do. In particular, given the term $(\lambda x.\lambda y.y)\Omega$, call-by-name reduction will find the normal form $(\lambda y.y)$ but call-by-value reduction will not terminate.

Our confluence proof is stronger. All reduction strategies will have the same effect. The intuitive reason why this is true is that unlike function application in the λ -calculus, our **par** construct is highly symmetric.

Overview

The full machine-verified proof of confluence is long and full of complicated details. Many of these relate to the propagation of \perp and the fact that induction over lazy structures must be admissible [Iga74]. We adopt a hybrid approach to describing the proof. We try to explain all the nitty-gritty technical problems encountered, while still never losing sight of the overall picture.

To try to give some structure to the proof, three important subsections of the proof each culminate with the proof of a key theorem:

- Theorem 4.1: If $c_l \diamond_{\mathfrak{s}} c_r$ and $\text{pre}_{\mathfrak{s}}$ then reduction of programs in context c_l and c_r , if possible in both contexts, is order-independent.
- Theorem 4.2: If $\text{PRE}_{\mathfrak{s}}$, then $w \Vdash m \uparrow^c$ implies $w \Vdash m \uparrow\uparrow^c$.
- Theorem 4.3 (Confluence): If $\text{PRE}_{\mathfrak{s}}$, then $w \Vdash m \Downarrow^c w' \Vdash m'$ implies $w \Vdash m \Downarrow^c w' \Vdash m'$.

4.2 Preliminaries

Non-deterministic single-step reduction “ \longrightarrow ” is a fine high-level notation for expressing how a world/program pair can change within a given context. Its downside, however, is that it hides some internal details of the reduction and its implementation, and these details are central to the machine-verified proof.

They include:

- The initial **Guess** which guides the search for a redex.
- The actions that were performed, if any.
- The whereabouts of the redex if one is eventually found.

The initial **Guess** is usually “hidden” by an existential quantification, but it is still explicit in the implementation. The other two are truly internal, however, and for this reason the implementation $\text{next}_{\mathfrak{s}}$ had to be rewritten as the interaction of three different functions. (Admissibility states that to obtain information about a lazy structure one must do so constructively and write a function which computes it. One is not able to just prove that it exists.)

$$\begin{aligned}
\text{nextR}_s &:: \text{Guess} \rightarrow \varsigma \rightarrow \omega \rightarrow \text{Prog } \nu \alpha \rho \rightarrow \text{Redex (Route, RxType } \alpha) \\
\text{advS}_s &:: \varsigma \rightarrow \text{Route} \rightarrow \text{Prog } \nu \alpha \rho \rightarrow \text{Prog } \nu \alpha \rho \\
\text{advA} &:: \nu \rightarrow \text{Route} \rightarrow \text{Prog } \nu \alpha \rho \rightarrow \text{Prog } \nu \alpha \rho \\
\text{next}'_s \ g \ c \ (w, m) &\triangleq \text{case (nextR}_s(g, c, w, m)) \text{ of} \\
&\quad \text{NOREDEX} \quad \quad \quad \rightarrow \text{CONVERGED} \\
&\quad \text{REDEX } (r, \text{SILENT}) \quad \rightarrow \text{REDUCT } (w, \text{advS}_s(c, r, m)) \\
&\quad \text{REDEX } (r, \text{ACTION } a) \quad \rightarrow \text{case (af } a \ w) \text{ of} \\
&\quad \quad (w_1, v) \rightarrow \text{REDUCT } (w_1, \text{advA}(v, r, m))
\end{aligned}$$

Figure 9: Definition of next'_s

Deconstructing next_s

The new definition of next_s and the types of the three new functions are in Figure 9.

The function nextR_s searches for a redex. If it finds one it indicates the **Route** to that redex and the action it will perform, if any. A **Route** is just a finite list of L/R values. When searching for a redex, each time a **par** is encountered the next element of the **Route** indicates whether to look to the left- or right-hand side.

If a redex doesn't perform an action we call it a **silent** redex. The function $\text{advA}(v, r, m)$ modifies the action at route r in program m by replacing it with the program **return** v . The program $\text{advS}_s(c, r, m)$ modifies the silent redex at route r in program m . The context information c is required so that a program of the form **test** $a \ m_1 \ m_2$ can determine which of the two programs must be executed.

We now prove that the two definitions of next_s are equivalent, thus allowing us to pick whichever is the more appropriate when proving a lemma.

Lemma 4.1. $\text{next}_s = \text{next}'_s$

Proof. A long but straightforward induction on program structure. We omit the details, but it is a proof that the three individual functions add up to the single original one. nextR_s does all the searching for a redex but never modifies either the world state or the program. The only reason it needs w at all is to check if an action is stalled. The four stages of the reduction of an action redex **action** a are (1) checking if the action is permitted (**ap** $c \ a = \text{TRUE}$), (2) checking if the action is stalled, (**wa** $a \ w = \text{FALSE}$) (3) performing the action (evaluating **af** $a \ w$ to some (w_1, v)) and (4) updating the program with the value v . Of these, the first two are performed by nextR_s , the third takes place in the “wiring” and the fourth is performed

$$\begin{aligned}
w \Vdash m \xrightarrow{g \mapsto r}^c_a w_1 \Vdash m_1 &\stackrel{\Delta}{=} \text{nextR}_s(g, c, w, m) = \text{REDEX}(r, \text{ACTION } a) \wedge \\
&\quad \text{af } a \ w = (w_1, v) \wedge m_1 = \text{advA}(v, r, m) \\
w \Vdash m \xrightarrow{g \mapsto r}^c_{\bullet} w_1 \Vdash m_1 &\stackrel{\Delta}{=} \text{nextR}_s(g, c, w, m) = \text{REDEX}(r, \text{SILENT}) \wedge \\
&\quad m_1 = \text{advS}_s(c, r, m) \wedge w = w_1 \\
w \Vdash m \xrightarrow{g \mapsto r} \uparrow_a^c &\stackrel{\Delta}{=} \text{nextR}_s(g, c, w, m) = \text{REDEX}(r, \text{ACTION } a) \wedge \\
&\quad \text{af } a \ w = \perp \\
w \Vdash m \xrightarrow{g \mapsto \perp} \uparrow^c &\stackrel{\Delta}{=} \text{nextR}_s(g, c, w, m) = \perp \\
w \Vdash m \downarrow^c &\stackrel{\Delta}{=} \text{nextR}_s(g, c, w, m) = \text{NOREDEX}
\end{aligned}$$

Figure 10: Annotating single-step reduction

by advA . With any silent redex, nextR_s just finds the redex returning the route, and advS_s modifies the program itself. \square

Annotating single-step reduction

Redexes are inherently slippery things to reason about since there is no obvious type or set which is isomorphic to them in any useful way. We therefore use *Guesses* to *quantify* over redexes when we’re looking for one and *Routes* to *identify* a redex when we have found one. The function from *Guesses* to redexes is onto; the function from redexes to *Routes* is one-to-one.

Guesses and *Routes* can be confusingly similar, at times. Although they are both implemented as lists of `Bool`, they are conceptually somewhat different. A *Guess* is best understood as always being infinite, and *Routes* as always being finite. So we can always successfully retrieve the head and tail of a *Guess*, whereas there is the possibility that a *Route* is empty. Also, we sometimes implicitly “cast” a *Route* to a *Guess* by padding the finite list to make it infinite.

The new annotated single-step reduction can be found in Figure 10. It should be clear that the new notation covers all possible cases and is a consistent extension to the old notation. If the guess, route or redex type is omitted this means there is an implicit existential quantification. $w \Vdash m \xrightarrow{c}_a w' \Vdash m'$ means there exists some *Guess* g and a route *Route* r such that $w \Vdash m \xrightarrow{g \mapsto r}^c_a w' \Vdash m'$. If an action a or \bullet is omitted from a reduction then we just don’t specify whether it was silent or performed an action.

Corollary 4.2. *If $w \Vdash m \xrightarrow{c}_{\bullet} w_1 \Vdash m_1$ then $w = w_1$.*

Proof. Immediate. \square

Corollary 4.3. *If $w \Vdash m \xrightarrow{c}_a w_1 \Vdash m_1$ then $\mathbf{af} a w = (w_1, v)$, for some v .*

Proof. Immediate. □

Corollary 4.4. *If $w \Vdash m \xrightarrow{g \mapsto r}_a \uparrow_a^c$ then $\mathbf{af} a w = \perp$.*

Proof. Immediate. □

A template for inductive proofs

A great many lemmas are proved by inducting over the recursive structure of $\mathbf{Prog} \nu \alpha \rho$. They are all, of course, different, but we can give a general shape to many of the proofs, and this will serve as a basic template.

Base programs (those of the form $\mathbf{return} v, \perp, \mathbf{return} v \gg= f, \mathbf{action} a, \mathbf{test} a m_t m_f$ or $\mathbf{return} v_l \parallel \parallel_*^p \mathbf{return} v_r$) are always entirely deterministic for a given context and world, and never requires any “deeper” knowledge, such as an inductive hypothesis. When inducting over $\mathbf{Prog} \nu \alpha \rho$, it is usually relatively easy to prove properties for these programs.

The behaviour of a program $m \gg= f$, where m isn't a value $\mathbf{return} v$, is solely determined by the behaviour of m . This is importance because when inducting over $\mathbf{Prog} \nu \alpha \rho$ no inductive hypothesis can be given for f . If m fails, diverges or is in normal form then the same will be true of $m \gg= f$.

Programs of the form $m_l \parallel \parallel_*^p m_r$ are usually the most troublesome to prove properties about. If m_l and m_r are both values then it is a base term. It always fails if $\mathbf{pf} p c = \perp, m_l = \perp$ or $m_r = \perp$, so we often just omit this simple case altogether. If neither of the above are true an inductive hypothesis will be needed for m_l or m_r (and occasionally both) and how they behave in their respective contexts. Proving lemmas inductively for programs like this is made easier if we can perform a simple case analysis on whether the left- or right-hand side was reduced. The three following lemmas show how by examining the resultant **Route** we can learn whether a redex lied on the left or right hand side. Doing case analysis on the **Guess** would not give this sort of information. This is our standard procedure for determining the side in which a reduction took place.

Lemma 4.5. *If $w \Vdash m_l \parallel \parallel_*^p m_r \xrightarrow{g \mapsto \square}_c w_1 \Vdash m_1$ then $\mathbf{pf} p c = (c_l, c_r)$ for some c_l, c_r and for some $v_l, v_r, m_l = \mathbf{return} v_l, m_r = \mathbf{return} v_r, w_1 = w$ and $m_1 = \mathbf{return} v_l * v_r$.*

Proof. Since the route is \square the redex cannot be within either m_l or m_r . Therefore m_l and m_r must be values. □

Lemma 4.6. *If m is not a value, $\mathbf{nextR}_s(g, c, w, m \gg= f) = \mathbf{nextR}_s(g, c, w, m)$.*

Proof. Immediate from implementation. □

Lemma 4.7. *If m is not a value, then $w \Vdash m \gg= f \xrightarrow{g \mapsto r} c w' \Vdash m' \gg= f$ if and only if $w \Vdash m \xrightarrow{g \mapsto r} c w' \Vdash m'$.*

Proof. Immediate from Lemma 4.6. \square

Lemma 4.8. *If $\text{nextR}_5([b:g], c, w, m_l \parallel \parallel_*^p m_r) = \text{REDEX}([L:r], x)$ then $\text{pf } p \ c = (c_l, c_r)$ for some c_l, c_r and $\text{nextR}_5(g, c_l, w, m_l) = \text{REDEX}(r, x)$.*

Proof. Regardless of the initial direction b of the **Guess**, if the resultant **Route** was L then the left-hand side will have been reduced. \square

Lemma 4.9. *If $\text{nextR}_5([b:g], c, w, m_l \parallel \parallel_*^p m_r) = \text{REDEX}([R:r], x)$ then $\text{pf } p \ c = (c_l, c_r)$ for some c_l, c_r and $\text{nextR}_5(g, c_r, w, m_r) = \text{REDEX}(r, x)$.*

Proof. Symmetric to proof of Lemma 4.9. \square

Lemma 4.10. *If $w \Vdash m_l \parallel \parallel_*^p m_r \xrightarrow{[b:g] \mapsto [L:r]} c w_1 \Vdash m_1$ then $\text{pf } p \ c = (c_l, c_r)$ for some c_l, c_r and there is a m'_l such that $w \Vdash m_l \xrightarrow{g \mapsto r} c_l w_1 \Vdash m'_l$ and $m_1 = m'_l \parallel \parallel_*^p m_r$.*

Proof. Immediate from Lemma 4.9. \square

Lemma 4.11. *If $w \Vdash m_l \parallel \parallel_*^p m_r \xrightarrow{[b:g] \mapsto [R:r]} c w_1 \Vdash m_1$ then $\text{pf } p \ c = (c_l, c_r)$ for some c_l, c_r and there is a m'_r such that $w \Vdash m_r \xrightarrow{g \mapsto r} c_r w_1 \Vdash m'_r$ and $m_1 = m_l \parallel \parallel_*^p m'_r$.*

Proof. Immediate from Lemma 4.9. \square

4.3 Initial results

Having established some new notation, in this subsection we are now ready to begin the confluence proof. We prove some important lemmas relating to how single-step reduction affects world state and what actions single-step reduction can perform.

Lemma 4.12. *If $\text{nextR}_5(g, c, w, m) = \text{NOREDEX}$ then for all g_1 , it is true that $\text{nextR}_5(g_1, c, w, m) = \text{NOREDEX}$.*

Proof. Induction on m . It is trivial that **return** v is always in normal form, and if $w \Vdash \text{action } a$ is in normal form it always will be. $m_l \parallel \parallel_*^p m_r$ is only in normal form if m_l and m_r are both in normal form and it's not the case that both are values. g_1 can be either $[L:g'_1]$ or $[R:g'_1]$, for some g'_1 , but in each case m_l and m_r will both be in normal form, regardless of g'_1 's value (IH), so the same will be true of $m_l \parallel \parallel_*^p m_r$. \square

Lemma 4.13. *If $w \Vdash m \downarrow^c$ then it is not the case that $w \Vdash m \longrightarrow^c w' \Vdash m'$ for some w', m' or that $w \Vdash m \uparrow^c$. (This is a proof of confluence for programs which require no reduction steps.)*

Proof. Immediate from Lemma 4.12. □

Lemma 4.14. *If $\text{nextR}_s(g, c, w, m) = \text{REDEX}(r, \text{ACTION } a)$ then $\text{wa } a \ w = \text{FALSE}$.*

Proof. Induction on m . The only base redex which isn't silent is **action** a , and if this reduces then $\text{wa } a \ w = \text{FALSE}$. This is true, by induction, for any reduction which performs an action. □

Lemma 4.15. *If $w \Vdash m \longrightarrow_a^c w' \Vdash m'$ then $\text{wa } a \ w = \text{FALSE}$.*

Proof. Immediate from Lemma 4.14. □

Lemma 4.16. *If pre_s and $\text{nextR}_s(g, c, w, m) = \text{REDEX}(r, \text{ACTION } a)$ then $\text{ap } c \ a = \text{TRUE}$.*

Proof. Induction on m . Trivial for **action** a , and true by contradiction for silent base redexes. If $m = m_1 \gg= f$ (m not a value), it is trivial from Lemma 4.6. If m is $m_l \mid \mid \mid_*^p m_r$ then do case analysis on the resultant route. If it is true that $\text{nextR}_s([b:g], c, w, m_l \mid \mid \mid_*^p m_r) = \text{REDEX}([L:r], x)$, then $\text{pf } p \ c = (c_l, c_r)$ for some c_l, c_r , and from Lemma 4.8, $\text{nextR}_s(g, c_l, w, m_l) = \text{REDEX}(r, x)$. From IH, conclude $\text{ap } c_l \ a = \text{TRUE}$, and by pre_s , $c_l \sqsubseteq_s c$, and therefore $\text{ap } c \ a = \text{TRUE}$. The proof for the right-hand side is symmetric. □

Lemma 4.17. *If pre_s and $w \Vdash m \longrightarrow_a^c w' \Vdash m'$ then $\text{ap } c \ a = \text{TRUE}$.*

Proof. Immediate from Lemma 4.16. □

The following lemmas relate successful single-step reduction of m on world w to reduction in a completely unrelated world w' . The key is that after reducing successfully we will have the correct route to that redex. The second time around this route is used as the new guess to guarantee that the same redex is actually found – some actions in m which were stalled in w (and therefore ignored) may have become unstalled in w' , and it necessary to supply the exact route to make sure that we never encounter these unstalled actions, or any other redex.

Lemma 4.18. *If $\text{nextR}_s(g, c, w, m) = \text{REDEX}(r, \text{SILENT})$ then for all w' , $\text{nextR}_s(g, c, w', m) = \text{REDEX}(r, \text{SILENT})$.*

Proof. Induction on m . Trivially true for any silent base redex, since they are deterministic. If $m = m_l \parallel \parallel_*^p m_r$ and $g = [b:g']$, then do case analysis on the route r . If $r = []$, it is a base redex. If $r = [L:r']$ then we can derive $\text{nextR}_s(g', c_l, w, m_l) = \text{REDEX}(r', \text{SILENT})$ with Lemma 4.8. With IH, prove $\text{nextR}_s(r', c_l, w', m_l) = \text{REDEX}(r', \text{SILENT})$, and therefore $\text{nextR}_s(r, c, w', m_l \parallel \parallel_*^p m_r) = \text{REDEX}(r, \text{SILENT})$. If the right-hand side is reduced, the proof is similar. Programs of the form $m_1 \gg= f$ are proved easily with induction. \square

Lemma 4.19. *If $w \Vdash m \xrightarrow{g \mapsto r}^c w \Vdash m'$ then for all w' , $w' \Vdash m \xrightarrow{r \mapsto r}^c w' \Vdash m'$.*

Proof. Immediate, from Lemma 4.18. \square

Lemma 4.20. *If $\text{nextR}_s(g, c, w, m) = \text{REDEX}(r, \text{ACTION } a)$ then for all w_1 , if $\text{wa } a \ w_1 = \text{FALSE}$ then $\text{nextR}_s(r, c, w_1, m) = \text{REDEX}(r, \text{ACTION } a)$.*

Proof. Induction on m . Similar to the proof of Lemma 4.18 except the only valid base program is **action** a – all the others reduce silently. Since the action was already performed successfully in the action's local context c_1 which won't have changed, we know that $\text{ap } c_1 \ a = \text{TRUE}$, and because we know $\text{wa } a \ w_1 = \text{FALSE}$ that action will definitely be returned as a legitimate, unstalled redex. \square

Lemma 4.21. *If $w \Vdash m \xrightarrow{g \mapsto r}^c_a w' \Vdash m'$, where $\text{af } a \ w = (w', v)$, then for all w_1 , if $\text{wa } a \ w_1 = \text{FALSE}$ and $\text{af } a \ w_1 = (w'_1, v)$, then $w_1 \Vdash m \xrightarrow{r \mapsto r}^c_a w'_1 \Vdash m'$.*

Proof. Immediate, using Lemma 4.20. \square

Lemma 4.22. *If $w \Vdash m \xrightarrow{g \mapsto r}^c_a w' \Vdash m'$, then for all w_1 , if $\text{wa } a \ w_1 = \text{FALSE}$ and $\text{af } a \ w_1 = \perp$, then $w_1 \Vdash m \xrightarrow{r \mapsto r} \uparrow_a^c$.*

Proof. Immediate, using Lemma 4.20. \square

Lemma 4.23. *If $w \Vdash m \xrightarrow{g \mapsto r}^c w' \Vdash m'$, then $w \Vdash m \xrightarrow{r \mapsto r}^c w' \Vdash m'$.*

Proof. If the reduction is silent, apply Lemma 4.19. If it performs an action a , apply Lemma 4.21 – we know from Lemma 4.15 that $\text{wa } a \ w = \text{FALSE}$. \square

Lemma 4.24. *If $w \Vdash m \xrightarrow{g \mapsto r} \uparrow_a^c$, then $w \Vdash m \xrightarrow{r \mapsto r} \uparrow_a^c$.*

Proof. Apply Lemma 4.22 – we know from Lemma 4.14 that $\text{wa } a \ w = \text{FALSE}$. \square

We can now prove that reduction in disjoint contexts is order independent.

Theorem 4.1. *Assuming pre_s and $c_l \diamond_s c_r$, then if both $w \Vdash m_l \xrightarrow{g_l \mapsto r_l}^{c_l} w_l \Vdash m'_l$ and $w \Vdash m_r \xrightarrow{g_r \mapsto r_r}^{c_r} w_r \Vdash m'_r$, then either*

- there exists some w_2 such that both

$$w_r \Vdash m_l \xrightarrow{r_l \mapsto r_l} c_l w_2 \Vdash m'_l$$

and

$$w_l \Vdash m_r \xrightarrow{r_r \mapsto r_r} c_r w_2 \Vdash m'_r$$

- or it is the case that both $w_r \Vdash m_l \xrightarrow{r_l \mapsto r_l} \uparrow^{c_l}$ and $w_l \Vdash m_r \xrightarrow{r_r \mapsto r_r} \uparrow^{c_r}$.

Proof. Case analysis on whether a reduction is silent.

- One is silent, say that of c_l : $w_l = w$, from Corollary 4.2, and both succeed. Let $w_2 = w_r$. Use Lemma 4.19 to prove that m_l still reduces to m'_l with world state w_r , and use Lemma 4.23 to prove that m_r will behave the same on world w with r_r as its guess instead of g_r . (The proof is symmetric if c_r 's redex is silent.)
- Both are actions, say a_l and a_r : This means (Corollary 4.3) that **af** $a_l w = (w_l, v_l)$ and **af** $a_r w = (w_r, v_r)$. From Lemma 4.17 we know **ap** $c_l a_l = \text{TRUE}$ and **ap** $c_r a_r = \text{TRUE}$, and from Lemma 4.15, we also know that **wa** $a_l w = \text{FALSE}$ and **wa** $a_r w = \text{FALSE}$. Now, since $c_l \diamond_{\mathfrak{s}} c_r$ holds this means $a_l \parallel_{\mathfrak{s}} a_r$, $\text{ally}_{\mathfrak{s}}(a_l, a_r)$ and $\text{ally}_{\mathfrak{s}}(a_r, a_l)$ are true. Using $\text{ally}_{\mathfrak{s}}(a_l, a_r)$ prove **wa** $a_r w_l = \text{FALSE}$, using $\text{ally}_{\mathfrak{s}}(a_r, a_l)$ prove **wa** $a_l w_r = \text{FALSE}$, and because $c_l \parallel_{\mathfrak{s}} c_r$, either:
 - There exists a w_2 such that **af** $a_l w_r = (w_2, v_l)$ and **af** $a_r w_l = (w_2, v_r)$. Apply Lemma 4.21 to both sides to prove that they both will succeed.
 - **af** $a_l w_r = \perp$ and **af** $a_r w_l = \perp$. Apply Lemma 4.22 to show they both fail.

□

4.4 Analysing Failure

The next important theorem is that given $\text{PRE}_{\mathfrak{s}}$, if $w \Vdash m \uparrow^c$ then $w \Vdash m \uparrow^c$. This is effectively a proof that denotational failure in the metalanguage always causes divergence in our language – if a single-step reduction can possibly fail then despite nondeterminism it is impossible that the program will ever converge to a normal form.

In previous subsections we defined two distinct types of failure:

- $w \Vdash m \xrightarrow{g \mapsto r} \uparrow_a^c$. A redex is found, action a , but the action a fails.
- $w \Vdash m \xrightarrow{g \mapsto \perp} \uparrow^c$. The search for a redex fails outright.

$$\begin{aligned}
\text{Tree } \beta &= \text{BRANCH } (\text{Tree } \beta) (\text{Tree } \beta) \mid \text{LEAF } \beta \\
\text{redexTree}_s &:: \varsigma \rightarrow \text{Prog } \nu \alpha \rho \rightarrow \text{Tree } (\text{Redex } (\text{Route}, (\varsigma, \text{RxType } \alpha))) \\
\text{redexList}_s &:: \text{Guess} \rightarrow \varsigma \rightarrow \text{Prog } \nu \alpha \rho \rightarrow [(\text{Redex } (\text{Route}, (\varsigma, \text{RxType } \alpha)))] \\
\text{check}_s &:: \omega \rightarrow \text{Redex } (\text{Route}, (\varsigma, \text{RxType } \alpha)) \rightarrow \text{Redex } (\text{Route}, \text{RxType } \alpha) \\
\text{shuffle} &:: \text{Guess} \rightarrow \text{Tree } \beta \rightarrow \text{Tree } \beta \\
\text{preorder} &:: \text{Tree } \beta \rightarrow [\beta] \\
\text{mapTree} &:: (\beta \rightarrow \gamma) \rightarrow \text{Tree } \beta \rightarrow \text{Tree } \gamma \\
\text{firstRx} &:: \text{Redex } \beta \rightarrow \text{Redex } \beta \rightarrow \text{Redex } \beta \\
\text{addL, addR} &:: \text{Redex } (\text{Route}, \beta) \rightarrow \text{Redex } (\text{Route}, \beta) \\
\\
\text{nextR}'_s(g, c, w, m) &\triangleq \text{firstRxW}_s w (\text{redexList}_s g c m) \\
\text{firstRxW}_s w &\triangleq \text{foldl firstRx NOREDEX} \circ \text{map } (\text{check}_s w) \\
\text{redexList}_s g c m &\triangleq \text{preorder } \$ \text{shuffle } g \$ \text{redexTree}_s c m
\end{aligned}$$

Figure 11: Definition of nextR_s

If a program fails in the first way, then proving it will diverge is relatively easy (Lemma 4.38). The second, however, is much more difficult since it forces us, like with next_s , to rewrite nextR_s , breaking the process of searching for a redex into more manageable chunks. This is required to find out exactly how and why searching for a redex might fail.

Deconstructing nextR_s

There are obvious tensions at play when trying to re-implement nextR_s . A tree structure perfectly describes the structure of redexes: a leaf node represents a definite redex or lack thereof, and a Branch indicates two concurrent programs, each of which may have their own redexes. Furthermore, the `Guess` used to coordinate the search for a redex is most “naturally” applied to a tree structure. On the other hand, we must search for a redex sequentially. A list is ideal for this, since each element in a list is indexed uniquely and sequentially by integers in a way that elements of a tree cannot be.

Our solution involved the construction of two temporary structures, a tree and a list, before the final redex is returned. By doing this we can then make use of existing theorems on lists and trees. An outline of the new function nextR'_s is given in Figure 11.

The first half is the function redexList_s . This constructs a lazy list of potential redexes and has three parts. $\text{redexTree}_s c m$ first builds a tree in which each leaf-node represents either `NOREDEX`, indicating that no

reduction can take place, or $\text{REDEX } (r, (c_1, x))$, indicating that there is a potential redex at Route r , where c_1 is the local context for that redex. x can be either SILENT or $\text{ACTION } a$ for some a . There are only two situations where NOREDEX will occur in the tree: (1) if m is a value at the top-level, and (2) if a sub-program of m is of the form $\text{return } v_l \mid \mid \mid_*^p m_r$ or $m_l \mid \mid \mid_*^p \text{return } v_r$ and m_r/m_l aren't values. The function $\text{shuffle } g$ then rearranges the tree according to $\text{Guess } g$ so that the left-to-right ordering of leaves respects the order in which nextR_s originally would have scanned for redexes. Pre-ordering this tree with preorder then results in a list of potential redexes in the correct order.

firstRxW_s is the second half and finds the first legitimate redex in this list. It has two parts. Mapping $\text{check}_s w$ across the list of potential redexes throws out any action redexes which (1) aren't permitted by their local context or (2) are permitted but are stalled in the current world w . Actions which aren't allowed result in \perp (nextR_s should only return valid actions – Lemma 4.16) but permitted actions which are currently stalled become NOREDEX . Having eliminated all spurious redexes from the list, we can simply scan the list looking for the first element which isn't NOREDEX . Folding firstRx across the list does exactly that.

The types of a few other internal functions are given in Figure 11. The function mapTree modifies each element of a tree with a given function and addL and addR prepend either a L or a R to the route contained within a redex.

It is useful to note that this re-implementation manages to separate the four arguments of nextR'_s into three different sequential transformations. The context c and program m are supplied to redexTree_s , the tree is shuffled with the $\text{Guess } g$, and the world value w is only needed by check_s to see if actions are stalled.

Unlike nextR_s , nextR'_s is not defined directly in a recursive manner. We therefore need to prove that a direct recursive relationship exists. The following lemmas are all required to show that the behaviour of nextR'_s for some program can be understood in terms of its behaviour for the sub-programs of that program.

Lemma 4.25.

- (i) $\text{length} \circ \text{preorder} = \text{length} \circ \text{preorder} \circ \text{shuffle } g$
- (ii) $\text{shuffle } g \circ \text{shuffle } g = \text{id}$
- (iii) $\text{shuffle } g_1 \circ \text{shuffle } g_2 = \text{shuffle } g_2 \circ \text{shuffle } g_1$
- (iv) $\text{shuffle } [L, L, \dots] = \text{id}$
- (v) $\text{shuffle } g \circ \text{mapTree } f = \text{mapTree } f \circ \text{shuffle } g$
- (vi) $\text{map } f \circ \text{preorder} = \text{preorder} \circ \text{mapTree } f$
- (vii) $(r1 \text{ 'firstRx' } r2) \text{ 'firstRx' } r3 = r1 \text{ 'firstRx' } (r2 \text{ 'firstRx' } r3)$
- (viii) $\text{firstRxW}_s w (r1+r2) = \text{firstRx } (\text{firstRxW}_s w r1) (\text{firstRxW}_s w r2)$

- (ix) $r = r \text{ 'firstRx' NOREDEX}$
- (x) $r = \text{NOREDEX 'firstRx' } r$
- (xi) $\text{check}_5 \omega \circ \text{addL} = \text{addL} \circ \text{check}_5 \omega$
- (xii) $\text{check}_5 \omega \circ \text{addR} = \text{addR} \circ \text{check}_5 \omega$
- (xiii) $\text{firstRxW}_5 w \circ \text{map addL} = \text{addL} \circ \text{firstRxW}_5 w$
- (xiv) $\text{firstRxW}_5 w \circ \text{map addR} = \text{addR} \circ \text{firstRxW}_5 w$

Proof. All by straightforward induction or case analysis. □

Result (i) in Lemma 4.25 is a proof that shuffling a redex tree does not change the number of elements in that tree.

The results (ii)-(iv) in Lemma 4.25 show that `shuffle g`, for all g , forms the elements of an Abelian (or commutative) group under the \circ operator. Function composition is always associative, the identity is `shuffle [L, L, ...]` and each element is its own inverse.

Results (vii), (ix) and (x) show that `firstRx` is a monoidal operator with identity `NOREDEX`. The final four results show that `firstRxW5` and `check` do not query the route at which a redex exists. Instead it is passed through these functions unchanged.

Lemma 4.26. *If m is not a value then $\text{redexList}_5 g c (m \gg= f) = \text{redexList}_5 g c m$.*

Proof. Immediate. □

Lemma 4.27. *If m_l, m_r are not both values and $\text{pf } p c = (c_l, c_r)$ then*

$$\begin{aligned} & \text{redexList}_5 [L:g] c (m_l \parallel_*^p m_r) = \\ & \text{map addL} (\text{redexList}_5 g c_l m_l) \text{ ++ map addR} (\text{redexList}_5 g c_r m_r) \end{aligned}$$

Proof. Equational.

$$\begin{aligned} & \text{redexList}_5 [L:g] c m_l \parallel_*^p m_r \\ & = (\text{redexList}_5, \text{defn.}) \\ & \text{preorder } \$ \text{shuffle } [L:g] \$ \text{redexTree}_5 c (m_l \parallel_*^p m_r) \\ & = (\text{redexTree}_5, \text{defn.}) \\ & \text{preorder } \$ \text{shuffle } [L:g] \$ \text{BRANCH} (\text{mapTree addL } \$ \text{redexTree}_5 c_l m_l) (\dots) \\ & = (\text{shuffle defn.}) \\ & \text{preorder } \$ \text{BRANCH} (\text{shuffle } g \$ \text{mapTree addL } \$ \text{redexTree}_5 c_l m_l) (\dots) \\ & = (\text{preorder defn.}) \\ & (\text{preorder } \$ \text{shuffle } g \$ \text{mapTree addL } \$ \text{redexTree}_5 c_l m_l) \text{ ++ } (\dots) \\ & = (\text{Lemma 4.25, (v)}) \\ & (\text{preorder } \$ \text{mapTree addL } \$ \text{shuffle } g \$ \text{redexTree}_5 c_l m_l) \text{ ++ } (\dots) \\ & = (\text{Lemma 4.25, (vi)}) \\ & (\text{map addL } \$ \text{preorder } \$ \text{shuffle } g \$ \text{redexTree}_5 c_l m_l) \text{ ++ } (\dots) \\ & = (\text{redexList}_5, \text{defn.}) \\ & \text{map addL} (\text{redexList}_5 g c_l m_l) \text{ ++ map addR} (\text{redexList}_5 g c_r m_r) \end{aligned}$$

□

Lemma 4.28. *If m_l, m_r are not both values and $\text{pf } p \ c = (c_l, c_r)$ then*

$$\begin{aligned} & \text{redexList}_s [R:g] \ c \ (m_l \ || \ |^p_* \ m_r) = \\ & \text{map } \text{addR} \ (\text{redexList}_s \ g \ c_r \ m_r) \ ++ \ \text{map } \text{addL} \ (\text{redexList}_s \ g \ c_l \ m_l) \end{aligned}$$

Proof. Symmetric to that of Lemma 4.27. \square

Lemma 4.29. *If m is not a value then $\text{nextR}'_s(g, c, w, m \gg= f) = \text{nextR}'_s(g, c, w, m)$*

Proof. Immediate from Lemma 4.26 \square

Lemma 4.30. *If $\text{pf } p \ c = (c_l, c_r)$ and m_l, m_r are not both values then*

$$\begin{aligned} & \text{nextR}'_s([L:g], c, w, m_l \ || \ |^p_* \ m_r) = \\ & \text{firstRx} \ (\text{addL} \ (\text{nextR}'_s(g, c_l, w, m_l))) \ (\text{addR} \ (\text{nextR}'_s(g, c_r, w, m_r))) \end{aligned}$$

This states that if the the initial Guess says to look on the left-hand side first, then do

Proof. Equational.

$$\begin{aligned} & \text{nextR}'_s([L:g], c, w, m_l \ || \ |^p_* \ m_r) \\ & = (\text{nextR}'_s \ \text{defn.}) \\ & \text{firstRxW}_s \ w \ (\text{redexList}_s [L:g] \ c \ (m_l \ || \ |^p_* \ m_r)) \\ & = (\text{Lemma 4.27}) \\ & \text{firstRxW}_s \ w \ (\text{map } \text{addL} \ (\text{redexList}_s \ g \ c_l \ m_l) \ ++ \ \text{map } \text{addR} \ (\text{redexList}_s \ g \ c_r \ m_r)) \\ & = (\text{Lemma } \text{firstRxW}_s \ w \ (rs_1 \ ++ \ rs_2)) \\ & \text{firstRx} \ (\text{firstRxW}_s \ w \ (\text{map } \text{addL} \ (\text{redexList}_s \ g \ c_l \ m_l))) \ (\dots) \\ & = (\text{Lemma 4.25, (xiii) and (xiv)}) \\ & \text{firstRx} \ (\text{addL} \ (\text{firstRxW}_s \ w \ (\text{redexList}_s \ g \ c_l \ m_l))) \ (\text{addR} \ (\dots)) \\ & = (\text{nextR}'_s \ \text{defn.}) \\ & \text{firstRx} \ (\text{addL} \ (\text{nextR}'_s(g, c_l, w, m_l))) \ (\text{addR} \ (\text{nextR}'_s(g, c_r, w, m_r))) \end{aligned}$$

\square

Lemma 4.31. *If $\text{pf } p \ c = (c_l, c_r)$ and m_l, m_r are not both values then*

$$\begin{aligned} & \text{nextR}'_s([R:g], c, w, m_l \ || \ |^p_* \ m_r) = \\ & \text{firstRx} \ (\text{addR} \ (\text{nextR}'_s(g, c_r, w, m_r))) \ (\text{addL} \ (\text{nextR}'_s(g, c_l, w, m_l))) \end{aligned}$$

Proof. Similar to proof of Lemma 4.30. \square

Lemma 4.32. $\text{nextR}_s = \text{nextR}'_s$.

Proof. Induction on program structure. All base terms are deterministic, so the Guess is irrelevant and it is a simple matter of showing that check_s and redexTree_s together behave the same as nextR_s . If the program is of the form $m \gg= f$, m not a value, then Lemma 4.6 and Lemma 4.26 together convert nextR_s and nextR'_s respectively to the form of the IH. If the program

$$\begin{aligned}
& \text{wf}_s : \varsigma \rightarrow (\text{Prog } \nu \alpha \rho) \rightarrow \mathbb{B} \\
& \text{wf}_s(c, m) \stackrel{\Delta}{=} \text{length}(\text{preorder}(\text{redexTree}_s c m)) \neq \perp \\
& \text{wf}_s(c, \text{return } v) \quad \text{wf}_s(c, \text{action } a) \quad \text{wf}_s(c, \text{test } a m_t m_f) \quad \neg \text{wf}_s(c, \perp) \\
& \text{wf}_s(c, m \gg= f) \iff \text{wf}_s(c, m) \\
& \text{if pf } p c = \perp, \text{ then } \neg \text{wf}_s(c, m_l \parallel_*^p m_r) \\
& \text{if pf } p c = (c_l, c_r), \text{ then } \text{wf}_s(c, m_l \parallel_*^p m_r) \iff (\text{wf}_s(c_l, m_l) \wedge \text{wf}_s(c_r, m_r))
\end{aligned}$$

Figure 12: Well-Formedness of Programs

is $m_l \parallel_*^p m_r$, m_l and m_r not both values, then depending on whether the `Guess` is of the form $[L:g]$ or $[R:g]$ use Lemma 4.27 or Lemma 4.28. With the IH for m_l and m_r one can then show that nextR'_s preserves (1) the order in which the implementation of nextR_s searches for redexes and (2) how the resultant route, if there is one, has either L or R prepended to it depending on the location of the redex. \square

Well-formedness of Programs

For the divergence proof we found it necessary to define a notion of a well-formed program with respect to some context. The purpose of this is to distinguish programs which fail as a result of their structure (programs which aren't well-formed) and those that fail because of actions and how they interact with world state (programs that are well-formed.) The key to the introduction of well-formedness is the proof that if a program is not well-formed then it will always diverge, *even though it's possible that it may never fail after a finite number of steps*. (Failure, as always, means immediate failure, \uparrow .)

If a program m is well-formed with respect to context c then this is written as $\text{wf}_s(c, m)$. The definition and some basic properties of well-formedness is given in Figure 12 – a program is well-formed if the list of potential redexes returned by redexTree_s is of defined length. If $\neg \text{wf}_s(c, m)$ then we know m is partial (contains \perp s), infinite or the context c was not split successfully in m . The properties are all simple consequences of the implementation, requiring just a few basic results such as

$$\text{length}(xs \# ys) = \perp \iff (\text{length } xs = \perp \vee \text{length } ys = \perp)$$

To explain why we need this new terminology one needs to understand why an attempt to prove the theorem relating failure to divergence

would fail without it. Divergence, \uparrow , expresses the property that for a given world/program pair no reduction sequence of finite length can result in convergence to some normal form. The way that one proves that a world/program pair diverges is to prove, by inducting over i , that if it converged to normal form after i steps one can show a contradiction.

It was originally thought that the following lemma would suffice: If PRE_5 and $w \Vdash m \longrightarrow^c w' \Vdash m'$, then $w \Vdash m \uparrow^c$ implies $w' \Vdash m' \uparrow^c$. In other words, if a program can possibly fail, then it cannot escape from that possibility by reducing to a different world/program pair. But this proof strategy doesn't hold for some programs which are badly-formed, and the following counter-example shows why: (to be run in model `lock` with any context; world state should be `TRUE` (the mutex is locked); `f` is irrelevant.)

```
let  waits = par () (action WAIT) waits f
in   par () (action UNLOCK) waits f
```

The program `waits` attempts to create an infinite number of processes each of which performs a single `WAIT` action. The entire program above places an `UNLOCK` in parallel with `waits`. If the mutex is locked and we first look for a redex on the right-hand side then it will always fail (\uparrow) – there are an infinite number of stalled actions within `waits` and the search for a redex simply won't terminate. After executing `(action UNLOCK)`, however, each of the infinite number of `WAITs` are released, and there will always be a redex, and thus no (immediate) failure can occur². By proving that badly-formed programs always diverge, we solve the problem with the `waits` counter-example. Since it always diverges, it is obviously true that if it can fail it will diverge.

It is worth noting that some partial or infinite programs *are* well-formed. Take the `reads` example in Section 3: `reads` itself isn't well-formed, but the program `(test READI reads (return (-1)))`, although “syntactically” infinite, *is* well-formed. It is also worth noting that well-formedness is a property of programs for a given context, not world/program pairs.

²One may ask, though, that if the single-step reduction function is given a `Guess` which tells it always to first try the right-hand side of a parallel reduction, it would also fail immediately, never observing an action. That might solve it for this example, but the problem runs much deeper. What's of real interest is that when we said it made the theorem unprovable, *we didn't mean it made it false!* The issue is that our denotational, domain theoretic model of computable programs contains denotations of non-existent programs - which is an example of the famous full abstraction problem [Ong95]. In the above example, to cause it to fail, the right-hand side must be constantly traversed. But what if the side with an infinite number of actions alternates from left to right? The domain of $\text{Prog } \nu \alpha \rho$ admits any infinite sequence of left/right, but, since an infinite list of `Bool` can encode any real number, and it's well known that certain real numbers are uncomputable, one cannot always compute a guess which would always pick the correct reduction path to guarantee failure. And the guess *must* be computable, because of the admissibility constraints mentioned.

4.5 Failure implies Divergence

We are now in a position to begin proving the following key theorem.

Theorem 4.2. *If PRE_s , then $w \Vdash m \uparrow^c$ implies $w \Vdash m \uparrow^c$.*

Proof. If $\neg \text{wf}_s(c, m)$, use Lemma 4.35. If $w \Vdash m \xrightarrow{g \rightarrow r} \uparrow_a^c$, use Lemma 4.38. Otherwise it must be true that $\text{wf}_s(c, m)$ and $w \Vdash m \xrightarrow{g \rightarrow \perp} \uparrow^c$, so use Lemma 4.51. \square

It's proof is effectively the merging of three separate lemmas, and these three lemmas are proved in the following three respective sub-subsections.

Badly-Formed Programs

We prove that badly-formed programs diverge by showing that a program cannot reduce out of a badly-formed state, and that no badly-formed program is in normal form.

Lemma 4.33. *If $w \Vdash m \downarrow^c$ then $\text{wf}_s(c, m)$.*

Proof. Induction on m . A value is well-formed, as is a stalled action (or any action, for that matter). The other base terms are not in normal form, so don't apply. This is easily proved by induction if $m = m' \gg= f$. When $m = m_l \mid \mid \mid_*^p m_r$ is in normal form, m_l and m_r must be as well. Therefore, by IH, $\text{wf}_s(c_l, m_l)$ and $\text{wf}_s(c_r, m_r)$ and since the concatenation of two finite lists is a finite list, $\text{wf}_s(c, m)$. \square

Lemma 4.34. *If $\neg \text{wf}_s(c, m)$ and $w \Vdash m \xrightarrow{c} w' \Vdash m'$, then $\neg \text{wf}_s(c, m')$.*

Proof. Induction on m .

- Base terms. The only base program which isn't well-formed is \perp , and it cannot reduce successfully, so doesn't apply.
- $m = m_1 \gg= f$: It must be true that $\neg \text{wf}_s(c, m_1)$ and $w \Vdash m_1 \xrightarrow{c} w'_1 \Vdash w'$ so, by IH, $\neg \text{wf}_s(c, m'_1)$. Therefore m'_1 cannot be a value, and $\neg \text{wf}_s(c, m_1 \gg= f)$.
- $m = m_l \mid \mid \mid_*^p m_r$: Since $\neg \text{wf}_s(c, m_l \mid \mid \mid_*^p m_r)$, that means at least one of the following is true: $\neg \text{wf}_s(c_l, m_l)$ or $\neg \text{wf}_s(c_r, m_r)$. Now, say m_l contains the reduced redex. This means $w \Vdash m_l \xrightarrow{c_l} w'_l \Vdash m'$ and $m' = m'_l \mid \mid \mid_*^p m_r$, for some m'_l . If it was m_r which was badly-formed then the resultant program will still be badly-formed. If it was m_l that was badly formed then, by IH, $\neg \text{wf}_s(c_l, m'_l)$ and the resultant program remains badly-formed. (The proof is symmetric if m_r was reduced, not m_l .)

□

Lemma 4.35. *If $\neg \text{wf}_s(c, m)$, then for all w , $w \Vdash m \uparrow^c$.*

Proof. Induction on the number of reduction steps it might take to reduce to normal form. Base case (0): from Lemma 4.33, since m is badly-formed it can't be in normal form. Inductive case: It can't converge to normal form after $i + 1$ steps because after one step it is still badly-formed (Lemma 4.34) and it can't converge after i steps (IH). □

Failure of Actions

We prove that the failure of an action implies $w \Vdash m \uparrow^c$ by showing that even if $w \Vdash m$ can also reduce successfully then it can never “escape” from the action that originally failed. This is the property we originally wanted to prove for all programs, before we showed a counter-example. That it is true for actions which fail is the result of the definition of $\|\|_s$. It preserves the fact that $\mathbf{af} a w = \perp$ after the world state is modified by another action.

Lemma 4.36. *Given $\text{pre}_s, c_l \diamond_s c_r$, then if $w \Vdash m_r \xrightarrow{g_r \mapsto r_r} \uparrow_{a_r}^{c_r}$ and $w \Vdash m_l \xrightarrow{c_l} w_l \Vdash m'_l$, then $w_l \Vdash m_r \xrightarrow{r_r \mapsto r_r} \uparrow_{a_r}^{c_r}$.*

Proof. We know that the action a_r failed: $\mathbf{af} a_r w = \perp$. If m_l reduces silently then $w_l = w$, and Lemma 4.24 can be used. If m_l performs some action a_l then, from Lemma 4.15, $\mathbf{wa} a_l w = \text{FALSE}$, from Lemma 4.17, $\mathbf{ap} c_l a = \text{TRUE}$, and from the definition, $\mathbf{af} a_l w = (w_l, v_l)$, for some v_l . With $c_l \diamond_s c_r$, then derive $a_l \|\|_s a_r$ and $\mathbf{ally}_s(a_l, a_r)$. Using these two facts we can prove $\mathbf{af} a_r w_l = \perp$ and $\mathbf{wa} a_r w_l = \text{FALSE}$ respectively. Apply Lemma 4.22 to show that a_r will fail when executed in world w_l . □

Lemma 4.37. *If $\text{PRE}_s, w \Vdash m \xrightarrow{g \mapsto r} \uparrow_a^c$ and $w \Vdash m \xrightarrow{c} w' \Vdash m'$, then $w' \Vdash m' \xrightarrow{g \mapsto r} \uparrow_a^c$.*

Proof. Induction on m . True by contradiction for all base programs – they are deterministic, and therefore cannot fail *and* successfully reduce. For $m_1 \gg= f$ (where m_1 isn't a value) apply IH. If m is of the form $m_l \|\|_*^p m_r$ (and m_l and m_r aren't both values) then do case analysis on whether the failure and the reduction were on the same side or different sides. If both failure and success are on the same side, apply the IH. If failure and success are on opposite sides, apply Lemma 4.36 to prove that the side that fails will still fail after a successful reduction on the opposite side. (Since reduction can succeed, $\mathbf{pf} p c = (c_l, c_r)$ for some c_l, c_r , and from PRE_s this means pre_s and $c_l \diamond_s c_r$). □

Lemma 4.38. *If PRE_s , then if $w \Vdash m \xrightarrow{g \mapsto r} \uparrow_a^c$, then $w \Vdash m \uparrow^c$.*

Proof. Induction on the number of reduction steps it might take to reach normal form. Base case (0): Since $w \Vdash m \uparrow^c$, by Lemma 4.13 it can't be in normal form. Inductive case: It can't reach normal form after $i + 1$ steps because, by Lemma 4.37, after one step it can still fail, and by IH it can't converge to normal form after i steps. \square

Well-formed failure of nextR_s

Finally, we must prove that, assuming PRE_s and $\text{wf}_s(c, m)$, if $w \Vdash m \xrightarrow{g} \perp \uparrow^c$ then $w \Vdash m \uparrow^c$. The structure to the proof is rather similar to our proof of Lemma 4.38 at a high level but there are considerably more technical details.

The well-formedness of m has a direct consequence: it means that there are a finite list of potential redexes, and we can therefore induct over the length of the list without admissibility constraints. One can therefore prove existential properties by induction, namely that if nextR_s results in \perp then it failed for some *specific* potential redex identified by an integer. Once we have an identifier for problematic redexes we can then prove admissible properties about a lazy structure like $\text{Prog } \nu \alpha \rho$.

This subsection of the proof is not pretty and requires us to get our hands dirty manipulating the elements of lists of redexes. (Our approach was influenced very much by the existence of other pre-proved theorems relating to lists.) Also, for the amount of effort required, its relevance is disappointingly small. As we show below, if a program fails in the above manner it is for one of two rather uninteresting reasons. We include the proof for absolute completeness and as a testament to the rigour a proof-assistant imposes on its user.

We begin by proving three lemmas relating lists of potential redexes of programs of the form $m_l \parallel \parallel_*^p m_r$ to that of m_l and m_r .

Lemma 4.39. *If $\text{pf } p \ c = (c_l, c_r)$, $\text{wf}_s(c, m_l \parallel \parallel_*^p m_r)$ and it is true that $\text{redexList}_s [b:g] \ c \ (m_l \parallel \parallel_*^p m_r) \ !! \ i = \text{REDEX } ([L:r], x)$ then for some i_l $\text{redexList}_s \ g \ c_l \ m_l \ !! \ i_l = \text{REDEX } (r, x)$.*

Proof. First, m_l and m_r cannot both be values since this would mean $\text{redexList}_s [b:g] \ c \ (m_l \parallel \parallel_*^p m_r) = [\text{REDEX } ([], (c, \text{SILENT}))]$, and $[L:r] \neq []$. Cases analysis on b :

- $b = L$: Apply Lemma 4.27. Because $\text{wf}_s(c, m_l \parallel \parallel_*^p m_r)$, there are a finite, defined number of potential redexes for both m_l and m_r . If $i < \text{length } (\text{redexList}_s \ g \ c_l \ m_l)$ then $\text{map addL } \$ \ \text{redexList}_s \ g \ c_l \ m_l \ !! \ i = \text{REDEX } ([L:r], x)$ and therefore $\text{redexList}_s \ g \ c_l \ m_l \ !! \ i = \text{REDEX } (r, x)$, so let $i_l = i$. If $i \geq \text{length } (\text{redexList}_s \ g \ c_l \ m_l)$ then it is true that $\text{map addR } \$ \ \text{redexList}_s \ g \ c_r \ m_r \ !! \ i = \text{REDEX } ([L:r], x)$, and this is a contradiction since addR cannot modify the resultant route so it becomes $[L:r]$.

- $b = R$: Apply Lemma 4.28. Similar to the above except m_r is searched first. If $i < \text{length}(\text{redexList}_5 g c_r m_r)$ then contradiction. Otherwise $i \geq \text{length}(\text{redexList}_5 g c_r m_r)$, so let $i_l = i - \text{length}(\text{redexList}_5 g c_r m_r)$.

□

Lemma 4.40. *If $\text{pf } p c = (c_l, c_r)$, $\text{wf}_5(c, m_l \parallel!^p_* m_r)$ and*

$$\text{redexList}_5 [b:g] c (m_l \parallel!^p_* m_r) !! i = \text{REDEX} ([R:r], x)$$

then for some i_r $\text{redexList}_5 g c_r m_r !! i_r = \text{REDEX} (r, x)$.

Proof. Similar to the proof of Lemma 4.39. □

Lemma 4.41. *If $\text{pf } p c = (c_l, c_r)$ and $\text{wf}_5(c, m_l \parallel!^p_* m_r)$ then*

$$\text{redexList}_5 g c (m_l \parallel!^p_* m_r) !! i = \text{REDEX} ([], x)$$

for some i if and only if m_l and m_r are both values.

Proof. If: By contradiction. If m_l and m_r were not both values then, depending on whether g is $[L:g']$ or $[R:g']$ for some g' , use either Lemma 4.27 or Lemma 4.28 to show that the resultant route must be either $[L:r']$ or $[R:r']$ for some r' . A route of value $[]$ is impossible.

Only if: both m_l and m_r are values, so by the implementation there is just one silent reduction at route $[]$. □

The next step is to identify how and why $w \Vdash m \xrightarrow{g \rightarrow \perp} \uparrow^c$. The following three lemmas together show that if $w \Vdash m \xrightarrow{g \rightarrow \perp} \uparrow^c$ then for some action a either

- $\text{ap } c a \neq \text{TRUE}$, or
- $\text{ap } c a = \text{TRUE}$ and $\text{wa } a w = \perp$

Lemma 4.42. *If a list rxs is finite and $\text{firstRxW}_5 w rxs = \perp$ then there exists some i , $i \geq 0$, $i < \text{length } rxs$ such that $\text{check}_5 w (rxs !! i) = \perp$.*

Proof. Induction on the length of rxs . If $rxs = []$, $\text{firstRxW}_5 w rxs \neq \perp$. If $rxs = [rx : rxs']$ then either $\text{check}_5 w rx = \perp$ (let $i = 0$), or $\text{firstRxW}_5 w rxs' = \perp$, in which case let $i = i' + 1$, where i' is the index from the IH. □

Lemma 4.43. *If $\text{wf}_5(c, m)$ and $w \Vdash m \xrightarrow{g \rightarrow \perp} \uparrow^c$ then there exists an integer i such that $(\text{redexList}_5 g c m) !! i = \text{REDEX} (r, (c_1, \text{ACTION } a))$ and $\text{check}_5 w (\text{REDEX} (r, (c_1, \text{ACTION } a))) = \perp$.*

Proof. $w \Vdash m \xrightarrow{g \rightarrow \perp} \uparrow^c$ is defined as $\text{nextR}_5(g, c, w, m) = \perp$, which is equivalent to $\text{firstRxW}_5 w (\text{redexList}_5 g c m) = \perp$. Since $\text{wf}_5(c, m)$ there are only a finite number of potential redexes, so apply Lemma 4.42 to prove that there is some i and such that $\text{check}_5 w (\text{redexList}_5 g c m !! i) = \perp$. $\text{redexList}_5 g c m !! i$ must be of the form $\text{REDEX } (r, (c_1, \text{ACTION } a))$ because check_5 never fails for silent redexes or NOREDEX . (There is also a separate and entirely uninteresting proof that redexList_5 cannot return a list with \perp as an element. We omit this entirely). \square

Lemma 4.44. $\text{check}_5 w (\text{REDEX } (r, (c, \text{ACTION } a))) = \perp$ if and only if either $\text{ap } c a \neq \text{TRUE}$ or $\text{ap } c a = \text{TRUE} \wedge \text{wa } a w = \perp$.

Proof. Immediate. \square

We temporarily replace any reference to divergence with properties of check_5 and redexList_5 and continue the proof in the style of previous subsections. Lemma 4.46, Lemma 4.47 and Lemma 4.48 are really just re-workings of Lemma 4.17, Lemma 4.36 and Lemma 4.37 respectively.

Lemma 4.45. If $\text{check}_5 w (\text{REDEX } (r, x)) = \perp$ then for all r_1 it is true that $\text{check}_5 w (\text{REDEX } (r_1, x)) = \perp$.

Proof. A direct consequence of the implementation. check_5 never examines the route r . \square

Lemma 4.46. If $\text{wf}_5(c, m)$ and pre_5 , then if $\text{redexList}_5 g c m !! i = \text{REDEX } (r, (c_1, \text{ACTION } a))$ then $\text{ap } c_1 a = \text{TRUE}$ implies $\text{ap } c a = \text{TRUE}$.

Proof. Induction on m . For base terms, only action a meets the precondition and $\text{ap } c a = \text{TRUE}$ implies $\text{ap } c a = \text{TRUE}$. If $m = m_1 \gg= f$ then apply Lemma 4.26 and IH. If $m = m_l \mid \mid \mid *_p^p m_r$ then apply either Lemma 4.39 or Lemma 4.40 depending on the route r . If $r = [L:r']$ we can prove that $\text{redexTree}_5 g' c_l m_l !! i_l = \text{REDEX } (r', (c_1, \text{ACTION } a))$ for some i_l . With IH, prove $\text{ap } c_l a = \text{TRUE}$ implies $\text{ap } c_l a = \text{TRUE}$ and with pre_5 prove that $c_l \sqsubseteq_5 c$ and therefore $\text{ap } c a = \text{TRUE}$. The proof for the right-hand side is similar. \square

Lemma 4.47. If pre_5 , $c_l \diamond_5 c_r$, $\text{wf}_5(c_r, m_r)$, $\text{redexTree}_5 g_r c_r m_r !! i = \text{REDEX } (r_r, (c_{r1}, \text{ACTION } a_r))$, $\text{check}_5 w (\text{REDEX } (r_r, (c_{r1}, \text{ACTION } a_r))) = \perp$ and $w \Vdash m_l \xrightarrow{c_l} w_l \Vdash m'_l$, then $\text{check}_5 w_l (\text{REDEX } (r_r, (c_{r1}, \text{ACTION } a_r))) = \perp$.

Proof. From Lemma 4.44, either $\text{ap } c_{r1} a_r \neq \text{TRUE}$ or $\text{ap } c_{r1} a_r = \text{TRUE} \wedge \text{wa } a_r w = \perp$. If it is the former then this will be unaffected by a reduced action, so regardless of w_l , check_5 will still fail. If $\text{ap } c_{r1} a_r = \text{TRUE}$ and $\text{wa } a_r w = \perp$, then apply Lemma 4.46 to prove $\text{ap } c a_r = \text{TRUE}$. Next examine the reduction of m_l . If it's silent, then $w_l = w$, so m_r can fail in

the same way. If reducing m_l performs an action a_l then apply Lemma 4.15 to prove $\mathbf{wa} a_l w = \text{FALSE}$ and, from Lemma 4.17, $\mathbf{ap} c_l a = \text{TRUE}$. With $c_l \diamond_s c_r$ derive $\mathbf{ally}_s(a_l, a_r)$, which guarantees that $\mathbf{wa} a_r w_l = \perp$. This, by Lemma 4.44, proves that \mathbf{check}_s will also fail with world w_l . \square

Lemma 4.48. *If $\mathbf{wf}_s(c, m)$ and PRE_s then*

$$\begin{aligned} \mathbf{redexTree}_s g_1 c m !! i &= \text{REDEX } (r_1, (c_1, \text{ACTION } a)) \\ \mathbf{check}_s w (\text{REDEX } (r_1, (c_1, \text{ACTION } a))) &= \perp \end{aligned}$$

and $w \Vdash m \xrightarrow{g_2 \mapsto r_2} c w' \Vdash m'$ implies

$$\mathbf{check}_s w' (\text{REDEX } (r_1, (c_1, \text{ACTION } a))) = \perp$$

Proof. Induction on m .

- True by contradiction for all base redexes, since, being deterministic, they cannot both reduce and fail to reduce.
- For $m = m_1 \gg= f$ (m_1 not a value), from Lemma 4.26

$$\mathbf{redexTree}_s g_1 c (m_1 \gg= f) = \mathbf{redexTree}_s g_1 c m_1$$

and from the language semantics we can derive $w \Vdash m_1 \xrightarrow{g_2 \mapsto r_2} c w' \Vdash m'_1$ and $m' = m'_1 \gg= f$ from $w \Vdash m_1 \gg= f \xrightarrow{g_2 \mapsto r_2} c w' \Vdash m'$. Apply IH.

- When $m = m_l |||_*^p m_r$, where m_l and m_r are not both values, perform case analysis on the routes r_1 and r_2 . Neither can be \square (Lemma 4.41, Lemma 4.5).

- If both redexes are on the same side, say the left-hand side, making $r_1 = [L:r'_1]$ and $r_2 = [L:r'_2]$, apply Lemma 4.39 and Lemma 4.10. With IH one can now prove that

$$\mathbf{check}_s w' (\text{REDEX } (r'_1, (c_1, \text{ACTION } a))) = \perp$$

Apply Lemma 4.45 to prove the same is true with route $[L:r'_1]$.

- If the redexes are on opposite sides then from $\mathbf{pf} p c = (c_l, c_r)$ and PRE_s we can prove \mathbf{pre}_s and $c_l \diamond_s c_r$ (and $c_r \diamond_s c_l$ by symmetry). Apply Lemma 4.47.

\square

We are finally in a position to “repackage” divergence. Lemma 4.49 shows that re-using a previous route as a **Guess** will mean the redex at that route will be the first to be chosen (a re-working of the proof of both Lemma 4.18 and Lemma 4.20). Lemma 4.50 uses this to tidy up Lemma 4.48, and Lemma 4.51 follows exactly the same procedure as Lemma 4.38 to prove failure implies divergence.

Lemma 4.49. *If $\text{wf}_s(c, m)$ and $\text{redexList}_s g c m !! i = \text{REDEX}(r, x)$ then $\text{redexList}_s r c m !! 0 = \text{REDEX}(r, x)$.*

Proof. Induction on m . For base terms, \perp is not well-formed and otherwise there is just one potential redex, the **Guess** is ignored and i must be 0 already. If $m = m_1 \gg f$ where m_1 is not a value, use Lemma 4.26 and IH. If $m = m_l |||_*^p m_r$, where m_l and m_r are not both values, then do case analysis on r . If $r = [L:r']$, apply Lemma 4.39 to prove that for some i_l , $\text{redexList}_s g' c_l m_l !! i_l = \text{REDEX}(r', x)$, where $g = [b:g']$. By IH, this implies $\text{redexList}_s r' c_l m_l !! 0 = \text{REDEX}(r', x)$, and by Lemma 4.27 this can be shown to imply $\text{redexList}_s [L:r'] c (m_l |||_*^p m_r) !! 0 = \text{REDEX}([L:r'], x)$. That is: if x is the first potential redex in m_l with guess r' then it will be the first potential redex in $m_l |||_*^p m_r$ with guess $[L:r']$, since it will be m_l that is searched first. If $r = [R:r']$ then the proof is similar. \square

Lemma 4.50. *If $\text{PRE}_s, \text{wf}_s(c, m)$, $w \Vdash m \xrightarrow{g \rightarrow \perp} \uparrow^c$ and $w \Vdash m \xrightarrow{c} w' \Vdash m'$, then for some g_1 $w' \Vdash m' \xrightarrow{g_1 \rightarrow \perp} \uparrow^c$.*

Proof. Apply Lemma 4.43 to prove that some i $\text{redexList}_s g c m !! i = \text{REDEX}(r, (c_1, \text{ACTION } a))$ and $\text{check}_s w \text{REDEX}(r, (c_1, \text{ACTION } a)) = \perp$. Use Lemma 4.48 to prove $\text{check}_s w' \text{REDEX}(r, (c_1, \text{ACTION } a)) = \perp$ (a can still cause failure in world w'). Apply 4.49 to prove $\text{redexList}_s r c m !! 0 = \text{REDEX}(r, (c_1, \text{ACTION } a))$ – if we chose **Guess** r the second time then the first redex to be chosen will be action a . It can be shown easily that this means $\text{nextR}_s(r, c, w', m') = \perp$, and therefore, letting $g_1 = r$, $w' \Vdash m' \xrightarrow{g_1 \rightarrow \perp} \uparrow^c$. \square

Lemma 4.51. *If PRE_s , then if $\text{wf}_s(c, m)$ and $w \Vdash m \xrightarrow{g \rightarrow \perp} \uparrow^c$, then $w \Vdash m \uparrow^c$.*

Proof. Induction on the number of reduction steps it might take to reach normal form. Base case (0): Since $w \Vdash m \uparrow^c$, by Lemma 4.13 it can't be in normal form. Inductive case: It can't converge to normal form after $i + 1$ steps because, by Lemma 4.50, after one step it can still fail, and by IH it can't converge to normal form after i steps. \square

4.6 Confluence of Reduction

The confluence result is finally proved in this subsection. As is standard for confluence proofs, this is achieved by proving a so-called ‘‘diamond property’’. This is a proof that if two different non-deterministic reductions are possible, then there exists a common reduct to which both sides can then reduce.

We assume for all of this sub-subsection that PRE_s holds.

Lemma 4.52. *If $w \Vdash m \stackrel{1}{\Downarrow^c} w' \Vdash m'$ (i.e. $w \Vdash m \longrightarrow^c w' \Vdash m'$ and $w' \Vdash m' \Downarrow^c$), then $w \Vdash m \stackrel{1}{\Downarrow^c} w' \Vdash m'$. (This is a proof of confluence for programs which require one reduction step.)*

Proof. Induction on m .

- All the base programs either can't reduce or reduce deterministically.
- With $m_1 \gg= f$ (m_1 not a value), we know that $w \Vdash m_1 \longrightarrow^c w' \Vdash m'_1$, $w' \Vdash m'_1 \Downarrow^c$ and $m' = m'_1 \gg= f$. Because m' is stalled, m'_1 can't be a value, so apply IH.
- With $m_l \parallel_*^p m_r$ first prove that neither side can fail (if one side did fail, then the whole program could fail, and, by Theorem 4.2, the whole program could never reduce to a normal form – but it does, after one step.) Next, prove that it's impossible for *both* sides to be successfully reduced (if both were, then by Theorem 4.1, after single-step reducing one particular side the program couldn't be in normal form – the opposite side could still either be reduced or fail.) So any two reductions must be both in m_l or both in m_r . Apply IH to prove that both reductions on that side will have the same outcome.

□

Lemma 4.53. *Diamond Property. This is the proof that if we know that $w \Vdash m \xrightarrow{g' \mapsto r'} w' \Vdash m'$ and $w \Vdash m \xrightarrow{g'' \mapsto r''} w'' \Vdash m''$ then either*

- *The same redex was reduced: $w' = w''$ and $m' = m''$.*
- *There is a common reduct: For some w''' and m''' , $w' \Vdash m' \xrightarrow{r'' \mapsto r'''} w''' \Vdash c'''$ and $w'' \Vdash m'' \xrightarrow{r' \mapsto r'''} w''' \Vdash c'''$.*
- *Both sides will diverge: $w' \Vdash m' \Uparrow^c$ and $w'' \Vdash m'' \Uparrow^c$.*

Actually, in the last case we prove something slightly more specific which always implies divergence. This because divergence is inadmissible, requiring existential quantification and negation. We prove instead that they either fail or reduce to a badly-formed program. Both of these imply divergence, and it is easy to prove that if a $\gg=$ or a par diverge in this special way, the whole program will also diverge in this way.

Proof. Induction on m .

- All base reduction rules either can't reduce at all (contradiction), or deterministically reduce to the same redex.

- $w \Vdash m_1 \gg = f \longrightarrow^c w' \Vdash m'_1 \gg = f$ and $w \Vdash m_1 \gg = f \longrightarrow^c w'' \Vdash m''_1 \gg = f$.
If m'_1 or m''_1 are stalled, then by Lemma 4.52, $w' = w''$ and $m'_1 = m''_1$. This case rules out either m'_1 or m''_1 being values, so the next reduction must be within m'_1 and m''_1 , and not the application of f to something. Apply IH.

- Both

$$w \Vdash m_l \parallel \parallel_*^p m_r \longrightarrow^c w' \Vdash m'_l \parallel \parallel_*^p m'_r$$

and

$$w \Vdash m_l \parallel \parallel_*^p m_r \longrightarrow^c w'' \Vdash m''_l \parallel \parallel_*^p m''_r$$

The two reductions may have reduced different sides, or the same side.

- If the same side was reduced, say that of m_l , use Lemma 4.52 to account for the possibility of m'_l or m''_l being stalled or values (which would cause reduction to flip to the other side the second time around.) If either side was, then both reductions would have been the same. Apply IH.
- If one reduction was in m_l and the other was in m_r , then apply Theorem 4.1. Either both reduction orders can then fail (which entails failure for the whole program and therefore divergence by Theorem 4.2), or both reduction orders can reduce once more to get the same result by reducing the opposite side to that initially reduced. However, if one side reduces to a program which is \perp (which is still a successful reduction), then depending on the ordering, the parallel execution may accidentally force the program's evaluation.

This is why we need divergence at all, and is a consequence of the side condition that $m_l \neq \perp$ and $m_r \neq \perp$ in the single-step semantics. As an example, consider the program

`par LEFTWR (action (WRITEI 6)) (return 2 >>= λ_.⊥) f`

The left-hand side reduces to `(return 0)` and the right-hand side reduces to \perp . Reducing the left side then the right side yields `par LEFTWR (return 0) ⊥ f`, a badly-formed program (which diverges), but reducing the right side and then the left will fail outright (which causes divergence.)

□

Theorem 4.3. *Confluence.* *If $w \Vdash m \stackrel{i}{\Downarrow}^c w_1 \Vdash m_1$ then $w \Vdash m \stackrel{i}{\Downarrow}^c w_1 \Vdash m_1$.*

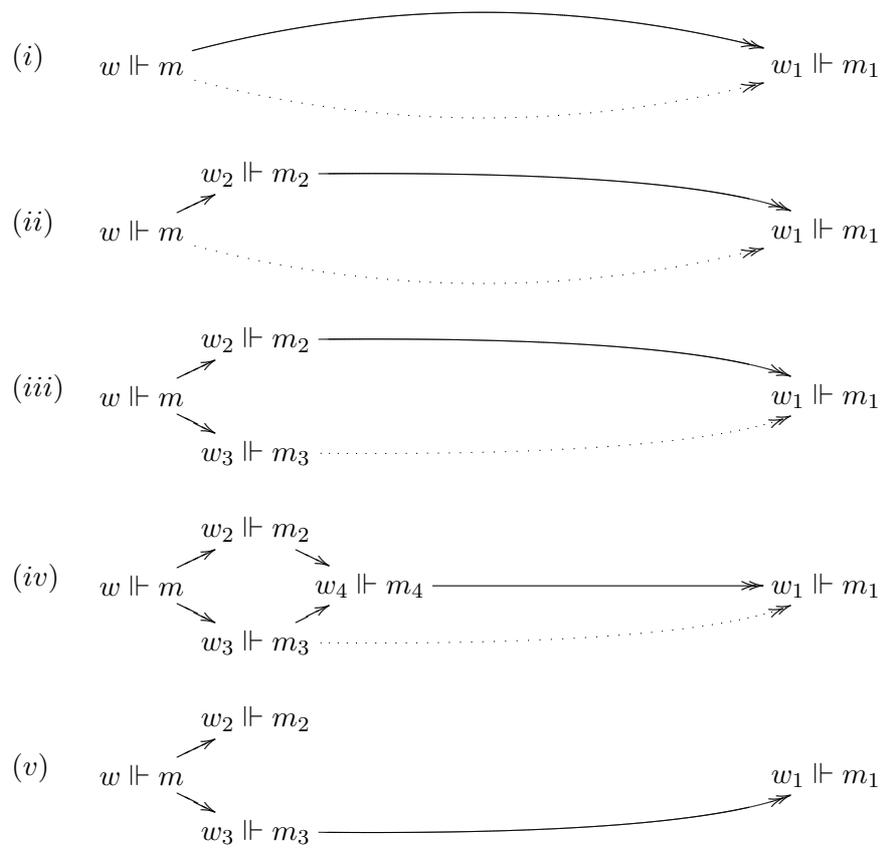


Figure 13: Proving Confluence using the Diamond Property

Proof. Induction on i . If $i = 0$, then apply Lemma 4.13. In the inductive case ($i = k + 1$), first prove for $k = 0$ using Lemma 4.52. Figure 13 demonstrates the sequence of steps required to prove the full inductive case. Diagram (i) is the initial proof obligation.

Assuming $k \geq 1$ (and $i > 1$), extract the first of the $k + 1$ reduction steps to yield $w \Vdash m \xrightarrow{c} w_2 \Vdash m_2$ and $w_2 \Vdash m_2 \overset{k}{\Downarrow^c} w_1 \Vdash m_1$, for some w_2, m_2 (diagram (ii)). To prove confluence, we must show that for any arbitrary reduction of $w \Vdash m$, it will eventually reduce to $w_1 \Vdash m_1$. First prove that despite the nondeterminism, $w \Vdash m$ cannot fail or be in normal form. (It can't be in normal form because it can reduce – Lemma 4.13; it can't fail because it does eventually converge to normal form – Theorem 4.2.) Therefore, the only alternative to reducing to $w_2 \Vdash m_2$ is that there is some other $w_3 \Vdash m_3$ such that $w \Vdash m \xrightarrow{c} w_3 \Vdash m_3$ (diagram (iii)). We must now show that $w_3 \Vdash m_3 \overset{k}{\Downarrow^c} w_1 \Vdash m_1$. Apply IH twice in a row: first in a forwards style to prove that we know $w_2 \Vdash m_2 \overset{k}{\Downarrow^c} w_1 \Vdash m_1$; second in a backwards style to show that we need only prove $w_3 \Vdash m_3 \overset{k}{\Downarrow^c} w_1 \Vdash m_1$. From the diamond property (Lemma 4.53), one of the following is true:

- $w_2 = w_3$ and $m_2 = m_3$: Trivial.
- There is a common reduct $w_4 \Vdash m_4$ (diagram (iv) in Figure 13) such that $w_2 \Vdash m_2 \xrightarrow{c} w_4 \Vdash m_4$ and $w_3 \Vdash m_3 \xrightarrow{c} w_4 \Vdash m_4$: From the confluence of reduction of $w_2 \Vdash m_2$, prove that $w_4 \Vdash m_4 \overset{k-1}{\Downarrow^c} w_1 \Vdash m_1$. Thus $w_3 \Vdash m_3 \overset{k}{\Downarrow^c} w_1 \Vdash m_1$.
- $w_2 \Vdash m_2 \Uparrow^c$ and $w_3 \Vdash m_3 \Uparrow^c$: Proof by contradiction (we know $w_2 \Vdash m_2$ can converge to normal form.)

□

Corollary 4.54. *Either $w \Vdash m \Uparrow^c$ or there exists a w', m' such that $w \Vdash m \Downarrow^c w' \Vdash m'$ (and not both.)*

Proof. The definition $w \Vdash m \Uparrow^c$ is that it's not the case that $w \Vdash m \Downarrow^c w' \Vdash m'$, and $w \Vdash m \Downarrow^c w' \Vdash m'$, by confluence, is equivalent to $w \Vdash m \Downarrow^c w' \Vdash m'$.

□

A Sparkle Proof Sections

The machine-readable form of the proofs may be obtained from the following URL:

http://www.cs.tcd.ie/research_groups/fmg/archive/CURIO-Proofs.html

It isn't always straightforward to simply list the Sparkle theorem which corresponds to each individual lemma. Almost all proofs which required induction over program structure needed to be proved in such a way that Sparkle could show it was admissible. These results – most of which have the letters `_adms` appended to their name – contain the guts of the proof. Usually a theorem was then proved (without the “adms”) which expressed the result in a more natural style. Also, many definedness results were omitted entirely in this document. These are uninteresting but still affect what was actually proved and our ability to form a direct link between presented lemmas and the Sparkle theorems.

Below is a select list of results in this document for which there are equivalent Sparkle theorems.

Result	Sparkle Section	Theorem Name
Lemma 2.1	<code>curio_examples</code>	<code>PRE_Buffer</code>
Lemma 2.2	<code>curio_examples</code>	<code>PRE_Lock</code>
Lemma 2.3	<code>curio_examples</code>	<code>PRE_IVar</code>
Lemma 3.1	<code>dountil</code>	<code>dountil_main_theorem</code>
R.E. Proof	<code>curio_rdce</code>	<code>rdce_run</code>
Lemma 4.1	<code>curio_prelude</code>	<code>next_separated</code>
Lemma 4.8	<code>curio_prelude</code>	<code>nextRedex_Par_False</code>
Lemma 4.9	<code>curio_prelude</code>	<code>nextRedex_Par_True</code>
Lemma 4.12	<code>curio_reduction</code>	<code>nextRedex_stalled_any_guess</code>
Lemma 4.14	<code>curio_reduction</code>	<code>nextRedex_wa</code>
Lemma 4.16	<code>curio_reduction</code>	<code>nextRedex_ap</code>
Lemma 4.18	<code>curio_reduction</code>	<code>nextRedex_interfere_Silent</code>
Lemma 4.20	<code>curio_reduction</code>	<code>nextRedex_interfere_Action</code>
Theorem 4.1	<code>curio_reduction</code>	<code>next_disjoint_with_nextRedex</code>
Lemma 4.25 (i)	<code>curio_failure_internals</code>	<code>length_preorder_shuffle</code>
Lemma 4.25 (v)	<code>curio_failure_internals</code>	<code>shuffle_mapTree</code>
Lemma 4.25 (vi)	<code>curio_failure_internals</code>	<code>preorder_mapTree</code>
Lemma 4.25 (viii)	<code>curio_failure_internals</code>	<code>firstRedex_++</code>
Lemma 4.32	<code>curio_failure</code>	<code>nextRedex_separated</code>
Theorem 4.2	<code>curio_failure</code>	<code>next_failure_divergence</code>
Lemma 4.34	<code>curio_failure</code>	<code>next_badlyformed_failure</code>
Lemma 4.35	<code>curio_failure</code>	<code>next_badlyformed_divergence</code>
Lemma 4.52	<code>curio_confluence</code>	<code>next_final_step</code>
Lemma 4.53	<code>curio_confluence</code>	<code>tidy_diamond</code>
Theorem 4.3	<code>curio_confluence</code>	<code>rdce_CONFLUENCE</code>

References

- [Bar84] H. P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, The Netherlands, Revised Edition, 1984.
- [CR36] A. Church and J. B. Rosser. Some properties of conversion. *Transactions of the American Mathematical Society*, 39:472–482, 1936.
- [DBvE05] Malcolm Dowse, Andrew Butterfield, and Marko van Eekelen. Reasoning about deterministic concurrent functional I/O. In Clemens Greck and Frank Huch, editors, *Proceedings of IFL 2004*, volume LNCS3474, pages 177–194. Springer-Verlag, 2005.
- [dMvEP01] Maarten de Mol, Marko van Eekelen, and Rinus Plasmeijer. Theorem proving for functional programmers. In Thomas Arts and Markus Mohnen, editors, *Proceedings of the 13th International Workshop, IFL2001*, number 2312 in LNCS, pages 55–71. Springer-Verlag, 2001.
- [Hue94] Gérard Huet. Residual theory in λ -calculus: A formal development. *Journal of Functional Programming*, 4(3):371–394, July 1994. Preliminary version available as INRIA Technical Report 2009, August 1993.
- [Iga74] Shigeru Igarashi. Admissibility of fixed-point induction in first-order logic of typed theories. In *Proceedings of the International Symposium on Theoretical Programming*, pages 344–383, London, UK, 1974. Springer-Verlag.
- [Nip96] Tobias Nipkow. More Church-Rosser proofs (in Isabelle/HOL). In M. McRobbie and J.K. Slaney, editors, *Automated Deduction — CADE-13*, volume 1104 of *LNCS*, pages 733–747. Springer, 1996.
- [Ong95] C.-H. L. Ong. Correspondence between operational and denotational semantics: the full abstraction problem for PCF. *Handbook of logic in computer science (vol. 4): semantic modelling*, pages 269–356, 1995.
- [Pau87] Lawrence C. Paulson. *Logic and computation: interactive proof with Cambridge LCF*. Cambridge University Press, New York, NY, USA, 1987.

- [Pey01] Simon Peyton Jones. Tackling the awkward squad – monadic input/output, concurrency, exceptions, and foreign language calls in Haskell. In CAR Hoare, M Broy, and R Stein-brueggen, editors, *Engineering theories of software construction, Marktoberdorf Summer School 2000*, pages 47–96. IOS Press, 2001.
- [PGF96] Simon Peyton Jones, Andrew Gordon, and Sigbjorn Finne. Concurrent Haskell. In ACM, editor, *POPL '96: Florida, 21–24 January 1996*, pages 295–308, New York, NY, USA, 1996. ACM Press.
- [PHWH03] Simon Peyton Jones, Paul Hudak, Philip Wadler, and John Hughes, editors. *Haskell 98 Language and Libraries, the Revised Report*. CUP, April 2003.
- [Plo81] Gordon Plotkin. A structural approach to operational semantics. Technical Report DAIMI FN-19, Computer Science Department, Aarhus University, Aarhus, Denmark, September 1981.
- [PvE01] Rinus Plasmeijer and Marko van Eekelen. Concurrent Clean version 2.0 language report. <http://www.cs.kun.nl/~clean/>, December 2001.
- [Ter03] Terese. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003.