

# Introduction to Cost Semantics for Parallel Languages

Adrien Guatto

MPRI 2.37.1  
2019–2020

## Abstract

In this short note, we present *cost semantics* for parallel  $\lambda$ -calculi. Such semantics specify not only *what* programs compute, but also *at what cost*. We first apply this approach to a pure call-by-value  $\lambda$ -calculus. Following [Blleloch and Greiner \[1995\]](#), we sketch a provably-efficient implementation of this calculus on top of a Parallel Random Access Machine. Then, we extend the calculus with operations acting upon some abstract *state* living in an arbitrary partial commutative monoid. We then instantiate this construction to obtain an idealized version of Cilk [\[Frigo et al., 1998\]](#).

## 1 Introduction

During the previous lectures, we studied the performance of Cilk programs in an analytical fashion, in terms of *work* and *span*. Rather than resorting to some low-level machine model, we reasoned informally at the program level. This reasoning, however, was informal.

The goal of this lecture is to show how this approach can be made rigorous using techniques inspired from programming languages semantics.

**Conventions.** By *list* we mean a finite sequence, possibly empty. We write  $(X^*, \cdot, \epsilon)$  for the free monoid over  $X$ , that is, for the set of lists of elements of  $X$ , equipped with the concatenation operator  $x \cdot y$ , as well as the empty list  $\epsilon$ .

We write  $X \rightarrow_{\text{fin}} Y$  for the set of finite partial maps from  $X$  to  $Y$ . Given such a finite partial map  $f$ , we write  $f[x \mapsto y]$  for the finite partial map sending  $x'$  to  $y$  when  $x = x'$  or to  $f(x')$  when  $x' \neq x$  and  $f(x')$  is defined. We write  $\text{dom}(f)$  for the subset of  $X$  for which  $f$  is defined. Thus,  $\text{dom}(f[x \mapsto y]) = \text{dom}(f) \cup \{x\}$ . We write  $\emptyset$  for an empty finite partial map.

We write  $\mathcal{M}_{\text{fin}}(X)$  for the set of finite multisets of elements of  $X$ . We write  $[x_1, \dots, x_n]$  for the multiset of size  $n$  whose elements are  $x_1, \dots, x_n$ . By definition, we may have  $x_i = x_j$ . Given two multisets  $A, B \in \mathcal{M}_{\text{fin}}(X)$ , we denote  $A + B$  their multiset union.

We assume given some disjoint countably infinite sets  $Var$  and  $Loc$ . We call *variables* the elements  $x, y, z \dots$  of  $Var$  and *locations* the elements  $\ell, \ell_1, \ell_2 \dots$  of  $Loc$ .

All the syntactic objects we consider are always identified up to  $\alpha$ -equivalence. Furthermore, we follow Barendregt's variable convention: in a given metatheoretical statement, no variable is both free and bound and all bound variables are distinct.

Terms	$M, N, S ::= x \mid \lambda x.M \mid M N \mid c$
Constants	$c ::= i \mid \text{add} \mid \text{add}_i \mid \text{sub} \mid \text{sub}_i \mid \dots$
Environments	$\gamma, \sigma \in \text{Var} \rightarrow_{\text{fin}} \text{Val}$
Values	$V, W ::= c \mid (x.M)\{\gamma\} \mid \text{nil}$

Figure 1: Syntax of PAL

$M; \gamma \Downarrow V$			
CONST	VAR	FUN	
$\frac{}{c; \gamma \Downarrow c}$	$\frac{}{x; \gamma \Downarrow \gamma(x)}$	$\frac{}{\lambda x.M; \gamma \Downarrow (x.M)\{\gamma\}}$	
$\frac{\text{APPCLO}}{M; \gamma \Downarrow (x.S)\{\sigma\}}$	$\frac{N; \gamma \Downarrow V \quad S; \sigma[x \mapsto V] \Downarrow W}{M N; \gamma \Downarrow W}$		$\frac{\text{APPCONST}}{M; \gamma \Downarrow c \quad N; \gamma \Downarrow V}$
		$M N; \gamma \Downarrow \delta_c(V)$	

Figure 2: Plain Semantics of PAL

## 2 A Stateless Calculus

### 2.1 Syntax and Plain Semantics

Our initial goal is to design a calculus for which a cost semantics is both interesting and useful. Such a calculus is subject to the following requirements.

1. It should be as simple as possible to make formal study doable.
2. It should form a reasonable basis for a full-fledged programming language.
3. It has to be implementable efficiently atop a parallel machine model.

The first two requirements steer us towards a variant of  $\lambda$ -calculus. However, to make  $\lambda$ -calculus into an idealized programming language, one should tame unrestricted  $\beta$ -reduction by fixing an evaluation strategy. The third requirement discourages the use of call-by-name and call-by-need, which have rather complicated cost models, especially in a parallel setting. Less importantly, it also suggests adding dedicated arithmetic facilities to our calculus.

For these reasons, we follow [Blleloch and Greiner \[1995\]](#) and study the Parallel Applicative Calculus (PAL), an *applicative* (call-by-value) calculus parameterized over a family of constants. Its syntax is given in Figure 1. The definition of terms is unremarkable. Constants  $c$  should include at least integer literals  $i$  and some arithmetic operations such as `add` and `addi`. We will explain the latter shortly. This grammar also includes the objects that, while not part of the syntax proper, are needed by our semantics. Values  $V, W$  are the results of execution, and can be either a constant  $c$  or a *closure*  $(x.M)\{\gamma\}$  bundling a  $\lambda$ -abstraction with an environment  $\gamma$  holding the values of the free variables of  $M$ . The special value `nil` plays a technical role in a later section, and can be ignored for now. An environment is a finite map from variables to values.

The cost semantics we will adopt for PAL constitutes an extension of the classic big-step formulation of call-by-value with environments and closures, so it is worth remembering its

definition. It takes the form of an *evaluation* judgment

$$M; \gamma \Downarrow V$$

stating that  $M$  evaluates to  $V$  in the environment  $\gamma$ . The rules of this judgment are given in Figure 2. A constant evaluates to itself. A variable  $x$  evaluates to the value assigned to it in  $\gamma$ . Its evaluation is not defined if  $x \notin \text{dom}(\gamma)$ . A  $\lambda$ -abstraction evaluates by closing over the environment. Finally, there are two rules for applications  $M N$  depending on whether the function  $M$  evaluates to a closure or to a constant. When  $M$  evaluates to a closure  $(x.S)\{\sigma\}$ , the final result is that of the evaluation of  $S$  in  $\sigma$  extended by binding the formal parameter  $x$  to the value  $W$  of  $N$ . When  $M$  evaluates to a constant  $c$ , the final result is determined according to a family of partial map  $(\delta_c : \text{Val} \rightarrow \text{Val})_{c \in \text{Const}}$ . This map is a parameter of PAL, as is the language of constants. As an example, it could contain the following clauses.

$$\delta_{\text{add}}(i) = \text{add}_i \quad \delta_{\text{add}_i}(j) = i + j \quad \delta_{\text{sub}}(i) = \text{sub}_i \quad \delta_{\text{sub}_i}(j) = i - j \quad \dots$$

This illustrates the purely technical role played by constants such as  $\text{add}_i$  and  $\text{sub}_i$ , which serve as intermediate steps when evaluating  $\text{add}$  and  $\text{sub}$ . An alternative would be to enrich PAL with pairs and formulate  $\text{add}$  and  $\text{sub}$  as uncurried operations.

## 2.2 Cost Semantics

For our purpose, a cost semantics for PAL should give enough information to analyze its performance on a parallel machine. To do so, we will associate to every evaluation a *cost graph* build according to the following grammar.

$$\text{Cost graph} \quad g ::= \mathbf{0} \mid \mathbf{1} \mid g_1 \oplus g_2 \mid g_1 \otimes g_2$$

Here,  $\mathbf{0}$  is the graph with no vertex,  $\mathbf{1}$  the graph with a single vertex and no edge,  $g_1 \oplus g_2$  the series composition of  $g_1$  and  $g_2$  and  $g_1 \otimes g_2$  their parallel composition. We write  $\mathbf{n}$  for  $\mathbf{1}$  composed sequentially with itself  $n$  times. The work  $\text{work}(g)$  and span  $\text{span}(g)$  of a cost graph  $g$  are defined as expected.

$$\begin{array}{ll} \text{work}(\mathbf{0}) & = 0 & \text{span}(\mathbf{0}) & = 0 \\ \text{work}(\mathbf{1}) & = 1 & \text{span}(\mathbf{1}) & = 1 \\ \text{work}(g_1 \oplus g_2) & = \text{work}(g_1) + \text{work}(g_2) & \text{span}(g_1 \oplus g_2) & = \text{span}(g_1) + \text{span}(g_2) \\ \text{work}(g_1 \otimes g_2) & = \text{work}(g_1) + \text{work}(g_2) & \text{span}(g_1 \otimes g_2) & = \max(\text{span}(g_1), \text{span}(g_2)) \end{array}$$

Now, we can restate our big-step judgment as  $M; \sigma \Downarrow V; g$ , where  $g$  describes the cost of evaluating  $M$ . Its rules are given in Figure 3. They compute the same result  $V$  as those of Figure 2. Since we are mostly interested in asymptotics, elementary operations are assumed to have unit costs (rules CONST, VAR, and FUN). Rule APPCLO is more interesting: it specifies that the function and its argument are evaluated in parallel. The overhead of the rule itself is  $\mathbf{2}$ . This choice of constant is needed to obtain a tight correspondance with the machine model described in the next section — one may replace it with  $\mathbf{1}$  without affecting asymptotic complexity results. Rule APPCONST is similar but simpler.

$$\boxed{M; \gamma \Downarrow V; g}$$

CONST	VAR	FUN
$\frac{}{c; \gamma \Downarrow c; \mathbf{1}}$	$\frac{}{x; \gamma \Downarrow \gamma(x); \mathbf{1}}$	$\frac{}{\lambda x.M; \gamma \Downarrow (x.M)\{\gamma\}; \mathbf{1}}$
$ \frac{\text{APPCLO} \quad M; \gamma \Downarrow (x.S)\{\sigma\}; g_1 \quad N; \gamma \Downarrow V; g_2 \quad S; \sigma[x \mapsto V] \Downarrow W; g_3}{M N; \gamma \Downarrow W; (g_1 \otimes g_2) \oplus g_3 \oplus \mathbf{2}} $		
$ \frac{\text{APPCONST} \quad M; \gamma \Downarrow c; g_1 \quad N; \gamma \Downarrow V; g_2}{M N; \gamma \Downarrow \delta_c(V); (g_1 \otimes g_2) \oplus \mathbf{2}} $		

Figure 3: Cost Semantics of PAL

## 2.3 Towards A Provable Implementation

We have defined a cost semantics for PAL, but how do we know that this semantics is sensible, in the sense that it can be used to reason about real-world performance, even when one only cares about asymptotics? After all, mathematics does not prevent us from associating evaluations with arbitrary cost graphs.

The solution is, as in the sequential case, to implement PAL on top of some widely-accepted, concrete parallel machine model, and then show that this implementation is sound with respect to our cost semantics. In other words, it should compute the same values as the cost semantics, and with the same work and span (up to constant factors).

We follow the exact same route as [Blleloch and Greiner \[1995\]](#), using an abstract machine as an intermediate point between PAL and the concrete machine model. We then implement this abstract machine onto the concrete machine. The only difference with their work is that we use a variant of the Krivine/CK machine rather than the SECD machine.

### 2.3.1 A Parallel Abstract Machine

The abstract machine schedules the execution of PAL programs onto several parallel *threads*. It also uses a global store to hold thread results. The grammar below describes its components.

Configuration	$C ::= T / R$
Thread pools	$T \in \mathcal{M}_{\text{fin}}(\text{Thr})$
Threads	$t ::= \langle V \mid - \mid K \rangle \mid \langle M \mid \sigma \mid K \rangle$
Stack	$K ::= \square_{\mathbf{t}}^\ell \mid \square_{\mathbf{a}}^\ell :: K \mid \square_{\mathbf{f}}^\ell :: K \mid (V \square) :: K$
Results	$R \in \text{Loc} \rightarrow_{\text{fin}} \text{Val}$

The global store  $R$  is a finite map from locations to values. A thread pool  $T$  is a bag of threads. Each thread  $t$  is a call-by-value Krivine machine, which pairs a stack  $K$  with either a value  $V$  or a closure  $(M|\sigma)$ . The empty stack  $\square_{\mathbf{t}}^\ell$  designates a final result to be stored in  $R(\ell)$ . The stack  $(V \square) :: K$  appears as is in the usual CK machine — it represents the evaluation context  $K[V \square]$ . The contexts  $\square_{\mathbf{a}}^\ell :: K$  and  $\square_{\mathbf{f}}^\ell :: K$  are used for synchronization, as will be made evident by machine transitions.

$$\begin{array}{c}
\langle c \mid \sigma \mid K \rangle \rightsquigarrow \langle c \mid - \mid K \rangle \qquad \text{SCONST} \\
\langle x \mid \sigma \mid K \rangle \rightsquigarrow \langle \sigma(x) \mid - \mid K \rangle \qquad \text{SVAR} \\
\langle \lambda x.M \mid \sigma \mid K \rangle \rightsquigarrow \langle (x.\sigma)\{M\} \mid - \mid K \rangle \qquad \text{SFUN} \\
\langle V \mid - \mid ((x.M)\{\sigma\} \square) :: K \rangle \rightsquigarrow \langle M \mid \sigma[x \mapsto V] \mid K \rangle \qquad \text{SAPPCLO} \\
\langle V \mid - \mid (c \square) :: K \rangle \rightsquigarrow \langle \delta(c, V) \mid - \mid K \rangle \qquad \text{SAPPCONST}
\end{array}$$
  

$$\begin{array}{c}
\text{ESTEP} \qquad \text{ESPAWN} \\
\frac{t \rightsquigarrow t'}{[t] / R \Rightarrow_{\mathbf{e}} [t'] / R} \qquad \frac{\ell \notin \text{dom}(R)}{[\langle MN \mid \sigma \mid K \rangle] / R \Rightarrow_{\mathbf{e}} [\langle M \mid \sigma \mid \square_{\mathbf{a}}^{\ell} :: K \rangle, \langle N \mid \sigma \mid \square_{\mathbf{f}}^{\ell} :: K \rangle] / R[\ell \hookrightarrow \text{nil}]}
\end{array}$$
  

$$\begin{array}{c}
\text{EIDLEA} \qquad \text{EIDLEF} \\
\frac{}{[\langle V \mid - \mid \square_{\mathbf{a}}^{\ell} :: K \rangle] / R \Rightarrow_{\mathbf{e}} [\langle V \mid - \mid \square_{\mathbf{a}}^{\ell} :: K \rangle] / R} \qquad \frac{}{[\langle V \mid - \mid \square_{\mathbf{f}}^{\ell} :: K \rangle] / R \Rightarrow_{\mathbf{e}} [\langle V \mid - \mid \square_{\mathbf{f}}^{\ell} :: K \rangle] / R}
\end{array}$$
  

$$\begin{array}{c}
\text{EFINISH} \qquad \text{EPAR} \\
\frac{}{[\langle V \mid - \mid \square_{\mathbf{f}}^{\ell} \rangle] / R \Rightarrow_{\mathbf{e}} [] / R[\ell \hookrightarrow V]} \qquad \frac{T_1 / R \Rightarrow_{\mathbf{e}} T'_1 / R' \quad T_2 / R' \Rightarrow_{\mathbf{e}} T'_2 / R''}{T_1 + T_2 / R \Rightarrow_{\mathbf{e}} T'_1 + T'_2 / R''}
\end{array}$$
  

$$\begin{array}{c}
\text{FSTORE} \qquad \text{FSYNC} \\
\frac{R(\ell) = \text{nil}}{[\langle V \mid - \mid \square_{\mathbf{f}}^{\ell} :: K \rangle] / R \Rightarrow_{\mathbf{f}} [] / R[\ell \hookrightarrow V]} \qquad \frac{R(\ell) \neq \text{nil}}{[\langle V \mid - \mid \square_{\mathbf{f}}^{\ell} :: K \rangle] / R \Rightarrow_{\mathbf{f}} [\langle R(\ell) \mid - \mid (V \square) :: K \rangle] / R}
\end{array}$$
  

$$\begin{array}{c}
\text{FIDLE} \qquad \text{FPAR} \\
\frac{t \neq \langle - \mid - \mid \square_{\mathbf{f}}^- :: - \rangle}{[t] / R \Rightarrow_{\mathbf{a}} [t] / R} \qquad \frac{T_1 / R \Rightarrow_{\mathbf{f}} T'_1 / R' \quad T_2 / R' \Rightarrow_{\mathbf{f}} T'_2 / R''}{T_1 + T_2 / R \Rightarrow_{\mathbf{f}} T'_1 + T'_2 / R''}
\end{array}$$
  

$$\begin{array}{c}
\text{ASTORE} \qquad \text{ASync} \\
\frac{R(\ell) = \text{nil}}{[\langle V \mid - \mid \square_{\mathbf{a}}^{\ell} :: K \rangle] / R \Rightarrow_{\mathbf{a}} [] / R[\ell \hookrightarrow V]} \qquad \frac{R(\ell) \neq \text{nil}}{[\langle V \mid - \mid \square_{\mathbf{a}}^{\ell} :: K \rangle] / R \Rightarrow_{\mathbf{a}} [\langle V \mid - \mid (R(\ell) \square) :: K \rangle] / R}
\end{array}$$
  

$$\begin{array}{c}
\text{AIDLE} \qquad \text{APAR} \\
\frac{t \neq \langle - \mid - \mid \square_{\mathbf{a}}^- :: - \rangle}{[t] / R \Rightarrow_{\mathbf{a}} [t] / R} \qquad \frac{T_1 / R \Rightarrow_{\mathbf{a}} T'_1 / R' \quad T_2 / R' \Rightarrow_{\mathbf{a}} T'_2 / R''}{T_1 + T_2 / R \Rightarrow_{\mathbf{a}} T'_1 + T'_2 / R''}
\end{array}$$
  

$$\text{STEP} \\
\frac{T / R \Rightarrow_{\mathbf{e}} T' / R' \quad T' / R' \Rightarrow_{\mathbf{f}} T'' / R'' \quad T'' / R' \Rightarrow_{\mathbf{a}} T''' / R'''}{T / R \Rightarrow T''' / R'''}$$

Figure 4: Parallel Abstract Machine for PAL

The transitions between threads and configurations are given in Figure 4. The relation  $\rightsquigarrow$  on threads is the usual one for the call-by-value Krivine machine, with most transitions related to applications missing. The main relation,  $\Rightarrow$ , operates on configurations. Its execution divides into three phases modeled by the relations  $\Rightarrow_e$ ,  $\Rightarrow_a$ , and  $\Rightarrow_f$ , as per rule STEP. All three perform exactly one transition per thread in the current thread pool — this form of parallelism arises from rules EPAR, APAR, and FPAR.

1. The *evaluation* relation  $\Rightarrow_e$  performs either a sequential step  $\rightsquigarrow$  (ESTEP), or computes  $MN$  by spawning two new threads to compute  $M$  and  $N$  in parallel (ESPAWN). In this case, we allocate a fresh location  $\ell$  to store the result of the first thread to terminate. This location is initialized to nil, so that a location is only allocated once.
2. The main rules of the *function synchronization* relation  $\Rightarrow_a$  are FSTORE and FSYNC. They deal with threads of the form  $\langle V \mid - \mid \square_a^\ell :: K \rangle$ , with the rule FIDLE ensuring that all other threads are left unchanged. Such a thread has finished computing the value  $V$  of the function term in some application that was spawned previously. What happens depends on whether the evaluation of the corresponding argument term has finished during the previous  $\Rightarrow_a$ -phase. If it is the case, then  $R(\ell)$  contains its result. Then, rule FSYNC proceeds by, so to speak, joining the two threads into  $\langle R(\ell) \mid - \mid (V \square) :: K \rangle$ , which will resume evaluating at the next  $\Rightarrow_e$ -phase. Otherwise,  $R(\ell)$  is nil, and rule FSTORE replaces it with  $V$ . This binding will be used once the argument term terminates during a subsequent  $\Rightarrow_a$ -phase.
3. The *argument synchronization* relation  $\Rightarrow_a$  is similar to the previous one, except that it deals with the argument terms spawned by applications, rather than the function ones.

To summarize, this semantics works by alternating phases that make sequential progress, spawning new threads as needed, and phases that check whether each spawned thread has finished running its subcomputation, in which case its result should be either stored or give rise to a new continuation thread.

As stated, our semantics is not deterministic, since new locations get allocated into the store in a nondeterministic manner by rule ESPAWN. However, this is the sole source of nondeterminism: since the semantics is parallel and greedy, every thread terminates as early as possible. In the following, we reason up to the renaming of locations by a bijection. This can of course be made formal, at the cost of heavier theorem statements.

In the remainder of this subsection, we prove that the parallel abstract machine is a *provably-efficient* implementation of the cost semantics. In other words, it computes correct results, and its performance matches the work and span specified by our big-step judgment. To do so, we need a number of technical results, such as the fact that the store  $R$  grows *monotonically*.

**Definition 1** (Store Extension).  $R'$  is an *extension* of  $R$ , denoted  $R \sqsubseteq R'$ , when  $\text{dom}(R) \subseteq \text{dom}(R')$  and for all  $\ell \in \text{dom}(R)$ ,  $R(\ell) \neq \text{nil}$  implies  $R'(\ell) = R(\ell)$ .

**Proposition 1** (Store Monotonicity). *If  $T / R \Rightarrow T' / R'$  then  $R \sqsubseteq R'$ .*

We can now state our main result.

**Theorem 1.** *If  $M; \emptyset \Downarrow V; g$  then there is  $R$  such that  $[\langle M \mid \emptyset \mid \square_t^\ell \rangle] / \emptyset[\ell \mapsto \text{nil}] \Rightarrow^{\text{span}(g)} [] / R$  with  $R(\ell) = V$ . Furthermore, this reduction sequence creates  $\text{work}(g)$  threads.*

*Proof sketch.* To prove the theorem, we need to generalize its statement to nonempty environments and general stacks: if  $M; \gamma \Downarrow V; g$  then, for all  $R$ ,  $\ell \notin \text{dom}(R)$ , and  $K$  of the form  $\square_{\mathbf{t}}^\ell$ ,  $\square_{\mathbf{a}}^\ell :: K'$ , or  $\square_{\mathbf{a}}^\ell :: K'$ , there exists  $R' \sqsupseteq R$  such that

$$T + [\langle M \mid \gamma \mid K \rangle] / R[\ell \hookrightarrow \text{nil}] \Rightarrow^{\text{span}(g)} T' / R'$$

with  $R'(\ell) = V$ . In addition, this reduction sequence creates one thread per work item.

We proceed by induction over the cost-semantic derivation. The only interesting rule is APPCLO, the case APPCONST being a simpler variant. By the induction hypothesis,  $M$  completes in  $\text{span}(g_1)$  steps and  $N$  in  $\text{span}(g_2)$  steps. Assume that  $\text{span}(g_1) < \text{span}(g_2)$ . Then,  $M$  terminates first, and  $N$  must terminate with ASYNC, which triggers the evaluation of the closure body. This takes  $\text{span}(g_2) + \text{span}(g_3) + 2 = \max(\text{span}(g_1), \text{span}(g_2)) + \text{span}(g_3) + 2$  steps, with the 2 additional steps coming from ESPAWN and ASYNC. The case of  $\text{span}(g_2) < \text{span}(g_1)$  is symmetric, with the argument term triggering the evaluation of the closure body. In the case of  $\text{span}(g_1) = \text{span}(g_2)$ , it is again the function term that triggers the evaluation of the closure, since  $\Rightarrow_{\mathbf{f}}$ -phases happen before  $\Rightarrow_{\mathbf{a}}$ -phases in the definition of  $\Rightarrow$ .  $\square$

### 2.3.2 From Abstract to Concrete

The structure of the parallel abstract machine makes it simple to implement on top of a low-level machine model. We choose the popular Parallel Random Access Machine (PRAM) model, in its Concurrent-Read Exclusive-Write (CREW) variant. A PRAM is a variant of the classic RAM machine, itself a convenient extension of the Turing machine. In a PRAM,  $p$  processors access a shared memory in parallel. Each step of the machine performs one step on each processor. The CREW variant is reminiscent of the *data-race-free* model studied in L. Maranget's part of the course: any number of processors may read the same memory cell during the same step, but two processors writing to the same cell raises an error.

As usual with PRAM arguments, we describe the simulation in an informal way. The multiset  $T$  of threads is represented in an array of size  $q \triangleq |T|$ . We fix an ordering of its threads. Processor  $i \in [1, p]$  is responsible for threads  $[iq/p, \dots, (i+1)q/p - 1]$ . In addition to the thread array, memory also holds the contents of  $R$ . This part of memory, which we will call  $R$ -cells, is the only one that is *logically* shared between processors; however, no data race may occur. The simulation proceeds in four phases, at the end of which all processors wait for one another.

1. Each processor evaluates the  $\Rightarrow_{\mathbf{e}}$  transitions on its segment of the thread array, updating it as needed. This requires no communication between processors, assuming we maintain a bound on the maximum memory cell used up to now, and use it in order to allocate fresh cells when simulating ESPAWN. This guarantees that the only write to an  $R$ -cell, which occurs in ESPAWN, cannot be concurrent with another write. Nor can it be concurrent to a read since those happen during later phases.
2. Each processor evaluates the  $\Rightarrow_{\mathbf{f}}$  transitions on its segment of the thread array. At most one processor will read or write to the same  $R$ -cell during this phase.
3. Similarly for  $\Rightarrow_{\mathbf{a}}$  transitions.
4. Finally, each thread has generated zero, one, or two threads to be processed during the next simulation step. They need to be flattened into a new thread array. This can

be accomplished using a prefix-sum computation to compute their offsets into the new array, as seen during the first lecture.

Of course, a more formal description would have to explain how to implement functional programming constructs on top of a (P)RAM machine. For example, we must assume the machine has unit-time operations implementing all the constants in PAL. The only non-trivial thing is the handling of environments, which can be implemented using balanced trees. This is why the following result features an overhead of  $v_M$ , where  $v_M$  is the logarithm of the number of variables in  $M$ .

**Lemma 1.** *A simulation step runs in  $kv_e(\lceil q/p \rceil + \log p)$  time, with  $k$  some positive constant.*

**Theorem 2.** *If  $M; \emptyset \Downarrow V; g$  then  $M$  is computed in  $kv_M(\text{work}(g)/p + \text{span}(g) \log p)$  time, with  $k$  some positive constant.*

*Proof.* By combining Theorem 1 and Lemma 1 as in the proof of Brent’s bound. □

[Blleloch and Greiner \[1995\]](#) also discuss the reverse direction, where one simulates an arbitrary PRAM algorithm onto the parallel abstract machine. This requires implementing the random-access memory in a purely functional way, which has some overhead.

### 3 A Stateful Calculus

While its simplicity makes it an appealing object of study, the calculus presented in the previous section is a bit too far from actual parallel programming languages. In particular, our experience writing Cilk code should have convinced us that state actually plays a crucial role when implementing real algorithms. For this reason, we would like to extend the calculus to deal with some global state.

One option would be to fix some notion of state, perhaps as a form of heap. We will rather construct a calculus parameterized by an arbitrary *partial commutative monoid* (PCM), a concept that arose in the study of separation logic. Its elements will provide the abstract notion of state, and its monoid law a way to split the state into distinct parts which can be acted upon independently.

We start by recalling the definition of a PCM before describing our calculus and its cost semantics. We then instantiate the construction to obtain an idealized variant of Cilk.

#### 3.1 Partial Commutative Monoids

We follow the modular algebraic approach to PCMs outlined by [Jung et al. \[2015\]](#).

**Definition 2.** A *partial commutative monoid*  $\mathbb{M}$  is a tuple  $(|\mathbb{M}|, \cdot_{\mathbb{M}}, \epsilon_{\mathbb{M}}, \perp_{\mathbb{M}})$  where  $(\mathbb{M}, \cdot_{\mathbb{M}}, \epsilon_{\mathbb{M}})$  is a commutative monoid and  $\perp_{\mathbb{M}} \in |\mathbb{M}|$  is absorbing for  $\cdot_{\mathbb{M}}$ .

As usual, we omit the subscripts and brackets  $| - |$  when the PCM is clear from the context. Given such a PCM, we write  $(-)\odot(=)$  for the partial function from  $\mathbb{M}^2$  to  $\mathbb{M}$  such that  $x \odot y \triangleq x \cdot y$  when  $x \cdot y \neq \perp$ , and is undefined otherwise. When defining the carrier of a PCM, we frequently omit the elements 1 and 0, and specify only  $\odot$ .

**Definition 3.** The *exclusive* PCM over a set  $X$ , denoted  $\text{Ex}(X)$ , is  $(X \uplus \{\epsilon, \perp\}, \cdot, \epsilon, \perp)$  with the monoid law defined only by  $x \odot \epsilon = \epsilon \odot x = x$  for all  $x \in X$ .



$$\begin{array}{c}
\boxed{M; \gamma; A \Downarrow V; B; g} \\
\\
\text{CONST} \qquad \qquad \text{VAR} \qquad \qquad \text{FUN} \\
\hline
c; \gamma; A \Downarrow c; A; \mathbf{1} \qquad \frac{\gamma(x) = V}{x; \gamma; A \Downarrow V; A; \mathbf{1}} \qquad \frac{}{\lambda x. M; \gamma; A \Downarrow (x.M)\{\gamma\}; A; \mathbf{1}} \\
\\
\text{APPCLO} \\
\frac{M; \gamma; A \Downarrow (x.S)\{\sigma\}; A'; g_1 \quad N; \gamma; B \Downarrow V; B'; g_2 \quad S; A' \otimes B'; \sigma[x \mapsto V] \Downarrow W; C; g_3}{M N; \gamma; A \otimes B \Downarrow W; C; (g_1 \otimes g_2) \oplus g_3 \oplus \mathbf{1}} \\
\\
\text{APPCONST} \\
\frac{M; \gamma; A \Downarrow c; A'; g_1 \quad N; \gamma; B \Downarrow V; B'; g_2 \quad (V, A' \otimes B') \delta_c (W, C)}{M N; \gamma; A \otimes B \Downarrow W; C; (g_1 \otimes g_2) \oplus \mathbf{1}}
\end{array}$$

Figure 5: Cost Semantics of SPAL $_{\mathbb{M}}$

Given a map  $f : X \rightarrow \mathbb{M}$ , we define  $\text{dom}(f) \triangleq \{x \mid f(x) \neq \perp\}$ .

**Definition 4.** The *finite map* PCM over a set  $X$  and a monoid  $\mathbb{M}$ , denoted  $\text{FMap}(X, \mathbb{M})$ , is defined as  $(\{f \in \prod_{x \in X} \mathbb{M} \mid \text{dom}(f) \text{ finite}\} \cup \perp, *, x \mapsto \epsilon_{\mathbb{M}}, \perp)$  where  $f \otimes g$  is defined as  $x \mapsto f(x) \cdot g(x)$  only when  $f(x) \cdot g(x) \neq \perp_{\mathbb{M}}$  for all  $x \in \text{dom}(f) \cup \text{dom}(g)$ .

**Definition 5.** The *plain heap* PCM over a set  $X$ , denoted  $\text{Heap}(X)$ , is  $\text{FMap}(\text{Loc}, \text{Ex}(X))$ .

It is worth unfolding the last definition. The set of non- $\perp$  elements of  $\text{Heap}(X)$  is isomorphic to maps  $\text{Loc} \rightarrow_{\text{fin}} X$ . The composite  $f \otimes g$  is defined only when  $\text{dom}(f) \cap \text{dom}(g) = \emptyset$ , in which case it sends  $\ell \in \text{dom}(f) \cup \text{dom}(g)$  to  $f(\ell)$  if  $\ell \in \text{dom}(g)$  and to  $g(\ell)$  otherwise.

We write  $\ell \hookrightarrow x$  for the element of  $\text{Heap}(X)$  sending  $\ell$  to  $x$  and  $\ell' \neq \ell$  to  $\epsilon$ .

### 3.2 The Generic Calculus

We now fix a PCM  $\mathbb{M}$  and use it to build our calculus, the Stateful Parallel Applicative Calculus over  $\mathbb{M}$ , which we abbreviate SPAL $_{\mathbb{M}}$ . We write  $A, B, C$  for the elements of  $\mathbb{M}$ .

This calculus and PAL share the same syntax. However, its language of constants will generally feature items related to  $\mathbb{M}$ , or contain operations able to act upon an element of the monoid. Its cost semantics, again expressed as a big-step judgment  $M; \sigma; A \Downarrow V; B; g$ , modifies that of PAL to express the relation between the initial state  $A$  and the final state  $B$  reached after evaluation has terminated. The rules of this semantics are given in Figure 5.

- Rules CONST, VAR, and FUN do not affect the state.
- Parallelism is handled via the monoid law: to apply rules APPCLO and APPCONST, one has to split the input state into distinct parts, but also to recombine the final states.
- Constants may alter the current state. It is convenient to generalize  $\delta$  from a partial function to a family of relations  $(\delta_c \subseteq (\text{Val} \times \mathbb{M})^2)_{c \in \text{Const}}$ .

This means that the only way to do a primitive state modification is via the proper constants. To illustrate this fact, let us describe a useful family of constants for SPAL $_{\text{Heap}(\mathbb{Z})}$ .

$$c ::= i \mid \dots \mid \ell \mid \text{alloc} \mid \text{read} \mid \text{write} \mid \text{write}_{\ell}$$

Their semantics is unsurprising.

$$\begin{aligned}
& (i, A) \delta_{\text{alloc}} (\ell, A \otimes \ell \hookrightarrow i) \\
& (\ell, A \otimes \ell \hookrightarrow i) \delta_{\text{read}} (i, A \otimes \ell \hookrightarrow i) \\
& (\ell, A) \delta_{\text{write}} (\text{write}_\ell, A) \\
& (i, A \otimes \ell \hookrightarrow j) \delta_{\text{write}_\ell} (i, A \otimes \ell \hookrightarrow i)
\end{aligned}$$

One should try to play with this definition. How does it differ from a language like Cilk?

### 3.3 A Cilk-like Instance

To allow for concurrent reads yet disallow races, we use *fractional* permissions, another well-known idea from separation logic [Boyland, 2003, Jung et al., 2015].

**Definition 6.** The PCM of *fractional permissions* over a set  $X$ , denoted  $\text{Frac}(X)$ , is  $((\mathbb{Q} \cap (0, 1] \times X) \cup \{\epsilon, \perp\}, \cdot, \epsilon, \perp)$ , where the monoid law is defined only by

$$(q, x) \odot (q', x') \triangleq (q + q', x) \text{ if } x = x' \text{ and } q + q' \leq 1.$$

**Definition 7.** The *fractional heap* PCM over a set  $X$ , denoted  $\text{Frac}(X)$ , is  $\text{FMap}(\text{Loc}, \text{Frac}(X))$ .

Again, let us unfold this definition. The set of non- $\perp$  elements of  $\text{Frac}(X)$  is isomorphic to maps  $\text{Loc} \rightarrow_{\text{fin}} (\mathbb{Q} \cap (0, 1]) \times X$ . The composite  $f \otimes g$  is defined only when for all  $\ell \in \text{dom}(f) \cup \text{dom}(g)$ ,  $f(\ell) = (q, x)$  and  $g(\ell) = (q', x)$  and  $q + q' \leq 1$ , in which case it sends  $\ell$  to  $(q + q', x)$ .

We write  $\ell \xrightarrow{q} x$  for the element of  $\text{Heap}(X)$  sending  $\ell$  to  $(q, x)$  and  $\ell' \neq \ell$  to  $\epsilon$ .

The semantics of constants only allows writing to a location  $\ell$  that is fully owned, in the sense that  $\ell \xrightarrow{1} i$  for some  $i$ . A location is fully owned by the code that created it.

$$\begin{aligned}
& (i, A) \delta_{\text{alloc}} (\ell, A \otimes \ell \xrightarrow{1} i) \\
& (\ell, A \otimes \ell \xrightarrow{q} i) \delta_{\text{read}} (i, A \otimes \ell \xrightarrow{q} i) \\
& (\ell, A) \delta_{\text{write}} (\text{write}_\ell, A) \\
& (i, A \otimes \ell \xrightarrow{1} j) \delta_{\text{write}_\ell} (i, A \otimes \ell \xrightarrow{1} i)
\end{aligned}$$

## References

- Guy Blelloch and John Greiner. Parallelism in Sequential Functional Languages. In *Functional Programming Languages and Computer Architecture (FPCA '95)*. ACM, 1995. URL <http://www.cs.cmu.edu/~blelloch/papers/BG95.pdf>.
- John Boyland. Checking interference with fractional permissions. Springer, 2003. URL <http://www.cs.uwm.edu/faculty/boyland/papers/permissions.pdf>.
- Matteo Frigo, Charles Leiserson, and Keith Randall. The Implementation of the Cilk-5 Multithreaded Language. In *Programming Language Design and Implementation (PLDI'98)*. ACM, 1998. URL <http://supertech.csail.mit.edu/papers/cilk5.pdf>.
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. ACM, 2015. URL <https://iris-project.org/pdfs/2015-popl-iris1-final.pdf>.