

# On the Soundness of Coroutines with Snapshots

Aleksandar Prokopec<sup>1</sup> and Fengyun Liu<sup>2</sup>

1 Oracle Labs, Switzerland

aleksandar.prokopec@oracle.com

2 École Polytechnique Fédérale de Lausanne, Switzerland

fengyun.liu@epfl.ch

---

## Abstract

*Coroutines* are a general control flow construct that can eliminate control flow fragmentation inherent in event-driven programs, which is still missing in many popular languages. Coroutines with snapshots are a first-class, type-safe, stackful coroutine model, which unifies many variants of suspendable computing, and is sufficiently general to express iterators, single-assignment variables, *async-await*, actors, event streams, backtracking, symmetric coroutines and continuations.

In this paper, we develop a formal model called  $\lambda_{\rightsquigarrow}$  that captures the essence of type-safe, stackful, delimited coroutines with snapshots. We prove the standard progress and preservation safety properties. Finally, we show a formal transformation from the  $\lambda_{\rightsquigarrow}$  calculus to the simply-typed lambda calculus with references.

**Keywords and phrases** coroutines, continuations, asynchronous programming, inversion of control

**Digital Object Identifier** 10.4230/LIPIcs.ECOOP.2018.23

## 1 Introduction

Asynchronous programming is becoming increasingly important, with applications ranging from actor systems [1, 14], futures and network programming [8, 15], user interfaces [21], to functional stream processing [23]. Traditionally, these programming models were realized either by blocking execution threads (which can be detrimental to performance [4]), or callback-style APIs [8, 15, 18], or with monads [53]. However, these approaches often feel unnatural, and the resulting programs can be hard to understand and maintain. *Coroutines* [9] overcome the need for blocking threads, callbacks and monads by allowing parts of the execution to pause at arbitrary points, and resuming that execution later.

In related work [38], we introduced *stackful coroutines with snapshots*, and showed that they are sufficiently general to express iterators, single-assignment variables, *async-await*, actors, event streams, backtracking, symmetric coroutines and continuations. Additionally, we provided an efficient implementation of coroutines with snapshots for Scala [31].

In this paper, we develop a formal model called  $\lambda_{\rightsquigarrow}$  that captures the essence of type-safe, stackful, delimited coroutines with snapshots. We prove the standard progress and preservation safety properties. Finally, we show a formal transformation from the  $\lambda_{\rightsquigarrow}$  calculus to the simply-typed lambda calculus with references.

## 2 Simply Typed Lambda Calculus with Coroutines ( $\lambda_{\rightsquigarrow}$ )

In addition to the standard abstraction, application and variable terms associated with the lambda calculus, the  $\lambda_{\rightsquigarrow}$  extension defines coroutines and the accompanying operations. A coroutine can be defined, called, resumed, suspended and copied. A coroutine definition is similar to a function definition, the main difference being that the body of the coroutine can



© Aleksandar Prokopec and Fengyun Liu;  
licensed under Creative Commons License CC-BY

Europe (ECOOP 2018).

Editors: John Q. Open and Joan R. Acces; Article No. 23; pp. 23:1–23:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

## 23:2 On the Soundness of Coroutines with Snapshots

suspend itself. Calling a coroutine creates an invocation value of that coroutine, which is initially paused. When resumed, that invocation runs until suspending, or completing. A coroutine instance is suspended whenever it evaluates the `yield` term (when this happens, we say that the evaluation *yields*), and can be subsequently resumed from the same point. The state of a coroutine instance can also be duplicated into a new instance, which is referred to as creating a *snapshot*.

A coroutine instance is represented not just by the respective term, but also by its suspended computation state (a coroutine *instance* is a *stateful* entity). Multiple variables can refer to (i.e. alias) the same coroutine instance, which holds some mutable state. For this reason, coroutine instances are represented with special instance labels  $i$ . Each instance label has a corresponding mapping in the *coroutine store*, as we will show.

Importantly, in the model that we will define, coroutines are *delimited*. This means that yielding (i.e. suspending a coroutine) must take place inside the program region that is specially marked as a coroutine, and only such regions of the program can get suspended. This region is *not* necessarily lexically scoped, since a coroutine definition can call into another coroutine definition (defined elsewhere). In other words, we model stackful delimited coroutines.

We start by defining the syntax for  $\lambda_{\rightsquigarrow}$ , which consists of the user terms, and the runtime terms (i.e. terms that arise only during evaluation).

**Definition 2.1** [Syntax] The  $\lambda_{\rightsquigarrow}$  *programming model* consists of the following terms, which can appear in user programs:

$t ::=$		user terms:
$(x:T) \Rightarrow t$		abstraction
$t(t)$		application
$x$		variable
$()$		unit value
$(x:T) \rightsquigarrow^T t$		coroutine
$\text{yield}(t)$		yielding
$\text{start}(t, t)$	coroutine instance creation	
$\text{resume}(t, v, v, v)$	resuming	
$\text{snapshot}(t)$	snapshot creation	
$\text{fix}(t)$	recursion	

The  $\lambda_{\rightsquigarrow}$  model defines the following types:

$T ::=$		types:
$T \Rightarrow T$		function type
$T \rightsquigarrow^T T$		coroutine type
$T \rightsquigarrow\rightsquigarrow T$	coroutine instance type	
$\text{Unit}$	unit type	
$\perp$	bottom type	

The  $\lambda_{\rightsquigarrow}$  *model of computation* additionally defines the following *runtime terms*, which cannot appear in a user program, but can appear during the evaluation of a program:

$r ::=$		runtime terms:
$i$		coroutine instance
$\langle t, v, v, v \rangle_i$		coroutine resumption
$\llbracket t \rrbracket_v$		suspension
$\emptyset$		empty term

The following subset of terms are considered *values*:

$v ::=$		values:
$(x:T) \Rightarrow t$		abstraction
$()$		unit value
$(x:T) \overset{T}{\rightsquigarrow} t$		coroutine
$i$		coroutine instance
$\emptyset$		empty term

We once more highlight the difference between a *coroutine*, which is akin to a function definition, and a *coroutine instance* which is akin to an invocation of a function. Since a coroutine instance, unlike a function invocation, can be suspended and resumed, it must be a first-class value that the program can refer to. We therefore distinguish between a coroutine type  $T_1 \overset{T_y}{\rightsquigarrow} T_2$  (where  $T_1$  is the input type,  $T_2$  is the return type, and  $T_y$  is the yield type) and the coroutine instance type  $T_y \rightsquigarrow T_2$  (where  $T_y$  is the yield type, and  $T_2$  is the return type).

Before defining the typing rules for  $\lambda_{\rightsquigarrow}$ , we first define the contexts in which the typing of a term takes place. There are two kinds of contexts that we need – the first is the standard typing context  $\Gamma$  used to track variable types, and the second is a *instance typing*  $\Sigma$ , used to track the types of the coroutine instances that exist at runtime.

**Definition 2.2** [Typing context] The *typing context*  $\Gamma$  is a sequence of variables and their respective types, where the comma operator  $(,)$  extends a typing context with a new binding:

$\Gamma ::=$		typing context:
$\emptyset$		empty context
$\Gamma, x:T$		variable binding

**Definition 2.3** [Instance typing] The *instance typing*  $\Sigma$  is a sequence of coroutine instance labels and their respective types, where the comma operator  $(,)$  extends an instance typing with a new binding:

$\Sigma ::=$		instance typing:
$\emptyset$		empty instance typing
$\Sigma, i:T$		new instance

Aside from tracking the type of each term, our typing rules will track the type of values that a term can yield. This allows typechecking coroutine declarations against the values yielded in their bodies. Therefore, our typing relation will be a five place relation between the typing context, instance typing, the term and its type, and the yield type.

**Definition 2.4** [Typing relation] The *typing relation*  $\Sigma|\Gamma \vdash t:T|T_y$  on  $\lambda_{\rightsquigarrow}$  is a relation between the instance typing  $\Sigma$ , the typing context  $\Gamma$ , the term  $t$ , the type of the term  $T$ , and the *yield type*  $T_y$ , where  $T_y$  denotes the type of values that may be yielded during the evaluation of the term  $t$ . The inductive definition of this typing relation is shown in Fig. 1 and Fig. 3.

We now briefly discuss the typing rules in Fig. 1. The rules T-ABS, T-APP, T-VAR and T-UNIT are standard in simply typed lambda calculus. In our case, we base the typing judgement on the instance typing  $\Sigma$  in addition to the typing context  $\Gamma$ . Furthermore, each typing judgement has the yield type as the last element. The rule T-APP allows the evaluation of the lambda  $\tau_1$  and the argument  $\tau_2$  to yield values of some type  $T_y$  (but the

$$\begin{array}{c}
\frac{\Sigma|\Gamma, \mathbf{x}:T_1 \vdash \mathbf{t}_2:T_2|\perp}{\Sigma|\Gamma \vdash (\mathbf{x}:T_1) \Rightarrow \mathbf{t}_2 : T_1 \Rightarrow T_2|\perp} \quad \text{(T-ABS)} \qquad \frac{\Sigma|\Gamma \vdash \mathbf{t}_1:T_2 \Rightarrow T_1|T_y}{\Sigma|\Gamma \vdash \mathbf{t}_1(\mathbf{t}_2) : T_1|T_y} \quad \text{(T-APP)} \qquad \frac{\mathbf{x}:T \in \Gamma}{\Sigma|\Gamma \vdash \mathbf{x}:T|\perp} \quad \text{(T-VAR)} \qquad \frac{\Sigma|\Gamma \vdash \mathbf{t}:T|\perp}{\Sigma|\Gamma \vdash \mathbf{t}:T|T_y} \quad \text{(T-CTX)} \\
\\
\Sigma|\Gamma \vdash () : \mathbf{Unit}|\perp \quad \text{(T-UNIT)} \qquad \frac{\Sigma|\Gamma, \mathbf{x}:T_1 \vdash \mathbf{t}_2:T_2|T_y}{\Sigma|\Gamma \vdash (\mathbf{x}:T_1) \overset{T_y}{\rightsquigarrow} \mathbf{t}_2 : T_1 \overset{T_y}{\rightsquigarrow} T_2|\perp} \quad \text{(T-COROUTINE)} \qquad \frac{\Sigma|\Gamma \vdash \mathbf{t}_1:T_1 \overset{T_y}{\rightsquigarrow} T_2|T_w}{\Sigma|\Gamma \vdash \mathbf{t}_2:T_1|T_w} \quad \text{(T-START)} \\
\\
\frac{\Sigma|\Gamma \vdash \mathbf{t}:T|T}{\Sigma|\Gamma \vdash \mathbf{yield}(\mathbf{t}) : \mathbf{Unit}|T} \quad \text{(T-YIELD)} \qquad \frac{\Sigma|\Gamma \vdash \mathbf{t}:T_y \rightsquigarrow T_2|T_w}{\Sigma|\Gamma \vdash \mathbf{snapshot}(\mathbf{t}) : T_y \rightsquigarrow T_2|T_w} \quad \text{(T-SNAPSHOT)} \qquad \frac{\Sigma|\Gamma \vdash \mathbf{t}:T \Rightarrow T|\perp}{\Sigma|\Gamma \vdash \mathbf{fix}(\mathbf{t}) : T|\perp} \quad \text{(T-FIX)} \\
\\
\frac{\Sigma|\Gamma \vdash \mathbf{t}_1:T_y \rightsquigarrow T_2|T_w \quad \Sigma|\Gamma \vdash \mathbf{t}_2:T_2 \overset{T_w}{\rightsquigarrow} T_R|T_w}{\Sigma|\Gamma \vdash \mathbf{t}_3:T_y \overset{T_w}{\rightsquigarrow} T_R|T_w} \quad \Sigma|\Gamma \vdash \mathbf{t}_4:\mathbf{Unit} \overset{T_w}{\rightsquigarrow} T_R|T_w \quad \Sigma|\Gamma \vdash \mathbf{t}_1:T_2 \overset{T_y}{\rightsquigarrow} T_1|T_y \quad \Sigma|\Gamma \vdash \mathbf{t}_2:T_2|T_y}{\Sigma|\Gamma \vdash \mathbf{resume}(\mathbf{t}_1, \mathbf{t}_2, \mathbf{t}_3, \mathbf{t}_4) : T_R|T_w} \quad \text{(T-RESUME)} \qquad \frac{\Sigma|\Gamma \vdash \mathbf{t}_1:T_2 \overset{T_y}{\rightsquigarrow} T_1|T_y \quad \Sigma|\Gamma \vdash \mathbf{t}_2:T_2|T_y}{\Sigma|\Gamma \vdash \mathbf{t}_1(\mathbf{t}_2) : T_1|T_y} \quad \text{(T-APPCOR)}
\end{array}$$

■ **Figure 1** Typing relation on terms

type  $T_y$  must be the same for both terms). The rule T-VAR assigns the bottom type  $\perp$  as the yield type of a variable, since this term does not yield any values.

Note that the rule T-ABS requires that the body of the lambda has the yield type  $\perp$ , that is, does not yield any values. The reason for this is that we model *delimited* coroutines. Yielding is only allowed from within a lexical scope of a coroutine definition, that is, a yielded value cannot cross a function boundary. In our concrete implementation, this means that we allow defining a coroutine that normally invokes a 3rd party library function, but the body of that 3rd party library function is not allowed to yield (unless the 3rd party library function is itself a coroutine defined using our transformation).

The rule T-CTX says that any term  $\mathbf{t}$  of type  $T$  and the yield type  $\perp$  can be assumed to have any yield type  $T_y$ . A typechecker may apply this rule when consolidating two terms one of which does not yield. For example, given a term  $\mathbf{t}_1$  whose yield type is  $T_y \neq \perp$ , and a term  $\mathbf{t}_2$  whose yield type is  $\perp$ , a typechecker must apply T-CTX before applying the rule T-APP.

The rule T-COROUTINE types a coroutine declaration. Similar to how T-ABS types a lambda, this rule types a coroutine using the return type and the yield type of its body. The yield type is “swallowed” by the coroutine, leaving  $\perp$  as the yield type of the resulting value.

Having inspected the rules for typing a coroutine declaration, we turn to coroutine operations. The rule T-START says that given a term of the coroutine type  $T_1 \overset{T_y}{\rightsquigarrow} T_2$ , and a term of type  $T_1$ , a **start** expression has the coroutine instance type  $T_y \rightsquigarrow T_2$ . The evaluation of the **start** term can itself yield a value of type  $T_w$ , independently of the yield type of the newly created coroutine instance – this means that the yielding context of the coroutine instance is separated from the context of the caller, where the caller is either the body of the enclosing coroutine, or the enclosing function, or the top-level program.

The T-YIELD rule says that if a term  $\mathbf{t}$  has a type  $T$ , and its yield type is also  $T$ , then **yield**( $\mathbf{t}$ ) has the type **Unit**, with the same yield type  $T$ . In other words, a coroutine can yield a value of type  $T$ , only if all the previous yields were of the same type  $T$  (or, if there were not previous yields, the typechecker can apply the T-CTX rule).

The T-RESUME rule describes the type of the **resume** expression. Consider resuming a coroutine instance of type  $T_y \rightsquigarrow T_2$ . Depending on the state of the instance, this has several

outcomes. First, the instance can complete and return a result value of type  $T_2$ . Second, the resumption of the instance can suspend itself and yield a value of type  $T_y$ . Finally, the instance could have already been completed when `resume` was called. Since  $\lambda_{\rightsquigarrow}$  does not have a variants or sum types to distinguish between these cases, the `resume` statement acts as a proverbial poor man's pattern matching. The values  $\mathfrak{t}_2$ ,  $\mathfrak{t}_3$  and  $\mathfrak{t}_4$  represent the code segments that deal with each of the above-described cases, and they return a result value of type  $T_R$ . Note that we could have modeled  $\mathfrak{t}_2$ ,  $\mathfrak{t}_3$  and  $\mathfrak{t}_4$  as function values. However, that would mean that the bodies of  $\mathfrak{t}_2$ ,  $\mathfrak{t}_3$  and  $\mathfrak{t}_4$  cannot yield values themselves (according to T-ABS). Hence, we model  $\mathfrak{t}_2$ ,  $\mathfrak{t}_3$  and  $\mathfrak{t}_4$  as coroutines with the yield type  $T_w$ , which corresponds to the yield type of the enclosing context.

The T-SNAPSHOT rule types the `snapshot` expression, which copies the given coroutine instance. The coroutine instance type and the yield type are preserved between the premise and the conclusion.

We mentioned that this model must describe stackful coroutines. To this end, a coroutine body must be able to invoke another coroutine as if it were a normal function. If that other function then yields, both the caller and the callee must be suspended. Note that coroutine invocation is syntactically equivalent to function application, but the callee is a coroutine, not a function. This is captured by the T-APPCOR rule, which additionally requires that the yield type  $T_y$  of the callee corresponds to the yield type of the current coroutine.

Finally, T-FIX is a standard typing rule for general recursion, which allows a coroutine or a lambda to refer to itself. We defer the discussion of the typing rules for runtime terms, shown in Fig. 3, until we cover the operational semantics of  $\lambda_{\rightsquigarrow}$ . Before we proceed, we define what it means for a program to be well-typed.

**Definition 2.5** [Well-typed program] A term  $\mathfrak{t}$  is *well-typed* if and only if  $\exists T, T_y$  and an instance typing  $\Sigma$  such that  $\Sigma|\emptyset \vdash \mathfrak{t} : T|T_y$ . Furthermore, a term  $\mathfrak{t}$  is a *well-typed user program* if  $\mathfrak{t}$  is well-typed and its yield type  $T_y = \perp$ . ◀

Coroutine instances are stateful – each coroutine instance maps to a term that it evaluates. Program evaluation is modeled not only as a transition between terms, but also between coroutine stores  $\mu$ , which we define next.

**Definition 2.6** [Coroutine store] A *coroutine store*  $\mu$  is a sequence of coroutine instance labels  $i$  bound to respective evaluation terms  $\mathfrak{t}$ , where the comma operator  $(,)$  extends the coroutine store with a new binding.

$$\begin{array}{ll} \mu ::= & \text{instance store:} \\ \emptyset & \text{empty instance store} \\ \mu, i \triangleright \mathfrak{t} & \text{instance binding} \end{array}$$

In what follows, we will use the convention that the evaluation term  $\mathfrak{t}$  in the coroutine store is suspended ( $\mathfrak{t} = \llbracket \mathfrak{t}_1 \rrbracket_{\emptyset}$ ) if and only if the coroutine is currently executing, or has completed altogether. Otherwise, if the coroutine can be resumed, the evaluation term  $\mathfrak{t}$  will not be suspended.

**Definition 2.7** [Well-typed coroutine store] A coroutine store  $\mu$  is well-typed with respect to the instance typing  $\Sigma$ , denoted  $\Sigma \vdash \mu$ , if and only if it is true that  $\forall i \in \text{dom}(\mu)$ ,  $\Sigma(i) = T_y \rightsquigarrow T_2 \Leftrightarrow \Sigma|\emptyset \vdash \mu(i) : T_2|T_y$ , and  $\text{dom}(\Sigma) = \text{dom}(\mu)$ . ◀

We next define the operational semantics of  $\lambda_{\rightsquigarrow}$ .

**Definition 2.8** [Transition relation] The *transition relation*  $\mathfrak{t}|\mu \rightarrow \mathfrak{t}'|\mu'$  is a four place relation between the source term  $\mathfrak{t}$  and source coroutine store  $\mu$  and the target term  $\mathfrak{t}'$  and the target coroutine store  $\mu'$ . The inductive definition of the transition relation is shown in Fig. 2. ◀

For simplicity, the evaluation rules are presented with the evaluation context  $E$  and suspension context  $P$ . The evaluation context  $E$  is standard, and it's used in E-CONTEXT. The suspension context  $P$  is used in E-PAUSE to simplify evaluations like the following:

$$\begin{aligned} \llbracket \mathfrak{t}_1 \rrbracket_v(\mathfrak{t}_2)|\mu &\rightarrow \llbracket \mathfrak{t}_1(\mathfrak{t}_2) \rrbracket_v|\mu \\ \mathfrak{t}_1(\llbracket \mathfrak{t}_2 \rrbracket_v)|\mu &\rightarrow \llbracket \mathfrak{t}_1(\mathfrak{t}_2) \rrbracket_v|\mu \\ \text{yield}(\llbracket \mathfrak{t} \rrbracket_v)|\mu &\rightarrow \llbracket \text{yield}(\mathfrak{t}) \rrbracket_v|\mu \end{aligned}$$

We briefly discuss the evaluation rules in Fig. 2. The rule E-APPABS is standard in lambda calculus. The only difference in our case is the addition of the coroutine store, which does not change in E-APPABS.

The rule E-START reduces a **start** expression with a coroutine and its argument to a coroutine instance with a fresh label  $i$ , and adds a binding from  $i$  to the coroutine body  $\mathfrak{t}$  in which occurrences of  $x$  are replaced with the argument  $v$ .

The rule E-YIELD reduces the **yield** expression to a *suspension* of the **Unit** value  $()$ , which has a pending yield of the value  $v$ . A term suspension is a runtime term of the form  $\llbracket \mathfrak{t} \rrbracket_v$ , where  $\mathfrak{t}$  represents the suspended computation, and  $v$  is the value that is about to be yielded. Once reduced from a **yield** expression, a suspension spreads through the program until reaching the limits of the enclosing coroutine resumption. To model this, we need to introduce evaluation rules that suspend all term shapes. The rule E-PAUSE exists for this purpose.

The expanding suspension is captured once it reaches the *coroutine resumption term*  $\langle \llbracket \mathfrak{t} \rrbracket_v, v_2, v_3, v_4 \rangle_i$ , as described by the rule E-CAPTURE. The suspended execution term  $\mathfrak{t}$  from the suspension  $\llbracket \mathfrak{t} \rrbracket_v$  is placed into the coroutine store binding of the coroutine instance  $i$ , and the yielded value  $v$  is passed to the yield-handling function  $v_3$ .

Consider the outcomes of resuming a coroutine. If a coroutine instance  $i$  is not terminated and not currently executing (that is, the binding for  $i$  in the coroutine store does not point to a suspended term  $\llbracket \mathfrak{t}_0 \rrbracket_\emptyset$ ), then a **resume** expression reduces to a coroutine resumption term, by the rule E-RESUME1. It is illegal to reduce a coroutine instance that is already completed or currently executing. A **resume** expression on such a coroutine instance reduces to an application of the fourth argument to a **Unit** term, by the rule E-RESUME2.

The rule E-TERMINATE states that if the term  $\mathfrak{t}$  reduces to a value  $v$ , the coroutine resumption reduces to an application of the second argument  $v_2$  to the reduced value  $v$ , leaving the binding for the instance  $i$  mapped to a suspended state  $\llbracket v \rrbracket_\emptyset$ .

Given an existing coroutine instance  $i_1$ , the E-SNAPSHOT rule reduces a **snapshot**( $i_1$ ) expression to a fresh coroutine instance  $i_2$ , and adds a copy of the  $i_1$ 's term  $\mathfrak{t}$  to the store as a binding for  $i_2$ . Finally, note that  $\lambda_{\dots}$  models stackful delimited coroutines, so a coroutine application must be allowed inside the body of a coroutine. This is shown in the rule E-APPCOR, which essentially describes beta reduction on coroutines.

Now that we saw how normal terms reduce to runtime terms, we can inspect the remaining typing derivations, shown in Fig. 3. A suspension must have the same type  $T$  as the term  $\mathfrak{t}$  it suspends, and the yield type  $T_y$  that corresponds to the yielded value, as stated by the

$$\begin{aligned}
E & ::= [\cdot] \mid E(t) \mid v(E) \mid \text{start}(E, t) \mid \text{start}(v, E) \mid \text{yield}(E) \mid \text{snapshot}(E) \mid \text{fix}(E) \mid \\
& \quad \text{resume}(E, t, t, t) \mid \text{resume}(v, E, t, t) \mid \text{resume}(v, v, E, t) \mid \text{resume}(v, v, v, E) \mid \\
& \quad \langle E, v, v, v \rangle_i \\
P & ::= [\cdot](t) \mid v([\cdot]) \mid \text{start}([\cdot], t) \mid \text{start}(v, [\cdot]) \mid \text{yield}([\cdot]) \mid \text{snapshot}([\cdot]) \mid \text{fix}([\cdot]) \mid \\
& \quad \text{resume}([\cdot], t, t, t) \mid \text{resume}(v, [\cdot], t, t) \mid \text{resume}(v, v, [\cdot], t) \mid \text{resume}(v, v, v, [\cdot])
\end{aligned}$$

$$\begin{aligned}
& \frac{t \mid \mu \rightarrow t' \mid \mu'}{E[t] \mid \mu \rightarrow E[t'] \mid \mu'} & \text{(E-CONTEXT)} \\
& P[[t]_v] \mid \mu \rightarrow [P[t]]_v \mid \mu & \text{(E-PAUSE)} \\
& ((x:T) \Rightarrow t)(v) \mid \mu \rightarrow [x \mapsto v]t \mid \mu & \text{(E-APPABS)} \\
& \frac{i \notin \text{dom}(\mu)}{\text{start}((x:T_1) \overset{T_y}{\rightsquigarrow} t, v) \mid \mu \rightarrow i \mid \mu, i \triangleright [x \mapsto v]t} & \text{(E-START)} \\
& \text{yield}(v) \mid \mu \rightarrow [()]_v \mid \mu & \text{(E-YIELD)} \\
& \frac{i_2 \notin \text{dom}(\mu) \quad i_1 \neq i_2}{\text{snapshot}(i_1) \mid \mu, i_1 \triangleright t \rightarrow i_2 \mid \mu, i_1 \triangleright t, i_2 \triangleright t} & \text{(E-SNAPSHOT)} \\
& ((x:T_1) \overset{T_y}{\rightsquigarrow} t_2)(v) \mid \mu \rightarrow [x \mapsto v]t_2 \mid \mu & \text{(E-APPCOR)} \\
& \text{fix}((x:T_1) \Rightarrow t_2) \mid \mu \rightarrow [x \mapsto \text{fix}((x:T_1) \Rightarrow t_2)]t_2 \mid \mu & \text{(E-FIX)} \\
& \frac{t \neq [[t_0]]_\emptyset}{\text{resume}(i, v_2, v_3, v_4) \mid \mu, i \triangleright t \rightarrow \langle t, v_2, v_3, v_4 \rangle_i \mid \mu, i \triangleright [[t]]_\emptyset} & \text{(E-RESUME1)} \\
& \text{resume}(i, v_2, v_3, v_4) \mid \mu, i \triangleright [[t_0]]_\emptyset \rightarrow v_4([\cdot]) \mid \mu, i \triangleright [[t_0]]_\emptyset & \text{(E-RESUME2)} \\
& \langle [[t_1]]_v, v_2, v_3, v_4 \rangle_i \mid \mu, i \triangleright [[t_0]]_{v'} \rightarrow v_3(v) \mid \mu, i \triangleright t_1 & \text{(E-CAPTURE)} \\
& \langle v, v_2, v_3, v_4 \rangle_i \mid \mu, i \triangleright [[t_0]]_{v'} \rightarrow v_2(v) \mid \mu, i \triangleright [[v]]_\emptyset & \text{(E-TERMINATE)}
\end{aligned}$$

■ **Figure 2** Transition relation

$$\begin{aligned}
& \frac{\Sigma \mid \Gamma \vdash t : T_y \quad \Sigma \mid \Gamma \vdash v : T_y \perp}{\Sigma \mid \Gamma \vdash [[t]]_v : T_y} & \text{(T-SUSPENSION)} \\
& \frac{\Sigma(i) = T_y \iff T_2}{\Sigma \mid \Gamma \vdash i : T_y \iff T_2 \perp} & \text{(T-INSTANCE)} \\
& \frac{\Sigma(i) = T_y \iff T_2 \quad \Sigma \mid \Gamma \vdash t_1 : T_2 \mid T_y \quad \Sigma \mid \Gamma \vdash v_2 : T_2 \overset{T_w}{\rightsquigarrow} T_R \perp}{\Sigma \mid \Gamma \vdash v_3 : T_y \overset{T_w}{\rightsquigarrow} T_R \perp \quad \Sigma \mid \Gamma \vdash v_4 : \text{Unit} \overset{T_w}{\rightsquigarrow} T_R \perp} & \text{(T-RESUMPTION)} \\
& \frac{\Sigma \mid \Gamma \vdash \langle t_1, v_2, v_3, v_4 \rangle_i : T_R \mid T_w}{\Sigma \mid \Gamma \vdash \emptyset : T \perp} & \text{(T-EMPTY)}
\end{aligned}$$

■ **Figure 3** Typing relation on runtime terms

rule T-SUSPENSION. The rule T-INSTANCE states that the type of the coroutine instance  $i$  is  $T_y \rightsquigarrow T_2$ , under the assumption that the instance typing  $\Sigma$  contains the corresponding binding for  $i$ . T-EMPTY states that one can assume that the empty term has any type. Finally, the rule T-RESUMPTION assigns a type to a coroutine resumption term, and is a direct equivalent of the rule T-RESUME from Fig. 1.

With the typing rules and the operational semantics in place, we can prove the basic safety properties of the  $\lambda_{\rightsquigarrow}$  model – progress and preservation. We start by establishing several helper lemmas.

**Lemma 2.9** [Inversion of the typing relation]

1. If  $\Sigma|\Gamma \vdash () : T|T_y$ , then  $T = \mathbf{Unit}$ .
2. If  $\Sigma|\Gamma \vdash x : T|T_y$ , then  $x : T \in \Gamma$ .
3. If  $\Sigma|\Gamma \vdash (x : T_1) \Rightarrow t_2 : T|T_y$ , then  $\exists T_2, T = T_1 \Rightarrow T_2, \Sigma|\Gamma, x : T_1 \vdash t_2 : T_2|\perp$ .
4. If  $\Sigma|\Gamma \vdash t_1(t_2) : T|T_y$ , then either  $\exists T_2$  such that  $\Sigma|\Gamma \vdash t_1 : T_1 \Rightarrow T|T_y$  and  $\Sigma|\Gamma \vdash t_2 : T_2|T_y$ , or  $\exists T_2$  such that  $\Sigma|\Gamma \vdash t_1 : T_1 \xrightarrow{T_y} T|T_y$  and  $\Sigma|\Gamma \vdash t_2 : T_2|T_y$ .
5. If  $\Sigma|\Gamma \vdash (x : T_1) \xrightarrow{T_y} t_2 : T|T_w$ , then  $\exists T_2, T = T_1 \xrightarrow{T_y} T_2, \Sigma|\Gamma, x : T_1 \vdash t_2 : T_2|T_y$ .
6. If  $\Sigma|\Gamma \vdash \mathbf{start}(t_1, t_2) : T|T_w$ , then  $\exists T_1, T_2, T_y$  such that  $T = T_y \rightsquigarrow T_2, \Sigma|\Gamma \vdash t_1 : T_1 \xrightarrow{T_y} T_2|T_w$ , and  $\Sigma|\Gamma \vdash t_2 : T_2|T_w$ .
7. If  $\Sigma|\Gamma \vdash \mathbf{yield}(t) : T|T_y$ , then  $T = \mathbf{Unit}$  and  $\Sigma|\Gamma \vdash t : T_y|T_y$ .
8. If  $\Sigma|\Gamma \vdash \mathbf{snapshot}(t) : T|T_w$ , then  $\exists T_2, T_y$  such that  $T = T_y \rightsquigarrow T_2$ , and  $\Sigma|\Gamma \vdash t : T_y \rightsquigarrow T_2|T_y$ .
9. If  $\Sigma|\Gamma \vdash \mathbf{resume}(t_1, t_2, t_3, t_4) : T|T_w$ , then  $\exists T_2, T_y$  such that the following holds:  $\Sigma|\Gamma \vdash t_1 : T_y \rightsquigarrow T_2|T_w$ , and  $\Sigma|\Gamma \vdash t_2 : T_2 \xrightarrow{T_w} T_R|T_w$ , and  $\Sigma|\Gamma \vdash t_3 : T_y \xrightarrow{T_w} T_R|T_w$ , and  $\Sigma|\Gamma \vdash t_4 : \mathbf{Unit} \xrightarrow{T_w} T_R|T_w$ .
10. If  $\Sigma|\Gamma \vdash \mathbf{fix}(t) : T|T_w$ , then  $\Sigma|\Gamma \vdash t : T \Rightarrow T|\perp$ .
11. If  $\Sigma|\Gamma \vdash \llbracket t \rrbracket_v : T|T_y$ , then  $\Sigma|\Gamma \vdash t : T|T_y$ , and  $\Sigma|\Gamma \vdash v : T_y|\perp$ .
12. If  $\Sigma|\Gamma \vdash i : T|T_w$ , then  $\exists T_y, T_2$  such that  $i : T_y \rightsquigarrow T_2 \in \Sigma$ .
13. If  $\Sigma|\Gamma \vdash \langle t_1, v_2, v_3, v_4 \rangle_i : T_R|T_w$ , then  $\exists T_2, T_y$  such that the following holds:  $\Sigma|\Gamma \vdash t_1 : T_y \rightsquigarrow T_2|T_w$ ,  $\Sigma|\Gamma \vdash v_2 : T_2 \xrightarrow{T_w} T_R|T_w$ ,  $\Sigma|\Gamma \vdash v_3 : T_y \xrightarrow{T_w} T_R|T_w$ , and  $\Sigma|\Gamma \vdash v_4 : \mathbf{Unit} \xrightarrow{T_w} T_R|T_w$ .

**Proof.** Follows immediately from the typing derivations in Definition 2.4. ◀

**Lemma 2.10** [Canonical forms]

1. If  $v$  is a value of type  $\mathbf{Unit}$ , then  $v$  is  $()$ .
2. If  $v$  is a value of type  $T_1 \Rightarrow T_2$ , then  $v = (x : T_1) \Rightarrow t_2$ .
3. If  $v$  is a value of type  $T_1 \xrightarrow{T_y} T_2$ , then  $v = (x : T_1) \xrightarrow{T_y} t_2$ .
4. If  $v$  is a value of type  $T_y \rightsquigarrow T_2$ , then  $v = i$ , where  $i$  is an instance label.

**Proof.** We consider the possible forms of values as per syntax from Definition 2.1, and rely on Lemma 2.9 to prove the claim.

For example, the value  $()$  immediately satisfies the claim, by T-UNIT. From Lemma 2.9, we see that other types of values (functions, coroutines and coroutine instances) never have the type  $\mathbf{Unit}$ . The remaining cases are proved in a similar way. ◀

We can now state the progress property of  $\lambda_{\rightsquigarrow}$ .



**Theorem 2.11** [Progress] Suppose that  $\mathfrak{t}$  is a closed, well-typed term for some  $T$  and  $\Sigma$ , as defined in Definition 2.5. Then, either  $\mathfrak{t}$  is a value, or  $\mathfrak{t}$  is a suspension term  $\llbracket \mathfrak{t} \rrbracket_v$ , or, for any store  $\mu$  such that  $\Sigma \vdash \mu$ , there is some term  $\mathfrak{t}'$  and store  $\mu'$  such that  $\mathfrak{t}|\mu \rightarrow \mathfrak{t}'|\mu'$ .

**Proof.** Since  $\mathfrak{t}$  is well-typed, we proceed casewise on the typing derivations. Cases T-ABS, T-UNIT, T-COROUTINE and T-INSTANCE follow directly, since  $\mathfrak{t}$  is a value. Cases T-VAR, T-CTX, T-SUSPENSION and T-EMPTY are trivial.

For the sake of simplicity, in most of the following cases when we use the induction hypothesis, we only consider the case where a subterm is a value, ignoring the case where the subterm is a suspension ( $\llbracket \mathfrak{t} \rrbracket_v$ ), and the case where the subterm can take a step. This is valid because of the rule E-PAUSE and E-CONTEXT, which can be used to make  $t$  take one step when the subterm is not a value.

*Case T-APP:  $\mathfrak{t} = \mathfrak{t}_1(\mathfrak{t}_2)$*

We only consider the case where both  $\mathfrak{t}_1$  and  $\mathfrak{t}_2$  are values, then Lemma 2.10 tells us that  $\mathfrak{t}_1 = (x:T_2) \Rightarrow \mathfrak{t}_{11}$ , so E-APPABS applies to  $\mathfrak{t}$ .

*Case T-START:  $\mathfrak{t} = \text{start}(\mathfrak{t}_1, \mathfrak{t}_2)$*

Similar to T-APP, but use E-START.

*Case T-YIELD:  $\mathfrak{t} = \text{yield}(\mathfrak{t}_1)$*

Similar to T-APP, but we rely on on E-YIELD.

*Case T-SNAPSHOT:  $\mathfrak{t} = \text{snapshot}(\mathfrak{t}_1)$*

Similar to T-YIELD, but use E-SNAPSHOT.

*Case T-RESUME:  $\mathfrak{t} = \text{resume}(\mathfrak{t}_1, \mathfrak{t}_2, \mathfrak{t}_3, \mathfrak{t}_4)$*

Only consider the case where all subterms are values. By Lemma 2.10,  $\mathfrak{t}_1$  must be an instance label  $i$ . In that case, there exists a store  $\mu$  that contains a binding for  $i$ , such that either E-RESUME1 or E-RESUME2 applies.

*Case T-APPCOR:  $\mathfrak{t} = \mathfrak{t}_1(\mathfrak{t}_2)$*

Similar to T-APP, but use E-APPCOR.

*Case T-FIX:  $\mathfrak{t} = \text{fix}(\mathfrak{t}_1)$*

Similar to T-APP, but use E-FIX.

*Case T-RESUMPTION:  $\mathfrak{t} = \langle \mathfrak{t}_1, v_2, v_3, v_4 \rangle_i$*

If  $\mathfrak{t}_1$  is a value, then E-TERMINATE applies. If  $\mathfrak{t}_1$  is a suspension, then E-CAPTURE applies. Otherwise,  $\mathfrak{t}_1$  reduces by the induction hypothesis, so  $\mathfrak{t}$  reduces by E-CONTEXT. ◀

Before we prove that types are preserved during evaluation, we state several standard helper lemmas.

**Lemma 2.12** [Permutation] Assume that  $\Sigma|\Gamma \vdash \mathfrak{t}:T|T_y$ , and that  $\Sigma'$  and  $\Gamma'$  are permutations of  $\Sigma$  and  $\Gamma$ , respectively. Then  $\Sigma'|\Gamma' \vdash \mathfrak{t}:T|T_y$ .

**Proof.** The proof is a straightforward induction on the typing rules. ◀

**Lemma 2.13** [Weakening] Assume that  $\Sigma|\Gamma \vdash \mathfrak{t}:T|T_y$ . Then it holds that for any  $\Sigma' \supseteq \Sigma, \Gamma' \supseteq \Gamma$ ,  $\Sigma'|\Gamma' \vdash \mathfrak{t}:T|T_y$ .

**Proof.** The proof is a straightforward induction on the typing rules. ◀

**Lemma 2.14** [Substitution] If  $\Sigma|\Gamma, x:S \vdash \mathfrak{t}:T|T_y$ , and  $\Sigma|\Gamma \vdash s:S|\perp$ , then  $\Sigma|\Gamma \vdash [x \mapsto s]\mathfrak{t}:T|T_y$ .

## 23:10 On the Soundness of Coroutines with Snapshots

**Proof.** The proof is a straightforward induction on the typing rules. The more interesting cases are T-ABS, T-VAR and T-COROUTINE, and they rely on Lemma 2.12 and Lemma 2.13. We show the case for T-COROUTINE, where  $\mathfrak{t} = (y:T_2) \overset{T_y}{\rightsquigarrow} \mathfrak{t}_1$ ,  $T = T_2 \overset{T_y}{\rightsquigarrow} T_1$ , and  $\Sigma|\Gamma, \mathbf{x}:\mathbf{S}, y:T_2 \vdash \mathfrak{t}_1:T_1|T_y$ . By applying permutation, we obtain  $\Sigma|\Gamma, y:T_2, \mathbf{x}:\mathbf{S} \vdash \mathfrak{t}_1:T_1|T_y$ . By applying weakening, we obtain  $\Sigma|\Gamma, y:T_2 \vdash \mathfrak{s}:\mathbf{S}|T_y$ . We use the last two results with the induction hypothesis to obtain  $\Sigma|\Gamma, y:T_2 \vdash [\mathbf{x} \mapsto \mathfrak{s}]\mathfrak{t}_1:T_1|T_y$ . Finally, from T-COROUTINE, we get  $\Sigma|\Gamma, y:T_2 \vdash (y:T_2) \overset{T_y}{\rightsquigarrow} [\mathbf{x} \mapsto \mathfrak{s}]\mathfrak{t}_1:T_1|\perp$ . Since  $\mathbf{x} \neq y$  (we can rename coroutine variables as needed), the result follows.  $\blacktriangleleft$

**Lemma 2.15** [Suspension] If  $\Sigma|\Gamma \vdash P[[\mathfrak{t}]]_v :T|T_y$ , then If  $\Sigma|\Gamma \vdash [P[\mathfrak{t}]]_v :T|T_y$ .

**Proof.** By cases analysis on the suspension context  $P$ . We only show the case for application, other cases are similar.

*Case  $P = [\cdot](\mathfrak{t}_1)$ :*

From the typing rule T-APP, we have  $\Sigma|\Gamma \vdash [[\mathfrak{t}]]_v :T_1 \Rightarrow T|T_y$  and  $\Sigma|\Gamma \vdash \mathfrak{t}_1:T_1|T_y$ . From the typing rule T-SUSPENSION, we have  $\Sigma|\Gamma \vdash v:T_y|\perp$  and  $\Sigma|\Gamma \vdash \mathfrak{t}:T_1 \Rightarrow T|T_y$ . Now, it's easy to apply the typing rules to show that  $\Sigma|\Gamma \vdash [[\mathfrak{t}(\mathfrak{t}_1)]]_v :T|T_y$ .  $\blacktriangleleft$

**Theorem 2.16** [Preservation] If a term and the coroutine store are well-typed, that is,  $\Sigma|\Gamma \vdash \mathfrak{t}:T|T_y$ , and  $\Sigma|\Gamma \vdash \mu$ , and if  $\mathfrak{t}|\mu \rightarrow \mathfrak{t}'|\mu'$ , then there exists  $\Sigma' \supseteq \Sigma$  such that  $\Sigma'|\Gamma \vdash \mathfrak{t}':T|T_y$  and  $\Sigma'|\Gamma \vdash \mu'$ .

**Proof.** We prove this by the induction on the typing derivations. Cases T-UNIT, T-ABS, T-COROUTINE, T-EMPTY and T-INSTANCE are straightforward, since  $\mathfrak{t}$  is a value and does not reduce.

Note that we only consider reduction rules, ignoring the rule E-CONTEXT and E-PAUSE. This is valid because induction hypothesis makes the case E-CONTEXT trivial, and Lemma 2.15 makes the case E-PAUSE trivial.

*Case T-APP:  $\mathfrak{t} = \mathfrak{t}_1(\mathfrak{t}_2)$*

Consider the evaluation rule E-APPABS. Both  $\mathfrak{t}_1$  and  $\mathfrak{t}_2$  are values, their yield type is  $T_y = \perp$ . Moreover, then  $\mathfrak{t}_1$  must have the form  $(\mathbf{x}:T_2) \Rightarrow \mathfrak{t}_{11}$ , so by E-APPABS  $\mathfrak{t}' = [\mathbf{x} \mapsto \mathfrak{t}_2]\mathfrak{t}_{11}$ . From Lemma 2.9, we know that  $\Sigma|\Gamma, \mathbf{x}:T_2 \vdash \mathfrak{t}_{11}:T_1|\perp$ . The claim about  $\mathfrak{t}'$  follows from Lemma 2.14, and  $\mu' = \mu$ .

*Case T-START:  $\mathfrak{t} = \mathbf{start}(\mathfrak{t}_1, \mathfrak{t}_2)$*

By the rule E-START,  $\mathfrak{t}_1$  and  $\mathfrak{t}_2$  are values, and  $\mathfrak{t}' = i$  such that  $i \notin \text{dom}(\mu)$ . Since  $\mu$  is by assumption well-typed, it follows that  $i \notin \text{dom}(\Sigma)$ . But there exists  $\Sigma' = \Sigma, i:T_y \rightsquigarrow T_2 \supseteq \Sigma$  such that  $\Sigma'|\Gamma \vdash i:T_y \rightsquigarrow T_2|T_w$  and  $\Sigma' \vdash \mu'$ .

*Case T-YIELD:  $\mathfrak{t} = \mathbf{yield}(\mathfrak{t}_1)$*

By the rule E-YIELD,  $\mathfrak{t}_1$  is a value, and  $\mathfrak{t}' = [()]_v$ . From T-YIELD, we know that  $\Sigma|\Gamma \vdash \mathfrak{t}_1:T_y|T_y$ . From T-SUSPENSION, we know that  $\Sigma|\Gamma \vdash \mathfrak{t}':\mathbf{Unit}|T_y$ . We also know that  $\mu' = \mu$ , which proves the claim.

*Case T-SNAPSHOT:  $\mathfrak{t} = \mathbf{snapshot}(\mathfrak{t}_1)$*

This case is similar to T-YIELD, but we rely on E-SNAPSHOT for the transition, and on T-INSTANCE to type the resulting term  $\mathfrak{t}' = i$ .

*Case T-RESUME:  $\mathfrak{t} = \mathbf{resume}(\mathfrak{t}_1, \mathfrak{t}_2, \mathfrak{t}_3, \mathfrak{t}_4)$*

This case is similar to T-YIELD, but we distinguish two cases – that the value  $\mathfrak{t}_1 = i$  is a terminated coroutine instance, in which case we rely on E-RESUME2 for the transition,

and that  $i$  is not terminated, in which case we rely on E-RESUME1 for the transition. We furthermore rely on T-APP and T-UNIT to prove the typing relation on  $\tau'$  in the former case, and on T-RESUMPTION in the latter. In both of these cases, we rely on T-EMPTY, T-SUSPENSION and T-RESUMPTION to establish that  $\Sigma \vdash \mu'$ .

*Case T-APPCOR:*  $\tau = \tau_1(\tau_2)$

By E-APPCOR, we know  $\tau_1 = (x:T_2) \xrightarrow{T_y} \tau_{11}$ ,  $\tau' = [x \mapsto \tau_2]\tau_{11}$ , and  $\mu' = \mu$ , so the result follows from Lemma 2.14.

*Case T-FIX:*  $\tau = \text{fix}(\tau_1)$

Trivial by E-FIX and Lemma 2.14.

*Case T-RESUMPTION:*  $\tau = \langle \tau_1, v_2, v_3, v_4 \rangle_i$

There are two subcases. (1) The reduction rule is E-TERMINATE. we have  $\tau' = \tau_2(\tau_1)$ . The claim then follows from T-APP, T-EMPTY and T-RESUMPTION. (2) The reduction rule is E-CAPTURE. The claim similarly follows.

*Case T-SUSPENSION:*  $\tau = \llbracket \tau_1 \rrbracket_v$

Trivial, as a suspension cannot take a step. ◀

We are now ready to state another safety property that follows directly from the preservation theorem. We want to show that if the program was typed such that the yield type is  $\perp$ , then the program will not yield a value outside of a coroutine resumption.

**Corollary 2.17** [Yield safety] A well-typed user program  $\tau_u$  never evaluates to a suspension term  $\llbracket \tau \rrbracket_v$ .

**Proof.** Assume that the program evaluates to the suspension term  $\llbracket \tau \rrbracket_v$ . Then, by T-SUSPENSION, it must be that  $\Sigma|\emptyset \vdash \llbracket \tau \rrbracket_v : T|T_y$ . Note that  $T_y \neq \perp$ , since there is no non-empty term  $v$  whose type is  $\perp$ , and the empty term cannot appear in the evaluation. Furthermore, by Theorem 2.16, the original program  $\tau_u$  must have the same yield type  $T_y \neq \perp$ . This is a contradiction, since such a program would not be a well-typed user program, by Definition 2.5. ◀

### 3 Formal Transformation of $\lambda_{\rightsquigarrow}$

The formal transformation translates  $\lambda_{\rightsquigarrow}$  programs to programs in a simply typed lambda calculus extended with references and restricted sums. Intuitively, references are necessary because coroutine instances are stateful entities – we use references to store the evaluation state. The key idea is to translate coroutines into functions that accept a *store function* as an argument, as explained shortly. We begin by introducing the target language.

**Definition 3.1** [Target language] The syntax of the target language is as follows:

$\tau ::=$	terms:	$T ::=$	types:
$(x:T) \Rightarrow \tau$	abstraction	$\text{Ret}(\tau)$	tagged return
$\tau(\tau)$	application	$\text{Yield}(\tau)$	tagged yield
$x$	variable	$\text{Term}$	tagged ended
$()$	unit value	$\text{t match } \{$	pattern match
$\text{ref}(\tau)$	new reference	$\text{case Ret}(x) \Rightarrow \tau;$	
$!\tau$	dereferencing	$\text{case Yield}(x) \Rightarrow \tau;$	
$\tau := \tau$	assignment	$\text{case Term} \Rightarrow \tau \}$	
		$\text{Unit}$	function
		$\text{Ref}[T]$	unit
		$\text{Out}[T, T]$	reference
			output

◀

Informally, references of type  $\text{Ref}[T]$  are created with the `ref` expression, assigned with `:=` and dereferenced with `!`. Output values of type  $\text{Out}[T_y, T_r]$  describe the result of resuming

a coroutine – either a `Yield(x)`, where  $x$  has the type  $T_y$ , indicating a yield; or a `Ret(x)`, where  $x$  has the type  $T_r$ , indicating a normal return; or a `Term`, indicating a termination. A pattern match reduces to the term in the respective case. We do not show the exact typing rules and the operational semantics for the target language, since this was already treated in-depth [17, 27]. We also skip runtime terms, as they are not used in the transformation.

**Translation approach.** The transformation applies only to terms that are lexically enclosed by a coroutine definition. These terms are transformed into a continuation-passing style (CPS) – the result of evaluating every term gets passed to a function that represents the remainder of the enclosing coroutine (*not the entire program*). For a term of type  $T$ , this function takes the value of type  $T$ , and returns either a `Yield` value or a `Ret` value.

**Definition 3.2** The *continuation type* is defined as  $\varkappa[T, T_y, T_r] \triangleq T \Rightarrow \text{Out}[T_y, T_r]$ .

Since the evaluation rules in Figure 2 relied on the instance store  $\mu$ , a coroutine instance must translate into a stateful entity. Concretely, a coroutine instance will become a reference that stores the continuation of the coroutine’s execution, typed  $\text{Ref}[\text{Unit} \Rightarrow \text{Out}[T_y, T_r]]$ . To yield is to modify this reference. To resume is to read it and run the continuation.

**Definition 3.3** The *evaluation state type* is defined as  $\rho[T_y, T_r] \triangleq \text{Ref}[\varkappa[\text{Unit}, T_y, T_r]]$ .

A coroutine definition translates into a lambda that takes two arguments. One argument must obviously correspond to the coroutine’s argument. The other argument must encode the runtime state of the coroutine instance. At first glance, it is tempting to model this state with the coroutine instance reference, which would make the coroutine type:

$$\text{Ref}[\varkappa[\text{Unit}, T_y, T_r]] \Rightarrow T \Rightarrow \text{Out}[T_y, T_r]$$

However, such a type would not allow modeling stackful coroutines. Recall that a coroutine can be either started with `start`, or called by another coroutine. In the latter case, a `yield` inside a coroutine must provide a continuation that captures not only the current coroutine, but also the caller coroutine (i.e. yielding captures the entire call stack). But, lexically speaking, a coroutine has no way of knowing what the continuation of its caller is. A coroutine can only create a continuation for its own scope, and pass that continuation fragment to its caller. The caller can then recursively extend the continuation with its own fragment. Once the bottom of the call stack is reached, the continuation is stored into the reference. Therefore, we need to abstract this with a separate *store function*.

**Definition 3.4** The *store function type* (i.e. a function that stores the continuation), is defined as  $\sigma[T_y, T_r] \triangleq \varkappa[\text{Unit}, T_y, T_r] \Rightarrow \text{Unit}$ . The *coroutine definition type* (i.e. the coroutine’s equivalent after transformation) is  $\gamma[T_1, T_y, T_r] \triangleq \sigma[T_y, T_r] \Rightarrow T_1 \Rightarrow \text{Out}[T_y, T_r]$ .

**Definition 3.5** The terms  $t_1; t_2$ ,  $\text{val } x:T=t_1; t_2$  and  $() \Rightarrow t$  are syntactic sugar:  $t_1; t_2 \triangleq ((u:\text{Unit}) \Rightarrow t_2)(t_1)$ , and  $\text{val } x:T=t_1; t_2 \triangleq ((x:T) \Rightarrow t_2)(t_1)$ , and  $() \Rightarrow t \triangleq (u:\text{Unit}) \Rightarrow t$ .

The translation scheme is as follows. A `start` term becomes the creation of a reference  $\rho[T_y, T_r]$ . A `resume` becomes a read from this reference, followed by a call to a continuation lambda. A `yield` calls the store function to store the current continuation lambda. Finally, when a coroutine calls another coroutine, a new store function is created, which takes the continuation of the callee, and chains it with the continuation of the callsite.

**Example.** Consider the `dup` coroutine, which simply returns the sum of its arguments, and is translated as follows:

$$(x:\text{Int}) \overset{\perp}{\rightsquigarrow} x+x \quad \rightarrow \quad (s:\sigma[\perp, \text{Int}]) \Rightarrow (x:\text{Int}) \Rightarrow s(() \Rightarrow \text{Term}); \text{Ret}(x+x)$$

The body of each coroutine undergoes a variant of the CPS transform [46, 43]. Our transform is particular in the sense that each term translates not to function that consumes a continuation, but to a function that takes a store function and a continuation. In the example above, the store function is invoked immediately before returning from the coroutine, in order to update the instance state.

**Example.** The following coroutine yields the argument once, and is translated as follows:

$$\begin{aligned} (x:\text{Int}) \stackrel{\text{Int}}{\rightsquigarrow} \text{yield}(x) &\rightarrow (s:\sigma[\text{Int}, \text{Unit}]) \Rightarrow \text{Int} \Rightarrow (y')(s)(q) \\ y' = (s:\sigma[\text{Int}, \text{Unit}]) \Rightarrow (k:\mathcal{K}[\text{Unit}, \text{Int}, \text{Unit}]) \Rightarrow s(k); \text{Yield}(x) \\ q = (y:\text{Unit}) \Rightarrow s(() \Rightarrow \text{Term}); \text{Ret}(y) \end{aligned}$$

The coroutine body is in continuation-passing style. To see that this transformation is correct, imagine that we passed the store function and an integer to the translated coroutine. By evolving the term, we would eventually pass the continuation to the store function  $s$ , and arrive at the `Yield` value. We define the precise transformation relation on terms next.

**Definition 3.6** The *type substitution*  $\tau$  replaces the types of a term as follows:

$$\tau(\mathbf{t}) \triangleq \forall T_1, T_y, T_2 [(x:T_1) \stackrel{T_y}{\rightsquigarrow} T_2 \mapsto \gamma[T_1, T_y, T_2]] [T_y \rightsquigarrow T_2 \mapsto \rho[T_y, T_2]] \mathbf{t}$$

The notation  $\xi$  is an abbreviation used to express transformed terms, and is defined as:

$$\xi(T_y, T_r, T, \mathbf{t}) \triangleq \tau((s:\sigma[T_y, T_r]) \Rightarrow (k:\mathcal{K}[T, T_y, T_r]) \Rightarrow \mathbf{t})$$

**Definition 3.7** [Transformation relation] The *transformation relation* is a five place relation  $\Gamma \vdash T_y | T_r | \mathbf{t} \rightarrow \mathbf{t}'$ , where  $\Gamma$  is the typing context,  $T_y$  is the yield type of the current term,  $T_r$  is the return type of the enclosing coroutine,  $\mathbf{t}$  is the term in the source language, and  $\mathbf{t}'$  is the term in the target language. This relation is inductively defined according to Figures 4 and 5.  $\blacktriangleleft$

The rules fall into two groups. The first is CPS-based and transforms terms inside coroutines (Fig. 4). The second group transforms terms outside of coroutines (Fig. 5). All rules in the first group assume that we are inside a body of some coroutine, which has the yield type  $T_y$  and a return type  $T_r$  (outside of a body of a coroutine, as we show shortly, there is no need for a CPS transform).

Consider the X-UNIT rule, which transforms `Unit` constants *inside* coroutines. A unit constant becomes a  $\xi$ -function that takes a store function of type  $\sigma[T_y, T_r]$ , where  $T_y$  and  $T_r$  are the yield and return type of the surrounding context. However, the transformed `Unit` constant does not yield, and hence does not need to invoke the store function. Instead, it just invokes the current continuation  $k$ , passing it the unit value.

On the other hand, the result of the X-YIELD rule must use the store function  $s$  to store the continuation before yielding. Given a term  $\mathbf{t}$  that is recursively translated to  $\mathbf{t}'$ , the `yield( $\mathbf{t}$ )` term is translated to a  $\xi$ -function that first evaluates  $\mathbf{t}'$  by passing it the store

function and the continuation. The continuation body stores the remainder of the continuation with  $s$ , and it itself reduces to `Yield( $x$ )`, where  $x$  is the result of evaluating  $\mathbf{t}'$ .

The X-VAR rule simply calls the continuation by passing the identifier  $x$  to the continuation. The X-APP rule assumes that the subterm  $\mathbf{t}_1$  evaluates to a function  $\mathbf{t}'_1$ , and  $\mathbf{t}_2$  evaluates to  $\mathbf{t}'_2$ . Consider what happens after both  $\mathbf{t}'_1$  and  $\mathbf{t}'_2$  produce values  $x_1$  and  $x_2$ , respectively. The next step in the evaluation is to simply apply  $x_1$  to  $x_2$ . However, the application  $x_1(x_2)$

$$\begin{array}{c} T_y \neq \perp \quad \Gamma \vdash \mathbf{t}:T \quad \Gamma \vdash T | T_r | \mathbf{t} \rightarrow \mathbf{t}' \\ \hline p = \mathbf{t}'(s)((x:T) \Rightarrow s(k); \text{Yield}(x)) \\ \Gamma \vdash T | T_r | \text{yield}(\mathbf{t}) \rightarrow \xi(T, T_r, \text{Unit}, p) \end{array} \quad (\text{X-YIELD})$$

is a continuation for the term  $\tau'_2$ , whose evaluation is itself a continuation for  $\tau'_1$ . The transformation of  $\tau_1(\tau_2)$  requires chaining these continuations together as described in the X-APP rule.

The X-COROUTINE rule relies on the assumption that the coroutine body  $\tau$  translates to a function  $\tau'$  under the extended typing context  $\Gamma, x:T_1$ . The transformed term has the  $\gamma$  type – it takes a store function and an argument. The store function is passed to  $\tau'$ , along with a continuation that stores a terminated continuation  $()\Rightarrow\text{Term}$ , and wraps the result of  $\tau'$  into a **Ret** value. The resulting  $\gamma$  value is then passed to the continuation of the definition site of the coroutine (which is itself assumed to be inside another coroutine).

A coroutine created this way is, by rule X-START, *called* by first allocating a reference  $x$  for the evaluation state (term  $q_1$ ), then assigning the continuation of the coroutine into the reference  $x$  (term  $q_2$ ), and finally, passing the reference  $x$  to the continuation  $k$  (term  $p$ ). Note that the abbreviation  $\psi$  is used to construct the store function – this is just a store function that assigns the continuation to the reference.

**Definition 3.8** The *output transformer*  $\Phi$  is a function that, for a given continuation value  $k$  of type  $\varkappa[T_r, T_y, T_q]$ , maps a value of type  $\text{Out}[T_y, T_r]$  to a value of type  $\text{Out}[T_y, T_q]$ , and is defined as follows:

$$\Phi(T_y, T_r, T_q, k) \triangleq \\ (x:\text{Out}[T_y, T_r])\Rightarrow x \text{ match}\{\text{case Ret}(x)\Rightarrow k(x); \text{case Yield}(x)\Rightarrow \text{Yield}(x); \text{case Term}\Rightarrow \text{Term}\}$$

The *store function constructor*  $\psi$  is a function that maps a reference  $x$  of type  $\rho[T_y, T_r]$  to a store function of type  $\sigma[T_y, T_r]$ , and is defined as follows:

$$\psi(T_y, T_r, x) \triangleq (x:\rho[T_y, T_r])\Rightarrow (k:\varkappa[\text{Unit}, T_y, T_r])\Rightarrow x:=k$$

◀

We are now ready to take a look at the X-APPCOR rule. A transformed coroutine application first constructs a mapping  $f$  between output type of the callee coroutine and the output type of the caller coroutine (term  $q_1$ ), whose type is defined as follows:  $\varphi[T_y, T_r, T_q] \triangleq \text{Out}[T_y, T_r]\Rightarrow\text{Out}[T_y, T_q]$ . This mapping is defined by the abbreviation  $\Phi$  (defined in Definition 3.8), which just forwards **Yield** values. However, if the callee returns a **Ret**( $x$ ) value, then the wrapped value  $x$  is passed to the continuation  $k$  of the callee (which itself must return the correctly typed **Ret** value).

Next, the coroutine application must create a new store function  $s'$  with the appropriate type (term  $q_2$ ). This new store function  $s'$  passes a modified continuation to the store function  $s$  of the callee, such that the modified continuation calls the callee's continuation  $k'$ , and then adapts the result using the mapping  $f$ . Finally, the transformed coroutine  $x_1$  is invoked with the new store function  $s'$ , and the result transformed using  $f$  (term  $q_3$ ).

Having seen X-APPCOR, rules X-SNAPSHOT and X-RESUME should be self-explanatory. The **snapshot** term translates to the cloning the  $\rho$  reference, and the **resume** term translates to dereferencing the continuation from the  $\rho$  reference, invoking it, and calling the correct handler, case-wise.

We deliberately left X-ABS as the last rule in our discussion. Recall that the standard abstraction term effectively *resets* the yield type back to  $\perp$  – the standard abstraction term *cannot* yield back to its callsite. For this reason, the X-ABS rule transforms the body of an abstraction term under the assumption that the enclosing yield and return types are both  $\perp$ . However, the fingerprint present in each transformation rule is that the  $T_y$  *must not* be  $\perp$  – the transformation rules shown so far cannot translate such a term!

The conclusion is that we need a second set of rules that govern the transformation *outside* of coroutines. Naturally, these rules do not require a CPS transform, and they leave most of

the program as is – they deal with coroutine definitions and handling coroutine instances. In particular, there is no analogue for the X-APPCOR and X-YIELD rules – a well-typed user program cannot yield or call coroutines outside of a lexical scope of a coroutine.

The rules in Fig. 5, whose name starts with X-FREE, transform non-yielding terms. They are analogous to the coroutine transformation rules. When the yield and return types are  $\perp$ , the transformation produces a normal term instead of a  $\xi$  expression – the logic inside each translated coroutine-related operation stays the same, but the value is not being passed to a continuation.

To switch to the CPS transform from Fig. 4, the transformation must apply the X-FREECOROUTINE rule, which transforms a coroutine definition.

$$\frac{\mathbf{v} = (\mathbf{x} : \mathbf{T}_1) \overset{\mathbf{T}_y}{\rightsquigarrow} \mathbf{t} \quad \Sigma | \Gamma \vdash \mathbf{v} : \mathbf{T}_1 \overset{\mathbf{T}_y}{\rightsquigarrow} \mathbf{T}_2 | \perp \quad \Gamma, \mathbf{x} : \mathbf{T}_1 \vdash \mathbf{T}_y | \mathbf{T}_2 | \mathbf{t} \rightarrow \mathbf{t}' \quad \mathbf{T}_y \neq \perp}{\Gamma \vdash \perp | \perp | \mathbf{v} \rightarrow \tau((\mathbf{s} : \sigma[\mathbf{T}_y, \mathbf{T}_2]) \Rightarrow (\mathbf{x} : \mathbf{T}_1) \Rightarrow \mathbf{t}'(\mathbf{s})((\mathbf{x}_2 : \mathbf{T}_2) \Rightarrow \mathbf{s}(\perp) \Rightarrow \mathbf{Term}); \mathbf{Ret}(\mathbf{x}_2)))} \text{(X-FREECOROUTINE)}$$

We read this as follows. Consider a coroutine definition  $\mathbf{v}$  whose body is the term  $\mathbf{t}$ , and whose type is  $\mathbf{T}_1 \overset{\mathbf{T}_y}{\rightsquigarrow} \mathbf{T}_2$ . If the term  $\mathbf{t}$  translates to a term  $\mathbf{t}'$ , then the coroutine definition translates to a function that takes a store function  $\mathbf{s}$  and the argument  $\mathbf{x}$ , and then invokes the transformed body  $\mathbf{t}'$  with the store function and the continuation. This continuation stores the **Term**-returning function (i.e. terminates), and returns the result **Ret**( $\mathbf{x}$ ). From this, we can conclude that the transformed body  $\mathbf{t}'$  must be a  $\xi$ -function from Definition 3.6.

Note, finally, that the only terms that are different after the transformation are those that mention elements of the  $\lambda_{\rightsquigarrow}$  calculus syntax. Any term that makes no mention of  $\lambda_{\rightsquigarrow}$ -specific terms is left unchanged. This shows that the transformation is both selective, and that it does not require the recompilation of legacy programs that use coroutines.

## 4 Related Work

We organize the related work on coroutines into two categories. We start with the origins and previous formalization approaches, and then discuss related work on continuations. We do not discuss related concepts like iterators or generators, and interested readers can refer to our main work [38]. Whereas coroutines have been studied extensively before, the main novelty in our work is to augment coroutines with snapshots. This allows for several new useful use-cases.

**Origins and formalizations.** The idea of coroutines dates back to Erdwinn and Conway’s work on a tape-based Cobol compiler and its separability into modules [9]. Although the original use-case is no longer relevant, other use-cases emerged. Coroutines were investigated on numerous occasions, and initially appeared in languages such as Modula-2 [54], Simula [6], and BCPL [24]. A detailed classification of coroutines is given by Moura and Ierusalimsky [25], along with a formalization of asymmetric coroutines through an operational semantics. Moura and Ierusalimsky observed that asymmetric first-class stackful coroutines have an equal expressive power as one-shot continuations, but did not investigate snapshots, which make coroutines equivalent to full continuations. Anton and Thiemann showed that it is possible to automatically derive type systems for symmetric and asymmetric coroutines by converting their reduction semantics into equivalent functional implementations, and then applying existing type systems for programs with continuations [2]. James and Sabry identified the input and output types of coroutines [45], where the output type corresponds to the *yield* type described in this paper. The input type ascribes the value passed to the coroutine when it is resumed. As a design tradeoff, we chose not to have explicit input values

in our model. First, the input type increases the verbosity of the coroutine type, which may have practical consequences. Second, as shown in [38], the input type can be simulated with the return type of another coroutine, which yields a writable location, and returns its value when resumed. Fischer et al. proposed a coroutine-based programming model for the Java programming language, along with the respective formal extension of Featherweight Java [12].

**Transformation-based continuations.** Continuations are closely related to coroutines, and with the addition of `snapshot` the two can express the same programs. Scheme supports programming with continuations via the `call/cc` operator, which has a similar role as `shift` in shift-reset delimited continuations [10, 3]. In several different contexts, it was shown that continuations subsume other control constructs such as exception handling, backtracking, and coroutines. Nonetheless, most programming languages do not support continuations. It is somewhat difficult to provide an efficient implementation of continuations, since the captured continuations must be callable more than once. One approach is to transform the program to continuation-passing style [51]. Scala’s continuations [44] implement delimited shift-reset continuations with a CPS transform. The downside of CPS is the risk of runtime stack overflows in the absence tail-call optimization, as is the case of JVM.

Optimizing compilers tend to be tailored to the workloads that appear in practice. For example, it was shown that optimizations such as inlining, escape analysis, loop unrolling and devirtualization make most collection programs run nearly optimally [37, 28, 42, 50, 35]. However, abstraction overheads associated with coroutines are somewhat new, and are not addressed by most compilers. For this reason, compile-time transformations of coroutine-heavy workloads typically produce slower programs compared to their runtime-based counterparts. We postulate that targeted high-level JIT optimizations could significantly narrow this gap.

**Runtime-based continuations.** There were several attempts to provide runtime continuation support for the JVM, ranging from Ovm implementations [11] based on `call/cc`, to JVM extensions [49], based on the `capture` and `resume` primitives. While runtime continuations are not delimited and can be made very efficient, maintenance pressure and portability requirements prevented these implementations from becoming a part of official JVM releases. An alternative, less demanding approach relies only on stack introspection facilities of the host runtime [26]. There exists a program transformation that relies on exception-handling to capture the stack [48]. Here, before calling the continuation, the saved state is used in method calls to rebuild the stack. This works well for continuations, where the stack must be copied anyway, but may be too costly for coroutine `resume`. Bruggeman et al. observed that many use cases call the continuation only once and can avoid the copying overhead, which lead to *one-shot continuations* [7]. One-shot continuations are akin to coroutines without snapshots.

**Domain-specific approaches.** One of the early coroutine applications was data structure traversal. Push-style traversal with `foreach` is easy, but the caller must relinquish control, and many applications cannot do this (e.g. the same-fringe benchmark, in which two trees are traversed pairwise simultaneously). Java-style iterators with `next` and `hasNext` are harder to implement than a `foreach` method, and coroutines bridge this gap.

Iterators in CLU [20] are essentially coroutines – program sections with `yield` statements that are converted into traversal objects. C# inherited this approach – its iterator type `IEnumerator` exposes `Current` and `MoveNext` methods. Since enumerator methods are not first class entities, it is somewhat harder to abstract suspendable code. C# enumerators are not stackful, so iterator definitions must be implemented inside a single method. Enumerators can be used for asynchronous programming, but they require exposing `yield` in user code.



Therefore, separately from enumerators, C# exposes async-await primitives. Some newer languages such as Dart similarly expose an async-await pair of primitives.

Async-Await in Scala [16] is implemented using Scala's metaprogramming facilities. Async-await programs can compose by expressing asynchronous components as first-class `Future` objects. The Async-Await model does not need to be stackful, since separate modules can be expressed as separate futures. However, reliance on futures and concurrency makes it hard to use Async-Await generically. For example, iterators implemented using futures have considerable performance overheads due to synchronization involved in creating future values.

There exist other domain-specific suspension models. For example, Erlang's `receive` statement effectively captures the program continuation when awaiting for the inbound message [52]. A model similar to Scala Async was devised to generate Rx's `Observable` values [13, 23], and the event stream composition [33, 36] in the reactor model [34, 32, 41, 30], as well as callbacks usages in asynchronous programming models based on futures and flow-pools [15, 40, 39, 29, 47] can be similarly simplified. Cilk's spawn-sync model [19] is similar to async-await, and it is implemented as a full program transformation. The Esterel language defines a `pause` statement that pauses the execution, and continues it in the next event propagation cycle [5]. Behaviour trees [22] are AI algorithms used to simulate agents – they essentially behave as AST interpreters with yield statements.

---

## References

---

- 1 Akka. Akka documentation, 2011. <http://akka.io/docs/>.
- 2 Konrad Anton and Peter Thiemann. *Towards Deriving Type Systems and Implementations for Coroutines*, pages 63–79. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. URL: [http://dx.doi.org/10.1007/978-3-642-17164-2\\_6](http://dx.doi.org/10.1007/978-3-642-17164-2_6), doi:10.1007/978-3-642-17164-2\_6.
- 3 Kenichi Asai and Chihiro Uehara. Selective cps transformation for shift and reset. In *Proceedings of the ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation*, PEPM '18, pages 40–52, New York, NY, USA, 2018. ACM. URL: <http://doi.acm.org/10.1145/3162069>, doi:10.1145/3162069.
- 4 V. Beltran, D. Carrera, J. Torres, and E. Ayguade. Evaluating the scalability of java event-driven web servers. In *International Conference on Parallel Processing, 2004. ICPP 2004.*, pages 134–142 vol.1, Aug 2004. doi:10.1109/ICPP.2004.1327913.
- 5 Gérard Berry and Georges Gonthier. The ESTEREL Synchronous Programming Language: Design, Semantics, Implementation. *Sci. Comput. Program.*, 19(2):87–152, November 1992. URL: [http://dx.doi.org/10.1016/0167-6423\(92\)90005-V](http://dx.doi.org/10.1016/0167-6423(92)90005-V), doi:10.1016/0167-6423(92)90005-V.
- 6 G.M. Birtwhistle, O.J. Dahl, B. Myhrhaug, and K. Nygaard. *Simula Begin*. Chartwell-Bratt Ltd, 1979.
- 7 Carl Bruggeman, Oscar Waddell, and R. Kent Dybvig. Representing control in the presence of one-shot continuations. In *Proceedings of the ACM SIGPLAN 1996 Conference on Programming Language Design and Implementation*, PLDI '96, pages 99–107, New York, NY, USA, 1996. ACM. URL: <http://doi.acm.org/10.1145/231379.231395>, doi:10.1145/231379.231395.
- 8 Mike Cantelon, Marc Harter, TJ Holowaychuk, and Nathan Rajlich. *Node.js in Action*. Manning Publications Co., Greenwich, CT, USA, 1st edition, 2013.
- 9 Melvin E. Conway. Design of a Separable Transition-Diagram Compiler. *Commun. ACM*, 6(7):396–408, July 1963. URL: <http://doi.acm.org/10.1145/366663.366704>, doi:10.1145/366663.366704.

- 10 Olivier Danvy and Andrzej Filinski. Abstracting Control. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming*, LFP '90, pages 151–160, New York, NY, USA, 1990. ACM. URL: <http://doi.acm.org/10.1145/91556.91622>, doi:10.1145/91556.91622.
- 11 Iulian Dragos, Antonio Cunej, and Jan Vitek. Continuations in the Java Virtual Machine. In *Second ECOOP Workshop on Implementation, Compilation, Optimization of Object-Oriented Languages, Programs and Systems (ICOOOLPS'2007)*, Berlin, 2007. Technische Universität Berlin.
- 12 Jeffrey Fischer, Rupak Majumdar, and Todd Millstein. Tasks: Language support for event-driven programming. In *Proceedings of the 2007 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation*, PEPM '07, pages 134–143, New York, NY, USA, 2007. ACM. URL: <http://doi.acm.org/10.1145/1244381.1244403>, doi:10.1145/1244381.1244403.
- 13 Philipp Haller and Heather Miller. RAY: Integrating Rx and Async for Direct-Style Reactive Streams. In *Workshop on Reactivity, Events and Modularity*, 2013.
- 14 Philipp Haller and Martin Odersky. Scala actors: Unifying thread-based and event-based programming. *Theor. Comput. Sci.*, 410(2-3):202–220, February 2009. URL: <http://dx.doi.org/10.1016/j.tcs.2008.09.019>, doi:10.1016/j.tcs.2008.09.019.
- 15 Philipp Haller, Aleksandar Prokopec, Heather Miller, Viktor Klang, Roland Kuhn, and Vojin Jovanovic. Scala improvement proposal: Futures and promises (SIP-14). 2012. URL: <http://docs.scala-lang.org/sips/pending/futures-promises.html>.
- 16 Philipp Haller and Jason Zaugg. Scala Async Repository, 2013. <https://github.com/scala/async>.
- 17 Robert Harper. A simplified account of polymorphic references. In *INFORMATION PROCESSING LETTERS*, pages 201–206, 1994.
- 18 Ralph E. Johnson and Brian Foote. Designing reusable classes. *Journal of Object-Oriented Programming*, 1(2):22–35, June/July 1988. URL: <http://www.laputan.org/drc.html>.
- 19 Charles E. Leiserson. *Programming irregular parallel applications in Cilk*, pages 61–71. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997. URL: [https://doi.org/10.1007/3-540-63138-0\\_6](https://doi.org/10.1007/3-540-63138-0_6), doi:10.1007/3-540-63138-0\_6.
- 20 Barbara Liskov, Alan Snyder, Russell Atkinson, and Craig Schaffert. Abstraction Mechanisms in CLU. *Commun. ACM*, 20(8):564–576, August 1977. URL: <http://doi.acm.org/10.1145/359763.359789>, doi:10.1145/359763.359789.
- 21 Ingo Maier and Martin Odersky. Deprecating the Observer Pattern with Scala.react. Technical report, 2012.
- 22 A. Marzotto, M. Colledanchise, C. Smith, and P. Ögren. Towards a Unified Behavior Trees Framework for Robot Control. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5420–5427, May 2014. doi:10.1109/ICRA.2014.6907656.
- 23 Erik Meijer. Your Mouse is a Database. *Commun. ACM*, 55(5):66–73, May 2012. URL: <http://doi.acm.org/10.1145/2160718.2160735>, doi:10.1145/2160718.2160735.
- 24 Ken Moody and Martin Richards. A coroutine mechanism for bcpl. *Softw., Pract. Exper.*, 10(10):765–771, 1980.
- 25 Ana Lúcia De Moura and Roberto Ierusalimschy. Revisiting Coroutines. *ACM Trans. Program. Lang. Syst.*, 31(2):6:1–6:31, February 2009. URL: <http://doi.acm.org/10.1145/1462166.1462167>, doi:10.1145/1462166.1462167.
- 26 Greg Pettyjohn, John Clements, Joe Marshall, Shriram Krishnamurthi, and Matthias Felleisen. Continuations from Generalized Stack Inspection. *SIGPLAN Not.*, 40(9):216–227, September 2005. URL: <http://doi.acm.org/10.1145/1090189.1086393>, doi:10.1145/1090189.1086393.

- 27 Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, Cambridge, MA, USA, 2002.
- 28 A. Prokopec, D. Petrashko, and M. Odersky. Efficient lock-free work-stealing iterators for data-parallel collections. In *2015 23rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing*, pages 248–252, March 2015. doi:10.1109/PDP.2015.65.
- 29 Aleksandar Prokopec. *Data Structures and Algorithms for Data-Parallel Computing in a Managed Runtime*. PhD thesis, IC, Lausanne, 2014. doi:10.5075/epfl-thesis-6264.
- 30 Aleksandar Prokopec. Pluggable scheduling for the reactor programming model. In *Proceedings of the 6th International Workshop on Programming Based on Actors, Agents, and Decentralized Control*, AGERE 2016, pages 41–50, New York, NY, USA, 2016. ACM. URL: <http://doi.acm.org/10.1145/3001886.3001891>, doi:10.1145/3001886.3001891.
- 31 Aleksandar Prokopec. Scala Coroutines Website, 2016. <https://storm-enroute/coroutines>.
- 32 Aleksandar Prokopec. Accelerating by idling: How speculative delays improve performance of message-oriented systems. In Francisco F. Rivera, Tomás F. Pena, and José C. Cabaleiro, editors, *Euro-Par 2017: Parallel Processing*, pages 177–191, Cham, 2017. Springer International Publishing.
- 33 Aleksandar Prokopec. Encoding the building blocks of communication. In *Proceedings of the 2017 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, Onward! 2017, pages 104–118, New York, NY, USA, 2017. ACM. URL: <http://doi.acm.org/10.1145/3133850.3133865>, doi:10.1145/3133850.3133865.
- 34 Aleksandar Prokopec. Reactors.io website, 2018. URL: <http://reactors.io>.
- 35 Aleksandar Prokopec, Phil Bagwell, Tiark Rompf, and Martin Odersky. A generic parallel collection framework. In *Proceedings of the 17th international conference on Parallel processing - Volume Part II*, Euro-Par’11, pages 136–147, Berlin, Heidelberg, 2011. Springer-Verlag. URL: <http://dl.acm.org/citation.cfm?id=2033408.2033425>.
- 36 Aleksandar Prokopec, Philipp Haller, and Martin Odersky. Containers and aggregates, mutators and isolates for reactive programming. In *Proceedings of the Fifth Annual Scala Workshop*, SCALA ’14, pages 51–61, New York, NY, USA, 2014. ACM. URL: <http://doi.acm.org/10.1145/2637647.2637656>, doi:10.1145/2637647.2637656.
- 37 Aleksandar Prokopec, David Leopoldseder, Gilles Duboscq, and Thomas Würthinger. Making collection operations optimal with aggressive JIT compilation. In *Proceedings of the 8th ACM SIGPLAN International Symposium on Scala*, SCALA 2017, pages 29–40, New York, NY, USA, 2017. ACM. URL: <http://doi.acm.org/10.1145/3136000.3136002>, doi:10.1145/3136000.3136002.
- 38 Aleksandar Prokopec and Fengyun Liu. Theory and practice of coroutines with snapshots. In *European Conference on Object-Oriented Programming*, 2018.
- 39 Aleksandar Prokopec, Heather Miller, Philipp Haller, Tobias Schlatter, and Martin Odersky. FlowPools: A Lock-Free Deterministic Concurrent Dataflow Abstraction, Proofs. Technical report, 2012.
- 40 Aleksandar Prokopec, Heather Miller, Tobias Schlatter, Philipp Haller, and Martin Odersky. Flowpools: A lock-free deterministic concurrent dataflow abstraction. In *LCPC*, pages 158–173, 2012.
- 41 Aleksandar Prokopec and Martin Odersky. Isolates, Channels, and Event Streams for Composable Distributed Programming. In *2015 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!)*, Onward! 2015, pages 171–182, New York, NY, USA, 2015. ACM. URL: <http://doi.acm.org/10.1145/2814228.2814245>, doi:10.1145/2814228.2814245.

- 42 Aleksandar Prokopec, Dmitry Petrashko, and Martin Odersky. On lock-free work-stealing iterators for parallel data structures. 2014.
- 43 John C. Reynolds. The discoveries of continuations. *LISP and Symbolic Computation*, 6(3):233–247, Nov 1993. URL: <https://doi.org/10.1007/BF01019459>, doi:10.1007/BF01019459.
- 44 Tiark Rompf, Ingo Maier, and Martin Odersky. Implementing First-Class Polymorphic Delimited Continuations by a Type-Directed Selective CPS-Transform. *SIGPLAN Not.*, 44(9):317–328, August 2009. URL: <http://doi.acm.org/10.1145/1631687.1596596>, doi:10.1145/1631687.1596596.
- 45 P. James Roshan and Amr Sabry. Yield: Mainstream delimited continuations. 1994.
- 46 Amr Sabry and Matthias Felleisen. Reasoning About Programs in Continuation-Passing Style. *SIGPLAN Lisp Pointers*, V(1):288–298, January 1992. URL: <http://doi.acm.org/10.1145/141478.141563>, doi:10.1145/141478.141563.
- 47 Tobias Schlatter, Aleksandar Prokopec, Heather Miller, Philipp Haller, and Martin Odersky. Multi-lane flowpools: A detailed look. *Tech Report*, 2012.
- 48 Tatsuro Sekiguchi, Takahiro Sakamoto, and Akinori Yonezawa. Portable implementation of continuation operators in imperative languages by exception handling. In *Advances in Exception Handling Techniques (the Book Grow out of a ECOOP 2000 Workshop)*, pages 217–233, London, UK, UK, 2001. Springer-Verlag. URL: <http://dl.acm.org/citation.cfm?id=647332.722736>.
- 49 Lukas Stadler, Christian Wimmer, Thomas Würthinger, Hanspeter Mössenböck, and John Rose. Lazy continuations for java virtual machines. In *Proceedings of the 7th International Conference on Principles and Practice of Programming in Java, PPPJ '09*, pages 143–152, New York, NY, USA, 2009. ACM. URL: <http://doi.acm.org/10.1145/1596655.1596679>, doi:10.1145/1596655.1596679.
- 50 Arvind K. Sujeeth, Tiark Rompf, Kevin J. Brown, HyoukJoong Lee, Hassan Chafi, Victoria Popic, Michael Wu, Aleksandar Prokopec, Vojin Jovanovic, Martin Odersky, and Kunle Olukotun. Composition and reuse with compiled domain-specific languages. In *Proceedings of the 27th European Conference on Object-Oriented Programming, ECOOP'13*, pages 52–78, Berlin, Heidelberg, 2013. Springer-Verlag. URL: [http://dx.doi.org/10.1007/978-3-642-39038-8\\_3](http://dx.doi.org/10.1007/978-3-642-39038-8_3), doi:10.1007/978-3-642-39038-8\_3.
- 51 Gerald Jay Sussman and Guy L. Steele, Jr. Scheme: A interpreter for extended lambda calculus. *Higher Order Symbol. Comput.*, 11(4):405–439, December 1998. URL: <http://dx.doi.org/10.1023/A:1010035624696>, doi:10.1023/A:1010035624696.
- 52 Robert Virding, Claes Wikström, and Mike Williams. *Concurrent Programming in ER-LANG (2nd Ed.)*. Prentice Hall International (UK) Ltd., Hertfordshire, UK, UK, 1996.
- 53 Philip Wadler. Monads for Functional Programming. In *Advanced Functional Programming, First International Spring School on Advanced Functional Programming Techniques-Tutorial Text*, pages 24–52, London, UK, UK, 1995. Springer-Verlag. URL: <http://dl.acm.org/citation.cfm?id=647698.734146>.
- 54 N. Wirth. *Programming in Modula-2*. Texts and Monographs in Computer Science. Springer-Verlag, 1985. URL: <https://books.google.ch/books?id=ZVaRXPrD1AoC>.

$$\begin{array}{c}
\frac{\Gamma \vdash T_y | T_r | x \rightarrow \xi(T_y, T_r, T, k(x)) \quad \Gamma \vdash T_y | T_r | () \rightarrow \xi(T_y, T_r, \text{Unit}, k(()))}{\Gamma \vdash T_y | T_r | x \rightarrow \xi(T_y, T_r, T, k(x))} \text{(X-VAR)} \quad \frac{\Gamma \vdash T_y | T_r | () \rightarrow \xi(T_y, T_r, \text{Unit}, k(()))}{\Gamma \vdash T_y | T_r | () \rightarrow \xi(T_y, T_r, \text{Unit}, k(()))} \text{(X-UNIT)} \\
\frac{\Gamma \vdash T_y | T_r | v \rightarrow \xi(T_y, T_r, T'_1 \Rightarrow T'_2, k(v'))}{\Gamma \vdash T_y | T_r | v \rightarrow \xi(T_y, T_r, T'_1 \Rightarrow T'_2, k(v'))} \text{(X-ABS)} \\
\frac{\Gamma \vdash T_y | T_r | t_1 \rightarrow t'_1 \quad \Gamma \vdash T_y | T_r | t_2 \rightarrow t'_2 \quad T_{21} = T_2 \Rightarrow T_1 \quad p = t'_1(s) ((x_1 : T_{21}) \Rightarrow t'_2(s) ((x_2 : T_2) \Rightarrow k(x_1(x_2))))}{\Gamma \vdash T_y | T_r | t_1(t_2) \rightarrow \xi(T_y, T_r, T_1, p)} \text{(X-APP)} \\
\frac{\Gamma, f : T \vdash t : T | \perp \quad \Gamma, f : T' \vdash t' : T' | \perp \quad T_y \neq \perp \quad \Gamma, f : T \vdash T_y | T_r | t \rightarrow t'}{\Gamma \vdash T_y | T_r | \text{fix}((f : T) \Rightarrow t) \rightarrow \xi(T_y, T_r, T' \Rightarrow T', k(\text{fix}((f : T') \Rightarrow t')))} \text{(X-FIX)} \\
\frac{v = (x : T_1) \xrightarrow{T_y} t \quad \Gamma \vdash v : T_1 \xrightarrow{T_y} T_2 | \perp \quad \Gamma, x : T_1 \vdash T_y | T_2 | t \rightarrow t' \quad T_y \neq \perp \quad q = (s : \sigma[T_y, T_2]) \Rightarrow (x : T_1) \Rightarrow p \quad p = t'_1(s) ((x_1 : \gamma[T_2, T_y, T_1]) \Rightarrow t'_2(s) ((x_2 : T_2) \Rightarrow k(q_1; q_2; x)))}{\Gamma \vdash T_w | T_r | v \rightarrow \xi(T_w, T_r, \gamma[T_1, T_y, T_2], k(q))} \text{(X-COROUTINE)} \\
\frac{\Gamma \vdash T_w | T_r | \text{start}(t_1, t_2) \rightarrow \xi(T_w, T_r, \rho(T_y, T_1), p)}{\Gamma \vdash T_w | T_r | \text{start}(t_1, t_2) \rightarrow \xi(T_w, T_r, \rho(T_y, T_1), p)} \text{(X-START)} \\
\frac{\Gamma \vdash T | T_r | t \rightarrow t' \quad p = t'(s) ((x : T) \Rightarrow s(k); \text{Yield}(x))}{\Gamma \vdash T | T_r | \text{yield}(t) \rightarrow \xi(T, T_r, \text{Unit}, p)} \text{(X-YIELD)} \\
\frac{\Gamma \vdash T | T_r | t \xrightarrow{T_y} T_2 | T_w \quad \Gamma \vdash T_w | T_r | t \rightarrow t' \quad p = t'(s) ((x : \rho(T_y, T_2)) \Rightarrow k(\text{ref}(!x)))}{\Gamma \vdash T_w | T_r | \text{snapshot}(t) \rightarrow \xi(T_w, T_r, \rho(T_y, T_2), p)} \text{(X-SNAPSHOT)} \\
\frac{\Gamma \vdash T_w | T_r | t_1 \rightarrow t'_1 \quad \Gamma \vdash T_w | T_r | t_2 \rightarrow t'_2 \quad \Gamma \vdash T_w | T_r | t_3 \rightarrow t'_3 \quad \Gamma \vdash T_w | T_r | t_4 \rightarrow t'_4 \quad \Gamma \vdash T_w | T_r | x_2(x) \rightarrow q'_2 \quad \Gamma \vdash T_w | T_r | x_3(x) \rightarrow q'_3 \quad \Gamma \vdash T_w | T_r | x_4(()) \rightarrow q'_4 \quad U = \text{Unit} \quad p = t'_1(s) ((x_1 : \rho[T_y, T_2]) \Rightarrow t'_2(s) ((x_2 : \gamma[T_2, T_w, T_R]) \Rightarrow t'_3(s) ((x_3 : \gamma[T_y, T_w, T_R]) \Rightarrow t'_4(s) ((x_4 : \gamma[U, T_w, T_R]) \Rightarrow q))) \quad q = (!x_1) (()) \text{ match } \{ \text{case Ret}(x) \Rightarrow q'_2; \text{case Yield}(x) \Rightarrow q'_3; \text{case Term} \Rightarrow q'_4 \}}}{\Gamma \vdash T_w | T_r | \text{resume}(t_1, t_2, t_3, t_4) \rightarrow \xi(T_w, T_r, T_R, p)} \text{(X-RESUME)}
\end{array}$$

■ **Figure 4** Transformation of coroutines in  $\lambda_{\rightarrow}$  to a simply typed lambda calculus with references

$$\begin{array}{c}
 \frac{x:T \in \Gamma}{\Gamma \vdash \perp | \perp | x \rightarrow x} \quad (X\text{-FREEVAR}) \\
 \\
 \frac{\Gamma \vdash \perp | \perp | () \rightarrow ()}{\Gamma \vdash \perp | \perp | () \rightarrow ()} \quad (X\text{-FREEUNIT}) \\
 \\
 \frac{v = (x:T_1) \Rightarrow t \quad v' = (x:T'_1) \Rightarrow t' \quad \Gamma, x:T_1 \vdash t:T_2 | \perp \quad \Gamma, x:T'_1 \vdash t':T'_2 | \perp \quad \Gamma, x:T_1 \vdash \perp | \perp | t \rightarrow t'}{\Gamma \vdash \perp | \perp | (x:T_1) \Rightarrow t \rightarrow (x:T'_1) \Rightarrow t'} \quad (X\text{-FREEABS}) \\
 \\
 \frac{\Gamma \vdash t_1:T_2 \Rightarrow T_1 | T_y \quad \Gamma \vdash t_2:T_2 | T_y \quad \Gamma \vdash \perp | \perp | t_1 \rightarrow t'_1 \quad \Gamma \vdash \perp | \perp | t_2 \rightarrow t'_2}{\Gamma \vdash \perp | \perp | t_1(t_2) \rightarrow t'_1(t'_2)} \quad (X\text{-FREEAPP}) \\
 \\
 \frac{\Gamma \vdash \perp | \perp | t \rightarrow t' \quad \Gamma \vdash t':T' | \perp}{\Gamma \vdash \perp | \perp | \text{fix}((f:T) \Rightarrow t) \rightarrow \text{fix}((f:T') \Rightarrow t')} \quad (X\text{-FREEFIX}) \\
 \\
 \frac{v = (x:T_1) \overset{T_y}{\rightsquigarrow} t \quad \Gamma \vdash v:T_1 \overset{T_y}{\rightsquigarrow} T_2 | \perp \quad \Gamma, x:T_1 \vdash T_y | T_2 | t \rightarrow t' \quad T_y \neq \perp \quad q = (s:\sigma[T_y, T_2]) \Rightarrow (x:T_1) \Rightarrow p}{\Gamma \vdash \perp | \perp | v \rightarrow q} \quad (X\text{-FREECOROUTINE}) \\
 \\
 \frac{\Gamma \vdash t_1:T_2 \overset{T_y}{\rightsquigarrow} T_1 | T_w \quad \Gamma \vdash t_2:T_2 | T_w \quad \Gamma \vdash \perp | \perp | t_1 \rightarrow t'_1 \quad \Gamma \vdash \perp | \perp | t_2 \rightarrow t'_2}{\Gamma \vdash \perp | \perp | \text{start}(t_1, t_2) \rightarrow q_1; q_2; x} \quad (X\text{-FREESTART}) \\
 \\
 \frac{\Gamma \vdash t:T_y \rightsquigarrow T_2 | T_w \quad \Gamma \vdash \perp | \perp | t \rightarrow t'}{\Gamma \vdash \perp | \perp | \text{snapshot}(t) \rightarrow \text{ref}(!t')} \quad (X\text{-FREESNAPSHOT}) \\
 \\
 \frac{\Gamma \vdash t_1:T_y \rightsquigarrow T_2 | T_w \quad \Gamma \vdash t_2:T_2 \overset{T_w}{\rightsquigarrow} T_R | T_w \quad \Gamma \vdash t_3:T_y \overset{T_w}{\rightsquigarrow} T_R | T_w \quad \Gamma \vdash t_4:\text{Unit} \overset{T_w}{\rightsquigarrow} T_R | T_w \quad \Gamma \vdash \perp | \perp | t_1 \rightarrow t'_1 \quad \Gamma \vdash \perp | \perp | t_2 \rightarrow t'_2 \quad \Gamma \vdash \perp | \perp | t_3 \rightarrow t'_3 \quad \Gamma \vdash \perp | \perp | t_4 \rightarrow t'_4}{\Gamma \vdash \perp | \perp | \text{resume}(t_1, t_2, t_3, t_4) \rightarrow q} \quad (X\text{-FREERESUME}) \\
 \\
 q = (!t'_1)() \text{ match } \{ \text{case Ret}(x) \Rightarrow t'_2(x); \text{case Yield}(x) \Rightarrow t'_3(x); \text{case Term} \Rightarrow t'_4() \}
 \end{array}$$

■ **Figure 5** Transformation of non-yielding terms in  $\lambda_{\rightsquigarrow}$  to a simply typed lambda calculus with references