

# Assignment #5, Update 5: Abstract Machines and Parallelism

15-312: Principles of Programming Languages

Out: Tuesday, April 8, 2014

Due: Thursday, April 24, 2014 1:29PM

## Introduction

In this assignment, we will develop a language with the parallel programming primitive of *sequences* and we will write programs in that language. The work will be divided into two stages. In the first stage, we will write a *control stack machine* that controls the evaluation of the sequential portion of the language. In the second stage, we will extend this to implement the parallel constructs of the language.

## Submission

We will collect *exactly* the following files from the `/afs/andrew/course/15/312/` directory:

```
handin/<yourandrewid>/assn5/assn5.pdf
handin/<yourandrewid>/assn5/dynamics.sml
handin/<yourandrewid>/assn5/scheduler.sml
handin/<yourandrewid>/assn5/mergesort.par
```

Make sure that your files have the right names (especially `assn5.pdf`!) and are in the correct directories.

## 1 A Data Parallel Language

The language of study in this assignment features a new type,  $\text{seq}(\tau)$ , for sequences containing elements of type  $\tau$ . When we create a new sequence, the contents of the resulting array are evaluated in parallel. This effectively allows the programmer to spawn an unbounded number of parallel tasks simultaneously. The remaining elements of the language are sequentially evaluated in left-to-right order.

Types and expressions are given in Figure 1. This language departs a bit from the familiar PCF language, introducing a number of integer operations as primitives and defining operations that work on those values. There are four operations for manipulating sequences:

- $\text{tab}[\tau](x.e; e_n)$  – If  $e_n$  evaluates to  $\text{num}[n]$  where  $n > 0$ , this expression creates a new sequence of length  $e_n$  by evaluating  $[0/x]e, \dots, [n-1/x]e$ .
- $\text{sub}(e; e_n)$  – If  $e_n$  evaluates to  $\bar{i}$  and  $e$  evaluates to  $\text{seq}(e_0, \dots, e_{n-1})$ , this expression evaluates to  $e_i$ .
- $\text{len}(e)$  – Evaluates to the length of the sequence that  $e$  evaluates to.

<b>Sort</b>		<b>Abstract Form</b>	<b>Concrete Form</b>
Type	$\tau ::=$	int bool parr( $\tau_1; \tau_2$ ) prod( $\tau_1; \tau_2$ ) seq( $\tau$ )	int bool $\tau_1 \rightarrow \tau_2$ $\tau_1 \times \tau_2$ seq( $\tau$ )
Comparisons	$cmp ::=$	eq neq leq geq	= $\neq$ $\leq$ $\geq$
Integer ops	$op ::=$	plus minus times div mod	+ - $\times$ $\div$ %
Exp	$e ::=$	$x$ let( $e_0; x.e_1$ ) lam[ $\tau$ ]( $x.e$ ) ap( $e_1; e_2$ ) fix[ $\tau$ ]( $x.e$ ) pair( $e_1; e_2$ ) pr[l]( $e$ ) pr[r]( $e$ )  num[ $n$ ] cmp[ $cmp$ ]( $e_1; e_2$ ) calc[ $op$ ]( $e_1; e_2$ ) true false if( $e; e_t; e_f$ )  seq( $e_0, \dots, e_{n-1}$ ) tab[ $\tau$ ]( $x.e; e_n$ ) sub( $e; e_n$ ) len( $e$ ) showt( $e; x.e_1; y_l.y_r.e_2$ )	$x$ let $x = e_0$ in $e_1$ fn ( $x:\tau$ ) $e$ $e_1$ ( $e_2$ ) fix $x:\tau$ is $e$ $\langle e_1, e_2 \rangle$ $e \cdot l$ $e \cdot r$  $\bar{n}$ $e_1$ $cmp$ $e_2$ $e_1$ $op$ $e_2$ true false if $e$ then $e_t$ else $e_f$  seq( $e_0, \dots, e_{n-1}$ ) $[e \mid x < e_n]$ $e[e_n]$ len( $e$ ) showt $e \{elt\ x \Rightarrow e_1 \mid node\ y_l\ y_r \Rightarrow e_2\}$

Figure 1: Expressions and types in our parallel language

- $\text{showt}(e; x.e_1; y_l.y_r.e_2)$  – If  $e$  evaluates to a sequence with one element  $v$ , then this evaluates  $[v/x]e_1$ . If  $e$  evaluates to a sequence with more than one element, then the sequence gets split into two equal sized arrays, the first half-sequence (which may be one element smaller than the second half-sequence) is substituted for  $y_l$  in  $e_2$ , and the second half-sequence is substituted for  $y_r$  in  $e_2$ .

To help you test your sequential implementation, we will allow you to explicitly write sequences in the dynamic semantics, but the statics of our language will require that sequences start out containing only values (more on this in a moment). Therefore, it would be better to think of them as part of the internal language, not available to the programmer. This is essentially how we thought of  $\text{num}[n]$  in Dynamic PCF.

## 1.1 Values

There's nothing particularly surprising about the values for this language, except that we're using lazy pairs this time around:

$$\frac{}{\text{lam}[\tau](x.e) \text{ val}} (\text{lam-val}) \quad \frac{}{\text{pair}(e_1; e_2) \text{ val}} (\text{pair-val}) \quad \frac{}{\text{num}[n] \text{ val}} (\text{num-val})$$

$$\frac{}{\text{true val}} (\text{true-val}) \quad \frac{}{\text{false val}} (\text{false-val}) \quad \frac{e_0 \text{ val} \quad \dots \quad e_{n-1} \text{ val}}{\text{seq}(e_0, \dots, e_{n-1}) \text{ val}} (\text{seq-val})$$

## 1.2 Statics

The static semantics for variables  $x$ ,  $\text{let}(e_0; x.e_1)$ ,  $\text{lam}[\tau](x.e)$ ,  $\text{ap}(e_1; e_2)$ , and  $\text{fix}[\tau](x.e)$ ,  $\text{pair}(e_1; e_2)$ ,  $\text{pr}[l](e)$ , and  $\text{pr}[r](e)$  remain the same as they have been in previous assignments, and in Homework 2 in particular. Rules for numbers and booleans are straightforward:

$$\frac{}{\Gamma \vdash \text{num}[n] : \text{int}} (\text{int-}I) \quad \frac{\Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash \text{cmp}[cmp](e_1; e_2) : \text{bool}} (\text{int-}E_1)$$

$$\frac{\Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash \text{calc}[op](e_1; e_2) : \text{int}} (\text{int-}E_2) \quad \frac{}{\Gamma \vdash \text{true} : \text{bool}} (\text{bool-}I_1) \quad \frac{}{\Gamma \vdash \text{false} : \text{bool}} (\text{bool-}I_1)$$

$$\frac{\Gamma \vdash e : \text{bool} \quad \Gamma \vdash e_t : \tau \quad \Gamma \vdash e_f : \tau}{\Gamma \vdash \text{if } e \text{ then } e_t \text{ else } e_f : \tau} (\text{bool-}E)$$

By this point in the semester, the following rules for dealing with sequences will hopefully also be straightforward, with the exception that the value judgment is now being referenced as part of the *statics* in rule  $\text{seq-}I_1$ . In fact, in the abstract machine dynamics we give in Section 2, we won't mention the value judgment at all: it moves completely to the statics in the abstract machine presentation used in this assignment.

$$\begin{array}{c}
\frac{\Gamma \vdash e_0 : \tau \quad e_0 \text{ val} \quad \dots \quad \Gamma \vdash e_{n-1} : \tau \quad e_{n-1} \text{ val}}{\Gamma \vdash \text{seq}(e_0, \dots, e_{n-1}) : \text{seq}(\tau)} (\text{seq-}I_1) \\
\\
\frac{\Gamma, x : \text{int} \vdash e : \tau \quad \Gamma \vdash e_n : \text{int}}{\Gamma \vdash \text{tab}[\tau](x.e; e_n) : \text{seq}(\tau)} (\text{seq-}I_2) \\
\\
\frac{\Gamma \vdash e : \text{seq}(\tau) \quad \Gamma \vdash e_n : \text{int}}{\Gamma \vdash \text{sub}(e; e_n) : \tau} (\text{seq-}E_1) \qquad \frac{\Gamma \vdash e : \text{seq}(\tau)}{\Gamma \vdash \text{len}(e) : \text{int}} (\text{seq-}E_2) \\
\\
\frac{\Gamma \vdash e : \text{seq}(\tau) \quad \Gamma, x : \tau \vdash e_1 : \tau' \quad \Gamma, y_1 : \text{seq}(\tau), y_2 : \text{seq}(\tau) \vdash e_2 : \tau'}{\Gamma \vdash \text{showt}(e; x.e_1; y_1.y_2.e_2) : \tau'} (\text{seq-}E_3)
\end{array}$$

## 2 Abstract Machines

Abstract machines give a picture of execution that is closer to a realistic implementation than the structural dynamics that have prevailed so far in this course.<sup>1</sup> We are using abstract machines in this assignment because they give us a convenient way of thinking about pausing a sequential process of evaluation, removing a subterm in order to evaluate that subterm in parallel, and then resuming the evaluation of the original term.

Abstract machines record the path from an outermost term to a subterm as a *stack*  $k$ , which is made up of a list of *frames*  $f$ .

$$\begin{array}{l}
k ::= \epsilon \mid k; f \\
f ::= \text{let}(\square; x.e_1) \mid \text{ap}(\square; e_2) \mid \text{ap}(v_1; \square) \mid \text{pr}[l](\square) \mid \text{pr}[r](\square) \mid \\
\quad \text{cmp}[cmp](\square; e_2) \mid \text{cmp}[cmp](v_1; \square) \mid \text{calc}[op](\square; e_2) \mid \text{calc}[op](v_1; \square) \mid \text{if}(\square; e_t; e_f) \mid \\
\quad \text{tab}[\tau](x.e; \square) \mid \text{sub}(\square; e_n) \mid \text{sub}(v; \square) \mid \text{len}(\square) \mid \text{showt}(\square; x.e_1; y_1.y_2.e_2)
\end{array}$$

There are three states of the abstract machine: a state  $s$  can be either an expression evaluating on top of a stack ( $k \triangleright e$ ), a value returning to a stack ( $k \triangleleft v$ ), or a failure state ( $k \blacktriangleleft$ ). We'll add a fourth, ( $k$  awaiting  $[a_0, \dots, a_{n-1}]$ ), in the next section.

### 2.1 Dynamics

The transition relation is an inductive definition, the same as we've seen before for structural dynamics or evaluation dynamics. Presented as an inductive definition, all the rules would look like this:

$$\begin{array}{c}
\frac{}{k \triangleright \text{let}(e_0; x.e_1) \mapsto k; \text{let}(\square; x.e_1) \triangleright e_0} (\text{let-}S_1) \\
\\
\frac{0 \leq i < n}{k; \text{sub}(\text{seq}(v_0, \dots, v_{n-1}); \square) \triangleleft \text{num}[i] \mapsto k \triangleleft v_i} (\text{sub-}S_3)
\end{array}$$

<sup>1</sup>Do remember, though, that this is still a *semantics*, a specification of how implementations should behave and not a guide to implementation. There are elements to this presentation, like the presence of substitution, that you wouldn't want in a realistic implementation. Additionally, the type annotation on the tabulate command  $\text{tab}[\tau](x.e; e_n)$  is present in our semantics but is not present in our SML implementation, where we will write  $\text{Exp.Tab}((x, e), en)$ . As with the type annotation on  $\text{dc1}[\tau](e; a.e')$  in the last assignment

This would take up a lot of space, and because almost none of the rules of the dynamics have any premises, we'll instead present the dynamics in a compressed format, writing premises off to the side where applicable.

### Core (PCF-like) language

$$\begin{array}{l}
(\text{let-}S_1) \quad k \triangleright \text{let}(e_0; x.e_1) \quad \mapsto \quad k; \text{let}(\square; x.e_1) \triangleright e_0 \\
(\text{let-}S_2) \quad k; \text{let}(\square; x.e_1) \triangleleft v_0 \quad \mapsto \quad k \triangleright [v_0/x]e_1 \\
(\text{lam-}S) \quad k \triangleright \text{lam}[\tau](x.e) \quad \mapsto \quad k \triangleleft \text{lam}[\tau](x.e) \\
(\text{ap-}S_1) \quad k \triangleright \text{ap}(e_1; e_2) \quad \mapsto \quad k; \text{ap}(\square; e_2) \triangleright e_1 \\
(\text{ap-}S_2) \quad k; \text{ap}(\square; e_2) \triangleleft v_1 \quad \mapsto \quad k; \text{ap}(v_1; \square) \triangleright e_2 \\
(\text{ap-}S_3) \quad k; \text{ap}((\text{fn}(x:\tau) e_0); \square) \triangleleft v_2 \quad \mapsto \quad k \triangleright [v_2/x]e_0 \\
(\text{fix-}S) \quad k \triangleright \text{fix}[\tau](x.e) \quad \mapsto \quad k \triangleright [\text{fix}[\tau](x.e)/x]e \\
(\text{pair-}S) \quad k \triangleright \text{pair}(e_1; e_2) \quad \mapsto \quad k \triangleleft \text{pair}(e_1; e_2) \\
(\text{prl-}S_1) \quad k \triangleright \text{pr}[l](e) \quad \mapsto \quad k; \text{pr}[l](\square) \triangleright e \\
(\text{prl-}S_2) \quad k; \text{pr}[l](\square) \triangleleft \langle e_1, e_2 \rangle \quad \mapsto \quad k \triangleright e_1 \\
(\text{prr-}S_1) \quad k \triangleright \text{pr}[r](e) \quad \mapsto \quad k; \text{pr}[r](\square) \triangleright e \\
(\text{prr-}S_2) \quad k; \text{pr}[r](\square) \triangleleft \langle e_1, e_2 \rangle \quad \mapsto \quad k \triangleright e_2
\end{array}$$

**Integers and Booleans** When we write  $(k \triangleleft \text{num}[n_1 \text{ op } n_2])$  or  $(k \triangleleft n_1 \text{ cmp } n_2)$ , we mean to perform the actual primitive operation on numbers. For example, the state  $(k; \text{calc}[\text{times}](\text{num}[3]; \square) \triangleleft \text{num}[4])$  will step to  $(k \triangleleft \text{num}[12])$  and the state  $(k; \text{cmp}[\text{neq}](\text{num}[3]; \square) \triangleleft \text{num}[4])$  will step to  $(k \triangleleft \text{true})$ .

$$\begin{array}{l}
(\text{num-}S) \quad k \triangleright \text{num}[n] \quad \mapsto \quad k \triangleleft \text{num}[n] \\
(\text{cmp-}S_1) \quad k \triangleright \text{cmp}[\text{cmp}](e_1; e_2) \quad \mapsto \quad k; \text{cmp}[\text{cmp}](\square; e_2) \triangleright e_1 \\
(\text{cmp-}S_2) \quad k; \text{cmp}[\text{cmp}](\square; e_2) \triangleleft v_1 \quad \mapsto \quad k; \text{cmp}[\text{cmp}](v_1; \square) \triangleright e_2 \\
(\text{cmp-}S_3) \quad k; \text{cmp}[\text{cmp}](\text{num}[n_1]; \square) \triangleleft \text{num}[n_2] \quad \mapsto \quad k \triangleleft n_1 \text{ cmp } n_2 \\
(\text{calc-}S_1) \quad k \triangleright \text{calc}[\text{op}](e_1; e_2) \quad \mapsto \quad k; \text{calc}[\text{op}](\square; e_2) \triangleright e_1 \\
(\text{calc-}S_2) \quad k; \text{calc}[\text{op}](\square; e_2) \triangleleft v_1 \quad \mapsto \quad k; \text{calc}[\text{op}](v_1; \square) \triangleright e_2 \\
(\text{calc-}S_3) \quad k; \text{calc}[\text{op}](\text{num}[n_1]; \square) \triangleleft \text{num}[n_2] \quad \mapsto \quad k \triangleleft \text{num}[n_1 \text{ op } n_2] \quad (\text{if } n_1 \text{ op } n_2 \text{ is defined}) \\
(\text{calc-}F) \quad k; \text{calc}[\text{op}](\text{num}[n_1]; \square) \triangleleft \text{num}[n_2] \quad \mapsto \quad k \blacktriangleleft \quad (\text{if } n_1 \text{ op } n_2 \text{ is not defined}) \\
(\text{true-}S) \quad k \triangleright \text{true} \quad \mapsto \quad k \triangleleft \text{true} \\
(\text{false-}S) \quad k \triangleright \text{false} \quad \mapsto \quad k \triangleleft \text{false} \\
(\text{if-}S_1) \quad k \triangleright \text{if}(e; e_t; e_f) \quad \mapsto \quad k; \text{if}(\square; e_t; e_f) \triangleright e \\
(\text{if-}S_2) \quad k; \text{if}(\square; e_t; e_f) \triangleleft \text{true} \quad \mapsto \quad k \triangleright e_t \\
(\text{if-}S_3) \quad k; \text{if}(\square; e_t; e_f) \triangleleft \text{false} \quad \mapsto \quad k \triangleright e_f
\end{array}$$

**Sequences** We are only using sequences to describe the *sequential* evaluation of our language; this means that if we look exclusively at this section then the semantics of sequences seem to be missing a critical rule for tabulation  $\text{tab}[\tau](x.e; e_n)$ .

$$\begin{array}{l}
(\text{seq-}S) \quad k \triangleright \text{seq}(v_0, \dots, v_{n-1}) \quad \mapsto \quad k \triangleleft \text{seq}(v_0, \dots, v_{n-1}) \\
(\text{tab-}S) \quad k \triangleright \text{tab}[\tau](x.e; e_n) \quad \mapsto \quad k; \text{tab}[\tau](x.e; \square) \triangleright e_n \\
(\text{tab-}F) \quad k; \text{tab}[\tau](x.e; \square) \triangleleft \text{num}[n] \quad \mapsto \quad k \blacktriangleleft \quad (\text{if } n \leq 0) \\
(\text{sub-}S_1) \quad k \triangleright \text{sub}(e; e_n) \quad \mapsto \quad k; \text{sub}(\square; e_n) \triangleright e \\
(\text{sub-}S_2) \quad k; \text{sub}(\square; e_n) \triangleleft v \quad \mapsto \quad k; \text{sub}(v; \square) \triangleright e_n \\
(\text{sub-}S_3) \quad k; \text{sub}(\text{seq}(v_0, \dots, v_{n-1}); \square) \triangleleft \text{num}[i] \quad \mapsto \quad k \triangleleft v_i \quad (\text{if } 0 \leq i < n)
\end{array}$$

$$\begin{array}{ll}
(\text{sub-}F) & k; \text{sub}(\text{seq}(v_0, \dots, v_{n-1}); \square) \triangleleft \text{num}[i] \quad \mapsto \quad k \blacktriangleleft \quad (\text{if } 0 > i \text{ or } i \geq n) \\
(\text{len-}S_1) & k \triangleright \text{len}(e) \quad \mapsto \quad k; \text{len}(\square) \triangleright e \\
(\text{len-}S_2) & k; \text{len}(\square) \triangleleft \text{seq}(e_0, \dots, e_{n-1}) \quad \mapsto \quad k \triangleleft \text{num}[n] \\
(\text{showt-}S_1) & k \triangleright \text{showt}(e; x.e_1; y_1.y_2.e_2) \quad \mapsto \quad k; \text{showt}(\square; x.e_1; y_1.y_2.e_2) \triangleright e \\
(\text{showt-}S_2) & k; \text{showt}(\square; x.e_1; y_1.y_2.e_2) \triangleleft \text{seq}(v) \quad \mapsto \quad k \triangleright [v/x]e_1 \\
(\text{showt-}S_3) & k; \text{showt}(\square; x.e_1; y_1.y_2.e_2) \triangleleft \text{seq}(v_0, \dots, v_{n-1}) \quad (\text{if } n > 1 \text{ and } k = \lfloor n/2 \rfloor) \\
& \mapsto \quad k \triangleright [\text{seq}(v_0, \dots, v_{k-1})/y_1][\text{seq}(v_k, \dots, v_{n-1})/y_2]e_2
\end{array}$$

## 2.2 Statics

There are three judgments giving the statics of states  $s$ . The judgment  $s : \tau$  expresses that, if the state  $s$  ever returns a value to the empty stack  $s = (\epsilon \triangleleft v)$ , then  $v$  will be a value of type  $\tau$ . In support of this judgment, we need two more judgments  $k : \tau \Rightarrow \tau'$  and  $f : \tau \Rightarrow \tau'$  which express that the stack  $k$  and the frame  $f$  transform values of type  $\tau$  to values of type  $\tau'$ .

$$\begin{array}{ccc}
\frac{k : \tau' \Rightarrow \tau \quad \emptyset \vdash e : \tau'}{(k \triangleright e) : \tau} (S_1) & \frac{k : \tau' \Rightarrow \tau \quad \emptyset \vdash v : \tau' \quad v \text{ val}}{(k \triangleleft v) : \tau} (S_2) & \frac{k : \tau' \Rightarrow \tau}{(k \blacktriangleleft) : \tau} (S_3) \\
\frac{}{\epsilon : \tau \Rightarrow \tau} (K_1) & \frac{k : \tau' \Rightarrow \tau_{fin} \quad f : \tau \Rightarrow \tau'}{k; f : \tau \Rightarrow \tau_{fin}} (K_2) &
\end{array}$$

The value judgment shows up in the typing rules for frames as it did in the statics rule for  $\text{seq}(e_0, \dots, e_{n-1})$ :

$$\begin{array}{c}
\frac{x : \tau \vdash e_1 : \tau'}{\text{let}(\square; x.e_1) : \tau \Rightarrow \tau'} (\text{let-}F) \\
\\
\frac{\emptyset \vdash e_2 : \tau'}{\text{ap}(\square; e_2) : \tau' \rightarrow \tau \Rightarrow \tau} (\text{apA-}F) \qquad \frac{\emptyset \vdash v_1 : \tau \rightarrow \tau' \quad v_1 \text{ val}}{\text{ap}(v_1; \square) : \tau \Rightarrow \tau'} (\text{apB-}F) \\
\\
\frac{}{\text{pr}[l](\square) : \tau_1 \times \tau_2 \Rightarrow \tau_1} (\text{prl-}F) \qquad \frac{}{\text{pr}[r](\square) : \tau_1 \times \tau_2 \Rightarrow \tau_2} (\text{prrr-}F) \\
\\
\frac{\emptyset \vdash e_2 : \text{int}}{\text{cmp}[cmp](\square; e_2) : \text{int} \Rightarrow \text{bool}} (\text{cmpA-}F) \qquad \frac{v_1 \text{ val} \quad \emptyset \vdash v_1 : \text{int}}{\text{cmp}[cmp](v_1; \square) : \text{int} \Rightarrow \text{bool}} (\text{cmpB-}F) \\
\\
\frac{\emptyset \vdash e_2 : \text{int}}{\text{calc}[op](\square; e_2) : \text{int} \Rightarrow \text{int}} (\text{calcA-}F) \qquad \frac{v_1 \text{ val} \quad \emptyset \vdash v_1 : \text{int}}{\text{calc}[op](v_1; \square) : \text{int} \Rightarrow \text{int}} (\text{calcB-}F) \\
\\
\frac{\emptyset \vdash e_t : \tau \quad \emptyset \vdash e_f : \tau}{\text{if}(\square; e_t; e_f) : \text{bool} \Rightarrow \tau} (\text{if-}F)
\end{array}$$

**Task 2.1** (5%). The statics for frames in the sequence-manipulating fragment of the language – that is, the frames  $\text{tab}[\tau](x.e; \square)$ ,  $\text{sub}(\square; e_n)$ ,  $\text{sub}(v; \square)$ ,  $\text{len}(\square)$ , and  $\text{showt}(\square; x.e_1; y_1.y_2.e_2)$  – have been omitted. Give these rules such that progress and preservation (discussed in the next section) will hold.  $\square$

**Task 2.2** (5%). We will not require you to implement the statics of frames and control machine states in this homework. Implementing these statics, however, gives you a powerful invariant for checking the well-formedness of intermediate states. This might help you find errors in the implementation of the dynamics (which will show up as failures of preservation).

In this task, we will ask you to think through one aspect of this optional implementation. If you do set `TypeChecker.canCheckStates` to `true` in `typechecker.sml`, then the top-level interpreter's `step` command will perform typechecking on intermediate states. In order to correctly implement `TypeChecker.checkstate: State.t -> Type.t`, you will need a helper function that handles the statics of frames. **[Bug fixing update: this function only needs to correctly handle states ( $k \triangleright e$ ) and ( $k \triangleleft v$ ), not failure states ( $k \blacktriangleleft$  ).]** For each of the following options, explain whether or not it would be implementable and usable as a helper function to the `TypeChecker.checkstate` function.

- `frametype: Frame.t -> Tp.t * Tp.t -> unit`  
where (`frametype f ( $\tau, \tau'$ )`) will run without error iff  $f : \tau \Rightarrow \tau'$
- `frametype: Frame.t -> Tp.t -> Tp.t`  
where (`frametype f  $\tau$` ) =  $\tau'$  iff  $f : \tau \Rightarrow \tau'$
- `frametype: Frame.t -> Tp.t -> Tp.t`  
where (`frametype f  $\tau'$` ) =  $\tau$  iff  $f : \tau \Rightarrow \tau'$
- `frametype: Frame.t -> Tp.t * Tp.t`  
where (`frametype f`) = ( $\tau, \tau'$ ) iff  $f : \tau \Rightarrow \tau'$

If you indicate that a particular version of `frametype` would work, then give pseudo-code or code for indicating how it would be used to check the type of states – describing the ( $k \triangleright e$ ) case will be sufficient.

### 2.3 Safety

We'd like to say that the standard progress and preservation formulation of our language holds:

**Theorem 1** (Preservation). *If  $s : \tau$  and  $s \mapsto s'$  then  $s' : \tau$ .*

**Theorem 2** (Progress). *If  $s : \tau$  then either there exists an  $s'$  such that  $s \mapsto s'$  or  $s$  final.*

However, progress only holds true if we carefully formulate the  $s$  final judgment.

**Task 2.3** (5%). Give a definition of the  $s$  final judgment such that the progress theorem holds. Make sure to account for all the well-formed states that cannot step according to the dynamics we've presented so far. Your final judgment shouldn't declare any states to be final if they can take a step, so you can't just write

$$\frac{}{s \text{ final}} \text{ (this is bogus)}$$

even though this would certainly ensure that the progress theorem holds.

**Task 2.4** (10%). Consider each of the following changes being made to our language definition *independently and one at a time, not cumulatively*.

1. Adding a  $e_1$  val premise to `pair-val`.

2. Removing the  $e_i$  val premises from  $\text{seq-}I_1$ .
3. Making the right-hand side of  $\text{pr1-}S_2$  ( $k \triangleleft e_1$ ) instead of ( $k \triangleright e_1$ ).
4. Making the right-hand side of  $\text{pr1-}S_2$  ( $k \triangleright e_2$ ) instead of ( $k \triangleright e_1$ ).
5. Removing the premise from  $S_3$ .
6. Removing the  $v_1$  val premise from  $\text{apB-}F$ .
7. Adding a  $e_2$  val premise to  $\text{cmpA-}F$ .

For each potential change, do one of the following:

- State that it does not break preservation and progress for the language. If you feel it breaks the language in some way that isn't captured by progress and preservation, you can state this, but it isn't necessary to do so.
- State that it breaks preservation and give an  $s$  such that  $s : \tau$  and  $s \mapsto s'$  but it is not the case  $s' : \tau$ . (If it will not be obvious to us why  $s : \tau$  or why  $s \mapsto s'$ , make this clear.)
- State that it breaks progress and give an  $s$  such that  $s : \tau$  but  $s$  neither steps nor takes a value. (If it will not be obvious to us why  $s : \tau$ , make this clear.)

## 2.4 Implementation

The top-level for sequential evaluation can be invoked by running `Seq.repl()`.

**Task 2.5 (15%).** Implement the sequential dynamics of control stacks in `dynamics.sml` as the function `trystep`. Your code should be clear and well-organized in addition to being correct. Running `trystep s` should return:

- `Steps s'` if  $s \mapsto s'$ ,
- `Value v` if  $s = (\epsilon \triangleleft v)$ ,
- `Tabulate (n, k, (x, e))` if  $s = (k; \text{tab}[\tau](x.e; \square) \triangleleft \text{num}[n])$ , and
- `Err` if  $s = (k \blacktriangleleft)$ . □

Here is an example of running the sequential abstract machine toplevel interpreter:

```
$ rlwrap sml -m sources.cm
- Seq.repl ();
->step 1+2 == 5;
Statics: exp has type bool
--> emp; Cmp[Eq] (-, Num[5])
      |> Calc[Plus] (Num[1], Num[2])
->step;
--> emp; Cmp[Eq] (-, Num[5]); Calc[Plus] (-, Num[2])
      |> Num[1]
->step;
--> emp; Cmp[Eq] (-, Num[5]); Calc[Plus] (-, Num[2])
```



```

    <| Num[1]
->eval;
  False VAL
->eval let val S = seq(8,6,4,2,0) in S[0] + S[3] end;
Statics: exp has type int
  Num[10] VAL

```

### 3 Parallel dynamics

We describe the global state of a parallel computation as a mapping from symbols  $a_i$  to abstract machines  $s_i$ . These abstract machines collectively capture the steps that need to happen to finish the global computation. These global states have the form  $\nu\Sigma\{\mu\}$ , where the signature  $\Sigma$  associates types to each symbol:

$$\Sigma = a_1 \sim \tau_1, \dots, a_n \sim \tau_n$$

and  $\mu$  associates an abstract machine with each symbol:

$$\mu = a_1 \hookrightarrow s_1 \otimes \dots \otimes a_n \hookrightarrow s_n$$

We treat the symbols  $a_i$  as being *bound* in  $\Sigma$  and therefore subject to renaming, though this renaming will not be visible in our implementation – as with variables (or assignables in Algol) we will make sure each newly generated symbol is given a fresh name when we call the `Sym.newsym` function.

A local transition deals only with a restricted portion of the memory. The easy case of a local transition is a state that is able to take a step according to sequential semantics of control stacks:

$$\frac{s \mapsto s'}{\nu a \sim \tau \{a \hookrightarrow s\} \mapsto_{loc} \nu a \sim \tau \{a \hookrightarrow s'\}} \text{(locstep-step)}$$

A different local transition finally gives us a way to consider the previously-unhandled abstract machine state  $(k; \text{tab}[\tau](x.e; \square) \triangleleft \text{num}[n])$ , which needs to spawn  $n$  computations that will be evaluated in parallel.

$$\frac{n > 0}{\nu a \sim \tau' \{a \hookrightarrow (k; \text{tab}[\tau](x.e; \square) \triangleleft \text{num}[n])\} \mapsto_{loc} \nu a \sim \tau', a_0 \sim \tau, \dots, a_{n-1} \sim \tau \{a \hookrightarrow (k \text{ awaiting}[a_0, \dots, a_{n-1}])\} \otimes \mu} \text{(locstep-fork)}$$

where  $\mu = a_0 \hookrightarrow (\epsilon \triangleright [\text{num}[0]/x]e) \otimes \dots \otimes a_{n-1} \hookrightarrow (\epsilon \triangleright [\text{num}[n-1]/x]e)$ . This rule also introduces a new form of state  $s = (k \text{ awaiting}[a_0, \dots, a_{n-1}])$ , which represents a stack that expects to receive a length- $n$  sequence.

A sequence can be created and returned to the stack when the computations associated with symbols  $a_0, \dots, a_{n-1}$  eventually produce values, which happens with our final local transition rule:

$$\frac{\nu a \sim \tau', a_0 \sim \tau, \dots, a_{n-1} \sim \tau \{a \hookrightarrow (k \text{ awaiting}[a_0, \dots, a_{n-1}])\} \otimes \mu}{\mapsto_{loc} \nu a \sim \tau' \{a \hookrightarrow (k \triangleleft \text{seq}(v_0, \dots, v_{n-1}))\}} \text{(locstep-join)}$$

where  $\mu = a_0 \hookrightarrow (\epsilon \triangleleft v_0) \otimes \dots \otimes a_{n-1} \hookrightarrow (\epsilon \triangleleft v_{n-1})$ .

Global transitions are made by partitioning the global state into pieces and performing some number of local transitions on those pieces:

$$\frac{\nu\Sigma_A\{\mu_A\} \mapsto_{loc} \nu\Sigma'_A\{\mu'_A\} \quad \dots \quad \nu\Sigma_Z\{\mu_Z\} \mapsto_{loc} \nu\Sigma'_Z\{\mu'_Z\}}{\nu\Sigma_A, \dots, \Sigma_Z, \Sigma\{\mu_A \otimes \dots \otimes \mu_Z \otimes \mu\} \mapsto_{glo} \nu\Sigma'_A, \dots, \Sigma'_Z, \Sigma\{\mu'_A \otimes \dots \otimes \mu'_Z \otimes \mu\}}$$

This global transition rule is highly nondeterministic in general.

**Task 3.1** (10%). Implement the parallel dynamics in `dynamics.sml`.

- The function `ready` should calculate all the symbols that are ready to take a local transition.

Given  $\mu = (a_1 \hookrightarrow s_1) \otimes \dots \otimes (a_n \hookrightarrow s_n)$ , we can say that  $a_i$  is *ready* if

- $s_i = (k \triangleright e)$ ,
- $s_i = (k \triangleleft v)$  and  $k \neq \epsilon$ , or if
- $s_i = (k \text{ awaiting}[b_0, \dots, b_{n-1}])$  and for each of the  $b_i$ , there is a mapping  $b_i \hookrightarrow (\epsilon \triangleleft v_i) \in \mu$

Every possible global transition step can be uniquely specified as some subset of the ready symbols.

- The function `trystep_par` should take a well-formed execution context and a list of ready symbols and should execute a step of the global transition.

Here is an example of running the parallel abstract machine toplevel interpreter:

```
$ rlwrap sml -m sources.cm
- Par.repl ();
->use "parfact.par";
Statics: exp has type int
--> {
    (id: 0)
    init@1 ~~~>
    emp
    |> Let(Fix[parr(int, int)](fact@11. Lam[int](x@12. If(Cmp[Eq](x@...
    })
    *** Local steps taken (work): 0
    *** Global steps taken: 0
->eval;
Num[3623758] VAL
    *** Local steps taken (work): 1329
    *** Global steps taken: 254
->eval let val S = [ x * x * x | x < 300 ] in S[12] end;
Statics: exp has type int
Num[1728] VAL
    *** Local steps taken (work): 2711
    *** Global steps taken: 20
->load [ x | x < 5 ];
Statics: exp has type seq(int)
--> {
    (id: 0)
    init@31 ~~~>
    emp
    |> Tab(x@386. x@386, Num[5])
    })
    *** Local steps taken (work): 0
    *** Global steps taken: 0
->step; step; step;
(...there's more in par-interpreter-session.txt...)
```

## 4 Parallel Programming

Different implementations of sequences can have different costs in terms of their work (the amount of time computation takes with one processor) and span (the amount of time computation takes with an unlimited number of processors). For the code in this section, which you will write in `mergesort.par`, we will assume the 15-210 `ArraySequence` implementation of sequences, meaning that getting a single element of an sequence with `e[en]`, finding out the length of an sequence with `len(e)`, and case analyzing on whether a sequence has length-1 or not with `showt e {elt x ⇒ e1 | node y1 y2 ⇒ e2}` are all constant-work and constant-depth operations. Sequence tabulation can be done entirely in parallel, so its work just the sum of the work needed to calculate each sequence element, and its span is the maximum of the spans needed to calculate each sequence element.

**Task 4.1** (5%). The `map` and `cat` operations described in *PFPL* can be implemented with appropriate work (linear) and span (constant) using the tabulation primitive. Implement `map` and `cat` for `seq(int)`.

**Task 4.2** (5%). Write a parallel reduce on `int` sequences. Consider a sequence of integers and an associative binary operation called  $\odot$ . Then, the expected behavior of reduce is as follows:

$$\text{reduce } (\odot) (\text{seq}(v_0, \dots, v_{n-1})) = \text{seq}(v_0) \odot \text{seq}(v_1) \odot \dots \odot \text{seq}(v_{n-2}) \odot \text{seq}(v_{n-1})$$

Since  $\odot$  is associative, we can parenthesize any way we want. In particular, we can parenthesize this as a balanced binary tree:

$$\left( \left( \left( \text{seq}(v_0) \odot \text{seq}(v_1) \right) \odot \left( \text{seq}(v_2) \odot \text{seq}(v_3) \right) \right) \odot \left( \dots \right) \right) \odot \left( \dots \right)$$

Then for each level of this tree, we can compute all the nodes at that level in parallel.

This is almost trivial to implement with `showt e {elt x ⇒ e1 | node y1 y2 ⇒ e2}`, though you have to take care to make sure evaluation actually happens in parallel. For full credit, implement `reduce` without `showt` (a correct `showt` solution will get 3 out of 5 points).

**Task 4.3** (5%). Assume the length of the sequence is a power of 2. What is the work and span of `reduce` if the associative binary operation has constant work and span? What if it has constant span but linear work (like `cat`)? What if it has  $O(\log n)$  span and linear work (like `merge`)?

**Task 4.4** (5%). Implement a semantically correct merge operation on sorted sequences. It is not known how to implement `merge` with  $O(n)$  work and  $O(\log n)$  span using the primitives and cost semantics we have established so far. We'll be looking for the best answers and may award extra credit (5-10 points) to particularly efficient solutions, but your solution doesn't have to be work efficient (though the work should be in  $O(n^2)$ ) or parallelizable. It's fine to use Google and outside sources (15-150 notes, Guy Blelloch's dissertation...) for reference and inspiration, but the usual policy on peer-to-peer collaboration applies.

**Task 4.5** (5%). Implement parallel mergesort using the other functions you have written (use `showt`). Make sure the two halves of a sequence are actually sorted in parallel according to the semantics.

## 5 Scheduling

Take a look at the file `tex/par-interpreter-session.txt` in the handout code. Running `step` or `eval` without any arguments always causes the maximum number of local transitions to accompany every global transition, simulating execution with an unlimited number of processors. But there is also an option for selecting exactly which local transitions should take a step.

```
->step <0,3,4>; (* Picks the local transitions identified by 0, 3, and 4 *)
```

By manually performing these selections such that at most  $p$  local transitions happen at any time, you would be, in effect, acting like a *scheduler* that is coordinating the available work among  $p$  different processors.

The overall progression of a parallel computation can be viewed as a *series-parallel* diagram like the one to the right, where the dots represent a *local* transition that must be made for the computation to complete. Any global state that could arise during evaluation corresponds to a cut through this graph. The cut labeled **1** to the right corresponds to a global state where there are four mappings from assignables to states. Three of these mappings (call them  $a_1 \mapsto s_1$ ,  $a_2 \mapsto s_2$ , and  $a_3 \mapsto s_3$ ) would correspond to the states that are ready to make local transitions labeled `o`, `g`, and `h`, respectively. The transitions `o` and `g` are sequential transitions by `(locstep-step)`, and the transition `h` is a forking transition (`locstep-fork`). The fourth mapping would be of the form  $a_0 \mapsto (k \text{ awaiting}[a_1, a_2, a_3])$ , as it is a state awaiting the termination of these three other computations; the next local step that this state will take corresponds to the node labeled `q`.

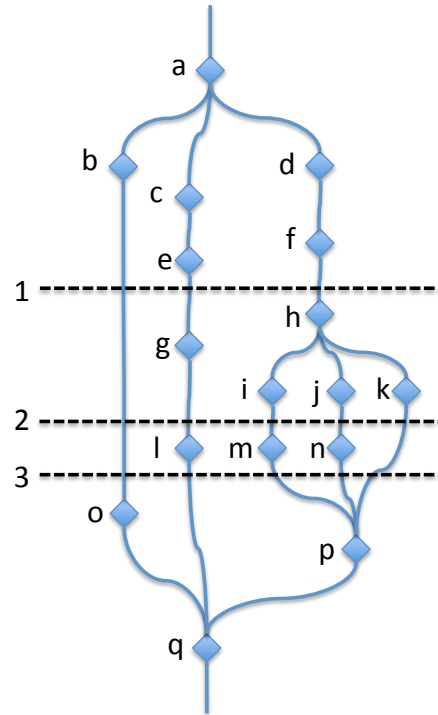


Figure 2: Series-parallel diagram

### 5.1 Scheduling strategies

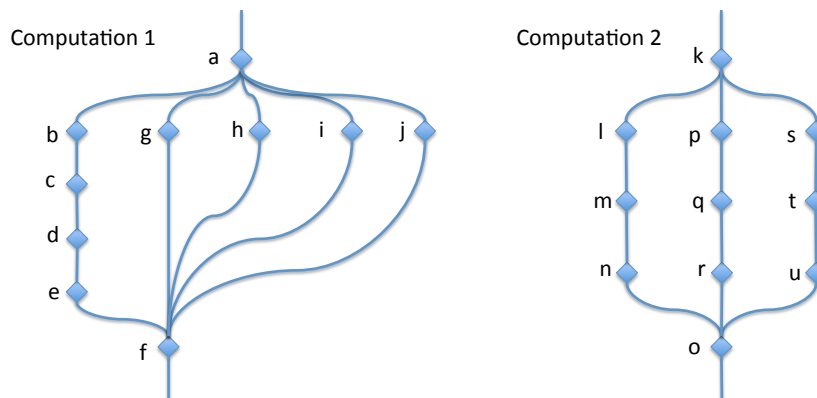
A scheduler provides some way of mapping the work that needs to be done onto a fixed number of processors that are available to do work. *Greedy* schedulers are ones that always keep processors working if there's available work to do. The most obvious greedy scheduling strategies that have been studied in practice are  $p$ -BFS and  $p$ -DFS (breadth-first and depth-first scheduling on  $p$  processors). Breadth-first scheduling is a simple strategy that can be implemented with a queue: if there's more available work than available parallelism, always prioritize one of the units of work that has been waiting longest to be scheduled. Depth-first scheduling always schedules the *leftmost* work from the perspective of the series-parallel diagram; visually this is rather clear, but the data structure to implement it is a bit trickier.

For the computation associated with the series-parallel diagram in Figure 2, Figure 3 gives an example of how scheduling would proceed under strict 2-DFS and 2-BFS scheduling.

Round #	2-DFS Scheduling		2-BFS Scheduling	
	Scheduled	Also ready	Scheduled	Also ready
1	<i>a</i>		<i>a</i>	
2	<i>b, c</i>	<i>d</i>	<i>b, c</i>	<i>d</i>
3	<i>o, e</i>	<i>d</i>	<i>d, o</i>	<i>e</i>
4	<i>g, d</i>		<i>e, f</i>	
5	<i>l, f</i>		<i>g, h</i>	
6	<i>h</i>		<i>l, i</i>	<i>j, k</i>
7	<i>i, j</i>	<i>k</i>	<i>j, k</i>	<i>m</i>
8	<i>m, n</i>	<i>k</i>	<i>m, n</i>	
9	<i>k</i>		<i>p</i>	
10	<i>p</i>		<i>q</i>	
11	<i>q</i>			

Figure 3: Scheduling the diagram in Figure 2

**Task 5.1** (5%). For both of these computations, give the order that work would be scheduled on 2 processors under breadth-first and depth-first scheduling.



This task shows that, BFS and DFS can have rather different performance depending on the shape of the computation; however, any greedy scheduler will take, at worst, twice the number of transitions that the best greedy scheduler would take. In practice, BFS and DFS scheduling (and more advanced variants, like *work stealing* schedulers, that we won't consider here) are important because they make a huge difference in space usage; unfortunately, this is outside the scope of this assignment.  $\square$

## 5.2 Implementation

In `scheduler.sml`, you will implement a centralized scheduler for  $p$  processors with processor IDs  $0, \dots, p - 1$ . Your implementation doesn't have to be any particular strategy ( $p$ -BFS,  $p$ -DFS, etc), but it should be greedy: if there's work to do and a processor is known to be idle, work should be scheduled on that processor.

Unlike our more theoretical treatments of scheduling, we won't necessarily try to reschedule everything at every single time-step. Consider our original series-parallel diagram example, for instance. If we'd

already assigned a processor to do work **c**, then we should probably just let it continue to run the sequential computation, doing work **e**, **g**, and **l**. The scheduler will only be called when a processor runs out of work to do, generating one of three messages:

- **Finished**  $v$  – the processor has finished sequentially evaluating the work it was given, and the result was a value.
- **Fork**  $(n, k, (x, e))$  – the processor reached the state  $(k; \text{tab}[\tau](x.e; \square) \triangleleft \text{num}[n])$  and now needs  $n$  sub-computations to be run.
- **NotFinished**  $s$  – the processor didn't finish sequentially evaluating the work it was given, but it was told to yield to the scheduler after taking some fixed number of steps. (You can also tell processors to run as long as necessary, in which case you'll never get this message.)

Whenever the scheduler is called, it can either declare the global computation finished and return a value, or it can return a list of processor IDs and states  $[(p_0, s_0), \dots, (p_k, s_k)]$  that should begin computing. The processors  $p_i$  in this returned list must currently be idle. In order to do this correctly, you'll need to hang on to three key pieces of state:

- Information about which processors are currently known to be idle. The starter code suggests storing this in an array `currently_working`. If `NONE = Array.sub(!currently_working, k)`, it means that the processor  $k$  is known to be idle.
- Information about which ready pieces of work are waiting to be scheduled (because at some point there enough processors). The starter code suggests storing this in a list `unscheduled_work`. If you follow this, it means that what you implement won't quite be  $p$ -DFS. You're welcome to modify this to implement something closer to  $p$ -BFS or true  $p$ -DFS.
- Information about *why* you're running, or waiting to run, a given computation. Every computation except for the initial one is trying to compute the  $i^{\text{th}}$  element of some sequence. Taking this into account, instead of associating every computation with a symbol, we will associate every computation with a *destination*, either `Initial` for the first computation or `WriteTo{loc, index = i}` for every other computation.

Whenever a fork spawns  $n$  new computations, the computation  $(\text{stack awaiting}[a_0, \dots, a_{n-1}])$  is represented by an object  $\text{loc} = \text{Mem}\{\text{dest}, \text{waiting\_on}, \text{stack}, \text{seq}\}$ , where  $\text{seq}$  is initially a length- $n$  array containing all `NONE` values and  $\text{waiting\_on}$  is a reference initially containing  $n$ . Whenever a sub-computation returns a value  $v$ , you should write `SOME v` to the appropriate index of the array  $\text{seq}$  and decrement  $\text{waiting\_on}$ . If  $\text{waiting\_on}$  has been decremented all the way to 0, then the array  $\text{seq}$  contains  $n$  values `SOME v0, ..., SOME vn-1`, and it is possible to immediately re-schedule the computation  $(\text{stack} \triangleleft \text{seq}(v_0, \dots, v_{n-1}))$  that has been waiting by adding it back to the `unscheduled_work` worklist. This newly-rescheduled computation's destination,  $\text{dest}$ , was stored as a part of  $\text{loc}$ .

**Task 5.2 (20%).** Implement the scheduler in `scheduler.sml`. Upon receiving a message from the processor  $k$ , the scheduler should:

1. Grab the destination associated with  $k$ 's computation from `currently_working[k]`, and then mark processor  $k$  as idle by setting `currently_working[k]` to `NONE`.

2. If  $k$  returned a value (`Finished`), store the resulting value in the right place – or terminate the computation if it's all finished. If storing the resulting value re-enables a suspended computation, add that suspended computation to the worklist.
3. If the processor returned with more work to do (`NotFinished` or `Fork`), add all the unfinished work to the worklist.
4. Reschedule: remove jobs from the worklist and assign them to idle processors until either you run out of available processors or until the worklist is empty.

You're welcome to modify the given data structures at your own risk as long as you do not change the interface or reference the sequential dynamic semantics. □

Rather than giving a REPL, the structure `Simulate` in `simulate.sml` provides facilities for running and testing the scheduler on the command line. The `step_limit` argument determines how long a processor will run a sequential computation before yielding to the scheduler with a `NotFinished` message, and `step_limit = NONE` means it will only yield when it is finished.

```
$ rlwrap sml -m sources.cm
- Simulate.evalFile { num_procs = 5, step_limit = NONE, filename = "parfact.par" };
one new task scheduled
<00000> => <X0000>
Processor 0 invoking scheduler (fork 11)... 5 new tasks scheduled
<!0000> => <XXXXXX>
Processor 4 invoking scheduler (done)... one new task scheduled
<XXXX!> => <XXXXXX>
Processor 3 invoking scheduler (done)... one new task scheduled
<XXX!X> => <XXXXXX>
Processor 2 invoking scheduler (done)... one new task scheduled
<XX!XX> => <XXXXXX>
Processor 1 invoking scheduler (done)... one new task scheduled
<X!XXX> => <XXXXXX>
Processor 0 invoking scheduler (done)... one new task scheduled
<!XXXX> => <XXXXXX>
Processor 0 invoking scheduler (done)... one new task scheduled
<!0000> => <X0000>
Processor 1 invoking scheduler (done)... no new tasks scheduled
<X!000> => <X0000>
Processor 2 invoking scheduler (done)... no new tasks scheduled
<X0!00> => <X0000>
Processor 3 invoking scheduler (done)... no new tasks scheduled
<X00!0> => <X0000>
Processor 4 invoking scheduler (done)... no new tasks scheduled
<X000!> => <X0000>
Processor 0 invoking scheduler (done)... one new task scheduled
<!0000> => <X0000>
Processor 0 invoking scheduler (done)... all done.
val it = {res=-,steps=304,work=1329} : {res:?.t, steps:int, work:int}
- print (Term.toString (#res it) ^ "\n");
Num[3623758]
```

By default, the simulator will print a message before and after every invocation of the scheduler, reporting on which processors are busy (X) and which are idle (O). To turn this off, you can set the flag `Simulate.verbose := false`.

The work (total number of local transitions) reported by this simulator should be the same as the work reported by the parallel semantics, but even if you make more than enough processors available, the number of global transitions may be slightly different. This is both because the scheduler now effectively performs the (locstep-join) local transitions as soon as they are possible and because finished computations need to wait for an additional global transition before they report back to the scheduler that they've finished.

## A Putting The Code Together

As usual, you can compile your files for using `CM.make "sources.cm"`, and examples of using the top-level interpreter to run sequential and parallel dynamics have been given above.

### A.1 Interpreter

The syntax for each term construct is as close as possible to the concrete syntax mentioned for it. (The concrete syntax is the second column in the table which introduces the syntax for the language.) We provide below the grammar that the interpreter accepts.

```
ident ::= (* a letter followed by alphanumeric characters, _, or ' *)
numeral ::= (* a series of digits *)
stringconst ::= (* a standard C-escaped string like "foo" or "why \"hello\"" *)

seqcommand ::=
  eval <exp>;
| eval;
| step <exp>;
| step;
| load <exp>;
| use <stringconst>;

parcommand ::=
  eval <exp>;
| eval;
| step < <exps> >; (* <exps> should be comma-separated numbers *)
| step; (* all possible local transitions happen *)
| load <exp>;
| use <stringconst>;

ty ::= int | bool | <ty> -> <ty> | <ty> * <ty> | seq <ty>

exp ::=
  let <decls> in <exp> end
| fn (<ident> : <ty>) <exp>
| <exp>(<exp>)
| fix <ident> : <ty> is <exp>
| <<exp>, <exps>>
| <exp>.l
```



```

| <exp>.r
| <numeral>
| <exp> == <exp>
| <exp> != <exp>
| <exp> >= <exp>
| <exp> <= <exp>
| <exp> + <exp>
| <exp> - <exp>
| <exp> * <exp>
| <exp> / <exp>
| <exp> % <exp>
| true
| false
| if <exp> then <exp> else <exp>
| seq(<exps>)
| [ <exp> | <ident> < <exp> ]
| <exp> [ <exp> ]
| len <exp>
| showt <exp> { elt <ident> => <exp> | node <ident> <ident> => <exp> }

exps ::= <exp> | <exp>, <exps>

decls ::= <decl> | <decl> <decls>

decl ::= val <ident> = <exp>

```

## A.2 Reference implementation

You can use the reference interpreter to test your code and help debug your interpreter.

The reference interpreter can be loaded like this:

```

$ rlwrap sml @SMLload=/afs/andrew/course/15/312/bin/ref_impl_par
- Seq.repl ();
- Par.repl ();

```