

Compiling Cooperative Task Management to Continuations

Keiko Nakata, Andri Saar

► **To cite this version:**

Keiko Nakata, Andri Saar. Compiling Cooperative Task Management to Continuations. Farhad Arbab; Marjan Sirjani. 5th International Conference on Fundamentals of Software Engineering (FSEN), Apr 2013, Tehran, Iran. Springer Berlin Heidelberg, Lecture Notes in Computer Science, LNCS-8161, pp.95-110, 2013, Fundamentals of Software Engineering. <10.1007/978-3-642-40213-5_7>. <hal-01514661>

HAL Id: hal-01514661

<https://hal.inria.fr/hal-01514661>

Submitted on 26 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Compiling cooperative task management to continuations

Keiko Nakata and Andri Saar

Institute of Cybernetics at Tallinn University of Technology
Akadeemia tee 21, 12618 Tallinn

Abstract. Although preemptive concurrency models are dominant for multi-threaded concurrency, they may be criticized for the complexity of reasoning because of the implicit context switches. The actor model and cooperative concurrency models have regained attention as they encapsulate the thread of control. In this paper, we formalize a continuation-based compilation of cooperative multitasking for a simple language and prove its correctness.

1 Introduction

In a preemptive concurrency model, threads may be suspended and activated at any time. While preemptive models are dominant for multi-threaded concurrency, they may be criticized for the complexity of reasoning because of the implicit context switches. The programmer often has to resort to low-level synchronization primitives, such as locks, to prevent unwanted context switches. Programs written in such a way tend to be error-prone and are not scalable. The actor model [2] addresses this issue. Actors encapsulate the thread of control and communicate with each other by sending messages. They are also able to call blocking operations such as sleep, await and receive, reminiscent of cooperative multi-tasking. Erlang and Scala actors support actor-based concurrency models.

Creol [8] and ABS [7] combine a message-passing concurrency model and a cooperative concurrency model. In Creol, each object encapsulates a thread of control and objects communicate with each other using asynchronous method calls. Asynchronous method calls, instead of messages-passing, provide a type-safe communication mechanism and are a good match for object-oriented languages [4,3]. ABS generalizes the concurrency model of Creol by introducing *concurrent object groups* [10] as the unit of concurrency. The concurrency model of ABS can be split in two: in one layer, we have local, synchronous and shared-memory communication¹ in one concurrent object group (COG) and on the second layer we have asynchronous message-based concurrency between different concurrent object groups as in Creol. The behavior of one COG is based on the cooperative multitasking of external method invocations and internal

¹ In ABS, different tasks originating from the same object may communicate with each other via fields of the object.

method activations, with concrete scheduling points where a different task may get scheduled. Between different COGs only asynchronous method calls may be used; different COGs have no shared object heap. The order of execution of asynchronous method calls is not specified. The result of an asynchronous call is a future; callers may decide at run-time when to synchronize with the reply from a call. Asynchronous calls may be seen as triggers that spawn new method activations (or tasks) within objects. Every object has a set of tasks that are to be executed (originating from method calls). Among these, at most one task of all the objects belonging to one COG is active; others are suspended and awaiting execution. The concurrency models of Creol and ABS are designed to be suitable in the distributed setting, where one COG executes on its own (virtual) processor in a single node and different COGs may be executed on different nodes in the network.

In this paper, we are interested in the compilation of cooperative multi-tasking into continuations, motivated to execute a cooperative multi-tasking model on the JVM platform, which employs a preemptive model. The basic idea of using continuations to manage the control behavior of the computation has been known from 80's [12,6], and is still considered as a viable technique [9,11,1]. This is particularly so, if the programming language supports first-class continuations, as in the case of Scala, and hence one can obviate manual stack management. The contribution of the paper is a correctness proof of such a compilation scheme. Namely, we create a simplified source language, by extending the While language with (synchronous) procedure calls and operations for cooperative multi-tasking (i.e., blocking operations and creation of new tasks) and define a compilation function from the source language into the target language, which extends While with continuation operations—the target language is sequential. We then prove that the compilation preserves the operational behavior from the source language to the target language.

The remainder of the paper is organized as follows. We define the source language and its operational semantics in the next section, and the target language and its operational semantics in Section 3. In Section 4, we present the compilation function from the source language to the target language and, in Section 5 we prove its correctness. We conclude in Section 6.

2 Source language

$$\begin{aligned}
 S ::= & \overline{x := e} \mid \overline{u := e} \mid \overline{\text{skip}} \mid \overline{S_1; S_2} \mid \overline{\text{if } e \text{ then } S_1 \text{ else } S_2} \mid \overline{\text{while } e \text{ do } S} \\
 & \mid \overline{\text{var } x = e} \mid \overline{f(\bar{e})} \mid \overline{\text{await } e} \mid \overline{\text{spawn } f \bar{e}} \mid \overline{\text{return}}
 \end{aligned}$$

Fig. 1. Syntax of the source language

In Figure 1, we define the syntax for the source language. We use the overline notation to denote sequences, with ϵ denoting an empty sequence. It is the

While language extended with (local) variable definitions, **var** $x = e$, procedure calls $f(\bar{e})$, the await statement **await** e , creation of a new task **spawn** $f \bar{e}$, and the return statement **return**. For simplicity, we syntactically distinguish local variable assignment $x := e$ and global variable assignment $u := e$. The statement **var** $x = e$ defines a (task) local variable x and initializes it with the value of e . The statement $f(\bar{e})$ invokes the procedure f with arguments \bar{e} , to be executed within the same task. The procedure does not return the result to the caller, but may store the result in a global variable. The statement **await** e suspends the execution of the current task, which can be resumed when the guard expression e evaluates to true. The statement **spawn** $f \bar{e}$ spawns a new task, which executes the body of f with arguments \bar{e} . (Hence, in contrast to procedure calls, which are synchronous, **spawn** $f \bar{e}$ is like an asynchronous procedure call executed in a new task.) The **return** statement is a runtime construct, not appearing in the source program, and will be explained later.

We assume disjoint supplies of local variables (ranged over by x), global variables (ranged over by u), and procedure names (ranged over by f). We assume a set of (pure) expressions, whose elements are ranged over by e . We assume the set of values to be the integers, non-zero integers counting as truth and zero as falsity. The metavariable v ranges over values. We have two kinds of states – local states and global states. A local state, ranged over by ρ , maps local variables to values; a global state, ranged over by σ , maps global variables to values. We denote by \emptyset an empty mapping, whose domain is empty. Communication between different tasks is achieved via global variables. For simplicity, we assume a fixed set of global variables. The notation $\rho[x \mapsto v]$ denotes the update of ρ with v at x , when x is in the domain of ρ . If x is not in the domain, it denotes a mapping extension. The notation $\sigma[u \mapsto v]$ denotes similar. We assume given an evaluation function $\llbracket e \rrbracket_{(\rho, \sigma)}$, which evaluates e in the local state ρ and the global state σ . We write $(\rho, \sigma) \models e$ and $(\rho, \sigma) \not\models e$ to denote that e is true, resp. false with respect to ρ and σ . A *stack* Π is a non-empty list of local states, whose elements are separated by semicolons. A stack grows leftward, i.e., the leftmost element is the topmost element.

A program P consists of a procedure environment env_F which maps procedure names to pairs of a formal argument list and a statement, and a global state which maps global variables to their initial values. The entry point of the program will be the procedure named **main**.

We define the operational semantics of the source language as a transition system on *configurations*, in the style of structural operational semantics. A configuration cfg consists of an active task identifier n , a global variable mapping σ and a set of tasks Θ . A task has an identifier and may be in one of the three forms: a triple $\langle e, S, \Pi \rangle$, representing a task that is awaiting to be scheduled, where e is the guard expression, S the statement and Π its stack; or, a pair $\langle S, \Pi \rangle$, representing the currently active task; or, a singleton $\langle \Pi \rangle$, representing a terminated task.

$$\begin{aligned} \text{Configuration } \text{cfg} &::= n, \sigma \triangleright \Theta \\ \text{Task sets } \Theta &::= n \langle e, S, \Pi \rangle \mid n \langle S, \Pi \rangle \mid n \langle \Pi \rangle \mid \Theta \parallel \Theta \end{aligned}$$

The order of tasks in the task set is irrelevant: the parallel operator \parallel is commutative and associative. Formally, we assume the following structural equivalence:

$$\theta \equiv \theta \quad \theta \parallel \theta' \equiv \theta' \parallel \theta \quad \theta \parallel (\theta' \parallel \theta'') \equiv (\theta \parallel \theta') \parallel \theta''$$

Transition rules in the semantics are in the form $\text{env}_F \vdash \text{cfg} \rightarrow \text{cfg}'$, shown in Figure 2. The first two rules (S-CONG and S-EQUIV) deal with congruence and structural equivalence. The rules for assignment, **skip**, if-then-else and while are self-explanatory. For instance, in the rule S-ASSIGN-LOCAL, the task is of the form $n\langle S, \Pi' \rangle$ where $S = x := e$ and $\Pi' = \rho; \Pi$. Note that the topmost element of the stack Π is the current local state. The rules for sequential composition may deserve some explanation. If the first statement S_1 suspends guarded by e in the stack Π' with the residual statement S'_1 to be run when resumed, then the entire statement $S_1; S_2$ suspends in $\langle e, S'_1; S_2, \Pi' \rangle$, where the residual statement now contains the second statement S_2 (S-SEQ-GRD). If S_1 terminates in Π' , then S_2 will run next in Π' (S-SEQ-FIN). Otherwise, S_1 transfers to S'_1 with the stack Π' , so that $S_1; S_2$ transfers to $S'_1; S_2$ with the same stack (S-SEQ-STEP). The **await** statement immediately suspends (S-AWAIT) the currently active task, enabling us to switch to some other task in accordance to the scheduling rules. An example of the **await** statement (and the scheduling rules) at work can be found in the example in Figure 3. The statement **spawn** $f \bar{e}$ creates a new task $n'(\text{true}, S, [\bar{x} \mapsto \bar{v}])$ with n' a fresh identifier (S-SPAWN). The caller task continues to be active. The newly created task is suspended, guarded by **true**, and may get scheduled at scheduling points by the scheduling rules (see below). Procedure invocation $f(\bar{e})$ evaluates the arguments \bar{e} in the current state, pushes into the stack the local state $[\bar{x} \mapsto \bar{v}]$, mapping the formal parameters to the actual arguments, and transfers to $S; \text{return}$, where S is the body of f (S-CALL). The **return** statement pops the topmost element from the stack (S-RETURN). The local variable definition **var** $x = e$ extends the current local state with the newly defined variable and initializes it with the value of e (S-VAR).

The last three rules deal with scheduling. If the current active task has terminated, then a new task whose guard evaluates to true is chosen to be active (S-SCHED-FIN). When the active task suspends, a scheduling point is reached. The rule (S-SCHED-SAME) considers the case in which the same task is scheduled; the rule (S-SCHED-OTHER) considers the case in which a different task is scheduled.

As an example, we will look at a program containing one global variable u with the initial value 0 and the following procedures:

$$\begin{aligned} & f \mapsto u := 1 \\ & \text{main} \mapsto u := 3; \text{spawn } f \ \epsilon; \text{await } u = 1; u := 2 \end{aligned}$$

A detailed step, showing the full derivation, can be seen in Figure 4. A full execution trace, showing all intermediate configurations, is shown in Figure 3.

$$\begin{array}{c}
\frac{\text{env}_F \vdash n, \sigma \triangleright \Theta' \rightarrow n', \sigma' \triangleright \Theta''}{\text{env}_F \vdash n, \sigma \triangleright \Theta \parallel \Theta' \rightarrow n', \sigma' \triangleright \Theta \parallel \Theta''} \text{S-CONG} \\
\\
\frac{\Theta \equiv \Theta' \quad \text{env}_F \vdash n, \sigma \triangleright \Theta' \rightarrow n', \sigma' \triangleright \Theta''' \quad \Theta'' \equiv \Theta'''}{\text{env}_F \vdash n, \sigma \triangleright \Theta \rightarrow n', \sigma' \triangleright \Theta''} \text{S-EQUIV} \\
\\
\frac{x \in \text{dom } \rho}{\text{env}_F \vdash n, \sigma \triangleright n \langle x := e, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle \rho[x \mapsto \llbracket e \rrbracket_{(\rho, \sigma)}]; \Pi \rangle} \text{S-ASSIGN-LOCAL} \\
\\
\frac{u \in \text{dom } \sigma}{\text{env}_F \vdash n, \sigma \triangleright n \langle u := e, \Pi \rangle \rightarrow n, \sigma[u \mapsto \llbracket e \rrbracket_{(\rho, \sigma)}] \triangleright n \langle \Pi \rangle} \text{S-ASSIGN-GLOBAL} \\
\\
\frac{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1, \Pi \rangle \rightarrow n', \sigma' \triangleright n \langle e, S'_1, \Pi' \rangle \parallel \Theta}{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1; S_2, \Pi \rangle \rightarrow n', \sigma' \triangleright n \langle e, S'_1; S_2, \Pi' \rangle \parallel \Theta} \text{S-SEQ-GRD} \\
\\
\frac{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1, \Pi \rangle \rightarrow n', \sigma' \triangleright n \langle \Pi' \rangle \parallel \Theta}{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1; S_2, \Pi \rangle \rightarrow n', \sigma' \triangleright n \langle S_2, \Pi' \rangle \parallel \Theta} \text{S-SEQ-FIN} \\
\\
\frac{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1, \Pi \rangle \rightarrow n', \sigma' \triangleright n \langle S'_1, \Pi' \rangle \parallel \Theta}{\text{env}_F \vdash n, \sigma \triangleright n \langle S_1; S_2, \Pi \rangle \rightarrow n', \sigma' \triangleright n \langle S'_1; S_2, \Pi' \rangle \parallel \Theta} \text{S-SEQ-STEP} \\
\\
\frac{(\rho, \sigma) \models e}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{if } e \text{ then } S_1 \text{ else } S_2, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle S_1, \rho; \Pi \rangle} \text{S-IF-TRUE} \\
\\
\frac{(\rho, \sigma) \not\models e}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{if } e \text{ then } S_1 \text{ else } S_2, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle S_2, \rho; \Pi \rangle} \text{S-IF-FALSE} \\
\\
\frac{(\rho, \sigma) \models e}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{while } e \text{ do } S, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle S; \text{while } e \text{ do } S, \rho; \Pi \rangle} \text{S-WHILE-TRUE} \\
\\
\frac{(\rho, \sigma) \not\models e}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{while } e \text{ do } S, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle \rho; \Pi \rangle} \text{S-WHILE-FALSE} \quad \frac{}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{skip}, \Pi \rangle \rightarrow n, \sigma \triangleright n \langle \Pi \rangle} \text{S-SKIP} \\
\\
\frac{\text{env}_F f = (\bar{x}, S) \quad \bar{v} = \llbracket \bar{e} \rrbracket_{(\rho, \sigma)} \quad n' \text{ is fresh}}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{spawn } f \bar{e}, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle \rho; \Pi \rangle \parallel n' \langle \text{true}, S, [\bar{x} \mapsto \bar{v}] \rangle} \text{S-SPAWN} \quad \frac{}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{await } e, \Pi \rangle \rightarrow n, \sigma \triangleright n \langle e, \text{skip}, \Pi \rangle} \text{S-AWAIT} \\
\\
\frac{x \notin \text{dom } \rho}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{var } x = e, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle \rho[x \mapsto \llbracket e \rrbracket_{(\rho, \sigma)}]; \Pi \rangle} \text{S-VAR} \quad \frac{\text{env}_F f = (\bar{x}, S) \quad \bar{v} = \llbracket \bar{e} \rrbracket_{(\rho, \sigma)}}{\text{env}_F \vdash n, \sigma \triangleright n \langle f(\bar{e}), \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle S; \text{return}, [\bar{x} \mapsto \bar{v}]; \rho; \Pi \rangle} \text{S-CALL} \\
\\
\frac{}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{return}, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle \Pi \rangle} \text{S-RETURN} \quad \frac{(\rho, \sigma) \models e}{\text{env}_F \vdash n, \sigma \triangleright n \langle \Pi \rangle \parallel n' \langle e, S, \rho; \Pi' \rangle \rightarrow n', \sigma \triangleright n' \langle S, \rho; \Pi' \rangle \parallel n \langle \Pi \rangle} \text{S-SCHED-FIN} \\
\\
\frac{(\rho, \sigma) \models e}{\text{env}_F \vdash n, \sigma \triangleright n \langle e, S, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle S, \rho; \Pi \rangle} \text{S-SCHED-SAME} \quad \frac{(\rho', \sigma) \models e'}{\text{env}_F \vdash n, \sigma \triangleright n \langle e, S, \rho \rangle \parallel n' \langle e', S', \rho' \rangle \rightarrow n', \sigma \triangleright n' \langle S', \rho' \rangle \parallel n \langle e, S, \rho \rangle} \text{S-SCHED-OTHER}
\end{array}$$

Fig. 2. Semantics of the source language

```

env_F ⊢ main, [u ↦ 0] ▷ main⟨u := 3; spawn f ε; await u = 1; u := 2, ∅⟩
→ main, [u ↦ 3] ▷ main⟨spawn f ε; await u = 1; u := 2, ∅⟩
→ main, [u ↦ 3] ▷ main⟨await u = 1; u := 2, ∅⟩ ∥ f⟨true, u := 1, ∅⟩
→ main, [u ↦ 3] ▷ main⟨u = 1, skip; u := 2, ∅⟩ ∥ f⟨true, u := 1, ∅⟩
→ f, [u ↦ 3] ▷ main⟨u = 1, skip; u := 2, ∅⟩ ∥ f⟨u := 1, ∅⟩
→ f, [u ↦ 1] ▷ main⟨u = 1, skip; u := 2, ∅⟩ ∥ f⟨∅⟩
→ main, [u ↦ 1] ▷ main⟨skip; u := 2, ∅⟩ ∥ f⟨∅⟩
→ main, [u ↦ 1] ▷ main⟨u := 2, ∅⟩ ∥ f⟨∅⟩
→ main, [u ↦ 2] ▷ main⟨∅⟩ ∥ f⟨∅⟩

```

Fig. 3. The full execution trace of the example program

$$\frac{\frac{\frac{\text{env}_F \vdash \text{main}, [u \mapsto 3] \triangleright \text{main}\langle \text{await } u = 1, \emptyset \rangle}{\rightarrow \text{main}, [u \mapsto 3] \triangleright \text{main}\langle u = 1, \text{skip}, \emptyset \rangle} \text{S-AWAIT}}{\text{env}_F \vdash \text{main}, [u \mapsto 3] \triangleright \text{main}\langle \text{await } u = 1; u := 2, \emptyset \rangle} \text{S-SEQ-GRD}}{\text{env}_F \vdash \text{main}, [u \mapsto 3] \triangleright \text{main}\langle \text{await } u = 1; u := 2, \emptyset \rangle \parallel \text{f}\langle \text{true}, u := 1, \emptyset \rangle} \text{S-CONG}}{\rightarrow \text{main}, [u \mapsto 3] \triangleright \text{main}\langle u = 1, \text{skip}; u := 2, \emptyset \rangle \parallel \text{f}\langle \text{true}, u := 1, \emptyset \rangle}$$

Fig. 4. Example of a derivation in the source language

3 Target language

We proceed to the target language. In Figure 5, we present the syntax of the target language. Expressions of the target language contain, besides pure expressions, *continuations* $[S, \Pi]$, which are pairs of a statement S and a stack Π , and support for *guarded (multi)sets*: collections which contain pairs of an expression and a value. The expression stored with each element is called a *guard expression*, and is evaluated when we query the set: only elements whose guard expressions hold may be returned. There are five expressions in the language to work with guarded sets: an empty set \emptyset , checking whether a set is empty (`isEmpty e_s`), adding an element (`add $e_s e_g e$`), fetching an element (`get e_s`) and removing an element (`del $e_s e$`).

Similar to the source language, for the target language we extend While with local variable definitions and procedure calls. We also add delimited control operators, `shift $k \{S\}$` , `reset $\{S\}$` , `invoke k` [5]. The statement `shift $k \{S\}$` captures the rest of the computation, or continuation, up to the closest surrounding `reset $\{\}$` , binds it to k , and proceeds to execute S . In `shift $k \{S\}$` , k is a binding occurrence, whose scope is S . Hence, the statement `reset $\{S\}$` delimits the captured continuation. The statement `invoke k` invokes, or jumps into, the continuation bound to the variable k . The statement `R $\{S\}$` , where R is a new constant, is a runtime construct to be explained later.

The target language is sequential and, unlike the source language, it contains no explicit support for parallelism. Instead, we provide building blocks – contin-

Expressions $e ::= \dots$
 $\quad | [S, II] | \emptyset | \text{isEmpty } e | \text{add } e_1 e_2 e_3 | \text{del } e_1 e_2 | \text{get } e$
Statements $S ::= x := e | u := e | \text{skip} | S_1; S_2 | \text{if } e \text{ then } S_1 \text{ else } S_2 | \text{while } e \text{ do } S$
 $\quad | \text{var } x = e | f(\bar{e}) | \text{reset } \{S\} | \text{shift } k \{S\} | \text{invoke } k$
 $\quad | \text{return} | \mathbf{R} \{S\}$

Fig. 5. Syntax of the target language

uations and guarded sets – that are used to switch between tasks and implement an explicit scheduler in Section 4.

We assume given an evaluation function $\llbracket e \rrbracket_{(II, \sigma)}$ for pure expressions, which evaluates e with respect to the stack II (the evaluation only looks at the current local state, which is the topmost element of II) and the global state σ . In Figure 6, the evaluation function is extended to operations for guarded sets.

We define an operational semantics for the target language as a reduction semantics over configurations, using evaluation contexts. A configuration $\langle S, II, \sigma \rangle$ is a triple of a statement S , a stack II and a global state σ .

Evaluation contexts E are statements with a hole, specifying where the next reduction may occur. They are defined by

$$E ::= [] | E; S | \mathbf{R} \{E\}$$

We denote by $E[S]$ the statement obtained by placing S in the hole of E .

We define basic reduction rules in Figure 6. **reset** $\{S\}$ inserts a marker \dagger into the stack, just below the current state, and reduces to $\mathbf{R} \{S\}$ to continue the execution of S (T-RESET). We use a marker to delimit the portion of the stack captured by **shift** and to align the stack when exiting from $\mathbf{R} \{ \}$. The runtime construct $\mathbf{R} \{ \}$ is used to record that the marker has been set. **shift** $k \{S\}$ captures the rest of the execution up to and including the closest surrounding $\mathbf{R} \{ \}$ together with the corresponding portion of the stack, binds it to a fresh variable k' in the local state, and continues with the statement S' obtained by substituting k by k' in S (T-SHIFT). The surrounding $\mathbf{R} \{ \}$ is kept intact. F is an evaluation context that does not intersect $\mathbf{R} \{ \}$, formally,

$$F ::= [] | F; S$$

Note that **shift** $k \{S\}$ captures the stack up to and including the topmost \dagger , which has been inserted by the closest surrounding **reset**. Once the body of $\mathbf{R} \{ \}$ terminates, i.e., reduces to **skip**, then we remove the $\mathbf{R} \{ \}$ and pop the stack until the topmost \dagger , but leaving the state just above \dagger in the stack (T-R). **invoke** k invokes the continuation bound to k (T-VOKE). Namely, if k is bound to $[S, II']$ in the local or global state, then the statement reduces to S ; **return** and the stack II' is pushed into the current stack. S must be necessarily of the form $\mathbf{R} \{S'\}$, where S' does not contain \mathbf{R} , and II' contains exactly one \dagger at the bottom. When exiting from the $\mathbf{R} \{ \}$, the state immediately above \dagger in II' will

$$\begin{aligned}
\llbracket \text{add } e_s \ e_g \ e \rrbracket_{(II, \sigma)} &= \llbracket e_s \rrbracket_{(II, \sigma)} \cup (e_g, \llbracket e \rrbracket_{(II, \sigma)}) \\
\llbracket \text{get } e_s \rrbracket_{(II, \sigma)} &= v \text{ when } \exists e_g \exists v. (e_g, v) \in \llbracket e_s \rrbracket_{(II, \sigma)} \wedge \llbracket e_g \rrbracket_{(II, \sigma)} = \text{true} \\
\llbracket \text{del } e_s \ e \rrbracket_{(II, \sigma)} &= \llbracket e_s \rrbracket_{(II, \sigma)} \setminus (e_g, \llbracket e \rrbracket_{(II, \sigma)}) \text{ when } \exists e_g. (e_g, \llbracket e \rrbracket_{(II, \sigma)}) \in \llbracket e_s \rrbracket_{(II, \sigma)} \\
\llbracket \text{isEmpty } e_s \rrbracket_{(II, \sigma)} &= \begin{cases} \text{true} & \text{if } \llbracket e_s \rrbracket_{(II, \sigma)} = \emptyset \\ \text{false} & \text{otherwise} \end{cases}
\end{aligned}$$

$$\frac{}{\text{env}_F \vdash \langle \text{reset } \{S\}, \rho; II, \sigma \rangle \rightarrow \langle \mathbb{R} \{S\}, \rho \dagger; II, \sigma \rangle} \text{T-RESET}$$

$$\frac{\text{II does not contain } \dagger \quad k' \text{ fresh} \quad S' \text{ is obtained from } S \text{ by replacing } k \text{ with } k'}{\text{env}_F \vdash \langle \mathbb{R} \{F[\text{shift } k \{S\}]\}, \rho; II \dagger; II', \sigma \rangle \rightarrow \langle \mathbb{R} \{S'\}, \rho[k' \mapsto \mathbb{R} \{F[\text{skip}]\}], \rho; II \dagger; II', \sigma \rangle} \text{T-SHIFT}$$

$$\frac{\text{II does not contain } \dagger}{\text{env}_F \vdash \langle \mathbb{R} \{\text{skip}\}, II; \rho \dagger; II', \sigma \rangle \mapsto \langle \text{skip}, \rho; II', \sigma \rangle} \text{T-R}$$

$$\frac{\llbracket k \rrbracket_{(II, \sigma)} = [S, II']}{\text{env}_F \vdash \langle \text{invoke } k, II, \sigma \rangle \rightarrow \langle S; \text{return}, II'; II, \sigma \rangle} \text{T-VOKE}$$

$$\frac{}{\text{env}_F \vdash \langle \text{skip}; S, II, \sigma \rangle \rightarrow \langle S, II, \sigma \rangle} \text{T-SKIP}$$

$$\frac{\text{env}_F \ f = (\bar{x}, S) \quad \bar{v} = \llbracket \bar{e} \rrbracket_{(II, \sigma)}}{\text{env}_F \vdash \langle f(\bar{e}), II, \sigma \rangle \rightarrow \langle S; \text{return}, [\bar{x} \mapsto \bar{v}]; II, \sigma \rangle} \text{T-CALL}$$

$$\frac{}{\text{env}_F \vdash \langle \text{return}, \rho; II, \sigma \rangle \rightarrow \langle \text{skip}, II, \sigma \rangle} \text{T-RETURN}$$

$$\frac{(II, \sigma) \models e}{\text{env}_F \vdash \langle \text{if } e \text{ then } S_1 \text{ else } S_2, II, \sigma \rangle \rightarrow \langle S_1, II, \sigma \rangle} \text{T-IF-TRUE}$$

$$\frac{(II, \sigma) \not\models e}{\text{env}_F \vdash \langle \text{if } e \text{ then } S_1 \text{ else } S_2, II, \sigma \rangle \rightarrow \langle S_2, II, \sigma \rangle} \text{T-IF-FALSE}$$

$$\frac{(II, \sigma) \models e}{\text{env}_F \vdash \langle \text{while } e \text{ do } S, II, \sigma \rangle \rightarrow \langle S; \text{while } e \text{ do } S, II, \sigma \rangle} \text{T-WHILE-TRUE}$$

$$\frac{(II, \sigma) \not\models e}{\text{env}_F \vdash \langle \text{while } e \text{ do } S, II, \sigma \rangle \rightarrow \langle \text{skip}, II, \sigma \rangle} \text{T-WHILE-FALSE}$$

$$\frac{x \notin \text{dom } \rho}{\text{env}_F \vdash \langle \text{var } x = e, \rho; II, \sigma \rangle \rightarrow \langle \text{skip}, \rho[x \mapsto \llbracket e \rrbracket_{(\rho; II, \sigma)}]; II, \sigma \rangle} \text{T-VAR}$$

$$\frac{x \in \text{dom } \rho}{\text{env}_F \vdash \langle x := e, \rho; II, \sigma \rangle \rightarrow \langle \text{skip}, \rho[x \mapsto \llbracket e \rrbracket_{(\rho; II, \sigma)}]; II, \sigma \rangle} \text{T-ASSIGN-LOCAL}$$

$$\frac{u \in \text{dom } \sigma}{\text{env}_F \vdash \langle u := e, II, \sigma \rangle \rightarrow \langle \text{skip}, II, \sigma[u \mapsto \llbracket e \rrbracket_{(II, \sigma)}] \rangle} \text{T-ASSIGN-GLOBAL}$$

Fig. 6. Semantics of the target language

$$\begin{aligned}
\text{env}_F \vdash & \langle \text{reset } \{u' := 1; \text{shift } k \{u := k\}; u' := 0\}; \text{invoke } u, \emptyset, [u \mapsto 0; u' \mapsto 0] \rangle \\
& \mapsto \langle \text{R } \{u' := 1; \text{shift } k \{u := k\}; u' := 0\}; \text{invoke } u, \emptyset \dagger, [u \mapsto 0; u' \mapsto 0] \rangle \\
& \mapsto \langle \text{R } \{\text{shift } k \{u := k\}; u' := 0\}; \text{invoke } u, \emptyset \dagger, [u \mapsto 0; u' \mapsto 1] \rangle \\
& \mapsto \langle \text{R } \{u := k'\}; \text{invoke } u, [k' \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]] \dagger, [u \mapsto 0; u' \mapsto 1] \rangle \\
& \mapsto \langle \text{R } \{\text{skip}\}; \text{invoke } u, [k' \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]] \dagger, [u \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]; u' \mapsto 1] \rangle \\
& \mapsto \langle \text{invoke } u, [k' \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]], [u \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]; u' \mapsto 1] \rangle \\
& \mapsto \langle \text{R } \{\text{skip}; u' := 0\}; \text{return}, \emptyset \dagger; \\
& \quad [k' \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]], [u \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]; u' \mapsto 1] \rangle \\
& \mapsto \langle \text{R } \{u' := 0\}; \text{return}, \emptyset \dagger; [k' \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]], [u \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]; u' \mapsto 1] \rangle \\
& \mapsto \langle \text{R } \{\text{skip}\}; \text{return}, \emptyset \dagger; [k' \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]], [u \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]; u' \mapsto 0] \rangle \\
& \mapsto \langle \text{return}, \emptyset; [k' \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]], [u \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]; u' \mapsto 0] \rangle \\
& \mapsto \langle \text{skip}, [k' \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]], [u \mapsto [\text{R } \{\text{skip}; u' := 0\}, \emptyset \dagger]; u' \mapsto 0] \rangle
\end{aligned}$$

Fig. 7. Capturing and invoking a continuation

be left in the stack, which is popped by the trailing **return**. An example of how to capture and invoke a continuation is shown in Figure 7. In the example, we assume that the variables u and u' are global.

Procedure call $f(\bar{e})$ reduces to $S; \text{return}$ where S is the body of the procedure f , and pushes a local state $[\bar{x} \mapsto \bar{v}]$, binding procedure's formal arguments to actual arguments, into the stack (T-CALL). The trailing **return** ensures that, once the execution of $S; \text{return}$ terminates, the stack is aligned to the original Π . **return** pops the topmost element from the stack (T-RETURN). The remaining rules are self-explanatory.

Given the basic reduction rules, we now define a standard reduction, denoted by \mapsto , by

$$\frac{\text{env}_F \vdash \langle S, \Pi, \sigma \rangle \rightarrow \langle S', \Pi', \sigma' \rangle}{\text{env}_F \vdash \langle E[S], \Pi, \sigma \rangle \mapsto \langle E[S'], \Pi', \sigma' \rangle}$$

stating that the configuration $\langle S, \Pi, \sigma \rangle$ standard reduces to $\langle S', \Pi', \sigma' \rangle$ if there exist an evaluation context E and statement S_0 and S'_0 such that $S = E[S_0]$ and $S' = E[S'_0]$ and $\text{env}_F \vdash \langle S_0, \Pi, \sigma \rangle \rightarrow \langle S'_0, \Pi', \sigma' \rangle$. The standard reduction is deterministic.

4 Compilation

When compiling a program P into the target language, we compile expressions and statements according to the scheme shown in Figure 8. Expressions are translated into the target language as-is, and statements that have a corresponding equivalent in the target language are also translated in a straightforward manner. The two statements that have no direct correspondence in the target language are **await** and **spawn**. We look at how these statements are translated and how they interact with the scheduler later.

$$\begin{aligned}
\llbracket e \rrbracket &= e \\
\llbracket x := e \rrbracket &= x := \llbracket e \rrbracket \\
\llbracket u := e \rrbracket &= u := \llbracket e \rrbracket \\
\llbracket \text{skip} \rrbracket &= \text{skip} \\
\llbracket S_1; S_2 \rrbracket &= \llbracket S_1 \rrbracket; \llbracket S_2 \rrbracket \\
\llbracket \text{if } e \text{ then } S_1 \text{ else } S_2 \rrbracket &= \text{if } \llbracket e \rrbracket \text{ then } \llbracket S_1 \rrbracket \text{ else } \llbracket S_2 \rrbracket \\
\llbracket \text{while } e \text{ do } S \rrbracket &= \text{while } \llbracket e \rrbracket \text{ do } \llbracket S \rrbracket \\
\llbracket \text{var } x = e \rrbracket &= \text{var } x = \llbracket e \rrbracket \\
\llbracket \text{await } e \rrbracket &= \text{shift } k \{ T := \text{add } T \llbracket e \rrbracket k \}; \text{skip} \\
\llbracket f(\bar{e}) \rrbracket &= f_S(\llbracket \bar{e} \rrbracket) \\
\llbracket \text{spawn } f \bar{e} \rrbracket &= f_A(\llbracket \bar{e} \rrbracket)
\end{aligned}$$

Fig. 8. Compilation of source programs

$$\text{Sched} = \text{while } (\neg \text{isEmpty } T) \text{ do reset } \{ k := \text{get } T; T := \text{del } T k; \text{invoke } k \}$$

Fig. 9. Scheduler

The central idea of the compilation scheme is to use continuations to handle the suspension of tasks, and have an explicit *scheduler* (for brevity, in the examples we use *Sched* to denote the scheduler), shown in Figure 9. The control will pass to the scheduler every time a task either suspends or finishes, and the scheduler will pick up a new task to execute. During runtime, we also use a global variable *T*, which we assume not to be used by the program to be compiled. The global variable *T* stores the *task set*, corresponding to Θ in the source semantics, that contains all the tasks in the system. The tasks are stored as continuations with guard expressions.

The scheduler loops until the task set is empty (all tasks have terminated), in each iteration picking a continuation from *T* where the guard expression evaluates to **true**, removing it from *T* and then invoking the continuation using **invoke** *k*. The body of the scheduler is wrapped in a **reset**, guaranteeing that when a task suspends, the capture will be limited to the end of the current task. After the execution is completed – either by suspension or by just finishing the work – the control comes back to the scheduler.

Suspension (**await** *e* in the source language) is compiled to a **shift** statement. When evaluating the statement, the original computation until the end of the enclosing **R** $\{ \}$ will be captured and stored in the continuation *k*, and the original program is replaced with the body of the **shift**. The enclosing **R** $\{ \}$ guarantees that we capture only the statements up until the end of the current task, thus providing a facility to proceed with the execution of the task later.

$$\begin{aligned}
\llbracket n, \sigma \triangleright n \langle \Pi \rangle \parallel \Theta \rrbracket &= \langle \text{Sched}, \emptyset, \sigma[T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
\llbracket n, \sigma \triangleright n \langle e, S, \Pi \rangle \parallel \Theta \rrbracket &= \langle \text{Sched}, \emptyset, \sigma[T \mapsto \llbracket \Theta \rrbracket_2 \cup \llbracket n \langle e, S, \rho \rangle \rrbracket_2] \rangle \\
\llbracket n, \sigma \triangleright n \langle S, \Pi \rangle \parallel \Theta \rrbracket &= \langle \mathbf{R} \{ \mathbf{R} \{ \llbracket S \rrbracket \}; \mathbf{return} \}; \text{Sched}, \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
\llbracket \Theta \parallel \Theta' \rrbracket_2 &= \llbracket \Theta \rrbracket_2 \cup \llbracket \Theta' \rrbracket_2 \\
\llbracket n \langle e, S, \Pi \rangle \rrbracket_2 &= \{ \langle \llbracket e \rrbracket, [\mathbf{R} \{ \mathbf{skip}; \llbracket S \rrbracket \}, \Pi \dagger] \rangle \} \\
\llbracket n \langle \Pi \rangle \rrbracket_2 &= \emptyset
\end{aligned}$$

Fig. 10. Compilation of configurations

The body of the `shift` statement simply takes the captured continuation, k , and adds it to the global task set, with the appropriate guard expression. After adding the continuation to the task set, the control passes back to the scheduler.

Procedures in env_F will get translated into two different procedures for synchronous and asynchronous calls, as follows:

$$\begin{aligned}
f_S &\mapsto \llbracket S \rrbracket \\
f_A &\mapsto \mathbf{reset} \{ \mathbf{shift} \ k \ \{ T := \mathbf{add} \ T \ \mathbf{true} \ k \}; \llbracket S \rrbracket \}
\end{aligned}$$

When making an asynchronous call, the body of the procedure will be immediately captured in a continuation, added to the global task set, and the control passes back to the invoker via the usual synchronous call mechanism.

The entry point of a program in the source language, `main`, is a regular procedure and will get translated according to the usual rules into two procedures, `mainA` and `mainS`. In the target language, we must invoke the scheduler, and thus we use a different entry point:

$$T := \emptyset; \text{main}_A(); \text{Sched}$$

After initializing the task set to be empty, the first statement will add an asynchronous call to the original entry point of the program, and passes control to the scheduler. As there is only one task in the task set – the task that will invoke the original entry point – the scheduler will immediately proceed with that.

5 Correctness

In this section, we prove that our compilation scheme is correct in the sense that it preserves the operational behavior from the source program into the (compiled) target program. Specifically, we prove that reductions in the source language are simulated by corresponding reductions in the target language. To do so, we extend the compilation scheme to configurations in Figure 10.

The compilation scheme for configurations follows the idea of the compilation scheme detailed in Section 4. We have two compilation functions: $\llbracket \cdot \rrbracket_2$, which

$$\frac{\rho \ x = [S, \Pi']}{\text{env}_F \vdash \langle \text{invoke } x, \rho; \Pi, \sigma \rangle \rightarrow \langle S; \text{return}, \Pi'; \rho \setminus x; \Pi, \sigma \rangle} \text{T-INVOKEOONCE}$$

Fig. 11. Alternative rule for `invoke`

generates a task set from Θ , and $\llbracket \cdot \rrbracket$, which generates a configuration in the target semantics.

Every suspended task in the task set Θ is compiled to a pair consisting of the compiled guard expression and a continuation that has been constructed from the original statement and stack. The statement is wrapped in a `R { }` block and we prepend a `skip` statement, just as it would happen when a continuation is captured in the target language.

If the active task is finished or is suspended (but no new task has been scheduled yet), the generated configuration will immediately contain the scheduler. If the task has suspended, the task is compiled according to the previously described scheme and appended to T . Active tasks are wrapped in two `R { }` blocks and the stack Π is concatenated on top of the local state of the scheduler.

When the scheduler invokes a continuation k , the continuation will stay in the local state of the scheduler until control comes back to the scheduler. This is unnecessary, as the value is never used after it has been invoked; furthermore, the variable is immediately assigned a new value after control passes back to the scheduler. Thus, as an optimization, we may switch to an alternative reduction rule for `invoke` k , which only allows a continuation to be used once, `T-INVOKEOONCE`, shown in Figure 11. Although the behavior of the program is equivalent under both versions, using the one-shot version also allows us to state the correctness theorem in a more concise and straightforward manner, as the local state of the scheduler will always be empty when we are currently executing some task. In the proof, we assume this rule to be used instead of the original `T-INVOKE` rule.

The following lemma states that the compilation of statements is compositional with respect to evaluation contexts, where evaluation contexts for the source language are defined inductively by

$$K := [] \mid K; S$$

Lemma 1.

$$\llbracket K[S] \rrbracket = \llbracket K \rrbracket [\llbracket S \rrbracket]$$

Proof. By induction on the structure of K . □

The correctness theorem below states that a one-step reduction in the source language is simulated by multiple-step reductions in the target language.

As an example, in Figure 12 we show the compiled form for both the initial and final configurations shown for the step in Figure 4 and in Figure 13, we show how to reach the compiled equivalent of the configuration in multiple steps in the target semantics.

$$\begin{aligned}
& \llbracket \text{main}, [u \mapsto 0] \triangleright \text{main}(\text{await } u = 1; u := 2, \emptyset) \parallel f(\text{true}, u := 1, \emptyset) \rrbracket \\
& = \langle \mathbb{R} \{ \mathbb{R} \{ \text{shift } k \{ T := \text{add } T(u = 1) k \}; \text{skip}; u := 2 \}; \text{return} \}; \text{Sched}, \emptyset \dagger; \emptyset \dagger, \\
& \quad [u \mapsto 0; T \mapsto \{ (\text{true}, [\mathbb{R} \{ \text{skip}; u := 1 \}, \emptyset \dagger]) \}] \rangle \\
& \llbracket \text{main}, [u \mapsto 0] \triangleright \text{main}(u = 1, \text{skip}; u := 2, \emptyset) \parallel f(\text{true}, u := 1, \emptyset) \rrbracket \\
& = \langle \text{Sched}, \emptyset, [u \mapsto 0; T \mapsto \{ (u = 1, [\mathbb{R} \{ \text{skip}; \text{skip}; u := 2 \}, \emptyset \dagger]), (\text{true}, [\mathbb{R} \{ \text{skip}; u := 1 \}, \emptyset \dagger]) \}] \rangle
\end{aligned}$$

Fig. 12. Example of compiling a configuration

$$\begin{aligned}
& \langle \mathbb{R} \{ \mathbb{R} \{ \text{shift } k \{ T := \text{add } T(u = 1) k \}; \text{skip}; u := 2 \}; \text{return} \}; \text{Sched}, \emptyset \dagger; \emptyset \dagger, \\
& \quad [u \mapsto 0; T \mapsto \{ (\text{true}, [\mathbb{R} \{ \text{skip}; u := 1 \}, \emptyset \dagger]) \}] \rangle \\
& \mapsto \langle \mathbb{R} \{ \mathbb{R} \{ T := \text{add } T(u = 1) k' \}; \text{return} \}; \text{Sched}, [k' \mapsto [\mathbb{R} \{ \text{skip}; \text{skip}; u := 2 \}, \emptyset \dagger]] \dagger; \emptyset \dagger, \\
& \quad [u \mapsto 0; T \mapsto \{ (\text{true}, [\mathbb{R} \{ \text{skip}; u := 1 \}, \emptyset \dagger]) \}] \rangle \\
& \mapsto \langle \mathbb{R} \{ \mathbb{R} \{ \text{skip} \}; \text{return} \}; \text{Sched}, [k' \mapsto [\mathbb{R} \{ \text{skip}; \text{skip}; u := 2 \}, \emptyset \dagger]] \dagger; \emptyset \dagger, \\
& \quad [u \mapsto 0; T \mapsto \{ (\text{true}, [\mathbb{R} \{ \text{skip}; u := 1 \}, \emptyset \dagger]), (u = 1, [\mathbb{R} \{ \text{skip}; \text{skip}; u := 2 \}, \emptyset \dagger]) \}] \rangle \\
& \mapsto \langle \mathbb{R} \{ \text{return} \}; \text{Sched}, [k' \mapsto [\mathbb{R} \{ \text{skip}; \text{skip}; u := 2 \}, \emptyset \dagger]] \dagger; \emptyset \dagger, \\
& \quad [u \mapsto 0; T \mapsto \{ (\text{true}, [\mathbb{R} \{ \text{skip}; u := 1 \}, \emptyset \dagger]), (u = 1, [\mathbb{R} \{ \text{skip}; \text{skip}; u := 2 \}, \emptyset \dagger]) \}] \rangle \\
& \mapsto \langle \mathbb{R} \{ \text{skip} \}; \text{Sched}, \emptyset \dagger, \\
& \quad [u \mapsto 0; T \mapsto \{ (\text{true}, [\mathbb{R} \{ \text{skip}; u := 1 \}, \emptyset \dagger]), (u = 1, [\mathbb{R} \{ \text{skip}; \text{skip}; u := 2 \}, \emptyset \dagger]) \}] \rangle \\
& \mapsto \langle \text{Sched}, \emptyset, [u \mapsto 0; T \mapsto \{ (\text{true}, [\mathbb{R} \{ \text{skip}; u := 1 \}, \emptyset \dagger]), (u = 1, [\mathbb{R} \{ \text{skip}; \text{skip}; u := 2 \}, \emptyset \dagger]) \}] \rangle
\end{aligned}$$

Fig. 13. Reduction of the compiled configuration

Theorem 1. For all configurations cfg_S and cfg'_S such that

$$\text{env}_F \vdash cfg_S \rightarrow cfg'_S$$

holds, then the following must also hold:

$$\llbracket \text{env}_F \rrbracket \vdash \llbracket cfg_S \rrbracket \mapsto^+ \llbracket cfg'_S \rrbracket$$

Proof. By induction over the derivation, analyzing the step taken. The possible steps have one of the following forms:

– Case

$$\text{env}_F \vdash n, \sigma \triangleright n \langle S, \Pi \rangle \parallel \Theta \rightarrow n', \sigma' \triangleright n \langle \Pi' \rangle \parallel \Theta'$$

Rules matching this pattern are S-ASSIGN-LOCAL, S-ASSIGN-GLOBAL, S-WHILE-FALSE, S-SPAWN, S-RETURN, S-VAR. As a representative example, we will look at S-SPAWN in detail.

$$\frac{\text{env}_F f = (\bar{x}, S) \quad \bar{v} = \llbracket \bar{e} \rrbracket_{(\rho, \sigma)} \quad n' \text{ is fresh}}{\text{env}_F \vdash n, \sigma \triangleright n \langle \text{spawn } f \bar{e}, \rho; \Pi \rangle \rightarrow n, \sigma \triangleright n \langle \rho; \Pi \rangle \parallel n' \langle \text{true}, S, [\bar{x} \mapsto \bar{v}] \rangle}$$

In this case, the source and target configurations are compiled to:

$$\begin{aligned}
\llbracket cfg_S \rrbracket & = \langle \mathbb{R} \{ \mathbb{R} \{ f_A(\bar{e}) \}; \text{return} \}; \text{Sched}, \rho; \Pi \dagger; \emptyset \dagger, \sigma [T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
\llbracket cfg'_S \rrbracket & = \langle \text{Sched}, \emptyset, \sigma [T \mapsto \llbracket \Theta \rrbracket_2 \cup \{ (\text{true}, [\mathbb{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v} \dagger]) \}] \rangle
\end{aligned}$$

Let the bottommost element of Π be ρ' , where $\rho = \rho'$ if Π is empty. The compiled source configuration will reduce as follows:

$$\begin{aligned}
& \llbracket \text{env}_F \rrbracket \vdash \langle \mathbf{R} \{ \mathbf{R} \{ f_A(\bar{e}) \}; \text{return} \}; \text{Sched}, \rho; \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
& \mapsto \langle \mathbf{R} \{ \mathbf{R} \{ \text{reset } \{ \text{shift } k \{ T := \text{add } T \text{ true } k \}; \llbracket S \rrbracket \}; \text{return} \}; \text{return} \}; \text{Sched}, \\
& \quad [\bar{x} \mapsto \bar{v}]; \rho; \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
& \mapsto \langle \mathbf{R} \{ \mathbf{R} \{ \text{shift } k \{ T := \text{add } T \text{ true } k \}; \llbracket S \rrbracket \}; \text{return} \}; \text{return} \}; \text{Sched}, \\
& \quad [\bar{x} \mapsto \bar{v}]; \rho; \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
& \mapsto \langle \mathbf{R} \{ \mathbf{R} \{ T := \text{add } T \text{ true } k \}; \text{return} \}; \text{return} \}; \text{Sched}, \\
& \quad [\bar{x} \mapsto \bar{v}, k \mapsto \mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]; \rho; \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
& \mapsto \langle \mathbf{R} \{ \mathbf{R} \{ \text{skip} \}; \text{return} \}; \text{return} \}; \text{Sched}, [\bar{x} \mapsto \bar{v}, k \mapsto \mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]; \rho; \Pi \dagger; \emptyset \dagger, \\
& \quad \sigma[T \mapsto \llbracket \Theta \rrbracket_2 \cup \{ (\text{true}, [\mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]) \}] \rangle \\
& \mapsto \langle \mathbf{R} \{ \text{return} \}; \text{return} \}; \text{Sched}, [\bar{x} \mapsto \bar{v}, k \mapsto \mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]; \rho; \Pi \dagger; \emptyset \dagger, \\
& \quad \sigma[T \mapsto \llbracket \Theta \rrbracket_2 \cup \{ (\text{true}, [\mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]) \}] \rangle \\
& \mapsto \langle \mathbf{R} \{ \text{skip} \}; \text{return} \}; \text{Sched}, \rho; \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2 \cup \{ (\text{true}, [\mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]) \}] \rangle \\
& \mapsto \langle \text{return} \}; \text{Sched}, \rho'; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2 \cup \{ (\text{true}, [\mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]) \}] \rangle \\
& \mapsto \langle \text{skip} \}; \text{Sched}, \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2 \cup \{ (\text{true}, [\mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]) \}] \rangle \\
& \mapsto \langle \text{Sched}, \emptyset, \sigma[T \mapsto \llbracket \Theta \rrbracket_2 \cup \{ (\text{true}, [\mathbf{R} \{ \text{skip}; \llbracket S \rrbracket \}, [\bar{x} \mapsto \bar{v}]) \}] \rangle
\end{aligned}$$

The configuration we obtain from evaluation is exactly equal to the compiled configuration, thus for this case our claim holds.

– Case

$$\text{env}_F \vdash n, \sigma \triangleright n \langle S, \Pi \rangle \parallel \Theta \rightarrow n', \sigma' \triangleright n \langle e, S', \Pi' \rangle \parallel \Theta'$$

There are only two possible rules: S-SEQ-GRD and S-AWAIT. In both cases, it must be that $\sigma = \sigma'$, $\Pi = \Pi'$, $\Theta \equiv \Theta'$ and there exists some K such that $S = K[\text{await } e]$ and $S' = K[\text{skip}]$. Therefore, taking into account Lemma 1, the source and target configurations are compiled to:

$$\begin{aligned}
\llbracket \text{cfg}_S \rrbracket &= \langle \mathbf{R} \{ \mathbf{R} \{ \llbracket K \rrbracket \llbracket \llbracket \text{await } e \rrbracket \rrbracket \}; \text{return} \}; \text{Sched}, \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
&= \langle \mathbf{R} \{ \mathbf{R} \{ \llbracket K \rrbracket \llbracket \text{shift } k \{ T := \text{add } T \llbracket e \rrbracket k \}; \text{skip} \}; \text{return} \}; \text{Sched}, \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \llbracket \Theta \rrbracket_2] \rangle \\
\llbracket \text{cfg}'_S \rrbracket &= \langle \text{Sched}, \emptyset, \sigma[T \mapsto \llbracket \Theta \rrbracket_2 \cup \{ (\llbracket e \rrbracket, \mathbf{R} \{ \llbracket K \rrbracket \llbracket \text{skip} \rrbracket \}, \Pi \dagger) \}] \rangle
\end{aligned}$$

An example of this reduction can be seen in Figure 13.

– Case

$$\text{env}_F \vdash n, \sigma \triangleright n \langle S, \Pi \rangle \parallel \Theta \rightarrow n', \sigma' \triangleright n \langle S', \Pi' \rangle \parallel \Theta'$$

Rules matching this pattern are S-SEQ-FIN, S-SEQ-STEP, S-IF-TRUE, S-IF-FALSE, S-WHILE-TRUE, S-CALL. In the case of S-SEQ-STEP, we know that $S = S_0; S_1$ and $S' = S'_0; S_1$. By induction hypothesis, we get that

$$\llbracket \text{env}_F \rrbracket \vdash \llbracket n, \sigma \triangleright n \langle S_0, \Pi \rangle \parallel \Theta \rangle \rightarrow \llbracket n', \sigma' \triangleright n \langle S'_0, \Pi' \rangle \parallel \Theta' \rrbracket$$

As by the definition of the compilation function $\llbracket S \rrbracket = \llbracket S_0 \rrbracket; \llbracket S_1 \rrbracket$ and $\llbracket S' \rrbracket = \llbracket S'_0 \rrbracket; \llbracket S_1 \rrbracket$, we obtain the needed result:

$$\llbracket \text{env}_F \rrbracket \vdash \llbracket n, \sigma \triangleright n \langle S_0; S_1, \Pi \rangle \parallel \Theta \rangle \rightarrow \llbracket n', \sigma' \triangleright n \langle S'_0; S_1, \Pi' \rangle \parallel \Theta' \rrbracket$$

For S-SEQ-FIN, we know that $S = S_0; S_1$ and $S' = S_1$. Then the case follows by analyzing the step taken to reduce S_0 .

The other cases are straightforward.

– One of the following three:

$$\begin{aligned} & \text{env}_F \vdash n, \sigma \triangleright n \langle \Pi' \rangle \parallel n' \langle e, S, \Pi \rangle \parallel \Theta \rightarrow n', \sigma \triangleright n \langle \Pi' \rangle \parallel n' \langle S, \Pi \rangle \parallel \Theta \\ & \text{env}_F \vdash n, \sigma \triangleright n \langle e', S', \Pi' \rangle \parallel n' \langle e, S, \Pi \rangle \parallel \Theta \rightarrow n', \sigma \triangleright n \langle e', S', \Pi' \rangle \parallel n' \langle S, \Pi \rangle \parallel \Theta \\ & \text{env}_F \vdash n, \sigma \triangleright n \langle e, S, \Pi \rangle \parallel \Theta \rightarrow n, \sigma \triangleright n \langle S, \Pi \rangle \parallel \Theta \end{aligned}$$

These three patterns match each of the scheduling rules. We will look only at the first one.

$$\frac{(\rho, \sigma) \models e}{\text{env}_F \vdash n, \sigma \triangleright n \langle \Pi' \rangle \parallel n' \langle e, S, \rho; \Pi \rangle \rightarrow n', \sigma \triangleright n' \langle S, \rho; \Pi \rangle \parallel n \langle \Pi' \rangle} \text{S-SCHED-FIN}$$

$$\begin{aligned} \llbracket \text{cfg}_S \rrbracket &= \langle \text{Sched}, \emptyset, \sigma[T \mapsto \{(\llbracket e \rrbracket, [\mathbf{R} \{\text{skip}; \llbracket S \rrbracket], \rho; \Pi \dagger\})\}] \rangle \\ \llbracket \text{cfg}'_S \rrbracket &= \langle \mathbf{R} \{\mathbf{R} \{\llbracket S \rrbracket\}; \text{return}\}; \text{Sched}, \rho; \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \emptyset] \rangle \end{aligned}$$

The initial configuration will reduce as (with some of the steps omitted):

$$\begin{aligned} & \llbracket \text{env}_F \rrbracket \vdash \langle \text{while } (\neg \text{isEmpty } T) \text{ do reset } \{k := \text{get } T; T := \text{del } T \ k; \text{invoke } k\}, \\ & \quad \emptyset, \sigma[T \mapsto \{(\llbracket e \rrbracket, [\mathbf{R} \{\text{skip}; \llbracket S \rrbracket], \rho; \Pi \dagger\})\}] \rangle \\ & \mapsto \langle \text{reset } \{k := \text{get } T; T := \text{del } T \ k; \text{invoke } k\}; \text{Sched}, \\ & \quad \emptyset, \sigma[T \mapsto \{(\llbracket e \rrbracket, [\mathbf{R} \{\text{skip}; \llbracket S \rrbracket], \rho; \Pi \dagger\})\}] \rangle \\ & \mapsto \langle \mathbf{R} \{k := \text{get } T; T := \text{del } T \ k; \text{invoke } k\}; \text{Sched}, \emptyset \dagger, \sigma[T \mapsto \{(\llbracket e \rrbracket, [\text{skip}; \mathbf{R} \{\llbracket S \rrbracket\}], \rho; \Pi \dagger\})\}] \rangle \\ & \mapsto^* \langle \mathbf{R} \{\text{invoke } k\}; \text{Sched}, [k \mapsto [\mathbf{R} \{\text{skip}; \llbracket S \rrbracket], \rho; \Pi \dagger\}] \dagger, \sigma[T \mapsto \emptyset] \rangle \\ & \mapsto \langle \mathbf{R} \{\mathbf{R} \{\text{skip}; \llbracket S \rrbracket\}; \text{return}\}; \text{Sched}, \rho; \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \emptyset] \rangle \\ & \mapsto \langle \mathbf{R} \{\mathbf{R} \{\llbracket S \rrbracket\}; \text{return}\}; \text{Sched}, \rho; \Pi \dagger; \emptyset \dagger, \sigma[T \mapsto \emptyset] \rangle \end{aligned}$$

6 Conclusion

In this paper, we formalized a compilation scheme for cooperative multi-tasking into delimited continuations. For the source language, we extend While with procedure calls and operations for blocking and creation of new tasks. The target language extends While with shift/reset—the target language is sequential. We then proved that the compilation scheme is correct: reductions in the source language are simulated by corresponding reductions in the target language. We have implemented this compilation scheme in our compiler from ABS to Scala. The compiler covers a much richer language than our source language, including object-oriented features, and employs the experimental continuations plugin for Scala. The compiler is integrated into the wider ABS Tool Suite, available at <http://tools.hats-project.eu/>. We are currently formalizing the results of the paper in the proof assistant Agda.

Acknowledgments This research was supported by the EU FP7 ICT project no. 231620 (HATS), the Estonian Centre of Excellence in Computer Science, EXCS, financed mainly by the European Regional Development Fund, ERDF, the Estonian Ministry of Education and Research target-financed research theme no. 0140007s12, and the Estonian Science Foundation grant no. 9398.

References

1. Adya, A., Howell, J., Theimer, M., Bolosky, W.J., Douceur, J.R.: Cooperative task management without manual stack management. In: ATEC '02: Proceedings of the General Track of the annual conference on USENIX Annual Technical Conference. pp. 289–302. USENIX (2002)
2. Agha, G.: Actors: a model of concurrent computation in distributed systems. MIT Press (1986)
3. de Boer, F.S., Clarke, D., Johnsen, E.B.: A complete guide to the future. In: ESOP'07 Proceedings of the 16th European conference on Programming. LNCS, vol. 4421, pp. 316–330. Springer (2007)
4. Caromel, D., Henrio, L., Serpette, B.P.: Asynchronous and deterministic objects. ACM SIGPLAN Notices - POPL'04 39(1), 123–134 (2004)
5. Danvy, O., Filinski, A.: Abstracting control. In: LFP '90: Proceedings of the 1990 ACM conference on LISP and functional programming. pp. 151–160. ACM (1990)
6. Haynes, C.T., Friedman, D.P., Wand, M.: Continuations and coroutines. In: LFP '84: Proceedings of the 1984 ACM Symposium on LISP and functional programming. pp. 293–298. ACM (1984)
7. Johnsen, E.B., Hähnle, R., Schäfer, J., Schlatte, R., Steffen, M.: ABS: a core language for abstract behavioral specification. In: FMCO'10 Proceedings of the 9th international conference on Formal Methods for Components and Objects. LNCS, vol. 6957, pp. 142–164. Springer (2012)
8. Johnsen, E.B., Owe, O., Yu, I.C.: Creol: A type-safe object-oriented model for distributed concurrent systems. Theoretical Computer Science 365(1-2), 23–66 (2006)
9. Karmani, R.K., Shali, A., Agha, G.: Actor frameworks for the JVM platform: a comparative analysis. In: PPPJ '09: Proceedings of the 7th International Conference on Principles and Practice of Programming in Java. pp. 11–20. ACM (2009)
10. Schäfer, J., Poetzsch-Heffter, A.: JCoBox: generalizing active objects to concurrent components. In: ECOOP'10: Proceedings of the 24th European conference on Object-oriented programming. LNCS, vol. 6183, pp. 275–299. Springer (2010)
11. Srinivasan, S., Mycroft, A.: Kilim: Isolation-typed actors for Java. In: ECOOP '08: Proceedings of the 22nd European conference on Object-Oriented Programming. LNCS, vol. 5142, pp. 104–128. Springer (2008)
12. Wand, M.: Continuation-based multiprocessing. In: LFP '80: Proceedings of the 1980 ACM conference on LISP and functional programming. pp. 19–28. ACM (1980)