

# A One-Pass CPS Transform with Simulation on the Nose

Pascal Y. Lasnier<sup>[0009–0009–9371–3467]</sup>, Jeremy Yallop<sup>[0009–0002–1650–6340]</sup><sup>1</sup>, and  
Magnus O. Myreen<sup>[0000–0002–9504–4107]</sup><sup>2,3</sup>

<sup>1</sup> University of Cambridge

<sup>2</sup> Chalmers University of Technology

<sup>3</sup> University of Gothenburg

**Abstract.** Danvy & Nielsen’s one-pass CPS transform has a straightforward definition, but clashes between the names of variables it introduces make it difficult to mechanically prove correct. Existing mechanical proofs either side-step the issue by using nameless representations, or rely on tedious  $\alpha$ -equivalence relations between target terms. This paper presents a new formulation of the transform using evaluation contexts that allows deterministic introduction of fresh names, eliminating the need to work up to  $\alpha$ -equivalence. We use our formulation to present a new and straightforward simulation proof of the correctness of the one-pass CPS transform, which we have mechanised in the HOL4 theorem prover.

## 1 Introduction


Continuation-passing style (CPS) transforms have many useful properties: they make evaluation order explicit, turn all calls into tail calls, name intermediate computations, and support simulation of non-standard control flow. These properties make the CPS transform a key technique in many applications, from defining control operator semantics [7] to intermediate representation in functional languages compilers [12] and languages with first-class continuations [2,13,22].

The classic CPS transform developed by Plotkin [18] introduces many *administrative redexes*, i.e. function applications which are trivial to evaluate and do not correspond to source program reductions. For example, the program fragment  $x_1 x_2 (x_3 x_4 x_5)$  CPS-transforms into the following:

$$\begin{aligned} & \lambda k. (\lambda k. (\lambda k. k x_1) (\lambda m. (\lambda k. k x_2) (\lambda n. m n k))) \\ & (\lambda m. (\lambda k. (\lambda k. (\lambda k. k x_3) (\lambda m. (\lambda k. k x_4) (\lambda n. m n k)))) \\ & (\lambda m. (\lambda k. k x_5) (\lambda n. m n k))) \\ & (\lambda n. m n k) \end{aligned}$$

Reducing applications of the form  $(\lambda k. k x)(\lambda y. e)$ , to  $[x/y]e$  (i.e.  $e$  with  $x$  substituted for  $y$ ) straightforwardly produces the following reduced term:

$$\lambda k. (\lambda k. x_1 x_2 k) (\lambda m. (\lambda k. (\lambda k. x_3 x_4 k) (\lambda m. m x_5 k)) (\lambda n. m n k))$$



**Fig. 1.** Simulation as a commutative diagram.

There are three further administrative redexes that can be eliminated in this program. However, reducing these redexes requires renaming one of the  $m$  variables to avoid clashes. After reduction and renaming, the result is as follows:

$$\lambda k.x_1 x_2 (\lambda m.x_3 x_4 (\lambda m'.m' x_5 (\lambda n.m n k)))$$

There are several optimised CPS transforms [4,9,21] that eliminate these administrative redexes and produce CPS programs such as the final version above. The state-of-the-art for the call-by-value lambda calculus is the first-order one-pass CPS transform by Danvy & Nielsen [9].

Proof of correctness of the CPS transform, which is typically formulated as a simulation proof, involves establishing a relation between source programs and CPS programs which the CPS transform satisfies, and proving that the relation is closed under source and target reductions. Fig. 1 portrays this statement as a commutative diagram. Simulation is simple to prove for the classic CPS transform, but the name clashes seen in the example above introduce challenges in the proofs for the optimised variants.

In their hand-written correctness proof, Danvy & Nielsen simply assert that introduced variables are unique [9]. In contrast, in a typical mechanised proof, binders require explicit reasoning. A naive fresh binder naming scheme will result in CPS programs that do not strictly reduce in the simulation diagram, but instead reduce to  $\alpha$ -equivalent programs [15], complicating the simulation proof with a new notion of equivalence.

Paraskevopoulou & Grover [16] summarise the accumulated literature on the variable binder issue for mechanised proofs of the optimised CPS transform. Proofs which represent variable binders with explicit names require establishing some notion of equivalence over programs in CPS, with Minamide & Okuma [15] using an  $\alpha$ -equivalence relation, and Paraskevopoulou & Grover [16] developing a more general logical relation over CPS terms. These approaches complicate verified implementation of the CPS transform, since constructing an equivalence relation and proving it for a target language involves significant additional verification work, especially if the target language is complex. Other developments use nameless variable binding system such as de Bruijn indices [10,19] or higher-order abstract syntax [3,23].

This paper contributes a new approach to the problem of name generation in an optimised CPS transform, with a mechanised simulation proof that both uses

named variable binders *and* completes the simulation diagram without requiring an additional notion of equivalence. Our approach to the problem consists of three components:

**Evaluation-context-based optimised CPS transform.** We refactor Danvy & Nielsen’s optimised CPS transform [9] using evaluation contexts, following Sabry & Felleisen’s non-compositional optimised transform [21], but preserving compositionality. Reframing the transform in this way highlights a critical feature of variable name introduction: name clashes are *local to the lexical scope* of source language terms. This observation is an instance of a general principle that freshness should not be treated as a global property, but determined relative to some local context [17]. We apply the principle to the challenge at hand, using it to design a name generation strategy that is deterministically dependent *only* on the evaluation context of the local scope.

**Modified CEK-machine with lexical scopes.** Using the insight that variable binder names may be allocated locally to the lexical scope of the source language, we modify the CEK-machine small-step semantic definition of the call-by-value lambda calculus by rephrasing evaluation contexts as a nested sequence of scopes, where each scope has a distinct, *local* environment and evaluation context. At the top level of the semantic state, the new abstract machine operates on a sequence of scopes, with the innermost evaluation context containing the control string in its hole.

**Simulation.** We construct a simulation relation between semantic states of the modified CEK-machine for the call-by-value lambda calculus to CPS terms through our evaluation-context-based optimised CPS transform, using the evaluation context of the local scope from the semantic state directly in the transform. Because variable binder names are deterministic from the local lexical scope, the exact variable name can be deduced directly from the simulation relation, and simulation can be easily proven without requiring a notion of equivalence.

We have mechanised our proof in the HOL4 theorem prover, and include the development as supplementary material.

## 2 Call-by-value lambda calculus CEK-machine

$$\begin{array}{ll}
 e ::= t \mid s & \in \text{Exp} \\
 t ::= x \mid \lambda x.e & \in \text{Triv} \\
 s ::= e_1 e_2 & \in \text{Comp} \\
 x & \in \text{Ide}
 \end{array}$$

**Fig. 2.** Lambda calculus grammar

We first present the lambda calculus which we use in our mechanised proof (Fig. 2). Terms  $e$  are either trivial terms  $t$ , which include variables  $x$  and abstractions  $\lambda x.e$ , or serious terms  $s$ , which are applications  $e_1 e_2$ .

$$\begin{array}{lll}
& \text{Cont[Ctrl]} \longrightarrow \text{Cont[Ctrl]} \\
v ::= \lambda x.\{\rho\}e & \in \text{Val} & E[\{\rho\}x] \longrightarrow E[\rho(x)] \\
\rho : \text{Ide} \rightarrow \text{Val} & & E[\{\rho\}\lambda x.e] \longrightarrow E[\lambda x.\{\rho\}e] \\
E ::= [] \mid E\{\rho\}e \mid vE & \in \text{Cont} & E[\{\rho\}e_1 e_2] \longrightarrow E[[\{\rho\}e_1]\{\rho\}e_2] \\
c ::= \{\rho\}e \mid v & \in \text{Ctrl} & E[[v]\{\rho\}e] \longrightarrow E[v[\{\rho\}e]] \\
& & E[(\lambda x.\{\rho\}e)[v]] \longrightarrow E[\{\rho[x \mapsto v]\}e]
\end{array}$$

**Fig. 3.** CEK-machine

Fig. 3 shows the call-by-value semantics for the calculus as a small-step CEK-machine semantics. The CEK-machine operates on control strings which may be either terms with an attached environment, which we denote similarly to substitution  $\{\rho\}e$ , or values  $v$ , which may only be closures which capture an environment  $\lambda x.\{\rho\}e$ . Environments  $\rho$  simply map variables to values.

The semantic state of the CEK-machine includes an evaluation context  $E$ , which is a nested sequence of stack frames. There are two types of stack frame: **fn** frames  $E\{\rho\}e$  enclose an unevaluated function term in an application, where the evaluation context holds the argument term and associated environment  $\{\rho\}e$ , and **arg** frames  $vE$  enclose an unevaluated argument term, where the evaluation context holds the previously evaluated function value  $v$ .

The complete definition of the CEK-machine is a relation between semantic states consisting of an evaluation context  $E$  with its hole filled by a control string  $c$ . We denote the type of a filled evaluation context by figuratively filling the hole of the  $\text{Cont}$  type, i.e. the type  $\text{Cont[Ctrl]}$  is that of evaluation contexts with their holes filled by control strings  $E[c]$ .

### 3 Revisiting the optimised CPS transform

#### 3.1 Classic transform

Fig. 4 shows Plotkin's classic CPS transform [18]. To distinguish the source and target domains, we write target terms in **bold**. The proof of Plotkin's classic simulation theorem is simple and well-understood, and has been mechanised by Minamide & Okuma [15] and Dargaye & Leroy [10].

A simulation proof technique which has proven quite effective is to relate CEK-machine source semantic states to CPS terms in the target domain. This relation is not arbitrary, but rather arises from the nature of small-step semantics

$$\begin{aligned}
\llbracket x \rrbracket &= \lambda \mathbf{k}. \mathbf{k} \ x \\
\llbracket \lambda x. e \rrbracket &= \lambda \mathbf{k}. \mathbf{k} \ (\lambda x. \llbracket e \rrbracket) \\
\llbracket e_1 \ e_2 \rrbracket &= \lambda \mathbf{k}. \llbracket e_1 \rrbracket \ (\lambda \mathbf{m}. \llbracket e_2 \rrbracket \ (\lambda \mathbf{n}. \mathbf{m} \ \mathbf{n} \ \mathbf{k}))
\end{aligned}$$

**Fig. 4.** Plotkin’s classic CPS transform

as a *defunctionalisation* of the continuation-passing style [1,5,6,8,20], and it underscores the nature of CPS as a *simulation* of control structures (i.e. evaluation contexts in a small-step semantics) by terms in a target domain. This relation generalises to larger source languages: Lasnier, Yallop, & Myreen [14] use the relation to prove simulation for Scheme, including first-class continuations.

### 3.2 One-pass transform

Fig. 5 shows Danvy & Nielsen’s first-order, compositional, one-pass formulation of the optimised CPS transform [9], with an explicit naming scheme we motivate in the following sections. The subscript  $\mathbf{K}$  in the expression transform  $\mathcal{E}$  corresponds to the top-level continuation of a given scope: it may be only either the identity continuation  $\mathbf{I}$  or a continuation variable  $\mathbf{k}$ . In contrast, in the serious transform  $\mathcal{S}$ , the continuation argument  $\mathbf{K}$  may be an arbitrary continuation. The complete transform of a program  $p$  is  $\mathcal{E}_{\mathbf{I}} \llbracket p \rrbracket$ .

$$\begin{aligned}
\mathcal{E}_{\mathbf{I}} \llbracket t \rrbracket &= \mathcal{T} \llbracket t \rrbracket & \mathcal{T} \llbracket x \rrbracket &= \mathbf{x} \\
\mathcal{E}_{\mathbf{k}} \llbracket t \rrbracket &= \mathbf{k} \ \mathcal{T} \llbracket t \rrbracket & \mathcal{T} \llbracket \lambda x. e \rrbracket &= \lambda \mathbf{x} \mathbf{k}. \mathcal{E}_{\mathbf{k}} \llbracket e \rrbracket \\
\mathcal{E}_{\mathbf{K}} \llbracket s \rrbracket &= \mathcal{S} \llbracket s \rrbracket_0 \ \mathbf{K}
\end{aligned}$$
  

$$\begin{aligned}
\mathcal{S} \llbracket t_1 \ t_2 \rrbracket_i \ \mathbf{K} &= \mathcal{T} \llbracket t_1 \rrbracket \ \mathcal{T} \llbracket t_2 \rrbracket \ \mathbf{K} \\
\mathcal{S} \llbracket t_1 \ s_2 \rrbracket_i \ \mathbf{K} &= \mathcal{S} \llbracket s_2 \rrbracket_i \ (\lambda \mathbf{n}. \mathcal{T} \llbracket t_1 \rrbracket \ \mathbf{n} \ \mathbf{K}) \\
\mathcal{S} \llbracket s_1 \ t_2 \rrbracket_i \ \mathbf{K} &= \mathcal{S} \llbracket s_1 \rrbracket_i \ (\lambda \mathbf{m}_i. \mathbf{m}_i \ \mathcal{T} \llbracket t_2 \rrbracket \ \mathbf{K}) \\
\mathcal{S} \llbracket s_1 \ s_2 \rrbracket_i \ \mathbf{K} &= \mathcal{S} \llbracket s_1 \rrbracket_i \ (\lambda \mathbf{m}_i. \mathcal{S} \llbracket s_2 \rrbracket_{i+1} \ (\lambda \mathbf{n}. \mathbf{m}_i \ \mathbf{n} \ \mathbf{K}))
\end{aligned}$$

**Fig. 5.** Danvy & Nielsen’s one-pass CPS transform

### 3.3 Clash of variable names and $\alpha$ -equivalence

Minamide & Okuma [15] observe that it is not possible to implement Danvy & Nielsen’s optimised transform without variable renaming, since it may introduce multiple conflicting instances of the variable  $\mathbf{m}$ . They highlight this property with the example transform for  $x_1 \ x_2 \ (x_3 \ x_4 \ x_5)$ , in which one of the  $\mathbf{m}$  variables must be renamed:

$$x_1 \ x_2 \ (\lambda \mathbf{m}. x_3 \ x_4 \ (\lambda \mathbf{m}'. \mathbf{m}' \ x_5 \ (\lambda \mathbf{n}. \mathbf{m} \ \mathbf{n} \ \mathbf{k})))$$

Implementations of the optimised transform must consequently generate fresh variable names. Minamide & Okuma [15] generate names via a counter  $i$  attached to the transform similarly to Fig. 5, and Paraskevopoulou & Grover [16] allocate globally fresh names using a state monad. However, Danvy & Nielsen’s simulation proof is complicated by these fresh variable name generation schemes, which only fit the simulation diagram up to  $\alpha$ -equivalence. Consider Danvy & Nielsen’s simulation theorem, which states that if a serious source term reduces to a new term, the CPS transform of the serious term must also reduce to the CPS transform of the new term. Using our definition from Fig. 5 and splitting cases for either a trivial or serious reduced source term, the theorem is formally stated as follows:

$$\begin{aligned} s \longrightarrow t &\implies \mathcal{S}[[s]]_i \mathbf{K} \longrightarrow^+ \mathbf{K} \mathcal{T}[[t]] \\ s \longrightarrow s' &\implies \mathcal{S}[[s]]_i \mathbf{K} \longrightarrow^+ \mathcal{S}[[s']]_i \mathbf{K} \end{aligned}$$

Decomposing  $s$  into an evaluation context and a redex  $E[t_1 t_2]$ , the proof [9] follows by induction over  $E$ . This proof deals with the lambda calculus using term rewriting with substitutions rather than an abstract machine with an environment, so the evaluation context does not contain environments  $\rho$ .

Let us consider the case for an  $\text{fn}$  stack frame  $E e$  where the argument expression is also a serious term, so that our source term may be written as  $s_1 s_2$ . The CPS transform for this term is  $\mathcal{S}[[s_1]]_i (\lambda \mathbf{m}_i. \mathcal{S}[[s_2]]_{i+1} (\lambda \mathbf{n}. \mathbf{m}_i \mathbf{n} \mathbf{K}))$ .

If  $s_1$  reduces to a trivial term  $s_1 \longrightarrow t$ , we may apply our induction hypothesis for the reduction to a trivial term, so our CPS term reduces to  $(\lambda \mathbf{m}_i. \mathcal{S}[[s_2]]_{i+1} (\lambda \mathbf{n}. \mathbf{m}_i \mathbf{n} \mathbf{K})) \mathcal{T}[[t]]$ , then further to  $\mathcal{S}[[s_2]]_{i+1} (\lambda \mathbf{n}. \mathcal{T}[[t]] \mathbf{n} \mathbf{K})$ . This CPS term, though indeed the CPS transform of the reduced source term  $t s_2$ , does not satisfy the simulation theorem, because the fresh variable name counter  $i$  has been incremented.

It would be incorrect to increment the counter  $i$  in the starting simulation theorem as a way to compensate for this mismatch, because applying the induction hypothesis to the  $s_1 s_2$  case where  $s_1$  reduces to another serious term  $s_1 \longrightarrow s'_1$  results in the reduced CPS term  $\mathcal{S}[[s'_1]]_{i+1} (\lambda \mathbf{m}_i. \mathcal{S}[[s_2]]_{i+1} (\lambda \mathbf{n}. \mathbf{m}_i \mathbf{n} \mathbf{K}))$ , which is strictly not the CPS transform of  $s'_1 s_2$  and breaks simulation.

It is therefore required, when proving simulation with explicit binder names, to establish  $\alpha$ -equivalence over the CPS transform for incrementing fresh variable counters  $\mathcal{S}[[s]]_i =_\alpha \mathcal{S}[[s]]_{i+1}$  as described by Minamide & Okuma [15].

### 3.4 Optimised CPS transform as static reduction

The core issue with existing implementations of the optimised transform is that fresh variable names of an arbitrary program fragment in simulation cannot be determined by the simulation relation from the fragment alone, instead depending on a floating variable counter  $i$  or some global fresh variable allocator. In order to eliminate the need for an equivalence relation, we seek a transform with a variable name generation method which is deterministic, whilst still avoiding clashes.

With this aim in mind, we reformulate the optimised transform to more closely correspond to the source semantics, with the goal of facilitating a suitable

simulation relation. Our starting point is the evaluation-context-based optimised transform presented by Sabry & Felleisen [21]. Drawing on the observation that the optimised transform reduces administrative redexes at the point of the transform, we construct the optimised transform as a kind of partial evaluator based on evaluation contexts like the CEK-machine source semantics, which is given in Fig. 6. Whereas evaluation contexts in the source semantics carry values  $v$  for **arg** frames, in the transform evaluation contexts carry trivial terms  $t$  in the target domain. We highlight this distinction by denoting evaluation contexts in the transform with a star  $E^*$ .

$$\begin{aligned}
\mathcal{E}_K[[E^*[x]]] &= \mathcal{C}_K[[E^*]] x \\
\mathcal{E}_K[[E^*[\lambda x.e]]] &= \mathcal{C}_K[[E^*]] \lambda x k. \mathcal{E}_k[[e]] \\
\mathcal{E}_K[[E^*[e_1 e_2]]] &= \mathcal{E}_K[[E^*[[e_1] e_2]]] \\
\\
\mathcal{C}_I[[[]] t] &= t \\
\mathcal{C}_k[[[]] t] &= k t \qquad \mathcal{K}_K[[[]] ] = K \\
\mathcal{C}_K[[E^*[[ ] e]] t] &= \mathcal{E}_K[[E^*[t [e]]]] \quad \mathcal{K}_K[[E^*[[ ] e]]] = \lambda m_i. \mathcal{E}_K[[E^*[m_i [e]]]] \quad i = \text{fresh } E^* \\
\mathcal{C}_K[[E^*[t_1 []]] t_2] &= t_1 t_2 \mathcal{K}_K[[E^*]] \quad \mathcal{K}_K[[E^*[t []]]] = \lambda n. t n \mathcal{K}_K[[E^*]]
\end{aligned}$$

**Fig. 6.** Our evaluation-context-based optimised CPS transform

In viewing the optimised CPS transform as partial evaluation, we conceptually split continuation applications into static and dynamic variants. In a static application  $\mathcal{C}_K[[E^*]] t$ , the argument is a trivial term in the target domain  $t$  rather than a value, and is directly substituted into the continuation bodies at the point of the transform. In a dynamic application  $k t$ ,  $k$  maps to a closure evaluated from  $\mathcal{K}_K[[E^*]]$ . The definitions of both these continuation transforms derive from the continuation steps of the CEK-machine from Fig. 3, i.e. steps from a value control string with a particular evaluation context.

The evaluation contexts used to form these static and dynamic continuations are grown from the shape of the program by the expression transform  $\mathcal{E}$ , which operates on an evaluation context filled with an expression, and whose definition is derived from the expression steps of the CEK-machine from Fig. 3.

One key distinction between our  $\mathcal{E}$  transform and the transform by Sabry & Felleisen [21] is that  $\mathcal{E}$  explicitly recurses into the evaluation context  $E^*$  of an expression, while Sabry & Felleisen’s transform does not. Sabry & Felleisen instead implicitly arrive at an evaluation context via unique decomposition [24], the same notion relating CEK-machines and term-rewriting systems [11], and in effect squash our  $\mathcal{E}$  and  $\mathcal{C}$  transforms into a single  $\mathcal{C}$  transform. A second distinction is that  $\mathcal{E}$  is compositional, preserving the adjustments Danvy & Nielsen [9] made to the transform.

**Benefits of the new transform.** As Section 5 shows, constructing the transform using evaluation contexts benefits the simulation proof by allowing us to directly relate the source language semantic evaluation state, including its evaluation context, to a CPS-transformed expression in the target language. We believe that this same construction can be applied to more complex languages with small-step semantics that use evaluation contexts, such as Scheme.

Constructing the transform as we have done also more clearly highlights the name clash issue encountered by Minamide & Okuma [15]. Note the continuation closure definition for an **fn** continuation  $\mathcal{K}_{\mathbf{k}}[E^*[\square] e]$ ; this closure introduces a variable  $\mathbf{m}$  containing the resultant value of application function term to fill the hole in the evaluation context  $E^*[\square] e$ . However, this variable is not immediately dispatched, and is instead stashed into the new evaluation context from which the next term  $e$  is transformed. Consequently, there can be multiple variables  $\mathbf{m}$  present in the evaluation context; for example the transform of the  $x_1 x_2 (x_3 x_4 x_5)$  expands as follows:

$$\begin{aligned} \mathcal{E}_{\mathbf{k}}[x_1 x_2 (x_3 x_4 x_5)] &= \mathbf{x}_1 \mathbf{x}_2 (\lambda \mathbf{m}. \mathcal{E}_{\mathbf{k}}[\mathbf{m} [x_3 x_4 x_5]]) \\ &= \mathbf{x}_1 \mathbf{x}_2 (\lambda \mathbf{m}. \mathbf{x}_3 \mathbf{x}_4 (\lambda \mathbf{m}. \mathcal{E}_{\mathbf{k}}[\mathbf{m} [x_5]])) \end{aligned}$$

The result is multiple  $\mathbf{m}$  variables present in the same evaluation context, which would clash in the result of the transform if not appropriately renamed.

As Minamide & Okuma state, this clash is not clear from the definition of the optimised CPS transform presented by Danvy & Nielsen; our transform makes it clear by how the introduced variables interact with continuations. This clarity is important, as it enables us to thoughtfully design fresh variable name generation which simplifies the simulation proof.

The transform of Fig. 6 introduces new variables in three places. The first, the lambda abstraction  $\lambda x.e$  case for  $\mathcal{E}$ , introduces  $\mathbf{k}$ , which never clashes because the transform only ever considers at most one  $\mathbf{k}$  at a time, in the subscript of the transform functions. The second, the **arg** continuation  $E^*[\mathbf{t} \square]$  case for  $\mathcal{K}$ , introduces  $\mathbf{n}$ , which never clashes because it is immediately dispatched after its introduction, rather than stashed in the transform parameters for later use. The third, the **fn** continuation  $E^*[\square] e$  case for  $\mathcal{K}$ , introduces  $\mathbf{m}$ , which is stashed in the evaluation context  $E^*[\mathbf{m} \square]$  passed to the next  $\mathcal{E}$  transform, and hence may only clash with other  $\mathbf{m}$  variables present in that evaluation context. What is important is that, when transforming a lambda abstraction  $\lambda x.e$ , we transform its body without the surrounding evaluation context of the abstraction. Intuitively, this means that  $\mathbf{m}$  variables may only clash within a particular lexical scope. Consequently, variable naming need only depend on the evaluation context of the current scope, and the names of these introduced variables may also be determined uniquely from a program fragment in simulation. We will use this fact to build a simulation relation that can determine the name of introduced variables without depending on an external variable counter or allocator.

The function **fresh** on the evaluation context  $E^*$  allocates names that are fresh for the context, the simplest implementation being to count the number of existing  $\mathbf{m}$  variables in the evaluation context  $E^*$ . Including fresh variable

naming, the transform of  $x_1 x_2 (x_3 x_4 x_5)$  becomes:

$$\mathbf{x}_1 \mathbf{x}_2 (\lambda \mathbf{m}_0. \mathbf{x}_3 \mathbf{x}_4 (\lambda \mathbf{m}_1. \mathbf{m}_1 \mathbf{x}_5 (\lambda \mathbf{n}. \mathbf{m}_0 \mathbf{n} \mathbf{k})))$$

With this particular implementation of **fresh**, our transform (Fig. 6) is the same as Danvy & Nielsen's with an equivalent naming scheme (Fig. 5), as seen by considering the expression transform on an expression in an empty evaluation context  $\mathcal{E}_{\mathbf{K}}[e]$ . We prove this property by induction over terms through the mutually recursive definitions of Danvy & Nielsen's transform functions  $\mathcal{E}$  and  $\mathcal{S}$ , with an additional induction lemma establishing that our transform for an application expression with an arbitrary evaluation context  $E^*$  is the same as Danvy & Nielsen's serious term  $\mathcal{S}$  transform with a continuation generated by our  $\mathcal{K}$  transform, and an appropriate name counter index:

$$\mathcal{E}_{\mathbf{K}}[E^*[e_1 e_2]] = \mathcal{S}[e_1 e_2]_i \mathcal{K}_{\mathbf{K}}[E^*] \quad i = \text{fresh } E^*$$

We have mechanised this proof alongside our simulation proof.

## 4 Adjusting the semantics

Having analysed the nature of introduced variable name clashes in the optimised transform, we now turn our attention back to the semantics which inspired the analysis. In order to enable a suitable simulation relation from the side of the source semantics, we divide the evaluation contexts of the source semantics abstract machine into *scopes*, where environments are attached to particular scopes rather than stack frames in an evaluation context.

This division reflects the structure of evaluation contexts in the optimised CPS transform, where each transform function operates on a scoped evaluation context  $E^*$  and an underlying continuation  $\mathbf{K}$  (which may be either the identity continuation  $\mathbf{I}$  or a continuation variable  $\mathbf{k}$ ). New scopes are introduced into CPS by the transform for a lambda abstraction  $\mathcal{E}_{\mathbf{K}}[E^*[\lambda x. e]] = \mathcal{C}_{\mathbf{K}}[E^*] \lambda x \mathbf{k}. \mathcal{E}_{\mathbf{k}}[e]$ , where the expression transform is applied to the lambda body  $e$  with an empty evaluation context, but with an underlying continuation captured by variable  $\mathbf{k}$  which will eventually map to some continuation closure evaluated from  $\mathcal{K}$ .

We hence may formulate evaluation contexts in the source semantics as a nested sequence of scopes, where each scope has a distinct environment  $\rho$  and evaluation context  $E^*$ . At the top level of the semantic state, the abstract machine operates on a sequence of scopes, with the innermost evaluation context containing the control string in its hole  $O(\{\rho\} E^*[c^*])$ . We formalise this notion of scopes and define the adjusted abstract machine semantics in Fig. 7. Each step in the semantics corresponds to a step in the original CEK-machine with the scopes flattened into one evaluation context, except for CBV-CONT-ID which is a CEK-machine no-op.

We equip our new, environment-less evaluation contexts  $E^*$  with a parameter  $\alpha$ , which corresponds to the class of terms that may reside in an **arg** stack frame  $a E^*$ , i.e.  $\alpha$  ranges over the term sets Val, Exp, etc. This parameterised

$$\begin{aligned}
E_\alpha^\star &::= []_\alpha \mid E_\alpha^\star e \mid a E_\alpha^\star && \in \text{Cont}_\alpha^\star, a \in \alpha \\
O &::= () \mid \{\rho\} E_{\text{Val}}^\star O && \in \text{Scope} \\
\\
p &::= \{\rho\} E_{\text{Val}}^\star && \in \text{EnvCont} \\
c^\star &::= e \mid v && \in \text{Ctrl}^\star \\
\\
\text{Scope}(\text{EnvCont}[\text{Ctrl}^\star]) &\longrightarrow^\star \text{Scope}(\text{EnvCont}[\text{Ctrl}^\star]) \\
\\
O(\{\rho\} E^\star[x]) &\longrightarrow^\star O(\{\rho\} E^\star[\rho(x)]) && (\text{CBV-EXP-VAR}) \\
O(\{\rho\} E^\star[\lambda x. e]) &\longrightarrow^\star O(\{\rho\} E^\star[\Lambda x. \{\rho\} e]) && (\text{CBV-EXP-LAM}) \\
O(\{\rho\} E^\star[e_1 e_2]) &\longrightarrow^\star O(\{\rho\} E^\star[[e_1] e_2]) && (\text{CBV-EXP-APP}) \\
\\
O(\{\rho\} E^\star[[v] e]) &\longrightarrow^\star O(\{\rho\} E^\star[v [e]]) && (\text{CBV-CONT-FN}) \\
O(\{\rho\} E^\star[\Lambda x. \{\rho'\} e [v]]) &\longrightarrow^\star O(\{\rho\} E^\star(\{\rho' [x \mapsto v]\} [e])) && (\text{CBV-CONT-ARG}) \\
\\
O(\{\rho'\} E^\star(\{\rho\} [v])) &\longrightarrow^\star O(\{\rho'\} E^\star[v]) && (\text{CBV-CONT-ID})
\end{aligned}$$

**Fig. 7.** Our adjusted semantics, defined on sequences of scopes

evaluation context allows us to share the same evaluation context construction with the evaluation-context-based optimised transform, by instantiating it with **Triv**, the set of trivial terms in the target domain, rather than **Val**, the set of simple values used in the new source semantics. We may hence preserve our optimised transform definition from Fig. 6 and assign the transform functions types based on **Cont**, with types representing terms from the target domain in **bold** to differentiate from source terms:

$$\begin{aligned}
\mathcal{E}_K &: \text{Cont}_{\mathbf{Triv}}^\star[\text{Exp}] \rightarrow \mathbf{Exp} \\
\mathcal{C}_K &: \text{Cont}_{\mathbf{Triv}}^\star \times \mathbf{Triv} \rightarrow \mathbf{Exp} \\
\mathcal{K}_K &: \text{Cont}_{\mathbf{Triv}}^\star \rightarrow \mathbf{Triv}
\end{aligned}$$

We will omit  $\alpha$  for the evaluation context  $E^\star$  when it is clear from the context.

## 5 Simulation relation

Using our adjusted semantics, we construct a simulation relation between semantic abstract machine states and CPS terms built from the optimised transform. The relation is defined inductively over components of the semantic state.

*Values and trivial terms.* The aspect of the optimised transform that enables it to eliminate administrative redexes is that it directly substitutes trivial terms into continuation bodies, which corresponds to statically applying a continuation

to a trivial term at the point of transform through  $\mathcal{C}_K[[E]]t$ . However, this static substitution amounts to deferring of evaluation of these trivial terms until they are used in a serious computation, whereas the corresponding reductions in the source semantics involve the value evaluated from these terms. We must therefore relate values from the source semantics to corresponding trivial terms in CPS to prove simulation.

$$\begin{array}{c}
\frac{\rho \sim \rho}{\Lambda x. \{\rho\}e \sim \Lambda x. \{\rho\} \lambda k. \mathcal{E}_k[[e]]} \text{CPS-VAL} \quad \frac{\forall x \in \text{dom}(\rho). \rho(x) \sim \rho(x)}{\rho \sim \rho} \text{CPS-ENV} \\
\\
\frac{v \sim \rho(x)}{v \rightsquigarrow \{\rho\}x} \text{SIM-TRIV-VAR} \quad \frac{\rho \sim \rho}{\Lambda x. \{\rho\}e \rightsquigarrow \{\rho\} \lambda x k. \mathcal{E}_k[[e]]} \text{SIM-TRIV-LAM}
\end{array}$$

**Fig. 8.** Value relations.

Fig. 8 defines our simulation relations for values from the source semantics. The  $\sim$  relation, which we call *CPS equivalence*, relates equivalent semantic constructs between the source and target semantics. The CPS-VAL rule relates closures to CPS closures, ensuring that their respective captured environments are also related. The corresponding environment relation rule CPS-ENV requires all values mapped to by variables in the source semantics to be related to an equivalently mapped value in the target semantics.

The simulation relation  $\rightsquigarrow$ , which may be read as ‘simulated by’, relates source semantics constructs to the target CPS terms that simulate those constructs. For values, the simulation relation is the relation described above between source semantic values and target semantic trivial terms. Whereas the CPS equivalence relation  $\sim$  on values is primarily used to ensure variables appropriately map to CPS values that have already been evaluated and bound, the simulation relation  $\rightsquigarrow$  is used for values in the control string. Depending on how the value was introduced, the relevant trivial term is either a variable by SIM-TRIV-VAR, or a literal lambda abstraction by SIM-TRIV-LAM. The simulation relation  $\rightsquigarrow$  includes the target semantics environment, in order to establish the conditions for the simulating terms to evaluate to the correct values. It is trivial to prove the lemma that these trivial terms evaluate under the included environment to values satisfying the CPS equivalence relation:

**Lemma 1 (Simulated values evaluate to CPS-equivalent values).**

$$\vdash \forall v \rho t v. v \rightsquigarrow \{\rho\}t \implies (\{\rho\}t \longrightarrow v \wedge v \sim v)$$

*Evaluation contexts.* Our eventual simulation relation on semantic states is built on the transform functions  $\mathcal{E}$ ,  $\mathcal{C}$ , and  $\mathcal{H}$  from Fig. 6, and these functions operate on evaluation contexts parameterised by trivial terms in CPS  $E_{\text{Triv}}^*$ . We must therefore define a relation between the value-parameterised evaluation contexts

$$\begin{array}{c}
\frac{}{\boxed{\ }_{\text{Val}} \rightsquigarrow \{\rho\} \boxed{\ }_{\text{Triv}}} \text{SIM-CONT-ID} \quad \frac{E_{\text{Val}}^* \rightsquigarrow \{\rho\} E_{\text{Triv}}^*}{E_{\text{Val}}^*[\boxed{\ } e] \rightsquigarrow \{\rho\} E_{\text{Triv}}^*[\boxed{\ } e]} \text{SIM-CONT-FN} \\
\\
\frac{E_{\text{Val}}^* \rightsquigarrow \{\rho\} E_{\text{Triv}}^* \quad v \rightsquigarrow \{\rho\} t}{E_{\text{Val}}^*[v \boxed{\ }] \rightsquigarrow \{\rho\} E_{\text{Triv}}^*[t \boxed{\ }]} \text{SIM-CONT-ARG-TRIV} \\
\\
\frac{E_{\text{Val}}^* \rightsquigarrow \{\rho\} E_{\text{Triv}}^* \quad v \rightsquigarrow \rho(\mathbf{m}_i) \quad i = \text{fresh } E_{\text{Triv}}^*}{E_{\text{Val}}^*[v \boxed{\ }] \rightsquigarrow \{\rho\} E_{\text{Triv}}^*[\mathbf{m}_i \boxed{\ }]} \text{SIM-CONT-ARG-COMP}
\end{array}$$

**Fig. 9.** Evaluation context relation.

of the source semantics  $E_{\text{Val}}^*$  and the trivial CPS term-parameterised evaluation contexts used by the transform (Fig. 9).

The empty context and **fn** stack frames  $E^*[\boxed{\ } e]$  trivially relate to their counterparts, as they do not contain any values. For **arg** stack frames  $E^*[v \boxed{\ }]$ , there are two possibilities for the stashed trivial term. The **SIM-CONT-ARG-TRIV** rule captures the case that the value was evaluated from a trivial term, and the resultant **arg** context  $E^*[v \boxed{\ }]$  arises from a static application of a **fn** continuation to a trivial term  $\mathcal{C}_{\mathbf{K}}[E^*[\boxed{\ } e]] t$ .

The **SIM-CONT-ARG-COMP** rule instead captures the case that the value resulted from a computation, i.e. an application. In this case, it is a dynamic application of a **fn** continuation given by the closure evaluated from  $\mathcal{K}_{\mathbf{K}}[E^*[\boxed{\ } e]]$  which gives rise to the resulting **arg** continuation. The computed value is bound to the variable  $\mathbf{m}_i$  and, importantly, only depends on the context of the local scope  $E^*$ , as per the definition of  $\mathcal{K}$ .

Note that, though the simulation relation  $\rightsquigarrow$  usually relates source semantic constructs to CPS terms with an attached environment, the relation on evaluation contexts only goes halfway, relating to similar evaluation contexts but with CPS terms in their value slots. It is the CPS transform applied to these evaluation contexts that produces the final simulating CPS terms, as we show in the following sections.

*Scopes.* Consider how the value bound to  $\mathbf{m}_i$  in the relevant case of the evaluation context relation above arises, as the result of a computation. In the optimised transform, such values are bound to introduced variables such as  $\mathbf{n}$  and  $\mathbf{m}_i$  in continuation closures evaluated from  $\mathcal{K}$ , and arise from the application of an underlying continuation  $\mathbf{k}$  being applied to a trivial term from the definition of  $\mathcal{C}$  for an empty evaluation context  $\mathcal{C}_{\mathbf{k}}[\boxed{\ } t = \mathbf{k} t$ . Note, then, that this application of an underlying continuation must correspond to the case in the source semantics of an empty evaluation context with an underlying scope:

$$O(\{\rho'\} E^*(\{\rho\}[v])) \longrightarrow^* O(\{\rho'\} E^*[v])$$

hence it is *changing scope* which results in a value becoming a computation. The **SIM-SCOPE-CONT** rule in Fig. 10 captures this correspondence by relating scopes to underlying continuation closures evaluated from  $\mathcal{K}$  which are bound to  $\mathbf{k}$ .

$$\begin{array}{c}
\overline{() \rightsquigarrow \{\rho\} \mathbf{I}} \quad \text{SIM-SCOPE-ID} \\
\\
\frac{O \rightsquigarrow \{\rho'\} K \quad \rho' \sim \rho' \quad \frac{E_{\text{Val}}^* \rightsquigarrow \{\rho'\} E_{\text{Triv}}^* \quad \{\rho'\} \mathcal{K}_K \llbracket E_{\text{Triv}}^* \rrbracket \longrightarrow \rho(k)}{\text{SIM-SCOPE-CONT}}}{O(\{\rho'\} E_{\text{Val}}^*(\cdot)) \rightsquigarrow \{\rho\} k}
\end{array}$$

**Fig. 10.** Scoping relation.

*Semantic state relation.* Combining the relations covered in this section, we arrive at a full simulation relation which relates source semantic states to CPS terms and an appropriate environment to simulate those semantic states. The relation is given in Fig. 11.

$$\begin{array}{c}
\frac{O \rightsquigarrow \{\rho\} K \quad \rho \sim \rho \quad \frac{E_{\text{Val}}^* \rightsquigarrow \{\rho\} E_{\text{Triv}}^*}{\text{SIM-EXP}}}{O(\{\rho\} E_{\text{Val}}^*[e]) \rightsquigarrow \{\rho\} \mathcal{E}_K \llbracket E_{\text{Triv}}^*[e] \rrbracket} \\
\\
\frac{O \rightsquigarrow \{\rho\} K \quad \rho \sim \rho \quad \frac{E_{\text{Val}}^* \rightsquigarrow \{\rho\} E_{\text{Triv}}^* \quad v \rightsquigarrow \{\rho\} t}{\text{SIM-TRIV}}}{O(\{\rho\} E_{\text{Val}}^*[v]) \rightsquigarrow \{\rho\} \mathcal{C}_K \llbracket E_{\text{Triv}}^* \rrbracket t} \\
\\
\frac{O(\{\rho'\} E_{\text{Val}}^*(\cdot)) \rightsquigarrow \{\rho\} k \quad v \rightsquigarrow \{\rho\} t}{\text{SIM-COMP}}}{O(\{\rho'\} E_{\text{Val}}^*[v]) \rightsquigarrow \{\rho\} k t}
\end{array}$$

**Fig. 11.** Semantic state relation.

The resultant CPS expressions of the simulation relation are constructed from the definitions of the optimised CPS transform that we began with, each rule corresponding to each transform function. The first, SIM-EXP, relates an expression to its CPS transform by  $\mathcal{E}$  in the corresponding evaluation context. The last two rules relate values to either their static application by  $\mathcal{C}$  (SIM-TRIV) or dynamic application (SIM-COMP) by the closure evaluated from  $\mathcal{K}$ .

We choose to include SIM-COMP rather than attempt to use  $\mathcal{C}$  with computed values as well as trivial ones, because it allows us to always relate the control string value to a directly related trivial term by the value simulation relation instead of having to also consider the possibility of binding the related value to an introduced variable such as  $\mathbf{n}$ . As a result, this also allows us to ignore introduced variables if they are not stashed in the evaluation context, i.e. we never need to consider the presence of  $\mathbf{n}$  in the relation.

## 6 Simulation proof

**Theorem 1 (Simulation).**

$$\begin{aligned}
& \vdash \forall O O' \rho \rho' E^* E^{*'} c^* c^{*'} \rho e. \\
& O(\{\rho\}E^*[c^*]) \longrightarrow^* O'(\{\rho'\}E^{*'}[c^{*'}]) \wedge O(\{\rho\}E^*[c^*]) \rightsquigarrow \{\rho\}e \\
& \implies \\
& \exists n \rho' e' . \{\rho\}e \longrightarrow^n \{\rho'\}e' \wedge O'(\{\rho'\}E^{*'}[c^{*'}]) \rightsquigarrow \{\rho'\}e'
\end{aligned}$$

*Static reductions.* The majority of the cases of Theorem 1 for small steps in the source semantics result in no reductions in the target semantics, which is consistent with the nature of the optimised transform. For example, the CBV-EXP-VAR and CBV-EXP-LAM steps for evaluation of trivial terms to values correspond only to expansions from the expression transform  $\mathcal{E}_{\mathbf{K}}[E^*[t]]$  to the static continuation application transform  $\mathcal{C}_{\mathbf{K}}[E^*]t$  where the source trivial term  $t$  evaluates to a value which is simulated by the target trivial term  $v \rightsquigarrow \{\rho\}t$ . In effect, these steps which do not involve dynamic reductions instead correspond to static reductions.

Another case where there are no reductions in the target semantics is for the CBV-CONT-ID reduction, which is the application to a value of an empty evaluation context with an underlying scope, as the simulation relation explicitly captures the introduction of a value as a computation by changing scope. The semantic state  $O(\{\rho'\}E^*(\{\rho\}[v]))$  relates either to  $\{\rho\}\mathcal{C}_{\mathbf{k}}[[]]t$  which simply expands  $\{\rho\}kt$  without reduction, or to  $\{\rho^+\}kt$  with a different environment which maps  $\mathbf{k}$  to the evaluation of  $\mathcal{K}_{\mathbf{k}}[[]]$  under  $\rho$ , which is again just whatever  $\mathbf{k}$  maps to under  $\rho$ , hence we similarly arrive at  $\{\rho\}kt$ .

*Computations.* The simulation proof becomes interesting when considering a dynamic application of a continuation to a value, i.e. the value was a computation. The starting point for such a step is from a semantic state under the SIM-COMP rule, where the continuation closure in  $\mathbf{k}$  is applied to term. Let us consider the case for the CBV-CONT-FN step, where the evaluation context is  $E^*[[]e]$ ; by applying the corresponding closure in  $\mathbf{k}$ , we arrive at the term  $\mathcal{E}[E^*[\mathbf{m}_i[e]]]$  where the value that the continuation was applied to is bound to  $\mathbf{m}_i$ .

To satisfy the resultant simulation relation, we must ensure that the environment with the newly bound variable  $\rho[\mathbf{m}_i \mapsto v]$  is consistent with the conditions on it imposed by the subrelations of the simulation relation, i.e. we need to prove that various relations on values, evaluation contexts, etc. are monotonic with respect to the additional binding to the introduced variable  $\mathbf{m}_i$ . For the environment CPS equivalence and value simulation relations, this is simple to prove if the transform ensures that variables from the original semantics are always distinct from introduced variables such as  $\mathbf{m}_i, \mathbf{k}$ , etc.

**Lemma 2 (Environment relation monotonicity).**

$$\vdash \forall \rho \rho x v . (\forall x . x \neq x) \implies \rho \sim \rho \implies \rho \sim \rho[x \mapsto v]$$

**Lemma 3 (Value simulation monotonicity).**

$$\vdash \forall v \rho t x v. (\forall x. x \neq x) \Rightarrow v \rightsquigarrow \{\rho\}t \Rightarrow v \rightsquigarrow \{\rho[x \mapsto v]\}t$$

*Proof.* Trivial.

Monotonicity for the scoping relation is similarly trivial, as that relation depends only on the binding of  $k$ .

**Lemma 4 (Scoping relation monotonicity).**

$$\vdash \forall O \rho K x v. x \neq k \Rightarrow O \rightsquigarrow \{\rho\}K \Rightarrow O \rightsquigarrow \{\rho[x \mapsto v]\}K$$

*Proof.* Trivial.

The interesting monotonicity lemma, and really the pin holding together simulation for an optimised transform, is the monotonicity lemma for evaluation contexts. Not only must introduced variables not clash with source variables, but they must also not clash with the *other* introduced variables in the evaluation contexts, i.e. other  $m_i$  variables as per the SIM-CONT-ARG-COMP rule. Using the implementation of *fresh* described in Section 3.4, which counts the number of  $m$  variables in  $E^*$ , the variable index must simply be at least *fresh*  $E^*$  to not clash. We then arrive at the evaluation context lemma of monotonicity:

**Lemma 5 (Evaluation context relation monotonicity).**

$$\begin{aligned} &\vdash \forall E_{\text{Val}}^* \rho E_{\text{Triv}}^* x v. \\ &\quad (\forall x. x \neq x) \wedge (\forall i < \text{fresh } E_{\text{Triv}}^*. x \neq m_i) \\ &\quad \Rightarrow E_{\text{Val}}^* \rightsquigarrow \{\rho\}E_{\text{Triv}}^* \Rightarrow E_{\text{Val}}^* \rightsquigarrow \{\rho[x \mapsto v]\}E_{\text{Triv}}^* \end{aligned}$$

*Proof.* By induction over the SIM-CONT rules.

This lemma is the crux of our proof. The deterministic nature of the introduced variable names in the simulation relation, as a result of the local evaluation-context-based *fresh* name generation in our reformulated transform, enables us to prove monotonicity of the simulation relation with respect to the binding of introduced variables, even with multiple  $m$  variables already introduced into the scope.

*Application.* The CBV-CONT-ARG step, on the evaluation context  $E^*[v \ ]]$ , always produces reductions in the target domain, regardless of if the continuation is applied to a trivial or computed term. Indeed, to simulate the source semantics, a completed reduction of an application expression in the source domain must correspond to a reduction of an application in the target domain; the transform produces the application  $t_1 t_2 \mathcal{K}_K[E^*]$  in CPS which must be reduced.

By the simulation relation,  $t_1$  will always evaluate to a closure  $\Lambda x. \{\rho\} \lambda k. \mathcal{E}_k[e]$ , hence the application leads to the expression  $\mathcal{E}_k[e]$  with  $x$  and  $k$  freshly bound to the passed value and underlying continuation evaluated from  $\mathcal{K}_K[E^*]$  respectively. The corresponding environment in the source semantics similarly has the variable  $x$  bound, and we may trivially prove that the environment relation is preserved when binding a source variable.

**Lemma 6 (Environment relation synchronisation).**

$$\vdash \forall \rho \boldsymbol{\rho} x v \boldsymbol{v}. v \sim \boldsymbol{v} \implies \rho \sim \boldsymbol{\rho} \implies \rho[x \mapsto v] \sim \boldsymbol{\rho}[\boldsymbol{x} \mapsto \boldsymbol{v}]$$

*Proof.* Cases on  $x \in \text{dom}(\rho)$ .

Before Lemma 6 may be applied, we must first apply the environment monotonicity lemma, Lemma 2, to disregard the binding of  $\boldsymbol{k}$ .

Note that, in the case that the continuation was applied to a computed value, the value is bound to  $\boldsymbol{n}$  then very quickly dispatched. We must then also apply the environment, scoping, and evaluation context monotonicity lemmas, Lemmas 2, 4, and 5, to disregard the binding of  $\boldsymbol{n}$ .

## 7 Conclusion

We have presented a reformulation of Danvy & Nielsen’s one-pass CPS transform [9] in the style of Sabry & Felleisen’s optimised transform based on evaluation contexts [21]. Our transform highlights the nature of the name clashes encountered in Minamide & Okuma’s mechanised correctness proof [15]: clashes are local to the lexical scope, and may be reasoned about by the inclusion of the offending variables in the local evaluation context. It is hence possible to use a fresh variable generation method that depends only on the local evaluation context, and allows introduced variable names to be uniquely determined by the simulation relation.

We have proposed a modified CEK-machine to represent the source semantics, with evaluation contexts rephrased as nested sequences of lexical scopes, each with its own local environment and evaluation context. In combination with our new transform, this semantics builds naturally into a simulation relation which can determine introduced variable names, that consequently enables a straightforward simulation proof, with no need to consider  $\alpha$ -equivalence. The modified CEK-machine represents only a small departure from a typical CEK-machine or alternative semantics and is easy to prove equivalent, hence we believe that it is a beneficial trade-off for the more difficult proof of  $\alpha$ -equivalence.

The nature of our transform to use evaluation contexts from the semantics should allow it to be applied more generally to richer languages. We would particularly like to investigate the use of an evaluation-context-based CPS transform to languages with control operators such as `call/cc` or delimited continuations, which would likely have interesting variable clashing and fresh naming challenges for captured continuation variables.

**Acknowledgments** We thank the PADL 2026 reviewers for helpful comments.

## References

1. Ager, M.S., Biernacki, D., Danvy, O., Midtgaard, J.: A functional correspondence between evaluators and abstract machines. In: Proceedings of the 5th ACM SIG-

- PLAN International Conference on Principles and Practice of Declarative Programming, p. 8–19. PPDP '03, Association for Computing Machinery, New York, NY, USA (2003). <https://doi.org/10.1145/888251.888254>
2. Baker, H.G.: Cons should not cons its arguments, part ii: Cheney on the M.T.A. SIGPLAN Not. **30**(9), 17–20 (Sep 1995). <https://doi.org/10.1145/214448.214454>
  3. Chlipala, A.: Parametric higher-order abstract syntax for mechanized semantics. In: Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming, p. 143–156. ICFP '08, Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1411204.1411226>
  4. Danvy, O., Filinski, A.: Representing control: a study of the CPS transformation. Mathematical Structures in Computer Science **2**(4), 361–391 (1992). <https://doi.org/10.1017/S0960129500001535>
  5. Danvy, O.: On evaluation contexts, continuations, and the rest of computation (02 2004)
  6. Danvy, O.: Defunctionalized interpreters for programming languages. In: Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming, p. 131–142. ICFP '08, Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1411204.1411206>
  7. Danvy, O., Filinski, A.: Abstracting control. In: Kahn, G. (ed.) Proceedings of the 1990 ACM Conference on LISP and Functional Programming, LFP 1990, Nice, France, 27–29 June 1990. pp. 151–160. ACM (1990). <https://doi.org/10.1145/91556.91622>, <https://doi.org/10.1145/91556.91622>
  8. Danvy, O., Nielsen, L.R.: Defunctionalization at work. In: Proceedings of the 3rd ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming, p. 162–174. PPDP '01, Association for Computing Machinery, New York, NY, USA (2001). <https://doi.org/10.1145/773184.773202>
  9. Danvy, O., Nielsen, L.R.: A first-order one-pass CPS transformation. Theoretical Computer Science **308**(1), 239–257 (2003). [https://doi.org/10.1016/S0304-3975\(02\)00733-8](https://doi.org/10.1016/S0304-3975(02)00733-8)
  10. Dargaye, Z., Leroy, X.: Mechanized verification of CPS transformations. In: Proceedings of the 14th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, p. 211–225. LPAR'07, Springer-Verlag, Berlin, Heidelberg (2007)
  11. Felleisen, M., Friedman, D.P.: Control operators, the SECD-machine, and the  $\lambda$ -calculus. In: Formal Description of Programming Concepts (1987), <https://api.semanticscholar.org/CorpusID:57760323>
  12. Kelsey, R.A.: A correspondence between continuation passing style and static single assignment form. In: Papers from the 1995 ACM SIGPLAN Workshop on Intermediate Representations, p. 13–22. IR '95, Association for Computing Machinery, New York, NY, USA (1995). <https://doi.org/10.1145/202529.202532>
  13. Kranz, D., Kelsey, R., Rees, J., Hudak, P., Philbin, J., Adams, N.: Orbit: an optimizing compiler for Scheme. In: Proceedings of the 1986 SIGPLAN Symposium on Compiler Construction, p. 219–233. SIGPLAN '86, Association for Computing Machinery, New York, NY, USA (1986). <https://doi.org/10.1145/12276.13333>
  14. Lasnier, P.Y., Yallop, J., Myreen, M.O.: BRACK: A verified compiler for Scheme via CakeML. In: Proceedings of the 15th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP '26, Association for Computing Machinery, New York, NY, USA (2026)

15. Minamide, Y., Okuma, K.: Verifying CPS transformations in Isabelle/HOL. In: Proceedings of the 2003 ACM SIGPLAN Workshop on Mechanized Reasoning about Languages with Variable Binding. p. 1–8. MERLIN '03, Association for Computing Machinery, New York, NY, USA (2003). <https://doi.org/10.1145/976571.976576>
16. Paraskevopoulou, Z., Grover, A.: Compiling with continuations, correctly. Proc. ACM Program. Lang. 5(OOPSLA) (Oct 2021). <https://doi.org/10.1145/3485491>
17. Pitts, A.M.: Nominal logic, a first order theory of names and binding. Inf. Comput. 186(2), 165–193 (2003). [https://doi.org/10.1016/S0890-5401\(03\)00138-X](https://doi.org/10.1016/S0890-5401(03)00138-X), [https://doi.org/10.1016/S0890-5401\(03\)00138-X](https://doi.org/10.1016/S0890-5401(03)00138-X)
18. Plotkin, G.: Call-by-name, call-by-value and the  $\lambda$ -calculus. Theoretical Computer Science 1(2), 125–159 (1975). [https://doi.org/10.1016/0304-3975\(75\)90017-1](https://doi.org/10.1016/0304-3975(75)90017-1)
19. Pottier, F.: Revisiting the CPS transformation and its implementation (2006)
20. Reynolds, J.C.: Definitional interpreters for higher-order programming languages. In: Proceedings of the ACM Annual Conference - Volume 2. p. 717–740. ACM '72, Association for Computing Machinery, New York, NY, USA (1972). <https://doi.org/10.1145/800194.805852>
21. Sabry, A., Felleisen, M.: Reasoning about programs in continuation-passing style. In: Proceedings of the 1992 ACM Conference on LISP and Functional Programming. p. 288–298. LFP '92, Association for Computing Machinery, New York, NY, USA (1992). <https://doi.org/10.1145/141471.141563>
22. Steele, G.L.: Rabbit: A compiler for Scheme. Tech. rep., USA (1978)
23. Tian, Y.H.: Mechanically verifying correctness of CPS compilation. In: Proceedings of the Twelfth Computing: The Australasian Theory Symposium - Volume 51. p. 41–51. CATS '06, Australian Computer Society, Inc., AUS (2006)
24. Xiao, Y., Sabry, A., Ariola, Z.M.: From syntactic theories to interpreters: Automating the proof of unique decomposition. Higher Order Symbol. Comput. 14(4), 387–409 (Dec 2001). <https://doi.org/10.1023/A:1014408032446>