

# Continuing Into the Future: the Return (Invited Paper)

Luc Moreau\*  
University of Southampton  
L.Moreau@ecs.soton.ac.uk

## Abstract

`future` is an annotation that indicates which expressions of a program may be evaluated in parallel. By definition, `future` is transparent, i.e. annotated programs return the same results as in the absence of annotations. In such a framework, the interaction of parallelism and first-class continuations has been considered as a delicate matter for a long time. Indeed, unrestricted parallelism and first-class continuations may lead to non-deterministic programs, which is contradictory to the notion of annotation. In this paper, we overview the formal semantics of `future` and first-class continuations. The semantics is an abstract machine that models a parallel computer with a shared memory.

## 1 Introduction

The continuation of an expression is defined as the computation that remains to be performed after evaluating the expression [22]. Some languages, like Scheme [21, 23] or Standard ML of New-Jersey [1] provide the programmer with *first-class* continuations; in these languages, continuations have the same status as numerical values, i.e. they can be passed in argument to or returned by functions, or stored in data structures. First-class continuations are useful to program control structures like backtracking, coroutines, engines, or exceptions [6, 8, 9].

Our interest is to design a language that would allow the programmer to build parallel applications easily. In the approach called “*parallelism by annotations*”, programming languages are extended with annotations by which the programmer points out the expressions that may be evaluated in parallel. By definition, annotations must be *transparent*, that is, annotated terms must return the same result as in the absence of annotations. This approach to parallelism is high-level because transparent annotations avoid the programmer to concentrate on parallelism-specific problems such as deadlocks, race conditions, and non-determinism.

The annotation `future`, initially proposed by Baker and Hewitt [2], is the construct providing parallelism in Halstead’s MultiLisp [7]. Intuitively, an expression `(future exp)` immediately returns a new value, called *placeholder*, which is a data structure with one slot; in parallel, a new task is created to evaluate the argument of `future` and to store its value in the placeholder. Hence, `future` creates a “producer-consumer” type of parallelism, where a producer task computes the value of `exp` in parallel with a task which continues the evaluation as if it already had obtained the value of `exp`.

However, providing transparent annotations is not a trivial task, especially in the presence of first-class continuations. Indeed, using first-class continuations, one can write programs whose final values depend on the evaluation order. If we add parallelism to the language in a uncontrolled way, the evaluation order changes, and the value of programs may become non-deterministic, which is contrary to the notion of transparent annotation.

The simple following program illustrates the problem:

```
(callcc (λ exit. (cons (future (exit 1)) (future (exit 2)))))
```

 (1)

Each `future` creates a new task to evaluate its argument, and returns a new placeholder as value. As a result, this expression leads to three tasks running in parallel and respectively evaluating `(exit 1)`, `(exit 2)`, and the `cons` primitive. Applying the continuation `exit` on a value has the effect of returning the value as a final result. Consequently, if parallelism is not controlled, three different results may be returned: 1, 2, or a pair.

---

\*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.

So far, implementation and efficiency questions [3, 7, 10, 11, 12, 13, 24] have mainly motivated research on the interaction between annotation-based parallelism and first-class continuations in Lisp-like. Recently only, the author of this paper defined a semantic framework for functional programs with first-class continuations and annotations for parallelism [14, 15, 16, 17, 18].

In their LFP'90 paper, entitled “Continuing Into the Future”, Katz and Weise [12] describe an implementation of future in the presence of first-class continuations. The goal of this paper is to present the formal semantics of future in this framework; its proof of correctness can be found in [17].

## 2 The Formal Semantics

Our purpose is to define the semantics of a functional language with first-class continuations. Its set of terms  $\Lambda_f$ , defined in Figure 1, is a  $\lambda$ -calculus extended with conditionals, constants, a primitive `callcc` to capture continuations, and a future construct. The semantics is formalised by the F-PCEKS-machine, which is an abstract machine that models a parallel computer with a shared memory in the tradition of MultiLisp systems. It extends Felleisen and Friedman’s CEK-machine [4].

The CEK-machine [4] is an abstract machine able to evaluate sequential functional programs with first-class continuations, i.e. programs of  $\Lambda_f$  without the construct `future`. Its configuration, called a *computational state*, can be either  $\text{Ev}\langle M, \rho, \kappa \rangle$  or  $\text{Ret}\langle V, \kappa \rangle$ . The first configuration models the evaluation of a term  $M$  in an environment  $\rho$ , with a continuation  $\kappa$ , while the second configuration designates the return of a value  $V$  to a continuation  $\kappa$ . Continuations, representing the rest of the computation, are encoded by a data structure called *continuation code*. A complete definition of the state space appears in Figure 1. Let us observe that closures are represented as  $\langle \text{cl } \lambda x.M, \rho \rangle$ , while first-class continuation are denoted by  $\langle \text{co } \kappa \rangle$ .

Transitions rules of the CEK-machine are displayed in the first part of Figure 2. Details about these transitions can be found in [4, 15]. In order to evaluate a program  $P$ , we start the CEK-machine in an initial configuration  $\text{Ev}\langle P, \emptyset, (\mathbf{init}) \rangle$  and end the computation if a configuration  $\text{Ret}\langle V, (\mathbf{init}) \rangle$  can be reached. The final answer is defined as  $\text{Unload}[V, \emptyset]$ . According to the function *Unload*, answers are values where closures and first-class continuations are replaced by a tag **procedure**.

The F-PCEKS-machine generalises the CEK-machine by providing parallelism. It is composed of several CEK-configurations that have access to a shared memory. Each CEK-configuration represents a task in the parallel machine; new tasks can be created with the future construct. The second part of Figure 2 displays the transitions related to parallelism in the F-PCEKS-machine.

In the introduction, we described future as a construct creating a producer-consumer type of parallelism. Synchronisations between the producer and consumer take place when the consumer requires the content of the placeholder, while its value is still being computed by the producer. Requiring the value of a placeholder is called *touching* the placeholder. It is performed by primitives like `+`, `*`, `car`, `cdr`, which can only be executed if they receive the actual value of their argument; these primitives are said to be *strict*. In our semantics, we suppose that a touch primitive is explicitly introduced for each strict operation by a translation  $\mathcal{X}$ . Let us observe that touch is also introduced in operator position of applications and in predicate position of conditionals.

A configuration, or *state*, of the F-PCEKS-machine is represented by a set of active tasks, a store, and a set of suspended tasks. Each task is a triple composed of a CEK-configuration, a legitimacy (to be explained), and a task name.

The semantics of the construct `future` is given by rule (*fork*). It allocates a new placeholder  $ph$ , which is a data structure represented by  $\langle \text{ph } \alpha \rangle$ , referring to a new location  $\alpha$  in the shared store. Then, it creates a new task  $\text{Ret}\langle ph, \kappa \rangle$ , which continues the evaluation with the placeholder as if it already had received the value of the argument of future. In parallel, the initial task evaluates the argument of future, with a continuation  $(\kappa \text{ det}(ph, \ell))$  which means that the value to be obtained should be stored in the placeholder. Parallelism is modelled in the machine by the possibility to evaluate any task in the set of active tasks.

When the argument of future gets evaluated, rule (*determine*) stores the value obtained into the location associated with the placeholder and removes the current task from the set of active tasks. This action is called *determining* the placeholder to the value. A placeholder whose associated location is empty is said to be *undetermined*.

Rules (*touch*) and (*touch suspend*) deal with the situation where the primitive touch is applied on a value  $V$ . Both call the auxiliary function touch on  $V$ . If  $V$  differs from a placeholder, then  $V$  is returned and evaluation proceeds on the value. Otherwise, if  $V$  is a determined placeholder, we recursively proceed on its content. Finally, if  $V$  is an undetermined placeholder, the associated producer task has not yet produced a value, and the touching

---

$M_u \in \Lambda_f$	::=	$V_u \mid (M_u M_u) \mid (\text{if } M_u M_u M_u) \mid (\text{future } M_u)$	(User Term)
$V_u \in \text{Value}_f$	::=	$a \mid f \mid x \mid (\lambda x.M_u)$	(User Value)
$\mathcal{M} \in \text{State}_{fpceks}$	::=	$\langle T, \sigma, S \rangle$	(State)
$t \in \text{Task}$	::=	$\langle C, \ell \rangle_\tau$	(Task)
$C \in \text{CoSt}$	::=	$\text{Ev}\langle M, \rho, \kappa \rangle \mid \text{Ret}\langle V, \kappa \rangle$	(Computational State)
$\rho \in \text{Env}$	::=	$\{(x_1 V_1) \dots (x_n V_n)\}$	(Environment)
$S$	::=	$\{t_1, \dots, t_n\}$	(Suspended Tasks)
$T$	=	$\{t_1, \dots, t_n\}$	(Active Tasks)
$M \in \Lambda_{fpceks}$	::=	$V_s \mid (M M) \mid (\text{if } M M M) \mid (\text{future } M)$	(Term)
$V_s \in \text{SValue}_{fpceks}$	::=	$c \mid x \mid (\lambda x.M)$	(Syntactic Value)
$W \in \text{PValue}_{fpceks}$	::=	$c \mid \langle \text{cl } \lambda x.M, \rho \rangle \mid \langle \text{co } \kappa \rangle \mid (\text{cons } V V) \mid f_c$	(Proper Value)
$V \in \text{Value}_{fpceks}$	::=	$W \mid ph$	(Runtime Value)
$\kappa \in \text{CCode}$	::=	$(\text{init}) \mid (\kappa \text{ fun } V) \mid (\kappa \text{ arg } M, \rho)$ $\mid (\kappa \text{ cond}(M, M, \rho)) \mid (\kappa \text{ det } (ph, \ell))$	(Continuation code)
$ph \in \text{Placeholder}$	::=	$\langle \text{ph } \alpha \rangle$	(Placeholder)
$\ell \in \text{Legitimacy}$	::=	$\langle \text{leg } \alpha \rangle$	(Legitimacy)
$\sigma \in \text{Store}_{fpceks}$	::=	$\{(\alpha_1 O_1) \dots (\alpha_n O_n)\}$	(Store)
$O \in \text{Contents}_{fpceks}$	::=	$\ell \mid \perp$	(Store Content)
$\alpha \in \text{Loc}$	=	$\{\alpha_0, \alpha_1, \dots\}$	(Location)
$\tau \in \text{Tid}$	=	$\{\tau_0, \tau_1, \dots\}$	(Task Identifier)
$c \in \text{Const}$	::=	$a \mid d \mid f \mid f_i$	(Constant)
$f_c \in \text{PApp}$	::=	$(\text{cons } V)$	(Partial Application)
$g \in \text{AValue}$	::=	$\langle \text{cl } \lambda x.M, \rho \rangle \mid f \mid f_i \mid f_c \mid \langle \text{co } \kappa \rangle$	(Applicable Value)
$f_n \in \text{NSTP}$	=	$\{\text{cons, callcc}\}$	(Non Strict Primitives)
$A \in \text{Answer}$	::=	$c \mid (\text{cons } A A) \mid \text{procedure}$	(Answer)
$x \in \text{Vars}$	=	$\{x, y, z \dots\}$	(User Variable)
$a \in \text{BConst}$	=	$\{\text{true, false, nil, 0, 1, \dots}\}$	(Basic Constant)
$f \in \text{FConst}$	=	$\{\text{cons, car, cdr, callcc}\}$	(Functional Constant)
$f_i \in \text{IFConst}$	=	$\{\text{touch}\}$	(Internal Functional Constant)
$d \in \text{DConst}$	=	$\{\text{error}\}$	(Distinguished Constant)
Touch function:		Initial Legitimacy: $\ell_i = \langle \text{leg } \alpha_i \rangle$	
touch	: $\text{Value}_{fpceks} \times \text{Store}_{fpceks}$ $\rightarrow \text{Value}_{fpceks}$	Initial Store: $\{(\alpha_i \perp)\}$	
touch $_\sigma(V)$	= $V$ if $V \neq ph$	Store extension:	
touch $_\sigma(\langle \text{ph } \alpha \rangle)$	= touch $_\sigma(\sigma(\alpha))$ if $\sigma(\alpha) \neq \perp$	$\sigma[\alpha \leftarrow V] = \sigma \cup \{(\alpha V)\}$ with $\alpha \notin \text{DOM}(\sigma)$	
touch $_\sigma(\langle \text{ph } \alpha \rangle)$	= $\langle \text{ph } \alpha \rangle$ if $\sigma(\alpha) = \perp$	Store update:	
Mandatory descendant:		$\sigma' = \sigma[b := V]$ is defined if $\exists (b V') \in \sigma$ and $\sigma' = (\sigma \setminus \{(b V')\}) \cup \{(b V)\}$	
$\ell_0 \rightsquigarrow_\sigma \ell_1$ if	$\ell_0 = \ell_1,$ or	Store Content:	
$\ell_0 \rightsquigarrow_\sigma \ell_1$ if $\ell_0 = \langle \text{leg } \alpha_0 \rangle$ and $\sigma(\alpha_0) \neq \perp$		$\sigma(\alpha) = V$ if $(\alpha V) \in \sigma$	
$\text{Unload}$	: $\text{Value}_{fpceks} \rightarrow \text{Answer}$		
$\text{Unload}[c, \sigma]$	= $c$		
$\text{Unload}[(\text{cons } V_1 V_2), \sigma]$	= $(\text{cons } \text{Unload}[V_1, \sigma] \text{Unload}[V_2, \sigma])$		
$\text{Unload}[(\text{cl } \lambda x.M, \rho), \sigma]$	= <b>procedure</b>		
$\text{Unload}[(\text{co } \kappa), \sigma]$	= <b>procedure</b>		
$\text{Unload}[\langle \text{ph } \alpha \rangle, \sigma]$	= $\text{Unload}[\sigma(\alpha), \sigma]$		
Explicit translation: $\mathcal{X} : \Lambda_f \rightarrow \Lambda_{fpceks}$		Environment Extension:	
$\mathcal{X}[\text{car}]$	= $\lambda x.(\text{car } (\text{touch } x))$	$\rho[X \leftarrow V] = \rho \cup \{(x V)\}$	
$\mathcal{X}[\text{cdr}]$	= $\lambda x.(\text{cdr } (\text{touch } x))$	Environment Access:	
$\mathcal{X}[(\text{future } M_u)]$	= $(\text{future } \mathcal{X}[M_u])$	$\rho(x) = V$ if $(x V) \in \rho$	
$\mathcal{X}[(M_{u1} M_{u2})]$	= $(\lambda m_1 m_2.((\text{touch } m_1) m_2)) \mathcal{X}[M_{u1}] \mathcal{X}[M_{u2}]$		
$\mathcal{X}[(\lambda x.M)]$	= $(\lambda x.\mathcal{X}[M])$		
$\mathcal{X}[(\text{if } M_{u1} M_{u2} M_{u3})]$	= $(\text{if } (\text{touch } \mathcal{X}[M_{u1}]) \mathcal{X}[M_{u2}] \mathcal{X}[M_{u3}])$		
$\mathcal{X}[x]$	= $x$ if $x \in \text{Vars} \cup \text{BConst} \cup \text{NSTP}$		

---

Figure 1: State space of the F-PCEKS-machine

---

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

---

$\langle\{ \langle C, \ell \rangle_\tau \} \cup T, \sigma, S\rangle$	$\rightarrow_{fpceks}$	$\langle\{ \langle C_1, \ell \rangle_\tau \} \cup T, \sigma, S\rangle$ if $C \rightarrow_{cek} C_1$	(sequential)
$\langle\{ \langle \text{Ev}\langle(\text{future } M), \rho, \kappa\rangle, \ell \rangle_\tau \} \cup T, \sigma, S\rangle$	$\rightarrow_{fpceks}$	$\langle\{ \langle \text{Ev}\langle M, \rho, (\kappa \text{ det } ph, \ell_1)\rangle, \ell \rangle_\tau, \langle \text{Ret}\langle ph, \kappa \rangle, \ell_1 \rangle_{\tau_1} \} \cup T, \sigma_1, S\rangle$ with $ph = \langle\text{ph } \alpha\rangle, \ell_1 = \langle\text{leg } \alpha_1\rangle, \alpha_1 \notin \text{DOM}(\sigma), \alpha \notin \text{DOM}(\sigma)$ with $\sigma_1 = \sigma[\alpha \leftarrow \perp][\alpha_1 \leftarrow \perp], \tau_1 \notin \text{FN}(T \cup S) \cup \{\tau\}$	(fork)
$\langle\{ \langle \text{Ret}\langle V, (\kappa \text{ det } \langle\text{ph } \alpha\rangle), \langle\text{leg } \alpha_1\rangle \rangle, \ell_2 \rangle_\tau \} \cup T, \sigma, S\rangle$	$\rightarrow_{fpceks}$	$\langle T_1, \sigma_1, S_1 \rangle$ , if $\sigma(\alpha) = \perp$ $T_1 = (T \cup T_2), S_1 = (S \setminus T_2), \sigma_1 = \sigma[\alpha_1 := \ell_2][\alpha := V]$ with $T_2 = \{ \langle \text{Ret}\langle \langle\text{ph } \alpha\rangle, (\kappa' \text{ fun touch}) \rangle, \ell_3 \rangle_{\tau_1} \in S \}$	(determine)
$\langle\{ \langle \text{Ret}\langle V, \kappa \rangle, \ell_2 \rangle_\tau \} \cup T, \sigma, S\rangle$	$\rightarrow_{fpceks}$	$\langle\{ \langle \text{Ret}\langle V, \kappa \rangle, \ell_2 \rangle_\tau \} \cup T, \sigma, S\rangle$ , if $\sigma(\alpha) \neq \perp$	(determinen)
$\langle\{ \langle \text{Ret}\langle V, (\kappa \text{ fun touch}) \rangle, \ell \rangle_\tau \} \cup T, \sigma, S\rangle$	$\rightarrow_{fpceks}$	$\langle\{ \langle \text{Ret}\langle \text{touch}_\sigma(V), \kappa \rangle, \ell \rangle_\tau \} \cup T, \sigma, S\rangle$ if $\text{touch}_\sigma(V) \in PValue$	(touch)
$\rightarrow_{fpceks}$	$\rightarrow_{fpceks}$	$\langle T, \sigma, \{ \langle \text{Ret}\langle \langle\text{ph } \alpha\rangle, (\kappa \text{ fun touch}) \rangle, \ell \rangle_\tau \} \cup S\rangle$ if $\text{touch}_\sigma(V) = \langle\text{ph } \alpha\rangle, \sigma(\alpha) = \perp$	(touch suspend)

---

$\langle T, \sigma \rangle \xrightarrow{1, a}_{fpceks} \langle T', \sigma' \rangle$  if  $\langle T, \sigma \rangle \rightarrow_{fpceks} \langle T', \sigma' \rangle$ ,  $a = 1$  if  $\ell \rightsquigarrow_\sigma \ell_i$ , and  $a = 0$  otherwise. (base)

$\mathcal{M} \xrightarrow{a+a', b+b'}_{fpceks} \mathcal{M}_2$  if  $\mathcal{M} \xrightarrow{a, b}_{fpceks} \mathcal{M}_1$  and  $\mathcal{M}_1 \xrightarrow{a', b'}_{fpceks} \mathcal{M}_2$ . (transitive)

$$\text{eval}_{fpceks}(P) = \begin{cases} \text{Unload}[V, \sigma] & \text{if } \langle \{ \langle \text{Ev}\langle \mathcal{X}[P], \emptyset, (\text{init}) \rangle, \ell_i \rangle_{\tau_0} \}, \sigma_i, \emptyset \rangle \xrightarrow{*}_{fpceks} \langle T, \sigma, S \rangle, \\ & \text{with } \langle \text{Ret}\langle V, (\text{init}) \rangle, \ell \rangle_\tau \in T, \text{ such that } \ell \rightsquigarrow_\sigma \ell_i \\ \perp & \text{if } \forall j \in \mathbf{IN}, \exists \mathcal{M}_j \in \text{State}_{fpceks}, n_j, m_j \in \mathbf{IN}, \text{ such that} \\ & \mathcal{M}_0 = \langle \{ \langle \text{Ev}\langle \mathcal{X}[P], \emptyset, (\text{init}) \rangle, \ell_i \rangle_{\tau_0} \}, \sigma_i, \emptyset \rangle, \\ & \mathcal{M}_j \xrightarrow{n_j, m_j}_{fpceks} \mathcal{M}_{j+1}, m_j > 0 \end{cases}$$

Figure 2: Evaluator specification of the F-PCEKS-machine

---

task should be suspended, by transferring it to the set of suspended task. A suspended task is awakened when the placeholder get determined by rule (*determine*).

So far, we deliberately avoided to talk about continuations, and did not comment on legitimacies and rule (*determinen*). Placeholders are data structures that can receive at most one value: once determined, a placeholder remains constant. This invariant is automatically satisfied in the absence of first-class continuations by the functional nature of the language. However, when using first-class continuations, an expression may return “several times”, i.e. several values may be passed to its continuation. Therefore, in order to preserve the property of placeholders, rule (*determine*) is allowed to be fired only if the placeholder is undetermined. (We assume here the atomicity of the transitions.) If the placeholder is already determined, rule (*determinen*) evaluates the continuation as if no future had existed (cfr. Katz and Weise [12]).

However, this does not guarantee the transparency of the future construct. Indeed, the program (1) can still return several final results. Following Katz and Weise, we use a notion of legitimacy to distinguish *mandatory* from *speculative* tasks. The mandatory task is the one that performs the transition that would be performed if evaluation had been sequential; all the other tasks are speculative.

A legitimacy ( $\text{leg } \alpha$ ), like a placeholder, is a data-structure that refers to a location in the shared store, but unlike a placeholder, it is not considered as a value because there exists no primitive to reify it to the status of value. We shall also use the terms “undetermined” and “determined” for legitimacies.

When starting the execution of a program, the initial task is given the initial legitimacy  $l_i$ . Whenever a task  $\tau$  evaluates a future, rule (*fork*) creates a task  $\tau'$ , and allocates a new legitimacy  $l_1$  in addition to the new placeholder. After transition (*fork*), task  $\tau$  still has the same legitimacy, but is now evaluating the future argument with a continuation ( $\kappa \text{ det}(ph, l_1)$ ), where  $ph$  is the placeholder to determine and  $l_1$  the legitimacy of  $\tau'$ . Meanwhile, the task  $\tau'$  begins to evaluate the continuation of future with the new placeholder  $ph$ . We know that task  $\tau'$  performs a speculative computation on behalf of  $\tau$  because the legitimacy of  $\tau'$  is  $l_1$ , and the continuation of task  $\tau$  contains a code ( $\kappa \text{ det}(ph, l_1)$ ), where  $l_1$  is explicit.

Rule (*determine*) keeps track of legitimacy as follows. If the placeholder  $\langle ph \ \alpha \rangle$  gets determined to a value  $V$ , the task that speculatively consumes this placeholder becomes dependent on the value  $V$ . This dependency is made explicit by giving the consuming task the legitimacy of the producing task. More precisely, the legitimacy of the consumer task is determined to the legitimacy of the producer task. So, legitimacy is passed between tasks as a token, whenever a placeholder gets determined.

When a legitimacy ( $\text{leg } \alpha$ ) gets determined, location  $\alpha$  receives a legitimacy, which might also be determined. Hence, as evaluation proceeds, chains of legitimacies get formed in memory. We define a relation  $l_1 \rightsquigarrow_\sigma l_2$  stating that there is a path from legitimacy  $l_1$  to legitimacy  $l_2$ , which means that control has flowed from a task with legitimacy  $l_2$  to a task with legitimacy  $l_1$ . Intuitively, legitimacy models the fact that a sequential implementation would have performed the evaluation done by the tasks with legitimacies  $l_2$  to  $l_1$ . The relation  $\rightsquigarrow$  is used to determine whether a final value, i.e. a value returned to the (**init**) continuation, is a valid answer. A valid answer is produced by the task whose legitimacy  $l$  is such that  $l \rightsquigarrow_\sigma l_i$ . The initial legitimacy  $l_i$  is a pre-allocated legitimacy which serves as a marker for the end of legitimacy chains. This legitimacy always remains undetermined because the initial program does not depend on any placeholder.

The third part of Figure 2 defines the evaluation relation of the F-PCEKS-machine. The evaluation of a program  $P$  begins with an initial task  $\langle \text{Ev}\langle P, \emptyset, (\mathbf{init}) \rangle, l_i \rangle_{\tau_0}$  and ends if a final configuration can be reached; a final configuration contains a task  $\langle \text{Ret}\langle V, (\mathbf{init}) \rangle, l \rangle_\tau$ , such that  $l \rightsquigarrow_\sigma l_i$ , i.e. the task returning the value is mandatory.

Let us note that there exist programs with a finite mandatory computation but a possible unbounded number of speculative transitions, like for instance  $(\text{callcc} \lambda k. ((\text{future } (k \ 1)) \ \Omega))$ , where  $\Omega$  is a sequential divergent term. Therefore, divergence should be defined with the greatest care: a program is said to diverge if it regularly often performs mandatory transitions. Hence, we introduce the relation  $\mathcal{M}_1 \rightarrow^{n,m} \mathcal{M}_2$  to denote a transition between two configurations  $\mathcal{M}_1$  and  $\mathcal{M}_2$ , involving  $n$  steps among which  $m$  were mandatory.

### 3 Related Work and Conclusion

This paper overviewed the semantics of future in the presence of first-class continuations. Details and proof of correctness can be found in [17, 18]. The same framework is also extended to deal with side-effects. Flanagan and Felleisen [5] defined the semantics of future in a purely functional language. Their goal was to derive a program optimisation, called the *touch optimisation*, which removed provably redundant touch operations. Such an optimisation would also be useful for the language we deal with.

Katz and Weise [12] introduced the notion of *legitimacy*. For a deep study of the performance of MultiLisp, we refer to Feeley's thesis [3]. Halstead [7, page 19] gave three criteria which should be satisfied by an implementation with transparent annotations for parallelism. Thanks to the proof of correctness, our semantics satisfies these criteria.

There exist other approaches to parallelism which do not preserve the sequential meaning of programs. Among others, let us cite Queinnec's ICSLAS [19, 20], and Ito's PAILISP [10, 11], which both define new semantics of continuations in a parallel framework. Let us observe that their goals differ from ours because their primitives for parallelism add expressiveness to the sequential core, like for instance Queinnec's `pcall` or Ito's *parallel-or*.

## References

- [1] Andrew W. Appel and David B. MacQueen. A Standard ML compiler. In *Functional Programming Languages and Computer Architecture*, number LNCS 274, pages 301–324. Springer-Verlag, 1987.
- [2] Henry Baker and Carl Hewitt. The Incremental Garbage Collection of Processes. Technical Report AI Memo 454, M.I.T., Cambridge, Massachusetts, March 1977.
- [3] Marc Feeley. *An Efficient and General Implementation of Futures on Large Scale Shared-Memory Multiprocessors*. PhD thesis, Brandeis University, 1993.
- [4] Matthias Felleisen and Daniel P. Friedman. Control Operators, the SECD-Machine and the  $\lambda$ -Calculus. In M. Wirsing, editor, *Formal Description of Programming Concepts III*, pages 193–217, Amsterdam, 1986. Elsevier Science Publishers B.V.
- [5] Cormac Flanagan and Matthias Felleisen. The Semantics of Future and Its Use in Program Optimization. In *Proceedings of the Twenty Second Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, January 1995. Technical Reports 238, 239, Rice University, 1994.
- [6] Daniel P. Friedman and Christopher T. Haynes. Constraining Control. In *Proceedings of the Twelfth Annual Symposium on Principles of Programming Languages*, pages 245–254, New Orleans, LA., January 1985. ACM.
- [7] Robert H. Halstead, Jr. New Ideas in Parallel Lisp : Language Design, Implementation. In T. Ito and Robert H. Halstead, editors, *Parallel Lisp : Languages and Systems*, LNCS 441, pages 2–57. Springer-Verlag, 1990.
- [8] Christopher T. Haynes and Daniel P. Friedman. Abstracting Timed Preemption with Engines. *Comput. Lang.*, 12(2):109–121, 1987.
- [9] Christopher T. Haynes, Daniel P. Friedman, and Mitchell Wand. Obtaining Coroutines with Continuations. *Comput. Lang.*, 11(3/4):143–153, 1986.
- [10] Takayasu Ito and Manabu Matsui. A Parallel Lisp Language Pailisp and its Kernel Specification. In T. Ito and Robert H. Halstead, editors, *Parallel Lisp : Languages and Systems*, LNCS 441, pages 58–100. Springer-Verlag, 1990.
- [11] Takayasu Ito and Tomohiro Seino. On Pailisp Continuation and its Implementation. In *Proceedings of the ACM SIGPLAN workshop on Continuations CW92*, pages 73–90, San Francisco, June 1992.
- [12] Morry Katz and Daniel Weise. Continuing Into the Future: On the Interaction of Futures and First-Class Continuations. In *Proceedings of the 1990 ACM Conference on Lisp and Functional Programming*, pages 176–184, June 1990.
- [13] James S. Miller. *MultiScheme : A Parallel Processing System Based on MIT Scheme*. PhD thesis, MIT, 1987.
- [14] Luc Moreau. The PCKS-machine. An Abstract Machine for Sound Evaluation of Parallel Functional Programs with First-Class Continuations. In *European Symposium on Programming (ESOP'94)*, LNCS 788, pages 424–438, Edinburgh, 1994.
- [15] Luc Moreau. *Sound Evaluation of Parallel Functional Programs with First-Class Continuations*. PhD thesis, University of Liège, Belgium, June 1994. Also available by anonymous ftp from `ftp.montefiore.ulg.ac.be` in directory `pub/moreau`.
- [16] Luc Moreau. Non-speculative and Upward Invocation of Continuations in a Parallel Language. In *International Joint Conference on Theory and Practice of Software Development (TAPSOFT/FASE'95)*, LNCS 915, pages 726–740, Aarhus, 1995.
- [17] Luc Moreau. The Semantics of Scheme with Future. Technical Report M95/7, University of Southampton, 1995. 56 pages.
- [18] Luc Moreau. The Semantics of Scheme with Future. In *In ACM SIGPLAN International Conference on Functional Programming (ICFP'96)*, Philadelphia, May 1996.
- [19] Christian Queinnec. Locality, Causality and Continuations. In *Proceedings of the 1994 ACM Conference on Lisp and Functional Programming*, Orlando, Florida, June 1994.
- [20] Christian Queinnec and David De Roure. Design of a Concurrent and Distributed Language. In A. Agarwal, R. H. Halstead, and Takayasu Ito, editors, *Proceedings of the Workshop on Parallel Symbolic Computing: Languages, Systems and Applications*, Boston, Massachusetts, October 1992.
- [21] Jonathan Rees and William Clinger, editors. Revised<sup>4</sup> Report on the Algorithmic Language Scheme. *Lisp Pointers*, 4(3):1–55, July-September 1991.
- [22] Christopher Strachey and Christopher P. Wadsworth. A Mathematical Semantics for Handling Full Jumps. Technical Monography PRG-11, Oxford University Computing Laboratory, Programming Research Group, Oxford, England, 1974.
- [23] Gerald Jay Sussman and Guy Lewis Steele, Jr. SCHEME: an Interpreter for Extended Lambda Calculus. Technical Report 349, MIT AI Lab, December 1975.
- [24] Pete Tinker and Morry Katz. Parallel Execution of Sequential Scheme with ParaTran. In *Proceedings of the 1988 ACM Conference on Lisp and Functional Programming*, pages 28–39, Snowbird, Utah, July 1988.