The Functional Machine Calculus

Willem Heijltjes[1]

Department of Computer Science
University of Bath
Bath, United Kingdom

Abstract

This paper presents the Functional Machine Calculus (FMC) as a simple model of higher-order computation with "reader/writer" effects: higher-order mutable store, input/output, and probabilistic and non-deterministic computation.

The FMC derives from the lambda-calculus by taking the standard operational perspective of a call–by–name stack machine as primary, and introducing two natural generalizations. One, "locations", introduces multiple stacks, which each may represent an effect and so enable effect operators to be encoded into the abstraction and application constructs of the calculus. The second, "sequencing", is known from kappa-calculus and concatenative programming languages, and introduces the imperative notions of "skip" and "sequence". This enables the encoding of reduction strategies, including call–by–value lambda-calculus and monadic constructs.

The encoding of effects into generalized abstraction and application means that standard results from the lambda-calculus may carry over to effects. The main result is confluence, which is possible because encoded effects reduce algebraically rather than operationally. Reduction generates the familiar algebraic laws for state, and unlike in the monadic setting, reader/writer effects combine seamlessly. A system of simple types confers termination of the machine.

Keywords: lambda-calculus, computational effects, confluence, concatenative programming

1 Introduction

Higher-order programming and computational effects are ubiquitous in modern programs. Understanding them, and in particular their potent combination, is therefore an important challenge to computer science. Higher-order functional programming enjoys an elegant foundational theory in the λ-calculus, where β-reduction gives rise not only to operational semantics—by imposing an evaluation strategy—but also to an equational theory which may be regarded as definitive for higher-order functions. For computational effects, however, there are many approaches and, as yet, no single definitive theory. Such a theory would ideally include a convenient syntax, expressing a natural and convincing semantics, and supporting reasoning tools and methods such as type-systems and compile-time optimizations, while remaining amenable to refinement, extension and variation.

The rich history of approaches to the problem of computational effects in a higher-order setting includes Landin's pioneering work [17], which cemented the central position of λ-calculus, and highlighted the difficulty of reconciling the drive for a complete theory with the practice of programming: Landin focussed on a call-by-value strategy and used *thunks* to delay evaluation where necessary. A more modular and

---

[1] Email: wbh22@bath.ac.uk

flexible account was provided by Moggi's use of *monads* [25], which has influenced not only theoretical work but also the design of the Haskell programming language. However, a fundamental problem with monads is that they don't compose, and in practice multiple effects are combined by building a stack of monad transformers, which can become unwieldy for the programmer. Many alternatives and refinements have been proposed, including uniqueness types [32], continuations [10], encodings in (intuitionistic) linear logic [3,22] and in process calculi [23,13], premonoidal categories [31,30,20], Call–By–Push–Value [18] and its variants in linear logic [8,9], Arrows [14], algebraic effects [27,1], and effect handlers [29].

This paper offers a new solution to this challenge: the *Functional Machine Calculus* (FMC). It includes the effects of state, input/output, and probabilistic and non-deterministic computation—here referred to collectively as *reader/writer effects*. The FMC consists of two independent generalizations of the λ-calculus, *locations* and *sequencing*, that individually give two fragments, the *poly-λ-calculus* and the *sequential λ-calculus*. It enjoys a clean equational theory supported by a confluent reduction semantics, and expresses both effects and higher-order features in the same terms, retaining the simplicity of the λ-calculus while being powerful enough to capture the reality of higher-order programming with multiple effects. We provide operational semantics in terms of an abstract machine, and a type system that confers termination of the machine. The remainder of this section will introduce both generalizations. Throughout the paper, proofs are omitted when they are straightforward.

1.1   Locations

The main objective of this work has been to preserve confluence, following the recent presentation of a confluent probabilistic λ-calculus by Dal Lago, Guerrieri, and Heijltjes [4]. This is perhaps surprising, as λ-calculi with effects are known to be non-confluent. The apparent contradiction disappears by disentangling the *operational* and the *algebraic* aspects of evaluation. In λ-calculus, β-reduction is algebraic (or more precisely, β-*equivalence* is), while stack machines such as Krivine's [15] give an operational semantics. For effects, looking up a global variable or generating a random value is operational, while effect operators may interact algebraically via the laws of Plotkin and Power [27].

	global, operational	**local, algebraic**
λ-calculus	stack machines	β-reduction
effects	update, lookup, read, write, random	algebraic effect equations

Our starting point is the observation that for both the λ-calculus and reader/writer effects, the operational side can be given by push and pop actions on global stacks or streams. In a simple stack machine for the λ-calculus, *application* $M\,N$ pushes its argument N to the stack and continues as M, and *abstraction* $\lambda x.\,M$ pops a term N from the stack and binds it to x, to continue as $\{N/x\}M$ (the capture-avoiding substitution of N for x in M). Then, the following effects are also modelled via stacks or streams:

- reading from *input* is a pop from an input stream;
- writing to *output* is a push to an output stream;
- a memory cell c is modelled by a stack of depth one, where
  - *update* $c := N$ pops from c, discarding the value, then pushes the new value N;
  - *lookup* $!c$ pops the value N from c, pushes N to reinstate c, and then returns N;
- *probabilistic* and *non-deterministic* generators can be modelled as separate input streams.

This idea is captured in the *poly-λ-calculus*: we introduce a set of *locations* to represent independent stacks or streams on the machine, and parameterize abstraction and application in this set to act as pop and push actions on the corresponding stack. Effect operators are then encoded in these constructs according to the above scheme. Beta-reduction, generalized to multiple locations, remains confluent, and for encoded effect operators it gives rise to the expected algebraic laws [27]. However, this encoding of effects forces their call–by–name semantics, while programming with effects requires control over when they are called. This is the purpose of the second generalization, *sequencing*.

1.2 Sequencing

The literature offers several ways to control reduction behaviour in higher-order languages, including continuation encodings between cbv and cbn [28], and call–by–push–value (cbpv) [18] which encodes both. To complement *locations*, our approach takes the stack machine as primary. Viewing the λ-calculus as a language of machine instruction sequences, the *sequencing* generalization extends it with composition and the empty sequence, analogous to imperative "sequence" and "skip"; not by introducing these as primitives, but again by generalizing the calculus in a subtle way so that they arise naturally. Such designs have arisen several times before: in the first-order κ-calculus of Hasegawa [11], generalized to higher-order in the context of premonoidal categories [30]; as the Λ_s-calculus in an analysis of compilers [7]; and in higher-order stack programming languages, also called *concatenative* languages [12], such as *Joy* [36], λ-*FORTH* [21], *Cat* [6], and closest to our design, *Factor* [26].

Douence and Fradet demonstrate how their Λ_s encodes Plotkin's cbv λ-calculus [28] as well as Moggi's monadic constructs [24,25], illustrating how this design gives control over reduction. We will recall these encodings in Sections 3.2 and 3.3, and demonstrate the encoding of cbpv and Arrows [14].

1.3 The Functional Machine Calculus

The FMC combines both generalisations, *locations* and *sequencing*, in a simple model of higher-order computation with multiple effects. Its design and solid foundations in semantics mean that several important properties of the λ-calculus are preserved. The aim of this paper is introductory: it presents the syntax, operational semantics, and fundamental ideas and results with an emphasis on explanation and examples. In a forthcoming paper we will deepen these results with a strong normalization theorem, a domain-theoretic semantics of the untyped calculus, and semantics of the typed calculus in premonoidal and in Cartesian closed categories. The main results for the FMC as presented here are the following:

Confluence Beta-reduction is confluent, with evaluation behaviour expressed in syntax.

Algebraic effects The algebraic laws for reader/writer effects arise from reduction.

Compositionality Reader/writer effects combine seamlessly, due to the use of independent locations.

Types Simple types cover (higher-order) effect operations and confer termination of the machine.

2 The poly-lambda-calculus

We introduce a set of *locations* A, ranged over by a, b, $c, \dots$, to indicate the different stacks or streams for each effect. Abstraction and application are parameterized in A to give the corresponding pop and push actions. We write application $M\,N$ as $[N].\,M$ to emphasize the operational reading (push N and continue as M), to easily attach a location a, and to give unique parsing—cf. De Bruijn [5]. Abstraction $\lambda x.\,N$ is written $\langle x \rangle.\,N$ to emphasize the duality with application.

Definition 2.1 The *poly-λ-calculus* is given by the grammar

$$M, N \ ::= \ x \ \mid \ [N]a.\,M \ \mid \ a\langle x \rangle.\,M$$

with from left to right a *variable*, an *application* or *push action* on location a with function M and argument N, and an *abstraction* or *pop action* on location a that binds x in M. Terms are considered modulo α-equivalence. The regular λ-calculus embeds via a dedicated *main* location $\lambda \in A$, omitted from terms for brevity; so we may write $\lambda x.\,M$ or $\langle x \rangle.\,M$ for $\lambda\langle x \rangle.\,M$, and $M\,N$ or $[N].\,M$ for $[N]\lambda.\,M$.

The *poly-stack machine* is given by the following data. A *stack* of terms S is written with the top element to the right; we define them inductively below left, but they should better be considered as coinductive, to include streams. A *memory* S_A is a family of stacks or streams in A, defined below left. We write $S_A; S_a$ to identify the stack for a in S_A. A *state* is a pair (S_A, M), and the *transitions* or *steps* are given as top–to–bottom rules below centre. A *run* of the machine is a sequence of steps, written as

$(S_A, M) \Downarrow (T_A, N)$ or with a double line as below right.

$$S \quad ::= \quad \varepsilon \ | \ S \cdot M \qquad\qquad \frac{(\ S_A\ ;\ S_a \qquad ,\ [N]a.\,M\ )}{(\ S_A\ ;\ S_a \cdot N\ , \qquad\quad M\ )} \qquad \frac{(\ S_A\ ;\ S_a \cdot N\ ,\quad a\langle x\rangle.\,M\ )}{(\ S_A\ ;\ S_a \qquad ,\ \{N/x\}M\ )} \qquad \frac{(\ S_A\ ,\quad M\ )}{(\ T_A\ ,\quad N\ )}$$

$$S_A ::= \{\, S_a \ | \ a \in A\,\}$$

2.1   Encoding effects

Consider the following λ-calculus with effects. We will encode it in the poly-λ-calculus, with its *lazy* or cbn semantics. At the end of this section we will consider what would be needed to encode its *eager* or cbv semantics. (We assume familiarity with the operational semantics of λ-calculus and of effects; for an introduction see e.g. Winskel [37].)

$$\begin{aligned}
M, N, P ::= \ & x \ | \ M\,N \ | \ \lambda x.M && \lambda\text{-calculus} \\
| \ & \mathsf{read} \ | \ \mathsf{write}\ N; M && input/output \\
| \ & c := N; M \ | \ !c && state\ update\ and\ lookup \\
| \ & N \oplus M \ | \ N + M && probabilistic\ and\ non\text{-}deterministic\ sum
\end{aligned}$$

The cbn-encoding will follow the description in the introduction. The constructs of the above language are introduced as defined constructs ("sugar") into the poly-λ-calculus.

Input/output: Input uses a dedicated *input* location $\mathsf{in} \in A$ and is encoded by $\mathsf{read} \triangleq \mathsf{in}\langle x\rangle.\,x$. The machine is initialized with a stream $S_{\mathsf{in}} = \cdots N_3 \cdot N_2 \cdot N_1$ (infinite to the left), and the pop transition gives the expected operational semantics, below left. Writing to output uses a dedicated *output* location $\mathsf{out} \in A$ and is encoded by $\mathsf{write}\ N; M \triangleq [N]\mathsf{out}.\,M$. Evaluation then generates an output stream $N_1, N_2, \ldots$ (finite at any step) by the *push* machine transition, below right.

$$\frac{(\ S_A\ ;\ S_{\mathsf{in}} \cdot N\ ,\ \mathsf{in}\langle x\rangle.\,x\ )}{(\ S_A\ ;\ S_{\mathsf{in}} \qquad ,\qquad N\ )} \qquad\qquad \frac{(\ S_A\ ;\ S_{\mathsf{out}} \qquad ,\ [N]\mathsf{out}.\,M\ )}{(\ S_A\ ;\ S_{\mathsf{out}} \cdot N\ , \qquad\qquad M\ )}$$

State: A memory cell is modelled by a location $c \in A$. The associated stack is expected to hold at most one value, which is preserved by the encoding of the operators, and not enforced externally. Update and lookup are encoded by $c := N\,;M \ \triangleq\ c\langle\_\rangle.\,[N]c.\,M$ and $!c \ \triangleq\ c\langle x\rangle.\,[x]c.\,x$ where the underscore ($\_$) represents a variable that does not occur in M or N. In the machine, the stack for each cell is initialized with a (dummy) value, and the transitions then give the expected operational semantics.

$$\frac{(\ S_A\ ;\ \varepsilon_c \cdot P\ ,\ c\langle\_\rangle.\,[N]c.\,M\ )}{\dfrac{(\ S_A\ ;\ \varepsilon_c \qquad ,\qquad [N]c.\,M\ )}{(\ S_A\ ;\ \varepsilon_c \cdot N\ , \qquad\qquad M\ )}} \qquad\qquad \frac{(\ S_A\ ;\ \varepsilon_c \cdot M\ ,\ c\langle x\rangle.\,[x]c.\,x\ )}{\dfrac{(\ S_A\ ;\ \varepsilon_c \qquad ,\qquad [M]c.\,M\ )}{(\ S_A\ ;\ \varepsilon_c \cdot M\ , \qquad\qquad M\ )}}$$

Probabilistic and non-deterministic sums: Following the probabilistic case [4], probabilistic and non-deterministic sums are included via dedicated locations $\mathsf{rnd}, \mathsf{nd} \in A$ by $N \oplus M \ \triangleq\ \mathsf{rnd}\langle x\rangle.\,x\,M\,N$ and $N + M \ \triangleq\ \mathsf{nd}\langle x\rangle.\,x\,M\,N$. The machine is initialized with the corresponding streams of Church-encoded Booleans $\lambda x.\,\lambda y.\,x$ and $\lambda x.\,\lambda y.\,y$, generated probabilistically for rnd and non-deterministically for nd. The machine steps are as expected.

Example 2.2 Consider the following example term and its cbn encoding in the poly-λ-calculus. (Numbers can be seen informally as constants, or as Church numerals.)

$$a := 2\,;\,(\lambda x.\,!a)\,(a := 3\,;\,0) \quad = \quad a\langle\_\rangle.\,[2]a.\,[a\langle\_\rangle.\,[3]a.\,0].\,\langle x\rangle.\,a\langle y\rangle.\,[y]a.\,y$$

Its cbn reduction gives 2. It evaluates in the machine as follows (where the cell a is initialized with zero).

$(\ \varepsilon_a{\cdot}0\ ;\ \varepsilon_\lambda$	$,\ a\langle\_\rangle.\,[2]a.\,[a\langle\_\rangle.\,[3]a.\,0].\,\langle x\rangle.\,a\langle y\rangle.\,[y]a.\,y\ )$	
$(\ \varepsilon_a\ \ \ ;\ \varepsilon_\lambda$	$,\ \ \ \ \ \ [2]a.\,[a\langle\_\rangle.\,[3]a.\,0].\,\langle x\rangle.\,a\langle y\rangle.\,[y]a.\,y\ )$	
$(\ \varepsilon_a{\cdot}2\ ;\ \varepsilon_\lambda$	$,\ \ \ \ \ \ \ \ [a\langle\_\rangle.\,[3]a.\,0].\,\langle x\rangle.\,a\langle y\rangle.\,[y]a.\,y\ )$	
$(\ \varepsilon_a{\cdot}2\ ;\ \varepsilon_\lambda{\cdot}a\langle\_\rangle.\,[3]a.\,0\ ,$	$\langle x\rangle.\,a\langle y\rangle.\,[y]a.\,y\ )$	
$(\ \varepsilon_a{\cdot}2\ ;\ \varepsilon_\lambda$	$,\ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ a\langle y\rangle.\,[y]a.\,y\ )$	
$(\ \varepsilon_a\ \ \ ;\ \varepsilon_\lambda$	$,\ [2]a.\,2\ )$	
$(\ \varepsilon_a{\cdot}2\ ;\ \varepsilon_\lambda$	$,\ 2\ )$	

2.2 Beta-reduction

In the λ-calculus, β-reduction lets successive push- and pop-actions interact. Generalizing to multiple locations, these must be actions on the same stack, while other stacks may be accessed in-between. The β-rule is then as below, where each X_i is an abstraction or application not on location a, and (if the former) not capturing in N. Reduction is closed under any context. (A formal definition is given in Section 4.)

$$[N]a.\,X_1\ldots X_n.\,a\langle x\rangle.\,M\ \twoheadrightarrow\ X_1\ldots X_n.\,\{N/x\}M$$

Example 2.3 The term from example 2.2 reduces as follows, with reduced redexes underlined.

$$a\langle\_\rangle.\,\underline{[2]a.\,[a\langle\_\rangle.\,[3]a.\,0].\,\langle x\rangle.\,a\langle y\rangle}.\,[y]a.\,y\ \rightarrow\ a\langle\_\rangle.\,\underline{[a\langle\_\rangle.\,[3]a.\,0].\,\langle x\rangle}.\,[2]a.\,2\ \rightarrow\ a\langle\_\rangle.\,[2]a.\,2\ =\ a:=2\,;2$$

Analogous to β-reduction, η-reduction is as below, where each X_i is an abstraction or application not on location a, and x does not occur free in any of the X_i nor in M.

$$a\langle x\rangle.\,X_1\ldots X_n.\,[x]a.\,M\ \twoheadrightarrow_\eta\ X_1\ldots X_n.\,M$$

As an alternative to these rule schemes, terms may be taken modulo an equivalence $\sim$ generated by the *permutations* below left, and β- and η-reduction defined only on adjacent operators, as below right. Observe that the machine semantics immediately validates these equivalences. We will use the below formulation to consider the relation with algebraic effects, but otherwise use the above formulation.

$$\begin{aligned}[M]a.\,[N]b.\,P\ &\sim\ [N]b.\,[M]a.\,P\\ \langle x\rangle a.\,[N]b.\,P\ &\sim\ [N]b.\,\langle x\rangle a.\,P \quad\ \text{if } x\notin \mathsf{fv}(N)\\ \langle x\rangle a.\,\langle y\rangle b.\,P\ &\sim\ \langle y\rangle b.\,\langle x\rangle a.\,P\end{aligned}$$

$$\begin{aligned}[N]a.\,a\langle x\rangle.\,M\ &\twoheadrightarrow_\beta\ \{N/x\}M\\ a\langle x\rangle.\,[x]a.\,M\ &\twoheadrightarrow_\eta\ M \quad\quad\quad \text{if } x\notin \mathsf{fv}(M)\end{aligned}$$

Beta-reduction is confluent, as will be shown more generally for the FMC in Section 4. This is possible because it follows the algebraic laws for effects [27] instead of their operational semantics. For instance, laws for the interaction of *lookup* and *update* correspond to the following reductions.

$$\begin{aligned}c:=M\,;\,c:=N\,;P\ &=\ c\langle\_\rangle.\,\underline{[M]c.\,c\langle\_\rangle}.\,[N]c.\,P\ \rightarrow\ c\langle\_\rangle.\,[N]c.\,P\ \ =\ c:=N\,;P\\ c:=M\,;\,!c\ &=\ c\langle\_\rangle.\,\underline{[M]c.\,c\langle x\rangle}.\,[x]c.\,x\ \rightarrow\ c\langle\_\rangle.\,[M]c.\,M\ =\ c:=M\,;\,M\end{aligned}$$

The seven algebraic laws for global state of Plotkin and Power [27, p. 348] arise in our setting from β/η-reduction and $\sim$. Their notation has *update* $u_{loc,v}(M)$ of location loc with value v in M, and *lookup* $l_{loc}(M)_v$ of the value v at location loc using the value v as a parameter in M. These are encoded in the poly-λ-calculus as below, using abstraction with x instead of parametrization in v in the lookup case. Values v may be taken as arbitrary poly-λ-terms.

$$u_{a,v}(M)\ \triangleq\ a\langle\_\rangle.\,[v]a.\,M \qquad l_a(M)_x\ \triangleq\ a\langle x\rangle.\,[x]a.\,M$$

Proposition 2.4 *The poly-λ-calculus with $\twoheadrightarrow_{\beta\eta} \cup \sim$ generates the algebraic laws for state.*

Proof. By the following equations, where $a \neq b$ and in equation 7, $x \notin \mathsf{fv}(v)$.

1. $l_a(u_{a,y}(x))_y$ $=$ $a\langle y\rangle.\,[y]a.\,a\langle\_\rangle.\,[y]a.\,x$ $\to_\beta$ $a\langle y\rangle.\,[y]a.\,x$ $\to_\eta$ x

2. $l_a(l_a(M_{x,y})_x)_y$ $=$ $a\langle y\rangle.\,[y]a.\,a\langle x\rangle.\,[x]a.\,M_{x,y}$ $\to_\beta$ $a\langle y\rangle.\,[y]a.\,\{y/x\}M_{x,y}$ $=$ $l_a(M_{y,y})_y$

3. $u_{a,v}(u_{a,v'}(x))$ $=$ $a\langle\_\rangle.\,[v]a.\,a\langle\_\rangle.\,[v']a.\,x$ $\to_\beta$ $a\langle\_\rangle.\,[v']a.\,x$ $=$ $u_{a,v'}(x)$

4. $u_{a,v}(l_a(M_x)_x)$ $=$ $a\langle\_\rangle.\,[v]a.\,a\langle x\rangle.\,[x]a.\,M_x$ $\to_\beta$ $a\langle\_\rangle.\,[v]a.\,\{v/x\}M_x$ $=$ $u_{a,v}(M_v)$

5. $l_a(l_b(M_{x,y})_y)_x$ $=$ $a\langle x\rangle.\,[x]a.\,b\langle y\rangle.\,[y]b.\,M_{x,y}$ $\sim$ $b\langle y\rangle.\,[y]b.\,a\langle x\rangle.\,[x]a.\,M_{x,y}$ $=$ $l_b(l_a(M_{x,y})_x)_y$

6. $u_{a,v}(u_{b,v'}(x))$ $=$ $a\langle\_\rangle.\,[v]a.\,b\langle\_\rangle.\,[v']b.\,x$ $\sim$ $b\langle\_\rangle.\,[v']b.\,a\langle\_\rangle.\,[v]a.\,x$ $=$ $u_{b,v'}(u_{a,v}(x))$

7. $u_{a,v}(l_b(M_x)_x)$ $=$ $a\langle\_\rangle.\,[v]a.\,b\langle x\rangle.\,[x]b.\,M_x$ $\sim$ $b\langle x\rangle.\,[x]b.\,a\langle\_\rangle.\,[v]a.\,M_x$ $=$ $l_b(u_{a,v}(M_x))_x$

$\square$

2.3   Poly-types

A simple type for a poly-term represents its expected inputs, taken from multiple independent locations. Correspondingly, the antecedent of an implication is parameterized in a location, and implications on distinct locations may permute.

Definition 2.5 *Simple poly-types* are given by the grammar below left, where o (omicron) is a *base* type and $a(\sigma)\to\tau$ an *arrow* type, and considered modulo the congruence $\sim$ given below right.

$$\rho,\,\sigma,\,\tau \;::=\; o \;\mid\; a(\sigma)\to\tau \qquad\qquad a(\rho)\to b(\sigma)\to\tau \;\sim\; b(\sigma)\to a(\rho)\to\tau \quad (a \neq b)\,.$$

The typing rules are as follows, where a *context* Γ is a finite function from variables to types.

$$\frac{}{\Gamma,\,x:\tau \vdash x:\tau} \qquad\qquad \frac{\Gamma,\,x:\sigma \vdash M:\tau}{\Gamma \vdash a\langle x\rangle.\,M:a(\sigma)\to\tau} \qquad\qquad \frac{\Gamma \vdash N:\sigma \qquad \Gamma \vdash M:a(\sigma)\to\tau}{\Gamma \vdash [N]a.\,M:\tau}$$

Observe that the congruence $\sim$ means that a term $M:a(\rho)\to b(\sigma)\to\tau$ may be prefixed by a push action $[N]a$ where $N:\rho$ or by one $[P]b$ where $P:\sigma$ (or both, in either order).

2.4   Towards encoding call–by–value semantics

The poly-λ-calculus gives control over when effects are called, as we demonstrate by the following example.

Example 2.6 Consider the following example term, which is a normal form with cbn semantics.

$$f\ (c := 2\,;0)\ (!c)\ (c := 3\,;1)$$

With cbv, the arguments may be evaluated left to right, reducing to $f\ 0\ 2\ 1$, or right to left, which gives $f\ 0\ 3\ 1$. The two readings are encoded as follows (using regular applications to f for readability).

$$c\langle\_\rangle.\,[2]c.\,c\langle x\rangle.\,[x]c.\,c\langle\_\rangle.\,[3]c.\,f\,0\,x\,1 \quad \twoheadrightarrow \quad c\langle\_\rangle.\,[3]c.\,f\,0\,2\,1$$

$$c\langle\_\rangle.\,[3]c.\,c\langle x\rangle.\,[x]c.\,c\langle\_\rangle.\,[2]c.\,f\,0\,x\,1 \quad \twoheadrightarrow \quad c\langle\_\rangle.\,[2]c.\,f\,0\,3\,1$$

The encodings rely on repositioning an update $c := 1$ as a prefix, and for a lookup $!c$, on separating the global actions $c\langle x\rangle.\,[x]c$ from the variable x where the value is used. (The latter idea gives the cbv semantics in the probabilistic case [4].) However, it is unlikely that an encoding that only repositions effect operations can encode the cbv semantics of λ-calculus with effects. Consider the following example.

Example 2.7 With a cbv semantics, the term

$$a := (\lambda x.\, b := 1 \,;\, x)\, 0 \,;\, !b$$

first reduces the redex, to give $a := (b := 1 \,;\, 0) \,;\, !b$, which then evaluates by updating $b := 1$, then $a := 0$, and reading $!b$ as 1. To obtain this semantics in the poly-λ-calculus by manipulaton of effect operations would require lifting $b := 1$ out of a redex—a process which is likely undecidable in general.

The poly-λ-calculus thus gives control over effects, but not evaluation behaviour in general. It is an open question whether this is sufficient for practical purposes—one we cannot answer here. Instead, we will consider *sequencing* as a natural way to include cbv semantics.

3 The sequential lambda-calculus

As an instruction sequence for a stack machine, a λ-term is a string of push and pop actions that must end in a variable. But the machine would naturally accept *any* sequence of actions and variables. Relaxing the variable restriction would further enable composition of sequences. This design of λ-calculus with sequential composition appears several times in the literature and in practice: as the calculus Λ_s [7], as the higher-order κ-calculus [30], and in concatenative programming languages such as Factor [26]. We call this generalization of the λ-calculus *sequencing*, and implement it by introducing a *skip* (or *nil*) construct and making the variable a prefix.

Definition 3.1 The *sequential λ-calculus* is given by the following grammar.

$$M,\, N,\, P \quad ::= \quad \star \;\mid\; x.\, M \;\mid\; [N].\, M \;\mid\; \langle x\rangle.\, M$$

We may omit the trailing $.\star$ from terms for readability. *Capture-avoiding composition* $N \,;\, M$ is given by

$$\star \,;\, M \;=\; M \qquad x.\, N \,;\, M \;=\; x.\,(N \,;\, M) \qquad [P].\, N \,;\, M \;=\; [P].\,(N \,;\, M) \qquad \langle y\rangle.\, N \,;\, M \;=\; \langle y\rangle.\,(N \,;\, M)$$

where in the last case y is not free in M. *Capture-avoiding substitution* $\{M/x\}N$ is as follows.

$$\{M/x\}\star \;=\; \star \qquad\qquad \{M/x\}[P]a.\, N \;=\; [\{M/x\}P]a.\, \{M/x\}N$$
$$\{M/x\}x.\, N \;=\; M \,;\, \{M/x\}N \qquad\qquad \{M/x\}a\langle x\rangle.\, N \;=\; a\langle x\rangle.\, N$$
$$\{M/x\}y.\, N \;=\; y.\, \{M/x\}N \;\; (x \neq y) \qquad \{M/x\}a\langle y\rangle.\, N \;=\; a\langle y\rangle.\, \{M/x\}N \;\; (x \neq y,\; y \notin \mathsf{fv}(M))$$

Beta-reduction is otherwise standard, by closing the rule below left under all contexts. The abstract machine has *states* (S, M) of a stack and a term, and the transitions below right.

$$[N].\, \langle x\rangle.\, M \;\to\; \{N/x\}M \qquad\qquad \frac{(\; S \quad,\; [N].\, M\;)}{(\; S{\cdot}N \;,\quad M\;)} \qquad \frac{(\; S{\cdot}N \;,\quad \langle x\rangle.\, M\;)}{(\; S \quad,\; \{N/x\}M\;)}$$

Example 3.2 Consider the following example terms.

$$\langle x\rangle.\, [x].\, [x] \qquad \langle x\rangle.\, \langle y\rangle \qquad [\langle x\rangle.\, [x]].\, \langle f\rangle.\, f.\, f.\, f$$

The first duplicates the top item on the stack; the second removes two items; the third pushes the term $\langle x\rangle.\, [x]$ (which picks up and returns an item), pops it as f, and runs it three times.

Observe that the changes to evaluation are absorbed by substitution and composition, while β-reduction and machine evaluation remain unchanged. There is nevertheless a change in perspective from the λ-calculus, in that the return values or outputs of a computation are pushed to the stack, rather than left as

the remainder of the term. Machine evaluation for a term M with input stack S is expected to terminate in $\star$ (just as imperative computation successfully terminates in a *skip* command) with an output stack T, i.e. $(S, M) \Downarrow (T, \star)$. Then $\star$ gives the identity run (of zero steps), and $M \,;N$ gives composition of runs.

Proposition 3.3 *If* $(R, M) \Downarrow (S, \star)$ *and* $(S, N) \Downarrow (T, \star)$ *then* $(R, M \,;N) \Downarrow (T, \star)$.

3.1 Sequential types

The type system for the sequential λ-calculus follows that of the κ-calculus [30]; similar type systems have also been studied for stack languages [33]. The type of a term describes the input/output behaviour of its machine evaluation: it consists of an implication between a vector of input types, one for each element consumed from the stack, and a vector of output types, one for each item returned to the stack.

Definition 3.4 *Sequential types* are defined by:

$$\rho, \sigma, \tau, \upsilon \;::=\; \vec{\sigma} \Rightarrow \vec{\tau} \qquad\qquad \vec{\tau} \;::=\; \tau_1 \ldots \tau_n$$

Vector concatenation is by juxtaposition, $\vec{\sigma}\vec{\tau}$, and the reverse of a vector $\vec{\tau} = \tau_1 \ldots \tau_n$ is $\overleftarrow{\tau} = \tau_n \ldots \tau_1$. Typing rules for the sequential λ-calculus are given below. A stack is typed by a type vector, where $\Gamma \vdash \varepsilon \cdot M_1 \cdots M_n : \tau_1 \ldots \tau_n$ if $\Gamma \vdash M_i : \tau_i$ for each $i \leq n$.

$$\frac{}{\Gamma \vdash \star : \overleftarrow{\tau} \Rightarrow \vec{\tau}} \;\star \qquad \frac{\Gamma, x : \overleftarrow{\rho} \Rightarrow \vec{\sigma} \vdash \quad M : \overleftarrow{\sigma}\,\overleftarrow{\tau} \Rightarrow \vec{\upsilon}}{\Gamma, x : \overleftarrow{\rho} \Rightarrow \vec{\sigma} \vdash x.\,M : \overleftarrow{\rho}\,\overleftarrow{\tau} \Rightarrow \vec{\upsilon}} \;\mathsf{var} \qquad \frac{\Gamma, x : \rho \vdash M : \overleftarrow{\sigma} \Rightarrow \vec{\tau}}{\Gamma \vdash \langle x \rangle.\,M : \rho\,\overleftarrow{\sigma} \Rightarrow \vec{\tau}} \;\mathsf{abs} \qquad \frac{\Gamma \vdash N : \rho \quad \Gamma \vdash M : \rho\,\overleftarrow{\sigma} \Rightarrow \vec{\tau}}{\Gamma \vdash [N].\,M : \overleftarrow{\sigma} \Rightarrow \vec{\tau}} \;\mathsf{app}$$

Example 3.5 The terms in Example 3.2 can be typed as follows.

$$\langle x \rangle.\,[x].\,[x] : \tau \Rightarrow \tau\tau \qquad \langle x \rangle.\,\langle y \rangle : \tau\sigma \Rightarrow \qquad [\langle x \rangle.\,[x]].\,\langle f \rangle.\,f.\,f.\,f : \tau \Rightarrow \tau$$

Observe that because stacks are last-in first-out, the identity function on the top two stack items is the term $\langle x \rangle.\,\langle y \rangle.\,[y].\,[x] : \tau\sigma \Rightarrow \sigma\tau$, whereas the function that swaps them is $\langle x \rangle.\,\langle y \rangle.\,[x].\,[y] : \tau\sigma \Rightarrow \tau\sigma$.

Example 3.6 The term $\lambda x.\,x\,x = \langle x \rangle.\,[x].\,x$ can be typed by assigning x a type that does not consume input, i.e. of the form $(\Rightarrow \vec{\tau})$. The self-application $x\,x = [x].\,x$ then has the type $\Rightarrow (\Rightarrow \vec{\tau})\vec{\tau}$, which reflects that the return values accumulate: if evaluating x returns the stack $T : \vec{\tau}$, then $[x].\,x$ returns the stack consisting of x prepended to T. The type derivation is below. Note that the term $(\lambda x.\,x\,x)(\lambda y.\,y\,y)$ is not typeable: the argument $\lambda y.\,y\,y$ needs a type that takes input, and then so should x.

$$\frac{\dfrac{\dfrac{}{x : \Rightarrow \vec{\tau} \vdash \star : \overleftarrow{\tau} \Rightarrow \vec{\tau}} \;\star}{x : \Rightarrow \vec{\tau} \vdash x.\,\star : \Rightarrow \vec{\tau}} \;\mathsf{var} \qquad \dfrac{\dfrac{}{x : \Rightarrow \vec{\tau} \vdash \star : \overleftarrow{\tau}(\Rightarrow \vec{\tau}) \Rightarrow (\Rightarrow \vec{\tau})\vec{\tau}} \;\star}{x : \Rightarrow \vec{\tau} \vdash x.\,\star : (\Rightarrow \vec{\tau}) \Rightarrow (\Rightarrow \vec{\tau})\vec{\tau}} \;\mathsf{var}}{\dfrac{x : \Rightarrow \vec{\tau} \vdash [x.\,\star].\,x.\,\star : \Rightarrow (\Rightarrow \vec{\tau})\vec{\tau}}{\vdash \langle x \rangle.\,[x.\,\star].\,x.\,\star : (\Rightarrow \vec{\tau}) \Rightarrow (\Rightarrow \vec{\tau})\vec{\tau}} \;\mathsf{abs}} \;\mathsf{app}$$

Remark 3.7 [Due to Chris Barrett] Observe that all sequential types τ are inhabited by at least the element $\bot_\tau : \tau$, defined below (note that the base case $n = m = 0$ gives $\star : (\Rightarrow)$). This is in contrast with the simply-typed λ-calculus, where not all types are inhabited due to the presence of the uninhabited base type o.

$$\bot_\tau = \langle x_n \rangle \ldots \langle x_1 \rangle.\,[\bot_{\tau_1}] \ldots [\bot_{\tau_m}] \quad \text{where} \quad \tau = \sigma_n \ldots \sigma_1 \Rightarrow \tau_1 \ldots \tau_m$$

3.2 Encodings of call–by–name calculi

The (regular, call–by–name) λ-calculus is included as a fragment of the sequential λ-calculus. We extend this embedding with types and with products, and to Moggi's *computational metalanguage* [25] following

Douence and Fradet [7]. The main observation is that implications embed as input-only sequential types, and products as output-only types, below left. The formal, inductive encoding is then below right.

$$\tau_1 \to \cdots \to \tau_n \to o \;=\; \tau_1 \ldots \tau_n \Rightarrow \qquad\qquad o \;\triangleq\; (\Rightarrow) \qquad \rho \to (\overleftarrow{\sigma} \Rightarrow \overrightarrow{\tau}) \;\triangleq\; \rho\overleftarrow{\sigma} \Rightarrow \overrightarrow{\tau}$$

$$\tau_1 \times \cdots \times \tau_n \;=\; \Rightarrow \tau_n \ldots \tau_1 \qquad\qquad 1 \;\triangleq\; (\Rightarrow) \qquad \sigma \times \tau \;\triangleq\; \Rightarrow \tau\sigma$$

Following the types, product terms encode as follows.

$$(M, N) \;\triangleq\; [N].\,[M] \qquad\qquad \pi_i(P) \;\triangleq\; P\,;\langle x_1\rangle.\,\langle x_2\rangle.\,x_i \qquad\qquad () \;\triangleq\; \star$$

The computational metalanguage extends the λ-calculus with monadic type formers $T(\sigma)$, and a *return* construct $[M]_T$ and a *let* construct $\mathtt{let}_T$ parameterized in T. In the interpretation in the sequential λ-calculus, the return value of a monadic function is pushed to the stack. A monadic function type is then interpreted as one with a single output, as below left. The language constructs are encoded as below right. It is easily verified that this extends correctly to type derivations and reductions.

$$\tau_1 \to \cdots \to \tau_n \to T(\sigma) \;=\; \tau_1 \ldots \tau_n \Rightarrow \sigma \qquad\qquad [M]_T \;\triangleq\; [M] \qquad \mathtt{let}_T \; x \Leftarrow N \text{ in } M \;\triangleq\; N\,;\langle x\rangle.\,M$$

3.3 Encodings of call–by–value calculi

The cbv λ-calculus [28] and the computational λ-calculus λ_c [24], which extends the former with a monadic type constructor T and with the term constructs *return* and *let*, have an encoding in the κ-calculus [7,30]. We recall this for the sequential λ-calculus, and observe that it naturally extends to types. The cbv-interpretion of types is naturally viewed in two stages. First, types in isolation are translated as follows.

$$o_v \;=\; (\Rightarrow) \qquad (\sigma \to \tau)_v \;=\; \sigma_v \Rightarrow \tau_v \qquad T(\sigma)_v \;=\; \Rightarrow \sigma_v$$

Evaluation of a λ_c-term returns a value, which in the encoding is pushed to the stack. A typed term $M : \tau$ will then translate as $M_v \colon \Rightarrow \tau_v$, with a single output type. Terms of the computational λ-calculus are then interpreted as below. Again it is easily verified that this extends correctly to type derivations and reduction. The machine behaviour of encoded terms can be seen to follow the SECD-machine [16].

$$x_v \;=\; [x] \qquad (\lambda x.M)_v \;=\; [\langle x\rangle.\,M_v] \qquad ([M]_T)_v \;=\; [M_v]$$

$$(M\,N)_v \;=\; N_v\,;M_v\,;\langle x\rangle.\,x \qquad (\mathtt{let}_T \; x \Leftarrow N \text{ in } M)_v \;=\; N_v\,;\langle x\rangle.\,M_v$$

3.4 Arrows, call–by–push–value, and kappa-calculus

The sequential λ-calculus may encode the related formalisms of Arrows, cbpv, and κ-calculus. For reasons of space, we will not recall these calculi in detail and only provide an outline to the interested reader.

Hughes's *Arrows* [14] take the λ-calculus with products and extend it with a second implication $\sigma \rightsquigarrow \tau$, which we interpret directly as that of the sequential λ-calculus, $\sigma \Rightarrow \tau$. Arrow terms have three constructors, encoded as below. The first lifts a regular term $M \colon \rho \to \sigma$ to an arrow term; the second composes two arrow terms $P \colon \rho \rightsquigarrow \sigma$ and $Q \colon \sigma \rightsquigarrow \tau$; and the third applies the arrow term P to the first element of a pair.

$$\begin{aligned}
\mathsf{arr}\; M \;:\; \rho \rightsquigarrow \sigma \quad &\triangleq\; \langle x\rangle.\,[[x].\,M] \;:\; \rho \Rightarrow \sigma \\
P \ggg Q \;:\; \rho \rightsquigarrow \tau \quad &\triangleq\; \quad P\,;Q \;:\; \rho \Rightarrow \tau \\
\mathsf{first}\; P \;:\; (\rho \times \tau) \rightsquigarrow (\sigma \times \tau) \;&\triangleq\; \langle x\rangle.\,[x.\,P] \;:\; (\Rightarrow \tau\rho) \Rightarrow (\Rightarrow \tau\sigma)
\end{aligned}$$

The perspective that emerges from this encoding is that the arrow calculus corresponds to a version of the sequential λ-calculus with binary products instead of stacks (which may be considered n-ary products).

Characteristic of cbpv [18,19], and also featured in κ-calculus, is the separation of *computations* and *values*. In the sequential λ-calculus this distinction is present, too, if implicitly: values live on the stack, and computations run the machine. To make it explicit, we may extend the calculus with *thunk* and *force* constructs $!M$ and $?V$, and their reduction rule, as below left. Term constructs of cbpv (without products or sums) then embed as below right.

$$V, W \quad ::= \quad x \mid !M$$
$$M, N \quad ::= \quad \star \mid ?V.\,M \mid [V].\,M \mid \langle x \rangle.\,M$$
$$!?N.\,M \quad \rightarrow \quad N\,;M$$

$$\text{thunk } M \triangleq !M \qquad \text{force } V \triangleq ?V$$
$$\text{return } V \triangleq [V] \qquad N \text{ to } x.\,M \triangleq N\,;\langle x \rangle.\,M$$
$$\lambda x.M \triangleq \langle x \rangle.\,M \qquad V\text{`}M \triangleq [V].\,M$$

Types for cbpv feature a monadic functor F and a *value* functor U. The latter could be introduced into sequential types as below left, though the structure of the arrow type $\Rightarrow$ makes it redundant. Types then further encode as call–by–name types, below right.

$$\sigma, \tau \quad ::= \quad \vec{\sigma} \Rightarrow \vec{\tau} \qquad \vec{\tau} \quad ::= \quad U\tau_1 \ldots U\tau_n \qquad \rho \rightarrow (\overleftarrow{\sigma} \Rightarrow \vec{\tau}) \triangleq \rho \overleftarrow{\sigma} \Rightarrow \vec{\tau} \qquad F\sigma \triangleq {} \Rightarrow \sigma$$

The higher-order κ-calculus [30] is closely related to the sequential λ-calculus. Types are the same as sequential types: an implication between type vectors. Terms omit the unit $\star$ and have composition $M\,;N$ as a primitive (rather than prefixing). The remaining constructs encode as follows.

$$\text{push } V \triangleq [V] \qquad \kappa x.M \triangleq \langle x \rangle.\,M \qquad \text{mkthunk } M \triangleq !M \qquad \text{apply} \triangleq \langle x \rangle.\,?x$$

3.5 String diagrams

We may view typed sequential λ-terms as string diagrams. A term $M: \rho_1 \ldots \rho_m \Rightarrow \sigma_n \ldots \sigma_1$ is rendered as below. The wires represent the input and output stacks, with the first element at the top.
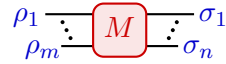
We will use these diagrams to illustrate how types compose. First, *strict composition* is the composition of terms $M: \overleftarrow{\rho} \Rightarrow \vec{\sigma}$ and $N: \overleftarrow{\sigma} \Rightarrow \vec{\tau}$ into $M\,;N: \overleftarrow{\rho} \Rightarrow \vec{\tau}$, below left. This does not give the most general form of composition. For that, we combine it with the following notion of *expansion*. If a term takes an input stack R to an output stack S, then when given a larger stack UR it returns US, with U untouched. Then if $M: \overleftarrow{\rho} \Rightarrow \vec{\sigma}$ also $M: \overleftarrow{\rho}\overleftarrow{v} \Rightarrow \vec{v}\vec{\sigma}$, illustrated below right.
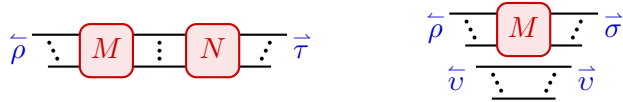
These constructions combine to give the general case, where in $M\,;N$ the type of M or N may be expanded. Note that in the regular λ-calculus only the first case arises, as the second case requires multiple outputs.
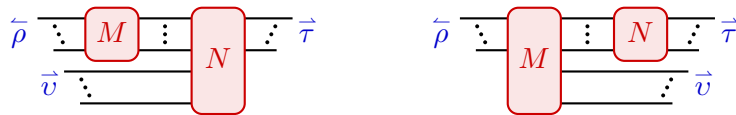
Definition 3.8 *Type composition* $\sigma \cdot \tau$ is the partial operation given below, and is undefined otherwise.

$$(\overleftarrow{\rho} \Rightarrow \vec{\sigma}) \cdot (\overleftarrow{\sigma}\,\overleftarrow{v} \Rightarrow \vec{\tau}) \quad = \quad (\overleftarrow{\rho}\,\overleftarrow{v} \Rightarrow \vec{\tau})$$
$$(\overleftarrow{\rho} \Rightarrow \vec{v}\,\vec{\sigma}) \cdot (\overleftarrow{\sigma} \Rightarrow \vec{\tau}) \quad = \quad (\overleftarrow{\rho} \Rightarrow \vec{v}\,\vec{\tau})$$

Observe that the first equation in Definition 3.8 corresponds to the left composition diagram above, and the second equation to the right diagram, where in both cases $\vec{\sigma}$ gives the types of the connecting wires between M and N. The following proposition establishes these basic properties, as well as the familiar *subject reduction*.

Proposition 3.9 *Typed terms satisfy the following properties:*

- *Strict composition: if $\Gamma \vdash M : \overleftarrow{\rho} \Rightarrow \vec{\sigma}$ and $\Gamma \vdash N : \overleftarrow{\sigma} \Rightarrow \vec{\tau}$ then $\Gamma \vdash M \mathbin{;} N : \overleftarrow{\rho} \Rightarrow \vec{\tau}$.*
- *Expansion: if $\Gamma \vdash M : \overleftarrow{\rho} \Rightarrow \vec{\sigma}$ then $\Gamma \vdash M : \overleftarrow{\rho}\,\overleftarrow{\upsilon} \Rightarrow \vec{\upsilon}\,\vec{\sigma}$.*
- *Composition: if $\Gamma \vdash M : \sigma$ and $\Gamma \vdash N : \tau$ and $\sigma \cdot \tau$ is defined, then $\Gamma \vdash M \mathbin{;} N : \sigma \cdot \tau$.*
- *Substitution: if $\Gamma \vdash M : \sigma$ and $\Gamma, x : \sigma \vdash N : \tau$ then $\Gamma \vdash \{M/x\}N : \tau$.*
- *Subject reduction: if $\Gamma \vdash N : \tau$ and $N \twoheadrightarrow M$ then $\Gamma \vdash M : \tau$.*

3.6 Machine termination

A remarkable aspect of the type system is how it gives a direct connection with termination of the machine. To expose this, we formalize the intuitive meaning of types as describing the initial and final stack of a run of the machine.

Definition 3.10 The set $\textsc{run}(\overleftarrow{\sigma} \Rightarrow \vec{\tau})$ is the set of terms M such that for any stack $S \in \textsc{run}(\vec{\sigma})$ there is a stack $T \in \textsc{run}(\vec{\tau})$ and a run of the machine $(S, M) \Downarrow (T, \star)$, where $\textsc{run}(\tau_1 \ldots \tau_n)$ is the set of stacks $\varepsilon \cdot N_1 \cdots N_n$ such that $N_i \in \textsc{run}(\tau_i)$.

Note that a successful run requires the term to be closed, so a set $\textsc{run}(\tau)$ contains only closed terms. The following lemma shows that $\textsc{run}(\tau)$ is always inhabited by the term $\bot_\tau$ (see Remark 3.7).

Lemma 3.11 *For any type τ the set $\textsc{run}(\tau)$ is inhabited: $\bot_\tau \in \textsc{run}(\tau)$.*

If $M : \tau$ implies $M \in \textsc{run}(\tau)$, then a type derivation *is* a termination proof of the machine. This is Theorem 3.13, and proving it gives a concrete *Tait-style* reducibility proof [34], where $\textsc{run}(\tau)$ takes the rôle of the reducibility set for τ. By using the properties of machine runs the proof is then a simple, direct induction on type derivations.

Vector notation is extended to variables, $\vec{x} = x_1 \ldots x_n$, to contexts as $\vec{x} : \vec{\tau} = x_1 : \tau_1, \ldots, x_n : \tau_n = \Gamma$, and to simultaneous substitutions: if $S = \varepsilon \cdot M_1 \cdots M_n$ then $\{S/\vec{x}\} = \{M_1/x_1, \ldots, M_n/x_n\}$. We write concatenation of stacks S and T by juxtaposition, ST.

Lemma 3.12 *If $\vec{w} : \vec{\omega} \vdash M : \tau$ then for any $W \in \textsc{run}(\vec{\omega})$, $\{W/\vec{w}\}M \in \textsc{run}(\tau)$.*

Proof. By induction on the type derivation. In each case, let $\Gamma = \vec{w} : \vec{\omega}$, let W be a stack in $\textsc{run}(\vec{\omega})$, and let $M' = \{W/\vec{w}\}M$.

- If the derivation is a $\star$-rule for $\Gamma \vdash \star : \overleftarrow{\tau} \Rightarrow \vec{\tau}$ then there is a trivial zero-step run $(T, \star) \Downarrow (T, \star)$.
- If the derivation ends in the variable rule below left, then for any $N \in \textsc{run}(\vec{\rho} \Rightarrow \vec{\sigma})$, the inductive hypothesis gives a run for $\{N/x\}M'$ from any $TS \in \textsc{run}(\vec{\tau}\,\vec{\sigma})$ to some $U \in \textsc{run}(\vec{\upsilon})$ (second item below). For N there is a run from any $R \in \textsc{run}(\vec{\rho})$ to some $S \in \textsc{run}(\vec{\sigma})$ (third item below). These runs compose into one for $\{W/\vec{w}, N/x\}x. M = N \mathbin{;} \{N/x\}M'$ as below right, expanding the stack on the run for N by T. Note that we may assume $x \notin \mathsf{fv}(W)$ (otherwise we rename x).

$$
\frac{\Gamma, x : \overleftarrow{\rho} \Rightarrow \vec{\sigma} \vdash \quad M : \overleftarrow{\sigma}\,\overleftarrow{\tau} \Rightarrow \vec{\upsilon}}{\Gamma, x : \overleftarrow{\rho} \Rightarrow \vec{\sigma} \vdash x. M : \overleftarrow{\rho}\,\overleftarrow{\tau} \Rightarrow \vec{\upsilon}} \text{ var}
\qquad
\frac{(\, TS \,,\, \{N/x\}M' \,)}{(\, U \quad,\qquad\qquad \star\,)}
\qquad
\frac{(\, R \,,\, N \,)}{(\, S \,,\, \star\,)}
\qquad
\frac{\dfrac{(\, TR \,,\, N \mathbin{;} \{N/x\}M' \,)}{(\, TS \,,\qquad \{N/x\}M' \,)}}{(\, U \quad,\qquad\qquad\qquad \star\,)}
$$

- If the derivation ends in the application rule below left, then by the inductive hypothesis for N we have $N' = \{W/\vec{w}\}N \in \textsc{run}(\rho)$, and for M we have a run from M' and any stack $S \in \textsc{run}(\vec{\sigma})$ with N'

added on top, to some $T \in \text{RUN}(\vec{\tau})$. This gives the run for $\{W/\vec{w}\}[N]. M = [N']. M'$ below right.

$$\frac{\Gamma \vdash N : \rho \qquad \Gamma \vdash M : \rho\, \overleftarrow{\sigma} \Rightarrow \vec{\tau}}{\Gamma \vdash [N]. M : \overleftarrow{\sigma} \Rightarrow \vec{\tau}} \text{ app} \qquad\qquad \frac{\dfrac{(\ S \qquad , \ [N']. M'\ )}{(\ S \cdot N' , \qquad M'\ )}}{(\ T \qquad , \qquad \star\ )}$$

- If the derivation ends in the abstraction rule below left, then for any $N \in \text{RUN}(\rho)$ and $S \in \text{RUN}(\vec{\sigma})$ the inductive hypothesis gives a run for $\{N/x\}M'$ to some $T \in \text{RUN}(\vec{\tau})$. This gives the run for $\{W/\vec{w}\}\langle x \rangle. M = \langle x \rangle. M'$ below right.

$$\frac{\Gamma, x : \rho \vdash M : \overleftarrow{\sigma} \Rightarrow \vec{\tau}}{\Gamma \vdash \langle x \rangle. M : \rho\, \overleftarrow{\sigma} \Rightarrow \vec{\tau}} \text{ abs} \qquad\qquad \frac{\dfrac{(\ S \cdot N , \quad \langle x \rangle. M'\ )}{(\ S \qquad , \ \{N/x\}M'\ )}}{(\ T \qquad , \qquad \star\ )}$$

$\square$

The following theorem is then immediate.

Theorem 3.13 *For a typed, closed term $M : \overleftarrow{\sigma} \Rightarrow \vec{\tau}$ and stack $S : \vec{\sigma}$ the machine terminates.*

4   The functional machine calculus

The combination of both generalizations, *locations* and *sequencing*, gives the Functional Machine Calculus.

Definition 4.1 The *Functional Machine Calculus* (FMC) is given by the below grammar, with from left to right the constructors *nil*, a *(sequential) variable*, an *application* or *push action* on the location a, and an *abstraction* or *pop action* on a which binds x in M. Terms are considered modulo α-equivalence.

$$M, N, P \quad ::= \quad \star \ \mid \ x. M \ \mid \ [N]a. M \ \mid \ a\langle x \rangle. M$$

Composition $N \,; M$ and substitution $\{N/x\}M$ are as for the sequential λ-calculus. The machine is as for the poly-λ-calculus: a *state* is a pair (S_A, M) of a memory and a term, and the *transitions* are:

$$\frac{(\ S_A \,; \ S_a \qquad , \ [N]a. M\ )}{(\ S_A \,; \ S_a \cdot N , \qquad M\ )} \qquad\qquad \frac{(\ S_A \,; \ S_a \cdot N , \quad a\langle x \rangle. M\ )}{(\ S_A \,; \ S_a \qquad , \ \{N/x\}M\ )}$$

Beta-reduction is as for the poly-λ-calculus: a redex consists of a successive application and abstraction *on the same location*, separated by any number of actions on other locations. We will now make this formal.

Head contexts H are defined as below left. The term obtained by replacing the hole $\{\}$ in H with M is denoted $H. M$, where a binder $a\langle x \rangle$ in H captures in M. The *binding variables* $\text{bv}(H)$ of H are those variables x where H is constructed over $a\langle x \rangle$. The set of *locations* used in a term or context is denoted $\text{loc}(M)$ respectively $\text{loc}(H)$. Then *Beta-reduction* is defined by the rewrite rule schema below right, where $a \notin \text{loc}(H)$ and $\text{bv}(H) \cap \text{fv}(N) = \varnothing$, and is closed under all contexts.

$$H \quad ::= \quad \{\} \ \mid \ [M]a. H \ \mid \ a\langle x \rangle. H \qquad\qquad [N]a. H. a\langle x \rangle. M \ \twoheadrightarrow \ H. \{N/x\}M$$

We will first consider the untyped calculus. We give the cbv encoding of effects and provide an intuition for programming in the FMC, then establish confluence and connect machine evaluation to β-reduction. We then consider simple types. Constants will be used informally, in examples.

4.1   Call–by–value with effects

We extend the encoding $(-)_v$ of the computational λ-calculus of Section 3.3 to effects as follows. (The case for $N \oplus M$ is the same as for cbn, as it expects a Church boolean for x.)

$$
\begin{aligned}
\mathsf{read}_v &= \mathsf{in}\langle x\rangle.\,[x] \\
(\mathsf{write}\ N; M)_v &= N_v\,;\,\langle x\rangle.\,[x]\mathsf{out}.\,M_v \\
(c := N; M)_v &= N_v\,;\,\langle x\rangle.\,c\langle\_\rangle.\,[x]c.\,M_v \\
!c_v &= c\langle x\rangle.\,[x]c.\,[x] \\
(N \oplus M)_v &= \mathsf{rnd}\langle x\rangle.\,[M_v].\,[N_v].\,x
\end{aligned}
$$

Example 4.2 The term from Example 2.2, below, is given a cbv interpretation as follows.

$$
a := 2\,;\,(\lambda x.\,!a)\,(a := 3\,;\,0)\ \to_{\mathsf{cbv}}\ 3
$$

Integers are values, and the translation will use $i_v = [i]$. An update with an integer then simplifies by:

$$
(a := i\,;\,M)_v\ =\ \underline{[i].\,\langle x\rangle.\,a\langle\_\rangle}.\,[x]a\,;\,M\ \to\ a\langle\_\rangle.\,[i]a\,;\,M
$$

The cbv-translated term, after applying this reduction to the two updates, further reduces as follows.

$$
a\langle\_\rangle.\,[2]a.\,a\langle\_\rangle.\,[3]a.\,[0].\,[\langle x\rangle.\,a\langle y\rangle.\,[y]a.\,[y]].\,\langle z\rangle.\,z\ \twoheadrightarrow\ a\langle\_\rangle.\,[3]a.\,[3]
$$

Example 4.3 The term from Example 2.7,

$$
a := (\lambda x.\,b := 1\,;\,x)\,0\,;\,!b\ \twoheadrightarrow_{\mathsf{cbv}}\ 1
$$

translates and reduces as follows (with the same shortcut for update as above).

$$
\begin{aligned}
&[0].\,[\langle x\rangle.\,b\langle\_\rangle.\,[1]b.\,[x]].\,\langle y\rangle.\,y.\,\langle z\rangle.\,a\langle\_\rangle.\,[z]a.\,b\langle u\rangle.\,[u]b.\,[u] \\
\to\ &\underline{[0].\,\langle x\rangle}.\,b\langle\_\rangle.\,[1]b.\,[x].\,\langle z\rangle.\,a\langle\_\rangle.\,[z]a.\,b\langle u\rangle.\,[u]b.\,[u] \\
\twoheadrightarrow\ &b\langle\_\rangle.\,\underline{[1]b.\,a\langle\_\rangle.\,[0]a.\,b\langle u\rangle}.\,[u]b.\,[u] \\
\to\ &b\langle\_\rangle.\,a\langle\_\rangle.\,[0]a.\,[1]b.\,[1] \\
\sim\ &a := 0\,;\,b := 1\,;\,1
\end{aligned}
$$

4.2   Programming in the FMC

As in the sequential λ-calculus, programming in the FMC naturally follows the concatenative paradigm. A term M is viewed as a function taking an input memory S_A to an output memory T_A by a run of the machine $(S_A, M) \Downarrow (T_A, \star)$. Functions standardly operate on the main stack λ, and it is then natural to consider effect operators that transfer values between the main stack and other locations, as the cbv translations of effect operators do. We introduce the following operations for input and output, a random generator, and a memory cell c. We further add *definitions* or *let*, as a redex.

$$
\begin{aligned}
\mathsf{print} &\triangleq \langle x\rangle.\,[x]\mathsf{out} & \mathsf{get}\ c &\triangleq c\langle x\rangle.\,[x]c.\,[x] \\
\mathsf{read} &\triangleq \mathsf{in}\langle x\rangle.\,[x] & \mathsf{set}\ c &\triangleq \langle x\rangle.\,c\langle\_\rangle.\,[x]c \\
\mathsf{rand} &\triangleq \mathsf{rnd}\langle x\rangle.\,[x] & (x{=}N)\,;\,M &\triangleq [N].\,\langle x\rangle.\,M
\end{aligned}
$$

$$[\mathsf{rnd}\langle x\rangle.\,\underline{[x]}.\,\langle y\rangle.\,c\langle \_\rangle.\,[y]c.\,c\langle z\rangle.\,[z]c.\,[z]].\,\langle f\rangle.\,f.\,f.\,+.\,\langle p\rangle.\,[p]out$$

$$[\mathsf{rnd}\langle x\rangle.\,c\langle \_\rangle.\,\underline{[x]c}.\,c\langle z\rangle.\,[z]c.\,[z]].\,\langle f\rangle.\,f.\,f.\,+.\,\langle p\rangle.\,[p]out$$

$$\underline{[\mathsf{rnd}\langle x\rangle.\,c\langle \_\rangle.\,\;[x]c.\,[x]]}.\,\langle f\rangle.\,f.\,f.\,+.\,\langle p\rangle.\,[p]out$$

$$\mathsf{rnd}\langle x\rangle.\,c\langle \_\rangle.\,\underline{[x]c}.\,[x]\;.\;\mathsf{rnd}\langle y\rangle.\,c\langle \_\rangle.\,[y]c.\,[y]\;.\,+.\,\langle p\rangle.\,[p]out$$

$$\mathsf{rnd}\langle x\rangle.\,c\langle \_\rangle.\,[x]\;.\;\mathsf{rnd}\langle y\rangle.\,[y]c.\,[y]\;.\,+.\,\langle p\rangle.\,[p]out$$

out	rnd	c	λ	
;	$6\cdot 7$;	$\star$;	,	$\mathsf{rnd}\langle x\rangle.\,c\langle \_\rangle.\,[x].\,\mathsf{rnd}\langle y\rangle.\,[y]c.\,[y].\,+.\,\langle p\rangle.\,[p]out$)
;	6 ;	$\star$;	,	$c\langle \_\rangle.\,[7].\,\mathsf{rnd}\langle y\rangle.\,[y]c.\,[y].\,+.\,\langle p\rangle.\,[p]out$)
;	6 ;	;	,	$[7].\,\mathsf{rnd}\langle y\rangle.\,[y]c.\,[y].\,+.\,\langle p\rangle.\,[p]out$)
;	6 ;	; 7	,	$\mathsf{rnd}\langle y\rangle.\,[y]c.\,[y].\,+.\,\langle p\rangle.\,[p]out$)
;	;	; 7	,	$[6]c.\,[6].\,+.\,\langle p\rangle.\,[p]out$)
;	;	6 ; 7	,	$[6].\,+.\,\langle p\rangle.\,[p]out$)
;	;	6 ; $7\cdot 6$	,	$+.\,\langle p\rangle.\,[p]out$)
;	;	6 ; 13	,	$\langle p\rangle.\,[p]out$)
;	;	6 ;	,	$[13]out$)
(13 ;	;	6 ;	,	$\star$)

Fig. 1. Reduction followed by machine evaluation of the term in Example 4.4

Constant operations such as the conditional **if** pop the required number of items from the main stack, and reinstate their result, as is standard for stack languages. For example:

$$\frac{(\;S_A;S_\lambda\cdot 2\cdot 3\;,\;\;+.\,M\;)}{(\;S_A;S_\lambda\cdot 5\quad,\qquad M\;)}\qquad\frac{(\;S_A;S_\lambda\cdot P\cdot N\cdot\bot\;,\;\mathsf{if}.\,M\;)}{(\;S_A;S_\lambda\cdot P\qquad,\qquad M\;)}$$

The FMC then operates similarly to a stack calculus for arithmetic: an expression $1+((2+3)\times 4)$ is given as a term $[4].\,[3].\,[2].\,+.\,\times.\,[1].\,+$ which indeed returns 21. This results in an imperative programming style similar to Haskell's *do*-notation, with the difference that terms may have *any* number of return values, and consume any number of previously returned values.

Example 4.4 Consider the following example, where **rnd** is taken to randomly sample natural numbers.

$$(f\;=\;\mathsf{rand}\;;\;\mathsf{set}\,c\;;\;\mathsf{get}\,c)\;;\;f\;;\;f\;;\;+\;;\;\mathsf{print}$$

The term assigns f to be the function that draws a random number, stores it in cell c, and reads the value at c again as its return value. It then executes f twice, sums the results, and sends that to output. The overall actions should be to take two random inputs i and j, to update the cell c with the last value j, and to output $i+j$. In Figure 1 the term is first interpreted as an FMC-term and reduced to normal form, where each line is a beta-step, and then evaluated on the machine, where the initial memory provides the two expected inputs on **rnd** and one on c. (For compactness we give locations as a header and show only necessary stack elements.)

4.3 Confluence

In demonstrating confluence for the λ-calculus, the difficulty is reduction inside an argument: duplicating or deleting it creates converging reductions of different length. By contrast, *spine reduction*, which reduces in every context *except* argument position, is *diamond* (peaks converge in one step).

While the two extensions of the FMC, *locations* and *sequencing*, do create new configurations of over-lapping redexes, the situation is fundamentally the same. Spine reduction, defined analogously as reduction in every context except argument position, is diamond, and the remaining problem is the same as for the λ-calculus. The calculus thus remains confluent, which will be proved by the standard *parallel reduction* technique of Tait and Martin-Löf (see [2]) and Takahashi [35].

The new configurations are the following. *Sequencing* introduces terms of the form $N\,;x.\,M$, with reduction in M or in N, where the latter may induce substitutions in $x.\,M$. But reduction in N may not duplicate a redex in M, and vice versa, so a peak of this kind converges immediately.

Locations create two new overlapping configurations,

$$\text{nested:} \qquad [N]a.\,[P]b.\,b\langle y\rangle.\,a\langle x\rangle.\,M$$

$$\text{interleaved:} \qquad [N]a.\,[P]b.\,a\langle x\rangle.\,b\langle y\rangle.\,M$$

but both resolve immediately: in each case the two reducts will converge in one step to $\{N/x\}\{P/y\}M$.

Because of this, the problem of confluence amounts to the problem of reduction in arguments, which may be duplicated or deleted, as it does in the regular λ-calculus. We formalize this observation in the following proposition (which is independent of the confluence result). *Spine reduction* is given by closing the β-step under all contexts except in argument position. That is, reduction in each of $x.\,M$, $[N]a.\,M$, and $\langle x\rangle.\,M$ may take place in M, but not in N.

Proposition 4.5 *Spine reduction is diamond.*

The full confluence proof is a standard application of parallel reduction, and is given in Appendix A.

Theorem 4.6 *Reduction $\twoheadrightarrow$ is confluent.*

4.4 Simple types for the FMC

As with sequential types, FMC types will represent the input/output behaviour of the machine. Since a memory is a family of stacks, types will use families of type vectors.

Definition 4.7 *FMC-types* ρ, σ, τ, υ are given by:

$$\tau \;::=\; \overleftarrow{\sigma}_A \Rightarrow \overrightarrow{\tau}_A \qquad \overrightarrow{\tau}_A \;::=\; \{\overrightarrow{\tau}_a \mid a \in A\} \qquad \overrightarrow{\tau} \;::=\; \tau_1 \ldots \tau_n$$

The typing rules for the FMC are in Figure 2. They use the following notation. *Concatenation* of two vector families is pointwise, $\overrightarrow{\sigma}_A \overrightarrow{\tau}_A = \{\overrightarrow{\sigma}_a \overrightarrow{\tau}_a \mid a \in A\}$, and he empty vector is denoted ε. A *slice* $\overrightarrow{\tau}_A|_a$ of a family $\overrightarrow{\tau}_A = \{\overrightarrow{\tau}_a \mid a \in A\}$ is the vector $\overrightarrow{\tau}_a$, and a slice τ_a of a type $\tau = \overleftarrow{\rho}_A \Rightarrow \overrightarrow{\sigma}_A$ is the type $\overleftarrow{\rho}|_a \Rightarrow \overrightarrow{\sigma}|_a$ (i.e. τ_a restricts τ to a single location a). *Composition* is slice-wise: $\sigma \cdot \tau = \{\sigma|_a \cdot \tau|_a \mid a \in A\}$. Finally, a *singleton* $a(\overrightarrow{\tau})$ is a type vector at a single location, with all other locations empty, defined by $a(\overrightarrow{\tau})|_a = \overrightarrow{\tau}$ and $a(\overrightarrow{\tau})|_b = \varepsilon$ for $a \neq b$. A singleton $\lambda(\overrightarrow{\tau})$ on the main location λ may be written as $\overrightarrow{\tau}$.

Poly-types embed as types of the form $\overleftarrow{\tau}_A \Rightarrow$ by the definitions $o \overset{\triangle}{=} (\Rightarrow)$ and $a(\sigma) \rightarrow (\overleftarrow{\tau}_A \Rightarrow) \overset{\triangle}{=} a(\sigma)\,\overleftarrow{\tau}_A \Rightarrow$, and sequential types embed directly as types over only the location λ. The properties proved of sequential types in Section 3 carry over straightforwardly: strict composition, expansion, composition, substitution, subject reduction, and termination of the machine.

Example 4.8 The *singleton* construct $a(\overrightarrow{\tau})$ gives a natural way of writing types in practice. The term from Example 4.4 may be typed as follows, where $\mathbb{Z}$ is a base type of integers.

$$(f = \text{rand} \;;\; \text{set } c \;;\; \text{get } c) \;;\; f \;;\; f \;;\; + \;;\; \text{print} \;:\; \text{rnd}(\mathbb{Z}\;\mathbb{Z})\, c(\mathbb{Z}) \Rightarrow c(\mathbb{Z})\, \text{out}(\mathbb{Z})$$

$$\frac{}{\Gamma \vdash \star \colon \overleftarrow{\tau}_A \Rightarrow \overrightarrow{\tau}_A} \; \star \qquad\qquad \frac{\Gamma \vdash N \colon \rho \qquad \Gamma \vdash M \colon a(\rho)\, \overleftarrow{\sigma}_A \Rightarrow \overrightarrow{\tau}_A}{\Gamma \vdash [N]a.\, M \colon \overleftarrow{\sigma}_A \Rightarrow \overrightarrow{\tau}_A} \; \text{app}$$

$$\frac{\Gamma,\, x \colon \overleftarrow{\rho}_A \Rightarrow \overrightarrow{\sigma}_A \vdash \quad M \colon \overleftarrow{\sigma}_A\, \overleftarrow{\tau}_A \Rightarrow \overrightarrow{\upsilon}_A}{\Gamma,\, x \colon \overleftarrow{\rho}_A \Rightarrow \overrightarrow{\sigma}_A \vdash x.\, M \colon \overleftarrow{\rho}_A\, \overleftarrow{\tau}_A \Rightarrow \overrightarrow{\upsilon}_A} \; \text{var} \qquad\qquad \frac{\Gamma,\, x \colon \rho \vdash M \colon \overleftarrow{\sigma}_A \Rightarrow \overrightarrow{\tau}_A}{\Gamma \vdash a\langle x\rangle.\, M \colon a(\rho)\, \overleftarrow{\sigma}_A \Rightarrow \overrightarrow{\tau}_A} \; \text{abs}$$

Fig. 2. Typing rules for the Functional Machine Calculus

The type expresses that the term pops two integers from rnd and one from c, and pushes one integer to c and one to out. (Note that there are other ways of writing the same type, since singleton types on different locations may permute.) Below left are the types of the defined subterms, and the full type is built up below right by composing these.

$$
\begin{aligned}
+ &: \mathbb{Z}\,\mathbb{Z} \Rightarrow \mathbb{Z} & \mathsf{rand}\,;\mathsf{set}\; c &: \mathsf{rnd}(\mathbb{Z})\; c(\mathbb{Z}) \Rightarrow c(\mathbb{Z}) \\
\mathsf{rand} = \quad \mathsf{rnd}\langle x\rangle.\,[x] &: \mathsf{rnd}(\mathbb{Z}) \Rightarrow \mathbb{Z} & \mathsf{rand}\,;\mathsf{set}\; c\,;\mathsf{get}\; c &: \mathsf{rnd}(\mathbb{Z})\; c(\mathbb{Z}) \Rightarrow c(\mathbb{Z})\, \mathbb{Z} \\
\mathsf{print} = \quad \langle x\rangle.\,[x]\mathsf{out} &: \mathbb{Z} \Rightarrow \mathsf{out}(\mathbb{Z}) & (f = \mathsf{rand}\,;\mathsf{set}\; c\,;\mathsf{get}\; c) &: (\Rightarrow) \\
\mathsf{set}\; c = \quad \langle x\rangle.\,c\langle\_\rangle.\,[x]c &: \mathbb{Z}\; c(\mathbb{Z}) \Rightarrow c(\mathbb{Z}) & (f = \ldots)\,;f\,;f &: \mathsf{rnd}(\mathbb{Z}\,\mathbb{Z})\; c(\mathbb{Z}) \Rightarrow c(\mathbb{Z})\, \mathbb{Z}\,\mathbb{Z} \\
\mathsf{get}\; c = \quad c\langle x\rangle.\,[x]c.\,[x] &: c(\mathbb{Z}) \Rightarrow c(\mathbb{Z})\, \mathbb{Z} & (f = \ldots)\,;f\,;f\,;\,+\,;\mathsf{print} &: \mathsf{rnd}(\mathbb{Z}\,\mathbb{Z})\; c(\mathbb{Z}) \Rightarrow c(\mathbb{Z})\, \mathsf{out}(\mathbb{Z})
\end{aligned}
$$

The third term is only a redex, which pushes and then pops a value, giving it an empty type $(\Rightarrow)$. The type of the fourth term is that of $f\,;f$. Separating the self-composition of the type of f for each location exposes how the inputs on rnd and outputs on λ accumulate, while the output and input on c interact:

$$
\begin{aligned}
\mathsf{rnd}(\mathbb{Z}) \Rightarrow &\qquad \cdot \quad \mathsf{rnd}(\mathbb{Z}) \Rightarrow &\quad = \mathsf{rnd}(\mathbb{Z}\,\mathbb{Z}) \Rightarrow \\
c(\mathbb{Z}) \Rightarrow c(\mathbb{Z})\; \cdot &\qquad c(\mathbb{Z}) \Rightarrow c(\mathbb{Z}) = &\quad c(\mathbb{Z}) \Rightarrow c(\mathbb{Z}) \\
\Rightarrow \mathbb{Z} &\qquad \cdot \qquad\quad \Rightarrow \mathbb{Z} \quad = &\quad \Rightarrow \mathbb{Z}\,\mathbb{Z}
\end{aligned}
$$

By way of illustration, in Haskell the same example may be written as follows.

```
example ::  RandT StdGen (StateT Int IO) ()
example = do
  let f = do x <- rand; lift (put x); lift get
  y <- f
  z <- f
  lift (lift (print (y+z)))
```

To combine the effects of I/O, state, and random generation, the Haskell example uses a stack of monad transformers. The effects are then layered in a fixed order, and to access each effect the function `lift` is used to move to the next layer. Thus, the `rand` action requires no lifting since the random generator is at the top of the transformer stack; the state actions `put` and `get` must be lifted once; and the `print` action must be lifted twice, since I/O is at the bottom of the transformer stack.

The example highlights the following differences between monad transformers and the FMC. Firstly, reader/writer effects in the FMC combine seamlessly, without requiring their organisation in a stack, and are accessed by their locations instead of a lifting function. Secondly, sequential composition in the FMC is a basic operation, with a close connection between the syntax and its execution, where Haskell `do`-notation is syntactic sugar for more involved monadic operations. Thirdly, the FMC type system accounts for the effectful operations that a term performs, where monad transformer types indicate which effects may be present, but not which operations they perform.

Of course, where monads are universal, the FMC is presently restricted to reader/writer effects. How to broaden the range of effects covered in the FMC is an important direction for future work.

5   Further work

We have given an exploratory overview of the Functional Machine Calculus with the most essential results: the natural capture of algebraic laws for effects by reductions and permutations; the encoding of related formalisms for controlling execution behaviour such as monadic constructs, cbpv, κ-calculus, and Arrows; confluence; and termination of the machine with simple types. A forthcoming paper will strengthen these results with domain-theoretic and categorical semantics, and strong normalization with simple types.

Present and future work aims to extend the FMC beyond reader/writer effects and with standard features. One direction is to include sum types, datatypes, and error handling, where it looks possible to capture all three in a uniform way. A second direction is to introduce parallel composition and explore the relation with process calculi, where our type system promises to give something closely related to session types. A third direction is local mutable store, by introducing a *new* construct for locations, and generalizing locations to *regions* to capture mutable data structures (arrays, graphs), which then leads naturally into an exploration of dependent types for the FMC. A fourth direction is to explore the connection with string diagrams, and to introduce constructs to capture diagrammatic reasoning, for example interaction nets or quantum diagrammatic systems.

An important theoretical challenge is type inference. This appears to be open also for the sequential λ-calculus: existing algorithms for concatenative languages are limited [33,6], and (we believe) unable to find the type for the self-application in Example 3.6.

Acknowledgements

References

[1] Ahman, D. and S. Staton, *Normalization by evaluation and algebraic effects*, Electronic Notes in Theoretical Computer Science **298**, pages 51–69 (2013).
https://doi.org/10.1016/j.entcs.2013.09.007

[2] Barendregt, H. P., *The Lambda Calculus – Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*, North-Holland (1984).
https://doi.org/10.1016/c2009-0-14341-6

[3] Benton, N. and P. Wadler, *Linear logic, monads and the lambda calculus*, in: *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 420–431, IEEE (1996).
https://doi.org/10.1109/LICS.1996.561458

[4] Dal Lago, U., G. Guerrieri and W. Heijltjes, *Decomposing probabilistic lambda-calculi*, in: J. Goubault-Larrecq and B. König, editors, *Foundations of Software Science and Computation Structures*, volume 12077 of *LNCS*, pages 136–156, Springer International Publishing, Cham (2020), ISBN 978-3-030-45231-5.
https://doi.org/10.1007/978-3-030-45231-5_8

[5] de Bruijn, N. G., *Algorithmic definition of lambda-typed lambda calculus*, in: G. Huet and G. Plotkin, editors, *Logical Environments*, pages 131–146, Cambridge University Press (1993).

[6] Diggins, C., *Simple type inference for higher-order stack-oriented languages*, Technical Report Cat-TR-2008-001 (2008).

[7] Douence, R. and P. Fradet, *A systematic study of functional language implementations*, ACM Transactions on Programming Languages and Systems **20**, pages 344–387 (1998).
https://doi.org/10.1145/276393.276397

[8] Egger, J., R. E. Møgelberg and A. Simpson, *The enriched effect calculus: syntax and semantics*, Journal of Logic and Computation **24**, pages 615–654 (2014).
https://doi.org/10.1093/logcom/exs025

[9] Ehrhard, T. and G. Guerrieri, *The bang calculus: An untyped lambda-calculus generalizing call-by-name and call-by-value*, in: *Proceedings of the 18th International Symposium on Principles and Practice of Declarative Programming (PPDP'16)*, pages 174–187 (2016).
https://doi.org/10.1145/2967973.2968608

[10] Filinski, A., *Controlling Effects*, Ph.D. thesis, School of Computer Science, Carnegie Mellon University (1996). Available at https://www.cs.cmu.edu/~rwh/students/filinski.pdf.

[11] Hasegawa, M., *Decomposing typed lambda-calculus into a couple of categorical programming languages*, in: *Category Theory and Computer Science (CTCS 1995)*, volume 953 of *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg (1995).
https://doi.org/10.1007/3-540-60164-3_28

[12] Herzberg, D. and T. Reichert, *Concatenative programming: An overlooked paradigm in functional programming*, in: *4th International Conference on Software and Data Technologies (ICSOFT)* (2009).

[13] Hirschkoff, D., E. Prebet and D. Sangiorgi, *On the representation of references in the pi-calculus*, in: *31st International Conference on Concurrency Theory (CONCUR)*, LIPIcs, pages 34:1–34:20, Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020).
https://doi.org/10.4230/LIPIcs.CONCUR.2020.34

[14] Hughes, J., *Generalising monads to arrows*, Science of Computer Programming **37**, pages 67–111 (2000).
https://doi.org/10.1016/S0167-6423(99)00023-4

[15] Krivine, J.-L., *A call-by-name lambda-calculus machine*, Higher-Order and Symbolic Computation **20**, pages 199–207 (2007).
https://doi.org/10.1007/s10990-007-9018-9

[16] Landin, P. J., *The mechanical evaluation of expressions*, The Computer Journal **6**, pages 308–320 (1964).
https://doi.org/10.1093/comjnl/6.4.308

[17] Landin, P. J., *Correspondence between Algol 60 and Church's lambda-notation: Part I*, Communications of the ACM **8**, pages 89–101 (1965).
https://doi.org/10.1145/363744.363749

[18] Levy, P. B., *Call-by-push-value: A functional/imperative synthesis*, volume 2 of *Semantic Structures in Computation*, Springer Netherlands (2003).
https://doi.org/10.1007/978-94-007-0954-6

[19] Levy, P. B., *Call-by-push-value: Decomposing call-by-value and call-by-name*, Higher-Order and Symbolic Computation **19**, pages 377–414 (2006).
https://doi.org/10.1007/s10990-006-0480-6

[20] Levy, P. B., J. Power and H. Thielecke, *Modelling environments in call–by–value programming languages*, Information and Computation **185**, pages 182–210 (2003).
https://doi.org/10.1016/S0890-5401(03)00088-9

[21] Lynas, A. R. and B. Stoddart, *Adding lambda expressions to Forth*, in: *22nd EuroForth Conference* (2006). Available at http://www.euroforth.org/ef06/lynas-stoddart06.pdf.

[22] Maraist, J., M. Odersky, D. N. Turner and P. Wadler, *Call-by-name, call-by-value, call-by-need, and the linear lambda calculus*, Theoretical Computer Science **228**, pages 175–210 (1999).
https://doi.org/10.1016/S0304-3975(98)00358-2

[23] Milner, R., *Functions as processes*, Mathematical Structures in Computer Science **2**, pages 119–141 (1992).
https://doi.org/10.1017/S0960129500001407

[24] Moggi, E., *Computational lambda-calculus and monads*, in: *Fourth Annual Symposium on Logic in Computer Science*, pages 14–23, IEEE (1989).
https://doi.org/10.1109/LICS.1989.39155

[25] Moggi, E., *Notions of computation and monads*, Information and Computation **93**, pages 55–92 (1991).
https://doi.org/10.1016/0890-5401(91)90052-4

[26] Pestov, S., D. Ehrenberg and J. Groff, *Factor: A dynamic stack-based programming language*, ACM SIGPLAN Notices **45**, pages 43–58 (2010).
https://doi.org/10.1145/1899661.1869637

[27] Plotkin, G. and J. Power, *Notions of computation determine monads*, in: *International Conference on Foundations of Software Science and Computation Structures (FoSSaCS)*, volume 2303 of *Lecture Notes in Computer Science*, pages 342–356, Springer, Berlin, Heidelberg (2002).
https://doi.org/10.1007/3-540-45931-6_24

[28] Plotkin, G. D., *Call-by-name, call-by-value and the λ-calculus*, Theoretical Computer Science **1**, pages 125–159 (1975).
https://doi.org/10.1016/0304-3975(75)90017-1

[29] Plotkin, G. D. and M. Pretnar, *Handlers of algebraic effects*, in: *European Symposium on Programming (ESOP)*, volume 5502 of *Lecture Notes in Computer Science*, pages 80–94, Springer, Berlin, Heidelberg (2009).
https://doi.org/10.1007/978-3-642-00590-9_7

[30] Power, A. and H. Thielecke, *Closed Freyd- and κ-categories*, in: *International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 1644 of *Lecture Notes in Computer Science*, pages 625–634, Springer (1999).
https://doi.org/10.1007/3-540-48523-6_59

[31] Power, J. and E. Robinson, *Premonoidal categories and notions of computation*, Mathematical Structures in Computer Science **7**, pages 453–468 (1997).
https://doi.org/10.1017/S0960129597002375

[32] Smetsers, J., E. Barendsen, M. v. Eekelen and M. Plasmeijer, *Guaranteeing safe destructive updates through a type system with uniqueness information for graphs*, in: *Workshop Graph Transformations in Computer Science*, volume 776 of *Lecture Notes in Computer Science*, Schloss Dagstuhl, Springer-Verlag, Berlin (1993).
https://doi.org/10.1007/3-540-57787-4_23

[33] Stoddart, B. and P. J. Knaggs, *Type inference in stack based languages*, Formal Aspects of Computing **5**, pages 289–298 (1993).
https://doi.org/10.1007/BF01212404

[34] Tait, W., *Intensional interpretations of functionals of finite type I*, The Journal of Symbolic Logic **32**, pages 198–212 (1967).
https://doi.org/10.2307/2271658

[35] Takahashi, M., *Parallel reductions in lambda-calculus*, Information and Computation **118**, pages 120–127 (1995).
https://doi.org/10.1006/inco.1995.1057

[36] von Thun, M., *Joy: Forth's functional cousin*, in: *Proc. 17th EuroForth Conference* (2001).

[37] Winskel, G., *The formal semantics of programming languages: An introduction*, MIT Press, Cambridge, Massachusetts (1993).
https://doi.org/10.7551/mitpress/3054.001.0001

A   The confluence proof

This section gives a complete proof of confluence using parallel reduction. The particular approach is to extend the syntax of terms with marked redexes, to define parallel reduction and to identify residuals.

Definition A.1 A *marked redex* is one whose application and abstraction are marked by the symbols (◂) and (▸), as follows.

$$[N]a\text{◂}.\,H.\,\text{▸}a\langle x\rangle.\,M \ .$$

A *redex-marked* or *marked* term $M^{\blacklozenge}$ is one where a selection of redexes is marked. For a marked term consisting of a head context and a term $(H.\,M)^{\blacklozenge}$, the marking of both components separately is indicated by $H^{\blacklozenge}$ and $M^{\blacklozenge}$. The *marked reduct* $(M^{\blacklozenge})_{\blacklozenge}$ of a marked term $M^{\blacklozenge}$ is defined as follows.

$$(\star)_{\blacklozenge} = \star \qquad ([N^{\blacklozenge}]a.\,M^{\blacklozenge})_{\blacklozenge} = [(N^{\blacklozenge})_{\blacklozenge}]a.\,(M^{\blacklozenge})_{\blacklozenge}$$

$$(x.\,M^{\blacklozenge})_{\blacklozenge} = x.\,(M^{\blacklozenge})_{\blacklozenge} \qquad (a\langle x\rangle.\,M^{\blacklozenge})_{\blacklozenge} = a\langle x\rangle.\,(M^{\blacklozenge})_{\blacklozenge}$$

$$([N^{\blacklozenge}]a\text{◂}.\,H^{\blacklozenge}.\,\text{▸}a\langle x\rangle.\,M^{\blacklozenge})_{\blacklozenge} = \{(N^{\blacklozenge})_{\blacklozenge}/x\}(H^{\blacklozenge}.\,M^{\blacklozenge})_{\blacklozenge}$$

A *parallel reduction step* $M^\blacklozenge \Rightarrow (M^\blacklozenge)_\bullet$ relates a marked term to its marked reduct, and an unmarked term to all its marked reducts: $M \Rightarrow (M^\blacklozenge)_\bullet$ for each marking $M^\blacklozenge$ of M.

To reduce clutter, a marked reduct $(M^\blacklozenge)_\bullet$ may be abbreviated as $M_\bullet$.

Proposition A.2 *A parallel reduction step is a beta-reduction:* $M \twoheadrightarrow (M^\blacklozenge)_\bullet$ *whenever* $M^\blacklozenge$ *is a marking of* M.

Proof. By induction on the size of $M^\blacklozenge$.

- Unit case: immediate by $(\star)_\bullet = \star$.
- Variable case: if $M^\blacklozenge \twoheadrightarrow M_\bullet$ then $x.\,M^\blacklozenge \twoheadrightarrow x.\,M_\bullet$.
- Unmarked application case: if $M^\blacklozenge \twoheadrightarrow M_\bullet$ and $N^\blacklozenge \twoheadrightarrow N_\bullet$ then $[N^\blacklozenge]a.\,M^\blacklozenge \twoheadrightarrow [N_\bullet]a.\,M_\bullet$.
- Unmarked abstraction case: if $M^\blacklozenge \twoheadrightarrow M_\bullet$ then $a\langle x\rangle.\,M^\blacklozenge \twoheadrightarrow a\langle x\rangle.\,M_\bullet$.
- Redex case: if $(H.\,M)^\blacklozenge \twoheadrightarrow (H.\,M)_\bullet$ and $N^\blacklozenge \twoheadrightarrow N_\bullet$ then

$$
\begin{aligned}
[N^\blacklozenge]a\triangleleft.\,H^\backprime.\triangleright a\langle x\rangle.\,M^\triangleright \ &\twoheadrightarrow\ [N_\bullet]a\triangleleft.\,H^\backprime.\triangleright a\langle x\rangle.\,M^\triangleright \\
&\twoheadrightarrow\ H^\backprime.\,\{N_\bullet/x\}M^\triangleright \\
&=\ \{N_\bullet/x\}(H^\backprime.\,M^\triangleright) \\
&\twoheadrightarrow\ \{N_\bullet/x\}(H.\,M)_\bullet\ .
\end{aligned}
$$

where in the second line, $a \in \mathsf{loc}(H)$ by the definition of a redex, and for the equality on the third line, H does not bind in N and x is not free in H, by α-equivalence.

$\square$

To show that parallel reduction is diamond, a term is reduced according to two markings, $\blacklozenge$ and $\diamond$. These are then applied simultaneously and interchangeably, and may be considered a single marking $\blacklozenge\diamond = \diamond\blacklozenge$:

$$
(M^\blacklozenge)^\diamond = M^{\blacklozenge\diamond} = (M^\diamond)^\blacklozenge
$$

If this term is reduced relative to one marking, the other is preserved: $(M^{\blacklozenge\diamond})_\diamond = ((M^\diamond)_\diamond)^\blacklozenge$. The proof then amounts to showing that reducing by each marking is commutative and the same as reducing along both markings simultaneously:

$$
(M_\bullet)_\diamond = M_{\bullet\diamond} = (M_\diamond)_\bullet
$$

Lemma A.3 *Parallel reduction commutes with composition:* $M_\bullet\,;N_\bullet = (M\,;N)_\bullet$.

Proof. By induction on the size of M.

- Unit case: $(\star)_\bullet\,;N_\bullet =\ \star\,;N_\bullet = N_\bullet = (\star\,;N)_\bullet$.
- Variable case: if $M_\bullet\,;N_\bullet = (M\,;N)_\bullet$ then

$$
(x.\,M)_\bullet\,;N_\bullet = x.\,(M_\bullet\,;N_\bullet) = x.\,(M\,;N)_\bullet = ((x.\,M)\,;N)_\bullet\ .
$$

- Unmarked application case: if $M_\bullet\,;N_\bullet = (M\,;N)_\bullet$ then

$$
\begin{aligned}
([P]a.\,M)_\bullet\,;N_\bullet \ &=\ [P_\bullet]a.\,(M_\bullet\,;N_\bullet) \\
&=\ [P_\bullet]a.\,(M\,;N)_\bullet \\
&=\ (([P]a.\,M)\,;N)_\bullet\ .
\end{aligned}
$$

- Unmarked abstraction case: if $M_\bullet \mathbin{;} N_\bullet = (M \mathbin{;} N)_\bullet$ then

$$
\begin{aligned}
(a\langle x\rangle.\, M)_\bullet \mathbin{;} N_\bullet &= a\langle x\rangle.\, (M_\bullet \mathbin{;} N_\bullet) \\
&= a\langle x\rangle.\, (M \mathbin{;} N)_\bullet \\
&= ((a\langle x\rangle.\, M) \mathbin{;} N)_\bullet\ .
\end{aligned}
$$

- Marked redex case: if $(H.\, M)_\bullet \mathbin{;} N_\bullet = ((H.\, M) \mathbin{;} N)_\bullet$ and $x \notin \mathsf{fv}(N)$ then

$$
\begin{aligned}
([P]a\!\shortmid.\, H.\!\blacktriangleright a\langle x\rangle.\, M)_\bullet \mathbin{;} N_\bullet &= (\{P_\bullet/x\}(H.\, M)_\bullet) \mathbin{;} N_\bullet \\
&= \{P_\bullet/x\}((H.\, M) \mathbin{;} N)_\bullet \\
&= (((\,[P]a\!\shortmid.\, H.\!\blacktriangleright a\langle x\rangle.\, M) \mathbin{;} N)_\bullet\ .
\end{aligned}
$$

$\square$

Lemma A.4 *Parallel reduction commutes with substitution:* $\{N_\bullet/x\}M_\bullet = (\{N/x\}M)_\bullet$.

Proof. By induction on the size of $M^\blacklozenge$.

- Unit case: immediate by $\{N/x\}\star = \star$.
- Variable case: if $\{N_\bullet/x\}M_\bullet = (\{N/x\}M)_\bullet$ then by Lemma A.3,

$$
\begin{aligned}
\{N_\bullet/x\}(x.\, M)_\bullet &= \{N_\bullet/x\}(x.\, M_\bullet) \\
&= N_\bullet \mathbin{;} \{N_\bullet/x\}M_\bullet \\
&= N_\bullet \mathbin{;} (\{N/x\}M)_\bullet \\
&= (N \mathbin{;} \{N/x\}M)_\bullet \\
&= (\{N/x\}(x.\, M))_\bullet\ .
\end{aligned}
$$

- Unmarked application case: if $\{N_\bullet/x\}M_\bullet = (\{N/x\}M)_\bullet$ and $\{N_\bullet/x\}P_\bullet = (\{N/x\}P)_\bullet$ then

$$
\begin{aligned}
\{N_\bullet/x\}([P]a.\, M)_\bullet &= \{N_\bullet/x\}([P_\bullet]a.\, M_\bullet) \\
&= [\{N_\bullet/x\}P_\bullet]a.\, \{N_\bullet/x\}M_\bullet \\
&= [(\{N/x\}P)_\bullet]a.\, (\{N/x\}M)_\bullet \\
&= ([\{N/x\}P]a.\, \{N/x\}M)_\bullet\ .
\end{aligned}
$$

- Unmarked abstraction case: if $\{N_\bullet/x\}M_\bullet = (\{N/x\}M)_\bullet$ then

$$
\begin{aligned}
\{N_\bullet/x\}(a\langle y\rangle.\, M)_\bullet &= \{N_\bullet/x\}(a\langle y\rangle.\, M_\bullet) \\
&= a\langle y\rangle.\, \{N_\bullet/x\}M_\bullet \\
&= a\langle y\rangle.\, (\{N/x\}M)_\bullet \\
&= (a\langle y\rangle.\, \{N/x\}M)_\bullet\ .
\end{aligned}
$$

- Marked redex case: if $\{N_\bullet/x\}(H.\,M)_\bullet = (\{N/x\}(H.\,M))_\bullet$ and $\{N_\bullet/x\}P_\bullet = (\{N/x\}P)_\bullet$ then

$$
\begin{aligned}
\{N_\bullet/x\}&([P]a\triangleleft.\,H.\,\triangleright a\langle y\rangle.\,M)_\bullet \\
&= \{N_\bullet/x\}\{P_\bullet/y\}(H.\,M)_\bullet \\
&= \{\{N_\bullet/x\}P_\bullet/y\}\{N_\bullet/x\}(H.\,M)_\bullet \\
&= \{(\{N/x\}P)_\bullet/y\}(\{N/x\}(H.\,M))_\bullet \\
&= \{(\{N/x\}P)_\bullet/y\}((\{N/x\}H).\,\{N/x\}M)_\bullet \\
&= ([\{N/x\}P]a\triangleleft.\,(\{N/x\}H).\,\triangleright a\langle y\rangle.\,\{N/x\}M)_\bullet \\
&= (\{N/x\}([P]a\triangleleft.\,H.\,\triangleright a\langle y\rangle.\,M))_\bullet\ .
\end{aligned}
$$

$\square$

Lemma A.5 *For a doubly marked term, $M^{\bullet\diamond}$, reducing each marking in turn gives the same result as reducing both simultaneously: $(M_\bullet)_\diamond = M_{\bullet\diamond}$.*

Proof. By induction on the size of M.

- Unit case: immediate by $((\star)_\bullet)_\diamond = \star = (\star)_{\bullet\diamond}$.
- Variable case: if $(M_\bullet)_\diamond = M_{\bullet\diamond}$ then

$$
((x.\,M)_\bullet)_\diamond = x.\,(M_\bullet)_\diamond = x.\,M_{\bullet\diamond} = (x.\,M)_{\bullet\diamond}\ .
$$

- Unmarked application case: if $(M_\bullet)_\diamond = M_{\bullet\diamond}$ and $(N_\bullet)_\diamond = N_{\bullet\diamond}$ then

$$
\begin{aligned}
(([N]a.\,M)_\bullet)_\diamond &= [(N_\bullet)_\diamond]a.\,(M_\bullet)_\diamond \\
&= [N_{\bullet\diamond}]a.\,M_{\bullet\diamond} \\
&= ([N]a.\,M)_{\bullet\diamond}\ .
\end{aligned}
$$

- Unmarked abstraction case: if $(M_\bullet)_\diamond = M_{\bullet\diamond}$ then

$$
((a\langle x\rangle.\,M)_\bullet)_\diamond = a\langle x\rangle.\,(M_\bullet)_\diamond = a\langle x\rangle.\,M_{\bullet\diamond} = (a\langle x\rangle.\,M)_{\bullet\diamond}\ .
$$

- Doubly-marked redex case: if $(H.\,M_\bullet)_\diamond = H.\,M_{\bullet\diamond}$ and $(N_\bullet)_\diamond = N_{\bullet\diamond}$ then

$$
\begin{aligned}
([N]a\triangleleft\!.\,H.\,\twoheadrightarrow a\langle x\rangle.\,M)_\bullet)_\diamond &= \{(N_\bullet)_\diamond/x\}((H.\,M)_\bullet)_\diamond \\
&= \{N_{\bullet\diamond}/x\}(H.\,M)_{\bullet\diamond} \\
&= ([N]a\triangleleft\!.\,H.\,\twoheadrightarrow a\langle x\rangle.\,M)_{\bullet\diamond}\ .
\end{aligned}
$$

- First singly-marked redex case: if $(H.\,M_\bullet)_\diamond = (H.\,M)_{\bullet\diamond}$ and $(N_\bullet)_\diamond = N_{\bullet\diamond}$ then, using Lemma A.4,

$$
\begin{aligned}
(([N]a\triangleleft.\,H.\,\triangleright a\langle x\rangle.\,M)_\bullet)_\diamond &= (\{N_\bullet/x\}(H.\,M)_\bullet)_\diamond \\
&= \{(N_\bullet)_\diamond/x\}((H.\,M)_\bullet)_\diamond \\
&= \{N_{\bullet\diamond}/x\}(H.\,M)_{\bullet\diamond} \\
&= ([N]a\triangleleft.\,H.\,\triangleright a\langle x\rangle.\,M)_{\bullet\diamond}\ .
\end{aligned}
$$

- Second singly-marked redex case: the redex considered is

$$
[N^{\bullet\diamond}]a\triangleleft.\,H^{\triangleleft\!\triangleleft}.\,\triangleright a\langle x\rangle.\,M^{\triangleright\!\triangleright}
$$

where H does not use the location a or a free variable x. First, it is established that $(H^{\triangleleft}.\triangleright a\langle x\rangle.M^{\triangleright})_{\blacklozenge}$ is of the form $K^{\triangleleft}.\triangleright a\langle x\rangle.N^{\triangleright}$ where $(H^{\triangleleft}.M^{\triangleright})_{\blacklozenge} = K^{\triangleleft}.N^{\triangleright}$ for some head context K and term N. For readability the marking $\diamond$ will be suppressed, except on the abstraction $\triangleright a\langle x\rangle$. The proof is by induction on the size of H.

· Unit case: if $H = \{\}$ then let $K = \{\}$ and $N = M_{\blacklozenge}$, which gives the following, as required.

$$
\begin{aligned}
(H.\triangleright a\langle x\rangle.M)_{\blacklozenge} &= (\triangleright a\langle x\rangle.M)_{\blacklozenge} & (H.M)_{\blacklozenge} &= M_{\blacklozenge} \\
&= \triangleright a\langle x\rangle.M_{\blacklozenge} & &= N \\
&= K.\triangleright a\langle x\rangle.N & &= K.N \ .
\end{aligned}
$$

· Unmarked application case: if $H = [P]b.H'$ then the inductive hypothesis for H' and M gives K' and N. Let $K = [P]b.H'$, which gives the following, as required.

$$
\begin{aligned}
([P]b.H'.\triangleright a\langle x\rangle.M)_{\blacklozenge} & & ([P]b.H'.M)_{\blacklozenge} \\
= [P_{\blacklozenge}]b.(H'.\triangleright a\langle x\rangle.M)_{\blacklozenge} & & = [P_{\blacklozenge}]b.(H'.M)_{\blacklozenge} \\
= [P_{\blacklozenge}]b.K'.\triangleright a\langle x\rangle.N & & = [P_{\blacklozenge}]b.K'.N \\
= K.\triangleright a\langle x\rangle.N & & = K.N \ .
\end{aligned}
$$

· Unmarked abstraction case: if $H = b\langle y\rangle.H'$ then the inductive hypothesis for H' and M gives K' and N. Let $K = b\langle y\rangle.H'$, which gives the following, as required.

$$
\begin{aligned}
(b\langle y\rangle.H'.\triangleright a\langle x\rangle.M)_{\blacklozenge} & & (b\langle y\rangle.H'.M)_{\blacklozenge} \\
= b\langle y\rangle.(H'.\triangleright a\langle x\rangle.M)_{\blacklozenge} & & = b\langle y\rangle.(H'.M)_{\blacklozenge} \\
= b\langle y\rangle.K'.\triangleright a\langle x\rangle.N & & = b\langle y\rangle.K'.N \\
= K.\triangleright a\langle x\rangle.N & & = K.N \ .
\end{aligned}
$$

· First marked redex case: if $H = [P]b\triangleleft.H'.\triangleright b\langle y\rangle.H''$ then the inductive hypothesis for $H'.H''$ and M gives K' and N'. Let $K = \{P_{\blacklozenge}/x\}K'$ and $N = \{P_{\blacklozenge}/x\}N'$, which gives the following, as required.

$$
\begin{aligned}
([P]b\triangleleft.H'.\triangleright b\langle y\rangle.H''.\triangleright a\langle x\rangle.M)_{\blacklozenge} & & ([P]b\triangleleft.H'.\triangleright b\langle y\rangle.H''.M)_{\blacklozenge} \\
= \{P_{\blacklozenge}/y\}(H'.H''.\triangleright a\langle x\rangle.M)_{\blacklozenge} & & = \{P_{\blacklozenge}/y\}(H'.H''.M)_{\blacklozenge} \\
= \{P_{\blacklozenge}/y\}(K'.\triangleright a\langle x\rangle.N') & & = \{P_{\blacklozenge}/y\}(K'.N') \\
= (\{P_{\blacklozenge}/y\}K').\triangleright a\langle x\rangle.\{P_{\blacklozenge}/y\}N' & & = (\{P_{\blacklozenge}/y\}K').\{P_{\blacklozenge}/y\}N' \\
= K.\triangleright a\langle x\rangle.N & & = K.N \ .
\end{aligned}
$$

· Second marked redex case: if $H = [P]b\triangleleft.H'$ and $M = H''.\triangleright b\langle y\rangle.M'$ then the inductive hypothesis for H' and $H''.M$ gives K' and N'. Let $K = \{P_{\blacklozenge}/x\}K'$ and $N = \{P_{\blacklozenge}/x\}N'$, which gives the

following, as required.

$$([P]b\triangleleft.H'.\triangleright a\langle x\rangle.H''.\blacktriangleright b\langle y\rangle.M)_\bullet \qquad\qquad ([P]b\triangleleft.H'.H''.\blacktriangleright b\langle y\rangle.M)_\bullet$$
$$= \{P_\bullet/x\}(H'.\triangleright a\langle x\rangle.H''.M)_\bullet \qquad\qquad = \{P_\bullet/x\}(H'.H''.M)_\bullet$$
$$= \{P_\bullet/x\}(K'.\triangleright a\langle x\rangle.N') \qquad\qquad = \{P_\bullet/x\}(K'.N')$$
$$= (\{P_\bullet/x\}K').\triangleright a\langle x\rangle.\{P_\bullet/x\}N') \qquad = (\{P_\bullet/x\}K').\{P_\bullet/x\}N')$$
$$= K.\triangleright a\langle x\rangle.N \qquad\qquad\qquad = K.N\ .$$

This concludes the subproof. Returning to the main proof, let $(H.\triangleright a\langle x\rangle.M)_\bullet = K.\triangleright a\langle x\rangle.N$ where $(H.M)_\bullet = K.N$. The case concludes as follows.

$$(([N]a\triangleleft.H.\triangleright a\langle x\rangle.M)_\bullet)_\diamond = ([N_\bullet]a\triangleleft.(H.\triangleright a\langle x\rangle.M)_\bullet)_\diamond$$
$$= ([N_\bullet]a\triangleleft.K.\triangleright a\langle x\rangle.N)_\diamond$$
$$= \{(N_\bullet)_\diamond/x\}(K.N)_\diamond$$
$$= \{N_{\bullet\diamond}/x\}(H.M)_{\bullet\diamond}$$
$$= ([N]a\triangleleft.H.\triangleright a\langle x\rangle.M)_{\bullet\diamond}$$

$\square$

Lemma A.6 *Parallel reduction is diamond: if $N \Longleftarrow M \Longrightarrow P$ then $N \Longrightarrow Q \Longleftarrow P$ for some Q.*

Proof. Let $M^\bullet \Longrightarrow M_\bullet = N$ and $M^\diamond \Longrightarrow M_\diamond = P$ for two separate markings $\bullet$ and $\diamond$ of M. Then with both markings on M the peak becomes

$$(M_\bullet)^\diamond \Longleftarrow M^{\bullet\diamond} \Longrightarrow (M_\diamond)^\bullet\ .$$

Let $Q = M_{\bullet\diamond}$ so that by Lemma A.5 the peak converges as

$$(M_\bullet)^\diamond \Longrightarrow M_{\bullet\diamond} \Longleftarrow (M_\diamond)^\bullet\ .$$

$\square$

Theorem 4.6 (restatement)   Reduction $\rightarrow$ is confluent.

Proof. A reduction step $M \rightarrow N$ is a parallel step $M^\bullet \Longrightarrow M_\bullet = N$ by marking only the reduced redex. A peak $M \twoheadleftarrow N \twoheadrightarrow P$ in regular reduction is then immediately one in parallel reduction, $M \Longleftarrow N \Longrightarrow P$. By the diamond property, Lemma A.6, this converges with parallel reductions, $M \Longrightarrow Q \Longleftarrow P$. By Proposition A.2 a parallel reduction ($\Longrightarrow$) is a regular reduction ($\twoheadrightarrow$), so that the peak converges as $M \twoheadrightarrow Q \twoheadleftarrow P$. $\square$