

The Semantics of Future

Cormac Flanagan
Matthias Felleisen

Rice COMP TR94-238

October 1994

Department of Computer Science
Rice University
P.O. Box 1892
Houston, TX 77251-1892

Copyright ©1994 by
Cormac Flanagan and Matthias Felleisen

The Semantics of Future

Cormac Flanagan* Matthias Felleisen*

Department of Computer Science,
Rice University,
Houston, TX 77251-1982

Abstract

The **future** annotation introduced by MultiLisp provides a simple method for taming the implicit parallelism of functional programs. Prior research on **futures** has concentrated on implementation and design issues, and has largely ignored the development of a semantic characterization of **futures**. This paper presents four operational semantics for an idealized functional language with **futures** with varying degrees of intensionality. The first semantics defines **future** to be a semantically-transparent annotation. The second semantics interprets a **future** expression as a potentially parallel task. The third semantics explicates the coordination of parallel tasks and the need for *touch* operations on placeholder-strict arguments to certain primitive operations by introducing placeholder objects. The fourth and last semantics is a low-level refinement of the third semantics, which explicates just enough information to permit the smooth derivation of program analyses. The paper includes proofs showing the equivalence of these semantics.

*Supported in part by NSF grant CCR 91-22518 and a sabbatical at Carnegie Mellon University.

Contents

1	Futures for Parallel Computation	1
2	A Functional Language with Futures	2
2.1	Syntax	2
2.2	Definitional Semantics	4
3	Parallel Operational Semantics	5
3.1	Specification of the $P(C)$ -machine	5
	State Space	5
	Transition Rules	7
	Indeterminism	8
	Evaluation	8
	Placeholder Transparency and Synchronization	9
3.2	Consistency of the $P(C)$ -Machine	11
3.3	Correctness of the $P(C)$ -Machine	17
4	Placeholder Object Semantics	18
4.1	Specification of the $P(C_{ph})$ -machine	19
4.2	Correctness of the $P(C_{ph})$ -machine	21
5	The Low-Level Operational Semantics	24
5.1	Specification of the $P(CEK)$ -machine	24
5.2	Correctness of the $P(CEK)$ -machine	27
6	Related Work	27
7	Applications of the Semantics	28
A	Correctness Proof for the $P(CEK)$-Machine	29

1 Futures for Parallel Computation

Programs in functional languages offer numerous opportunities for executing program components in parallel. In a call-by-value language, for example, the evaluation of every function application could spawn a parallel thread for each argument expression. However, if such a strategy were applied indiscriminately, the execution of a program would generate far too many parallel threads. The overhead of managing these threads would clearly outweigh any benefits from parallel execution.

The **future** annotations of MultiLisp [1, 11] and its Scheme successors provide a simple method for taming the implicit parallelism of functional programs. If a programmer believes that the parallel evaluation of some expression outweighs the overhead of creating a separate task, he may annotate the expression with the keyword **future**. An annotated functional program has the same observable behavior as the original program, but the run-time system may choose to evaluate the **future** expression in parallel to the rest of the program.

While past research on **futures** has concentrated on implementation and design issues [4, 13, 20, 16, 18, 15, 19], this technical report focuses on the semantics of **futures**. Specifically, it presents a series of semantics with varying degrees of intensionality. The first semantics equates the value and behavior of **(future e)** as that of e , which is the simplest way to interpret **future** as an annotation. The second semantics interprets an annotated expression as a potentially parallel task. The third semantics explicates the coordination of parallel tasks with the introduction of explicit placeholder objects. Since all of these semantics are formulated as program rewriting systems, the equivalence proofs rely on standard techniques like bisimulation and a diamond lemma. The fourth and last semantics is a low-level refinement of the third semantics, which explicates just enough information to permit the smooth derivation of program analyses.

A companion paper describes the use of our semantics in program optimization [7]. It focuses on the derivation of a provably correct program analysis algorithm from the last semantics. The paper also describes a optimization algorithm for lowering the overhead of task coordination in a language with **futures**, and the effects of implementing this optimization for the Gambit compiler [4].

The presentation of our semantics for **future** proceeds as follows. The second section introduces the language we consider and its definitional semantics, which interprets **futures** as annotations. The third section presents a parallel program rewriting semantics of **futures**, and the fourth section introduces placeholder objects and *touch* operations. The fifth section contains the low-level refinement. The sixth section compares our approach to related work concerning the semantics of parallel functional languages, and the last section discusses some applications for the our semantics. Appendix A contains the correctness proof for the low-level semantics.

Notation We use $f : A \longrightarrow B$ to denote that f is a total function from A to B , and similarly use $f : A \longrightarrow_f B$ to mean that f is a finite map from A to B .

2 A Functional Language with Futures

2.1 Syntax

Our motivation in developing a semantics for **future** is to provide a sound basis for the compilation of programs with **futures**. Compilers typically convert source programs into a simple intermediate representation, and then proceed to compile and optimize the intermediate representation of the program. Therefore, we formulate the semantics of **futures** for an intermediate representation of an idealized functional language. Specifically, we use the subset of A -normal forms [9] of a λ -calculus-like language with a **let** construct, a **future** construct, a conditional and an explicit **apply** primitive: see Figure 1. The language also includes the primitives **cons**, **car**, and **cdr** for list manipulation, which will serve to illustrate the treatment of primitive operations, and an unspecified set of basic primitive constants that includes at least numbers and the empty list **nil**.

The key property of terms in A -normal form is that each intermediate value is explicitly named and where the order of execution follows the lexical nesting of **let**-expressions. The use of A -normal forms facilitates the compile-time analysis of programs since every intermediate value is named [22], and it simplifies the definition of abstract machines [9], a fact that we exploit in the development of abstract machines below.

$M \in \Lambda_a$	$::=$	x $ (\text{let } (x \ V) \ M)$ $ (\text{let } (x \ (\text{future } M)) \ M)$ $ (\text{let } (x \ (\text{car } y)) \ M)$ $ (\text{let } (x \ (\text{cdr } y)) \ M)$ $ (\text{let } (x \ (\text{if } y \ M \ M)) \ M)$ $ (\text{let } (x \ (\text{apply } y \ z)) \ M)$	(Terms)
$V \in \text{Value}$	$::=$	$c \mid x \mid (\lambda x. M) \mid (\text{cons } x \ y)$	(Values)
$x \in \text{Vars}$	$=$	$\{x, y, z, \dots\}$	(Variables)
$c \in \text{Const}$	$=$	$\{\text{nil}, 0, 1, \dots\}$	(Constants)

FIGURE 1: The A -normalized Language Λ_a

A variable occurrence is *free* if it is not bound by an enclosing λ -expression or **let**-expression. A term is *closed* if it contains no free variables. We identify terms that differ only by consistent renaming of bound variables. The operation $M[x \leftarrow V]$ denotes the capture-free substitution of V for all free occurrences of x within M . We use X^0 to denote the set of closed terms of type X (terms, values).

Evaluator:

$$\begin{aligned}
 eval_c : \Lambda_a^0 &\longrightarrow Answers \cup \{\mathbf{error}, \perp\} \\
 eval_c(P) &= \begin{cases} unload_c[V] & \text{if } P \mapsto_c^* V \\ \mathbf{error} & \text{if } P \mapsto_c^* \mathbf{error} \\ \perp & \text{if } \forall i \in \mathbf{N} \exists M_i \in State_c \text{ such that} \\ & P = M_0 \text{ and } M_i \mapsto_c M_{i+1} \end{cases}
 \end{aligned}$$

Data Specifications:

$$\begin{aligned}
 S \in State_c &::= M \mid \mathbf{error} && \text{(States)} \\
 M \in \Lambda_c &::= V && \text{(Run-time Language)} \\
 & \mid (\mathbf{let} (x V) M) \\
 & \mid (\mathbf{let} (x (\mathbf{future} M)) M) \\
 & \mid (\mathbf{let} (x (\mathbf{car} V)) M) \\
 & \mid (\mathbf{let} (x (\mathbf{cdr} V)) M) \\
 & \mid (\mathbf{let} (x (\mathbf{if} V M M)) M) \\
 & \mid (\mathbf{let} (x (\mathbf{apply} V V)) M) \\
 & \mid (\mathbf{let} (x M) M) \\
 V \in Value_c &::= c \mid x \mid (\lambda x. M) \mid (\mathbf{cons} V V) && \text{(Run-time Values)} \\
 F \in FinalState_c &::= V \mid \mathbf{error} && \text{(Final States)} \\
 A \in Answers &::= c \mid \mathbf{procedure} \mid (\mathbf{cons} A A) && \text{(Answers)} \\
 \mathcal{E} \in EvalCtx &::= [] && \text{(Evaluation Contexts)} \\
 & \mid (\mathbf{let} (x \mathcal{E}) M) \\
 & \mid (\mathbf{let} (x (\mathbf{future} \mathcal{E})) M)
 \end{aligned}$$

Unload Function:

$$\begin{aligned}
 unload_c : Value_c^0 &\longrightarrow Answers \\
 unload_c[c] &= c \\
 unload_c[(\lambda x. M)] &= \mathbf{procedure} \\
 unload_c[(\mathbf{cons} V_1 V_2)] &= (\mathbf{cons} unload_c[V_1] unload_c[V_2])
 \end{aligned}$$

Transition Function:

$$\begin{aligned}
 \mathcal{E}[(\mathbf{let} (x V) M)] &\mapsto_c \mathcal{E}[M[x \leftarrow V]] && (bind) \\
 \mathcal{E}[(\mathbf{let} (x (\mathbf{future} V)) M)] &\mapsto_c \mathcal{E}[M[x \leftarrow V]] && (future-id) \\
 \mathcal{E}[(\mathbf{let} (x (\mathbf{car} V)) M)] &\mapsto_c \begin{cases} \mathcal{E}[M[x \leftarrow V_1]] & \text{if } V = (\mathbf{cons} V_1 V_2) \\ \mathbf{error} & \text{if } V \neq (\mathbf{cons} V_1 V_2) \end{cases} && (car) \\
 \mathcal{E}[(\mathbf{let} (x (\mathbf{cdr} V)) M)] &\mapsto_c \begin{cases} \mathcal{E}[M[x \leftarrow V_2]] & \text{if } V = (\mathbf{cons} V_1 V_2) \\ \mathbf{error} & \text{if } V \neq (\mathbf{cons} V_1 V_2) \end{cases} && (cdr) \\
 \mathcal{E}[(\mathbf{let} (x (\mathbf{if} V M_1 M_2)) M)] &\mapsto_c \begin{cases} \mathcal{E}[(\mathbf{let} (x M_1) M)] & \text{if } V \neq \mathbf{nil} \\ \mathcal{E}[(\mathbf{let} (x M_2) M)] & \text{if } V = \mathbf{nil} \end{cases} && (if) \\
 \mathcal{E}[(\mathbf{let} (x (\mathbf{apply} V_1 V_2)) M)] &\mapsto_c \begin{cases} \mathcal{E}[(\mathbf{let} (x N[y \leftarrow V_2]) M)] & \text{if } V_1 = (\lambda y. N) \\ \mathbf{error} & \text{if } V_1 \neq (\lambda y. N) \end{cases} && (apply)
 \end{aligned}$$

FIGURE 2: The sequential C -machine

2.2 Definitional Semantics

The semantics of the A -normalized language is a function from Λ_a^0 to results. A result is either an answer, which is a closed value with all λ -expression replaced by the tag **procedure**,¹ or **error**, indicating that some program operation was misapplied, or \perp , if the program diverges. We specify the definitional semantics for Λ_a^0 using a sequential abstract machine called the C -machine (see Figure 2), whose states are either closed terms in the run-time language Λ_c or the special state **error**, and whose deterministic transition rules are the typical leftmost-outermost reductions of the lambda calculus [5]. Each transition rule also specifies the error semantics of a particular class of expressions. For example, the transition rule for **car** defines that if the argument to **car** is a pair, then the transition rule extracts the first element of the pair. If the argument is not a pair, then the transition rule produces the state **error**.

The only unusual transition rule is the one for **future** expressions. This rule pretends that **future** is the identity operation. It demands that the body of a **future** expression is first reduced to a value, and then replaces the name for the **future** expression with this value.

The definition of the transition function relies on the notion of *evaluation contexts*. An evaluation context \mathcal{E} is a term with a hole $[]$ in place of the next sub-term to be evaluated; *e.g.*, in the term **(let** ($x M_1$) M_2), the next sub-term to be evaluated is M_1 , and thus the definition of evaluation contexts includes **(let** ($x \mathcal{E}$) M).

A machine state is a *final state* if it is either a value or the special state **error**. No transitions are possible from a final state, and for any state that is not a final state, there is a unique transition step from that state to its successor state.

Lemma 2.1 (Uniform Evaluation Theorem) *Let $M \in State_c$.*

1. *If $M \notin FinalState_c$, there exists a unique term M' such that $M \mapsto_c M'$.*
2. *If $M \in FinalState_c$, there is no term M' such that $M \mapsto_c M'$.*

Proof: We use structural induction on $State_c$ to prove that any term that is not a final state matches the lefthand side of exactly one of the clauses for \mapsto_c . Since each of these clauses defines a function on Λ_c , the relation \mapsto_c relates each term that is not a final state to a unique successor term. The proof that a final state has no successor is straightforward. ■

The Unique Evaluation Theorem implies that the relation $eval_c$ is a well-defined total function: Either the transition sequence for a program P terminates in a final state, in which case $eval_c(P)$ is an answer or **error**, or else the transition sequence is infinite, in which case $eval_c(P) = \perp$.

Theorem 2.2 *The evaluator $eval_c$ is a total function.*

Since the evaluator $eval_c$ of the C -machine obviously agrees with the sequential semantics of the underlying functional language, **future** is clearly nothing but an annotation.

¹We remove λ -expressions from answers so that the observable behavior of programs does not depend on the terms themselves, but only on their meaning.

3 Parallel Operational Semantics

The sequential C -machine defines **future** as an annotation, and ignores the *intension* of **future** as an advisory instruction concerning parallel evaluations. To understand this intensional aspect of **future** annotations, we need a semantics of **future** that models the concurrent evaluation of **future** expressions. For this purpose, we reformulate the C machine to account for parallelism. The result is the $P(C)$ -machine.

Evaluator:

$$\begin{aligned} eval_{pc} : \Lambda_a^0 &\longrightarrow Answers \cup \{\mathbf{error}, \perp\} \\ eval_{pc}(P) &= \begin{cases} unload_{pc}[V] & \text{if } P \xrightarrow{pc}^* V \\ \mathbf{error} & \text{if } P \xrightarrow{pc}^* \mathbf{error} \\ \perp & \text{if } \forall i \in \mathbf{N} \exists S_i \in State_{pc}, n_i, m_i \in \mathbf{N} \text{ such that} \\ & m_i > 0, P = S_0 \text{ and } S_i \xrightarrow{pc}^{n_i, m_i} S_{i+1} \end{cases} \end{aligned}$$

Data Specifications:

$$\begin{aligned} S \in State_{pc} &::= M \mid \mathbf{error} \mid (\mathbf{f-let} (p S) S) && \text{(States)} \\ M \in \Lambda_{pc} &::= V \mid (\mathbf{let} (x V) M) \mid \dots && \text{(As for } \Lambda_c) \\ V \in Value_{pc} &::= PValue_{pc} \mid p && \text{(Run-time Values)} \\ &PValue_{pc} ::= c \mid x \mid (\lambda x. M) \mid (\mathbf{cons} V V) && \text{(Proper Values)} \\ p \in Ph-Vars &::= \{p_1, p_2, p_3, \dots\} && \text{(Placeholder Variables)} \\ &Ph-Vars \cap Vars = \emptyset \\ F \in FinalState_{pc} &::= V \mid \mathbf{error} && \text{(Final States)} \end{aligned}$$

Unload Function:

$$\begin{aligned} unload_{pc} : Value_{pc}^0 &\longrightarrow Answers \\ unload_{pc}[W] &= unload_c[W] \end{aligned}$$

Placeholder Substitution $S[p \leftarrow V]$:

$$\begin{aligned} M[p \leftarrow V] &= M \text{ with all free occurrences of } p \text{ replaced by } V \\ \mathbf{error}[p \leftarrow V] &= \mathbf{error} \\ (\mathbf{f-let} (p' S_1) S_2)[p \leftarrow V] &= \begin{cases} (\mathbf{f-let} (p' S_1[p \leftarrow V]) S_2) & \text{if } p = p' \\ (\mathbf{f-let} (p' S_1[p \leftarrow V]) S_2[p \leftarrow V]) & \text{if } p \neq p' \end{cases} \end{aligned}$$

Free Placeholder Function FP :

$$\begin{aligned} FP : State_{pc} &\longrightarrow P(Ph-Vars) \\ FP(M) &= \{p \mid p \text{ occurs in } M\} \\ FP(\mathbf{error}) &= \emptyset \\ FP((\mathbf{f-let} (p S_1) S_2)) &= FP(S_1) \cup (FP(S_2) \setminus \{p\}) \end{aligned}$$

FIGURE 3: The parallel $P(C)$ -machine: Evaluator and Data Specifications

3.1 Specification of the $P(C)$ -machine

State Space The state space of the $P(C)$ -machine is defined in Figure 3. The set of $P(C)$ values includes the values of the sequential C -machine (constants, variables, closures

Transition Rules:

$$\begin{array}{ll}
\mathcal{E}[(\mathbf{let} (x V) M)] \mapsto_{pc}^{1,1} \mathcal{E}[M[x \leftarrow V]] & (\mathit{bind}) \\
\mathcal{E}[(\mathbf{let} (x (\mathbf{future} V)) M)] \mapsto_{pc}^{1,1} \mathcal{E}[M[x \leftarrow V]] & (\mathit{future-id}) \\
\mathcal{E}[(\mathbf{let} (x (\mathbf{car} V)) M)] \mapsto_{pc}^{1,1} \begin{cases} \mathcal{E}[M[x \leftarrow V_1]] & \text{if } V = (\mathbf{cons} V_1 V_2) \\ \mathbf{error} & \text{if } V \neq (\mathbf{cons} V_1 V_2), V \neq p \end{cases} & (\mathit{car}) \\
\mathcal{E}[(\mathbf{let} (x (\mathbf{cdr} V)) M)] \mapsto_{pc}^{1,1} \begin{cases} \mathcal{E}[M[x \leftarrow V_2]] & \text{if } V = (\mathbf{cons} V_1 V_2) \\ \mathbf{error} & \text{if } V \neq (\mathbf{cons} V_1 V_2), V \neq p \end{cases} & (\mathit{cdr}) \\
\mathcal{E}[(\mathbf{let} (x (\mathbf{if} V M_1 M_2)) M)] \mapsto_{pc}^{1,1} \begin{cases} \mathcal{E}[(\mathbf{let} (x M_1) M)] & \text{if } V \neq \mathbf{nil}, V \neq p \\ \mathcal{E}[(\mathbf{let} (x M_2) M)] & \text{if } V = \mathbf{nil} \end{cases} & (\mathit{if}) \\
\mathcal{E}[(\mathbf{let} (x (\mathbf{apply} V_1 V_2)) M)] \mapsto_{pc}^{1,1} \begin{cases} \mathcal{E}[(\mathbf{let} (x N[y \leftarrow V_2]) M)] & \text{if } V_1 = (\lambda y. N) \\ \mathbf{error} & \text{if } V_1 \neq (\lambda y. N), V_1 \neq p \end{cases} & (\mathit{apply}) \\
\mathcal{E}[(\mathbf{let} (x (\mathbf{future} N)) M)] \mapsto_{pc}^{1,0} (\mathbf{f-let} (p N) \mathcal{E}[M[x \leftarrow p]]) & (\mathit{fork}) \\
& p \notin FP(\mathcal{E}) \cup FP(M) \\
(\mathbf{f-let} (p V) S) \mapsto_{pc}^{1,1} S[p \leftarrow V] & (\mathit{join}) \\
(\mathbf{f-let} (p \mathbf{error}) S) \mapsto_{pc}^{1,1} \mathbf{error} & (\mathit{join-error}) \\
(\mathbf{f-let} (p_2 (\mathbf{f-let} (p_1 S_1) S_2)) S_2) \mapsto_{pc}^{1,1} (\mathbf{f-let} (p_1 S_1) (\mathbf{f-let} (p_2 S_2) S_3)) & (\mathit{lift}) \\
& p_1 \notin FP(S_3) \\
(\mathbf{f-let} (p S_1) S_2) \mapsto_{pc}^{a+c, b} (\mathbf{f-let} (p S'_1) S'_2) & (\mathit{parallel}) \\
& \text{if } S_1 \mapsto_{pc}^{a, b} S'_1, S_2 \mapsto_{pc}^{c, d} S'_2 \\
S \mapsto_{pc}^{0, 0} S & (\mathit{reflexive}) \\
S \mapsto_{pc}^{a+c, b+d} S'' & (\mathit{transitive}) \\
& \text{if } S \mapsto_{pc}^{a, b} S', S' \mapsto_{pc}^{c, d} S'' \text{ and } a, c > 0
\end{array}$$

FIGURE 4: The parallel $P(C)$ -machine: Transition Rules

and pairs), which we refer to as *proper* values. To model the implementation of futures, the $P(C)$ -machine also includes a new class of values called *placeholder variables*. A placeholder variable p represents the result of a computation that is in progress. Once the computation terminates, all occurrences of the placeholder are replaced by the value returned by the computation.

Each state of the C -machine represents a single thread of control or *task*. To model the concurrent evaluation of **future** expressions by parallel threads, the $P(C)$ -machine includes additional states of the form $(\mathbf{f}\text{-let } (p \ S_1) \ S_2)$. The *primary* sub-state S_1 is initially the body of the **future** expression, and the *secondary* sub-state S_2 is initially the context surrounding the **future** expression. The placeholder p represents the result of S_1 in S_2 . The usual conventions for binding constructs like λ and **let** apply to **f-let**. We use $S[p \leftarrow V]$ and $\mathcal{E}[p \leftarrow W]$ to denote the capture-free substitution of V for all free occurrences of p within a state S and an evaluation context \mathcal{E} respectively. A state is *closed* if it contains no free variables or free placeholders. The evaluation of S_1 is considered *mandatory*, since it is guaranteed to contribute to the completion of the computation. The evaluation of S_2 is *speculative*, since such work may not be required for the termination of the program. In particular, if S_1 raises an error signal, then the evaluator discards the state S_2 , and any effort invested in the evaluation of S_2 is wasted. The distinction between mandatory and speculative steps is crucial for ensuring a sound definition of an evaluator and is incorporated into the definition of the transition relation.

Transition Rules The transition relation of the $P(C)$ -machine is specified as a quadruple: see Figure 4. If $S \xrightarrow[p_c]{n,m} S'$ holds, then the index n is the number of steps involved in the transition from S to S' , and the index $m \leq n$ is the number of these steps that are *mandatory*.

The transition relation is formulated as a collection of transition rules. The rules (*bind*), (*future-id*), (*car*), (*cdr*), (*if*) and (*apply*) are simply the transition rules of the C -machine, appropriately modified to allow for undetermined placeholders. An application of one of these rules counts as a mandatory step.

The transition rule (*fork*) initiates parallel evaluation. This rule may be applied whenever the current term includes a **future** expression within an evaluation context, *i.e.*,

$$\mathcal{E}[(\mathbf{let} \ (x \ (\mathbf{future} \ N)) \ M)]$$

The **future** annotation allows the expression N to be evaluated in parallel with the enclosing context $\mathcal{E}[(\mathbf{let} \ (x \ [\]) \ M)]$. The machine creates a new placeholder p to represent the result of N , and initiates parallel evaluation of N and $\mathcal{E}[(\mathbf{let} \ (x \ p) \ M)]$.

The transition rule (*parallel*) permits concurrent evaluation of both sub-states of a parallel state $(\mathbf{f}\text{-let } (p \ S_1) \ S_2)$.

The transition rules (*join*) and (*join-error*) merge distinct threads of evaluation. When the primary sub-state S_1 of a parallel state $(\mathbf{f}\text{-let } (p \ S_1) \ S_2)$ returns a value V , then the rule (*join*) replaces all occurrences of the placeholder p within S_2 by that value. If the primary sub-state S_1 evaluates to **error**, then the rule (*join-error*) discards the secondary sub-state S_2 and returns **error** as the result of the parallel state.

The transition rule (*lift*) restructures nested parallel states, and thus exposes additional parallelism in certain cases. Consider the state $(\mathbf{f}\text{-let } (p_2 \ (\mathbf{f}\text{-let } (p_1 \ S_1) \ V)) \ S_3)$. The rule (*lift*) allows the value V to be returned to the sub-state S_3 (via a subsequent (*join*) transition), without having to wait on the termination of S_1 .²

²The addition of this optimization rule is the result of proving the soundness of the semantics via the standard diamond lemma. Without this rule, the machine does not possess the diamond property (but is probably still consistent).

The rules (*reflexive*) and (*transitive*) close the relation under reflexivity and transitivity.

We write $S \mapsto_{pc}^* S'$ if $S \mapsto_{pc}^{n,m} S'$ for some $n, m \in \mathbf{N}$. A state S is in *normal form* if there is no state S' such that $S \mapsto_{pc}^{n,m} S'$ with $n > 0$. A state is a *final state* if it is either a value, or the state **error**, and a state is *blocked* if it is in normal form but not a final state.

Indeterminism Unlike the *functional C-machine*, which maps each state to a unique successor state, the transition *relation* of the $P(C)$ -machine has an important degree of freedom. The definition of the $P(C)$ -machine does not specify when the transition rule (*fork*) applies. For example, consider the state $\mathcal{E}[(\text{let } (x (\text{future } N)) M)]$. Since the context surrounding N is an evaluation context, an implementation of the machine may proceed either by evaluating N , or by creating a new task via a (*fork*) transition. The choice of whether or not to apply the transition rule (*fork*) is entirely up to the implementation of the machine. An implementation may *immediately* apply this rule whenever a **future** expression is encountered, realizing a task creation strategy called *eager task creation* [16, 23, 10]. Alternatively, an implementation may *never* invoke the rule (*fork*), resulting in a purely sequential evaluation. In between these two extremes lie a range of strategies where new tasks are created according to some implementation-dependent and possibly load-dependent algorithm. A particularly efficient strategy is *lazy task creation* [4, 19], where new tasks are created via *fork* transitions only when the additional parallelism can exploit idle computing resources.

A second source of indeterminism in the specification of the $P(C)$ -machine is the transition rule (*parallel*). This rule does not specify the number of steps that parallel sub-states must perform before they synchronize. An implementation of the machine can use almost any scheduling strategy for allocating processors to tasks. The only constraint, as specified in the definition of $eval_{pc}$, is that the implementation must perform mandatory computation steps on a regular basis.

Evaluation In general, the evaluation of a program can proceed via many different transition sequences. Some of these transition sequences may be infinite, even if the program terminates according to the sequential semantics. Consider:

$$P = (\text{let } (x (\text{future } E)) \Omega)$$

where E is a term that raises a run-time error, and Ω is some diverging sequential term such that $\Omega \mapsto_{pc}^{1,1} \Omega_1 \mapsto_{pc}^{1,1} \Omega_2 \mapsto_{pc}^{1,1} \dots$. The sequential evaluator never executes Ω because P 's result is **error**. In contrast, P admits the following infinite parallel transition sequence:

$$\begin{array}{ll} P \mapsto_{pc}^{1,0} (\text{f-let } (p E) \Omega) & \text{via } (\text{fork}) \\ \mapsto_{pc}^{1,0} (\text{f-let } (p E) \Omega_1) & \text{since } \Omega \mapsto_{pc}^{1,1} \Omega_1 \\ \mapsto_{pc}^{1,0} (\text{f-let } (p E) \Omega_2) & \text{since } \Omega_1 \mapsto_{pc}^{1,1} \Omega_2 \\ \mapsto_{pc}^{1,0} \dots & \end{array}$$

This “evaluation” diverges because it exclusively consists of speculative transition steps and does not include any mandatory transition steps that contribute to the sequential evaluation of the program.

The evaluator for the $P(C)$ -machine excludes these *excessively speculative* transition sequences, and only admits transition sequences that regularly includes mandatory transition steps.³ For a terminating transition sequence, the number of speculative steps performed is implicitly bounded. For non-terminating sequences, the definition of the evaluator explicitly requires that mandatory transition steps are performed on a regular basis. This constraint implies that an implementation of the machine must keep track of the mandatory thread and must ensure that this mandatory thread is regularly executed.

In summary, the $P(C)$ -machine arbitrarily chooses any transition sequence that regularly performs mandatory computation, and reports on the behavior of that sequence. If the chosen transition sequence produces either a value V or **error**, then $eval_{pc}$ returns $unload_{pc}[V]$ or **error** respectively. If the chosen transition sequence does not terminate, then $eval_{pc}$ returns \perp . As we will prove below, the evaluator relation $eval_{pc}$ is a total function and agrees with the sequential evaluator $eval_c$.

Placeholder Transparency and Synchronization We say that a program operation is *placeholder-strict* in a position if it needs specific information about the value of the corresponding argument. For example, the operations **car** and **cdr** are placeholder-strict, and **if** and **apply** are placeholder-strict in their first position. Whenever an undetermined placeholder appears in a placeholder-strict argument position of one of these operations, then that operation must *block* until the placeholder is determined and specific information about the value of the argument is known. We model this behavior in the $P(C)$ -machine via side-conditions associated with the transition rules (*car*), (*cdr*), (*if*) and (*apply*). These side-conditions ensure that if a placeholder-strict argument is an undetermined placeholder, then the transition rule cannot fire.

For a brief illustration of this idea, consider the following transition sequence for the following program P :

$$P = (\text{let } (y \text{ (future (let } (x \text{ 1) (cons } x \text{ x)))) } (\text{let } (z \text{ (car } y)) z)) \quad (1)$$

$$\xrightarrow{1,0}_{pc} (\text{f-let } (p \text{ (let } (x \text{ 1) (cons } x \text{ x)))) } (\text{let } (z \text{ (car } p)) z)) \quad (2)$$

$$\xrightarrow{1,1}_{pc} (\text{f-let } (p \text{ (cons 1 1)}) (\text{let } (z \text{ (car } p)) z)) \quad (3)$$

$$\xrightarrow{1,1}_{pc} (\text{let } (z \text{ (car (cons 1 1)))) z) \quad (4)$$

$$\xrightarrow{1,1}_{pc} 1 \quad (5)$$

The first transition in this sequence creates a new task for the evaluation of the **future** expression via a (*fork*) transition. After task creation (line 2), no transition steps are possible from the secondary sub-state $(\text{let } (z \text{ (car } p)) z)$. The transition rule (*car*) cannot fire since the argument to the operation (**car**) is a placeholder variable. Evaluation of the primary substate (highlighted with a box) proceeds unhindered. Once the primary substate produces a value (line 3), the transition rule (*join*) synchronizes the separate threads of computation by replacing all occurrences of p by that value. After synchronization (line 4),

³The concept of a mandatory step is closely related to the notion of legitimacy introduced by Katz and Weise [15].

the operation **car** applies to the new argument (**cons** 1 1), and execution continues with the program returning the answer 1. Since program operations block whenever an argument in a placeholder-strict position is undetermined, the $P(C)$ -machine never performs a transition *before* a placeholder is determined that it would perform differently *after* the placeholder is determined. Hence the transition relation of the machine exhibits a *substitutivity* property: the transition relation commutes with substitution of values for placeholders. The proof of this property relies on a Substitution Lemma for placeholders.

Lemma 3.1 (Substitution Lemma for $\cdot [\cdot \leftarrow \cdot]$) *If $M \in \Lambda_{pc}$ and $V, W \in \text{Value}_{pc}$, then*

$$(M[p \leftarrow W])[x \leftarrow V[p \leftarrow W]] = (M[x \leftarrow V])[p \leftarrow W].$$

Proof: The proof proceeds by induction on the structure of M . ■

Lemma 3.2 (Substitutivity; Placeholder Transparency) *If $S_1 \xrightarrow{pc}^{n,m} S_2$, then for any placeholder p and any $W \in \text{Value}_{pc}$, $S_1[p \leftarrow W] \xrightarrow{pc}^{n,m} S_2[p \leftarrow W]$.*

Proof: The proof is by lexicographic induction on n and on the size of S_1 , and proceeds by case analysis on the last step in the transition $S_1 \xrightarrow{pc}^{n,m} S_2$.

- Suppose $S_1 \xrightarrow{pc}^{1,1} S_2$ via the rule (*car*) with $S_1 = \mathcal{E}[(\text{let } (x (\text{car } V)) M)]$. Let $\mathcal{E}' = \mathcal{E}[p \leftarrow W]$, $V' = V[p \leftarrow W]$ and $M' = M[p \leftarrow W]$. We consider three possibilities for V :
 - Suppose $V = (\text{cons } V_1 V_2)$. Then $S_2 = \mathcal{E}[M[x \leftarrow V_1]]$. Let $V'_1 = V_1[p \leftarrow W]$ and $V'_2 = V_2[p \leftarrow W]$.

$$\begin{aligned} S_1[p \leftarrow W] &= \mathcal{E}'[(\text{let } (x (\text{car } (\text{cons } V'_1 V'_2))) M')] \\ &\xrightarrow{pc}^{1,1} \mathcal{E}'[M'[x \leftarrow V'_1]] \\ &= S_2[p \leftarrow W] \quad \text{by Lemma 3.1} \end{aligned}$$

- Suppose V is neither a pair nor a placeholder. Then $S_2 = \text{error}$, and since V' is neither a pair nor a placeholder,

$$\begin{aligned} S_1[p \leftarrow W] &= \mathcal{E}'[(\text{let } (x (\text{car } V')) M')] \\ &\xrightarrow{pc}^{1,1} \text{error} \\ &= S_2[p \leftarrow W] \end{aligned}$$

- The case where V is a placeholder is impossible, since no transitions are possible from such a state.

- The reasoning for the transition rules (*bind*), (*future-id*), (*cdr*), (*if*), (*apply*), (*fork*), (*join*), (*join-error*), (*lift*), (*reflexive*) and (*transitive*) is similar.
- We consider the case where $S_1 \xrightarrow{pc}^{n,m} S_2$ via the rule (*parallel*) in order to justify our use of lexicographic induction. For this case:

$$\begin{array}{ll} S_1 &= (\text{f-let } (p' S'_1) S''_1) & S'_1 &\xrightarrow{pc}^{a,m} S'_2 \\ S_2 &= (\text{f-let } (p' S'_2) S''_2) & S''_1 &\xrightarrow{pc}^{c,d} S''_2 \end{array}$$

and $n = a + c$. Since S'_1 and S''_1 are substates of S_1 , by the inductive hypothesis we have:

$$\begin{aligned} S'_1[p \leftarrow W] &\mapsto_{pc}^{a,m} S'_2[p \leftarrow W] \\ S''_1[p \leftarrow W] &\mapsto_{pc}^{c,d} S''_2[p \leftarrow W] \end{aligned}$$

Hence:

$$\begin{aligned} S_1[p \leftarrow W] &= (\mathbf{f}\text{-let } (p' S'_1[p \leftarrow W]) S''_1[p \leftarrow W]) \\ &\mapsto_{pc}^{n,m} (\mathbf{f}\text{-let } (p' S'_2[p \leftarrow W]) S''_2[p \leftarrow W]) \quad \text{via } (\textit{parallel}) \\ &= S_2[p \leftarrow W] \end{aligned}$$

■

The Substitutivity Lemma shows that undetermined placeholders represent results of parallel computations in a transparent manner. This property is crucial in proving the correctness of the machine.

3.2 Consistency of the $P(C)$ -Machine

The *observable* behavior of the $P(C)$ -machine on a given program is deterministic, despite its indeterminate *internal* behavior. We prove this consistency in the traditional manner, using a modified form of the Diamond Lemma. The proof of the Diamond Lemma relies on the following two obvious properties of the relation \mapsto_{pc} .

Lemma 3.3 (1) If $S \mapsto_{pc}^{n,m} S'$, then $n \geq m$. (2) If $S \mapsto_{pc}^{0,0} S'$, then $S = S'$.

Proof: By induction of n , and case analysis of $S \mapsto_{pc}^{n,m} S'$. ■

The Modified Diamond Lemma states that if we reduce an initial state S_1 by two alternative transitions, producing respectively states S_2 and S_3 , then there is some state S_4 that is reachable from both S_2 and S_3 . Furthermore, the number of *mandatory* steps on the transition from S_1 to S_4 via S_2 is bounded by twice the *total* number of steps on the transition from S_1 to S_4 via S_3 , and vice-versa. This bound implies that all transition sequences for a given program exhibit the same termination behavior.

Lemma 3.4 (Modified Diamond Lemma) Let $S_1, S_2, S_3 \in \text{State}_{pc}$. If $S_1 \mapsto_{pc}^{n_1, m_1} S_2$ and $S_1 \mapsto_{pc}^{n_2, m_2} S_3$, then there exists $S_4 \in \text{State}_{pc}$ and $n_3, m_3, n_4, m_4 \in \mathbf{N}$ such that $S_2 \mapsto_{pc}^{n_3, m_3} S_4$ and $S_3 \mapsto_{pc}^{n_4, m_4} S_4$. Furthermore $m_1 + m_3 \leq 2(n_2 + n_4)$ and $m_2 + m_4 \leq 2(n_1 + n_3)$.

Proof: The proof is by lexicographic induction on n_1, n_2 and the size of S_1 . We define the transition relation (*seq*) to be the union of the transition relations (*bind*), (*car*), (*cdr*), (*if*) and (*apply*). Then the transition relation of the $P(C)$ -machine is the union of the eight relations (*seq*), (*future-id*), (*fork*), (*join*), (*join-error*), (*lift*), (*parallel*), (*reflexive*) and (*transitive*). The proof proceeds by case analysis of the pair of transition rules used for $S_1 \mapsto_{pc}^{n_1} S_2$ and $S_1 \mapsto_{pc}^{n_2} S_3$.

The following table enumerates the possible combinations after symmetry considerations, and annotates each case with a reference to the argument used to prove that case.

	<i>(seq)</i>	<i>(f-id)</i>	<i>(fork)</i>	<i>(join)</i>	<i>(join-error)</i>	<i>(lift)</i>	<i>(parallel)</i>	<i>(refl.)</i>	<i>(trans.)</i>
<i>(seq)</i>	\equiv	\times	3	\times	\times	\times	\times	1	2
<i>(future-id)</i>		\equiv	4	\times	\times	\times	\times	1	2
<i>(fork)</i>			5	\times	\times	\times	\times	1	2
<i>join</i>				\equiv	\times	\times	6	1	2
<i>(join-error)</i>					\equiv	\times	7	1	2
<i>(lift)</i>						\equiv	9	1	2
<i>(parallel)</i>							8	1	2
<i>(reflexive)</i>								1	1
<i>(transitive)</i>									2

The following arguments show that the lemma holds in each of the above cases.

- (\times) The cases marked by the symbol \times are impossible, since the domains of the respective relations are disjoint.
- (\equiv) The cases marked by the symbol \equiv hold, since the relation in question is a partial function, and hence $S_2 = S_3$.
- (1) For the case where $S_1 \xrightarrow[p_c]{0,0} S_2$ via the rule *reflexive*, take $S_4 = S_2$.
- (2) For the case where $S_1 \xrightarrow[p_c]{n_1, m_1} S_2$ via the rule *transitive*, we have that $S_1 \xrightarrow[p_c]{a_1, b_1} S_5 \xrightarrow[p_c]{a_2, b_2} S_2$. The following diagram outlines our proof technique for this case.

$$\begin{array}{ccccc}
 S_1 & \xrightarrow[p_c]{a_1, b_1} & S_5 & \xrightarrow[p_c]{a_2, b_2} & S_2 \\
 \downarrow n_2, m_2 & & \vdots a_3, b_3 & & \vdots n_3, m_3 \\
 S_3 & \dots a_4, b_4 & S_6 & \dots a_5, b_5 & S_4
 \end{array}$$

Since $a_1 < n_1$, by the inductive hypothesis there exists $S_6 \in \text{State}_{p_c}$ and $a_3, b_3, a_4, b_4 \in \mathbb{N}$ such that

$$\begin{array}{ll}
 S_5 \xrightarrow[p_c]{a_3, b_3} S_6 & b_1 + b_3 \leq 2(n_2 + a_4) \\
 S_3 \xrightarrow[p_c]{a_4, b_4} S_6 & m_2 + b_4 \leq 2(a_1 + a_3)
 \end{array}$$

Since $S_5 \xrightarrow[p_c]{a_2, b_2} S_2$, $S_5 \xrightarrow[p_c]{a_3, b_3} S_6$ and $a_2 < n_1$, by the inductive hypothesis there exists $S_4 \in \text{State}_{p_c}$ and $n_3, m_3, a_5, b_5 \in \mathbb{N}$ such that

$$\begin{array}{ll}
 S_2 \xrightarrow[p_c]{n_3, m_3} S_4 & b_2 + m_3 \leq 2(a_3 + a_5) \\
 S_6 \xrightarrow[p_c]{a_5, b_5} S_4 & b_3 + b_5 \leq 2(a_2 + n_3)
 \end{array}$$

Let $n_4 = a_4 + a_5$ and $m_4 = b_4 + b_5$, and we have that $S_2 \xrightarrow{pc}^{n_3, m_3} S_4$ and $S_3 \xrightarrow{pc}^{n_4, m_4} S_4$. Furthermore,

$$\begin{aligned}
 m_1 + m_3 &= b_1 + b_2 + m_3 \\
 &= (b_1 + b_3) + (b_2 + m_3) - b_3 \\
 &\leq 2(n_2 + a_4) + 2(a_3 + a_5) - b_3 \quad (I.H.) \\
 &= 2 \times n_2 + 2(a_4 + a_5) + (2 \times a_3 - b_3) \\
 &\leq 2(n_2 + n_4)
 \end{aligned}$$

Similarly, $m_2 + m_4 \leq 2(n_1 + n_3)$.

- (3) Suppose $S_1 \xrightarrow{pc}^{1,1} S_2$ via the rule (*seq*), and $S_1 \xrightarrow{pc}^{1,0} S_3$ via the rule (*fork*). Then $n_1 = m_1 = n_2 = 1$, $m_2 = 0$ and

$$\begin{aligned}
 S_1 &= \mathcal{E}[(\mathbf{let} (x (\mathbf{future} \mathcal{E}'[N])) M)] \\
 S_3 &= (\mathbf{f-let} (p \mathcal{E}'[N]) \mathcal{E}[M[x \leftarrow p]])
 \end{aligned}$$

We consider two cases for $S_1 \xrightarrow{pc}^{1,1} S_2$ separately.

- Suppose $\mathcal{E}'[N] \xrightarrow{pc}^{1,1} \mathcal{E}'[N']$, and $S_2 = (\mathbf{let} (x (\mathbf{future} \mathcal{E}'[N'])) M)$. Set $S_4 = (\mathbf{f-let} (p \mathcal{E}'[N']) \mathcal{E}[M[x \leftarrow p]])$. Clearly $S_2 \xrightarrow{pc}^{1,0} S_4$ via (*fork*), $S_3 \xrightarrow{pc}^{1,1} S_4$ via (*seq*), and the indices satisfy the inequalities.
- Alternatively, suppose $S_2 = \mathbf{error}$. Then let $S_4 = \mathbf{error}$. Observe that

$$\begin{aligned}
 S_3 &\xrightarrow{pc}^{1,1} (\mathbf{f-let} (p \mathbf{error}) \mathcal{E}[M[x \leftarrow p]]) \quad \text{via } (parallel) \text{ and } (seq) \\
 &\xrightarrow{pc}^{1,1} \mathbf{error} \quad \text{via } (join-error)
 \end{aligned}$$

Also $n_3 = m_3 = 0$ and $m_4 = m_4 = 2$. Hence $m_2 + m_4 = 2$ and $n_1 + n_3 = 1$. This case is the reason for the potential factor of two difference between $m_2 + m_4$ and $n_1 + n_3$.

- (4) Suppose $S_1 \xrightarrow{pc}^{1,1} S_2$ via the rule (*future-id*), and $S_1 \xrightarrow{pc}^{1,0} S_3$ via the rule (*fork*). There are two sub-cases to consider. Suppose the two transitions operate on the same **future** expression. Then,

$$\begin{aligned}
 S_1 &= \mathcal{E}[(\mathbf{let} (x (\mathbf{future} V)) M)] \\
 S_2 &= \mathcal{E}[M[x \leftarrow V]] \\
 S_3 &= (\mathbf{f-let} (p V) \mathcal{E}[M[x \leftarrow p]])
 \end{aligned}$$

For this case, choose $S_4 = S_2$:

$$\begin{aligned}
 S_2 &\xrightarrow{pc}^{0,0} S_4 \\
 S_3 &\xrightarrow{pc}^{1,1} (\mathcal{E}[M[x \leftarrow p]])[p \leftarrow V] \\
 &= \mathcal{E}[M[x \leftarrow V]] = S_4 \\
 &\quad \text{since } p \notin FP(\mathcal{E}) \cup FP(M)
 \end{aligned}$$

The indices clearly satisfy the inequalities.

Alternatively, suppose that the transitions operate on different **future** expressions. Then

$$\begin{aligned} S_1 &= \mathcal{E}[(\text{let } (x (\text{future } \mathcal{E}'[(\text{let } (y (\text{future } V)) N]) M))] \\ S_2 &= \mathcal{E}[(\text{let } (x (\text{future } \mathcal{E}'[N[y \leftarrow V]]) M)] \\ S_3 &= (\mathbf{f}\text{-let } (p \mathcal{E}'[(\text{let } (y (\text{future } V)) N]) \\ &\quad \mathcal{E}[M[x \leftarrow p]]) \\ &\quad \text{where } p \notin FP(\mathcal{E}) \cup FP(M) \end{aligned}$$

Let

$$S_4 = (\mathbf{f}\text{-let } (p \mathcal{E}'[N[y \leftarrow V]]) \mathcal{E}[M[x \leftarrow p]])$$

and we have

$$S_2 \xrightarrow{1,0}_{pc} S_4 \quad \text{via } (fork)$$

Also, since $\mathcal{E}'[(\text{let } (y (\text{future } V)) N)] \xrightarrow{1,1}_{pc} \mathcal{E}'[N[y \leftarrow V]]$ via *(future-id)*, we have (by using the rule *(parallel)* for compatible closure)

$$S_3 \xrightarrow{1,1}_{pc} S_4 \quad \text{via } (parallel)$$

The indices clearly satisfy the inequalities.

- (5) Suppose $S_1 \xrightarrow{1,1}_{pc} S_2$ via the rule *(fork)*, and $S_1 \xrightarrow{1,0}_{pc} S_3$ by the rule *(fork)*. If the two transition operate on the same **future** expression, then $S_2 = S_3$. Otherwise,

$$\begin{aligned} S_1 &= \mathcal{E}[(\text{let } (x (\text{future } \mathcal{E}'[(\text{let } (y (\text{future } N')) N]) M))] \\ S_1 \xrightarrow{1,1}_{pc} S_2 &= (\mathbf{f}\text{-let } (p_1 \mathcal{E}'[(\text{let } (y (\text{future } N')) N]) \\ &\quad \mathcal{E}[M[x \leftarrow p_1]]) \\ S_1 \xrightarrow{1,1}_{pc} S_3 &= (\mathbf{f}\text{-let } (p_2 N') \\ &\quad \mathcal{E}[(\text{let } (x (\text{future } \mathcal{E}'[N[y \leftarrow p_2]]) M)]) \end{aligned}$$

where $p_1 \notin FP(\mathcal{E}) \cup FP(M)$ and $p_2 \notin FP(\mathcal{E}) \cup FP(\mathcal{E}') \cup FP(N) \cup FP(M)$.

Let

$$\begin{aligned} S_4 &= (\mathbf{f}\text{-let } (p_2 N') \\ &\quad (\mathbf{f}\text{-let } (p_1 \mathcal{E}'[N[y \leftarrow p_2]]) \mathcal{E}[M[x \leftarrow p_1]])) \end{aligned}$$

Then

$$\begin{aligned} S_2 &\xrightarrow{1,1}_{pc} (\mathbf{f}\text{-let } (p_1 (\mathbf{f}\text{-let } (p_2 N') \mathcal{E}'[N[y \leftarrow p_2]])) \\ &\quad \mathcal{E}'[M[x \leftarrow p_1]]) \quad \text{via } (fork) \\ &\xrightarrow{1,0}_{pc} S_4 \quad \text{via } (lift) \\ S_3 &\xrightarrow{1,1}_{pc} S_4 \quad \text{via } (fork) \end{aligned}$$

and the indices clearly satisfy the inequalities.

- (6) For the case where $S_1 \mapsto_{pc}^{1,1} S_2$ via the rule (*join*), and $S_1 \mapsto_{pc}^{n,0} S_3$ by the rule (*parallel*), we have

$$\begin{aligned} S_1 &= (\mathbf{f}\text{-let } (p \ V) \ S) \\ S_2 &= S[p \leftarrow V] \\ S_3 &= (\mathbf{f}\text{-let } (p \ V) \ S') \quad \text{where } S \mapsto_{pc}^{n,m} S' \end{aligned}$$

Pick $S_4 = S'[p \leftarrow V]$, and $S_2 \mapsto_{pc}^{n,m} S_4$ by the Placeholder Transparency Lemma (3.2), $S_3 \mapsto_{pc}^{1,1} S_4$ via the rule (*join*), and the indices clearly satisfy the inequalities.

- (7) Suppose $S_1 \mapsto_{pc}^{1,1} S_2$ via the rule (*join-error*), and $S_1 \mapsto_{pc}^{n,0} S_3$ by the rule (*parallel*). Then

$$\begin{aligned} S_1 &= (\mathbf{f}\text{-let } (p \ \mathbf{error}) \ S) \\ S_2 &= \mathbf{error} \\ S_3 &= (\mathbf{f}\text{-let } (p \ \mathbf{error}) \ S') \quad \text{where } S \mapsto_{pc}^{n,m} S' \end{aligned}$$

Pick $S_4 = \mathbf{error}$. Then $S_2 \mapsto_{pc}^{0,0} S_4$, $S_3 \mapsto_{pc}^{1,1} S_4$ by the rule (*join-error*), and the indices clearly satisfy the inequalities.

- (8) Suppose both the transitions $S_1 \mapsto_{pc}^{n_1,m_1} S_2$ and $S_1 \mapsto_{pc}^{n_2,m_2} S_3$ are via the rule (*parallel*). Then

$$\begin{aligned} S_1 &= (\mathbf{f}\text{-let } (p \ S'_1) \ S''_1) \\ S_2 &= (\mathbf{f}\text{-let } (p \ S'_2) \ S''_2) \\ &\quad \text{where } S'_1 \mapsto_{pc}^{a_1,m_1} S'_2, S'_1 \mapsto_{pc}^{c_1,d_1} S''_2 \text{ and } n_1 = a_1 + c_1 \\ S_3 &= (\mathbf{f}\text{-let } (p \ S'_3) \ S''_3) \\ &\quad \text{where } S'_1 \mapsto_{pc}^{a_2,m_2} S'_3, S'_1 \mapsto_{pc}^{c_2,d_2} S''_3 \text{ and } n_2 = a_2 + c_2 \end{aligned}$$

Since $a_1 \leq n_1$, $a_2 \leq n_2$ and S'_1 is strictly smaller than S_1 , by the inductive hypothesis there exists S'_4 such that

$$\begin{array}{lll} S'_2 & \xrightarrow{a_3,b_3}_{pc} & S'_4 \\ S'_3 & \xrightarrow{a_4,b_4}_{pc} & S'_4 \end{array} \quad \begin{array}{ll} m_1 + b_3 & \leq 2(a_2 + a_4) \\ m_2 + b_4 & \leq 2(a_1 + a_3) \end{array}$$

Similarly, there exists S''_4 such that

$$\begin{array}{lll} S''_2 & \xrightarrow{c_3,d_3}_{pc} & S''_4 \\ S''_3 & \xrightarrow{c_4,d_4}_{pc} & S''_4 \end{array} \quad \begin{array}{ll} d_1 + d_3 & \leq 2(c_2 + c_4) \\ d_2 + d_4 & \leq 2(c_1 + c_3) \end{array}$$

Letting $S_4 = (\mathbf{f}\text{-let } (p \ S'_4) \ S''_4)$, we have

$$\begin{aligned} S_2 &\xrightarrow{pc}^{n_3,m_3} S_4 \\ S_3 &\xrightarrow{pc}^{n_4,m_4} S_4 \end{aligned}$$

where $n_3 = a_3 + c_3$, $m_3 = b_3$, $n_4 = a_4 + c_4$, $m_4 = b_4$. Furthermore,

$$\begin{aligned} m_1 + m_3 &= m_1 + b_3 \leq 2(a_2 + a_4) \leq 2(n_2 + n_4) \\ m_2 + m_4 &= m_2 + b_4 \leq 2(a_1 + a_3) \leq 2(n_1 + n_3). \end{aligned}$$

- (9) Suppose $S_1 \xrightarrow{1,0}_{pc} S_2$ via the rule (*lift*), and $S_1 \xrightarrow{n_2, m_2}_{pc} S_3$ by the rule (*parallel*). Then

$$\begin{aligned} S_1 &= (\mathbf{f}\text{-let } (p_2 (\mathbf{f}\text{-let } (p_1 S_a) S_b)) S_c) \\ S_1 \xrightarrow{1,0}_{pc} S_2 &= (\mathbf{f}\text{-let } (p_1 S_a) (\mathbf{f}\text{-let } (p_2 S_b) S_c)) \quad p_1 \notin FV(S_c) \\ S_1 \xrightarrow{n_2, m_2}_{pc} S_3 &= (\mathbf{f}\text{-let } (p_2 S'_3) S''_3) \\ &\quad \text{where } (\mathbf{f}\text{-let } (p_1 S_a) S_b) \xrightarrow{a, m_2}_{pc} S'_3 \\ &\quad \text{and } S_c \xrightarrow{c, d}_{pc} S''_3, n_2 = a + c \end{aligned}$$

We proceed by sub-case analysis of the last step in $(\mathbf{f}\text{-let } (p_1 S_a) S_b) \xrightarrow{a, m_2}_{pc} S'_3$.

- There is no S'_3 such that $(\mathbf{f}\text{-let } (p_1 S_a) S_b) \xrightarrow{a, m_2}_{pc} S'_3$ via any of the rules (*bind*), (*future-id*), (*car*), (*cdr*), (*if*), (*apply*) or (*fork*).
- Suppose $(\mathbf{f}\text{-let } (p_1 S_a) S_b) \xrightarrow{a, m_2}_{pc} S'_3$ via (*join*). Then

$$\begin{aligned} S_1 &= (\mathbf{f}\text{-let } (p_2 (\mathbf{f}\text{-let } (p_1 V) S_b)) S_c) \\ S_1 \xrightarrow{1,0}_{pc} S_2 &= (\mathbf{f}\text{-let } (p_1 V) (\mathbf{f}\text{-let } (p_2 S_b) S_c)) \quad p_1 \notin FV(S_c) \\ S_1 \xrightarrow{n_2, 1}_{pc} S_3 &= (\mathbf{f}\text{-let } (p_2 S_b[p_1 \leftarrow V]) S''_3) \quad n_2 = 1 + c \end{aligned}$$

Let $S_4 = S_3$:

$$\begin{aligned} S_2 &\xrightarrow{1,1}_{pc} (\mathbf{f}\text{-let } (p_2 S_b[p_1 \leftarrow V]) S_c) && \text{via } (\textit{join}) \\ &\xrightarrow{c, d}_{pc} (\mathbf{f}\text{-let } (p_2 S_b[p_1 \leftarrow V]) S''_3) = S_4 && \text{via } (\textit{parallel}) \end{aligned}$$

The indices clearly satisfy the inequalities.

- Suppose $(\mathbf{f}\text{-let } (p_1 S_a) S_b) \xrightarrow{a, m_2}_{pc} S'_3$ via (*join-error*). Then

$$\begin{aligned} S_1 &= (\mathbf{f}\text{-let } (p_2 (\mathbf{f}\text{-let } (p_1 \text{error}) S_b)) S_c) \\ S_1 \xrightarrow{1,0}_{pc} S_2 &= (\mathbf{f}\text{-let } (p_1 \text{error}) (\mathbf{f}\text{-let } (p_2 S_b) S_c)) \quad p_1 \notin FV(S_c) \\ S_1 \xrightarrow{n_2, 0}_{pc} S_3 &= (\mathbf{f}\text{-let } (p_2 \text{error}) S''_3) \quad n_2 = 1 + c \end{aligned}$$

Let $S_4 = \text{error}$:

$$\begin{aligned} S_2 &\xrightarrow{1,0}_{pc} S_4 \\ S_3 &\xrightarrow{1,0}_{pc} S_4 \end{aligned}$$

The indices clearly satisfy the inequalities.

- The cases where $(\mathbf{f}\text{-let } (p_1 S_a) S_b) \xrightarrow{a, m_2}_{pc} S'_3$ via the rules (*reflexive*), (*transitive*) or (*parallel*) are straightforward.

■

The Modified Diamond Lemma implies that all transition sequences exhibit the same observable behavior.

Lemma 3.5 (Consistency of Transitions) *Let P be a program. If $P \mapsto_{pc}^* S$, where S is in normal form, then*

1. *For all normal form S' such that $P \mapsto_{pc}^* S'$, $S' = S$.*
2. *It is impossible that for all $i \in \mathbb{N}$ there exists $S_i \in \text{State}_{pc}$ and $n_i, m_i \in \mathbb{N}$ such that $m_i > 0$, $P = S_0$ and $S_i \mapsto_{pc}^{n_i, m_i} S_{i+1}$.*

Proof:

1. Suppose $P \mapsto_{pc}^* S$ and $P \mapsto_{pc}^* S'$, where S, S' are in normal form. By the Modified Diamond Lemma (3.4), there exists some S_4 such that $S \mapsto_{pc}^* S_4$ and $S' \mapsto_{pc}^* S_4$. But since S, S' are in normal form, we have that $S = S_4 = S'$.
2. We prove part 2 by contradiction. Assume that $P \mapsto_{pc}^{n, m} S$ where S is in normal form, and that there exists some sequence of states $S_i \in \text{State}_{pc}$ and $n_i, m_i \in \mathbb{N}$ such that $m_i > 0$, $P = S_0$ and $S_i \mapsto_{pc}^{n_i, m_i} S_{i+1}$.
Pick an integer $k > 2n$. Then $P \mapsto_{pc}^{a, b} S_{k+1}$, where $a = \sum_{i=0}^k n_i$ and $b = \sum_{i=0}^k m_i$. By the Modified Diamond Lemma (3.4), $S_{k+1} \mapsto_{pc}^{c, d} S$ for some $c, d \in \mathbb{N}$, with $2n \geq b + d$. But $b + d \geq b = \sum_{i=0}^k m_i \geq k + 1 > k$, since $m_i > 0$, producing the contradiction $2n > k > 2n$.

■

Since all transition sequences for a given program exhibit the same observable behavior, the evaluator $eval_{pc}$ for the $P(C)$ -machine is a well-defined function.

Theorem 3.6 (Consistency of $eval_{pc}$) *The relation $eval_{pc}$ is a function.*

Proof: Follows from Lemma 3.5. ■

3.3 Correctness of the $P(C)$ -Machine

Since each sequential transition rule of the $P(C)$ -machine subsumes the corresponding transition rule of the C -machine, every transition of the C -machine is also a transition of the $P(C)$ -machine.

Lemma 3.7 *Suppose $S \mapsto_c S'$, for $S, S' \in \text{State}_c$. Then $S \mapsto_{pc}^{1,1} S'$.*

The correspondence between transitions of the two machines implies that their respective evaluators are equivalent.

Theorem 3.8 (Correctness of $eval_{pc}$) $eval_c = eval_{pc}$.

Proof: Let P be any program. We proceed by case analysis of the definition of $eval_c$:

- Suppose $eval_c(P) = unload_c[V]$ because

$$P \equiv S_0 \mapsto_c S_1 \mapsto_c \dots \mapsto_c S_n \equiv V$$

Then $P \mapsto_{pc}^{n,n} V$, and hence $eval_{pc}(P) = unload_{pc}[V] = unload_c[V]$.

- Alternatively, suppose $eval_c(P) = \mathbf{error}$ because

$$P \equiv S_0 \mapsto_c S_1 \mapsto_c \dots \mapsto_c S_n \equiv \mathbf{error}$$

Then $P \mapsto_{pc}^{n,n} \mathbf{error}$, and hence $eval_{pc}(P) = \mathbf{error}$.

- Finally, suppose $eval_c(P) = \perp$ via the infinite sequence

$$P \equiv S_0 \mapsto_c S_1 \mapsto_c \dots \mapsto_c S_n \mapsto_c \dots$$

Then we have that

$$P \equiv S_0 \mapsto_{pc}^{1,1} S_1 \mapsto_{pc}^{1,1} \dots \mapsto_{pc}^{1,1} S_n \mapsto_{pc}^{1,1} \dots$$

and hence $eval_{pc}(P) = \perp$.

Hence, for all programs P , $eval_c(P) = eval_{pc}(P)$. ■

The equivalence of the two evaluators implies that $eval_{pc}$ is defined for all programs.

Theorem 3.9 (Completeness of $eval_{pc}$) *The relation $eval_{pc}$ is a total function.*

In summary, the $P(C)$ -machine is a correct implementation of the C -machine in that both define the same semantics for the source language. Hence, the interpretation of **future** as a task creation construct, with implicit task coordination, is entirely consistent with the definitional semantics of **future** as an annotation.

4 Placeholder Object Semantics

The $P(C)$ -machine specifies the parallel execution behavior of programs with **futures** at a fairly high level. It thus hides certain low-level operations that are required in the implementation of **futures**. In particular, implementations typically represent placeholders using *placeholder objects*, and avoid the need for an expensive substitution operation on placeholders (cmp. *join*) by imperative updating these placeholder objects instead. This technique requires *touch* operations within placeholder-strict primitives to dereference placeholder objects whenever necessary. Since we plan to use the semantics of **future** to prove the correctness of an algorithm for removing redundant *touches*, we reformulate the $P(C)$ -machine to expose these placeholder objects and the associated *touch* operations. The result is the $P(C_{ph})$ -machine.

Evaluator:

$$\begin{aligned}
 eval_{pc'} : \Lambda_a^0 &\longrightarrow Answers \cup \{\mathbf{error}, \perp\} \\
 eval_{pc'}(P) &= \begin{cases} unload_{pc'}[V] & \text{if } P \xrightarrow{*}_{pc'} V \\ \mathbf{error} & \text{if } P \xrightarrow{*}_{pc'} \mathbf{error} \\ \perp & \text{if } \forall i \in \mathbf{N} \exists S_i \in State_{pc'}, n_i, m_i \in \mathbf{N} \text{ such that} \\ & m_i > 0, P = S_0 \text{ and } S_i \xrightarrow{n_i, m_i}_{pc'} S_{i+1} \end{cases}
 \end{aligned}$$

Data Specifications:

$$\begin{aligned}
 S \in State_{pc'} &::= M \mid \mathbf{error} \mid (\mathbf{f-let} (p \ S) \ S) && \text{(States)} \\
 M \in \Lambda_{pc'} &::= V \mid (\mathbf{let} (x \ V) \ M) \mid \dots && \text{(As for } \Lambda_c) \\
 V \in Value_{pc'} &::= PValue_{pc'} \mid \langle \mathbf{ph} \ p \ \circ \rangle \mid \langle \mathbf{ph} \ p \ V \rangle && \text{(Run-time Values)} \\
 PValue_{pc'} &::= c \mid x \mid Cl_{pc'} \mid Pair_{pc'} && \text{(Proper Values)} \\
 Cl_{pc'} &::= (\lambda x. M) && \text{(Closures)} \\
 Pair_{pc'} &::= (\mathbf{cons} \ V \ V) && \text{(Pairs)}
 \end{aligned}$$

Auxiliary Functions:

$$\begin{aligned}
 unload_{pc'} : Value_{pc'}^0 &\longrightarrow Answers & touch_{pc'} : Value_{pc'} &\longrightarrow PValue_{pc} \cup \{\circ\} \\
 unload_{pc'}[c] &= c & touch_{pc'}[\langle \mathbf{ph} \ p \ \circ \rangle] &= \circ \\
 unload_{pc'}[(\lambda x. M)] &= \mathbf{procedure} & touch_{pc'}[\langle \mathbf{ph} \ p \ V \rangle] &= touch_{pc'}[V] \\
 unload_{pc'}[(\mathbf{cons} \ V_1 \ V_2)] &= (\mathbf{cons} \ V_1' \ V_2') & touch_{pc'}[V] &= V \text{ if } V \in PValue_{pc'} \\
 &V_i' = unload_{pc'}[V_i] & & \\
 unload_{pc'}[\langle \mathbf{ph} \ p \ V \rangle] &= unload_{pc'}[V] & &
 \end{aligned}$$

Placeholder Substitution $S[p := V]$:

$$\begin{aligned}
 M[p := V] &= M \text{ with all free occurrences of } \langle \mathbf{ph} \ p \ \circ \rangle \text{ replaced by } \langle \mathbf{ph} \ p \ V \rangle \\
 \mathbf{error}[p := V] &= \mathbf{error} \\
 (\mathbf{f-let} (p' \ S_1) \ S_2)[p := V] &= \begin{cases} (\mathbf{f-let} (p' \ S_1[p := V]) \ S_2) & \text{if } p = p' \\ (\mathbf{f-let} (p' \ S_1[p := V]) \ S_2[p := V]) & \text{if } p \neq p' \end{cases}
 \end{aligned}$$

FIGURE 5: The parallel $P(C_{ph})$ -machine: Evaluator and Data Specifications

4.1 Specification of the $P(C_{ph})$ -machine

The state space of the $P(C_{ph})$ -machine is a minor revision of that of the $P(C)$ -machine: see Figure 5. Instead of plain variables, placeholders are now tagged objects. An *undetermined placeholder object* (written $\langle \mathbf{ph} \ p \ \circ \rangle$) is created to represent the result of each parallel task (cmp. (*fork*)). The symbol \circ indicates that the result of the computation is unknown, and p is used to distinguish distinct placeholder objects. When the computation associated with the placeholder object terminates, producing a value V , then the undetermined placeholder object is replaced by the *determined placeholder object* $\langle \mathbf{ph} \ p \ V \rangle$.⁴

The correctness of this technique requires that a determined placeholder object $\langle \mathbf{ph} \ p \ V \rangle$ representing the value V is observationally equivalent to the value V itself. To ensure

⁴The explicit replacement of undetermined placeholders by determined placeholders is still unrealistic, but suffices for our purposes.

Transition Rules:

$$\begin{array}{lll}
\mathcal{E}[(\text{let } (x \ V) \ M)] & \xrightarrow{1,1}_{pc'} & \mathcal{E}[M[x \leftarrow V]] \quad (\text{bind}) \\
\mathcal{E}[(\text{let } (x \ (\text{future } V)) \ M)] & \xrightarrow{1,1}_{pc'} & \mathcal{E}[M[x \leftarrow V]] \quad (\text{future-id}) \\
\mathcal{E}[(\text{let } (x \ (\text{car } V)) \ M)] & \xrightarrow{1,1}_{pc'} & \begin{cases} \mathcal{E}[M[x \leftarrow V_1]] & \text{if } \text{touch}_{pc'}[V] = (\text{cons } V_1 \ V_2) \\ \text{error} & \text{if } \text{touch}_{pc'}[V] \notin \text{Pair}_{pc'} \cup \{\circ\} \end{cases} \quad (\text{car}) \\
\mathcal{E}[(\text{let } (x \ (\text{cdr } V)) \ M)] & \xrightarrow{1,1}_{pc'} & \begin{cases} \mathcal{E}[M[x \leftarrow V_2]] & \text{if } \text{touch}_{pc'}[V] = (\text{cons } V_1 \ V_2) \\ \text{error} & \text{if } \text{touch}_{pc'}[V] \notin \text{Pair}_{pc'} \cup \{\circ\} \end{cases} \quad (\text{cdr}) \\
\mathcal{E}[(\text{let } (x \ (\text{if } V \ M_1 \ M_2)) \ M)] & \xrightarrow{1,1}_{pc'} & \begin{cases} \mathcal{E}[(\text{let } (x \ M_1) \ M)] & \text{if } \text{touch}_{pc'}[V] \notin \{\text{nil}, \circ\} \\ \mathcal{E}[(\text{let } (x \ M_2) \ M)] & \text{if } \text{touch}_{pc'}[V] = \text{nil} \end{cases} \quad (\text{if}) \\
\mathcal{E}[(\text{let } (x \ (\text{apply } V_1 \ V_2)) \ M)] & \xrightarrow{1,1}_{pc'} & \begin{cases} \mathcal{E}[(\text{let } (x \ N[y \leftarrow V_2]) \ M)] & \text{if } \text{touch}_{pc'}[V_1] = (\lambda y. N) \\ \text{error} & \text{if } \text{touch}_{pc'}[V_1] \notin \text{Cl}_{pc'} \cup \{\circ\} \end{cases} \quad (\text{apply}) \\
\mathcal{E}[(\text{let } (x \ (\text{future } N)) \ M)] & \xrightarrow{1,0}_{pc'} & \begin{array}{l} (\mathbf{f}\text{-let } (p \ N) \ \mathcal{E}[M[x \leftarrow (\mathbf{ph} \ p \ \circ)])] \\ p \notin FP(\mathcal{E}) \cup FP(M) \end{array} \quad (\text{fork}) \\
(\mathbf{f}\text{-let } (p \ V) \ S) & \xrightarrow{1,1}_{pc'} & S[p := V] \quad (\text{join}) \\
(\mathbf{f}\text{-let } (p \ \text{error}) \ S) & \xrightarrow{1,1}_{pc'} & \text{error} \quad (\text{join-error}) \\
(\mathbf{f}\text{-let } (p_2 \ (\mathbf{f}\text{-let } (p_1 \ S_1) \ S_2)) \ S_2) & \xrightarrow{1,1}_{pc'} & \begin{array}{l} (\mathbf{f}\text{-let } (p_1 \ S_1) \ (\mathbf{f}\text{-let } (p_2 \ S_2) \ S_3)) \\ p_1 \notin FP(S_3) \end{array} \quad (\text{lift}) \\
(\mathbf{f}\text{-let } (p \ S_1) \ S_2) & \xrightarrow{a+c, b}_{pc'} & \begin{array}{l} (\mathbf{f}\text{-let } (p \ S'_1) \ S'_2) \\ \text{if } S_1 \xrightarrow{a, b}_{pc} S'_1, S_2 \xrightarrow{c, d}_{pc} S'_2 \end{array} \quad (\text{parallel}) \\
S & \xrightarrow{0,0}_{pc'} & S \quad (\text{reflexive}) \\
S & \xrightarrow{a+c, b+d}_{pc'} & S'' \quad (\text{transitive}) \\
& & \text{if } S \xrightarrow{a, b}_{pc} S', S' \xrightarrow{c, d}_{pc} S'', a, c > 0
\end{array}$$

FIGURE 6: The parallel $P(C_{ph})$ -machine: Transition Rules

this behavior, the primitives **car**, **cdr**, **if** and **apply** perform *touch* operations on their placeholder-strict arguments. A *touch* operation behaves as the identity operation on proper values, but when applied to a placeholder object, the *touch* operation dereferences the placeholder object to retrieve the value that it represents.

An occurrence of an undetermined placeholder object $\langle \mathbf{ph} \ p \ \circ \rangle$ is *free* if p is not bound by an enclosing **flet** expression, and a state S is *closed* if it does not contain any free variables or free undetermined placeholder objects.

4.2 Correctness of the $P(C_{ph})$ -machine

The state space of the $P(C_{ph})$ -machine is identical to that of the $P(C)$ -machine, apart from the use of placeholder objects. Therefore, each $P(C_{ph})$ -machine state corresponds directly to a $P(C)$ -machine state. The translation function Θ from $P(C_{ph})$ -states to $P(C)$ -states simply replaces determined placeholder objects by the appropriate value, and replaces undetermined placeholder objects by the corresponding placeholder variable.

Definition 4.1. (Θ)

$$\begin{aligned}
 \Theta : State_{pc'} &\longrightarrow State_{pc} \\
 \Theta[\langle \mathbf{ph} \ p \ \circ \rangle] &= p \\
 \Theta[\langle \mathbf{ph} \ p \ V \rangle] &= \Theta[V] \\
 \Theta[c] &= c \\
 \Theta[x] &= x \\
 \Theta[(\lambda x. M)] &= (\lambda x. \Theta[M]) \\
 \dots &= \dots \text{ similar clauses for other } \Lambda_{pc'} \text{ terms} \\
 \Theta[\mathbf{error}] &= \mathbf{error} \\
 \Theta[(\mathbf{f-let} \ (p \ S_1) \ S_1)] &= (\mathbf{f-let} \ (p \ \Theta[S_1]) \ \Theta[S_2])
 \end{aligned}$$

Θ extends in a natural manner to evaluation contexts: $\Theta[\] = [\], \dots$ ■

The function Θ is a *bisimulation relation*, i.e., every transition of the $P(C_{ph})$ -machine corresponds to a transition of the $P(C)$ -machine, and vice versa. The proof of this property relies on the following lemma concerning the translation Θ .

Lemma 4.2

1. For all $V \in Value_{pc'}$, if $touch_{pc'}[V] \neq \circ$ then $\Theta[touch_{pc'}[V]] = \Theta[V]$.
2. For all $V \in Value_{pc'}$, $unload_{pc'}[V] = unload_{pc}[\Theta[V]]$.
3. $\Theta[\mathcal{E}][\Theta[M]] = \Theta[\mathcal{E}[M]]$
4. $\Theta[M][x \leftarrow \Theta[V]] = \Theta[M[x \leftarrow V]]$

Proof: The four facts follow from straightforward structural induction. ■

Theorem 4.3 (Bisimulation Theorem) Let $S_1 \in State_{pc}$ and $S'_1 \in State_{pc'}$ such that $\Theta[S'_1] = S_1$.

1. If $S'_1 \xrightarrow{pc'}^{n,m} S'_2$, then $S_1 \xrightarrow{pc}^{n,m} \Theta[S'_2]$.
2. If $S_1 \xrightarrow{pc}^{n,m} S_2$, then there exists $S'_2 \in State_{pc'}$ such that $S'_1 \xrightarrow{pc'}^{n,m} S'_2$ and $\Theta[S'_2] = S_2$.

Proof:

1. The proof of the first part proceeds by lexicographic induction on n and on the size of S'_1 , and by case analysis of the last step in $S'_1 \xrightarrow{pc'}^{n,m} S'_2$. The interesting cases concern the placeholder-strict operations. We examine the case for the transition rule (*car*) in detail.

Suppose $S'_1 \xrightarrow{pc'}^{1,1} S'_2$ via the transition rule (*car*). Then

$$\begin{aligned} S'_1 &= \mathcal{E}[\text{let } (x \text{ (car } V)) M] \\ S_1 &= \Theta[\mathcal{E}] [(\text{let } (x \text{ (car } \Theta[V])) \Theta[M])] \end{aligned}$$

We consider three cases for V separately:

- Suppose $\text{touch}_{pc'}[V] = (\text{cons } V_1 V_2)$. Then

$$\begin{aligned} S'_2 &= \mathcal{E}[M[x \leftarrow V_1]] \\ \Theta[V] &= (\text{cons } \Theta[V_1] \Theta[V_2]) \\ S_1 &\xrightarrow{pc'}^{1,1} \Theta[\mathcal{E}] [\Theta[M][x \leftarrow \Theta[V_1]]] \\ &= \Theta[\mathcal{E}] [\Theta[M[x \leftarrow V_1]]] \\ &= \Theta[\mathcal{E}[M[x \leftarrow V_1]]] \\ &= \Theta[S'_2] \end{aligned}$$

- Suppose $\text{touch}_{pc'}[V] \notin \text{Pair}_{pc'} \cup \{\circ\}$, i.e., $\text{touch}_{pc'}[V] \in \text{Const} \cup \text{Vars} \cup \text{Cl}_{pc'}$. Then

$$\begin{aligned} S'_2 &= \text{error} \\ \Theta[V] &\in \text{Const} \cup \text{Vars} \cup \{(\lambda x. M)\} \\ S_1 &\xrightarrow{pc'}^{1,1} \text{error} \end{aligned}$$

- The case where $\text{touch}_{pc'}[V] = \circ$ is impossible, since S'_1 would not be a (*car*)-redex.

Thus the case for the transition rule (*car*) holds. The analysis of the transition rules (*cdr*), (*if*) and (*apply*) is similar. We consider in detail the case for the (*parallel*) rule, in order to justify our use of lexicographic induction.

Suppose $S'_1 \xrightarrow{pc'}^{n,m} S'_2$ via the rule (*parallel*). Then there exists some states $X'_1, Y'_1, X'_2, Y'_2 \in \text{State}_{pc'}$ such that:

$$\begin{aligned} S'_1 &= (\text{f-let } (p' X'_1) Y'_1) \quad n & X'_1 &\xrightarrow{pc'}^{a,m} X'_2 \\ S'_2 &= (\text{f-let } (p' X'_2) Y'_2) \quad n & Y'_1 &\xrightarrow{pc'}^{c,d} Y'_2 \end{aligned}$$

where $n = a + c$. Since X'_1 and Y'_1 are substates of S'_1 , by the inductive hypothesis, we have:

$$\begin{aligned} \Theta[X'_1] &\xrightarrow{pc'}^{a,m} \Theta[X'_2] \\ \Theta[Y'_1] &\xrightarrow{pc'}^{c,d} \Theta[Y'_2] \end{aligned}$$

Hence:

$$\begin{aligned}
S_1 &= (\mathbf{f}\text{-let } (p' \Theta[X'_1]) \Theta[Y'_1]) \\
&\xrightarrow[n,m]{pc} (\mathbf{f}\text{-let } (p' \Theta[X'_2]) \Theta[Y'_2]) && \text{via } (parallel) \\
&= \Theta[S'_2]
\end{aligned}$$

The proofs for the remaining cases are straightforward.

2. We prove the second part by lexicographic induction on n and on the size of S_1 , and by case analysis of the last step in the transition $S_1 \xrightarrow[n,m]{pc} S_2$. Again, the interesting cases concern placeholder-strict operations, and we consider the transition rule (*car*) in detail.

Suppose $S_1 \xrightarrow{1,1}{pc} S_2$ via the transition rule (*car*). Then

$$S_1 = \mathcal{E}[(\mathbf{let } (x (\mathbf{car } V)) M)]$$

An examination of the definition of Θ shows that S'_1 must be

$$\mathcal{E}'[(\mathbf{let } (x (\mathbf{car } V')) M')]$$

where $\Theta[\mathcal{E}'] = \mathcal{E}$, $\Theta[V'] = V$ and $\Theta[M'] = M$. We consider three cases for V :

- Suppose $V = (\mathbf{cons } V_1 V_2)$. It follows that $S_2 = \mathcal{E}[M[x \leftarrow V_1]]$ and $V' = (\mathbf{cons } V'_1 V'_2)$, where $\Theta[V'_i] = V_i$ for $i = 1, 2$. Choosing $S'_2 = \mathcal{E}'[M'[x \leftarrow V'_1]]$, we have that $S'_1 \xrightarrow{1,1}{pc'} S'_2$ and $\Theta[S'_2] = S_2$.
- Suppose V is neither a pair nor a placeholder. Then $V \in \mathbf{Const} \cup \mathbf{Vars} \cup \{(\lambda x. M)\}$, and hence $V' \in \mathbf{Const} \cup \mathbf{Vars} \cup \{(\lambda x. M)\}$. Therefore $S_2 = S'_2 = \mathbf{error}$, and the conclusion holds for case, too.
- The case where V is a placeholder is impossible, since S_1 would not be a (*car*)-redex.

Thus the theorem holds if $S_1 \xrightarrow[n,m]{pc} S_2$ via the transition rule (*car*). The analysis of the transition rules (*cdr*), (*if*) and (*apply*) is similar. We also examine in detail the case for the (*parallel*) rule, in order to show why we choose to use lexicographic induction.

Suppose $S_1 \xrightarrow[n,m]{pc} S_2$ via the rule (*parallel*). Then there exists some states $X_1, Y_1, X_2, Y_2 \in \mathbf{State}_{pc}$ such that:

$$\begin{aligned}
S_1 &= (\mathbf{f}\text{-let } (p X_1) Y_1) && X_1 \xrightarrow[a,m]{pc} X_2 \\
S_2 &= (\mathbf{f}\text{-let } (p X_2) Y_2) && Y_1 \xrightarrow[c,d]{pc} Y_2
\end{aligned}$$

with $n = a + c$. Examining the definition of Θ shows that S'_1 must be:

$$(\mathbf{f}\text{-let } (p X'_1) Y'_1)$$

for some $X'_1, Y'_1 \in \mathbf{State}_{pc'}$ such that $\Theta[X'_1] = X_1$ and $\Theta[Y'_1] = Y_1$. Since X_1 and Y_1 are substates of S_1 , by induction there exists states $X'_2, Y'_2 \in \mathbf{State}_{pc'}$ with:

$$\begin{aligned}
X'_1 &\xrightarrow[a,m]{pc} X'_2 && \Theta[X'_2] &= X_2 \\
Y'_1 &\xrightarrow[c,d]{pc} Y'_2 && \Theta[Y'_2] &= Y_2
\end{aligned}$$

Taking $S'_2 = (\mathbf{f}\text{-let } (p \ X'_2) \ Y'_2)$, we have that $\Theta[S'_2] = S_2$ and that:

$$\begin{aligned} S'_1 &= (\mathbf{f}\text{-let } (p \ X'_1) \ Y'_1) \\ &\xrightarrow[n, m]{pc} (\mathbf{f}\text{-let } (p \ X'_2) \ Y'_2) \\ &= S'_2 \end{aligned}$$

The proofs for the remaining cases are straightforward.

■

The Bisimulation Theorem implies the equivalence of the $P(C_{ph})$ -machine and the $P(C)$ -machine.

Theorem 4.4 (Correctness of $eval_{pc'}$) $eval_{pc'} = eval_{pc}$.

Proof: Both directions of the theorem are straightforward consequences of the Bisimulation Theorem. ■

Since the $P(C)$ -machine is consistent and complete (from Theorems 3.6 and 3.9), the equivalence of machines implies that the $P(C_{ph})$ -machine is also consistent and complete.

Theorem 4.5 (Consistency and Completeness of $eval_{pc'}$) *The relation $eval_{pc'}$ is a total function.*

In summary, the $P(C_{ph})$ -machine is a correct implementation of the $P(C)$ -machine in that both define the same semantics for the source language. Hence the use of placeholder objects, combined with *touch* operations for placeholder-strict primitives, is a valid technique for coordinating parallel tasks.

5 The Low-Level Operational Semantics

Since optimizations heavily rely on static information about the values that variables can assume, the $P(C_{ph})$ -machine is ill-suited for correctness proofs of appropriate analysis algorithms.⁵ The states of the $P(C_{ph})$ -machine contain no binding information relating program variables and values. Instead, the machine relies on substitution for making progress. To address this problem, we refine the $P(C_{ph})$ -machine to the $P(CEK)$ -machine (see Figures 7 and 8) using standard techniques [5, 9].

5.1 Specification of the $P(CEK)$ -machine

The substitution operation is replaced by an *environment* in the usual manner. An environment E is a finite mapping from variables to run-time values. The empty environment is denoted by \emptyset , and the operation $E[x \leftarrow V]$ extends the environment E to map the variable x to the value V . Using function notation, we write $E(x)$ to denote the value bound to x

⁵The machine is also far too abstract for the derivation of an implementation. This problem is also addressed by the following development.

Evaluator:

$$\begin{aligned}
eval_{pcek} : \Lambda_a^0 &\longrightarrow Answers \cup \{\mathbf{error}, \perp\} \\
eval_{pcek}(P) &= \begin{cases} unload_{pcek}[E(x)] & \text{if } \langle P, \emptyset, \epsilon \rangle \mapsto_{pcek}^* \langle x, E, \epsilon \rangle \\ \mathbf{error} & \text{if } \langle P, \emptyset, \epsilon \rangle \mapsto_{pcek}^* \mathbf{error} \\ \perp & \text{if } \forall i \in \mathbf{N} \exists S_i \in State_{pcek}, n_i, m_i \in N \text{ such that} \\ & m_i > 0, S_0 = \langle P, \emptyset, \epsilon \rangle \text{ and } S_i \mapsto_{pcek}^{n_i, m_i} S_{i+1} \end{cases}
\end{aligned}$$

Data Specifications:

$$\begin{array}{lll}
S \in State_{pcek} & ::= & \langle M, E, K \rangle \mid \mathbf{error} \mid (\mathbf{f-let} (p S) S) & (\text{States}) \\
M \in \Lambda_a & & & (A\text{-nf Language}) \\
E \in Env_{pcek} & ::= & Vars \longrightarrow_f Value_{pcek} & (\text{Environments}) \\
V \in Value_{pcek} & ::= & PValue_{pcek} \mid Ph\text{-}Obj_{pcek} & (\text{Run-Time Values}) \\
PValue_{pcek} & ::= & c \mid x \mid Cl_{pcek} \mid Pair_{pcek} & (\text{Proper Values}) \\
Cl_{pcek} & ::= & \langle (\lambda x. M), E \rangle & (\text{Closures}) \\
Pair_{pcek} & ::= & (\mathbf{cons} V V) & (\text{Pairs}) \\
Ph\text{-}Obj_{pcek} & ::= & \langle \mathbf{ph} p \circ \rangle \mid \langle \mathbf{ph} p V \rangle & (\text{Placeholder Values}) \\
K \in Cont_{pcek} & ::= & \epsilon \mid \langle \mathbf{ar} x, M, E \rangle.K \mid \langle \mathbf{ar}^\dagger x, M, E \rangle.K & (\text{Continuations}) \\
F \in FinalState_{pcek} & ::= & \langle x, E, \epsilon \rangle \mid \mathbf{error} & (\text{Final States})
\end{array}$$

Auxiliary Functions:

$$\begin{aligned}
unload_{pcek} : Value_{pcek}^0 &\longrightarrow Answers & touch_{pcek} : Value_{pcek} &\longrightarrow PValue_{pcek} \cup \{\circ\} \\
unload_{pcek}[c] &= c & touch_{pcek}[\langle \mathbf{ph} p \circ \rangle] &= \circ \\
unload_{pcek}[\langle (\lambda x. M), E \rangle] &= \mathbf{procedure} & touch_{pcek}[\langle \mathbf{ph} p V \rangle] &= touch_{pcek}[V] \\
unload_{pcek}[(\mathbf{cons} V_1 V_2)] &= (\mathbf{cons} V'_1 V'_2) & touch_{pcek}[V] &= V \text{ if } V \in PValue_{pcek} \\
&\text{where } V'_i = unload_{pcek}[V_i] & & \\
unload_{pcek}[\langle \mathbf{ph} p V \rangle] &= unload_{pcek}[V] & &
\end{aligned}$$

Placeholder Substitution $S[p := V], E[p := V], K[p := V]$:

$$\begin{aligned}
(\mathbf{f-let} (p' S_1) S_2)[p := V] &= \begin{cases} (\mathbf{f-let} (p' S_1[p := V]) S_2) & \text{if } p = p' \\ (\mathbf{f-let} (p' S_1[p := V]) S_2[p := V]) & \text{if } p \neq p' \end{cases} \\
\langle M, E, K \rangle[p := V] &= \langle M, E[p := V], K[p := V] \rangle \\
\mathbf{error}[p := V] &= \mathbf{error} \\
E[p := V] &= E' \text{ such that } E'(x) = \begin{cases} \langle \mathbf{ph} p V \rangle & \text{if } E(x) = \langle \mathbf{ph} p \circ \rangle \\ E(x) & \text{otherwise} \end{cases} \\
\epsilon[p := V] &= \epsilon \\
\langle \mathbf{ar} x, M, E \rangle.K[p := V] &= \langle \mathbf{ar} x, M, E[p := V] \rangle.K[p := V] \\
\langle \mathbf{ar}^\dagger x, M, E \rangle.K[p := V] &= \langle \mathbf{ar}^\dagger x, M, E[p := V] \rangle.K[p := V]
\end{aligned}$$

FIGURE 7: The $P(CEK)$ -machine: Evaluator and Data Specifications

in E , $dom(E)$ to denote the domain of E , and $E|_{FV[(\lambda x. M)]}$ to denote the restriction of E to the free variables of $(\lambda x. M)$.

An evaluation context, which represents the control stack, is now represented by a *continuation*. A continuation consists of a sequence of *activation records*, which are similar

Transition Rules:

$$\begin{array}{llll}
\langle (\text{let } (x \ c) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \langle M, E[x \leftarrow c], K \rangle & (\text{bind-const}) \\
\langle (\text{let } (x \ y) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \langle M, E[x \leftarrow E(y)], K \rangle & (\text{bind-var}) \\
\langle (\text{let } (x \ (\lambda y. N)) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \langle M, E[x \leftarrow \langle (\lambda y. N), E' \rangle], K \rangle & (\text{bind-lam}) \\
& & \text{where } E' = E|_{FV[(\lambda y. N)]} & \\
\langle (\text{let } (x \ (\text{cons } y \ z)) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \langle M, E[x \leftarrow (\text{cons } E(y) \ E(z))], K \rangle & (\text{bind-cons}) \\
\langle x, E, \langle \text{ar } y, M, E' \rangle. K \rangle & \xrightarrow[pcek]{1,1} & \langle M, E'[y \leftarrow E(x)], K \rangle & (\text{return}) \\
\langle (\text{let } (x \ (\text{car } y)) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \begin{cases} \langle M, E[x \leftarrow V_1], K \rangle & \text{if } touch_{pcek}[E(y)] = (\text{cons } V_1 \ V_2) \\ \text{error} & \text{if } touch_{pcek}[E(y)] \notin Pair_{pcek} \cup \{\circ\} \end{cases} & (\text{car}) \\
\langle (\text{let } (x \ (\text{cdr } y)) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \text{analogous to } car & (\text{cdr}) \\
\langle (\text{let } (x \ (\text{if } y \ M_1 \ M_2)) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \begin{cases} \langle M_1, E, \langle \text{ar } x, M, E \rangle. K \rangle & \text{if } touch_{pcek}[E(y)] \notin \{\text{nil}, \circ\} \\ \langle M_2, E, \langle \text{ar } x, M, E \rangle. K \rangle & \text{if } touch_{pcek}[E(y)] = \text{nil} \end{cases} & (\text{if}) \\
\langle (\text{let } (x \ (\text{apply } y \ z)) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \begin{cases} \langle N, E'[x' \leftarrow E(z)], \langle \text{ar } x, M, E \rangle. K \rangle & \text{if } touch_{pcek}[E(y)] = \langle (\lambda x'. N), E' \rangle \\ \text{error} & \text{if } touch_{pcek}[E(y)] \notin Cl_{pcek} \cup \{\circ\} \end{cases} & (\text{apply}) \\
\langle (\text{let } (x \ (\text{future } N)) \ M), E, K \rangle & \xrightarrow[pcek]{1,1} & \langle N, E, \langle \text{ar}^\dagger x, M, E \rangle. K \rangle & (\text{future}) \\
\langle x, E, \langle \text{ar}^\dagger y, M, E' \rangle. K \rangle & \xrightarrow[pcek]{1,1} & \langle M, E'[y \leftarrow E(x)], K \rangle & (\text{future-id}) \\
\langle M, E, K_1, \langle \text{ar}^\dagger x, N, E' \rangle. K_2 \rangle & \xrightarrow[pcek]{1,0} & \langle N, E'[x \leftarrow \langle \text{ph } p \circ \rangle], K_2 \rangle & \text{if } p \notin FP(E') \cup FP(K_2) \quad (\text{fork}) \\
\langle \text{f-let } (p \ \langle x, E, \epsilon \rangle) \ S \rangle & \xrightarrow[pcek]{1,1} & S[p := E(x)] & (\text{join}) \\
\langle \text{f-let } (p \ \text{error}) \ S \rangle & \xrightarrow[pcek]{1,1} & \text{error} & (\text{join-error}) \\
\langle \text{f-let } (p_2 \ (\text{f-let } (p_1 \ S_1) \ S_2)) \ S_2 \rangle & \xrightarrow[pcek]{1,1} & \langle \text{f-let } (p_1 \ S_1) \ (\text{f-let } (p_2 \ S_2) \ S_3) \rangle & (\text{lift}) \\
& & p_1 \notin FP(S_3) & \\
\langle \text{f-let } (p \ S_1) \ S_2 \rangle & \xrightarrow[pcek]{a+c, b} & \langle \text{f-let } (p \ S'_1) \ S'_2 \rangle & (\text{parallel}) \\
& & \text{if } S_1 \xrightarrow[pcek]{a, b} S'_1 \xrightarrow[pcek]{c, d} S'_2 & \\
S & \xrightarrow[pcek]{0,0} & S & (\text{reflexive}) \\
S & \xrightarrow[pcek]{a+c, b+d} & S'' & (\text{transitive}) \\
& & \text{if } S \xrightarrow[pcek]{a, b} S', S' \xrightarrow[pcek]{c, d} S'', a, c > 0 &
\end{array}$$

FIGURE 8: The $P(CEK)$ -machine: Transition Rules

to closures. Evaluation of an **apply** expression produces an activation record of the form $\langle \text{ar } x, M, E \rangle$ to record the calling context. Evaluation of a **future** expression produces a *tagged* activation record of the form $\langle \text{ar}^\dagger x, M, E \rangle$. A tagged activation record represents

a point where the continuation can be split into separate tasks (`cmp.` (`fork`)).

The transition relation \longrightarrow_{pcek} of the $P(CEK)$ -machine is a reformulation of the relation $\longrightarrow_{pc'}$ that takes into account the change of state representation. We write $S \longrightarrow_{pcek}^* S'$ if $S \longrightarrow_{pcek}^{n,m} S'$ for some $n, m \in \mathbb{N}$. A state $S \in \text{State}_{pcek}$ is in *normal-form* if there is no non-trivial transition from S . States of the form $\langle x, E, \epsilon \rangle$ or **error** are called *final states*, and a state *blocked* if it is in normal form but not a final state.

5.2 Correctness of the $P(CEK)$ -machine

The proof of correctness of the $P(CEK)$ -machine uses standard proof techniques from [5], appropriately modified to account for parallel evaluation.

Theorem 5.1 (Correctness of $eval_{pcek}$) $eval_{pcek} = eval_{pc'}$

Proof: The proof of this theorem is included in appendix A. ■

The equivalence of machines implies that the $P(CEK)$ -machine is also consistent and complete, since the $P(C_{ph})$ -machine is consistent and complete.

Corollary 5.2 (Consistency and Completeness of $eval_{pcek}$) *The relation $eval_{pcek}$ is a total function.*

The four evaluators we present, $eval_c$, $eval_{pc}$, $eval_{pc'}$ and $eval_{pcek}$, corresponding to the four abstract machines, are provably equivalent. The parallel execution behavior of programs on the $P(CEK)$ -machine is indistinguishable from their behavior on the C -machine, showing that the $P(CEK)$ -machine is a correct implementation of **futures** according to the definitional semantics of $eval_c$.

6 Related Work

The literature on programming languages contains a number of descriptions of the semantics of parallel Scheme-like languages. The only one that directly deals with parallelism based on transparent annotations is Moreau's Ph.D. thesis [20]. Moreau studies the functional core of Scheme extended with **pcall** (a construct for evaluating function and argument expressions of an application in parallel) and first-class continuations. His primary goal is to design a semantics for the language that treats **pcall** as a pure annotation, and to derive a reasonably efficient implementation. The semantics is an extension of Felleisen and Friedman's control calculus [5]; the implementation is a parallel version of the CESK machine [6] that implements placeholders as globally accessible reference cells. The equivalence proof establishes that both evaluators define the same observational equivalence relation via the construction of a number of intermediate calculi and machines. It is far more complicated than our diamond and bisimulation techniques, possibly due to the inclusion of continuations.

Independently, Reppy [21] and Leroy [17] define a formal operational semantics for an ML-like language with first-class synchronization operations. Reppy's language, Concurrent ML, can provide the **future** mechanism as an abstraction over the given primitives. The semantics is a two-level rewriting system. The first level, also a program rewriting system

in the tradition of Felleisen and Friedman [5, 6], accounts for the sequential behavior; the second-level specifies the behavior of *sets* of parallel tasks and the task communication mechanisms. Reppy uses the semantics to prove a type soundness theorem for the extended semantics; he does not construct a low-level semantics that can serve as the basis of an implementation or a program analysis tool. Leroy formulates a semantics for a subset of CML in the traditional “natural” semantics framework. He also uses his semantics to prove the type soundness of the complete language. No attempt is made to exploit the semantics for the derivation of an analysis algorithm or a compiler optimization.

Jaganathan and Weeks [14] define an operational semantics for a simple function language extended with the **spawn** construct. They also show how the **future** annotation can be implemented using **spawn**. Since their primary goal is the derivation of a semantically well-founded abstract interpretation (in the spirit of Cousot and Cousot [2]), they extend Deutsch’s transition semantics [3] to their language. The transition semantics requires the assignments of a unique label to each sub-expression of a program and expresses computation as the movement of a token from label to label. An auxiliary label on each sub-expression is used to collect information about the values of the expression. The semantics is well-suited for deriving traditional abstract interpretations, but is inappropriate for specifying a user-level semantics.

Finally, Wand [24] recently extended his work on correctness proofs for sequential compilers to parallel languages. In his prior work on the correctness of sequential compilers, he derived compilers from the semantic mappings that translate syntax into λ -calculus expressions. Such a derivation consists of a staging process that separates the run-time portion of the semantic mapping from the compile-time portion. To prove the correctness of the compiler, it suffices to prove that the “composition” of the two functions yields the semantic mapping. The extension of this work to parallel compilers starts from a semantic mapping that translates a Scheme-like language with process creation and communication constructs into a higher-order calculus of communication and computation. After separating the compiler from the “machine”, the correctness proof is a combination of (a stronger version of) the sequential correctness proof and a correctness proof for the parallel portion of the language. The proof techniques are related to the ones we used to prove the equivalence of the $P(CEK)$ -machine and the $P(C)$ -machine.

7 Applications of the Semantics

We believe that the low-level semantics of the $P(CEK)$ -machine is a sound basis for further research into the optimization and implementation of **futures**. We have used the semantics to derive a provably-correct program optimization algorithm that removes redundant *touch* operations from programs [7, 8]. Experiments with the Gambit compiler [4] show that this optimization substantially reduces program execution times on a standard set of benchmarks.

Another important application of the low-level semantics is as a basis from which to derive a parallel and distributed implementation of Scheme with **futures**. We plan to reformulate the low-level semantics in a *message-passing framework* and to implement it in the near future.

A Correctness Proof for the $P(CEK)$ -Machine

Each $P(CEK)$ -machine state corresponds to a $P(C_{ph})$ -machine state according to the translation Φ_s : see Figure 9. For example, the $P(C_{ph})$ -machine state $\langle C, E, K \rangle$ corresponds to the $P(C_{ph})$ -machine state $\mathcal{E}[M]$, where $\mathcal{E} = \Phi_k \llbracket K \rrbracket$ is the evaluation context represented by the continuation K , and the term $M = \Phi_m[C, E]$ is the result of applying the substitutions recorded in E to the term C .

The transition relation of the $P(CEK)$ -machine is simply a reformulation of the transition relation of the $P(C_{ph})$ -machine. Therefore, each transition of the $P(CEK)$ -machine corresponds to a transition of the $P(C_{ph})$ -machine, although the number of steps involved in corresponding transitions is not identical. In particular, the $P(C_{ph})$ -machine assumes an automatic division of each sequential term into an evaluation context and a redex. Since evaluation contexts may include **future** contexts of the form $(\mathbf{let} (x \text{ (future } [])) M)$, the $P(C_{ph})$ -machine can immediately initiate evaluation inside a **future** context. In contrast, the lower-level $P(CEK)$ -machine must first convert all enclosing **future** contexts into tagged activation records via the transition rule (*future*). Thus, certain mandatory transitions on the $P(CEK)$ -machine correspond to an identity transition on the $P(C_{ph})$ -machine. The following theorem formalizes this correspondence, and furthermore proves that there is no infinite sequence of mandatory $P(CEK)$ -machine transitions that corresponds to an identity transition on the $P(C_{ph})$ -machine.

$$\begin{array}{ll}
\Phi_m : \Lambda_a \times Env_{pcek} & \longrightarrow \Lambda_{pc'} \\
\Phi_m[M, \{x_1 \leftarrow V_1, \dots, x_n \leftarrow V_n\}] & = M[x_1 \leftarrow \Phi_v[V_1]] \dots [x_n \leftarrow \Phi_v[V_n]] \\
\\
\Phi_v : Value_{pcek} & \longrightarrow Value_{pc'} \\
\Phi_v[c] & = c \\
\Phi_v[\langle (\lambda x. M), E \rangle] & = \Phi_m[(\lambda x. M), E] \\
\Phi_v[\langle \mathbf{cons} V_1 V_2 \rangle] & = \langle \mathbf{cons} \Phi_v[V_1] \Phi_v[V_2] \rangle \\
\Phi_v[\langle \mathbf{ph} p \circ \rangle] & = \langle \mathbf{ph} p \circ \rangle \\
\Phi_v[\langle \mathbf{ph} p V \rangle] & = \langle \mathbf{ph} p \Phi_v[V] \rangle \\
\\
\Phi_k : Cont_{pcek} & \longrightarrow EvalCtx \\
\Phi_k \llbracket \epsilon \rrbracket & = [] \\
\Phi_k \llbracket \langle \mathbf{ar} x, M, E \rangle . K \rrbracket & = \Phi_k \llbracket K \rrbracket [(\mathbf{let} (x []) \Phi_m[M, E])] \\
\Phi_k \llbracket \langle \mathbf{ar}^\dagger x, M, E \rangle . K \rrbracket & = \Phi_k \llbracket K \rrbracket [(\mathbf{let} (x \text{ (future } [])) \Phi_m[M, E])] \\
\\
\Phi_s : State_{pcek} & \longrightarrow State_{pc'} \\
\Phi_s[\langle M, E, K \rangle] & = \Phi_k \llbracket K \rrbracket [\Phi_m[M, E]] \\
\Phi_s[\mathbf{error}] & = \mathbf{error} \\
\Phi_s[(\mathbf{f-let} (p S_1) S_2)] & = (\mathbf{f-let} (p \Phi_s[S_1]) \Phi_s[S_2])
\end{array}$$

FIGURE 9: The Functions Φ_m , Φ_v , Φ_k and Φ_s .

Lemma A.1 (Weak Bisimulation Theorem) *Let $S_1 \in State_{pcek}$ and $S'_1 \in State_{pc'}$ such that $\Phi_s[S_1] = S'_1$.*

1. If $S'_1 \mapsto_{pc'}^{n',m'} S'_2$ then there exists $S_2 \in \text{State}_{pcek}$ such that $S_1 \mapsto_{pcek}^{n,m} S_2$, $\Phi_s[S_2] = S'_2$ and $m \geq m'$.
2. If $S_1 \mapsto_{pcek}^{n,m} S_2$ then $S'_1 \mapsto_{pc'}^{n',m'} \Phi_s[S_2]$.
3. In part 2, for each S_1 there exists $m_0 \in \mathbb{N}$ such that $m > m_0$ implies $m' > 0$.

Proof:

1. The proof of the first part is by case analysis of the last step in $S'_1 \mapsto_{pc'}^{n',m'} S'_2$, based on a similar proof in [5] for sequential programs.
2. We prove the second part by case analysis of the last step in the transition $S_1 \mapsto_{pcek}^{n,m} S_2$.

- Suppose $S_1 \mapsto_{pcek}^{1,1} S_2$ via the rule (*bind-const*). Then

$$\begin{aligned}
S_1 &= \langle (\text{let } (x \ c) \ M), E, K \rangle \\
S_2 &= \langle M, E[x \leftarrow c], K \rangle \\
\Phi_s[S_1] &= \mathcal{E} [\Phi_m[(\text{let } (x \ c) \ M), E]] \quad \text{where } \mathcal{E} = \Phi_k \llbracket K \rrbracket \\
&= \mathcal{E} [(\text{let } (x \ c) \ \Phi_m[M, E])] \quad \text{renaming so } x \notin \text{dom}(E) \\
&\mapsto_{pc'}^{1,1} \mathcal{E} [\Phi_m[M, E][x \leftarrow c]] \\
&= \mathcal{E} [\Phi_m[M, E[x \leftarrow c]]] \quad \text{by definition of } \Phi_m \\
&= \Phi_s[S_2]
\end{aligned}$$

- The analysis for the rules (*bind-var*), (*bind-lam*) and (*bind-cons*) is similar.
- Suppose $S_1 \mapsto_{pcek}^{1,1} S_2$ via the rule (*future*). Then

$$\begin{aligned}
S_1 &= \langle (\text{let } (x \ (\text{future } N)) \ M), E, K \rangle \\
S_2 &= \langle N, E, \langle \text{ar}^\dagger x, M, E \rangle . K \rangle \\
\Phi_s[S_1] &= \Phi_k \llbracket K \rrbracket [\Phi_m[(\text{let } (x \ (\text{future } N)) \ M), E]] \\
&= \Phi_k \llbracket K \rrbracket [(\text{let } (x \ (\text{future } \Phi_m[N, E])) \ \Phi_m[M, E])] \\
&\quad \text{renaming so } x \notin \text{dom}(E) \\
&= \Phi_k \llbracket \langle \text{ar}^\dagger x, M, E \rangle . K \rrbracket [\Phi_m[N, E]] \\
&= \Phi_s[S_2]
\end{aligned}$$

- Suppose $S_1 \mapsto_{pcek}^{1,1} S_2$ via the rule (*return*). Then

$$\begin{aligned}
S_1 &= \langle x, E, \langle \text{ar } y, M, E' \rangle . K \rangle \\
S_2 &= \langle M, E'[y \leftarrow E(x)], K \rangle \\
\Phi_s[S_1] &= \Phi_k \llbracket K \rrbracket [(\text{let } (y \ E(x)) \ \Phi_m[M, E'])] \\
&\mapsto_{pc'}^{1,1} \Phi_k \llbracket K \rrbracket [\Phi_m[M, E'][y \leftarrow E(x)]] \\
&= \Phi_k \llbracket K \rrbracket [\Phi_m[M, E'[y \leftarrow E(x)]]] \\
&= \Phi_s[S_2]
\end{aligned}$$

- The analysis for the rule (*future-id*) is similar.
- Suppose $S_1 \xrightarrow{1,1}_{pcek} S_2$ via the rule (*car*). We analyze the error and non-error cases separately.
 - Suppose $S_2 \neq \mathbf{error}$. Then

$$\begin{aligned}
 S_1 &= \langle (\mathbf{let} (x (\mathbf{car} y)) M), E, K \rangle \\
 S_2 &= \langle M, E[x \leftarrow V_1], K \rangle \\
 &\quad \text{where } touch_{pcek}[E(y)] = (\mathbf{cons} V_1 V_2) \\
 \Phi_s[S_1] &= \Phi_k[K] [(\mathbf{let} (x (\mathbf{car} (\mathbf{cons} \Phi_v V_1 \Phi_v V_2))) \Phi_m[M, E])] \\
 &\xrightarrow{1,1}_{pc'} \Phi_k[K] [\Phi_m[M, E][x \leftarrow \Phi_v V_1]] \\
 &= \Phi_k[K] [\Phi_m[M, E[x \leftarrow V_1]]] \\
 &= \Phi_s[S_2]
 \end{aligned}$$

- Suppose $S_2 = \mathbf{error}$. Then

$$\begin{aligned}
 S_1 &= \langle (\mathbf{let} (x (\mathbf{car} y)) M), E, K \rangle \\
 S_2 &= \mathbf{error} \\
 &\quad \text{where } touch_{pcek}[E(y)] \notin Pair \cup \{\circ\} \\
 \Phi_s[S_1] &= \Phi_k[K] [(\mathbf{let} (x (\mathbf{car} V)) \Phi_m[M, E])] \\
 &\quad \text{where } V \neq (\mathbf{cons} V_1 V_2), V \neq p \\
 &\xrightarrow{1,1}_{pc'} \mathbf{error} \\
 &= \Phi_s[S_2]
 \end{aligned}$$

- The analysis for the rules (*cdr*), (*if*) and (*apply*) is similar.
- Suppose $S_1 \xrightarrow{1,1}_{pcek} S_2$ via the rule (*fork*). Then

$$\begin{aligned}
 S_1 &= \langle M, E, K_1. \langle \mathbf{ar} \dagger x, N, E' \rangle. K_2 \rangle \\
 S_2 &= (\mathbf{f-let} (p \langle M, E, K_1 \rangle) \langle N, E'[x \leftarrow \langle \mathbf{ph} p \circ \rangle], K_2 \rangle) \\
 &\quad \text{where } p \notin FP(E') \cup FP(K_2) \\
 \Phi_s[S_1] &= \Phi_k[K_2] [(\mathbf{let} (x (\mathbf{future} \Phi_k[K_1][\Phi_m[M, E]])) \Phi_m[N, E',])] \\
 &\xrightarrow{1,1}_{pc'} (\mathbf{f-let} (p \Phi_k[K_1][\Phi_m[M, E]]) \\
 &\quad \Phi_k[K_2] [\Phi_m[N, E'][x \leftarrow p]]) \\
 &= \Phi_s[S_2]
 \end{aligned}$$

- Suppose $S_1 \xrightarrow{1,1}_{pcek} S_2$ via the rule (*join*). Then

$$\begin{aligned}
 S_1 &= (\mathbf{f-let} (p \langle x, E, \epsilon \rangle) S) \\
 S_2 &= S[p := E(x)] \\
 \Phi_s[S_1] &= (\mathbf{f-let} (p \Phi_v[E(x)]) \Phi_s[S]) \\
 &\xrightarrow{1,1}_{pc'} \Phi_s[S][p := \Phi_v[E(x)]] \\
 &= \Phi_s[S[p := E(x)]] \\
 &= \Phi_s[S_2]
 \end{aligned}$$

- The analysis of for the rule (*join-error*) is similar.
 - The cases for the rules (*reflexive*), (*transitive*) and (*parallel*) are straightforward.
3. Suppose that $S_1 \xrightarrow[n, m]{pcek} S_2$ with $m > 0$, and $\Phi_s[S_1] \xrightarrow[n', m']{pc'} \Phi_s[S_2]$ with $m' = 0$. From the proof of part 2, we see that the only possibility for this case is via the rule (*future*), or via the (*transitive*) or (*parallel*) closure of the (*future*) rule. The number of consecutive such transitions from S_1 is bounded by $\theta[S_1]$, where θ is the following function:

$$\begin{aligned}
 \theta : State &\longrightarrow \mathbb{N} \\
 \theta[\langle M, E, K \rangle] &= \begin{cases} 1 + \theta[\langle N_1, E, K \rangle] & \text{if } M = (\mathbf{let} (x (\mathbf{future} N_1)) N_2) \\ 0 & \text{otherwise} \end{cases} \\
 \theta[\mathbf{error}] &= 0 \\
 \theta[(\mathbf{f-let} (p S_1) S_2)] &= \theta[S_1]
 \end{aligned}$$

If we pick $m_0 = \theta[S_1]$ and $m > m_0$, it is impossible that $\Phi_s[S_1] \xrightarrow[n', m']{pc'} \Phi_s[S_2]$ with $m' = 0$. Hence $m' > 0$.

■

The Correspondence Lemma implies the equivalence of the $P(C_{ph})$ -machine and the $P(CEK)$ -machines.

Lemma A.2 $unload_{pcek} = unload_{pc'} \circ \Phi_v$.

Theorem 5.1 (Correctness of $eval_{pcek}$) $eval_{pcek} = eval_{pc'}$.

Proof: The proof of the left-to-right direction proceeds by case analysis on the definitions of $eval_{pcek}$.

- Suppose $eval_{pcek}(P) = unload_{pcek}[E(x)]$ because $\langle P, \emptyset, \epsilon \rangle \xrightarrow{*}_{pcek} \langle x, E, \epsilon \rangle$. By the Weak Bisimulation Theorem (A.1), $P \xrightarrow{*}_{pc'} V$, where $V \equiv \Phi_v[E(x)]$. Hence

$$eval_{pc'}(P) = unload_{pc'}[\Phi_v[E(x)]] = unload_{pcek}[E(x)] = eval_{pcek}(P).$$

- The case where $eval_{pcek}(P) = \mathbf{error}$ is similar.
- Suppose $eval_{pcek}(P) = \perp$ because there exists $S_i \in State_{pcek}$ and $n_i, m_i \in \mathbb{N}$ such that $m_i > 0$, $S_0 = \langle P, \emptyset, \epsilon \rangle$ and $S_i \xrightarrow[n_i, m_i]{pcek} S_{i+1}$. By the Weak Bisimulation Theorem (A.1), there exists $S'_i \in State_{pc'}$ and $n'_i, m'_i \in \mathbb{N}$ such that $m'_i > 0$, $P = S_0$ and $S'_i \xrightarrow[n'_i, m'_i]{pc'} S'_{i+1}$. Hence $eval_{pc'}(P) = \perp$.

We prove right-to-left direction by case analysis of the definition of $eval_{pc'}$.

- Suppose $eval_{pc'}(P) = unload_{pc'}[V]$ via the reduction $P \mapsto_{pc'}^* V$. By the Weak Bisimulation Theorem (A.1), $\langle P, \emptyset, \epsilon \rangle \mapsto_{pcek}^* \langle x, E, \epsilon \rangle$, where $\Phi_v[E(x)] = V$. Hence

$$\begin{aligned} eval_{pcek}(P) &= unload_{pcek}[E(x)] = unload_{pc'}[\Phi_v[E(x)]] = unload_{pc'}[V] \\ &= eval_{pc'}(P). \end{aligned}$$

- The case where $eval_{pc'}(P) = \mathbf{error}$ is similar.
- Suppose $eval_{pc'}(P) = \perp$ because there exists $S_i \in State_{pc'}$ and $n_i, m_i \in \mathbf{N}$ such that $m_i > 0$, $S_0 = P$ and $S_i \mapsto_{pc'}^{n_i, m_i} S_{i+1}$. By the Weak Bisimulation Theorem (A.1), there exists $S'_i \in State_{pcek}$ and $n'_i, m'_i \in \mathbf{N}$ such that $m'_i > 0$, $S_0 = \langle P, \emptyset, \epsilon \rangle$ and $S'_i \mapsto_{pcek}^{n'_i, m'_i} S'_{i+1}$. Hence $eval_{pcek}(P) = \perp$.

Hence, for all programs P , $eval_c(P) = eval_{pc}(P)$. ■

References

- [1] BAKER, H., AND HEWITT, C. The incremental garbage collection of processes. In *Proceedings of the Symposium on Artificial Intelligence and Programming Languages* (1977), vol. 12(8), pp. 55–59.
- [2] COUSOT, P., AND COUSOT, R. Abstract interpretation: A unified lattice model for static analyses of programs by construction or approximation of fixpoints. In *POPL* (1977), pp. 238–252.
- [3] DEUTSCH, A. *Modèles Opérationnels de Langage de Programmation et Représentations de Relations sue des Langues Rationnels avec Application a la Détermination Statique de Propriétés de Partages Dynamiques de Données*. PhD thesis, Université Paris VI, 1992.
- [4] FEELEY, M. *An Efficient and General Implementation of Futures on Large Scale Shared-Memory Multiprocessors*. PhD thesis, Department of Computer Science, Brandeis University, 1993.
- [5] FELLEISEN, M., AND FRIEDMAN, D. P. Control operators, the SECD-machine, and the lambda-calculus. In *3rd Working Conference on the Formal Description of Programming Concepts* (Aug. 1986), pp. 193–219.
- [6] FELLEISEN, M., AND FRIEDMAN, D. P. A calculus for assignments in higher-order languages. In *Conference Record of the 14th Annual ACM Symposium on Principles of Programming Languages* (Munich, West Germany, Jan. 1987), pp. 314–345.
- [7] FLANAGAN, C., AND FELLEISEN, M. Well-founded touch optimization for futures. Rice University Computer Science TR94-239.
- [8] FLANAGAN, C., AND FELLEISEN, M. The semantics of future and its use in program optimizations. In *POPL* (Jan. 1995).

- [9] FLANAGAN, C., SABRY, A., DUBA, B. F., AND FELLEISEN, M. The essence of compiling with continuations. In *PLDI* (1993), pp. 237–247.
- [10] GABRIEL, R., AND MCCARTHY, J. Qlisp. *Parallel Computation and Computers for Artificial Intelligence* (1988), 63–89.
- [11] HALSTEAD, R. Multilisp: A language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems* 7, 4 (1985), 501–538.
- [12] ITO, T., AND HALSTEAD, R., Eds. *Parallel Lisp: Languages and Systems*. Springer-Verlag *Lecture Notes in Computer Science* 441, 1989.
- [13] ITO, T., AND MATSUI, M. A parallel lisp language: Pailisp and its kernel specification. [12:58–100].
- [14] JAGANNATHAN, S., AND WEEKS, S. Analyzing stores and references in a parallel symbolic language. In *LFP* (1994), pp. 294–305.
- [15] KATZ, M., AND WEISE, D. Continuing into the future: on the interaction of futures and first-class continuations. In *LFP* (1990).
- [16] KRANZ, D., HALSTEAD, R., AND MOHR, E. Mul-T: A high-performance parallel lisp. In *PLDI* (1989), pp. 81–90.
- [17] LEROY, X. *Typage polymorphe d’un langage algorithmique*. PhD thesis, Université Paris 7, 1992.
- [18] MILLER, J. *MultiScheme: A Parallel Processing System*. PhD thesis, MIT, 1987.
- [19] MOHR, E., KRANZ, R., AND HALSTEAD, R. Lazy task creation: A technique for increasing the granularity of parallel programs. In *LFP* (1990).
- [20] MOREAU, L. *Sound Evaluation of Parallel Functional Programs with First-Class Continuations*. PhD thesis, Université de Liège, 1994.
- [21] REPPY, J.H. *Higher-Order Concurrency*. PhD thesis, Cornell University, Jan. 1992.
- [22] SABRY, A., AND FELLEISEN, M. Is continuation-passing useful for data flow analysis. In *PLDI* (1994), pp. 1–12.
- [23] SWANSON, M., KESSLER, R., AND LINDSTROM, G. An implementation of portable standard lisp on the BBN butterfly. In *LFP* (1988), pp. 132–142.
- [24] WAND, M. Compiler correctness for parallel languages. Unpublished manuscript, 1995.