CS 312 Problem Set 5: Concurrent Language Interpreter

Assigned: March 29, 2004   Due: 11:59PM, April 15, 2004

Partner sign-up deadline: 11:59PM, April 3, 2004

1   Introduction

You have just won a contract to build self-replicating robots. They will be able to move, sense the world around them, and communicate with each other. One robot can even construct another robot and download into it code that will control its subsequent behavior. Thus, multiple robots may be operating simultaneously.

These robots will be controlled by programs written in a simple robot control language called RCL. The first step in fulfilling your contract is to build an RCL interpreter that can simulate the simultaneous operation of multiple robots, each running its own control program at the same time as the others. These programs are largely separate but can interact with each other through a global memory (presumably implemented using wireless communication). Thus, the RCL language is a concurrent programming language in which there can be multiple simultaneous threads of execution.

For Problem Set 5, you will implement this interpreter. In Problem Set 6, you will use your interpreter to implement a game that uses the robots.

2   Changes to problem set

The following changes have been made since it was released:

- Part 5, part (a) has been made a little easier. This was mentioned on the newsgroup earlier.

- Removed references to `hd` and `tl`. These operators do not exist in RCL.

- Specifications for `if` corrected.

- static.sml `isValue` recognizes pairs.

3   Instructions

You will do this problem set by modifying the source files provided in CMS, and submitting the program that results. The program that you submit 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.

Several parts of this assignment require that you implement a specification (and in some cases finish designing the specification). For each part, we expect you to provide the following documentation in an attached file for each part, as well as any additional information required for that part.

- **Specification changes.** If any changes to or refinements of the specifications given in the problem set are necessary, you should describe the changes you have made and justify them. Changes that are made with inadequate justification will result in point deductions even if the changes are reasonable.

- **Validation strategy.** Report how you validated your implementation. You will want to explain and justify your testing strategy, describing test cases or test harnesses that you wrote and reporting on the results from these tests.

- **Additional comments.** We welcome (but do not require) any additional feedback you would like to add; for example, tell us how much time the problem took, what was hard or easy about the problem, suggestions for what would make it a better problem.

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 half 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 it behooves you to start early on PS5 and understand the code given to you early. Solutions to PS5 will ***not*** be given out. PS5 and PS6 are also partner assignments. You are expected to find a partner to do this assignment with by April 1, and to sign up in CMS with that partner. If you do not have a partner and cannot find one, please e-mail the course staff and we will pair you up with someone else in the class.

4 The RCL language

The RCL language has some interesting features. First, it is a concurrent language in which multiple processes can be executing simultaneously. Second, it is an imperative language with memory heaps that can be updated. Third, the robots can perform input and output and interact with an external world.

To support concurrency, we will distinguish between different memories that the robots will affect. 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 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 interact with the external world by performing 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.

Another action is the expression `spawn` e, which launches another robot. The expression e provides the program that the newly created robot executes. Again, this action involves informing the external world so it can, for example, provide some information both to the old and to the new robot.

The behavior of the external world is not specified by the RCL 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

An RCL program for a single robot can contain the following expressions:

n	An integer constant, as in SML. Examples: ~ 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: $+, -, *, /, \mathtt{mod}, <, =$. 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.
let $id = e_1$ in e_2	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.
fn id => e	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.
id	Identifier. Must be contained inside a let or fn expression with the same identifier name, otherwise unbound identifier error will occur.
$e_1\ e_2$	Function application. Evaluates expression e_1 to a function fn id => e, expression e_2 to a value v_2, binds v_2 to the identifier id and uses the binding to evaluate e.
if e then e_1 else e_2	Similar to the ML if/then/else expression except that the result of expression e is tested for being greater than 0 (there are no booleans in RCL). Examples: if 1 then 1 else 2 returns 1, if 4<3 then 1 else 2 returns 2.
split e of e_1 else e_2	First evaluates expression e. If the result is a pair (v, v') then evaluates expression e_1 to a function f_1 and returns the result of function applications $((f_1\ v)\ v')$. If the result of e is a value v that is not a pair then evaluates e_2 and returns the result v_2.
	This operation is useful for emulating lists in RCL using pairs. Like pattern matching in ML, it gives the ability to treat the end of a list and a middle of a list differently.

typeof e	Evaluates e to a value v then returns an integer which indicates the type of v. A return value of 0 indicates that v is a function. Integers are 1; pairs are 2; references are 3.
lref e	Similar to the ML operation `ref`. 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.
gref e	Similar to `lref` 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).
! e	Evaluates expression e to location *loc* and returns the value stored at this location.
e_1 := e_2	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).
acquire e	This expression evaluates e to a location, *loc*, and, except as noted below, returns *loc*. If *loc* is in local memory, there is no side effect. If *loc* is in global memory and is not already locked, than the current process acquires a lock. If any other process already has a lock, this returns acquire v, and the process will continue to attempt the operation until the old lock is removed. All other cases are runtime errors.
release e	This expression evaluates e to a location, *loc*, and returns *loc*. If *loc* is in local memory, then there is no side effect. If *loc* is in global memory, and is locked by the current process, then *loc* is unlocked. All other situations are runtime errors.
do e	This expression is an example of an *action*. This is the main way for 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	Evaluates expression e to a function f, asks the external world for an expression e' and then starts a new robot by calling $f\ e'$. This gives the external world a chance to provide some information to the new robot.

In addition, there are some expressions that cannot be present in a robot source program, but can occur during evaluation of the program:

loc	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 be initially generated only by `lref` and `gref` expressions.
aid	Action identifier. Can initially be generated only by the external world as a result of evaluating an action (do, spawn or another action identifier). The

evaluation of aid depends entirely on the external world. The main purpose of this expression is to allow the external world to suspend the execution of a robot after performing an action.

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
- Pairs (v_1, v_2) (provided that v_1 and v_2 are values)
- Functions `fn` id `=>` e
- Locations loc

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 σ 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).
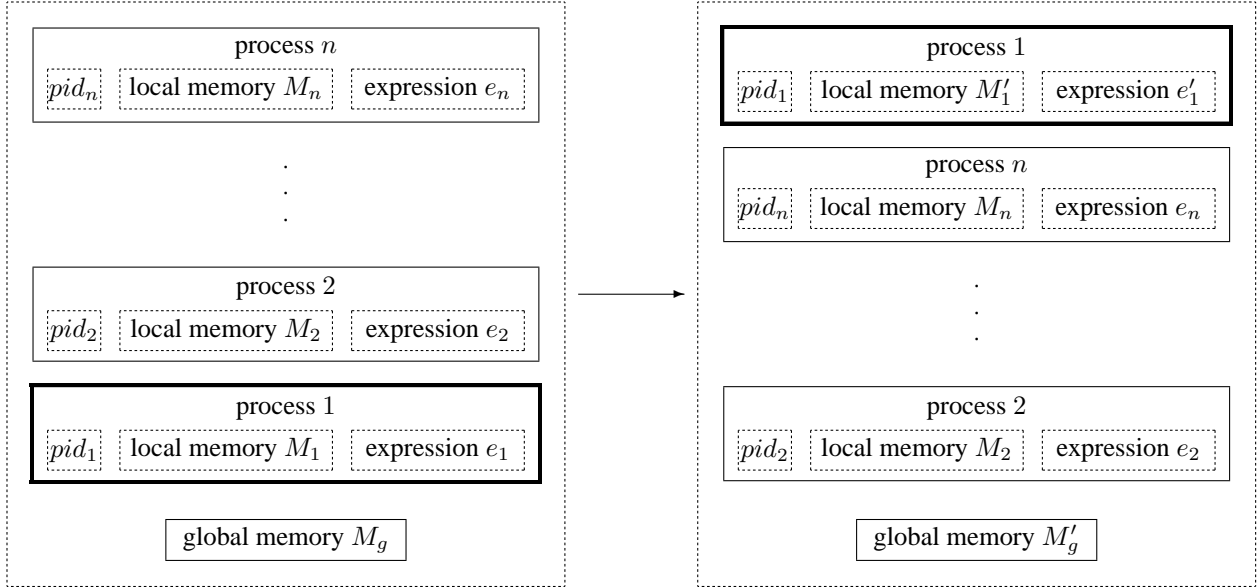
Figure 1: Single step of the interpreter on process 1. Expression e' is the result of a single evaluation step on e. Possible side effects include modifying local memory M_1 and global memory M_g

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. 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 memory or the global memory), 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.

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 `if` v `then` e_1 `else` e_2, `split` v `of` e_1 `else` e_2, `let` $id =$

v in e, fn id => 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 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.

$$
\begin{aligned}
\textit{unop } v &\longrightarrow v' & &\text{where } v' = \underline{\textit{unop } v} \\
v_0 \textit{ binop } v_1 &\longrightarrow v' & &\text{where } v' = \underline{v_0 \textit{ binop } v_1} \\
v; \ e &\longrightarrow e & & \\
\texttt{let } id = v \texttt{ in } e &\longrightarrow e\{v/id\} & & \\
(\texttt{fn } id \texttt{ => } e) \ v &\longrightarrow e\{v/id\} & & \\
\texttt{if } v \texttt{ then } e_1 \texttt{ else } e_2 &\longrightarrow e_1 & &v \in \{1,2,3\dots\} \\
\texttt{if } v \texttt{ then } e_1 \texttt{ else } e_2 &\longrightarrow e_2 & &\text{all other } v \\
\texttt{split } (v, v') \texttt{ of } e_1 \texttt{ else } e_2 &\longrightarrow (e_1 \ v) \ v' & & \\
\texttt{split } v \texttt{ of } e_1 \texttt{ else } e_2 &\longrightarrow e_2 & &\text{where } v \text{ is not a pair} \\
\texttt{typeof } v &\longrightarrow v' & &\text{where } v' \in \{0,1,2,3\} \text{ (see section 4.1)} \\
!loc &\longrightarrow v & &\text{where } loc \text{ is a location in the process local}
\end{aligned}
$$

memory or in the global memory, and v is the value stored at this location

$e\{v/id\}$ stands for the result of substitution of value v for all occurrences of identifier id in expression e.

With the exception of split, these reductions are similar to the reductions you have learned for SML. The split expression is like a simple version of case. In typical usage the expression e_1 would have the form fn x => fn y => e where e is an expression to be evaluated with x and y bound to the two components of the pair.

Now consider simple reductions which have side effects:

$\texttt{lref } v \longrightarrow loc$ — 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

$\texttt{gref } v \longrightarrow loc$ — 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

$\texttt{acquire } loc \longrightarrow loc$ — where loc is locked if it is a location in global memory that is not currently locked
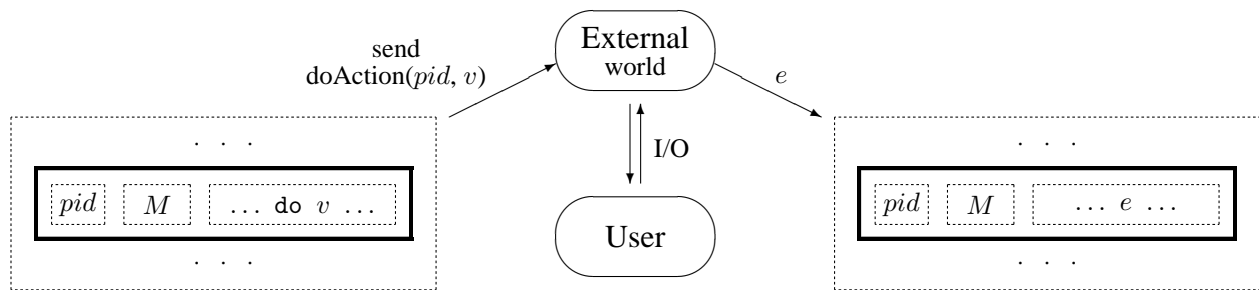
Figure 2: Evaluation of the do v expression

acquire loc $\longrightarrow$ acquire loc	where loc is locked if it is a location in global memory that is currently locked by another process	
release loc $\longrightarrow$ loc	where loc is released if it is a location in global memory that is locked by the current process	
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)	
spawn v $\longrightarrow$ e	where e is the expression returned by the external world Side effects: (1) select a fresh process identifier pid' (2) send spawn(pid, pid') to the external world (which will return two expressions (e, e')). As far as the interpreter is concerned, e and e' are arbitrary. (3) launch a new process with the process identifier pid', expression v e' and a copy of the process local memory (see Figure 3)	
aid $\longrightarrow$ e	where e is the expression returned by the external world Side effect: send actionID(pid, aid) to the external world (which will return an expression e)	

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

$$e_1 \; binop \; e_2 \longrightarrow v_1 \; binop \; e_2 \longrightarrow v_1 \; binop \; v_2 \longrightarrow v$$

4.6 Actions and the external world

Evaluating actions (that is, expressions do e, spawn e and aid) will cause sending certain events to the external world (Figures 2,3. We have provided an implementation of the external world in world/action.sml.

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:
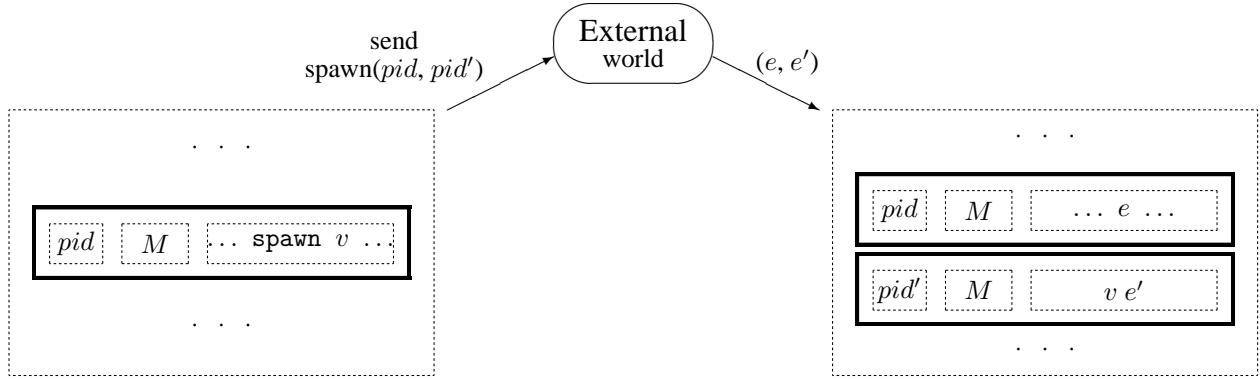
8

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

- `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 $c_1, \ldots, c_n$. 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.

In the current implementation of the external world, the `spawn` action prints a debugging message and returns two process identifiers.

Also, note that in the current implementation the external world never returns an action identifier. Therefore, the execution of processes will resume after `do` and `spawn` actions. If we wanted to delay processes, we would pick a fresh action identifier aid and return it as a result of `doAction` or `spawn`. The interpreter would then keep calling `actionID` for the corresponding process until the external world decides to resume it by returning an expression different from aid.

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, and a global memory that is shared by all the processes.

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:

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

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

9

$e_1 \longrightarrow e_1'$, with side effects that change the local memory M_1 to M_1' and the global memory M_g to M_g'. Then the effect of this step on the configuration as a whole is this:

$$\langle M_g, \langle pid_1, M_1, e_1 \rangle, \langle pid_2, M_2, e_2 \rangle, \ldots, \langle pid_n, M_n, e_n \rangle \rangle$$
$$\longrightarrow \quad \langle M_g', \langle pid_2, M_2, e_2 \rangle, \ldots, \langle pid_n, M_n, e_n \rangle, \langle pid_1, M_1', e_1' \rangle \rangle$$

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:

$$\langle M_g, \langle pid_1, M_1, v_1 \rangle, \langle pid_2, M_2, e_2 \rangle, \ldots, \langle pid_n, M_n, e_n \rangle \rangle$$
$$\longrightarrow \quad \langle M_g', \langle pid_2, M_2, e_2 \rangle, \ldots, \langle pid_n, M_n, e_n \rangle \rangle$$

Here M_g' is the global memory with all locks belonging to pid_1 released. This can be formalized with

$$\forall loc, M_g'(loc) \mapsto (v, p) \text{ iff } M_g(loc) \mapsto (v, p) \text{ and } p \neq pid_1$$

and

$$\forall loc, M_g'(loc) \mapsto (v, \emptyset) \text{ iff } M_g(loc) \mapsto (v, p) \text{ and } p \in \{pid_1, \emptyset\}$$

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)`). 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.

4.9 Formal description of the RCL language

This section has given an informal English description of the RCL language. Also available is a more formal mathematical description of the RCL language semantics, which is the canonical reference for the RCL language. The RCL language semantics may help if you are confused with certain parts of this section, however it assumes a lot of knowledge. Feel free to ask the course staff for help in understanding the formalisms used in the RCL language semantics.

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/configuration.sml`: definition of the configuration type (`Configuration.configuration`)

- `eval/evaluation.sml`: a single step of the main interpreter loop (`Evaluation.stepConfig`)

- `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/static.sig`, `static.sml`: well-formedness and consistency checking for expressions, processes and memories. Useful when debugging.

- `compat/compatibility07.sml`, `compatibility44.sml`: compatibility layer for different versions of the SML library. Useful for users of the working version.

- `rcl/tests/*.rcl`, `rcl/*.rch`: sample RCL programs and libraries

5.2 Running RCL code

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

$$\texttt{Debug.debug} \text{ "Your RCL program"}$$

which starts a single robot. 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 an RCL 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 RCL files

Complete RCL programs are stored in files with an `.rcl` extension. RCL programs and libraries may import libraries (extension `.rch`) with the `include` keyword. A library file contains RCL code which is incomplete but can be completed by appending a single expression. The files `rcl/logic.rch` and `rcl/lists.h` are examples.

5.4 String Literals

Although strings are not part of RCL 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, and there are three bugs in the code that is already written.

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/static.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 `doc1.txt`, so that we know where to look when we are grading.

Part 2: Robot design (7 pts)

Implement the following simple programs in RCL:

Write a RCL library that gives us functions that would be useful for manipulating binary search trees that map integers to integers. Each node of the BST should contain a key, the value that the key maps to, and two children (left and right). You must implement at least the following:

- `empty`: a value representing the empty tree

- `empty?`: a function returning 0 if the input is the empty tree, or 1 if the input is a non-empty tree

- `add`: a function that takes in a key, a piece of data, and a tree, and returns a new tree with the key-data mapping inside

- `lookup`: a function that takes in a key and a tree, and returns the data corresponding to the key.

- `rlookup`: a function that takes in a value and a tree, and searches for a key that maps to that value.

Write your own specifications for these functions based on the description above. All of these functions should execute in robot time of at most $O(h)$ where h is the height of the tree. For one of these functions, this time bound will require spawning robots that work in parallel to get the job done. The example `parfibo.rcl` may be a useful model.

You are allowed to use as many RCL helper functions that you want, as long as they are in the same file. You should also write implementation documentation that gives your abstraction function and rep invariant.

To Submit: A file `bst.rch` that implements the functions above and appropriate documentation that will convince us that your implementation works.

Part 3: Memory synchronization(10 pts)

We would like you to modify `memory.sml` to implement the functions `acquire`, `release` and `releaseAll`, each of which are specified in `memory.sig`.

To Submit: A completed copy of `memory.sml`.

Part 4: The garbage collector (13 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 RCL language. Occasionally the garbage collector will be used to clean up memory. For the purpose of RCL, 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: Provide an implementation of the signature `gc.sig` (do not change the signature) in the file `gc.sml`.

Part 5: Complexity analysis (10 pts)

A sorted array (or vector) is an appealing data structure for storing ordered data, because it offers the same $O(\lg n)$ lookup time as a balanced binary tree but has a compact representation and a good constant factor in front of $\lg n$. Unfortunately it doesn't support fast insertion.

Mort Ization has a idea for how to build a faster mutable ordered set abstraction. Instead of storing all the elements in the sorted array, he will maintain a separate short linked list of up to m elements, where m is some function of n yet to be determined.

```
type set = { sorted:  element array ref,
             recent:  element list ref }
```

When the data structure is searched, both the list `recent` (of length m) and the array `sorted` (of length n) are traversed. When an element is added to the data structure, it is appended to the list in constant time. If the `recent` list becomes longer than m elements, the m elements are sorted using a mergesort and then merged in linear time with the n elements, which are already in order.

a. Mort wants to get $O(\lg n)$ insertion, so he sets $m = \lg n$. What is the amortized asymptotic complexity for a series of n lookups? What about m inserts? Show your work. (For the insert part you may assume that the list `recent` is initially empty.)

b. How should Mort set the value of m (as a function of n) to get the best possible asymptotic amortized complexity for a series of n inserts and n lookups in any order? What is that resulting complexity? Explain.

To Submit: Turn in a file `complexity.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.

Optional Part 6: Karma

(Only good karma is received for doing these problems.)

- There are many improvements that could be done to the debugger. Options include but are not limited to: changing memory limits, better printing of AST's, more control over printing of the configurations, breakpoints, etc.

- Write a string library in RCL (see `hello-world.rcl` for an example of using strings) and add a new action to let the user input a string for use by the RCL program. Write an interesting program that does string manipulation.

- Implement red-black trees in RCL.

- Add vectors to RCL and implement hash tables.