

## 5.7 The Challenge Language

This section contains the complete K semantic definition of CHALLENGE, a non-trivial experimental programming language that we propose here in order to challenge existing or future language definitional frameworks. All language constructs of CHALLENGE can be found in existing programming languages, possibly in a more general form and possibly using a syntax slightly different from ours. In short, none of the CHALLENGE features is artificially crafted and the language, as a whole, is plausible and powerful.

We chose features that already exist in popular programming languages to avoid criticisms, from proponents of frameworks that cannot handle the proposed features, that those features are useless and thus do not deserve attention. Our standpoint is that the very fact that a language feature exists in a mainstream programming language is sufficient evidence that the feature is practical and useful, so an ideal semantic framework should unarguably support it: while a language designer may subjectively chose to include such a feature in her language or not, a language-design framework designer aiming at an ideal framework, as we do, has no choice and should simply support it. The CHALLENGE features, and particularly their combination, have been carefully chosen to satisfy two criteria:

1. To capture *the essence* of several conventional language concepts; if a semantic framework cannot handle a certain CHALLENGE feature then it is clear that that framework cannot handle a potentially wide range of desirable language features;
2. To illustrate how one can define a complex language *modularly* in K: each new feature is added to the language without changing any of the already defined features. We here assume that one knows upfront the entire language and that at each step one takes the globally best design decisions. In Section 5.8 we discuss a more complex language design scenario.

```

Exp ::= Bool | Int | Real | VarId
      | Exp + Exp [strict]
      | Exp * Exp [strict]
      | Exp / Exp [strict]
      | Exp <= Exp [seqstrict]
      | not Exp [strict]
      | Exp and Exp [strict(1)]
      | ++ VarId
      | Stmt ; Exp
      | malloc Exp [strict]
      | & VarId
      | * Exp [strict]
      | λList(0){ VarId } . Exp
      | Exp List(0){ Exp } [strict]
      | μ VarId . Exp
      | callcc Exp [strict]
      | randomBool
      | new-agent Stmt
      | me
      | parent
      | receive
      | receive-from Exp [strict]
      | quote Exp
      | unquote Exp
      | eval Exp [strict]

Stmt ::=
      | Stmt ; Stmt
      | { vars List{ VarId } ; Stmt }
      | if Exp then Stmt else Stmt [strict(1)]
      | while Exp do Stmt
      | output Exp [strict]
      | Exp := Exp [strict(2)]
      | aspect Stmt
      | spawn Stmt
      | acquire Exp [strict]
      | free Exp [strict]
      | release Exp [strict]
      | rv Exp [strict]
      | send-async Exp Exp [strict]
      | send-synch Exp Exp [strict]
      | halt-thread
      | halt-agent
      | halt-system

```

Figure 5.13: K annotated CHALLENGE syntax

### 5.7.1 Challenge Syntax

Figure 5.13 shows the syntax of CHALLENGE, which is split into two syntactic categories: expressions (*Exp*) and statements (*Stmt*). To save space and give the reader an early intuition on the evaluation strategies of the various constructs, the syntax of CHALLENGE in Figure 5.13 is already annotated with K strictness attributes (Section 5.5). The language starts with simple arithmetic expressions, then gradually grows in complexity by adding expressions with side effects, memory allocation and pointers, functions, recursion, call with current continuation, non-determinism, multi-threading with shared memory and thread synchronization, agents and both synchronous and asynchronous inter-agent communication, and in the end code generation. We next informally explain the CHALLENGE language constructs and briefly give our reasons for including them in CHALLENGE. Section 5.8 discusses in depth how the various semantic frameworks in Chapter 3 encounter difficulties in supporting some of the CHALLENGE features, and Chapters 10 and 12 elaborate on the meaning of these features. We here assume the reader familiar with them.

Expressions can be both arithmetic and boolean, and include both integer and real numbers. The arithmetic/boolean expression constructs and their evaluation strategies are similar to those of IMP, which are discussed in Section 5.2. However, we keep the integer and real numbers separate and all the arithmetic and relational operations are overloaded to work with both. Together with the top-level split of syntax into expressions and statements, we believe that the above make CHALLENGE stress relatively well the capabilities of the semantic framework to deal with multiple syntactic categories, both disjoint and included in each other, as well as with operator overloading.

Expressions can have side effects, both minimal (with variable increment, e.g., “ $++x$ ”) and arbitrarily complex (“ $s; e$ ” means “evaluate statement  $s$ , then evaluate expression  $e$  in the resulting state and return its value”). Side effects are very common in most programming languages and their inclusion may require non-modular changes when one uses certain semantic frameworks.

CHALLENGE allows to allocate memory with a `malloc` expression construct and allows to deallocate memory with a `free` statement construct. Like in C, `malloc` evaluates to the location where the allocated block starts, and `free` can only be called on a location which has been previously allocated with a `malloc` (so that the number of locations to free is known). Like C, CHALLENGE also allows unrestricted pointer arithmetic, allows accessing the address at which a variable is allocated with the expression construct `&x`, allows dereferencing locations with the expression construct `*p`, and allows writing values to locations using statements of the form `*p := e`. Unlike in C, we prefer to use `:=` instead of `=` for assignment, to avoid notational confusion with the equality symbol of equational logic that we use many places in this book. Not all programming languages allow as explicit and direct memory access as CHALLENGE. Nevertheless, a semantic framework able to naturally deal with the memory operations of CHALLENGE can likely support arbitrarily complex memory operations in programming language semantic definitions.

Any semantic framework worth its salt must support definitions of languages with functions. Since higher-order functions tend to be more complex than other functions and since there are many higher-order functional languages, it suffices to include only higher-order functions in CHALLENGE. We do it by means of  $\lambda$ -expressions (Section 4.4). To stress the framework’s capability to syntactically and semantically deal with lists of terms, we allow  $\lambda$ -abstractions taking arbitrary lists of parameters ( $\lambda \mathbf{List}^0\{VarId\}.Exp$ ) and  $\lambda$ -applications taking arbitrary lists of arguments ( $Exp \mathbf{List}^0\{Exp\}$ ) in CHALLENGE, with  $()$  the unit of these list constructs. For example,  $(\lambda().e)()$ ,  $(\lambda x.e)e'$ , and  $(\lambda(x_1, x_2).e)(e_1, e_2)$  are all well-formed CHALLENGE expressions when  $x, x_1, x_2 \in VarId$  and  $e, e', e_1, e_2$  are well-formed expressions. Note that we defined the  $\lambda$ -application to be strict both in its

first argument (expected to evaluate to a function) and in its second argument; its second argument is a list of expressions, so the application is implicitly strict in all these expressions in the list (i.e., each of the expressions in the list will be evaluated before the function is applied; how this works is explained in Section 5.5); therefore, CHALLENGE is call-by-value (see Chapter 12).

Since recursion is a fundamental design and semantic concept in programming languages, and since recursion can have many apparently different facets, CHALLENGE allows three different means to achieve recursion: (1) through implicit recursion obtained using conventional  $\lambda$ -calculus encodings of fixed-point combinators (see Section 4.4), which only work in untyped settings; (2) through iteration using the provided `while` loop statement construct (`while Exp do Stmt`); and (3) through explicit use of a provided fixed-point  $\mu$  construct ( `$\mu$  VarId. Exp`). Even though in theory only one approach to recursion suffices, we strongly believe that a language definitional framework should naturally support all three approaches in practice, to avoid artificial encodings of one into another.

Many languages have control-intensive constructs, such as abrupt termination, exceptions, break/continue of loops, and even `callcc` (from “call with current continuation”). Therefore, any worthwhile language definitional framework must support language constructs that abruptly change the execution flow. Since `callcc` is as complex a control-intensive construct as it can be, we included it as part of the CHALLENGE language (`callcc Exp`). We also included several abrupt termination statements (for threads, agents and program), which we discuss shortly. Other control constructs, such as exceptions and break/continue of loops, are defined and discussed in Chapters 10 and 12.

Some semantic frameworks, particularly those of a denotational nature, cannot handle non-determinism and/or concurrency satisfactorily: while they still allow in some cases to define a collecting semantics (i.e., a semantics that collects all behaviors), such a semantics is inappropriate for obtaining an interpreter for the defined language when executing it. To illustrate a framework’s capability to deal with non-determinism, CHALLENGE includes the simplest imaginable non-deterministic construct, namely a `randomBool` boolean constant which non-deterministically evaluates to `true` or `false` each time it is evaluated.

Concurrency and distribution of software in general and of programming languages in particular is not anymore the exception, but the norm. Thus, any practical language definitional framework should provide good support for concurrency. Therefore, CHALLENGE contains language constructs that stress the underlying framework in several ways with respect to concurrency. It allows to dynamically create and terminate both agents (`new-agent Stmt`) and threads (`spawn Stmt`). The threads live inside agents and the threads inside an agent communicate and synchronize with each other by means of shared memory, re-entrant acquisition (`acquire Exp`) and releasing (`release Exp`) of locks (any value can serve as a lock), and by rendez-vous barriers (`rv Exp`). Agents communicate with each other by means of synchronous (`send-synch Exp Exp`) and asynchronous (`send-asynch Exp Exp`) message-send commands (the first argument is the receiving agent’s name and the second argument is value being passed), and by targeted (`receive-from Exp`, where the argument is the sending agent’s name) and non-targeted (`receive`) expression constructs. If a language semantic framework supports the above common concurrency features, then it likely can support any degree of concurrency that a language designer may want to include in her language.

To keep a tight connection between the concurrency features and the rest of the language, CHALLENGE includes several other features that do not dictate the semantics of the concurrent ones but interact with them. For example, since there are three types of executing entities, namely threads, agents, and the entire system, it makes sense to have three abrupt termination statements, one for the current thread (`halt-thread`), one for the current agent (`halt-agent`) and one for

the entire system (`halt-system`); the former also releases the resources (i.e., locks) that the thread held. We also included in `CHALLENGE` an output that is shared by all threads and all agents (via the `output` statement construct). The previously discussed constructs `malloc`/`free` act on the memory of the current agent only (so the agent's threads compete on memory), and `callcc` acts on the computation/continuation of the current thread only.

Aspect-oriented programming is an important increasingly growing trend in programming languages. There are many different approaches to aspects (e.g., dynamic versus static aspects) and we cannot possibly include all of them in `CHALLENGE`. We only include one aspect-oriented feature, namely executing an aspect statement whenever a function is invoked. To make it more interesting, we combine static with dynamic semantics of aspects in that we statically weave the current aspect before the body of a function at function declaration time, but we allow one to dynamically change the aspect code. This can be done by executing a statement “`aspect s`”, which sets the argument statement `s` as the aspect of the current thread. Many variations are possible, particularly in combination with concurrency, but we prefer to keep `CHALLENGE` minimal. We added the particular aspect feature above for two reasons: first, executing code when a function is called is one of the most common aspects; second, it may be inconvenient to give it semantics in purely syntactic approaches (one needs to distinguish between weaved and not-yet-weaved functions) or in approaches which reduce desired language features to their “builtin” similar features (e.g., desired functions to their builtin functions, desired recursion to their fixed-points, etc.).

Finally, there are several programming languages providing support for dynamic code generation and execution of it (e.g., Lisp, Scheme, Jumbo, etc.). We therefore add support for code generation to `CHALLENGE`. We borrow our terminology from Scheme. The expression “`quote(e)`” freezes expression `e` into a special code value, and “`eval(e)`” first evaluates `e` to a code value, say `c`, and then unfreezes `c` and evaluates it in the current environment. To allow for arbitrarily complex code generation and reuse, one can use expressions like “`unquote(e)`” inside a quoted expression, with the following semantics: evaluate `e` in the current environment to a code value, then plug that code value into the quoted expression. Quoting of code can be nested and an expression must be unquoted as many times as quoted in order to be evaluated. Generation of code can be challenging or impossible to define modularly in most frameworks, because they drive the semantics according to the structure of the syntax, where the individuality of each language construct is important. Thus, most likely one needs to define quoting/unquoting for each language construct, which is non-modular. In `K`, we can define the semantics of `quote`/`unquote`/`eval` with 13 language independent rules.

## 5.7.2 Challenge Configuration

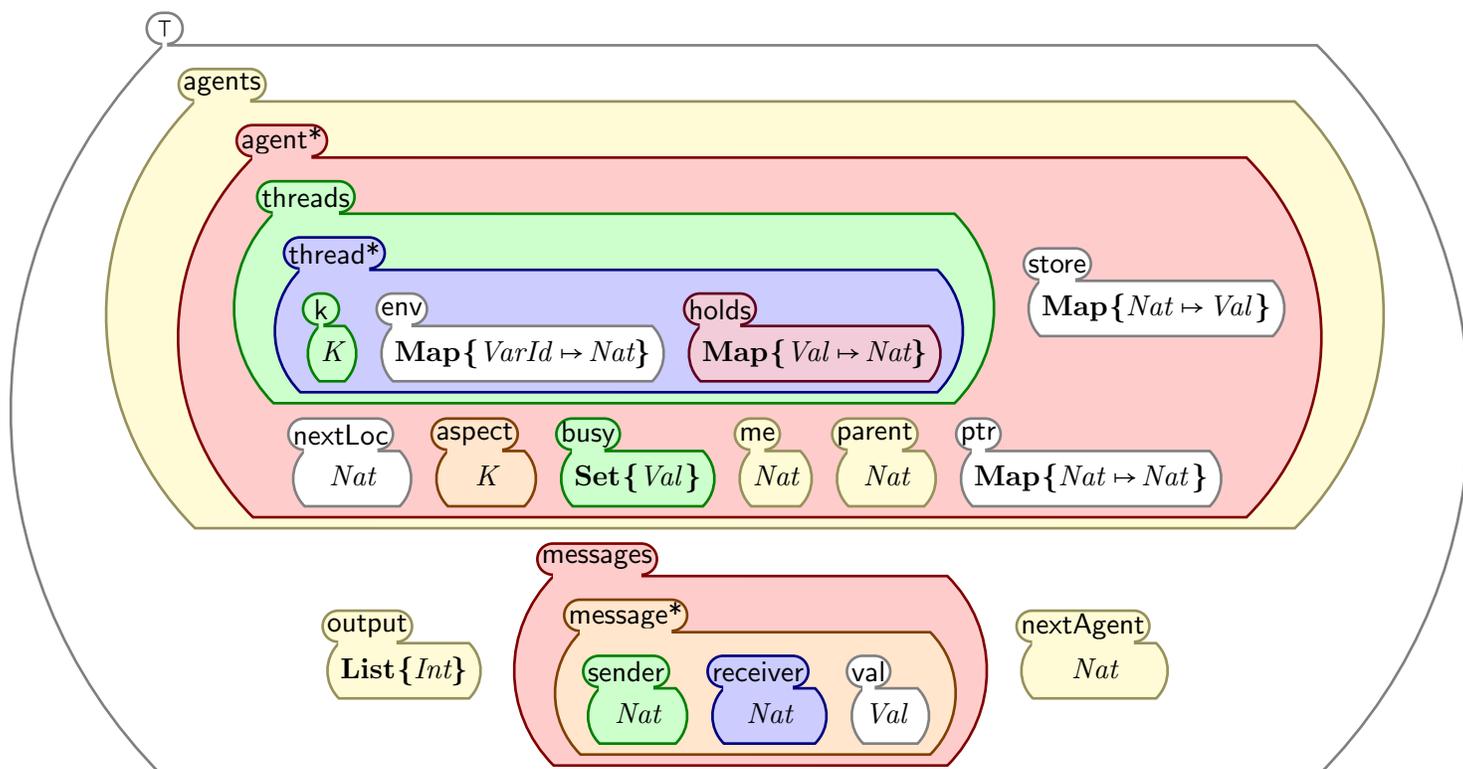
Figure 5.14 shows the complete `K`-configuration of `CHALLENGE`, both in term representation and in graphical representation. The cell contents is self-explanatory, but more detail will be given as we define the various language constructs in Section 5.7.3

## 5.7.3 Challenge Semantics

We will only give here rules for constructs having a different behavior than the one already presented for `IMP++`.

**Reading from the store** When a name is matched at the top of computation, it can be replaced by value `V`, provided that the mapping of `X` to location `L` can be matched in the environment and

$$\begin{aligned}
\text{Configuration}_{\text{CHALLENGE}} &\equiv \langle \text{Agents}_{\text{CHALLENGE}} \langle \text{List}\{Int\} \rangle_{\text{output}} \text{Messages}_{\text{CHALLENGE}} \langle Nat \rangle_{\text{nextAgent}} \rangle_{\top} \\
\text{Agents}_{\text{CHALLENGE}} &\equiv \left\langle \left\langle \left\langle \left\langle \left\langle Nat \right\rangle_{\text{nextLoc}} \langle K \rangle_{\text{aspect}} \langle \text{Set}\{Val\} \rangle_{\text{busy}} \right\rangle_{\text{agent}^*} \right\rangle_{\text{threads}} \right\rangle_{\text{agent}^*} \right\rangle_{\text{agents}} \\
\text{Threads}_{\text{CHALLENGE}} &\equiv \langle \langle \langle K \rangle_k \langle \text{Map}\{VarId \mapsto Nat\} \rangle_{\text{env}} \langle \text{Map}\{Val \mapsto Nat\} \rangle_{\text{holds}} \rangle_{\text{thread}^*} \rangle_{\text{thread}} \\
\text{Messages}_{\text{CHALLENGE}} &\equiv \langle \langle \langle Nat \rangle_{\text{sender}} \langle Nat \rangle_{\text{receiver}} \langle Val \rangle_{\text{val}} \rangle_{\text{message}^*} \rangle_{\text{messages}}
\end{aligned}$$



- |  |  |
|--|--|
| $\langle \rangle_{\top}$ : top cell, holding everything                | $\langle \rangle_{\text{busy}}$ : holds agent's busy locks               |
| $\langle \rangle_{\text{agents}}$ : holds all the agents               | $\langle \rangle_{\text{me}}$ : holds agent's id                         |
| $\langle \rangle_{\text{agent}^*}$ : holds one agent, can multiply     | $\langle \rangle_{\text{parent}}$ : holds agent's parent if              |
| $\langle \rangle_{\text{threads}}$ : holds all agent's threads         | $\langle \rangle_{\text{ptr}}$ : holds agent's store allocation map      |
| $\langle \rangle_{\text{thread}^*}$ : holds one thread, can multiply   | $\langle \rangle_{\text{output}}$ : holds the system output              |
| $\langle \rangle_k$ : holds thread's computation                       | $\langle \rangle_{\text{messages}}$ : holds all the pending messages     |
| $\langle \rangle_{\text{env}}$ : holds thread's environment            | $\langle \rangle_{\text{message}^*}$ : holds one message, can multiply   |
| $\langle \rangle_{\text{holds}}$ : holds thread's locks                | $\langle \rangle_{\text{sender}}$ : holds the message sender's id        |
| $\langle \rangle_{\text{store}}$ : holds agent's store                 | $\langle \rangle_{\text{receiver}}$ : holds the message receiver's id    |
| $\langle \rangle_{\text{nextLoc}}$ : holds the next available location | $\langle \rangle_{\text{val}}$ : holds the message's value               |
| $\langle \rangle_{\text{aspect}}$ : holds agent's aspect               | $\langle \rangle_{\text{nextAgent}}$ : holds the next available agent id |

Figure 5.14: The configuration of the CHALLENGE language in both term (top) and graphical (middle) representation, with short explanation of cell contents (bottom)

the mapping of L to V can be matched in the store:

$$\frac{\langle x \dots \rangle_k \langle \dots x \mapsto l \dots \rangle_{\text{env}} \langle \dots l \mapsto v \dots \rangle_{\text{store}}}{v}$$

**Statement Block with variable declaration** When variables are declared, new locations are allocated in the environment. We use the ‘nextLoc’ cell to hold the next available location; also, we use  $\text{Rho}[X \leftarrow L]$  to map X to L in Rho. Once the executions of the statements in the block is completed, the old environment must be resumed.

$$\frac{\langle \{\text{vars } xl; s\} \dots \rangle_k \langle \frac{\rho}{\rho[n..n +_{Int} |xl| -_{Int} 1/xl]} \rangle_{\text{env}} \langle \frac{\sigma}{\sigma[0/n..n +_{Int} |xl| -_{Int} 1]} \rangle_{\text{store}} \langle \frac{n}{n +_{Int} |xl|} \rangle_{\text{nextLoc}}}{s \rightsquigarrow \text{env}(\rho)}$$

where

$$\frac{\langle \text{env}(\rho) \dots \rangle_k \langle \_ \rangle_{\text{env}}}{\rho}$$

**Side effects: increment**

$$\frac{\langle ++x \dots \rangle_k \langle \dots x \mapsto l \dots \rangle_{\text{env}} \langle \dots l \mapsto \frac{i}{i +_{Int} 1} \dots \rangle_{\text{store}}}{i +_{Int} 1}$$

**References** Obtaining the reference location of a name:

$$\frac{\langle \&x \dots \rangle_k \langle \dots x \mapsto l \dots \rangle_{\text{env}}}{l}$$

Dereferencing a location:

$$\frac{\langle *l \dots \rangle_k \langle \dots l \mapsto v \dots \rangle_{\text{store}}}{v}$$

Memory allocation and free:

$$\frac{\langle \text{malloc}(n) \dots \rangle_k \langle \dots \frac{\cdot}{l \mapsto n} \dots \rangle_{\text{ptr}} \langle \frac{\sigma}{\sigma[0/l..l +_{Int} n -_{Int} 1]} \rangle_{\text{store}} \langle \frac{l}{l +_{Int} n} \rangle_{\text{nextLoc}}}{l}$$

$$\frac{\langle \text{free}(l) \dots \rangle_k \langle \dots \frac{l}{l \mapsto n} \dots \rangle_{\text{ptr}} \langle \frac{\sigma}{\sigma[1/l..l +_{Int} n -_{Int} 1]} \rangle_{\text{store}}}{\cdot}$$

**Exercise 176.** *Evaluating an expression and returning a reference to it:*

$$\text{ref} \equiv \lambda x. (\lambda p. *p := x; p)(\text{malloc}(1))$$

$$\frac{\langle \text{ref } v \dots \rangle_k \langle \frac{\sigma}{\sigma[v/l]} \rangle_{\text{store}} \langle \frac{l}{l +_{Int} 1} \rangle_{\text{nextLoc}}}{l}$$

**Assignment** The right hand side of an assignment is evaluated to a value, specified by ‘strict(2)’ in the operator declaration, while the left-hand-side is evaluated to a l-value, here either a Name, or the dereferencing of a location (the ‘strict(K)’ in the K-context declaration).

$$\frac{\langle x := v \dots \rangle_k \langle \dots x \mapsto l \dots \rangle_{\text{env}} \langle \dots l \mapsto \frac{\_}{v} \dots \rangle_{\text{store}}}{\cdot}$$

context:  $*e := e' [strict(e)]$

$$\frac{\langle *l := v \dots \rangle_k \langle \dots l \mapsto \frac{\_}{v} \dots \rangle_{\text{store}}}{\cdot}$$

**Function and aspects** Since our language allows usage of pointer and references, it is more convenient to work with environments and to represent functions as closures. In addition, the Challenge language allows the specification of aspects, whose semantics is that they affect each function defined after the ‘aspect’ declaration statement to execute the statement provided as its argument prior to execution their body, in the environment of the function. This is achieved by using a special cell to hold the current aspect, and by including the computation contained in that cell in the closure of each function being evaluated.

$$\frac{\langle \text{aspect } s \dots \rangle_k \langle \_ \rangle_{\text{aspect}}}{s}$$

$$\frac{\langle \lambda xl. e \dots \rangle_k \langle \rho \rangle_{\text{env}} \langle s \rangle_{\text{aspect}}}{\text{closure}(xl, s \rightsquigarrow e, \rho)}$$

**Exercise 177.** Define a different semantics for aspects, or a new aspect construct, that would weave the aspect at function call time instead of at function declaration time.

**Function application** Upon function call, the evaluated arguments are bound to the formal parameters, then the body of the function is executed in the environment saved by the closure, and finally the calling environment must be restored.

$$\frac{\langle \text{closure}(xl, e, \rho) vl \dots \rangle_k \langle \frac{\varrho}{\rho[n \dots n +_{Int} |xl| -_{Int} 1/xl]} \rangle_{\text{env}} \langle \frac{\sigma}{\sigma[vl/n \dots n +_{Int} |xl| -_{Int} 1]} \rangle_{\text{store}} \langle \frac{n}{n +_{Int} |xl|} \rangle_{\text{nextLoc}}}{e \rightsquigarrow \text{env}(\varrho)}$$

Restoring the original environment, when the body of the function is completely evaluated:

$$\frac{\langle v \rightsquigarrow \text{env}(\rho) \dots \rangle_k \langle \_ \rangle_{\text{env}}}{\rho}$$

Tail call optimization:

$$\text{env}(\_ ) \rightsquigarrow \text{env}(\_ )$$

**Exercise 178.** Prove that the tail call optimization is sound.

**Recursion** The semantics of  $\mu X.E$  is given by evaluating the expression  $E$  in the environment/store where  $X$  is bound to a  $\mu X.E$ .

$$\frac{\langle \frac{\mu x . e}{e \rightsquigarrow \text{env}(\rho)} \dots \rangle_k \langle \frac{\rho}{\rho[n/x]} \rangle_{\text{env}} \langle \dots \frac{\cdot}{n \mapsto \mu x . e} \dots \rangle_{\text{store}} \langle \frac{n}{n +_{\text{Int}} 1} \rangle_{\text{nextLoc}}}$$

**Call with current continuation (call/cc)** The semantics of call/cc is that of packaging the remainder of the computation as a value and passing it to the function passed as argument. When the thus packaged computation is applied on a value, the entire computation at the time is replaced by the value being put on top of the packed computation. Note that the environment needs to be packed together with the computation, same as for function closures.

$$\frac{\langle \text{callcc } v \rightsquigarrow k \rangle_k \langle \rho \rangle_{\text{env}}}{v \text{ cc}(k, \rho)}$$

$$\frac{\langle \text{cc}(k, \rho) v \dots \rangle_k \langle \_ \rangle_{\text{env}}}{v \rightsquigarrow k \quad \rho}$$

**Sequential non-determinism** ‘randomBool’ non-deterministically evaluates to either ‘true’ or ‘false’.

$$\frac{\langle \text{randomBool } \dots \rangle_k}{\text{true}}$$

$$\frac{\langle \text{randomBool } \dots \rangle_k}{\text{false}}$$

**Threads** Threads are characterized by independent control flow, but shared memory. In our definition, a ‘thread’ cell is used to group together all (computation, environment, acquired resources) cells related to a thread. The ‘spawn’ command creates a new thread which is initialized with the statement passed as argument as its computation, and the environment of the thread creating it.

$$\frac{\langle \dots \langle \text{spawn } s \dots \rangle_k \langle \rho \rangle_{\text{env}} \dots \rangle_{\text{thread}} \cdot}{\langle \langle s \rangle_k \langle \rho \rangle_{\text{env}} \langle \cdot \rangle_{\text{holds}} \rangle_{\text{thread}}}$$

When the computation of a thread has completed, it can be dissolved and its held resources be released. ‘busy’ is a shared cell holding the names of the resources acquired by any of the threads.

$$\frac{\langle \dots \langle \cdot \rangle_k \langle h \rangle_{\text{holds}} \dots \rangle_{\text{thread}} \cdot}{\langle \frac{\text{busy}}{\text{busy} - \text{set } \text{dom}_{\text{map}}(h)} \rangle_{\text{busy}}}$$

**Thread synchronization** To attain mutual exclusion, threads can be synchronized by means of locks. In this language, one can lock on any value. A lock can only be acquired if it is not busy. Same thread can acquire same lock multiple times, therefore multiplicities for each lock are maintained in the ‘holds’ cell.

$$\begin{array}{c}
\langle \underline{\text{acquire } v \ \dots} \rangle_k \langle \dots v \mapsto \frac{n}{s_{Nat}(n)} \dots \rangle_{\text{holds}} \\
\cdot \\
\langle \underline{\text{acquire } v \ \dots} \rangle_k \langle \dots \frac{\cdot}{v \mapsto 0} \dots \rangle_{\text{holds}} \langle \underline{\text{busy } \cdot} \rangle_{\text{busy}} \quad \text{when } \neg_{Bool}(v \in_{set} \text{busy}) \\
\cdot \\
\langle \underline{\text{release } v \ \dots} \rangle_k \langle \dots v \mapsto \frac{s_{Nat}(n)}{n} \dots \rangle_{\text{holds}} \\
\cdot \\
\langle \underline{\text{release } v \ \dots} \rangle_k \langle \dots v \mapsto 0 \dots \rangle_{\text{holds}} \langle \dots v \ \dots \rangle_{\text{busy}} \\
\cdot
\end{array}$$

**Rendez-vous synchronization** Two threads can only pass together a barrier specified by a lock V. This is harder to achieve in all definitional frameworks we know of, because it requires the ability to access and reduce redexes of two computational units simultaneously.

$$\langle \underline{\text{rv } v \ \dots} \rangle_k \langle \underline{\text{rv } v \ \dots} \rangle_k$$

**Agents** An agent is here a collection of threads, grouped in an ‘agent’ cell identified by an id held by the ‘me’ cell, and holding in the ‘parent’ cell a reference id to its creating agent. ‘nextAgent’ is used for providing fresh ids for agents.

$$\langle \underline{\text{new-agent } s \ \dots} \rangle_k \langle m \rangle_{\text{me}} \langle \frac{n}{n +_{Int} 1} \rangle_{\text{nextAgent}} \frac{\cdot}{NewAgent}$$

where *NewAgent* stands for:

$$\langle \langle \langle s \rangle_k \langle \cdot \rangle_{\text{env}} \langle \cdot \rangle_{\text{holds}} \rangle_{\text{thread}} \langle \cdot \rangle_{\text{busy}} \langle n \rangle_{\text{me}} \langle m \rangle_{\text{parent}} \langle \cdot \rangle_{\text{store}} \langle 0 \rangle_{\text{nextLoc}} \langle \cdot \rangle_{\text{aspect}} \rangle_{\text{agent}}$$

When all threads inside an agent have completed, the agent can be dissolved.

$$\langle \dots \langle \cdot \rangle_{\text{threads}} \dots \rangle_{\text{agent}} \rightarrow \cdot$$

An agent can send any value (including agents ids) to other agents (provided it knows their id). To model asynchronous communication, each value sent is wrapped in a ‘message’ cell identifying both the sender and the intended receiver.

$$\begin{array}{c}
\langle \underline{\text{me } \dots} \rangle_k \langle m \rangle_{\text{me}} \\
\frac{\cdot}{m} \\
\langle \underline{\text{parent } \dots} \rangle_k \langle n \rangle_{\text{parent}} \\
\frac{\cdot}{n} \\
\langle \underline{\text{send-asynch } n \ v \ \dots} \rangle_k \langle m \rangle_{\text{me}} \frac{\cdot}{\langle \langle m \rangle_{\text{sender}} \langle n \rangle_{\text{receiver}} \langle v \rangle_{\text{val}} \rangle_{\text{message}}}
\end{array}$$

An agent can request to receive a message from a certain agent, or from any agent.

$$\frac{\langle \text{receive-from } n \dots \rangle_k \langle m \rangle_{\text{me}} \langle \langle n \rangle_{\text{sender}} \langle m \rangle_{\text{receiver}} \langle v \rangle_{\text{val}} \rangle_{\text{message}}}{v}$$

$$\frac{\langle \text{receive } \dots \rangle_k \langle m \rangle_{\text{me}} \langle \dots \langle m \rangle_{\text{receiver}} \langle v \rangle_{\text{val}} \dots \rangle_{\text{message}}}{v}$$

The message can be sent synchronously, in which case, two agents need to be matched together for the exchange to occur.

$$\frac{\langle \dots \langle \text{send-synch } n \ v \dots \rangle_k \langle m \rangle_{\text{me}} \dots \rangle_{\text{agent}} \langle \dots \langle \text{receive-from } m \ \langle n \rangle_{\text{me}} \dots \rangle_k \dots \rangle_{\text{agent}}}{v}$$

$$\frac{\langle \text{send-synch } n \ v \dots \rangle_k \langle \dots \langle \text{receive } \dots \rangle_k \langle n \rangle_{\text{me}} \dots \rangle_{\text{agent}}}{v}$$

**Exercise 179.** While the rules for asynchronous sending allow the sender and the receiver to be same agent, the two rules for synchronous sending assume that the sender and the receiver agents are different. Add two other rules for synchronous sending that extend the semantics above to allow agents to communicate synchronously with themselves (they would have to send from one thread and receive from another thread).

**Abrupt termination** Execution of ‘halt’ in one of the threads of an agent dissolves that agent.

$$\langle \text{halt-thread } \dots \rangle_k \rightarrow \langle \cdot \rangle_k$$

$$\langle \dots \langle \text{halt-agent } \dots \rangle_k \dots \rangle_{\text{threads}} \rightarrow \langle \cdot \rangle_{\text{threads}}$$

$$\langle \dots \langle \text{halt-system } \dots \rangle_k \dots \rangle_{\text{agents}} \rightarrow \langle \cdot \rangle_{\text{agents}}$$

**Code Generation** The equations below implement the quote/unquote mechanism of runtime code generation. First, the syntax of computations is extended to include “boxed” versions of the K arrow and K list constructors, as well as the helping operator quote, which tracks the level of nesting of quoted expressions. Similarly, boxed versions are introduced for all K labels (and thus for all language constructs), and a new Label code is added to hold values of type code

$$\begin{aligned} KLabel & ::= \boxed{\curvearrowright} \text{ [strict]} \mid \boxed{KLabel} \text{ [strict]} \mid \boxed{\_} \text{ [strict]} \mid \text{quote}[Nat] \\ KResultLabel & ::= \text{code} \end{aligned}$$

Initially, a quoted expression starts at the quoting level 0, and acts as a morphism on all K constructors, transforming them into their boxed version, except for the quote/unquote labels. Each of these boxed items is supposed to evaluate to a code fragment; once this happens, the fragments are glued together to a bigger code fragment.

$\langle \text{quote}(k) \dots \rangle_k$   
 $\text{quote}[0]$

$\text{quote}[n](\text{quote}(k)) \rightarrow \boxed{\text{quote}}(\text{quote}[s_{Nat}(n)](k))$   
 $\text{quote}[0](\text{unquote}(k)) \rightarrow k$   
 $\text{quote}[s_{Nat}(n)](\text{unquote}(k)) \rightarrow \boxed{\text{unquote}}(\text{quote}[n](k))$

$\text{quote}[n](\cdot) \rightarrow \text{code}(\cdot)$   
 $\text{quote}[n](k_1 \rightsquigarrow k_2) \rightarrow \text{quote}[n](k_1) \boxed{\rightsquigarrow} \text{quote}[n](k_2) \quad \text{when } k_1 \neq \cdot, k_2 \neq \cdot$   
 $\text{code}(k_1) \boxed{\rightsquigarrow} \text{code}(k_2) \rightarrow \text{code}(k_1 \rightsquigarrow k_2)$

$\text{quote}[n](\text{label}(kl)) \rightarrow \boxed{\text{label}}(\text{quote}[n](kl)) \quad \text{when } \text{label} \neq \text{quote}, \text{label} \neq \text{unquote}$   
 $\boxed{\text{label}}(\text{code}(kl)) \rightarrow \text{code}(\text{label}(kl))$

$\text{quote}[n](\_) \rightarrow \text{code}(\_)$   
 $\text{quote}[n](kl_1, kl_2) \rightarrow \text{quote}[n](kl_1) \boxed{,} \text{quote}[n](kl_2) \quad \text{when } kl_1 \neq \_, kl_2 \neq \_$   
 $\text{code}(kl_1) \boxed{,} \text{code}(kl_2) \rightarrow \text{code}(kl_1, kl_2)$

$\text{eval}(\text{code}(k)) \rightarrow k$