

# CS 312 Problem Set 5: Concurrent Language Interpreter

Due: 11:00 PM, April 14, 2005

---

## 1 Introduction

In this assignment you will build an interpreter for a functional language called CL, with concurrency and imperative features. A CL program has multiple processes executing at the same time. You can think of the concurrent processes as robots. Each robot executes its own CL code, and has its own local memory. Robots can communicate to each other through a global shared memory. They can also start off other robots, or wait for the spawned robots to finish. Finally, there is an outside world that provides additional functionality (for instance, I/O support), and robots can request services from this outside world.

For Problem Set 5, you will implement an interpreter for programs written in CL. More precisely, you will implement the evaluation of CL expressions, including the concurrent constructs. You will also implement a garbage collector to manage your local and global memories. In the next assignment, you will use your interpreter to implement a game that uses the robots.

## 2 Changes to problem set

- [Apr 1] Small change in the specification for spawn trees: a) you need a function `singleton` that builds a tree with the root process (instead of empty); and b) when a process terminates, its child processes should not be removed from the spawn tree.

## 3 Instructions

You will do this problem set by modifying the source files provided in CMS, and submitting the program that results. As before, your programs must compile without any warnings. Programs that do not compile or compile with warnings will receive an automatic zero. All files submitted should *not* have any lines longer than 80 characters, and ideally all lines should be less than 78 characters long.

We will be evaluating your problem set on several different criteria: the specifications you write (where appropriate), the correctness of your implementation, code style, efficiency, and validation strategy. Correctness is worth about two thirds of the total score, and the importance of the other criteria varies from part to part.

Note that you will be building on your PS5 solution for PS6, so we strongly recommend you to start early on PS5 and understand the code given to you early. PS5 and PS6 are also partner assignments. You are expected to find a partner to do this assignment with by April 1. You can continue working with the partner from PS4, or find somebody else. In either case, you must sign up in CMS with that partner. If CMS shows that you don't have a partner by April 1, we will automatically pair you up with someone else in the class at that time.

## 4 The CL language

The CL language has some interesting features. First, it is a concurrent language in which multiple processes can be executing simultaneously. Second, it has imperative features that allow updating memory locations. Third, processes can get additional services (including input/output operations) from an external world.

A robot can launch another robot using the expression `spawn e`. The expression  $e$  provides the program that the newly created robot is supposed to execute. A robot can then wait for all its spawned children to finish using `sync`.

There are two different kinds of memories that the robots can read or write. Each robot has its own *local memory*, which can only be used by that robot. Local memory is allocated with `lref e` expressions. In addition, there is a *global memory* that is global and shared by all the robots. Robots can communicate with each other by modifying locations in the global memory. Global memory is allocated with `gref e` expressions.

Robots can request the external world to perform actions, usually using an expression of the form `do e`. This expression is evaluated by sending the result of  $e$  to the external world. Different possible values of  $e$  are interpreted as requests to perform different actions. In this problem set, the `do e` expression will be used for I/O operations. For example, the expression `do 0` causes the external world to ask the user to input a number, which is returned as a result of the expression.

The behavior of the external world is not specified by the CL language. We have given you one possible implementation of the external world, but it will be modified in PS6 to allow robots to sense and interact with their environment in many more ways.

### 4.1 Expressions

A CL program for a single robot can contain the following expressions:

|  |   |
|--|---|
| $n$  | An integer constant, as in SML. Examples: $\sim 3, 0, 2$ .  |
| $(e_1, e_2)$   | A pair. Evaluates to the value $(v_1, v_2)$ where $v_1$ and $v_2$ are the respective results of evaluating the expressions $e_1$ and $e_2$ .  |
| $unop\ e$  | Returns $unop$ applied to the result of evaluation of $e$ . $unop$ is one of the following unary operators: $\sim$ (negates an integer), and $rand$ (returns a random number between 1 and $n$ where $n$ is the result of evaluation of $e$ ).  |
| $e_1\ binop\ e_2$  | Applies binary operator $binop$ to the results of evaluations of the two expressions. Both $e_1$ and $e_2$ must evaluate to an integer. $binop$ is one of the following operators: $+, -, *, /, \text{mod}, <, =$ . For the last two operators the result will be 1 if the comparison is true, and 0 otherwise. |
| $e_1 ; e_2$  | A sequence of expressions. It is evaluated similarly to an ML sequence. First expression $e_1$ is evaluated, possibly creating side effects (modifying memories). After that the result of $e_1$ is thrown away and expression $e_2$ is evaluated.  |
| <code>let <math>id = e_1</math> in <math>e_2</math></code> | Binds the result of evaluating $e_1$ to the identifier $id$ and uses the binding to evaluate $e_2$ . Identifiers start with a letter and consist of letters, underscores, and primes.   |

|   |   |
|---|---|
| <code>fn <math>id \Rightarrow e</math></code>   | Anonymous function with the argument $id$ and the body $e$ . Note that functions are values, so the body $e$ is not evaluated until an argument is supplied to the function.  |
| <code><math>id</math></code>  | Identifier. Must be contained inside a <code>let</code> or <code>fn</code> expression with the same identifier name, otherwise unbound identifier error will occur.   |
| <code><math>e_1 e_2</math></code>   | Function application. Evaluates expression $e_1$ to a function <code>fn <math>id \Rightarrow e</math></code> , expression $e_2$ to a value $v_2$ , binds $v_2$ to the identifier $id$ and uses the binding to evaluate $e$ .  |
| <code>if <math>e</math> then <math>e_1</math> else <math>e_2</math></code>  | Similar to the ML <code>if/then/else</code> expression except that the result of expression $e$ is tested for being greater than 0 (there are no booleans in CL). Examples: <code>if 1 then 1 else 2</code> returns 1, <code>if 4&lt;3 then 1 else 2</code> returns 2.  |
| <code>typecase <math>e</math> of</code><br><code>  (<math>id, id'</math>) <math>\Rightarrow e_1</math></code><br><code>    int <math>id \Rightarrow e_2</math></code><br><code>    loc <math>id \Rightarrow e_3</math></code><br><code>    fun <math>id \Rightarrow e_4</math></code><br><code>    any <math>id \Rightarrow e_5</math></code> | First evaluates expression $e$ to a value. If the result is a pair, it binds the elements of the pair to $id$ and $id'$ in the case for pairs. Otherwise, it binds the the result to $id$ in the appropriate case. The case <code>any</code> matches any value. It then evaluates the expression $e_i$ of the matched case.<br><br>Each of the cases is optional and can occur at most once. The case for <code>any</code> is allowed only if at least two of the other cases are missing. As in ML, all cases must be covered. The expression “ <code>typecase <math>e_1</math> of any <math>id \Rightarrow e_2</math></code> ” is equivalent to “ <code>let <math>id = e_1</math> in <math>e_2</math></code> ”. |
| <code>delay <math>e</math> by <math>n</math></code>   | Note that lists can be emulated in CL using pairs of pairs. Like pattern matching in ML, the <code>typecase</code> construct gives the ability to distinguish between the head and the tail of a list.<br>Delays the evaluation of $e$ by $n$ evaluation steps. The number $n$ must be an integer constant greater or equal to 1. At each evaluation step, $n$ is decreased; when it reaches 1, the expression reduces to $e$ .   |
| <code>lref <math>e</math></code>  | Similar to the ML operation <code>ref</code> . First expression $e$ is evaluated to a value $v$ . After that a new location $loc$ is allocated in the robot’s local memory and value $v$ is stored at this location. The return result of the expression is location $loc$ which can be viewed as a memory address.   |
| <code>gref <math>e</math></code>  | Similar to <code>lref</code> except that the new location is allocated in the global shared memory. Before allocating the location the result of $e$ is checked to ensure that it satisfies the “global memory invariant” (see section 4.3).  |
| <code>! <math>e</math></code>   | Evaluates expression $e$ to location $loc$ and returns the value stored at this location.   |
| <code><math>e_1 := e_2</math></code>  | Evaluates expression $e_1$ to a location $loc_1$ and expression $e_2$ to a value $v_2$ . After that replaces the value at the location $loc_1$ with $v_2$ . The return result of this expression is $v_2$ . If $loc_1$ is a location in the global memory, then the value $v_2$ is checked for the “global memory invariant” before assigning (see section 4.3).  |
| <code>lock <math>e_1 e_2</math></code>  | This expression evaluates $e_1$ to a location, $loc$ , and, except as noted below, returns $loc$ . If $loc$ is in local memory, the program proceeds with the evaluation of $e_2$ . If $loc$ is in global memory and is not already locked, then the  |

current process acquires a lock for *loc*. If any other process already has the lock, the process will continue to attempt the operation until the old lock is removed. All other cases are runtime errors.

Once the current process has grabbed the lock, the expression reduces to a new expression of the form `locked loc e2`. Then it evaluates *e<sub>2</sub>*, while maintaining the lock. When the evaluation of *e<sub>2</sub>* finishes with a value *v*, the lock is released and the value *v* is returned.

`do e`

This allows a robot to interact with the external world. First expression *e* is evaluated to a value *v* which is then sent to the external world. The return result of this expression can be arbitrary (it is specified by the external world). The list of actions currently recognized by the external world is given in section 4.6.

`spawn e`

This launches a new robot. The code of the spawned robot is *e*.

`sync`

The current robot waits (i.e., does nothing) until all of its spawned children have finished executing. Once that happens, if *v<sub>1</sub>, ... v<sub>n</sub>* are the result values of the children (in the order they have been spawned), then `sync` evaluates to  $(v_1, (v_2, (\dots(v_n, 0))))$ .

There two expressions that never appear in the source of a CL program, but can occur during the evaluation:

- *loc*, a memory location. A location can be viewed as a pair  $(scope, addr)$  where *scope* identifies whether it is in the local or global memory and *addr* is a memory address. A location can only be generated using `lref` and `gref` expressions.
- `locked loc e`. This occurs during the evaluation of a lock expression, after the lock for *loc* has been acquired.

We have provided for you an implementation of the expression type as `AbSyn.exp` in the file `absyn/absyn.sml`.

## 4.2 Values

Some of the expressions described above are values (i.e. they cannot be evaluated any further). Here is the list of possible values:

- Integer constants *n*;
- Locations *loc*;
- Functions `fn id => e`;
- Pairs  $(v_1, v_2)$ , provided that *v<sub>1</sub>* and *v<sub>2</sub>* are values.

Note that there is no special type for values in our implementation; it is up to the programmer to identify which expressions are values.

### 4.3 Local and global memories

A memory  $\sigma$  can be viewed as a mapping from locations (or addresses) to values. Each robot has its own local memory that cannot be accessed by other robots. In addition, there is a global memory shared among all robots.

A difference between a local and the global memories can be illustrated with the following example:

```
let
  r = lref 0
in
  spawn (fn x => (r := 1));
  !r
```

This robot (let's call it "A") allocates a location (call it  $loc$ ) for an integer 0 and then launches another robot (let's call it "B"). The local memory of  $A$  is copied to the local memory of  $B$ , so local memories of  $A$  and  $B$  will contain two different locations storing value 0.

After some reductions robot  $A$  evaluates to expression  $!loc$  and robot  $B$  to expression  $(loc := 1)$ . Robot  $B$  then modifies its own copy of location  $loc$  to 1; memory of robot  $A$  is unchanged. Thus, robot  $A$  will return 0.

Now consider the same code where `lref` is replaced with `gref`. Then location  $loc$  will be allocated in the global memory, so after launching  $B$  locations  $loc$  in both robots will point to the same place. Therefore, depending on the order of executions of  $A$  and  $B$ , robot  $A$  will return either 0 (if  $A$  is executed before  $B$ ) or 1 (if  $A$  is executed after  $B$ ).

To make sure that the local memory of a robot cannot be accessed by other robots we need to maintain the following *global memory invariant*: values stored in the global memory do not contain locations from local memories. Thus, each modification of the global memory (i.e. expressions `gref v` and `loc := v` where  $loc$  is a location in the global memory) must be checked before evaluation: if value  $v$  contains references to local memories, then a run-time error will occur. An example of an invalid expression is `gref (lref 0, 0)`. A robot trying to execute such an expression should be terminated.

### 4.4 Evaluation

A process (that is, a single robot) is represented by a unique process identifier  $pid$ , local memory  $M$  and expression  $e$ . A current state of the interpreter is described by a queue of processes, as well as a global memory  $M_g$ , and a tree  $T$  of spawned processes. A process  $p$  is the child of  $p'$  in the spawn tree  $T$  if  $p'$  has been directly spawned by  $p$ .

The interpreter repeatedly performs the following operation: it takes the process at the head of the queue, performs a single evaluation step on its expression (possibly modifying the process local and global memory or the spawn tree), and places the modified process at the end of the queue. A single step is illustrated in Figure 1.

It is important that robot programs execute one step at time. If we evaluated a program down to a value all at once, the system would not be concurrent because only that robot would be able to run. Therefore, we must evaluate in steps.

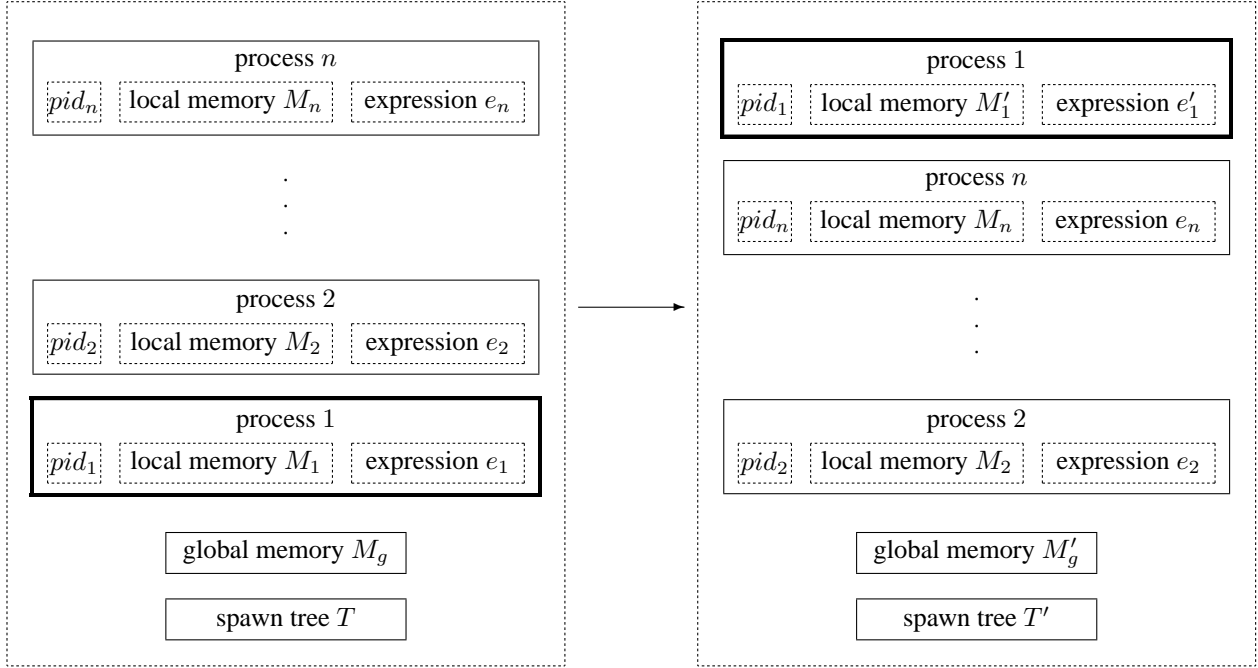


Figure 1: Single step of the interpreter on process 1. Expression  $e'$  is the result of a single evaluation step on  $e$ . Possible side effects include modifying local memory  $M_1$  and global memory  $M_g$ , and modifying the spawn tree  $T$ .

Given an expression, the evaluator finds the leftmost subexpression that can be reduced, and reduces this subexpression.

Note that just as in ML, some expressions do not evaluate some of their subexpressions before doing a reduction. These expressions are `let  $id = v$  in  $e$` , `if  $v$  then  $e_1$  else  $e_2$` , `fn  $id \Rightarrow e$` , `delay  $e$  by  $n$` , `typecase  $v$  of ( $id, id'$ )  $\Rightarrow e_1 \mid \dots$` , `spawn  $e$` , `lock  $loc$   $e$` , and  `$v; e$` . The  $v$ 's indicate subexpressions that must be fully evaluated before the expression can be reduced, and the  $e$ 's indicate subexpressions that are not evaluated until after the reduction of the whole expression.

#### 4.5 Reductions

The list of possible reductions that can be performed during evaluation is given below. First we consider reductions that do not change local or global memories. Letters  $v$  stand for values, and letters  $e$  for expressions which may or may not be values.

$$\begin{array}{ll}
 unop\ v \longrightarrow v' & \text{where } v' = unop\ v \\
 v_0\ binop\ v_1 \longrightarrow v' & \text{where } v' = v_0\ binop\ v_1 \\
 v; e \longrightarrow e & \\
 \text{let } id = v \text{ in } e \longrightarrow e\{v/id\} & \\
 (\text{fn } id \Rightarrow e) v \longrightarrow e\{v/id\} & \\
 \text{if } v \text{ then } e_1 \text{ else } e_2 \longrightarrow e_1 & v \in \{1, 2, 3, \dots\} \\
 \text{if } v \text{ then } e_1 \text{ else } e_2 \longrightarrow e_2 & \text{all other } v \\
 \text{delay } e \text{ by } n \longrightarrow \text{delay } e \text{ by } n' & \text{where } n' = n - 1, \text{ if } n > 1 \\
 \text{delay } e \text{ by } 1 \longrightarrow e & 
 \end{array}$$

$\text{typecase } (v, v') \text{ of } \dots (id, id') \Rightarrow e \dots \longrightarrow e\{v/id, v'/id'\}$   
 $\text{typecase } v \text{ of } \text{lab } id \Rightarrow e \dots \longrightarrow e\{v/id\}$       where *lab* is one of the cases *int*, *loc*,  
*fun*, or *any*, which matches *v*

$e\{v/id\}$  stands for the result of substitution of value *v* for all occurrences of identifier *id* in expression *e*. These reductions are similar to the reductions you have learned for SML. The rules for the memory accesses are as follows:

$!loc \longrightarrow v$       where *loc* is a location in the process local memory or in the global memory, and *v* is the value stored at this location  
 $\text{lref } v \longrightarrow loc$       where *loc* is a new location in the process local memory  
    Side effect: a location *loc* is allocated in the memory, its content is initialized with *v*  
 $\text{gref } v \longrightarrow loc$       where *loc* is a new location in the global memory  
    Checks: *v* satisfies the global memory invariant (Section 4.3)  
    Side effect: a location *loc* is allocated in the memory, with its contents initialized to *v*  
 $loc := v \longrightarrow v$       where *loc* is a location in the process local memory or in the global memory  
    Checks: *v* satisfies global memory invariant (if *loc* is global)  
    Side effect: content of the location *loc* is replaced with *v*

Finally, the reductions for concurrent constructs are:

$\text{lock } loc \ e \longrightarrow e$       where *loc* is a location in local memory  
 $\text{lock } loc \ e \longrightarrow \text{lock } loc \ e$       where *loc* is a location in global memory that is currently locked by another process  
 $\text{lock } loc \ e \longrightarrow \text{locked } loc \ e$       where *loc* is a location in global memory and is not currently locked. Effects: location *loc* is locked by the current process  
 $\text{locked } loc \ v \longrightarrow v$       where *loc* is a location in global memory that is locked by the current process. Effects: the lock for *loc* is released  
 $\text{do } v \longrightarrow e$       where *e* is the expression returned by the external world  
    Side effect: send `doAction(pid, v)` to the external world (which will return an expression *e*) where *pid* is the process identifier of the robot (see Figure 2)  
 $\text{spawn } e \longrightarrow 0$       Side effects: (1) select a fresh process identifier *pid'*  
    (2) update the spawn tree *T*; and (3) launch a new process with the process identifier *pid'* expression *e*, and a copy of the memory of the current process. (see Figure 3)  
 $\text{sync} \longrightarrow \text{sync}$       where not all the spawned children have finished  
 $\text{sync} \longrightarrow (v_1, \dots, (v_n, 0))$       where  $v_1, \dots, v_n$  are the values that each of the spawned children have evaluated to, in the order in which the kid robots were launched.

Notice that because expressions may have side effects, it is critical that expressions are evaluated left to right. For example,  $e_1 \text{ binop } e_2$  must be evaluated as

$$e_1 \text{ binop } e_2 \longrightarrow v_1 \text{ binop } e_2 \longrightarrow v_1 \text{ binop } v_2 \longrightarrow v$$

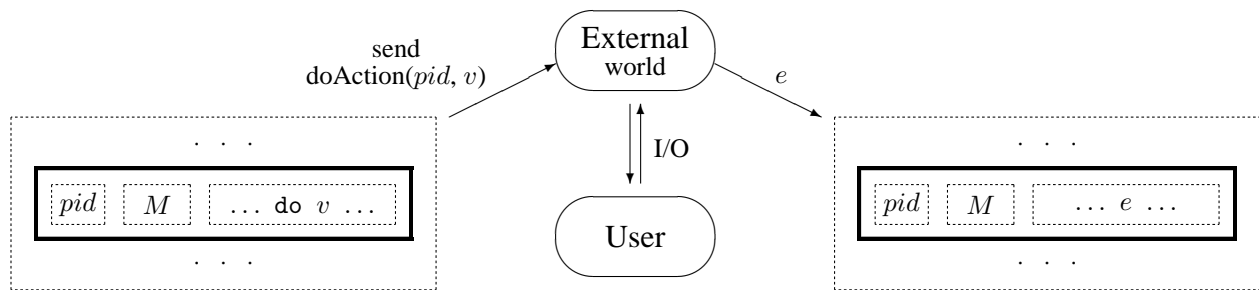


Figure 2: Evaluation of the `do v` expression




Figure 3: Evaluation of the `spawn v` expression. Before sending an event to the external world the interpreter picks a fresh process identifier  $pid'$

#### 4.6 The external world

Currently the `do` action performs simple I/O operations, though in PS6 it will be a general mechanism for interacting with the world. The following actions are currently provided:

- `do 0` : reads a number from the input, returns it to the interpreter
- `do (1, v)` : prints the value  $v$  to the output and returns  $v$ .
- `do (2, (c1, (c2, (c3, (... (cn, 0))))))` : prints the characters  $c_1, \dots, c_n$ . Returns 1 if well-formatted, 0 otherwise.
- `do (3, v)` : if value  $v$  is well formed, prints  $v$  and returns 1, otherwise prints undefined text and returns 0. Here  $v$  is considered well formed if it only contains pair and integer expressions.



## 4.7 Configurations

A *configuration* is the state of the entire interpreter at a particular point during execution. The configuration consists of a set of processes, each of which has a currently executing expression and local memory, a global memory that is shared by all the processes, and a spawn tree.

We can describe a single process as a triple  $\langle pid, M, e \rangle$ . The entire interpreter configuration is a tuple containing the global memory  $M_g$  and the current queue of processes:

$$\langle T, M_g, \langle pid_1, M_1, e_1 \rangle, \dots, \langle pid_n, M_n, e_n \rangle \rangle$$

The process at the head of the queue, process 1, is the one that will take the next evaluation step and be pushed to the end of the queue. Suppose that this process takes the evaluation step  $e_1 \rightarrow e'_1$ , with side effects that change the local memory  $M_1$  to  $M'_1$ , the global memory  $M_g$  to  $M'_g$ , and the spawn tree from  $T$  to  $T'$ . Then the effect of this step on the configuration as a whole is this:

$$\begin{aligned} & \langle T, M_g, \langle pid_1, M_1, e_1 \rangle, \langle pid_2, M_2, e_2 \rangle, \dots, \langle pid_n, M_n, e_n \rangle \rangle \\ \longrightarrow & \langle T', M'_g, \langle pid_2, M_2, e_2 \rangle, \dots, \langle pid_n, M_n, e_n \rangle, \langle pid_1, M'_1, e'_1 \rangle \rangle \end{aligned}$$

The type for configurations `Configuration.configuration` is implemented in `eval/configuration.sml`. A single step of the interpreter is performed by the function `Evaluation.stepConfig` in `eval/evaluation.sml`.

## 4.8 Creating and terminating robots

Robots can create other robots by calling `spawn e`. As a result, a new process will be added to the list of processes. The new process will have a copy of the old process local memory. The two processes will be able to communicate with each other if the old process had allocated locations in the global memory before spawning.

If a process has evaluated to a value, then it *terminates*—it is deleted from the list of processes. Thus, we have the following evaluation rule:

$$\begin{aligned} & \langle T, M_g, \langle pid_1, M_1, v_1 \rangle, \langle pid_2, M_2, e_2 \rangle, \dots, \langle pid_n, M_n, e_n \rangle \rangle \\ \longrightarrow & \langle T', M'_g, \langle pid_2, M_2, e_2 \rangle, \dots, \langle pid_n, M_n, e_n \rangle \rangle \end{aligned}$$

Here,  $M'_g$  is the global memory with all locks belonging to  $pid_1$  released. The updated spawn tree  $T'$ , is the same as  $T$ , but records the result value  $v_1$  for terminated process  $pid_1$ .

A process should also be terminated if it causes a run-time error such as a type error (e.g. `!0`) or a violation of the global memory invariant (e.g. `gref (lref 0)`). A process that is terminated due to a run-time error yields a result of `-1`. These run-time errors correspond to processes for which there is no legal reduction. Note that such errors should terminate the process encountering an error but should not affect other running processes.

# 5 Using the interpreter

## 5.1 File structure

The code is structured as follows:

- `absyn/absyn.sml`: definitions of basic types (`AbSyn.exp`, `AbSyn.pid`, `Absyn.action`)
- `eval/memory.sig`, `memory.sml`: definition of the memory type (`'a Memory.memory`) and associated operations
- `eval/spawntree.sig`, `spawntree.sml`: the spawn tree structure
- `eval/configuration.sml`: definition of the configuration type
- `eval/evaluation.sml`: a single step of the main interpreter loop
- `eval/gc.sig`, `gc.sml`: garbage collector
- `world/action.sig`: interface for interaction with the external world
- `debug/debug-loop.sml`: interface for debugging
- `eval/check.sig`, `check.sml`: well-formedness and consistency checking for expressions, processes and memories. Useful when debugging.
- `cl/*.cl`, a few sample CL programs

## 5.2 Running CL code

After compiling the code (`CM.make()`) you can enter the debugging mode using the command

```
Debug.debug "a string representing a CL program"
```

You will see a prompt (`>`). You can get the list of available commands by typing `"help"`. These are some commands for quick start:

- `step`: steps one step and shows the new stepped expression
- `run`: runs until the end
- `l file`: resets the interpreter and loads a file with a CL program
- `h`: gives you the help message and shows you many more commands
- `q`: quits the debugger

There are many other helpful functions and debugger commands; see `debug/debug-loop.sml` for more details. If you feel that the debugging tools implemented are inadequate, feel free to modify them.

## 5.3 String Literals

Although strings are not part of CL the parser will convert string literals into lists of integers. For example, `"hello"` parses as `(104, (101, (108, (108, (111, 0))))`.

## 6 Your task

### Part 1: Evaluator (60 pts)

Parts of the single-step evaluator are currently written, but there are holes in the implementation. Also, the implementation has not been tested fully, so it is your job to fix any problems you may encounter.

Your task is to finish the single-step evaluator. You will have to make changes to the following files:

- `eval/evaluation.sml`
- `eval/reductions.sml`

To help in your task, we have also implemented some functions in `eval/check.sml` that can be used to check whether expression, processes, and memories are well formed. These functions will be useful in checking that your interpreter is implemented correctly.

**To Submit:** Completed versions of `eval/evaluation.sml` and `eval/reductions.sml`. Also submit a summary of your changes in an ASCII file `eval.txt`, so that we know where to look when we are grading.

### Part 2: Memory Locks and The Spawn Tree (13 pts)

Finish the implementation of memory synchronization operations, and provide an implementation for spawn trees. You must modify the following files:

- In `eval/memory.sml`, provide implementations for `acquire`, `release` and `releaseAll`;
- In `eval/spawntree.sml`, provide an implementation for the spawn tree abstraction. Provide a concrete type `tree`, and fill in the all of the functions in this file.

**To Submit:** Completed copies of `memory.sml` and `spawntree.sml`.

### Part 3: The garbage collector (12 pts)

Garbage is data in local or global memory that is not reachable by following any chain of references from a running process. These locations should be periodically reclaimed and used for subsequent allocation requests. The process of reclaiming unreachable locations is known as *garbage collection*.

The signature `gc.sig` describes an automatic garbage collector for the CL language. Occasionally the garbage collector will be used to clean up memory. For the purpose of CL, two kinds of garbage collection are defined: local garbage collection and global garbage collection. Local garbage collection cleans up the local memory of a particular robot. Global garbage collection cleans the local memory of all robots as well as the shared global memory in a configuration.

Implement global and local garbage collection using the mark-and-sweep algorithm described in class. As implied by `gc.sig`, the `malloc` function should try to reuse locations that the garbage collector has reclaimed.

To help you test your garbage collector, the `localGC` and `globalGC` commands in debug mode will force garbage collections to take place immediately.

**To Submit:** An implementation of the signature `gc.sig` (do not change the signature) in the file `gc.sml`.

#### Part 4: Amortized Complexity Analysis (15 pts)

In class, you saw *static* hashing, where the only way to deal with overfull tables was to do an explicit table resize. There are many sophisticated *dynamic* hashing techniques, which are designed to handle growth in data far more gracefully.

This question is about a simple dynamic hashing technique known as *extendible hashing*. The main idea is that the hash buckets will be accessed through a *directory*; if a particular bucket becomes overfull, we will not modify the entire table, but we will only split the one offending bucket, and use the directory to reflect this change to the outside world.

The extendible hashing description below makes the following assumptions:

- each bucket has maximum occupancy  $k$ , where  $k$  is a constant
- finding an item in a bucket takes only constant time

Neither of these assumptions are essential to the functioning of extendible hashing, but they make the description clearer.

##### 6.1 Description of extendible hashing

As mentioned, the hash buckets will be accessed through a directory of pointers. The directory entries will be the first  $c$  bits of the hash for a particular value.  $c$  may change throughout the algorithm, as we choose to keep more or less information in the directory. For example, suppose that  $c = 2$ , and that the hash function  $h$  maps some values to the following numbers:  $h(x) = 10110$ ,  $h(y) = 10010$ ,  $h(z) = 11001$ ,  $h(w) = 11100$

When  $c = 2$ , the directory has four entries, 00, 01, 10 and 11. Suppose we insert the values  $x$ ,  $y$ ,  $z$  and  $w$  into the hash table. The table will now look as follows:

```
00
01
10 → [x, y]
11 → [z, w]
```

Where the directory entries are listed on the left,  $\rightarrow$  represents a pointer, and  $[x, y]$  represents a bucket containing the values  $x$  and  $y$ .

Now, suppose the maximum occupancy  $k$  for a bucket is 2, and suppose we want to insert  $a$  into the table, where  $h(a) = 10011$ . We must double the size of the directory, and we must increase  $c$  to 3. The new table will look as follows:

000  
 001  
 010  
 011  
 100  $\rightarrow [a, y]$   
 101  $\rightarrow [x]$   
 110  $\rightarrow [w, z]$   
 111  $\rightarrow [w, z]$

Note a few important things:

- There was a bucket split, creating two new buckets  $[a, y]$  and  $[x]$ .
- We did NOT split the  $[w, z]$  bucket; both entry 110 and 111 still point to the same bucket. This bucket did not need to be split, and we want to avoid unnecessary work. To access that bucket, we are only really using the two first bits of the directory entry.

Now, continuing with the same example, suppose we want to insert  $b$ , where  $h(b) = 11101$ . We see that we need to split the  $[w, z]$  bucket. But do we also need to double the directory? No! We do need to move from two-bit indexing to three-bit indexing, but our directory already allows for three-bit indexing. Thus, all we need to do is to split the bucket itself. The new table now looks like this:

000  
 001  
 010  
 011  
 100  $\rightarrow [a, y]$   
 101  $\rightarrow [x]$   
 110  $\rightarrow [z]$   
 111  $\rightarrow [w, b]$

This gives you the main idea behind extendible hashing. There are a few more technical details (for instance, we need some bookkeeping so that we know when we need to double the directory rather than just doing a bucket split). Deletions are basically the reverse of insertions, in that a bucket is merged with another if it becomes empty or underfull, and this may trigger a directory halving in some cases.

## 6.2 Questions

1. Warmup: Explain in words what a worst-case scenario for inserting would look like under extendible hashing. Also, give a small example to illustrate your scenario - show a few values being inserted into the hashtable, drawing pictures to show the directory and occupied buckets after each insertion.
2. For the above worst-case scenario, what is the complexity of a sequence of  $n$  inserts? Give a formal proof. You may assume that there are no collisions (no two elements hash to the same value). Hint: what is the largest that a directory can get with  $n$  inserts?.

3. Now, assume that when a directory doubles in size, the cost of this doubling is constant. Under this assumption, show that even in the worst-case scenario above, insertion takes  $O(1)$  amortized time. Again, you may assume no collisions. Prove formally, using induction, that a sequence of  $n$  worst-case inserts takes  $O(n)$  time. State very clearly any assumptions you make. Remember to take into account both the cost of the bucket splitting and the cost of the directory doubling.
4. Implementing the directory: The above constant-time doubling cannot be achieved with an array implementation of a directory. Briefly describe an implementation that allows constant-time doubling. If this implementation sacrifices the run-time complexity of some other operation which previously took constant time, mention what are the changes.

**To Submit:** Turn in a file `hashtable.txt` in simple ASCII format containing the solution to this problem. Note: this is a good problem to do as a warm-up for Prelim 2.