

# Robot Control Language Semantics

April 18, 2004

This is the PS5 language specification.

## 1 Syntax

$n \in \mathbf{Z}$	=	$\{\dots, -2, -1, 0, 1, 2, \dots\}$
$loc \in \mathbf{Loc}$		Locations in memory—appear only in programs being evaluated
$aid \in \mathbf{Aid}$		Action identifiers—appear only in programs being evaluated
$v \in \mathbf{Value}$	::=	$n \mid (v_1, v_2) \mid \mathbf{fn} \ id \Rightarrow \ e \mid \ loc$
$e \in \mathbf{Expr}$	::=	$v \mid \ id \mid \ unop \ e_0 \mid \ e_0 \ binop \ e_1 \mid \ e_0 ; \ e_1 \mid \ action \mid \ \mathbf{lref} \ e$ $\mid \ \mathbf{gref} \ e \mid \ !e \mid \ e_0 := \ e_1 \mid \ (e_0, e_1) \mid \ \mathbf{let} \ id = e_0 \ \mathbf{in} \ e_1 \mid \ \mathbf{if} \ e_0 \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2$ $\mid \ e_1 \ e_2 \mid \ \mathbf{split} \ e_0 \ \mathbf{of} \ e_1 \ \mathbf{else} \ e_2 \mid \ aid$
$binop$	::=	$+ \mid - \mid * \mid / \mid \mathbf{mod} \mid < \mid =$
$unop$	::=	$\sim \mid \mathbf{rand} \mid \mathbf{acquire} \mid \mathbf{release} \mid \mathbf{typeof}$
$action$	::=	$\mathbf{spawn} \ e \mid \ \mathbf{do} \ e$

The operation `do` interacts with the external world; essentially it is the single I/O operation. The significance of particular inputs and outputs will be defined in PS6.

## 2 Operational semantics

### 2.1 Configuration

A memory  $\sigma$  is a partial function from locations to values and process ids.

$$\sigma \in \mathbf{Memory} = \mathbf{Loc} \rightarrow \mathbf{Value} \times (\mathbf{Pid} \cup \emptyset)$$

The symbol  $\rightarrow$  indicates a partial function that is defined only on some subset of  $\mathbf{Loc}$  denoted by  $\text{dom}(\sigma)$ . We write  $\sigma[loc \mapsto (v, p)]$  to mean the function that is exactly like  $\sigma$  except that it maps  $loc$  to  $(v, p)$ , possibly extending the domain to include  $loc$ . The process idea stored at each memory location identifies which RCL program is currently holding a lock on that location. Locks are useful for controlling access to locations in memories shared by multiple RCL processes.

A configuration for a single process of the machine (that is, a single robot) is a pair  $\langle e, \sigma \rangle$ , and the configuration for the entire machine running  $n$  concurrent processes is as follows:

$$\langle \sigma_g, \langle e_1, \sigma_1 \rangle, \dots, \langle e_n, \sigma_n \rangle \rangle$$

where  $\sigma_g$  is the shared memory between the processes, and  $\sigma_m$  is the local memory for process  $m$ .

Every local memory has a special location  $pid$  where it stores the unique *process identifier* for the robot running the process. The process identifier is a positive integer. The process identifier is not the same thing as the index  $m$ , because the index of a process changes in the configuration as it evaluates.

## 2.2 Simple reductions

The following simple expression reductions do not involve memory and are similar to the evaluation reductions given for SML:

$$\begin{array}{ll}
unop\ v \longrightarrow v' & \text{where } v' = \underline{unop\ v} \\
v_0\ binop\ v_1 \longrightarrow v' & \text{where } v' = \underline{v_0\ binop\ v_1} \\
v; e \longrightarrow e & \\
\text{let } id = v \text{ in } e \longrightarrow e\{v/id\} & \\
\text{if } v \text{ then } e_1 \text{ else } e_2 \longrightarrow e_1 & \text{where } v \in Z^+ \\
\text{if } v \text{ then } e_1 \text{ else } e_2 \longrightarrow e_2 & v \notin Z^+ \\
\text{fn } id \Rightarrow e(v) \longrightarrow e\{v/id\} & \\
\text{split } (v_1, v_2) \text{ of } e_1 \text{ else } e_2 \longrightarrow (e_1\ v_1)\ v_2 & \\
\text{split } v \text{ of } e_1 \text{ else } e_2 \longrightarrow e_2 & \text{where } v \text{ is not a pair} \\
\text{typeof } (\text{fn } id \Rightarrow e) \longrightarrow 0 & \\
\text{typeof } v \longrightarrow 1 & \text{where } v \in Z \\
\text{typeof } (v_1, v_2) \longrightarrow 2 & \\
\text{typeof } loc \longrightarrow 3 &
\end{array}$$

If a simple reduction  $e \longrightarrow e'$  is legal, then a process  $\langle e, \sigma \rangle$  can step to  $\langle e', \sigma \rangle$ ; the local memory is unchanged by the step.

In addition, a process  $\langle e, \sigma \rangle$  with process id,  $p$ , can perform the following reductions that involve memory:

$$\begin{array}{ll}
\langle \sigma_g, \langle \text{lref } v, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g, \langle loc, \sigma[loc \mapsto (v, \emptyset)] \rangle \dots \rangle & loc \notin \text{dom}(\sigma) \\
\langle \sigma_g, \langle !loc, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g, \langle v, \sigma \rangle \dots \rangle & \sigma(loc) = (v, \emptyset) \text{ or } \sigma_g(loc) = (v, p') \\
\langle \sigma_g, \langle loc := v, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g, \langle v, \sigma[loc \mapsto (v, \emptyset)] \rangle \dots \rangle & loc \in \text{dom}(\sigma) \\
\langle \sigma_g, \langle loc := v, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g[loc \mapsto (v, p')], \langle v, \sigma \rangle \dots \rangle & \sigma_g(loc) \mapsto (v', p') \text{ and } p' \in \{\emptyset, p\} \\
\langle \sigma_g, \langle \text{acquire } loc, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g, \langle loc, \sigma \rangle \dots \rangle & loc \in \text{dom}(\sigma) \\
\langle \sigma_g, \langle \text{acquire } loc, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g[loc \mapsto (v, p)], \langle loc, \sigma \rangle \dots \rangle & \sigma_g(loc) \mapsto (v, \emptyset) \\
\langle \sigma_g, \langle \text{acquire } loc, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g, \langle \text{acquire } loc, \sigma \rangle \dots \rangle & \sigma_g(loc) \mapsto (v, p') \text{ and } p' \neq p \\
\langle \sigma_g, \langle \text{release } loc, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g, \langle loc, \sigma \rangle \dots \rangle & loc \in \text{dom}(\sigma) \\
\langle \sigma_g, \langle \text{release } loc, \sigma \rangle \dots \rangle \longrightarrow \langle \sigma_g[loc \mapsto (v, \emptyset)], \langle loc, \sigma \rangle \dots \rangle & \sigma_g(loc) \mapsto (v, p)
\end{array}$$

## 2.3 Evaluation contexts

The following grammar defines where a reduction can occur during evaluation. It generates language terms in which there is a single hole  $[\cdot]$  at the place where a reduction can be performed.

$$\begin{array}{l}
\mathcal{C} ::= [\cdot] \mid unop\ \mathcal{C} \mid \mathcal{C}\ binop\ e \mid v\ binop\ \mathcal{C} \mid \mathcal{C}; e \mid \text{do } \mathcal{C} \mid \text{spawn } \mathcal{C} \\
\mid \text{lref } \mathcal{C} \mid \text{gref } \mathcal{C} \mid !\mathcal{C} \mid \mathcal{C} := e \mid loc := \mathcal{C} \\
\mid \text{let } id = \mathcal{C} \text{ in } e \mid (\mathcal{C}, e) \mid (v, \mathcal{C}) \mid \text{if } \mathcal{C} \text{ then } e_1 \text{ else } e_2 \\
\mid \mathcal{C}\ e \mid v\ \mathcal{C} \mid \text{split } \mathcal{C} \text{ of } e_1 \text{ else } e_2
\end{array}$$

Given a reduction  $e \longrightarrow e'$ , a single process can perform the evaluation step  $\mathcal{C}[e] \longrightarrow \mathcal{C}[e']$ , where  $\mathcal{C}$  is a single evaluation context as described by the above grammar and  $\mathcal{C}[e]$  represents that evaluation context with the term  $e$  substituted for the single hole appearing in  $\mathcal{C}$ .

For example, the above grammar can derive the term with a hole  $\text{let } x = [\cdot] \text{ in } (\text{fn } (y) \Rightarrow y)$ , and  $2 + 2 \longrightarrow 4$  is a legal reduction; therefore, the following is a legal single-process evaluation step:

$$\text{let } x = 2 + 2 \text{ in fn } (y) \Rightarrow y \quad \longrightarrow \quad \text{let } x = 4 \text{ in } (\text{fn } (y) \Rightarrow y)$$

## 2.4 Concurrency

As discussed earlier, the state of all the robot programs is represented by a global configuration

$$\langle \sigma_g, \langle e_1, \sigma_1 \rangle, \dots, \langle e_n, \sigma_n \rangle \rangle$$

Steps of the machine therefore have the form

$$\langle \sigma_g, \langle e_1, \sigma_1 \rangle, \dots, \langle e_n, \sigma_n \rangle \rangle \longrightarrow \langle \sigma'_g, \langle e'_1, \sigma'_1 \rangle, \dots, \langle e'_n, \sigma'_n \rangle \rangle$$

To ensure round-robin scheduling, only the leftmost process takes a step, and the resulting process is always placed at the *end* of the list of processes. Thus, all of the single-process reduction rules can be lifted to apply to the entire machine using the following rule:

$$\frac{\langle e, \sigma_1 \rangle \longrightarrow \langle e', \sigma'_1 \rangle}{\langle \sigma_g, \langle \mathcal{C}[e], \sigma_1 \rangle, \langle e_2, \sigma_2 \rangle, \dots, \langle e_n, \sigma_n \rangle \rangle \longrightarrow \langle \sigma_g, \langle e_2, \sigma_2 \rangle, \dots, \langle e_n, \sigma_n \rangle, \langle \mathcal{C}[e'], \sigma'_1 \rangle \rangle}$$

A few more rules are needed for evaluation of expressions that access the that create or destroy a process:

$$\begin{aligned} \langle \sigma_g, \langle \mathcal{C}[\text{spawn } v_s], \sigma_1 \rangle, \langle e_2, \sigma_2 \rangle, \dots, \langle e_n, \sigma_n \rangle \rangle &\longrightarrow \langle \sigma_g, \langle e_2, \sigma_2 \rangle, \dots, \langle e_n, \sigma_n \rangle, \langle \mathcal{C}[e'_1], \sigma_1 \rangle, \langle v_s \ e'_2, \sigma_1 [pid \mapsto p] \rangle \rangle \\ &\text{where } p \text{ is a fresh process identifier and where } (e'_1, \\ &e'_2) \text{ is the response of the external world to the} \\ &\text{request } \text{Spawn}(\sigma_1(pid), p). \\ \langle \sigma_g, \langle v_1, \sigma_1 \rangle, \langle e_2, \sigma_2 \rangle, \dots, \langle e_n, \sigma_n \rangle \rangle &\longrightarrow \langle \sigma_g, \langle e_2, \sigma_2 \rangle, \dots, \langle e_n, \sigma_n \rangle \rangle \\ &\text{Effect: send } \text{Halted}(\sigma_1(pid)) \text{ to external world} \end{aligned}$$

## 3 Interaction with the external world

The last rule says that a process that has evaluated to a value simply disappears (it has halted). This rule and the rules for **spawn** and **do** require some communication with external world to inform it that a robot has been created, destroyed, or has taken some action. It is then up to the external world to give the robot(s) involved in the action an expression that is used in subsequent evaluation.

For example, in the **spawn** expression one natural choice for the expressions  $e'_1$  and  $e'_2$  is the process identifiers of the other robot. However, the external world may provide different information.

It may be necessary for the external world to delay execution of a process that has performed an action. It can achieve this by returning a value that is an action identifier *aid*. When asked to evaluate an action identifier, the interpreter always asks the external world to provide a value for it. This value then replaces the action identifier and evaluation proceeds. The external world can delay the interpreter as long as necessary by simply providing the same action identifier, causing the process to be placed on the back of the process queue unchanged.

$$\begin{aligned} \langle \text{do } v, \sigma \rangle &\longrightarrow \langle e, \sigma \rangle && \text{where } e \text{ is the result of } \text{Do}(\sigma(pid), v) \\ \langle \text{aid}, \sigma \rangle &\longrightarrow \langle \text{aid}, \sigma \rangle && \text{if external world says no step may be taken} \\ \langle \text{aid}, \sigma \rangle &\longrightarrow \langle e, \sigma \rangle && \text{if external world responds to request } \text{ActionId}(\sigma(pid), \text{aid}) \text{ with result } e. \end{aligned}$$

Because an action identifier *aid* is a legal expression, the second rule actually subsumes the first one.

It is up to the external world to keep track of the various action identifiers and what processes and actions they correspond to. The interpreter has no control over this aspect of evaluation.