

Correctness of a Distributed-Memory Model for Scheme

Luc Moreau*

University of Southampton
L.Moreau@ecs.soton.ac.uk

Abstract. We propose a high-level approach to program distributed applications; it is based on the annotation **future** by which the programmer specifies which expressions may be evaluated remotely in parallel. We present the CEKDS-Machine, an abstract machine with a distributed memory, able to evaluate Scheme-like **future**-based programs. In this paper, we focus on the issue of task migration and prove that task migration is transparent to the user, i.e. task migration does not change the observable behaviour of programs.

1 Introduction

Distributed systems are omnipresent: local area networks and the success of the Internet in the past years are particular illustrations of the ubiquity of distributed computing. A major research focus in this area has been the design of new languages or programming paradigms to develop distributed applications, like e.g. PVM [10], MPI [8], Nexus [9], Cilk [1]. We argue that those systems were designed to build high-performance distributed applications, and that they favour efficiency over ease of programming. Therefore, these languages or paradigms overwhelm the programmer with the burden of dealing with the complexity of distribution. Some approaches even impose programming styles, with which the programmer may not be familiar; e.g. Cilk [1] demands programs written in continuation-passing style.

Mostly-functional languages like Scheme and SML have traditionally provided the programmer with abstraction, expressiveness, first-class citizenship of objects, and automatic garbage collection. We believe that there is a niche for a high-level approach to distributed computing. Following Halstead's work on MultiLisp [11], we extend a Scheme-like language with an annotation **future** by which the programmer specifies which expressions may be evaluated in parallel, possibly remotely. By definition, annotations must be *transparent*, i.e. annotated programs return the same result as in the absence of annotations. This approach is abstract because it hides the intricacies of distribution by giving the programmer the illusion that a distributed system is programmable as a sequential one.

We consider the idealised Scheme-like language defined in Figure 1. It is a purely functional language, extended with a primitive **makeref** to create boxes, with primitives **deref**, **setref!** to read and modify them, and with a primitive **callcc** to capture first-class continuations. In addition, there is a construct **(future M)** to create a producer-consumer type of parallelism [11]. Intuitively, the evaluation of **(future M)** immediately returns an object called placeholder, while another task evaluates the argument M in parallel. The purpose of the latter task, called the *producer*, is to compute and then store the value of M in the placeholder. The task using the placeholder is called the *consumer*.

For a long time, this approach has been characterised by a lack of formal semantics due to the difficulty of providing transparent annotations for parallelism in the presence of first-class continuations and side-effects. Recently, the author [20] defined the semantics of **future** for the language of Figure 1. The goal of this paper is to extend this semantics to a distributed framework (the proof of its correctness is available in a technical report [21]).

More specifically, the contributions of this paper are the following. *i)* We define a distributed architecture able to evaluate **future**-based programs; *ii)* We prove that

* This research was supported in part by the Engineering and Physical Sciences Research Council, grant GR/K30773. Author's address: Department of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ. United Kingdom.

$P \in \Lambda_f^0$		(Program)
$M \in \Lambda_f$	$::= V \mid (M \ M) \mid (\text{if } M \ M \ M) \mid (\text{future } M)$	(Term)
$V \in \text{Value}_f$	$::= p \mid f \mid x \mid (\lambda x. M)$	(Syntactic Value)
$p \in BConst$	$= \{\text{true}, \text{false}, \text{nil}, 0, 1, \dots\}$	(Basic Constant)
$f \in FConst$	$= \{\text{cons}, \text{car}, \text{cdr}, \text{makeref}, \text{deref}, \text{setref!}, \text{callcc}\}$	(Functional Constant)
$x \in Vars$	$= \{x, y, z, \dots\}$	(Variable)

Fig. 1. Syntax of Λ_f

task migration is transparent to the programmer, i.e. task migration does not change the observable behaviour of programs; *iii*) This result is the reference semantics that we can use to design and prove the correctness of program optimisations. The architecture is described in Section 2. Section 3 discusses related work and is followed by concluding remarks.

2 The Architecture

In this section, we present the CEKDS-Machine, an abstract machine with a distributed store, which extends Felleisen and Friedman’s CEK and CEKS machines [4, 5], and the F-PCKS-machine [20]; its state space is formally described in Figure 2.

Some operations like **deref**, i.e. reading the content of a box, are rather complex. Indeed, as **deref** is *strict*², it touches its argument, checks whether it is legitimate to access the content of the box received in argument, and finally, reads the box content. In order to distinguish these three operations, we add two primitives **touch** and **sync** to Λ_{ceks} , the language accepted by the CEKDS-machine; besides, we translate every program of Λ_f into a term of Λ_{ceks} by the function \mathcal{X} of Figure 2, which makes the **touch** and **sync** operations explicit.

In our distributed architecture, computational resources are called *sites* and are uniquely identified by *site names*. A site has its own memory and can run several tasks that share the site memory. A *world* is the set of sites that can be used to evaluate a program; sites in a world communicate by exchanging messages. More specifically, in a site, we distinguish active tasks, i.e. tasks that can be run, from suspended tasks, i.e. tasks that wait for a message or a synchronisation. As far as communications are concerned, a site is equipped with two spools of messages: the input spool contains pending input messages, while the output spool contains the messages that remain to be transmitted.

A store is a finite function, also represented as a set of pairs, associating locations with store contents. In our distributed architecture, each site has its own procedure of memory allocation, and its proper task naming mechanism. Hence, we use the notion of *qualified* location or task name, to unambiguously refer to a location or a task in the world. We now appreciate how site-names uniqueness is important to define qualified locations or task names.

We abstract a task by a triple composed of a computational state, a *legitimacy* [15, 20] used to implement first-class continuations and side-effects, and a name. A computational state is a configuration of the CEK-machine [4], which can be either $\text{Ev}\langle M, \rho, \kappa \rangle$ representing the evaluation of a term M in an environment ρ with a continuation κ , or $\text{Ret}\langle V, \kappa \rangle$ meaning the return of a value to a continuation. The continuation, implemented as a data structure called *continuation code*, represents what remains to do after evaluating M in $\text{Ev}\langle M, \rho, \kappa \rangle$. In conventional languages, the continuation is nothing else but the evaluation stack. The environment is a finite function mapping variables to values.

In Figure 3, transitions between computational states specify how to evaluate the purely functional and sequential subset of the language extended with first-class continuations; details can be found in [4, 20, 22, 19]. Figure 4 shows the transitions that involve a collaboration of a task with its site. According to (*fork*), the evaluation of a **future** allocates a new placeholder ph and creates a new task, with a name τ_1 , which speculatively evaluates the

² A strict function applied to a placeholder, accesses the value that the producer task has stored in the placeholder; this operation, called *touching* the placeholder, can suspend the current task when the placeholder has not received a value yet.

\mathcal{W}	$::= \{m_1, \dots, m_n\}$	(World)	Explicit translation: $\mathcal{X} : A_f \rightarrow A_{ceksds}$
$m \in \mathcal{M}$	$::= \langle T, \theta, s, S, I, O \rangle$	(Site)	$\mathcal{X}[x] = x$ if $x \in \text{Vars} \cup \text{BConst} \cup \text{NStP}$
$t \in \text{Task}$	$::= \langle C, \ell, \tau \rangle$	(Task)	$\mathcal{X}[\text{car}] = \lambda x. (\text{car } (\text{touch } x))$
$\theta \in \text{Store}$	$::= \{(\alpha_1 X_1) \dots (\alpha_n X_n)\}$	(Store)	$\mathcal{X}[\text{cdr}] = \lambda x. (\text{cdr } (\text{touch } x))$
X	$::= V \mid \perp \mid \ell$	(Store Content)	$\mathcal{X}[\text{deref}] = \lambda x. (\text{deref } (\text{sync } (\text{touch } x)))$
$s \in \mathcal{S}$	$ = \{s_1, s_2, \dots\}$	(Site Name)	$\mathcal{X}[\text{setref!}] = \lambda x_1 x_2. (\text{setref! } (\text{sync } (\text{touch } x_1)) x_2)$
$\tau \in \mathcal{T}$	$ = \{\tau_1, \tau_2, \dots\}$	(Task Name)	$\mathcal{X}[(\text{future } M)] = (\text{future } \mathcal{X}[M])$
u	$::= \langle \tau, s \rangle$	(Qualified Task Name)	$\mathcal{X}[(M_1 M_2)] = (\lambda m_1 m_2. (\text{touch } m_1) m_2) \mathcal{X}[M_1] \mathcal{X}[M_2]$
a	$::= \langle \alpha, s \rangle$	(Qualified Location)	$\mathcal{X}[(\lambda x. M)] = (\lambda x. \mathcal{X}[M])$
			$\mathcal{X}[(\text{if } M_1 M_2 M_3)] = (\text{if } (\text{touch } \mathcal{X}[M_1]) \mathcal{X}[M_2] \mathcal{X}[M_3])$
$C \in \text{CoSt}$	$::= \text{Ev} \langle M, \rho, \kappa \rangle$	(Computational State)	$\text{Unload}[c, \mathcal{W}] = c$
	$ \mid \text{Ret} \langle V, \kappa \rangle$		$\text{Unload}[(\text{cons } V_1 V_2), \mathcal{W}] = (\text{cons } \text{Unload}[V_1, \mathcal{W}]$
$\rho \in \text{Env}$	$::= \{(x_1 V_1) \dots (x_n V_n)\}$	(Environment)	$\text{Unload}[V_2, \mathcal{W}]$
I	$ = \{q_1, \dots, q_n\}$	(Input Spool)	$\text{Unload}[\lambda x. M, \mathcal{W}] = \text{procedure}$
O	$ = \{q_1, \dots, q_n\}$	(Output Spool)	$\text{Unload}[b, \mathcal{W}] = \text{box}$
q	$::= R \mid A$	(Message)	$\text{Unload}[f_c, \mathcal{W}] = \text{procedure}$
T	$ = \{t_1, \dots, t_n\}$	(Active Tasks Set)	$\text{Unload}[(\text{co } \kappa), \mathcal{W}] = \text{cont}$
S	$ = \{t_1, \dots, t_n\}$	(Suspend Set)	$\text{Unload}[(\text{ph } \alpha, s), \mathcal{W}] = \text{Unload}[\mathcal{W}[(\alpha, s)], \mathcal{W}]$
	$ \cup \{R_1, \dots, R_m\}$		
$M \in A_{ceksds}$	$::= V_s \mid (M M)$	(Term)	Free Task Name:
	$ \mid (\text{if } M M M) \mid (\text{future } M)$		$FN(T) = \{\tau, \langle C, \ell, \tau \rangle \in T\}$
$V_s \in \text{SValue}$	$::= c \mid x \mid (\lambda x. M)$	(Syntactic Value)	
$W \in \text{PValue}$	$::= c \mid \langle \text{cl } \lambda x. M, \rho \rangle \mid f_c$	(Proper Value)	Store Operations:
	$ \mid (\text{cons } V V) \mid \langle \text{co } \kappa \rangle \mid b$		$\theta \uplus \{(\alpha V)\} = \theta \cup \{(\alpha V)\}$
$V \in \text{Value}$	$::= W \mid \text{ph}$	(Runtime Value)	with $\alpha \notin \text{DOM}(\theta)$
ph	$::= \langle \text{ph } \alpha, s \rangle$	(Placeholder)	$\theta(\alpha) = V$ if $(\alpha V) \in \theta$
$b \in \text{Box}$	$::= \langle \text{bx } \alpha, s, \ell \rangle$	(Box)	$\theta[\alpha := V] = (\theta \setminus \{(\alpha \theta(\alpha))\}) \cup \{(\alpha V)\}$
$\ell \in \text{Leg}$	$::= \langle \text{leg } \alpha, s \rangle$	(Legitimacy)	
$c \in \text{Const}$	$::= p \mid f$	(Constant)	
$f_c \in \text{PApp}$	$::= (\text{cons } V) \mid (\text{setref! } V)$	(Partial Application)	Global reference:
$g \in \text{AValue}$	$::= \langle \text{cl } \lambda x. M, \rho \rangle \mid f \mid f_c \mid \langle \text{co } \kappa \rangle$	(Applicable Val.)	$\mathcal{W}[(\alpha, s)] = \theta(s)$
$p \in \text{BConst}$	$ = \{\text{true}, \text{false}, \text{nil}, 0, 1, \dots, \text{void}\}$	(Basic Constant)	if $\exists \langle T, s, \theta, S, I, O \rangle \in \mathcal{W}$
$f \in \text{FConst}$	$ = \{\text{cons}, \text{car}, \text{cdr}, \text{makeref}, \text{deref}, \text{setref!}, \text{callcc}, \text{sync}, \text{touch}\}$	(Func. Cstnt)	Environment Operations:
$f_n \in \text{NStP}$	$ = \{\text{cons}, \text{makeref}, \text{callcc}\}$	(Non Strict Primitives)	$\rho(x) = V$ if $(x V) \in \rho$
$x \in \text{Vars}$	$ = \{x, y, z, \dots\}$	(User Variable)	$\rho[x \leftarrow V] = (\rho \setminus \{(x V')\}) \cup \{(x V)\}$
$\kappa \in \text{CCode}$	$::= (\text{init}) \mid (\kappa \text{ fun } V)$	(Continuation code)	if $(x V') \in \rho$
	$ \mid (\kappa \text{ arg } M \rho) \mid (\kappa \text{ cond } (M, M, \rho))$		$\rho[x \leftarrow V] = \rho \cup \{(x V)\}$
	$ \mid (\kappa \text{ det } (\text{ph}, \ell)) \mid (\kappa \text{ leg } (\ell, \ell))$		if $x \notin \text{DOM}(\rho)$
$R \in \text{Req}$	$::= \text{Req}(s, u, rc)$	(Request)	
$rc \in \text{RC}$	$::= \text{rtouch}(\alpha) \mid \text{rdet}(\alpha, V, \alpha, \ell)$	(Request Contents)	
	$ \mid \text{deref}(\alpha) \mid \text{rset}(\alpha, V) \mid \text{rleg}(\ell, \ell)$		
$A \in \text{Ans}$	$::= \text{Ans}(u, ac)$	(Answer)	
$ac \in \text{AC}$	$::= \text{rtouch}(V) \mid \text{rdet}(X)$	(Answer Contents)	
	$ \mid \text{rset}(X) \mid \text{deref}(V) \mid \text{rleg}$		
	$\text{Itouch}(W, \theta, s) = W$		$\ell \rightsquigarrow_\theta^s \ell$
	$\text{Itouch}(\langle \text{ph } \alpha, s \rangle, \theta, s) = \text{Itouch}(\theta(\alpha), \theta, s)$		if $\theta(\alpha) \neq \perp$
	$\text{Itouch}(\langle \text{ph } \alpha, s \rangle, \theta, s) = \langle \text{ph } \alpha, s \rangle$		if $\theta(\alpha) = \perp$
	$\text{Itouch}(\langle \text{ph } \alpha, s_1 \rangle, \theta, s) = \langle \text{ph } \alpha, s_1 \rangle$		if $s_1 \neq s$
			$\ell \rightsquigarrow_{\mathcal{W}} \ell$
			$\langle \text{leg } \alpha, s_1 \rangle \rightsquigarrow_{\mathcal{W}} \ell$ if $\theta_{s_1}(\alpha) \neq \perp$
			and $\mathcal{W}[(\alpha, s_1)] \rightsquigarrow_{\mathcal{W}} \ell$

Fig. 2. State Space of the CEKDS-machine

$\text{Ev}\langle(M\ N), \rho, \kappa\rangle \rightarrow_{cek} \text{Ev}\langle M, \rho, (\kappa\ \mathbf{arg}\ N, \rho)\rangle$	(operator)
$\text{Ev}\langle\lambda x.M, \rho, \kappa\rangle \rightarrow_{cek} \text{Ret}\langle\langle\mathbf{cl}\ \lambda x.M, \rho\rangle, \kappa\rangle$	(lambda)
$\text{Ev}\langle c, \rho, \kappa\rangle \rightarrow_{cek} \text{Ret}\langle c, \kappa\rangle$	(constant)
$\text{Ev}\langle x, \rho, \kappa\rangle \rightarrow_{cek} \text{Ret}\langle\rho(x), \kappa\rangle$	(variable)
$\text{Ret}\langle V, (\kappa\ \mathbf{arg}\ N, \rho)\rangle \rightarrow_{cek} \text{Ev}\langle N, \rho, (\kappa\ \mathbf{fun}\ V)\rangle$	(operand)
$\text{Ret}\langle V, (\kappa\ \mathbf{fun}\ \langle\mathbf{cl}\ \lambda x.M, \rho\rangle)\rangle \rightarrow_{cek} \text{Ev}\langle M, \rho[x \leftarrow V], \kappa\rangle$	(apply)
$\text{Ev}\langle\langle\mathbf{if}\ M\ M_1\ M_2\rangle, \rho, \kappa\rangle \rightarrow_{cek} \text{Ev}\langle M, \rho, (\kappa\ \mathbf{cond}\ (M_1, M_2, \rho))\rangle$	(predicate)
$\text{Ret}\langle V, (\kappa\ \mathbf{cond}\ (M_1, M_2, \rho))\rangle \rightarrow_{cek} \text{Ev}\langle M_2, \rho, \kappa\rangle$ if $V = \mathbf{false}$	(if else)
$\rightarrow_{cek} \text{Ev}\langle M_1, \rho, \kappa\rangle$ if $V \neq \mathbf{false}$	(if then)
$\text{Ret}\langle V, (\kappa\ \mathbf{fun}\ \mathbf{callcc})\rangle \rightarrow_{cek} \text{Ret}\langle\langle\mathbf{co}\ \kappa\rangle, (\kappa\ \mathbf{fun}\ V)\rangle$	(capture)
$\text{Ret}\langle V, (\kappa'\ \mathbf{fun}\ \langle\mathbf{co}\ \kappa\rangle)\rangle \rightarrow_{cek} \text{Ret}\langle V, \kappa\rangle$	(invoke)
$\text{Ret}\langle V, (\kappa\ \mathbf{fun}\ V_1)\rangle \rightarrow_{cek} \text{Ret}\langle(V_1\ V), \kappa\rangle$ if $V_1 \in PApp$	(partial apply)
$\text{Ret}\langle V, (\kappa\ \mathbf{fun}\ (\mathbf{cons}\ V_1))\rangle \rightarrow_{cek} \text{Ret}\langle(\mathbf{cons}\ V_1\ V), \kappa\rangle$	(cons)
$\text{Ret}\langle(\mathbf{cons}\ V_1\ V_2), (\kappa\ \mathbf{fun}\ \mathbf{car})\rangle \rightarrow_{cek} \text{Ret}\langle V_1, \kappa\rangle$	(car)
$\text{Ret}\langle(\mathbf{cons}\ V_1\ V_2), (\kappa\ \mathbf{fun}\ \mathbf{cdr})\rangle \rightarrow_{cek} \text{Ret}\langle V_2, \kappa\rangle$	(cdr)
$\text{Ret}\langle V, (\kappa\ \mathbf{fun}\ f)\rangle \rightarrow_{cek} \text{Ret}\langle\delta(f, V), \kappa\rangle$	(δ)
$\text{Ret}\langle V, (\kappa\ \mathbf{fun}\ V_1)\rangle \rightarrow_{cek} \text{Ret}\langle\mathbf{error}, (\mathbf{init})\rangle$ if $V_1 \notin AValue$	(apply error)

Fig. 3. Transitions between computational states

continuation of **future** with the placeholder *ph*. After transition, the initial task τ evaluates the argument of **future**. We shall explain later the purpose of the legitimacies ℓ and ℓ_1 .

A new box can be created by applying the functional constant **makeref** on a value. As a result, a new location is allocated in the local store, and a new box object, which refers to the new location and the site name, i.e. the qualified location, is returned.

In order to illustrate the behaviour of the machine, we consider the operation of reading a box, which will lead us to explain some rules of Figures 4, 5, 6, and 7; similar comments apply to box modification. A task considers that a box is local if the site name held in the box is the name of the current site. According to (*deref local*) in Figure 4, if the box is local, the value contained in the local store at the given location can be returned. Otherwise, rule (*deref remote*) adds to the output spool a *request* addressed to the site that allocated the box; in addition, the task is suspended, which is modelled by its transfer from the set T of runnable tasks to the set S of suspended tasks. We represent requests as triples composed of the destination site, the qualified name of the requesting task, and the message itself describing the type of the request. In the present case, the message **deref**(α) means that the distant site is asked to supply the content of location α .

Sites communicate according to the rules of Figure 5. In rule (*migrate request*), two sites exchange requests by moving them from the output spool of the source site to the input spool of the destination site. Figure 6 shows how a site handles incoming requests. In the case of rule (*request deref*), the content of the location is packaged up into an *answer* which must return back to the task that initiated the request; we again can see the interest of the qualified task name which indicates the name of the site that emitted the request. The answer is entered in the output spool and is migrated, like a request, by rule (*migrate answer*). Figure 7 shows the rules that handle incoming answers. The arrival of the answer **deref**(V) awakens the task waiting for this answer by transferring it back to the set of runnable tasks, with the value V as the content of the box.

Rule (*fork*) allows us to create new tasks on the current site. In Figure 5, rule (*migrate task*) shows that a task may be migrated from a site 1 running more than one active task to a site 2 without any active task. We see that a task is migrated by transferring its computational state, i.e. among others its continuation, and its legitimacy.

$$\begin{aligned}
& \langle \{ \langle C, \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \\
& \rightarrow_s \langle \{ \langle C_1, \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \quad \text{if } C \rightarrow_{cek} C_1 \quad (\text{sequential}) \\
& \langle \{ \langle \text{Ev}(\langle \text{future } M \rangle, \rho, \kappa), \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \\
& \rightarrow_s \langle \{ \langle \text{Ev}(M, \rho, (\kappa \text{ det } ph, \ell_1)), \ell, \tau \rangle, \langle \text{Ret}(ph, \kappa), \ell_1, \tau_1 \rangle \} \cup T, \theta_1, s, S, I, O \rangle \quad (\text{fork}) \\
& \quad \text{with } ph = \langle \text{ph } \alpha, s \rangle, \ell_1 = \langle \text{leg } \alpha_1, s \rangle, \theta_1 = \theta \uplus \{ (\alpha \perp) (\alpha_1 \perp) \}, \tau_1 \notin FN(T \cup S) \cup \{ \tau \} \\
& \langle \{ \langle \text{Ret}(V, (\kappa \text{ fun makeref})), \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \\
& \rightarrow_s \langle \{ \langle \text{Ret}(b, \kappa), \ell, \tau \rangle \} \cup T, \theta_1, s, S, I, O \rangle \text{ with } b = \langle \text{bx } \alpha, s, \ell \rangle, \theta_1 = \theta \uplus \{ (\alpha V) \} \quad (\text{makeref}) \\
& \langle \{ \langle \text{Ret}(\langle \text{bx } \alpha, s, \ell \rangle, (\kappa \text{ fun deref})), \ell_1, \tau \rangle \} \cup T, \theta, s_1, S, I, O \rangle \\
& \rightarrow_s \langle \{ \langle \text{Ret}(\theta(\alpha), \kappa), \ell_1, \tau \rangle \} \cup T, \theta, s_1, S, I, O \rangle \quad \text{if } s_1 = s \quad (\text{deref local}) \\
& \rightarrow_s \langle T, \theta, s_1, \{ \langle \text{Ret}(\langle \text{bx } \alpha, s, \ell \rangle, (\kappa \text{ fun deref})), \ell_1, \tau \rangle \} \cup S, I, O_1 \rangle \quad \text{if } s_1 \neq s \quad (\text{deref remote}) \\
& \quad \text{with } O_1 = \{ \text{Req}(s, \langle \tau, s_1 \rangle, \text{deref}(\alpha)) \} \cup O \\
& \langle \{ \langle \text{Ret}(V, (\kappa \text{ fun setref } \langle \text{bx } \alpha, s, \ell \rangle)), \ell_1, \tau \rangle \} \cup T, \theta, s_1, S, I, O \rangle \\
& \rightarrow_s \langle \{ \langle \text{Ret}(\text{void}, \kappa), \ell_1, \tau \rangle \} \cup T, \theta[\alpha := V], s_1, S, I, O \rangle \quad \text{if } s_1 = s \quad (\text{setref local}) \\
& \rightarrow_s \langle T, \theta, s_1, \{ \langle \text{Ret}(V, (\kappa \text{ fun setref } \langle \text{bx } \alpha, s, \ell \rangle)), \ell_1, \tau \rangle \} \cup S, I, O_1 \rangle \quad (\text{setref remote}) \\
& \quad \text{if } s_1 \neq s, \text{ with } O_1 = \{ \text{Req}(s, \langle \tau, s_1 \rangle, \text{rset}(\alpha, V)) \} \cup O \\
& \langle \{ \langle \text{Ret}(V, (\kappa \text{ det } \langle \text{ph } \alpha, s \rangle, \langle \text{leg } \alpha_1, s \rangle)), \ell_2, \tau \rangle \} \cup T, \theta, s_2, S, I, O \rangle \\
& \rightarrow_s \langle T_1, \theta_1, s_2, S_1, I_1, O \rangle \quad \text{if } s_2 = s, \theta(\alpha) = \perp \quad (\text{determine local}) \\
& \quad \text{with } I_1 = (I \cup I_2), T_1 = (T \cup T_2), S_1 = (S \setminus (I_2 \cup T_2)), \theta_1 = \theta[\alpha_1 := \ell_2][\alpha := V] \\
& \quad \text{with } I_2 = \{ \text{Req}(s_3, \langle \tau_1, s_4 \rangle, \text{rtouch}(\alpha)), \text{Req}(s_3, \langle \tau_1, s_4 \rangle, \text{rleg}(\langle \text{leg } \alpha_1, s \rangle, \ell_4)) \in S \} \\
& \quad \text{with } T_2 = \{ \langle \text{Ret}(\langle \text{ph } \alpha, s \rangle, (\kappa' \text{ fun touch})), \ell_3, \tau_1 \rangle \in S \} \cup \\
& \quad \quad \{ \langle \text{Ret}(V_1, (\kappa' \text{ leg } \langle \langle \text{leg } \alpha_1, s \rangle, \ell_4) \rangle)), \ell_5, \tau_1 \rangle \in S \} \\
& \rightarrow_s \langle \{ \langle \text{Ret}(V, \kappa), \ell_2, \tau \rangle \} \cup T, \theta, s_2, S, I, O \rangle \quad \text{if } s_2 = s, \theta(\alpha) \neq \perp \quad (\text{determine localn}) \\
& \rightarrow_s \langle T, \theta, s_2, \{ \langle \text{Ret}(V, (\kappa \text{ det } \langle \text{ph } \alpha, s \rangle, \langle \text{leg } \alpha_1, s \rangle)), \ell_2, \tau \rangle \} \cup S, I, O_1 \rangle \quad (\text{determine remote}) \\
& \quad \text{if } s_2 \neq s \text{ with } O_1 = \{ \text{Req}(s, \langle \tau, s_2 \rangle, \text{rdet}(\alpha, V, \alpha_1, \ell_2)) \} \cup O \\
& \langle \{ \langle \text{Ret}(V, (\kappa \text{ fun touch})), \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \\
& \rightarrow_s \langle \{ \langle \text{Ret}(\text{ltouch}(V, \theta, s), \kappa), \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \quad (\text{touch local}) \\
& \quad \text{if } \text{ltouch}(V, \theta, s) \in P\text{Value} \\
& \rightarrow_s \langle T, \theta, s, \{ \langle \text{Ret}(\langle \text{ph } \alpha, s_1 \rangle, (\kappa \text{ fun touch})), \ell, \tau \rangle \} \cup S, I, O_1 \rangle \quad (\text{touch remote}) \\
& \quad \text{if } \text{ltouch}(V, \theta, s) = \langle \text{ph } \alpha, s_1 \rangle, s \neq s_1, \text{ with } O_1 = \{ \text{Req}(s_1, \langle \tau, s \rangle, \text{rtouch}(\alpha)) \} \cup O \\
& \rightarrow_s \langle T, \theta, s, \{ \langle \text{Ret}(\langle \text{ph } \alpha, s_1 \rangle, (\kappa \text{ fun touch})), \ell, \tau \rangle \} \cup S, I, O \rangle \quad (\text{touch suspend}) \\
& \quad \text{if } \text{ltouch}(V, \theta, s) = \langle \text{ph } \alpha, s_1 \rangle, s = s_1, \theta(\alpha) = \perp \\
& \langle \{ \langle \text{Ret}(V, (\kappa \text{ fun sync})), \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \\
& \rightarrow_s \langle \{ \langle \text{Ret}(V, (\kappa \text{ leg } \langle \ell, \ell_1 \rangle)), \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \quad \text{if } V = \langle \text{bx } \alpha, s_1, \ell_1 \rangle \quad (\text{synchronise}) \\
& \rightarrow_s \langle \{ \langle \text{Ret}(\text{error}, \langle \text{init} \rangle), \ell, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \quad \text{if } V \notin \text{Box} \quad (\text{synchronise error}) \\
& \langle \{ \langle \text{Ret}(V, (\kappa \text{ leg } \langle \ell, \ell_1 \rangle)), \ell_2, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \\
& \rightarrow_s \langle \{ \langle \text{Ret}(V, \kappa), \ell_2, \tau \rangle \} \cup T, \theta, s, S, I, O \rangle \quad \text{if } \ell \rightsquigarrow_\theta^s \ell_1 \quad (\text{leg local}) \\
& \rightarrow_s \langle T, \theta, s, \{ \langle \text{Ret}(V, (\kappa \text{ leg } \langle \langle \text{leg } \alpha, s_1 \rangle, \ell_1 \rangle)), \ell_2, \tau \rangle \} \cup S, I, O_1 \rangle \quad (\text{leg remote}) \\
& \quad \text{if } \ell \rightsquigarrow_\theta^s \langle \text{leg } \alpha, s_1 \rangle, s_1 \neq s, \text{ with } O_1 = \{ \text{Req}(s_1, \langle \tau, s \rangle, \text{rleg}(\langle \text{leg } \alpha, s_1 \rangle, \ell_1)) \} \cup O \\
& \rightarrow_s \langle T, \theta, s, \{ \langle \text{Ret}(V, (\kappa \text{ leg } \langle \langle \text{leg } \alpha, s \rangle, \ell_1 \rangle)), \ell_2, \tau \rangle \} \cup S, I, O \rangle \quad (\text{leg suspend}) \\
& \quad \text{if } \ell \rightsquigarrow_\theta^s \langle \text{leg } \alpha, s \rangle, \ell \not\rightsquigarrow_\theta^s \ell_1, \theta(\alpha) = \perp
\end{aligned}$$

Fig. 4. Site Transitions

$$\begin{aligned}
& \{ \langle \langle C_1, \ell_1, \tau_1 \rangle \rangle \cup T_1, \theta, s_1, S_1, I_1, O_1 \rangle, \langle \emptyset, \theta_2, s_2, S_2, I_2, O_2 \rangle \} \cup \mathcal{W} & (\text{migrate task}) \\
& \rightarrow_c \{ \langle T_1, \theta_1, s_1, S_1, I_1, O_1 \rangle, \langle \langle C_1, \ell_1, \tau_2 \rangle \rangle, \theta_2, s_2, S_2, I_2, O_2 \rangle \} \cup \mathcal{W} \\
& \quad \text{if } T_1 \neq \emptyset, \text{ with } \tau_2 \notin FN(S_2) \\
& \{ \langle T_1, \theta_1, s_1, S_1, I_1, \{ \text{Req}(s_2, \langle \tau_3, s_3 \rangle, rc) \} \cup O_1 \rangle, \langle T_2, \theta_2, s_2, S_2, I_2, O_2 \rangle \} \cup \mathcal{W} & (\text{migrate request}) \\
& \rightarrow_c \{ \langle T_1, \theta_1, s_1, S_1, I_1, O_1 \rangle, \langle T_2, \theta_2, s_2, S_2, \{ \text{Req}(s_2, \langle \tau_3, s_3 \rangle, rc) \} \cup I_2, O_2 \rangle \} \cup \mathcal{W} \\
& \{ \langle T_1, \theta_1, s_1, S_1, I_1, \{ \text{Ans}(\tau_2, s_2, ac) \} \cup O_1 \rangle, \langle T_2, \theta_2, s_2, S_2, I_2, O_2 \rangle \} \cup \mathcal{W} & (\text{migrate answer}) \\
& \rightarrow_c \{ \langle T_1, \theta_1, s_1, S_1, I_1, O_1 \rangle, \langle T_2, \theta_2, s_2, S_2, \{ \text{Ans}(\tau_2, s_2, ac) \} \cup I_2, O_2 \rangle \} \cup \mathcal{W}
\end{aligned}$$

Fig. 5. Communications between sites

$$\begin{aligned}
& \langle T, \theta, s, S, \{ \text{Req}(s, \langle \tau, s_2 \rangle, \text{deref}(\alpha)) \} \cup I, O \rangle \\
& \rightarrow_s \langle T, \theta, s, S, I, \{ \text{Ans}(\tau, s_2, \text{deref}(\theta(\alpha))) \} \cup O \rangle & (\text{request deref}) \\
& \langle T, \theta, s, S, \{ \text{Req}(s, \langle \tau, s_2 \rangle, \text{rset}(\alpha, V)) \} \cup I, O \rangle \\
& \rightarrow_s \langle T, \theta_1, s, S, I, \{ \text{Ans}(\tau, s_2, \text{rset}) \} \cup O \rangle \text{ with } \theta_1 = \theta[\alpha := V] & (\text{request set}) \\
& \langle T, \theta, s, S, \{ \text{Req}(s, \langle \tau, s_2 \rangle, \text{rdet}(\alpha, V, \alpha_1, \ell)) \} \cup I, O \rangle \\
& \rightarrow_s \langle T_1, \theta_1, s, S_1, I_1, \{ \text{Ans}(\tau, s_2, \text{rdet}(\theta(\alpha))) \} \cup O \rangle \text{ if } \theta(\alpha) = \perp & (\text{request det first}) \\
& \quad \text{with } I_1 = (I \cup I_2), T_1 = (T \cup T_2), S_1 = (S \setminus (I_2 \cup T_2)), \theta_1 = \theta[\alpha := V][\alpha_1 := \ell] \\
& \quad \text{with } I_2 = \{ \text{Req}(s_3, \langle \tau_1, s_4 \rangle, \text{rtouch}(\alpha)), \text{Req}(s_3, \langle \tau_1, s_4 \rangle, \text{rleg}(\langle \text{leg } \alpha_1, s \rangle, \ell_1)) \in S \} \\
& \quad \text{with } T_2 = \{ \langle \text{Ret}(\langle \text{ph } \alpha, s \rangle, (\kappa \text{ fun touch})), \ell_1, \tau_1 \rangle \in S \} \cup \\
& \quad \quad \{ \langle \text{Ret}(V_2, (\kappa \text{ leg } (\langle \text{leg } \alpha_1, s \rangle, \ell_1))), \ell_2, \tau_1 \rangle \in S \} \\
& \rightarrow_s \langle T, \theta, s, S, I, \{ \text{Ans}(\tau, s_2, \text{rdet}(\theta(\alpha))) \} \cup O \rangle \text{ if } \theta(\alpha) \neq \perp & (\text{request det mult}) \\
& \langle T, \theta, s, S, \{ \text{Req}(s, \langle \tau, s_2 \rangle, \text{rtouch}(\alpha)) \} \cup I, O \rangle \\
& \rightarrow_s \langle T, \theta, s, S, I, \{ \text{Ans}(\tau, s_2, \text{rtouch}(\text{ltouch}(\langle \text{ph } \alpha, s \rangle, \theta, s))) \} \cup O \rangle & (\text{request touch local}) \\
& \quad \text{if } \text{ltouch}(\langle \text{ph } \alpha, s \rangle, \theta, s) \in P\text{Value}, s_2 \neq s \\
& \rightarrow_s \langle T, \theta, s, S, \{ \text{Ans}(\tau, s_2, \text{rtouch}(\text{ltouch}(\langle \text{ph } \alpha, s \rangle, \theta, s))) \} \cup I, O \rangle & (\text{request touch local'}) \\
& \quad \text{if } \text{ltouch}(\langle \text{ph } \alpha, s \rangle, \theta, s) \in P\text{Value}, s_2 = s \\
& \rightarrow_s \langle T, \theta, s, S, I, \{ \text{Req}(s_3, \langle \tau, s_2 \rangle, \text{rtouch}(\alpha_1)) \} \cup O \rangle & (\text{request touch remote}) \\
& \quad \text{if } \text{ltouch}(\langle \text{ph } \alpha, s \rangle, \theta, s) = \langle \text{ph } \alpha_1, s_3 \rangle, s \neq s_3 \\
& \rightarrow_s \langle T, \theta, s, \{ \text{Req}(s, \langle \tau, s_2 \rangle, \text{rtouch}(\alpha_1)) \} \cup S, I, O \rangle & (\text{request touch suspend}) \\
& \quad \text{if } \text{ltouch}(\langle \text{ph } \alpha, s \rangle, \theta, s) = \langle \text{ph } \alpha_1, s \rangle, \theta(\alpha_1) = \perp \\
& \langle T, \theta, s, S, \{ \text{Req}(s, \langle \tau, s_2 \rangle, \text{rleg}(\ell, \ell_1)) \} \cup I, O \rangle \\
& \rightarrow_s \langle T, \theta, s, S, I, \{ \text{Ans}(\tau, s_2, \text{rleg}) \} \cup O \rangle \text{ if } \ell \rightsquigarrow_{\theta}^s \ell_1, s_2 \neq s & (\text{request leg local}) \\
& \rightarrow_s \langle T, \theta, s, S, \{ \text{Ans}(\tau, s_2, \text{rleg}) \} \cup I, O \rangle \text{ if } \ell \rightsquigarrow_{\theta}^s \ell_1, s_2 = s & (\text{request leg local'}) \\
& \rightarrow_s \langle T, \theta, s, S, I, \{ \text{Req}(s_3, \langle \tau, s_2 \rangle, \text{rleg}(\langle \text{leg } \alpha, s_3 \rangle, \ell_1)) \} \cup O \rangle & (\text{request leg remote}) \\
& \quad \text{if } \ell \rightsquigarrow_{\theta}^s \langle \text{leg } \alpha, s_3 \rangle, s \neq s_3 \\
& \rightarrow_s \langle T, \theta, s, \{ \text{Req}(s, \langle \tau, s_2 \rangle, \text{rleg}(\langle \text{leg } \alpha, s_3 \rangle, \ell_1)) \} \cup S, I, O \rangle & (\text{request leg suspend}) \\
& \quad \text{if } \ell \rightsquigarrow_{\theta}^s \langle \text{leg } \alpha, s \rangle, \theta(\alpha) = \perp, \ell \not\rightsquigarrow_{\theta}^s \ell_1
\end{aligned}$$

Fig. 6. Handling of Requests

According to (*fork*), the effect of evaluating (*future* M) is to allocate a placeholder that a new task speculatively passes to its continuation. The original task, which evaluates M , acts as a producer for the placeholder value, while the new task acts as a consumer.

A producer task has obtained the value V of a *future* argument, when V is returned to a continuation code of the form $(\kappa \text{ det } \text{ph}, \ell)$; the producer task is then expected to store V

$\langle T, \theta, s, \{ \text{Ret}\langle V, (\kappa \text{ fun deref}) \rangle, \ell, \tau \} \cup S, \{ \text{Ans}\langle \tau, s, \text{deref}(V_1) \rangle \} \cup I, O \rangle$	
$\rightarrow_s \langle \{ \text{Ret}\langle V_1, \kappa \rangle, \ell, \tau \} \cup T, \theta, s, S, I, O \rangle$	(answer deref)
$\langle T, \theta, s, \{ \text{Ret}\langle V, (\kappa \text{ fun (setref } b)) \rangle, \ell, \tau \} \cup S, \{ \text{Ans}\langle \tau, s, \text{rset} \rangle \} \cup I, O \rangle$	
$\rightarrow_s \langle \{ \text{Ret}\langle \text{void}, \kappa \rangle, \ell, \tau \} \cup T, \theta, s, S, I, O \rangle$	(answer rset)
$\langle T, \theta, s, \{ \text{Ret}\langle V, (\kappa \text{ det } ph, \ell) \rangle, \ell_1, \tau \} \cup S, \{ \text{Ans}\langle \tau, s, \text{rdet}(X) \rangle \} \cup I, O \rangle$	
$\rightarrow_s \langle T, \theta, s, S, I, O \rangle$ if $X = \perp$	(answer det first)
$\rightarrow_s \langle \{ \text{Ret}\langle V, \kappa \rangle, \ell_1, \tau \} \cup T, \theta, s, S, I, O \rangle$ if $X \neq \perp$	(answer det mult)
$\langle T, \theta, s, \{ \text{Ret}\langle ph, (\kappa \text{ fun touch}) \rangle, \ell, \tau \} \cup S, \{ \text{Ans}\langle \tau, s, \text{rtouch}(V) \rangle \} \cup I, O \rangle$	
$\rightarrow_s \langle \{ \text{Ret}\langle V, \kappa \rangle, \ell, \tau \} \cup T, \theta, s, S, I, O \rangle$	(answer touch)
$\langle T, \theta, s, \{ \text{Ret}\langle V, (\kappa \text{ leg } (\ell, \ell_1)) \rangle, \ell_2, \tau \} \cup S, \{ \text{Ans}\langle \tau, s, \text{rleg} \rangle \} \cup I, O \rangle$	
$\rightarrow_s \langle \{ \text{Ret}\langle V, \kappa \rangle, \ell_2, \tau \} \cup T, \theta, s, S, I, O \rangle$	(answer leg)

Fig. 7. Handling of Answers

in the placeholder ph ; this operation is called *determining* the placeholder. Depending on whether the current task is running on the site where the placeholder was allocated, rules (*determine local*) or (*determine remote*) take care of assigning the value V to the placeholder ph . However, placeholders are not boxes because they are defined as datastructures that can receive one value *at most* [11]; placeholders are like single-assignment variables in CC+ [2] and PCN [7]. As opposed to conventional languages, the language Λ_f has first-class continuations which allow the programmer to write expressions that “return” multiple values; in other words, in Λ_f , different values can be passed to the same continuation. As a result, we distinguish the case where a placeholder is not assigned, in rules (*determine local*) and (*request det first*), from the case where it is already assigned, rules (*determine localn*) and (*request det mult*). In the former case, the placeholder is updated and the producer task ends its evaluation. In the latter case, the value is returned to the continuation of *future*, as if no *future* had existed, following Katz and Weise’s implementation [15, 3, 20]. Let us observe that transitions are atomically executed in order to ensure a sound behaviour of (*determine local*) and (*determine localn*).

Strict primitives introduce synchronisations between the consumer task of a placeholder and its producer task: they require their arguments to be *proper values*, i.e. values different from placeholders; strict primitives are said to *touch* their argument. The translation \mathcal{X} makes the touch action explicit by the call to *touch*, whose purpose is to return a *proper value*. We use an auxiliary function *ltouch*, displayed in Figure 2, which touches a value with respect to a local store θ of a site s . The function *ltouch* can return three results: a proper value, an undetermined placeholder that was allocated on site s , or a placeholder that was allocated on a different site. In the first case, the *touch* operation succeeds (*touch local*); in the second case, the task is suspended as long as the placeholder remains undetermined (*touch suspend*); in the third case, a request *rtouch*(α) is sent to the remote site (*touch remote*). The remote site behaves similarly: it can return a proper value, suspend the request, or pass it to another site. Tasks or requests that are suspended when touching a placeholder are reactivated when this placeholder gets determined, cfr. (*determine local*) or (*request det first*). Let us observe that the *touch* operation can initiate exchanges of messages between sites; as soon as a proper value is found, it is directly returned back to the site that started the operation, thanks to the qualified task name.

So far, our explanations have ignored legitimacies. Following Katz and Weise [15], we use a notion of legitimacy to keep track of the control flow that would exist if evaluation was sequential. An initial legitimacy is allocated when we start to evaluate a program, and each new task is given a new legitimacy. Legitimacies, like placeholders, are datastructures whose only slot can receive one value at most; unlike placeholders, legitimacies are not first-class values. When a placeholder gets determined, the consumer task becomes dependent on the value of the placeholder; hence, the legitimacy of the producer task, recorded in the

continuation ($\kappa \text{ det } ph \ell$), is stored into the legitimacy of the consumer task. As evaluation proceeds, chains of legitimacies get formed into memory. The relation $\ell_1 \sim_\theta^s \ell_2$ states that there is a path from legitimacy ℓ_1 to legitimacy ℓ_2 in the local store θ of site s , which means that control has flowed from a task with legitimacy ℓ_2 to a task with legitimacy ℓ_1 .

As we want **future** to be an annotation, every program should return the result that it would produce when evaluated sequentially in the absence of **future**. The solution adopted in our semantics is to perform causally-dependent [24] box accesses in the same order as in a sequential implementation; the solution relies on legitimacies. The translation of the primitive **deref**, $\lambda x.(\text{deref } (\text{sync } (\text{touch } x)))$, touches and then applies **sync** on its argument. The primitive **sync** behaves as the identity function if the legitimacy of the current task leads to the legitimacy associated with the box. In other words, **sync** acts as a synchronisation barrier by ensuring that all accesses to the box (read or write) that a sequential implementation would have performed before the current access are actually done in the parallel machine, and all accesses that a sequential implementation would perform after the current one remain to be performed by the parallel machine. The primitive **sync** suspends a task that illegitimately tries to access a box; it will be reactivated by (*determine local*) or (*request det first*).

In order to determine when a computation ends, the initial configuration contains a box aimed at receiving the final value. Consistent box accesses guarantee that the box will receive the legitimate final value (if there exists one).

It should be observed that using legitimacies to synchronise box accesses does not impose a *total* order on those operations, but a *partial* order. This property ensures that parallelism can exist for programs written in a mostly-functional style, where one generally considers that side-effects are performed locally in different modules or functions.

$$\begin{aligned}
& \mathcal{W}_1 \rightarrow_{ds}^{1,m} \mathcal{W}_2 \text{ if } \mathcal{W}_1 \rightarrow_c \mathcal{W}_2, \text{ or } \mathcal{W}_1 = \{m_1\} \cup \mathcal{W}, \mathcal{W}_2 = \{m_2\} \cup \mathcal{W}, m_1 \rightarrow_s m_2, \\
& \quad \text{with } m = 1 \text{ if } \ell \sim_{\mathcal{W}_1} \ell_0, \text{ with } \ell \text{ the legitimacy of the task related to the transition} \\
& \quad m = 0 \text{ otherwise} \\
& \mathcal{W}_1 \rightarrow_{ds}^{n+n',m+m'} \mathcal{W}_2 \text{ if } \mathcal{W}_1 \rightarrow_{ds}^{n,m} \mathcal{W}_3 \text{ and } \mathcal{W}_3 \rightarrow_{ds}^{n',m'} \mathcal{W}_2 \quad (\text{transitive}) \\
& \text{Conventions: } \mathcal{W}_1 \rightarrow_{ds} \mathcal{W}_2 \text{ if } \mathcal{W}_1 \rightarrow_{ds}^{1,m} \mathcal{W}_2; \mathcal{W}_1 \rightarrow_{ds}^* \mathcal{W}_2 \text{ if } \mathcal{W}_1 \rightarrow_{ds}^{n,m} \mathcal{W}_2, n \geq 0; \\
& \quad \mathcal{W}_1 \rightarrow_{ds}^+ \mathcal{W}_2 \text{ if } \mathcal{W}_1 \rightarrow_{ds}^{n,m} \mathcal{W}_2, n > 0. \\
& \text{Initial world for } P, \text{InitWorld}[P] = \{m_0, m_1, \dots, m_n\}, \text{ with} \\
& \quad \text{Initial Store: } \theta_0 = \{(\alpha_0 \perp) (\alpha_1 \perp)\} \quad \text{Initial Legitimacy: } \ell_0 = \langle \text{leg } \alpha_1, s_0 \rangle \\
& \quad \text{Initial Box: } b_0 = \langle \text{bx } \alpha_0, s_0, \ell_0 \rangle \quad \text{Initial Environment: } \rho_0 = \{(x \ b_0)\} \\
& \quad m_0 = \langle \{ \langle \text{Ev}((\lambda v.(\text{setref! } (\text{sync } x) v)) \ \mathcal{A}[P]) \rangle, \rho_0, (\text{init}), \ell_0, \tau_0 \} \rangle, s_0, \theta_0, \emptyset, \emptyset, \emptyset \rangle \\
& \quad \text{Empty Sites: } m_i = \langle \emptyset, s_i, \emptyset, \emptyset, \emptyset, \emptyset \rangle \ i = 1, \dots, n \\
& \text{Final World: } \text{Final}[\mathcal{W}_f] \text{ if } \mathcal{W}_f[\langle \alpha_0, s_0 \rangle] \neq \perp.
\end{aligned}$$

Fig. 8. Evaluation Relation

We now have all the components to define an evaluation relation that associates programs with their observable behaviour. The function *Unload* replaces each function, box, or continuation by a tag; in addition, *Unload* touches every placeholder appearing in the result. As values can be spread over different sites, *Unload* takes the world in argument.

Divergence should be defined with the greatest care, because **future** has the ability to create new tasks, but the scheduler may elect to evaluate any of them. One task only is *mandatory*; all the others are *speculative*. A task with legitimacy ℓ is mandatory if ℓ leads

to the initial legitimacy ℓ_0 in the current world \mathcal{W} , which is written $\ell \rightsquigarrow_{\mathcal{W}} \ell_0$. Figure 8 defines the relation $\rightarrow_{ds}^{n,m}$ [6] to denote reductions that involve n steps among which m are mandatory. According to the evaluation relation eval_{ds} , a program is said to be divergent, i.e. its value is \perp , if it leads to an infinite transition sequence that regularly often contains mandatory transitions.

The soundness of the CEKDS-machine is established by proving that its evaluation relation eval_{ds} is equal to the sequential evaluation function of the CEK-Machine.

Theorem 1 (Soundness) $\text{eval}_{ds} = \text{eval}_{ss}$ \square

3 Discussion and Related Work

This paper builds upon previous work about annotations for parallelism in functional languages. For a long time, research has focused on implementation issues and efficient designs [15, 3, 28, 14, 13, 11, 18]. Parallelism by annotations has been formalised recently only. Flanagan and Felleisen [6] have defined the semantics of **future** in a purely functional language. The author [19, 20] has proposed a semantic framework for continuations and side-effects in a language with the annotations **pcall**, **fork**, and **future**. This paper is the first to present a formal semantics for **futures**, side-effect, first-class continuations, and distribution. Our research is part of a project which aims at building a Virtual Multicomputer [23], which provides a soft architecture to support distributed applications, transcending the details of hardware architecture. The “distribution by annotation” paradigm is our contribution to this virtual multiprocessor, which provides the user with the view that a distributed network of computers is programmable as a sequential processor.

Both compile-time and runtime improvements could boost the performance of our architecture. Our semantics is a dynamic semantics, and there are opportunities to improve it using static analysis. Flanagan and Felleisen’s [6] *touch analysis* remove provably redundant **touch** operations in purely functional **future**-based programs, using a set-based analysis [12]; extending their analysis to side-effects and first-class continuations would be desirable for the CEKDS-machine. Similarly, an analysis removing unnecessary **sync** operations would greatly reduce the cost of synchronisations associated with side-effects.

As far as the runtime system is concerned, a realistic implementation needs a distributed garbage collection; the approach “garbage collecting the world” [16] appears to be a suitable candidate. Similarly, we have to address the issue of process collection. Miller’s MultiScheme [17] task collection is done during garbage collection: a task can be reclaimed if the placeholder that it determines is accessible from the gc roots.

According to Figure 4, every **future** creates a new task on the current site. This process creation strategy is referred to as *eager task creation* [11, 18, 3]. However, **future**-based programs can generate far more tasks than the number of sites in a CEKDS-world. In order to avoid the expensive cost of task creation, a *lazy task creation* [18, 3] strategy can be used: it postpones the creation of a task until a processor is ready to run it. A simple modification of our rules could make this strategy explicit. Though rule (*migrate task*) does not enforce any migration strategy, we think that task stealing [1, 3] would be appropriate; according to this strategy, a processor that becomes idle steals a task from a heavily loaded processor.

Queinnec’s ICSLA [27, 24, 26, 25] is a dialect of Lisp offering primitives for parallelism, transparent migration of objects, and maintenance of their cache coherence over the network. Queinnec’s purpose is different from ours: as he does not rely on transparent annotations, he does not preserve the sequential meaning of programs. However, he proposes a caching mechanism which is certainly lacking in our CEKDS machine. Although his notion of coherency is not suitable to the CEKDS machine because it is not relative to the sequential evaluation, the protocol that he proposes with lazy propagation of updated values is an interesting technique that would be worth investigating for our semantics.

4 Conclusion

Traditional approaches to distributed computing favour high-performance over ease of programming. We believe that there is a need for a high-level paradigm to program distributed systems. By supplying transparent annotations to create remote computations, we provide

the programmer with the illusion that a distributed system is programmable like a sequential one, because the runtime system itself takes care of task and data migrations, race conditions, or critical sections.

This paper is the first step in this direction: we propose an abstract machine with a distributed store and prove that task migration is transparent to the user. Work is under way to refine this semantics and to precisely investigate the issue of data migration and distributed garbage collection within this framework.

References

1. Robert D. Blumofe, Christopher F. Joerg, Bradley C. Kuszmaul, C. E. Leiserson, Keith H. Randall, and Y. Zhou. Cilk: an Efficient Multithreaded Runtime System. In *PPOPP'95*, 1995.
2. K. M. Chandy and Kesselman C. CC++: A Declarative, Concurrent, Object Oriented Programming Notation. Technical Report CS-92-01, California Institute of Technology, 1992.
3. Marc Feeley. *An Efficient and General Implementation of Futures on Large Scale Shared-Memory Multiprocessors*. PhD thesis, Brandeis University, 1993.
4. M. Felleisen and D. Friedman. Control Operators, the SECD-Machine and the λ -Calculus. In *Formal Description of Programming Concepts III*, pages 193–217, 1986. Elsevier Pub.
5. Matthias Felleisen and Daniel P. Friedman. A Reduction Semantics for Imperative Higher-Order Languages. In *PARLE'87*, LNCS 259, pages 206–223. Springer-Verlag, 1987.
6. Cormac Flanagan and Matthias Felleisen. The Semantics of Future and Its Use in Program Optimization. In *POPL'95*. Also in Technical Reports 238, 239, Rice University, 1994.
7. I. Foster, R. Olson, and S. Tuecke. Productive Parallel Programming: The PCN Approach. *Scientific Programming*, 1(1):51–66, 1992.
8. Message Passing Interface Forum. A Message-Passing Interface Standard. Technical report, University of Tennessee, Knoxville, Tennessee, June 1995.
9. I. Foster, C. Kesselman, and S. Tuecke. The Nexus Approach to Integrating Multithreading and Communications. Math. and Comp. Sci. Division, Argonne National Laboratory, 1995.
10. Al. Geist and al. PVM 3 User's Guide and Reference Manual. Technical report, Oak Ridge National Laboratory, Knoxville, Tennessee, May 1993.
11. Robert H. Halstead, Jr. New Ideas in Parallel Lisp : Language Design, Implementation. In *Parallel Lisp : Languages and Systems*, LNCS 441, pages 2–57, 1990.
12. Nevin Heintze. Set-Based Analysis of ML Programs. In *Proceedings of the 1994 ACM Conference on Lisp and Functional Programming*, pages 306–317, Orlando, Florida, June 1994.
13. Takayasu Ito and Manabu Matsui. A Parallel Lisp Language Pailisp and its Kernel Specification. In *Parallel Lisp : Languages and Systems*, LNCS 441, pages 58–100, 1990.
14. Takayasu Ito and Tomohiro Seino. On Pailisp Continuation and its Implementation. In *Proceedings of the ACM SIGPLAN workshop on Continuations CW92*, pages 73–90, 1992.
15. Morry Katz and Daniel Weise. Continuing Into the Future: On the Interaction of Futures and First-Class Continuations. In *LFP'90*, pages 176–184, June 1990.
16. Bernard Lang, Christian Queinnec, and José Piquet. Garbage Collecting the World. In *POPL'92*, pages 39–50, Albuquerque, New Mexico, 1992.
17. James S. Miller. *MultiScheme : A Parallel Processing System Based on MIT Scheme*. PhD thesis, MIT, 1987.
18. Eric Mohr, David A. Kranz, and Robert H. Halstead. Lazy Task Creation : a Technique for Increasing the Granularity of Parallel Programs. In *LFP'90*, pages 185–197, June 1990.
19. Luc Moreau. *Sound Evaluation of Parallel Functional Programs with First-Class Continuations*. PhD thesis, University of Liège, June 1994. Also available by anonymous ftp from <ftp.montefiore.ulg.ac.be> in directory pub/moreau.
20. Luc Moreau. The Semantics of Scheme with Future. In *In ACM SIGPLAN International Conference on Functional Programming (ICFP'96)*, Philadelphia, May 1996.
21. Luc Moreau. Correctness of a Distributed-Memory Model for Scheme. Technical report M96/3, University of Southampton, 1996.
22. Luc Moreau and Daniel Ribbens. The Semantics of pcall and fork. In *PSLS 95 – Parallel Symbolic Languages and Systems*, LNCS 1068, Beaune, France, October 1995.
23. Julian Padget. Controlling (Virtual) Multicomputers. In *Massively Parallel Computer Systems (MPCS'94)*, pages 102–112. IEEE Computer Society Press, 1994.
24. Christian Queinnec. Locality, Causality and Continuations. In *Proceedings of the 1994 ACM Conference on Lisp and Functional Programming*, Orlando, Florida, June 1994.
25. Christian Queinnec. Sharing mutable objects and controlling groups of tasks in a concurrent and distributed language. In *TPPP'94*, LNCS 700, pages 70–93, Sendai (Japan), 1994.
26. Christian Queinnec. DMEROON: a Distributed Class-based Causally-coherent Data Model: Preliminary Report. In *Parallel Symbolic Languages and Systems.*, LNCS 1068, 1995.
27. C. Queinnec and D. De Roure. Design of a Concurrent and Distributed Language. In *Parallel Symbolic Computing: Languages, Systems and Applications*, LNCS'748, p. 234–259, 1992.
28. Pete Tinker and Morry Katz. Parallel Execution of Sequential Scheme with ParaTran. In *LFP'88*, pages 28–39, 1988.

This article was processed using the L^AT_EX macro package with LLNCS style