JFP **35**, e22, 68 pages, 2025. © The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

doi:10.1017/S0956796825100117

Tail recursion modulo context: An equational approach (extended version)

DAAN LEIJEN
Microsoft Research, Redmond, WA 98052, USA
(*e-mail:* daan@microsoft.com)

ANTON FELIX LORENZEN
University of Edinburgh, Edinburgh EH8 9YL, UK
(*e-mail:* anton.lorenzen@ed.ac.uk)

Abstract

The tail recursion modulo *cons* transformation can rewrite functions that are not quite tail-recursive into a tail-recursive form that can be executed efficiently. In this article, we generalize tail recursion modulo *cons* (TRMc) to modulo *context* (TRMC) and calculate a general TRMC algorithm from its specification. We can instantiate our general algorithm by providing an implementation of application and composition on abstract contexts and showing that our *context laws* hold. We provide some known instantiations of TRMC, namely modulo *evaluation contexts* (CPS), and *associative operations*, and further instantiations not so commonly associated with TRMC, such as *defunctionalized* evaluation contexts, *monoids*, *semirings*, *exponents*, and *fields*. We study the modulo *cons* instantiation in particular and prove that an instantiation using Minamide's hole calculus is sound. We also calculate a second instantiation in terms of the Perceus heap semantics to precisely reason about the soundness of in-place update. While all previous approaches to TRMc fail in the presence of nonlinear control (e.g., induced by call/cc, shift/reset, or algebraic effect handlers), we can elegantly extend the heap semantics to a hybrid approach which dynamically adapts to nonlinear control flow. We have a full implementation of hybrid TRMc in the Koka language, and our benchmark shows the TRMc transformed functions are always as fast or faster than using manual alternatives.

1 Introduction

The tail recursion modulo *cons* (TRMc) transformation can rewrite functions that are not quite tail-recursive into a tail-recursive form that can be executed efficiently. This transformation was described already in the early 70s by Risch (1973) and Friedman & Wise (1975) and more recently studied by Bour *et al*. (2021) in the context of OCaml. A prototypical example of a function that can be transformed this way is `map`, which applies a function to every element of a list:

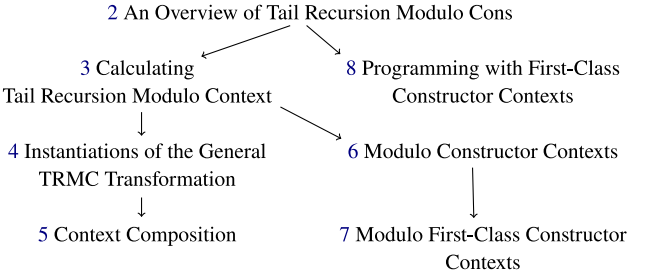
```
fun map( xs : list<a>, f : a -> b ) : list<b>
  match xs
    Cons(x,xx) -> Cons( f(x), map(xx,f) )
    Nil        -> Nil
```

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

We can see that the recursive call to `map` is behind a constructor, and thus `map` as written is not tail-recursive and uses stack space linear in the length of the list. Of course, it is well known that we can rewrite `map` by hand into a tail-recursive form by using an extra accumulating argument, but this comes at the cost of losing the simplicity of the original definition.

The TRMc transformation can automatically transform a function like `map` to a tail-recursive variant but also improves on the efficiency of the manual version by using in-place updates on the accumulation argument.

In the following sections, we formalize our calculus and calculate a general tail recursion modulo *contexts* algorithm (Section 3) that we then instantiate to various use cases (Section 4 and Section 5), and in particular we study the efficient modulo *cons* instantiation (Section 6), its extension to nonlinear control (Section 7), the user-facing *first-class constructor context* feature and finally conclude with benchmarks (Section 9) and related work. Readers may choose to read this article selectively in the following order:



1.1 Contributions

This paper is the extended version of Leijen & Lorenzen (2023). In previous work (Risch, 1973; Friedman & Wise, 1975; Bour *et al*., 2021), TRMc algorithms are given but all fall short of showing why these are correct, or provide particular insight in what other transformations may be possible. In this article, we generalize tail recursion modulo *cons* (TRMc) to modulo *context* (TRMC) and try to bring the general principles out of the shadows of particular implementations and into the light of equational reasoning.

- Inspired by the elegance of program calculation as pioneered by Bird (1984), Gibbons (2022), Hutton (2021), Meertens (1986), and many others, we take an equational approach where we *calculate* a general tail recursion modulo context transformation from its specification and two general *context laws*. The resulting generic algorithm is concise and independent of any particular instantiation of the abstract contexts as long as their operations satisfy the context laws (Section 3).
- We can instantiate the algorithm by providing an implementation of application and composition on abstract contexts and show that these satisfy the context laws. In Section 4, we provide known instantiations of TRMC, namely modulo *evaluation contexts* (CPS) and modulo *associative operations*, and show that those instances satisfy the context laws. We then proceed to show various instantiations not so commonly associated with TRMC that arise naturally in our generic approach, namely modulo

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

defunctionalized evaluation contexts, modulo *monoids*, modulo *semirings*, and modulo *exponents*.

- In Section 6, we turn to the most important instance in practice, modulo *cons*. We show how we can instantiate our operations to the hole calculus of Minamide (1998), and that this satisfies the context laws and the imposed linear typing discipline. This gives us an elegant and sound in-place updating characterization of TRMc where the in-place update is hidden behind a purely functional (linear) interface.
- This is still somewhat unsatisfying as it does not provide insight in the actual in-place mutation as such implementation is only alluded to in prose (Minamide, 1998). We proceed by giving a second instantiation of modulo *cons* where we target the heap semantics of Xie *et al*. (2021) to be able to reason explicitly about the heap and in-place mutation. Just like we could calculate the generic TRMC translation from its specification, we again *calculate* the efficient in-place updating versions for context application and composition from the abstract context laws. These calculated reductions are exactly the implementation as used in our Koka compiler.
- A well-known problem with the modulo *cons* transformation is that the efficient in-place mutating implementation fails if the semantics is extended with non-local control operations, like `call/cc`, `shift/reset` (Danvy & Filinski, 1990; Sitaram & Felleisen, 1990; Shan, 2007), or general algebraic effect handlers (Plotkin & Power, 2003; Plotkin & Pretnar, 2009), where one can resume more than once. This is in particular troublesome for a language like Koka which relies foundationally on algebraic effect handlers (Leijen, 2017; Xie & Leijen, 2021). In Section 7, we show two novel solutions to this: The general approach generates two versions for each TRMc translation and chooses at runtime the appropriate version depending on whether nonlinear control is possible. This duplicates code though and may be too pessimistic where the slow version is used even if no nonlinear control actually occurs. Suggested by our heap semantics, we can do better though – in the *hybrid* approach we rely on the precise reference counts (Xie *et al*., 2021), together with runtime support for *context paths*. This way we can efficiently detect at runtime if a context is unique and fall back to copying only if required due to nonlinear control.
- We have fully implemented the hybrid TRMc approach in the Koka compiler, and our benchmarks show that this approach can be very efficient. We measure various variants of modulo *cons* recursive functions and for linear control the TRMc transformed version is always faster than alternative approaches (Section 9).

In this version, we make the following contributions over the conference version:

- We extend the TRMC algorithm to ensure that (when instantiated to general evaluation contexts) it can optimize all recursive calls that are not under a lambda (Section 3). In contrast, the algorithm presented in the conference paper could only achieve this if the source program was in *A-normal form* (Flanagan *et al*., 1993). Our new algorithm extends the previous algorithm to perform the necessary A-normalizations on-demand.
- We describe a method for composing context instantiations (Section 5). This is especially useful for programs where fast instantiations (like semiring contexts) are not quite good enough to make the program tail-recursive. In that case, we can use the fast instantiation where it applies and use a slower instantiation like defunctionalized contexts for the rest. We use this insight to derive a tail-recursive evaluator for an arithmetic expression evaluator on fields.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

- We include a detailed description of Koka's implementation of constructor contexts. We discuss a snippet of the assembly code generated by the Koka and explain the optimizations that make the implementation efficient (Section 7.3). Furthermore, we describe in detail another implementation strategy proposed by Lorenzen *et al*. (2024), which does not rely on reference counting and contrast it with our implementation.
- Constructor contexts were special in the conference version of this paper, since they were the only contexts for which the transformation could not be done manually. However, Lorenzen *et al*. (2024) show that the hybrid approach can be used to make constructor contexts first-class values. This gives programmers the ability to make their functions tail-recursive manually. We include several practical examples of programming with first-class constructor contexts, where it would be hard to achieve a tail-recursive version fully automatically, but a manual solution is evident.

The new content in this version supersedes several sections of the conference paper. We no longer include "Improving Constructor Contexts" (which is now covered by the extended algorithm in Section 3), "Modulo Cons Products" (which can be achieved more easily using first-class constructor contexts), and "Fall Back to General Evaluation Contexts" (which is less efficient than the implementation proposed in Section 7.4).

2 An overview of tail recursion modulo cons

As shown in the introduction, the prototypical example of a function that can be transformed by TRMc is the `map` function. One way to rewrite the `map` function manually to be tail-recursive is to use CPS where we add a continuation parameter *k*:

```
fun mapk( xs : list<a>, f : a -> b, k : list<b> -> list<b> ) : list<b>
  match xs
    Cons(x,xx) -> val y = f(x) in mapk(xx,f,compose(k, fn(ys) Cons(y,ys)))
    Nil        -> apply(k,Nil)

fun map( xs : list<a>, f : a -> b ) : list<b>
  mapk(xs,f,id)
```

where we have to evaluate `f(x)` before allocating the closure `fn(ys) Cons(y,ys)` since `f` may have an observable (side) effect. The function `id` is the identity function, and `apply` and `compose` regular function application and composition:

```
fun compose( f : b -> c, g : a -> b ) : (a -> c) = fn(x) f(g(x))
fun apply( f : a -> b, x : a ) : b = f(x)
fun id( x : a ) : a = x
```

All our examples use the Koka language (Leijen, 2021) since it has a full implementation of TRMc using the design in this paper, including support for nonlinear control (which cannot be handled by previous TRMc techniques).

We would like to stress though that the described techniques are not restricted to Koka as such and apply generally to any strict programming language (and particular instances can already be found in various compilers, including GCC, see Section 4.6). Some techniques, like the hybrid approach in Section 7.2, may require particular runtime support (like precise reference counts), but this is again independent of the particular language.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

2.1 Continuation style TRMc

Our new tail-recursive version of `map` may not consume any extra stack space, but it achieves this at the cost of allocating many intermediate closures in the heap, that each allocate a `Cons` node for the final result list. The TRMc translation is based on the insight that for many contexts around a tail-recursive call, we can often use more efficient implementations than function composition.

In this paper, we are going to abstract over particular constructor contexts and instead represent abstract program contexts as `ctx<a>` with three operations. First, the `ctx body` expression creates such contexts which can contain a single hole denoted as □; for example, `ctx Cons(1,Cons(2,□)) : ctx<list<int>>`. We can see here that the context type `ctx<a>` is parameterized by the type of the hole `a`, which for our purposes must match the result type as well. Furthermore, we can compose and apply these abstract contexts as:

```
fun comp( k1 : ctx<a>, k2 : ctx<a> ) : ctx<a>
fun app( k : ctx<a>, x : a ) : a
```

Our general TRMC translation can convert a function like `map` automatically to a tail-recursive version by recognizing that each recursive invocation to `map` is under a constant constructor context (Section 6), leading to:

```
fun mapk( xs : list<a>, f : a -> b, k : ctx<list<b>> ) : list<b>
  match xs
    Cons(x,xx) -> val y = f(x) in mapk(xx,f,comp(k,ctx Cons(y,□)))
    Nil        -> app(k,Nil)

fun map( xs : list<a>, f : a -> b ) : list<b>
  mapk(xs,f,ctx □)
```

This is essentially equivalent to our manually translated CPS-style `map` function where we replaced function application and composition with context application and context composition, and the identity function with `ctx` □.

Thus, an obvious way to give semantics to our abstract contexts `ctx<a>` is to represent them as functions `a -> a`, where a context expression is interpreted as a function with a single parameter for the hole, for example, `ctx Cons(1,Cons(2,□))` = `fn(x) Cons(1,Cons(2,x))` (and therefore `ctx` □ = `fn(x) x` = `id`). Context application and composition then map directly onto function application and composition:

```
alias ctx<a> = a -> a
fun comp( k1 : ctx<a>, k2 : ctx<a> ) : ctx<a> = compose(k1,k2)
fun app( k : ctx<a>, x : a ) : a              = apply(k,x)
```

Of course, using such semantics is equivalent to our original manual implementation and does not improve efficiency.

2.2 Linear continuation style

The insight of Risch (1973) and Friedman & Wise (1975) that leads to increased efficiency is to observe that the transformation always uses the abstract context k in a linear way, and we can implement the composition and application by updating the context holes *in-place*. Following the implementation strategy of Minamide (1998) for their hole-calculus, we can represent our abstract contexts as a *Minamide tuple* with a `res` field pointing to the

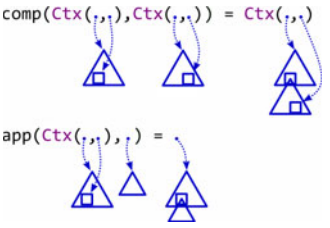
https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

final result object, and a `hole` field which points directly at the field containing the hole inside the result object. Assuming an assignment primitive `(:=)`, we can then implement composition and application efficiently as:

```
value type ctx<a>
  Id
  Ctx( res : a, hole : ptr<a> )

fun comp( k1 : ctx<a>, k2 : ctx<a> ) : ctx<a>
  Ctx( app(k1,k2.res), k2.hole)

fun app( k : ctx<a>, x : a ) : a
  match k
    Id -> x
    Ctx(res,hole) -> { hole := x; res }
```



where the empty `ctx` □ is represented as `Id` (since we do not yet have an address for the `Ctx.hole` field). If we inline these definitions in the `mapk` function, we can see that we end up with a very efficient implementation where each new `Cons` cell is directly appended to the partially build final result list. In our actual implementation, we optimize a bit more by defining the `ctx` type as a value type with only the `Ctx` constructor where we represent the `Id` case with a `hole` field containing a null pointer. Such a tuple is passed at runtime in two registers and leads to efficient code where the `match` in the `app` function, for example, just zero-compares a register (see Section 7.3). Section 9 shows detailed performance figures that show that the TRMc transformation always outperforms alternative implementations (for linear control flow).

2.3 First-class constructor contexts

The constructor context can be implemented using in-place updates since the TRMC conversion guarantees it is used linearly. However, it turns out this also makes the conversion invalid if the host language has advanced control primitives like `call/cc` or general algebraic effect handlers (as in Koka). For example, for the `map` function, the passed in function `f` might resume multiple times. We describe this issue in detail in Section 7.1. As we show in this paper, we can compile constructor contexts in such a way that we are able to detect if a context is used nonlinearly and, in such case, fall back to copying the *context path* at runtime to maintain correct semantics (Section 7.2). This still gives us the in-place update efficiency in the usual case but now also handles nonlinear usage.

Furthermore, this fully encapsulates the imperative implementation of constructor contexts behind a purely functional interface, and we can expose these contexts as first-class values in the language. This allows us to write the result of the TRMc transformation manually.

In Koka, first-class constructor contexts are written as a `ctx` K expression with a single hole denoted by an underscore `_` (Lorenzen *et al.*, 2024). For example, we can write a list constructor context as `ctx Cons(1,_)` or a binary tree constructor context as `ctx Node(Node(Leaf,1,Leaf),2,_)`. The composition operation is written as `(++)`, while application is written as `(++.)`. For example, the expression `(ctx Cons(1,_)) ++ (ctx Cons(2,_)) ++. Nil` evaluates to `(ctx Cons(1,Cons(2,_))) ++. Nil` and then to `[1,2]`. Using these context expressions, we can directly implement the TRMC transformed `map` function in Koka:

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

```
fun map-trmc'( xs : list<a>, f : a -> b, k : ctx<list<b>> ) : list<b>
  match xs
    Cons(x,xx) -> map-trmc'(xx,f, k ++ ctx Cons(f(x),_))
    Nil        -> k ++. Nil

fun map-trmc( xs : list<a>, f : a -> b ) : list<b>
  map-trmc'(xs,f,ctx _)
```

These constructor context are quite efficient (see Lorenzen *et al.*, 2024 and Section 7.3), and the Koka compiler implements the automatic TRMc conversion as a source-to-source transformation using the first-class constructor contexts directly. We give further examples of the programming with first-class constructor contexts in Section 8.

3 Calculating tail recursion modulo context

In order to reason precisely about our transformation, we define a small calculus in Figure 1. The calculus is mostly standard with expressions e consisting of values v, application $e_1\, e_2$, let-bindings, and pattern matches. We assume well-typed programs that cannot go wrong, and where pattern matches are always complete and cannot get stuck. Since we reason in particular over recursive definitions, we add a special environment F of named recursive functions f. We could have encoded this using a fix combinator, but using explicitly named definitions is more convenient for our purposes.

Following the approach of Wright & Felleisen (1994), we define applicative order evaluation contexts E. Generally, contexts are expressions with one subexpression denoted as a hole $\Box$. We write $\mathsf{E}[v]$ for the substitution $\mathsf{E}[\Box := v]$ (which binds tighter than function application). The definition of E ensures a single reduction order where we never evaluate under a lambda. Dually, we define tail contexts T (Abelson *et al.*, 1998). They describe the last term to be evaluated in an expression. Unlike evaluation contexts, there can be several holes in a tail context (say, in different branches of a match-statement) and the substitution $\mathsf{T}[\Box := v]$ is assumed to be capture-avoiding. We also define expression contexts X which match any expression that is not under a lambda.

The operational semantics can now be given using small step reduction rules of the form $e_1 \longrightarrow e_2$ together with the (*step*) rule to reduce in any evaluation context $\mathsf{E}[e_1] \longmapsto \mathsf{E}[e_2]$ (and in essence, an E context is an abstraction of the program stack and registers). We write $\longmapsto^*$ for the reflexive and transitive closure of the $\longmapsto$ reduction relation. The small step operational rules are standard, except for the (*fun*) rule that assumes a global F environment of recursive function definitions.

When $e \longmapsto^* v$, we call e *terminating* (also called *valuable* Harper, 2012). When an evaluation does not terminate, we write $e\Uparrow$. For closed terms, we write $e_1 \cong e_2$ if e_1 and e_2 are *extensionally equivalent*: either $e_1 \longmapsto^* v$ and $e_2 \longmapsto^* v$, or both $e_1\Uparrow$ and $e_2\Uparrow$. For open terms, we write $e_1 \cong e_2$ if $\sigma(e_1) \cong \sigma(e_2)$ for all substitutions σ of free variables by values. During reasoning, we sometimes use the following equalities:

(*letid*)	$\mathsf{let}\, x = e\, \mathsf{in}\, x$	$\cong$	e	
(*letapp*)	$(\lambda x.\, e_1)\, e_2$	$\cong$	$\mathsf{let}\, x = e_2\, \mathsf{in}\, e_1$	
(*letfloat*)	$\mathsf{E}[\mathsf{let}\, x = e_1\, \mathsf{in}\, e_2]$	$\cong$	$\mathsf{let}\, x = e_1\, \mathsf{in}\, \mathsf{E}[e_2]$	if $x \notin \mathsf{fv}(\mathsf{E})$
(*matchfloat*)	$\mathsf{E}[\mathsf{match}\, e\, \{\overline{p_i \to e_i}\}]$	$\cong$	$\mathsf{match}\, e\, \{\overline{p_i \to \mathsf{E}[e_i]}\}$	if $\mathsf{fv}(p_i) \not\cap \mathsf{fv}(\mathsf{E})$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Expressions:

$$
\begin{array}{lcll}
e & ::= & v & \text{(value)} \\
 & | & e\,e & \text{(application)} \\
 & | & \mathsf{let}\, x = e\, \mathsf{in}\, e & \text{(let binding)} \\
 & | & \mathsf{match}\, e\, \{\overline{p_i \mapsto e_i}\} & \text{(matching, } i \geqslant 1\text{)} \\
p & ::= & C^k\, x_1 \ldots x_k & \text{(pattern)}
\end{array}
$$

$$
\begin{array}{lcll}
v & ::= & x, y & \text{(variables)} \\
 & | & f & \text{(recursive functions)} \\
 & | & \lambda x.\, e & \text{(functions)} \\
 & | & C^k\, v_1 \ldots v_n & \text{(constructor of arity } k\text{, with } k \geqslant 0 \text{ and } n \leqslant k\text{)}
\end{array}
$$

$$
F ::= \{\overline{f_i = \lambda x.\, e_i}\} \qquad \text{(recursive definitions)}
$$

Syntax:

$$
\begin{array}{lcl}
f\, x_1 \ldots x_n = e & \doteq & f = \lambda x_1 \ldots x_n.\, e \\
\lambda x_1 \ldots x_n.\, e & \doteq & \lambda x_1 \ldots . \lambda x_n.\, e
\end{array}
$$

Evaluation Contexts:

$$
\mathsf{E} ::= \Box \mid \mathsf{E}\, e \mid v\, \mathsf{E} \mid \mathsf{let}\, x = \mathsf{E}\, \mathsf{in}\, e \mid \mathsf{match}\, \mathsf{E}\, \{\overline{p_i \mapsto e_i}\} \quad \text{(strict, left-to-right)}
$$

Tail Contexts:

$$
\mathsf{T} ::= \Box \mid e\, \mathsf{T} \mid \mathsf{let}\, x = e\, \mathsf{in}\, \mathsf{T} \mid \mathsf{match}\, e\, \{\overline{p_i \mapsto \mathsf{T}_i}\} \quad \text{(tail context)}
$$

Expression Contexts (= Tail Context + Evaluation Context):

$$
\begin{array}{lcl}
\mathsf{X} & ::= & \Box \mid \mathsf{X}\, e \mid e\, \mathsf{X} \mid \mathsf{let}\, x = \mathsf{X}\, \mathsf{in}\, e \mid \mathsf{let}\, x = e\, \mathsf{in}\, \mathsf{X} \\
 & | & \mathsf{match}\, \mathsf{X}\, \{\overline{p_i \mapsto e_i}\} \mid \mathsf{match}\, e\, \{\overline{p_i \mapsto \mathsf{X}_i}\}
\end{array}
$$

Operational Semantics:

$$
\begin{array}{llcl}
(\mathit{let}) & \mathsf{let}\, x = v\, \mathsf{in}\, e & \longrightarrow & e[x{:=}v] \\
(\mathit{beta}) & (\lambda x.\, e)\, v & \longrightarrow & e[x{:=}v] \\
(\mathit{fun}) & f\, v & \longrightarrow & e[x{:=}v] \text{ with } f = \lambda x.\, e \in F \\
(\mathit{match}) & \mathsf{match}\, (C^k\, v_1 \ldots v_k)\, \{\overline{p_i \mapsto e_i}\} & \longrightarrow & e_i[x_1{:=}v_1, \ldots, x_k{:=}v_k] \\
 & & & \text{with } p_i = C^k\, x_1 \ldots x_k
\end{array}
$$

$$
\frac{e_1 \longrightarrow e_2}{\mathsf{E}[e_1] \longmapsto \mathsf{E}[e_2]}\ [\textsc{step}]
$$

Fig. 1. Syntax and operational semantics.

Since the hole in an evaluation context marks the first term to be evaluated, rewriting along these equalities does not change the evaluation order. Rewrites using the last equality may lead to code duplication since they push the evaluation context into each branch (but this can be avoided by inserting join points Maurer *et al.*, 2017).

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

It is straightforward to show these equalities hold, for example, if e_2 is terminating:

$$
\begin{array}{lll}
 & (\lambda x.\, e_1)\, e_2 & \\
\longmapsto^* & (\lambda x.\, e_1)\, v & \{\ e_2 \textit{ terminating}\ \} \\
\longmapsto & e_1[x{:=}v] & \{\ (\textit{beta})\ \} \\
\longleftarrow & \mathsf{let}\, x = v\, \mathsf{in}\, e_1 & \{\ (\textit{let})\ \} \\
{}^*\!\longleftarrow & \mathsf{let}\, x = e_2\, \mathsf{in}\, e_1 & \{\ e_2 \textit{ terminating}\ \}
\end{array}
$$

(and if $e_2 \Uparrow$, then both do not terminate).

3.1 Abstract contexts

Before we start calculating our general TRMC transformation, we first define *abstract contexts* as an abstract type $\mathit{ctx}\,\tau$ in our calculus. There are three context operations: creation (as ctx), application (as app), and composition (as $(\bullet)$). These are not available to the user but instead are only generated as the target calculus of our TRMC translation. We extend the calculus as follows:

$$v \quad ::= \quad \ldots \mid \mathsf{ctx}\,\mathsf{E} \mid _ \bullet _ \mid \mathsf{app}$$

where we assume that the abstract context operations are always terminating. In order to reason about contexts as an abstract type, we assume two context laws. The first one relates the application with the construction of a context:

$$(\textit{appctx}) \quad \mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\, e \qquad = \mathsf{E}[e]$$

The second law states that composition of contexts is equivalent to a composition of applications:

$$(\textit{appcomp}) \quad \mathsf{app}\,(k_1 \bullet k_2)\, e \qquad = \mathsf{app}\, k_1\, (\mathsf{app}\, k_2\, e)$$

When we instantiate to a particular implementation context, we need to show the context laws are satisfied. In such case, we only need to show this for terminating expressions e, since if $e \Uparrow$, the laws hold by definition. In particular, for (*appctx*) it follows directly that $\mathsf{app}\, k\, e \Uparrow$ and $\mathsf{E}[e] \Uparrow$. Of particular note is that the latter only holds for E contexts since they guarantee that the hole is first in the evaluation order and that is one reason why evaluation contexts are the *maximum* context possible for our TRMC translation. Similarly, for (*appcomp*) it follows directly that $(\mathit{app}\,(k_1 \bullet k_2)\, e) \Uparrow$ and $\mathsf{app}\, k_1\, (\mathsf{app}\, k_2\, e) \Uparrow$.

3.2 Calculating a general tail-recursion-modulo-contexts algorithm

In this section, we are going to calculate a general TRMC translation algorithm from its specification. The algorithm is calculated assuming an abstract context where the context laws hold. Eventually, the algorithm needs to be instantiated in the compiler to particular contexts (like constructor contexts), with a particular implementation of context application and composition. We show many such instantiations in Sections 4 and 6.

For clarity, we use single parameter functions for proofs and derivations (but of course the results extend straightforwardly to multiple parameter functions). Now consider a function $f\, x = e_f$ with its TRMC transformed version denoted as f':

$$f'\, x\, k = [\![e_f]\!]_{f,k} \quad (k \notin \mathsf{fv}(e_f))$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Our goal is to calculate the static TRMC transformation algorithm $[\![_]\!]_{f,k}$ from its specification, where f is the function we intend to transform and k is a fresh variable representing the context. The first question is then how we should even specify the intended behavior of such function?

We can follow the standard approach for reasoning about continuation passing style (CPS) here. For example, Gibbons (2022) calculates the CPS version of the factorial function, called fact', from its specification as: $k\,(\mathit{fact}\,n) \cong \mathit{fact}'\,n\,k$, and similarly, Hutton (2021) calculates the CPS version of an evaluator from its specification as: $\mathsf{exec}\,k\,(\mathit{eval}\,e) \cong \mathit{eval}'\,e\,k$. Following that approach, we use $\mathsf{app}\,k\,(f\,e) \cong f'\,e\,k$ **(a)** as our initial specification. This seems a good start since it implies:

$$
\begin{array}{lll}
 & f\,e & \\
= & \Box[f\,e] & \{\ \mathit{context}\ \} \\
= & \mathsf{app}\,(\mathsf{ctx}\,\Box)\,(f\,e) & \{\ (\mathit{appctx})\ \} \\
\cong & f'\,e\,(\mathsf{ctx}\,\Box) & \{\ \mathit{specification}\,(a)\ \}
\end{array}
$$

and we can thus replace any applications of $f\,e$ in the program with applications to the TRMC translated f' instead as $f'\,e\,(\mathsf{ctx}\,\Box)$.

Unfortunately, the specification is not yet specific enough to calculate with as it does not include the translation function $[\![_]\!]_{f,k}$ itself which limits what we can derive. Can we change this? Let's start by deriving how we can satisfy our initial specification (a):

$$
\begin{array}{lll}
 & \mathsf{app}\,k\,(f\,e) & \\
\cong & \mathsf{app}\,k\,(\mathsf{let}\,x = e\,\mathsf{in}\,e_f) & \{\ (\mathit{letapp})\ \} \\
\cong & \mathsf{let}\,x = e\,\mathsf{in}\,\mathsf{app}\,k\,e_f & \{\ (\mathit{letfloat}),\ x \notin \mathsf{fv}(k)\ \} \\
= & \mathsf{let}\,x = e\,\mathsf{in}\,[\![e_f]\!]_{f,k} & \{\ \mathit{define\ specification}\,(b)\,\mathit{below}\ \} \\
= & \mathsf{let}\,x = e\,\mathsf{in}\,f'\,x\,k & \{\ \mathit{def}.\ \} \\
\cong & f'\,(\mathsf{let}\,x = e\,\mathsf{in}\,x)\,k & \{\ (\mathit{letfloat}),\ x \notin \mathsf{fv}(f', k)\ \} \\
\cong & f'\,e\,k & \{\ (\mathit{letid})\ \}
\end{array}
$$

This suggests a more general specification as $\mathsf{app}\,k\,e \cong [\![e]\!]_{f,k}$ **(b)** (for any e) which both implies our original specification but also includes the translation function now. The improved specification directly gives us a trivial solution for the translation as:

$$(\mathit{base}) \quad [\![e]\!]_{f,k} = \mathsf{app}\,k\,e$$

That is not quite what we need for general TRMC though since this does not translate any tail calls modulo a context in e at all. However, we can be more specific by matching on the shape of e. In particular, we can match on general tail-modulo-context calls in the expression e if it has the shape $\mathsf{E}[f\,e_1]$. We can then calculate[1]:

$$
\begin{array}{lll}
 & \mathsf{app}\,k\,\mathsf{E}[f\,e_1] & \\
= & \mathsf{app}\,k\,(\mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\,(f\,e_1)) & \{\ (\mathit{appctx})\ \} \\
= & \mathsf{app}\,(k \bullet \mathsf{ctx}\,\mathsf{E})\,(f\,e_1) & \{\ (\mathit{appcomp})\ \} \\
\cong & f'\,e_1\,(k \bullet \mathsf{ctx}\,\mathsf{E}) & \{\ \mathit{specification}\,(a)\ \} \\
= & [\![\mathsf{E}[f\,e_1]]\!]_{f,k} & \{\ \mathit{define}\ \}
\end{array}
$$

[1] The reader may note that the use specification (a) in this derivation is not directly inductive on the argument. In all the other derivations in this section, the use of our specification is always on a smaller term (and similarly in related work by Hutton, 2021 or Gibbons, 2022). Only in the directly tail-recursive case here, we use the specification on the original term. However, we show in Appendix B.1 that since we apply it after unfolding the original function, we can use Löb induction to show the use of our specification here is still sound (Theorem 1).

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Effectively, we replace all direct tail-recursive calls in the original function f to tail-recursive calls in our translated function f' by just extending the continuation parameter k. Together with our earlier equation, we now have an initial specification of our transation function:

$$\begin{array}{lll} (\mathit{tail}) & [\![\mathsf{E}[f\ e]]\!]_{f,k} & = f'\ e\ (k \bullet \mathsf{ctx}\ \mathsf{E}) \quad \mathsf{iff}\ (\star) \\ (\mathit{base}) & [\![e]\!]_{f,k} & = \mathsf{app}\ k\ e \quad\quad\quad \mathsf{otherwise} \end{array}$$

Note that the equations overlap – for a particular instance of the algorithm we generally constrain the (*tail*) rule to only apply for certain contexts E constrained by some particular ($\star$) condition (e.g., constructor contexts), falling back to (*base*) otherwise. Similarly, the (*tail*) case allows a choice in where to apply the tail call for expressions like $f\ (f\ e)$ for example and a particular instantiation of ($\star$) should disambiguate for an actual algorithm. By default, we assume that any instantiation matches on the innermost application of f (for reasons discussed in Section 4.2).

This is still a bit constrained, as these equations do not consider any evaluation contexts E where the recursive call is under a let or match expression. We can again match on these specific forms of e. For example, $\mathsf{let}\, x = e_0\, \mathsf{in}\, e_1$ where $e_0 \neq \mathsf{E}[f\ e']$ (so it does not overlap with E contexts):

$$\begin{array}{lll} & \mathsf{app}\ k\ (\mathsf{let}\, x = e_0\, \mathsf{in}\, e_1) & \\ \cong & \mathsf{let}\, x = e_0\, \mathsf{in}\, \mathsf{app}\ k\ e_1 & \{\ (\mathit{letfloat}),\ x \notin \mathsf{fv}(k)\ \} \\ \cong & \mathsf{let}\, x = e_0\, \mathsf{in}\, [\![e_1]\!]_{f,k} & \{\ \mathit{specification}\ \} \\ = & [\![\mathsf{let}\, x = e_0\, \mathsf{in}\, e_1]\!]_{f,k} & \{\ \mathit{define}\ \} \end{array}$$

Unfortunately, this rule is still too restrictive in general as it does not apply when the let-statement is itself under a context E. For example, we might encounter an expression like:

$$\mathsf{let}\, x = (\mathsf{let}\, y = e_0\, \mathsf{in}\, f\ x\ y)\, \mathsf{in}\, e_1$$

Here, the recursive call is under the let-binding of x (as $\mathsf{E}[\mathsf{let}\, y = e_0\, \mathsf{in}\, f\ x\ y]$), but the $y = e_0$ binding prevents the recursive call $f\ x\ y$ to be the focus of the evaluation context. This situation occurs whenever an expression is not in *A-normal form* (Flanagan *et al.*, 1993), and these cannot be optimized by the rules outlined so far (and neither by the rules as presented in earlier work Leijen & Lorenzen, 2023). Instead, we need to consider the general case where the let-binding appears under a context E. Assuming that the variables bound in let-bindings and matches are fresh, we can calculate:

$$\begin{array}{lll} & \mathsf{app}\ k\ \mathsf{E}[\mathsf{let}\, x = e_0\, \mathsf{in}\, e_1] & \\ \cong & \mathsf{app}\ k\ (\mathsf{app}\ (\mathsf{ctx}\ \mathsf{E})\ (\mathsf{let}\, x = e_0\, \mathsf{in}\, e_1)) & \{\ (\mathit{appctx})\ \} \\ \cong & \mathsf{let}\, x = e_0\, \mathsf{in}\, \mathsf{app}\ k\ (\mathsf{app}\ (\mathsf{ctx}\ \mathsf{E})\ e_1) & \{\ (\mathit{letfloat}),\ x \notin \mathsf{fv}(\mathsf{E}, k)\ \} \\ \cong & \mathsf{let}\, x = e_0\, \mathsf{in}\, \mathsf{app}\ k\ \mathsf{E}[e_1] & \{\ (\mathit{appctx})\ \} \\ \cong & \mathsf{let}\, x = e_0\, \mathsf{in}\, [\![\ \mathsf{E}[e_1]\ ]\!]_{f,k} & \{\ (\mathit{specification})\ \} \\ = & [\![\ \mathsf{E}[\mathsf{let}\, x = e_0\, \mathsf{in}\, e_1]\ ]\!]_{f,k} & \{\ \mathit{define}\ \} \end{array}$$

Effectively we have lifted out the let-binding from the evaluation context. We can do the same for matches:

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

$$
\begin{array}{lll}
 & \mathsf{app}\,k\,\mathsf{E}[\mathsf{match}\,e_0\,\{\overline{p_i \to e_i}\}] & \\
\cong & \mathsf{app}\,k\,(\mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\,(\mathsf{match}\,e_0\,\{\overline{p_i \to e_i}\})) & \{\ (\mathit{appctx})\ \} \\
\cong & \mathsf{match}\,e_0\,\{\overline{p_i \to \mathsf{app}\,k\,(\mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\,e_i)}\} & \{(\mathit{matchfloat}),\ p_i = C_i\,x_1 \ldots x_n,\ x_j \notin \mathsf{fv}(\mathsf{E}, k)\} \\
\cong & \mathsf{match}\,e_0\,\{\overline{p_i \to \mathsf{app}\,k\,\mathsf{E}[e_i]}\} & \{\ (\mathit{appctx})\ \} \\
\cong & \mathsf{match}\,e_0\,\{\overline{p_i \to [\![\,\mathsf{E}[e_i]\,]\!]_{f,k}}\} & \{\ (\mathit{specification})\ \} \\
= & [\![\,\mathsf{E}[\mathsf{match}\,e_0\,\{\overline{p_i \to e_i}\}]\,]\!]_{f,k} & \{\ \mathit{define}\ \}
\end{array}
$$

This form of specification essentially performs A-normalization whenever necessary to create further opportunities to match on tail-recursive calls. Our presentation follows the approach of Maurer *et al*. (2017), who describe the positions in a term that occur last in an evaluation order as *tail contexts* T. They show that A-normalization can be achieved by commuting the E and T contexts whenever possible. This is exactly the approach taken here, where we commute single let-bindings and matches under E contexts to the front of the term. A potential drawback of the match normalization is that it duplicates the evaluation context E in each of the branches. Maurer *et al*. (2017) also show how *join points* can be used to avoid code duplication in such case.

This leaves one last expression form to consider: the application of a function to an argument. Using the intuition of commuting tail contexts, we might define $[\![\,\mathsf{E}[e_0\,e_1]\,]\!]_{f,k} = e_0\,[\![\,\mathsf{E}[e_1]\,]\!]_{f,k}$. However, while e_0 can now be evaluated early, the application itself depends on the result of our transformation. Thus, we need to be a bit more careful and instead calculate:

$$
\begin{array}{lll}
 & \mathsf{app}\,k\,\mathsf{E}[e_0\,e_1] & \\
\cong & \mathsf{app}\,k\,(\mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\,(e_0\,e_1)) & \{\ (\mathit{appctx})\ \} \\
\cong & \mathsf{app}\,k\,(\mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\,((\mathsf{let}\,g = e_0\,\mathsf{in}\,g)\,e_1)) & \{\ (\mathit{letid}),\ \mathit{for fresh}\,g\ \} \\
\cong & \mathsf{let}\,g = e_0\,\mathsf{in}\,\mathsf{app}\,k\,(\mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\,(g\,e_1)) & \{\ (\mathit{letfloat}),\ \mathit{fresh}\,g\ \} \\
\cong & \mathsf{let}\,g = e_0\,\mathsf{in}\,\mathsf{app}\,k\,\mathsf{E}[g\,e_1] & \{\ (\mathit{appctx})\ \} \\
\cong & \mathsf{let}\,g = e_0\,\mathsf{in}\,[\![\,\mathsf{E}[g\,e_1]\,]\!]_{f,k} & \{\ (\mathit{specification})\ \} \\
= & [\![\,\mathsf{E}[e_0\,e_1]\,]\!]_{f,k} & \{\ \mathit{define}\ \}
\end{array}
$$

In this general form, we need to strengthen our requirement that $e_0 \neq \mathsf{E}[f\,e']$ to ensure that it does not overlap with our newly calculated rules. We write $e_0 \neq \mathsf{X}[f\,e']$ to mean that e_0 cannot have a recursive call under an *expression context* X. The expression context $\mathsf{X}[e]$ matches all possible expressions that contain e, unless e occurs in $\mathsf{X}[e]$ exclusively under lambdas.

3.3 The tail-recursion-modulo-contexts algorithm

Figure 2 shows all five of the calculated equations for our generic tail recursion modulo contexts transformation (extended to multiple parameters). We can instantiate this algorithm by defining the context type $\mathit{ctx}\,\alpha$, the context construction (ctx), composition ($\bullet$), and application (app) operations, and finally the ($\star$) condition constrains the allowed context E to fit the particular context type.

As remarked in Section 3.2, our equationally derived algorithm relied on a (single) non-inductive step, but we can still show formally that the algorithm is indeed sound:

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

$$
\begin{array}{lllll}
(\mathit{tlet}) & [\![\, \mathsf{E}[\mathsf{let}\, x = e_0\, \mathsf{in}\, e] \,]\!]_{f,k} & = \mathsf{let}\, x = e_0\, \mathsf{in}\, [\![\, \mathsf{E}[e] \,]\!]_{f,k} & \\
(\mathit{tmatch}) & [\![\, \mathsf{E}[\mathsf{match}\, e_0\, \{\overline{p_i \to e_i}\}] \,]\!]_{f,k} & = \mathsf{match}\, e_0\, \{\overline{p_i \to [\![\, \mathsf{E}[e_i] \,]\!]_{f,k}}\} & \\
(\mathit{tapp}) & [\![\, \mathsf{E}[e_0\, e] \,]\!]_{f,k} & = \mathsf{let}\, g = e_0\, \mathsf{in}\, [\![\, \mathsf{E}[g\, e] \,]\!]_{f,k} & \text{with } g \text{ fresh} \\
(\mathit{tail}) & [\![\, \mathsf{E}[f\, e_0] \,]\!]_{f,k} & = f'\, e_0\, (k \bullet (\mathsf{ctx}\, \mathsf{E})) & \text{iff } (\star) \\
(\mathit{base}) & [\![e]\!]_{f,k} & = \mathsf{app}\, k\, e & \text{otherwise}
\end{array}
$$

where $e_0 \neq \mathsf{X}[f\, e']$

Fig. 2. Calculated algorithm for general selective tail recursion modulo context transformation. It is parameterized by the ($\star$) condition, the composition ($\bullet$), and application (app) operations.

Theorem 1. (*The TRMC translation is sound*)
Let $f\, x = e_f$ and $f'\, x\, k = [\![e_f]\!]_{f,k}$, then $\mathit{app}\, k\, (f\, x) \cong f'\, x\, k$.

See Appendix B.1 for the proof.

Thanks to the A-normalization, we can also show that the TRMC algorithm exhaustively optimizes recursive calls:

Theorem 2. (*Matching all recursive calls*)
For any transformed expression $e' = [\![e]\!]_{f,k}$ with ($\star$) unconstrained, we have $e' \neq \mathsf{X}[f\, e_0]$.

There are two types of recursive calls that cannot be optimized by our algorithm: First, we do not optimize recursive calls inside lambdas. This is a necessary restriction, since it is impossible in general to push the accumulated context k under a lambda. Second, our algorithm will only optimize the *first* recursive call(s) in the evaluation order. If those are followed by further recursive calls, the evaluation context E stored as ctx E in the (*tail*) rule may still contain unoptimized recursive calls. We will revisit this problem in Section 4.2.

To see our algorithm in action, let's consider the *map* function again:

$$
\begin{array}{l}
\mathit{map}\, \mathit{xs}\, f\ =\ \mathsf{match}\, \mathit{xs}\, \{\ \mathit{Nil} \to \mathit{Nil} \\
\qquad\qquad\qquad\qquad \mathit{Cons}\, x\, \mathit{xx} \to \mathit{Cons}\, (f\, x)\, (\mathit{map}\, \mathit{xx}\, f)\ \}
\end{array}
$$

When translating this function, we first use the (*tmatch*) rule with $\mathsf{E} = \square$ to descend into the branches of the match. In the *Nil* branch, the (*base*) rule applies. In the *Cons* branch, we use the (*tapp*) rule (again with $\mathsf{E} = \square$) to bind the call to $f\, x$. We then use the (*tail*) rule to optimize the recursive call to `map xx f`:

$$
\begin{array}{l}
\mathit{map}'\, \mathit{xs}\, f\, k\ =\ \mathsf{match}\, \mathit{xs}\, \{\ \mathit{Nil} \to \mathsf{app}\, k\, \mathit{Nil} \\
\qquad\qquad\qquad\qquad \mathit{Cons}\, x\, \mathit{xx} \to \mathsf{let}\, c\ =\ \mathit{Cons}\, (f\, x)\, \mathsf{in}\, \mathit{map}'\, \mathit{xx}\, f\, (k \bullet (\mathsf{ctx}\, (c\, \square)))\ \}
\end{array}
$$

However, for Constructor Contexts (Section 6), it is useful to keep the *Cons* constructor in the context passed to *map*. In our pratical implementation, we therefore modify the (*tapp*) rule slightly to extract the arguments instead of the entire partially applied function. Our final transformation for *map* is then:

$$
\begin{array}{l}
\mathit{map}'\, \mathit{xs}\, f\, k\ =\ \mathsf{match}\, \mathit{xs}\, \{\ \mathit{Nil} \to \mathsf{app}\, k\, \mathit{Nil} \\
\qquad\qquad\qquad\qquad \mathit{Cons}\, x\, \mathit{xx} \to \mathsf{let}\, y\ =\ f\, x\, \mathsf{in}\, \mathit{map}'\, \mathit{xx}\, f\, (k \bullet (\mathsf{ctx}\, (\mathit{Cons}\, y\, \square)))\ \}
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

4 Instantiations of the general TRMC transformation

With the general TRMC transformation in hand, we discuss various instantiations in this section. In the next section, we look at the update-in-place modulo cons (TRMc) instantiation in detail.

4.1 Modulo evaluation contexts

If we use *true* for the ($\star$) condition, we can translate any recursive tail modulo evaluation context functions. Representing our abstract context directly as an E context is usually not possible though as E contexts generally contain *code*. The usual way to represent an arbitrary evaluation context E is simply as a (continuation) function $\lambda x.\, \mathsf{E}[x]$ with a context type $\mathit{ctx}\,\alpha = \alpha \to \alpha$:

$$
\begin{array}{lll}
(\mathit{ectx}) & \mathsf{ctx}\,\mathsf{E} & = \lambda x.\, \mathsf{E}[x] \qquad (x \notin \mathsf{fv}(\mathsf{E})) \\
(\mathit{ecomp}) & k_1 \bullet k_2 & = k_1 \circ k_2 \\
(\mathit{eapp}) & \mathsf{app}\, k\, e & = k\, e
\end{array}
$$

This is an intuitive definition where $\mathsf{ctx}\,\Box$ corresponds to the identity function and context composition to function composition. If we apply the TRMC translation, we are essentially performing a selective CPS translation where the context E is represented as the continuation function. We can verify that the context laws hold for this instantiation (where we can assume e is terminating):

Composition:

$$
\begin{array}{lll}
 & \mathsf{app}\,(k_1 \bullet k_2)\, e & \\
= & \mathsf{app}\,(k_1 \circ k_2)\, e & \{\ (\mathit{ecomp})\ \} \\
= & \mathsf{app}\,(\lambda x.\, k_1\,(k_2\, x))\, e & \{\ \mathit{def}\ \circ\ \} \\
= & (\lambda x.\, k_1\,(k_2\, x))\, e & \{\ (\mathit{eapp})\ \} \\
\cong & \mathsf{let}\, x = e\, \mathsf{in}\, k_1\,(k_2\, x) & \{\ (\mathit{letapp})\ \} \\
\cong & k_1\,(k_2\,(\mathsf{let}\, x = e\, \mathsf{in}\, x)) & \{\ (\mathit{letfloat})\ \} \\
\cong & k_1\,(k_2\, e) & \{\ (\mathit{letid})\ \} \\
= & k_1\,(\mathsf{app}\, k_2\, e) & \{\ (\mathit{eapp})\ \} \\
= & \mathsf{app}\, k_1\,(\mathsf{app}\, k_2\, e) & \{\ (\mathit{eapp})\ \}
\end{array}
$$

and application:

$$
\begin{array}{lll}
 & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\, e & \\
= & \mathsf{app}\,(\lambda x.\, \mathsf{E}[x])\, e & \{\ (\mathit{ecomp})\ \} \\
= & (\lambda x.\, \mathsf{E}[x])\, e & \{\ (\mathit{eapp})\ \} \\
\cong & \mathsf{let}\, x = e\, \mathsf{in}\, \mathsf{E}[x] & \{\ (\mathit{letapp})\ \} \\
\cong & \mathsf{E}[\mathsf{let}\, x = e\, \mathsf{in}\, x] & \{\ (\mathit{letfloat})\ \} \\
= & \mathsf{E}[e] & \{\ (\mathit{letid})\ \}
\end{array}
$$

As a concrete example, let's apply the modulo evaluation context to the *map* function:

$$
\begin{array}{l}
\mathit{map}\, \mathit{xs}\, f\ =\ \mathsf{match}\, \mathit{xs}\, \{\ \mathit{Nil} \to \mathit{Nil} \\
\qquad\qquad\qquad\qquad \mathit{Cons}\, x\, \mathit{xx} \to \mathsf{let}\, y = f\, x\, \mathsf{in}\, \mathit{Cons}\, y\,(\mathit{map}\, \mathit{xx}\, f)\ \}
\end{array}
$$

which translates to:

$$
\begin{array}{l}
\mathit{map}'\, \mathit{xs}\, f\, k\ =\ \mathsf{match}\, \mathit{xs}\, \{\ \mathit{Nil} \to \mathsf{app}\, k\, \mathit{Nil} \\
\qquad\qquad\qquad\qquad\qquad\qquad \mathit{Cons}\, x\, \mathit{xx} \to \mathsf{let}\, y = f\, x\, \mathsf{in}\, \mathit{map}'\, \mathit{xx}\, f\, (k \bullet (\mathsf{ctx}\,(\mathit{Cons}\, y\, \Box)))\ \}
\end{array}
$$

and which the compiler can further simplify into:

$$
\begin{array}{l}
\mathit{map}'\, \mathit{xs}\, f\, k\ =\ \mathsf{match}\, \mathit{xs}\, \{\ \mathit{Nil} \to k\, \mathit{Nil} \\
\qquad\qquad\qquad\qquad\qquad\qquad \mathit{Cons}\, x\, \mathit{xx} \to \mathsf{let}\, y = f\, x\, \mathsf{in}\, \mathit{map}'\, \mathit{xx}\, f\, (\lambda x.\, k\,(\mathit{Cons}\, y\, x))\ \}
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

where we derived exactly the standard CPS style version of `map` as shown in Section 2. A general evaluation context transformation creates more opportunities for tail-recursive calls, but this also happens at the cost of potentially heap allocating continuation closures. As such, it is not common for strict languages to use this instantiation. The exception would be languages like scheme that always guarantee tail calls, but in that case the modulo evaluation contexts instantiation is already subsumed by general CPS conversion.

4.2 Nested translation of modulo evaluation contexts

The current instantiation is already very general as it applies to any E context, but we can do a little better. While the innermost non-tail call $\mathsf{E}[f\ e]$ becomes $f'\ e\ (k \bullet \mathsf{ctx}\ \mathsf{E})$, the context E may contain itself further recursive calls to f. Since k is just a variable this allocates a closure for each composition ($\bullet$) and invokes every nested call $f\ e$ with an empty context as $f'\ e\ (\mathsf{ctx}\ \Box)$ before composing with k. This is not ideal, and in the classic CPS translation this is avoided by passing k itself into the closure for $\mathsf{ctx}\ \mathsf{E}$ directly. Fortunately, we can achieve the same by *specializing* the compose function using the specification (b):

$$
\begin{array}{lll}
 & k \bullet (ctx\ \mathsf{E}) & \\
= & \lambda x.\, k\ ((\mathsf{ctx}\ \mathsf{E})\ x) & \{\ (ecomp),\ (\bullet)\ \} \\
\cong & \lambda x.\, k\ \mathsf{E}[x] & \{\ (ectx),\ (beta)\ \} \\
= & \lambda x.\, \mathsf{app}\ k\ \mathsf{E}[x] & \{\ (eapp)\ \} \\
\cong & \lambda x.\, [\![\mathsf{E}[x]]\!]_{f,k} & \{\ specification\ (b)\ \}
\end{array}
$$

That is, in the compiler, instead of generating $k \bullet (\mathsf{ctx}\ \mathsf{E})$, we invoke the TRMC translation recursively in the (*tail*) case and generate $\lambda x.\, [\![\mathsf{E}[x]]\!]_{f,k}$ instead. This avoids the allocation of function composition closures and directly passes the continuation k to any nested recursive calls.

4.3 Modulo defunctionalized evaluation contexts

In order to better understand the shapes that evaluation contexts can take, we want to consider the *defunctionalization* (Reynolds, 1972; Danvy & Nielsen, 2001) of the general evaluation context transformation. It turns out that this yields an interesting context in its own right. First, we observe that in any recursive function the evaluation context can only take a finite number of shapes depending on the number of recursive calls. We write this as:

$$\mathsf{E} ::= \Box \mid \mathsf{E}_1 \mid \ldots \mid \mathsf{E}_n$$

We define an *accumulator* datatype by creating a constructor H for the $\Box$ context and for each E_i a constructor A_i that carries the free variables of E_i. The compiler then generates an `app` function where we interpret A_i by evaluating E_i with the stored free variables:

$$
\begin{array}{lllll}
(dctx) & \mathsf{ctx}\ \mathsf{E}_i & = \mathsf{A}_i\ x_1 \ldots x_m\ \mathsf{H} & \text{where } x_1, \ldots, x_m = \mathsf{fv}(\mathsf{E}_i) \\
(dcomp) & k_1 \bullet \mathsf{H} & = k_1 & \\
(dcomp) & k_1 \bullet (\mathsf{A}_i\ x_1 \ldots x_m\ k_2) & = \mathsf{A}_i\ x_1 \ldots x_m\ (k_1 \bullet k_2) & \\
(dapp) & \mathsf{app}\ \mathsf{H}\ e & = e & \\
(dapp) & \mathsf{app}\ (\mathsf{A}_i\ x_1 \ldots x_m\ k)\ e & = [\![\mathsf{E}_i[e\,,\ x_1 \ldots x_m]]\!]_{f,k} &
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Just as we saw in Section 4.2, we need to use the translated evaluation context in the definition of app to translate nested calls. The context laws now follow by induction – see Appendix B.2 in the supplement for the derivations. Applying this instantiation to the `map` function, we obtain:

$$\begin{aligned}&\mathsf{type}\ ctx\,\alpha\ =\ \mathsf{H}\mid \mathsf{A}_1\,\alpha\,(ctx\,\alpha)\\&map'\,xs\,f\,k\ =\ \mathsf{match}\,xs\,\{\,Nil\to app\,k\,Nil;\ Cons(x,xx)\to \mathsf{let}\,y=f\,x\,\mathsf{in}\,map'\,xx\,f\,(\mathsf{A}_1\,y\,k)\,\}\end{aligned}$$

In the *Cons* branch, we have inlined $k \bullet (\mathsf{A}_1\,y\,\mathsf{H})$. The *app* function interprets A_1 by calling itself recursively on the stored evaluation context:

$$app\,k\,xs\ =\ \mathsf{match}\,k\,\{\,\mathsf{H}\to xs;\ \mathsf{A}_1(y,\ k')\to app\,k'\,(Cons\,y\,xs)\,\}$$

As we can see, using the modulo defunctionalized evaluation context translation, we derived exactly the accumulator version of the `map` function that reverses the accumulated list in the end (where `app` is `reverse`)! In particular, for the special case where all evaluation contexts are constructor contexts $C^m\,x_1\ldots(f\ldots)\ldots x_m$ (as is the case for `map`), the accumulator datatype stores a path into the data structure we are building and thus essentially becomes a zipper structure (Huet, 1997).

This defunctionalized approach might resemble general closure conversion at first (Appel, 1991): In both approaches, we store the free variables in a datatype. However, in closure conversion the datatype typically also contains a machine code pointer and one jumps to the code by calling this pointer, while in our case we match on the specialized constructors (similar to the approach of Tolmach & Oliva, 1998).

4.3.1 Reuse

As the defunctionalization makes the evaluation context explicit, we can optimize it further. As Sobel & Friedman (1998) note, the defunctionalized closure is only applied once and we can reuse its memory for other allocations. This can happen automatically in languages with reuse analysis such as Koka (Lorenzen & Leijen, 2022), Lean (Ullrich & de Moura, 2019), or OPAL (Didrich *et al.*, 1994). In particular, in the `app` function, the match:

$$\mathsf{A}_1\,y\,k'\to app\,k'\,Cons(y,\ xs)$$

can reuse the A_1 in-place to allocate the *Cons* node if the A_1 is unique at runtime. In our case, the context is actually always unique (we show this formally in Section 6.1), and the A_1 nodes are always reused! Even better, if the initial list is unique, we also reuse the initial *Cons* cell for the A_1 accumulator itself in *map′* and no allocation takes place at all – the program is functional but in-place (Xie *et al.*, 2021; Lorenzen *et al.*, 2023).

4.4 Modulo associative operator contexts

In the previous instantiations, we considered general evaluation contexts. However, we can often derive more efficient instantiations by considering more restricted contexts. A particularly nice example are monoidal contexts. For any monoid with an associative operator $\odot : \tau\to\tau\to\tau$ and a unit value $unit : \tau$, we can define a restricted operator context as:

$$\mathsf{A} ::= \Box \mid v \odot \mathsf{A}$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

For a concrete example, consider the *length* function defined as:

$$\mathit{length}\,\mathit{xs} = \mathsf{match}\,\mathit{xs}\,\{\,\mathit{Cons}\,x\,\mathit{xx} \to 1 + \mathit{length}\,\mathit{xx};\ \mathit{Nil} \to 0\,\}$$

which applies for integer addition ($\odot = +$, $\mathit{unit} = 0$). The idea is now to define a compile time *fold* function $(\!|_|\!)$ over a context A to always reduce the context to a single element of type τ:

$$\begin{array}{ll} (\!|\Box|\!) & = \mathit{unit} \\ (\!|v \odot \mathsf{A}|\!) & = v \odot (\!|\mathsf{A}|\!) \end{array}$$

We can now instantiate the abstract contexts by defining the $(\star)$ condition to constrain the E context to A, and the context type to $\mathit{ctx}\,\tau = \tau$, where we use the fold operation to represent contexts always as a single element of type τ:

$$\begin{array}{lll} (\mathit{lctx}) & \mathsf{ctx}\,\mathsf{A} & = (\!|\mathsf{A}|\!) \\ (\mathit{lcomp}) & k_1 \bullet k_2 & = k_1 \odot k_2 \\ (\mathit{lapp}) & \mathsf{app}\,k\,e & = k \odot e \end{array}$$

The context laws hold for this definition. For composition, we can derive:

$$\begin{array}{lll} & \mathsf{app}\,(k_1 \bullet k_2)\,e & \\ = & \mathsf{app}\,(k_1 \odot k_2)\,e & \{\ (\mathit{lcomp})\ \} \\ = & (k_1 \odot k_2) \odot e & \{\ (\mathit{lapp})\ \} \\ = & k_1 \odot (k_2 \odot e) & \{\ \mathit{assoc.}\ \} \\ = & \mathsf{app}\,k_1\,(\mathsf{app}\,k_2\,e) & \{\ (\mathit{lapp})\ \} \end{array}$$

and for context application we have:

$$\begin{array}{lll} & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{A})\,e & \\ = & \mathsf{app}\,(\!|\mathsf{A}|\!)\,e & \{\ (\mathit{lctx})\ \} \\ = & (\!|\mathsf{A}|\!) \odot e & \{\ (\mathit{lapp})\ \} \end{array}$$

We proceed by induction over A.

Case $\mathsf{A} = \Box$:

$$\begin{array}{lll} = & (\!|\Box|\!) \odot e & \\ = & \mathit{unit} \odot e & \{\ \mathit{fold}\ \} \\ = & e & \{\ \mathit{unit}\ \} \\ = & \Box[e] & \{\ \Box\ \} \end{array}$$

and the case $\mathsf{A} = v \odot \mathsf{A}'$:

$$\begin{array}{lll} = & (\!|v \odot \mathsf{A}'|\!) \odot e & \\ = & (v \odot (\!|\mathsf{A}'|\!)) \odot e & \{\ \mathit{fold}\ \} \\ = & v \odot ((\!|\mathsf{A}'|\!) \odot e) & \{\ \mathit{assoc.}\ \} \\ = & v \odot \mathsf{A}'[e] & \{\ \mathit{induction\ hyp.}\ \} \\ = & \mathsf{A}[e] & \{\ \mathsf{A}\ \mathit{context}\ \} \end{array}$$

Common instantiations include integer addition ($\odot = +$, $\mathit{unit} = 0$) and integer multiplication ($\odot = \times$, $\mathit{unit} = 1$). The TRMC algorithm with A contexts instantiated with integer addition translates the previous *length* function to the following tail-recursive version:

$$\mathit{length}'\,\mathit{xs}\,k = \mathsf{match}\,\mathit{xs}\,\{\,\mathit{Cons}\,x\,\mathit{xx} \to \mathit{length}'\,\mathit{xx}\,(k \bullet (\mathsf{ctx}\,(1 + \Box)));\ \mathit{Nil} \to \mathsf{app}\,k\,0\,\}$$

The intention is that the fold function is performed by the compiler, and the compiler can simplify this further as:

$$k \bullet (\mathsf{ctx}\,(1 + \Box)) = k + (\mathsf{ctx}\,(1 + \Box)) = k + (\!|1 + \Box|\!) = k + 1$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

such that we end up with:

$$length'\ xs\ k\ =\ \mathsf{match}\ xs\ \{\ Cons\ x\ xx \to length'\ xx\ (k\ +\ 1);\ Nil \to k\ \}$$

This time we derived exactly the text book accumulator version of *length*.

4.4.1 Using right biased contexts

Our defined context only allows the recursive call on the left, but we can also define a right biased context:

$$\mathsf{A}\ ::= \Box \mid \mathsf{A} \odot v$$

with the fold defined as:

$$\begin{aligned} (\!|\Box|\!) &= unit \\ (\!|\mathsf{A} \odot v|\!) &= (\!|\mathsf{A}|\!) \odot v \end{aligned}$$

We can now compose in the opposite order:

$$\begin{aligned} &(rctx) && \mathsf{ctx}\ \mathsf{A} && = (\!|\mathsf{A}|\!) \\ &(rcomp) && k_1 \bullet k_2 && = k_2 \odot k_1 \\ &(rapp) && \mathsf{app}\ k\ e && = e \odot k \end{aligned}$$

We can again show that the context laws hold for this definition (see Appendix B.3 in the supplement). As an example, we can instantiate $\odot$ as list append $+\!\!+$ with the empty list as the unit element to transform the *reverse* function:

$$reverse\ xs\ =\ \mathsf{match}\ xs\ \{\ Cons\ x\ xx \to reverse\ xx +\!\!+ [x];\ Nil \to []\ \}$$

First, our TRMC algorithm transforms it into:

$$reverse'\ xs\ k\ =\ \mathsf{match}\ xs\ \{\ Cons\ x\ xx \to reverse'\ xx\ (k \bullet (\mathsf{ctx}\ (\Box +\!\!+ [x])));\ Nil \to \mathsf{app}\ k\ []\ \}$$

and with our instantiated context, this simplifies to:

$$reverse'\ xs\ k\ =\ \mathsf{match}\ xs\ \{\ Cons\ x\ xx \to reverse'\ xx\ ([x] +\!\!+ k);\ Nil \to [] +\!\!+ k\ \}$$

Using right-biased contexts, we derived the text book accumulator version of reverse. This shows that our general TRMC algorithm can be instantiated to eliminate `append` calls automatically as first proposed by Hughes (1986) and Wadler (1987).

4.5 Modulo monoid contexts

To handle general monoids, we need to consider recursive calls on both sides of the associative operation:

$$\mathsf{A}\ ::= \Box \mid v \odot \mathsf{A} \mid \mathsf{A} \odot v$$

This context A expresses arbitrarily nested applications of $\odot$. As monoid operations may not be commutative, we cannot use a single element to represent the context. Instead, we need to use a *product context* where we accumulate the left and right context separately:

$$\begin{aligned} (\!|\Box|\!) &= (unit,\ unit) \\ (\!|v \odot \mathsf{A}|\!) &= (v \odot l,\ r) && \text{where}\ (l,\ r) = (\!|\mathsf{A}|\!) \\ (\!|\mathsf{A} \odot v|\!) &= (l,\ r \odot v) && \text{where}\ (l,\ r) = (\!|\mathsf{A}|\!) \end{aligned}$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

which we compose as:

$$\begin{array}{lll}
(actx) & \mathsf{ctx}\,\mathsf{A} & = (\!|\mathsf{A}|\!) \\
(acomp) & (l_1, r_1) \bullet (l_2, r_2) & = (l_1 \odot l_2, r_2 \odot r_1) \\
(aapp) & \mathsf{app}\,(l, r)\,e & = l \odot e \odot r
\end{array}$$

We can again show that the context laws hold for this definition (see Appendix B.4 in the supplement).

4.6 Modulo semiring contexts

We can also combine the associative operators of two monoids, as long as one distributes over the other. This is the case for semirings in particular (although we do not need commutativity of $+$). Semiring contexts are relatively common in practice. For example, consider the following hashing function for a list of integers as shown by Bloch (2008):

$$\mathit{hash}\,\mathit{xs} = \mathsf{match}\,\mathit{xs}\,\{\,\mathit{Cons}\,x\,\mathit{xx} \to x + 31 * (\mathit{hash}\,\mathit{xx})\,;\ \mathit{Nil} \to 17\,\}$$

Implementing modulo semiring contexts in a compiler may be worthwhile as deriving a tail-recursive version manually for such contexts is not always straightforward (and the interested reader may want to pause here and try to rewrite the *hash* function in a tail-recursive way before reading on).

We can define a general context for semirings as:

$$\mathsf{A} ::= \Box \mid v + \mathsf{A} \mid v * \mathsf{A} \mid \mathsf{A} + v \mid \mathsf{A} * v$$

For simplicity, we assume we have a commutative semiring where both addition and multiplication commute. This allows us to use again a product representation at runtime where we accumulate the additions and multiplications separately (and without commutativity we need a quadruple instead). In the definition of the fold, we take into account that the multiplication distributes over the addition:

$$\begin{array}{lll}
(\!|\Box|\!) & = (\mathit{unit}^{+}, \mathit{unit}^{*}) & \\
(\!|v + \mathsf{A}|\!) & = (v + l, r) & \text{where}\,(l, r) = (\!|\mathsf{A}|\!) \\
(\!|v * \mathsf{A}|\!) & = (v * l, v * r) & \text{where}\,(l, r) = (\!|\mathsf{A}|\!) \\
(\!|\mathsf{A} + v|\!) & = (\!|v + \mathsf{A}|\!) & (+\,\mathit{commutes}) \\
(\!|\mathsf{A} * v|\!) & = (\!|v * \mathsf{A}|\!) & (*\,\mathit{commutes})
\end{array}$$

Finally, to compose the contexts we need to use distributivity again. Note how the (*scomp*) rule mirrors the definition of $(\!|\mathsf{A}|\!)$ above:

$$\begin{array}{lll}
(sctx) & \mathsf{ctx}\,\mathsf{A} & = (\!|\mathsf{A}|\!) \\
(scomp) & (l_1, r_1) \bullet (l_2, r_2) & = (l_1 + (r_1 * l_2), r_1 * r_2) \\
(sapp) & \mathsf{app}\,(l, r)\,e & = l + r * e
\end{array}$$

We can show the context laws hold for these definitions:

$$\begin{array}{lll}
 & \mathsf{app}\,((l_1, r_1) \bullet (l_2, r_2))\,e & \\
= & \mathsf{app}\,(l_1 + (r_1 * l_2), r_1 * r_2)\,e & \{\ (scomp)\ \} \\
= & (l_1 + (r_1 * l_2)) + (r_1 * r_2) * e & \{\ (sapp)\ \} \\
= & l_1 + r_1 * (l_2 + r_2 * e) & \{\ \mathit{assoc\ and\ distr.}\ \} \\
= & \mathsf{app}\,(l_1, r_1)\,(\mathsf{app}\,(l_2, r_2)\,e) & \{\ (sapp)\ \}
\end{array}$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

and

$$\begin{array}{lll} & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{A})\,e & \\ = & \mathsf{app}\,(\!|\mathsf{A}|\!)\,e & \{\ (sctx)\ \} \\ = & l + r * e & \{\ (sapp), for\,(l, r) = (\!|\mathsf{A}|\!)\ \} \end{array}$$

We proceed by induction over A (where we compress some cases for brevity):

case $\mathsf{A} = \square$:

$$\begin{array}{lll} = & l + r * e & \{\ (\!|\square|\!) = (l, r)\ \} \\ = & unit^{+} + unit^{*} * e & \{\ fold\ \} \\ = & e & \{\ unit\ \} \\ = & \square[e] & \{\ \square\ \} \end{array}$$

and $\mathsf{A} = v_1 + v_2 * \mathsf{A}'$:

$$\begin{array}{lll} = & l + r * e & \{\ (\!|v_1 + v_2 * \mathsf{A}'|\!) = (l, r)\ \} \\ = & (v_1 + v_2 * l') + (v_2 * r') * e & \{\ (\!|\mathsf{A}'|\!) = (l', r')\ \} \\ = & v_1 + v_2 * (l' + r' * e) & \{\ assoc.\,and\,distr\ \} \\ = & v_1 + v_2 * \mathsf{A}'[e] & \{\ induction\,hyp.\ \} \\ = & \mathsf{A}[e] & \{\ \mathsf{A}\,context\ \} \end{array}$$

When we apply this to the `hash` function, we derive the tail-recursive version as:

$$hash'\,xs\,k = \mathsf{match}\,xs\,\{\,Cons\,x\,xx \to hash'\,xx\,(k \bullet (\mathsf{ctx}\,(x + 31 * \square)));\ Nil \to \mathsf{app}\,k\,17\,\}$$

which further simplifies to:

$$hash'\,xs\,(l, r) = \mathsf{match}\,xs\,\{\,Cons\,x\,xx \to hash'\,xx\,(l + r * x,\ r * 31);\ Nil \to l + r * 17\,\}$$

The final definition may not be quite so obvious, and we argue that the modulo *semiring* instantiation may be a nice addition to any optimizing compiler. Indeed, it turns out that GCC implements this optimization (Dvořák, 2004) for integers and floating point numbers (if `-fast-math` is enabled to allow the assumption of associativity). This implementation specifically creates two local accumulators for addition and multiplication and uses a direct while loop to compile the tail-recursive calls.

4.7 Modulo exponent contexts

As a final example of an efficient representation of contexts, we consider *exponent* contexts that consist of a sequence of calls to a function g:

$$\mathsf{E} ::= \square \mid g\,\mathsf{E}$$

If we use a defunctionalized evaluation context from Section 4.3, we derive a datatype that is isomorphic to the peano-encoded natural numbers: the continuation *counts* how often we still have to apply g. As such, we can represent it more efficiently by an integer, where we fold an evaluation context into a count:

$$\begin{array}{ll} (\!|\square|\!) & = 0 \\ (\!|g\,\mathsf{A}|\!) & = (\!|\mathsf{A}|\!) + 1 \end{array}$$

We can define the primitive operations as:

$$\begin{array}{lll} (xctx) & \mathsf{ctx}\,\mathsf{A} & = (\!|\mathsf{A}|\!) \\ (xcomp) & k_1 \bullet k_2 & = k_1 + k_2 \\ (xapp) & \mathsf{app}\,0\,e & = e \\ (xapp) & \mathsf{app}\,(k+1)\,e & = \mathsf{app}\,k\,(g\,e) \end{array}$$

where $\mathsf{app}\,k\,e$ applies the function g to its argument k times. See Appendix B.5 in the supplement for the derivations that show the context laws hold for this definition.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Note that if g is the enclosing function f, then the (*xapp*) specification is not tail-recursive. In that case, we can again use specification (b) to replace $\mathsf{app}\,k\,(g\,e)$ by $[\![g\,e]\!]_{f,k}$ at compile time (as shown in Section 4.2). A nice example of such an exponent context is given by Wand (1980) who considers McCarthy's 91-function:

$$g\,x = \mathsf{if}\,x > 100\,\mathsf{then}\,x - 10\,\mathsf{else}\,g\,(g\,(x+11))$$

Using the exponent context with the recursive (*xapp*), we obtain a mutually tail-recursive version:

$$\begin{aligned} g'\,x\,k &= \mathsf{if}\,x > 100\,\mathsf{then}\,app\,k\,(x-10)\,\mathsf{else}\,g'\,(x+11)\,(k+1) \\ app\,k\,e &= \mathsf{if}\,k = 0\,\mathsf{then}\,e\,\mathsf{else}\,g'\,e\,(k-1) \end{aligned}$$

5 Context composition

While the contexts we have defined so far are useful when they apply, they can fall short if they only match *some* of the recursive calls. This makes them fragile when a new recursive call is added to a function, as the context may no longer apply. In this section, we remove this restriction by showing how fast but restricted contexts can be composed with slower more general ones. This is not implemented in Koka though.

5.1 A basic expression evaluator

To motivate the composition of contexts, we consider a basic arithmetic expression evaluator in the style of Hutton (2021):

```
type expr
  Lit(lit : int)
  Add(e1 : expr, e2 : expr)

fun eval(e)
  match e
    Add(e1,e2) -> eval(e1) + eval(e2)
    Lit(n)     -> n
```

The `+` suggests the use of a monoid context. However, this does not apply directly, since we have two recursive calls to `eval` instead of just one. The best we can do is to ignore the first recursive call and treat it as a regular value. Then we would obtain:

```
fun eval(e, addacc)
  match e
    Add(e1,e2) -> eval(e2, eval(e1,addacc))
    Lit(n)     -> addacc + n
```

However, we have not quite achieved a tail-recursive version yet. Like Hutton (2021), we can achieve this by using defunctionalized evaluation contexts:

```
type accum
  Hole
  AddL(k : accum, e : expr)
  AddR(addacc : int, k : accum)
```

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

```
fun eval(e, acc)
  match e
    Add(e1,e2) -> eval(e1, AddL(acc,e2))
    Lit(n)     -> app(acc, n)

fun app(acc, result)
  match acc
    Hole          -> result
    AddL(k,e)      -> eval(e, AddR(result,k))
    AddR(addacc,k) -> app(k, addacc + result)
```

This version is now tail-recursive, but it is also more complex than the original version and involves the allocation of `AddL` and `AddR` constructors. In particular, the `AddR` constructor seems superfluous, as it corresponds to the context `app(k,addacc + eval(e2))`, which we optimized using the monoid contexts earlier. In this section, we want to combine the two approaches to obtain a more efficient version, where we use *both* an accumulator and a monoid context:

```
type accum
  Hole
  AddL(k : accum, e : expr)

fun eval(e, acc, addacc)
  match e
    Add(e1,e2) -> eval(e1, AddL(acc,e2), addacc)
    Lit(n)     -> app(acc, addacc + n)

fun app(acc, result)
  match acc
    Hole      -> result
    AddL(k,e) -> eval(e,k,result)
```

This version has the best of both worlds: it is fully tail-recursive and only needs to allocate a defunctionalized continuation for the left recursive call (where we need to keep track of the expression `e2`), while the right recursive call is efficiently handled by the monoid context.

5.2 Swapping contexts

To achieve this transformation in general, we need to be able to compose two contexts. For two contexts E_1 and E_2, we can define their product context, which consists of tuples of the two contexts. We can apply a product context to an expression by applying each context in turn:

$$(papp) \quad \mathsf{app}\,(l,\, r)\,e = \mathsf{app}\,l\,(\mathsf{app}\,r\,e)$$

But how would we compose two product contexts? We would like to turn a composition of tuples into a tuple of compositions as $(l_1,\, r_1) \bullet (l_2,\, r_2) = (l_1 \bullet l_2,\, r_1 \bullet r_2)$. We can try to calculate this directly:

$$\begin{array}{lll} & \mathsf{app}\,((l_1,\, r_1) \bullet (l_2,\, r_2))\,e & \\ = & \mathsf{app}\,(l_1,\, r_1)\,(\mathsf{app}\,(l_2,\, r_2)\,e) & \{\ (appcomp)\ \} \\ = & \mathsf{app}\,l_1\,(\mathsf{app}\,r_1\,(\mathsf{app}\,l_2\,(\mathsf{app}\,r_2\,e))) & \{\ (papp)\ \} \end{array}$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

... but now we are stuck. Here, l_1, l_2 belong to the context E_1 and r_1, r_2 to E_2. In order to make progress, we have to *swap* the inner contexts r_1 and l_2. But this is not always going to be possible! Instead, we need to parameterize the product context with a swap operation:

$$(\mathit{appswap}) \quad \mathsf{app}\, l\, (\mathsf{app}\, r\, e) \quad = \mathsf{app}\, r'\, (\mathsf{app}\, l'\, e) \quad \mathsf{where}\, (l', r') = \mathsf{swap}(l, r)$$

If the contexts are connected in this sense, we can continue to calculate their composition:

$$\begin{array}{lll} & \mathsf{app}\, l_1\, (\mathsf{app}\, r_1\, (\mathsf{app}\, l_2\, (\mathsf{app}\, r_2\, e))) & \\ = & \mathsf{app}\, l_1\, (\mathsf{app}\, l_2'\, (\mathsf{app}\, r_1'\, (\mathsf{app}\, r_2\, e))) & \{\ \mathit{swap}(r_1, l_2) = (l_2', r_1')\ \} \\ = & \mathsf{app}\, (l_1 \bullet l_2')\, (\mathsf{app}\, (r_1' \bullet r_2)\, e) & \{\ (\mathit{appcomp})\ \} \\ = & \mathsf{app}\, (l_1 \bullet l_2', r_1' \bullet r_2)\, e & \{\ (\mathit{papp})\ \} \end{array}$$

This gives us a definition for product contexts: we can fold any context $\mathsf{E} = \mathsf{E}_1 \mid \mathsf{E}_2$ by composing the folds of E_1 and E_2:

$$\begin{array}{ll} (\!|\Box|\!) & = (\mathsf{ctx}\, \Box,\ \mathsf{ctx}\, \Box) \\ (\!|\mathsf{E}_1[e]|\!) & = (\mathsf{ctx}\, \mathsf{E}_1,\ \mathsf{ctx}\, \Box) \bullet (\!|e|\!) \\ (\!|\mathsf{E}_2[e]|\!) & = (\mathsf{ctx}\, \Box,\ \mathsf{ctx}\, \mathsf{E}_2) \bullet (\!|e|\!) \end{array}$$

$$\begin{array}{lll} (\mathit{pctx}) & \mathsf{ctx}\, \mathsf{E} & = (\!|\mathsf{E}|\!) \\ (\mathit{pcomp}) & (l_1, r_1) \bullet (l_2, r_2) & = (l_1 \bullet l_2', r_1' \bullet r_2) \quad \mathsf{where}\, (l_2', r_1') = \mathsf{swap}(r_1, l_2) \\ (\mathit{papp}) & \mathsf{app}\, (k_1, k_2)\, e & = \mathsf{app}\, k_1\, (\mathsf{app}\, k_2\, e) \end{array}$$

With this definition in hand, we can now derive several contexts from the previous section from the more basic contexts we defined earlier.

5.2.1 Modulo monoid contexts

We motivated the Modulo Monoid Contexts in Section 4.5 as the composition of a left-biased and a right-biased context. In fact, we can now *derive* this context as the product context of a left-biased and right-biased contexts, with $\mathit{swap}(l, r) = (r, l)$. This follows the swap law since:

$$\begin{array}{ll} & \mathsf{app}\, l\, (\mathsf{app}\, r\, e) \\ = & l \odot (e \odot r) \\ = & (l \odot e) \odot r \\ = & \mathsf{app}\, r\, (\mathsf{app}\, l\, e) \end{array}$$

With this, we get exactly the previous definition of (*acomp*) of Modulo Monoid Contexts.

5.2.2 Modulo semiring contexts

Similarly, we can derive the semiring context (Section 4.6) as the composition of two left-biased contexts for its addition (l) and multiplication (r). Here, the swap operation is given by $\mathit{swap}(r, l) = (r * l, r)$:

$$\begin{array}{ll} & \mathsf{app}_*\, r\, (\mathsf{app}_+\, l\, e) \\ = & r * (l + e) \\ = & (r * l) + (r * e) \\ = & \mathsf{app}_+\, (r * l)\, (\mathsf{app}_*\, r\, e) \end{array}$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

With this, we get exactly the previous definition of (*scomp*) and it our new context corresponds to the semiring contexts we defined earlier. For this definition to work, it is important though that the left-biased context for the addition is in the first component of the tuple with multiplication in the second. That allows us to define a swap operation that uses the distributivity of the semiring to swap the contexts. We could not define a swap operation if multiplication is in the first component, since this would require us to move an addition under a multiplication, which is only possible if the semiring has multiplicative inverses.

5.3 Composing (defunctionalized) evaluation contexts

(Defunctionalized) evaluation contexts are the only contexts introduced in the last section that can reliably make all recursive calls tail-recursive. For this reason, they are particularly attractive for composition with other contexts that lead to faster code in practice but only apply in more limited cases. Thankfully, this is easily possible, since we can swap an arbitrary context r with a general evaluation context l by storing it in a closure:

$$\mathit{swap}\,(r,\, l) = ((\lambda x.\, \mathsf{app}\, r\, x) \bullet l,\, \mathsf{ctx}\,\Box)$$

We can verify that this definition fulfills the swap law:

$$\begin{array}{lll}
 & \mathsf{app}\, r\, (\mathsf{app}\, l\, e) & \\
= & (\lambda x.\, \mathsf{app}\, r\, x)\, (\mathsf{app}\, l\, e) & \{\ \textit{eta expansion}\ \} \\
= & \mathsf{app}\, (\lambda x.\, \mathsf{app}\, r\, x)\, (\mathsf{app}\, l\, e) & \{\ (\textit{eapp})\ \} \\
= & \mathsf{app}\, ((\lambda x.\, \mathsf{app}\, r\, x) \bullet l)\, e & \{\ (\textit{appcomp})\ \} \\
= & \mathsf{app}\, ((\lambda x.\, \mathsf{app}\, r\, x) \bullet l)\, \Box[e] & \\
= & \mathsf{app}\, ((\lambda x.\, \mathsf{app}\, r\, x) \bullet l)\, (\mathsf{app}\, (\mathsf{ctx}\,\Box)\, e) & \{\ (\textit{appctx})\ \}
\end{array}$$

The same approach can also be used for defunctionalized evaluation contexts. Analogous to creating a fresh closure, we could create a special constructor to store an application of the other context. However, to avoid allocations and to enable a nested translation (Section 4.2), we integrate the restricted context into the constructors.

We define the extended accumulator datatype by creating a constructor H for the $\Box$ context and for each E_i a constructor A_i that carries the free variables of E_i *and the inner context* k'. The compiler then generates an app function where we interpret A_i by evaluating E_i with the stored free variables:

$$\begin{array}{lll}
(\textit{dctx}) & \mathsf{ctx}\,\mathsf{E}_i & = \mathsf{A}_i\, x_1 \ldots x_m\, (\mathsf{ctx}\,\Box)\, \mathsf{H} \quad \text{where}\, x_1, \ldots, x_m = \mathsf{fv}(\mathsf{E}_i) \\
(\textit{dcomp}) & k_1 \bullet \mathsf{H} & = k_1 \\
(\textit{dcomp}) & k_1 \bullet (\mathsf{A}_i\, x_1 \ldots x_m\, k'\, k_2) & = \mathsf{A}_i\, x_1 \ldots x_m\, k'\, (k_1 \bullet k_2) \\
(\textit{dapp}) & \mathsf{app}\, \mathsf{H}\, e & = e \\
(\textit{dapp}) & \mathsf{app}\, (\mathsf{A}_i\, x_1 \ldots x_m\, k'\, k)\, e & = [\![\mathsf{E}_i[e,\, x_1 \ldots x_m]]\!]_{f(k,k')}
\end{array}$$

Then we can define the swap operation as:

$$\begin{array}{l}
\mathit{swap}\,(k_1,\, \mathsf{H}) = (\mathsf{H},\, k_1) \\
\mathit{swap}\,(k_1,\, \mathsf{A}_i\, x_1 \ldots x_m\, k'\, k_2) = (\mathsf{A}_i\, x_1 \ldots x_m\, (k_1 \bullet k')\, k_2,\, \mathsf{ctx}\,\Box)
\end{array}$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

This definition again fulfills the swap law. Case H:

$$\begin{array}{lll} & \mathsf{app}\, r\, (\mathsf{app}\, \mathsf{H}\, e) & \\ = & \mathsf{app}\, r\, e & \{\ (\mathit{dapp})\ \} \\ = & \mathsf{app}\, \mathsf{H}\, (\mathsf{app}\, r\, e) & \{\ (\mathit{dapp})\ \} \end{array}$$

Case $\mathsf{A}_i\, x_1 \ldots x_m\, k'\, k_2$:

$$\begin{array}{lll} & \mathsf{app}\, r\, (\mathsf{app}\, (\mathsf{A}_i\, x_1 \ldots x_m\, r'\, k_2)\, e) & \\ = & \mathsf{app}\, r\, (\llbracket \mathsf{E}_i[e,\, x_1 \ldots x_m] \rrbracket_{f(r',\, k_2)}) & \{\ (\mathit{dapp})\ \} \\ = & \mathsf{app}\, r\, (\mathsf{app}\, (r',\, k_2)\, (\mathsf{E}_i[e,\, x_1 \ldots x_m])) & \{\ \mathit{specification}\, (b)\ \} \\ = & \mathsf{app}\, r\, (\mathsf{app}\, r'\, (\mathsf{app}\, k_2\, (\mathsf{E}_i[e,\, x_1 \ldots x_m]))) & \{\ (\mathit{papp})\ \} \\ = & \mathsf{app}\, (r \bullet r')\, (\mathsf{app}\, k_2\, (\mathsf{E}_i[e,\, x_1 \ldots x_m])) & \{\ (\mathit{appcomp})\ \} \\ = & \mathsf{app}\, (r \bullet r',\, k_2)\, (\mathsf{E}_i[e,\, x_1 \ldots x_m]) & \{\ (\mathit{papp})\ \} \\ = & \llbracket \mathsf{E}_i[e,\, x_1 \ldots x_m] \rrbracket_{f(k_2,\, r \bullet r')} & \{\ \mathit{specification}\, (b)\ \} \\ = & \mathsf{app}\, (\mathsf{A}_i\, x_1 \ldots x_m\, (r \bullet r')\, k_2)\, e & \{\ (\mathit{dapp})\ \} \\ = & \mathsf{app}\, (\mathsf{A}_i\, x_1 \ldots x_m\, (r \bullet r')\, k_2)\, \Box[e] & \\ = & \mathsf{app}\, (\mathsf{A}_i\, x_1 \ldots x_m\, (r \bullet r')\, k_2)\, (\mathsf{app}\, (\mathsf{ctx}\, \Box)\, e) & \{\ (\mathit{appctx})\ \} \end{array}$$

5.4 Extending the expression evaluator

We can use this insight to derive a tail-recursive expression evaluator which supports multiplication as well, where we compose a defunctionalized evaluation context with a semiring context. First, we add a new constructor `Mul(e1,e2)` to our expression datatype which encodes the multiplication `eval(e1) * eval(e2)`. We then create a datatype `accum` which stores the defunctionalized evaluation contexts when descending into the first expression `e1`. These constructors contain both the second expression `e2` and the semiring context `(a, m)`. When descending into `e1`, we store the current semiring context in the constructor and continue with the semiring context $\mathsf{ctx}\, \Box = (0,\, 1)$:

$$\begin{array}{lll} & \mathsf{app}\, (\mathit{acc},\, (a,\, m))\, (\mathsf{eval}\, e_1 \,+\, \mathsf{eval}\, e_2) & \\ = & \mathsf{eval}'\, ((\mathit{acc},\, (a,\, m)) \bullet (\mathsf{AddL}\, 0\, 1\, \mathsf{Hole}\, e_2,\, (0,\, 1)))\, e_1 & \{\ (\mathit{tail})\ \} \\ = & \mathsf{eval}'\, (\mathit{acc} \bullet \mathsf{AddL}\, a\, m\, \mathsf{Hole}\, e_2,\, (0,\, 1) \bullet (0,\, 1))\, e_1 & \{\ (\mathit{pcomp})\ \} \\ = & \mathsf{eval}'\, (\mathsf{AddL}\, a\, m\, \mathit{acc}\, e_2,\, (0,\, 1))\, e_1 & \{\ (\mathit{dcomp})\ \} \end{array}$$

This calculation directly follows the recipe for composing with defunctionalized evaluation contexts and can thus be derived algorithmically. Our full implementation becomes:

```
type expr
  Lit(lit : int)
  Add(e1 : expr, e2 : expr)
  Mul(e1 : expr, e2 : expr)

type accum
  Hole
  AddL(a : int, m : int, k : accum, e : expr)
  MulL(a : int, m : int, k : accum, e : expr)

fun eval(e, acc, a, m)
  match e
    Add(e1,e2) -> eval(e1, AddL(a,m,acc,e2), 0, 1)
    Mul(e1,e2) -> eval(e1, MulL(a,m,acc,e2), 0, 1)
    Lit(n)     -> app(acc, a + m * n)
```

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

```
fun app(acc, n)
  match acc
    Hole          -> n
    AddL(a,m,k,e) -> eval(e,k, a + m * n, m)
    MulL(a,m,k,e) -> eval(e,k,a, m * n)
```

In contrast, a translation using just defunctionalized evaluation contexts would require two more constructors `AddR` and `MulR` (just as in the basic example). Unfortunately though, now that we store the semiring context in `accum`, our constructors carry a few more elements than the constructors of `expr`. In a language like Koka, which can reuse constructors of equal size (Xie *et al*., 2021; Lorenzen & Leijen, 2022), it would be preferable to obtain constructors of the same size as `expr`, since we could then hope to avoid the allocation of `AddL` and `MulL` by reusing the memory of `Add` and `Mul`. This is often possible when using non-composed defunctionalized evaluation contexts and Lorenzen *et al*. (2023) show that it is guaranteed to work if the original function has the shape of a map or fold (like `eval`). Alas, the same is not true for composed contexts, since we need to store the additional semiring context. However, using composed contexts like here can still avoid allocations in languages that lack reuse analysis.

Finally, we can reduce the number of elements stored in the constructors and obtain a more natural version of the evaluator by using the distributivity law to push the semiring context into the expression. For the `Add` case, we calculate:

$$
\begin{array}{lll}
& \mathsf{app}\,(acc,\ (a,\ m))\,(\mathsf{eval}\,e_1\ +\ \mathsf{eval}\,e_2) & \\
= & \mathsf{app}\,acc\,(\mathsf{app}\,(a,\ m)\,(\mathsf{eval}\,e_1\ +\ \mathsf{eval}\,e_2)) & \{\ (papp)\ \} \\
= & \mathsf{app}\,acc\,(a\ +\ m\ *\ (\mathsf{eval}\,e_1\ +\ \mathsf{eval}\,e_2)) & \{\ (sapp)\ \} \\
= & \mathsf{app}\,acc\,(a\ +\ m\ *\ \mathsf{eval}\,e_1\ +\ m\ *\ \mathsf{eval}\,e_2) & \{\ distributivity\ \}
\end{array}
$$

At this point, the recursive call to $\mathsf{eval}\,e_1$ is under a semiring context (a, m) and an evaluation context $\square\ +\ m\ *\ \mathsf{eval}\,e_2$. We thus have to store an extra `m` in our `AddL` constructor:

$$
\begin{array}{lll}
& \mathsf{eval}'\,((acc,\ (0,1))\bullet(\mathsf{AddL}\ 0\ 1\ \mathsf{Hole}\,m\,e_2,\ (a,m)))\,e_1 & \{\ (tail)\ \} \\
= & \mathsf{eval}'\,(acc\bullet\mathsf{AddL}\ 0\ 1\ \mathsf{Hole}\,m\,e_2,\ (0,1)\bullet(a,\ m))\,e_1 & \{\ swap\ \} \\
= & \mathsf{eval}'\,(\mathsf{AddL}\ 0\ 1\ acc\,m\,e_2,\ (a,\ m))\,e_1 & \{\ (dcomp)\ and\ (scomp)\ \}
\end{array}
$$

It turns out that in this version, the semiring context stored in the accumulated datatype is always going to be $(0, 1)$, so we can simplify the definition by omitting it. We thus obtain the implementation:

```
type accum
  Hole
  AddL(m : int, k : accum, e : expr)
  MulL(a : int, k : accum, e : expr)

fun eval(e, acc, a, m)
  match e
    Add(e1,e2) -> eval(e1, AddL(m,acc,e2), a, m)
    Mul(e1,e2) -> eval(e1, MulL(a,acc,e2), 0, m)
    Lit(n)     -> app(acc, a + m * n)

fun app(acc, n)
  match acc
    Hole        -> n
    AddL(m,k,e) -> eval(e,k,n,m)
    MulL(a,k,e) -> eval(e,k,a,n)
```

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

This version is slightly more efficient than the previous one, but the constructors are still too big for reuse analysis to apply. Furthermore, it is unclear whether we can derive this algorithmically as well. Instead, we will stick with the more general version that can be derived directly from the composition of contexts and extend it to derive an evaluator that can also handle subtraction and division.

To extend the expression evaluator to support division, we might try to add a new context for division and show how to compose it with the semiring context. However, this is not straightforward, since $(a + \Box)^{-1}$ cannot be simplified to $a' + \Box^{-1}$ for any other a': the inverse of the sum depends on the $\Box$, which is not yet known, and there is no general rule for exchanging the inverse operation with addition. Instead, we need to use an idea from the theory of continued fractions.

5.5 Aside: Continued fractions

Continued fractions are a representation of the rational (or real) numbers that arises from the Euclidean algorithm. They consist of a sequence of nested additions and fractions with numerator 1. For example, we can calculate the continued fraction of 4.24 as:

$$\begin{aligned} 4.24 &= 4 + \tfrac{24}{100} \\ &= 4 + \frac{1}{\frac{100}{24}} = 4 + \frac{1}{4 + \frac{4}{24}} = 4 + \frac{1}{4 + \frac{1}{6}} = 4 + \frac{1}{4 + \frac{1}{\frac{6}{1}}} = 4 + \frac{1}{4 + \frac{1}{5 + \frac{1}{1}}} \end{aligned}$$

We can write such a (long-form) continued fraction (with the final '1' left implicit) as [4, 4, 5]. Then we can compute its floating point representation with a simple recursive algorithm:

```
fun frac(xs)
  match xs
    Nil        -> 1
    Cons(a,xx) -> a + (1/frac(xx))
```

This algorithm is not tail-recursive, and it might be quite difficult to make tail-recursive without further insight (and without resorting to general evaluation contexts). However, it is well known that continued fractions can be calculated by their *convergents*, which is a sequence h_n, k_n with $\mathsf{frac}([a_0, \ldots, a_n]) = h_n / k_n$. The convergents start with $h_{-2} = 0$, $h_{-1} = 1$, $k_{-2} = 1$, and $k_{-1} = 0$ and are further calculated by:

$$\frac{h_n}{k_n} = \frac{a_n * h_{n-1} + h_{n-2}}{a_n * k_{n-1} + k_{n-2}}$$

This gives us a tail-recursive algorithm to calculate the continued fraction (where we write `h1` for h_{n-1}, `h2` for h_{n-2} and equivalent for `k1` and `k2`):

```
fun frac'(xs, h1, k1, h2, k2)
  match xs
    Nil        -> (1*h1 + h2) / (1*k1 + k2)
    Cons(a,xx) -> frac'(xx, a*h1 + h2, h1, a*k1 + k2, k1)

fun frac(xs)
  frac'(xs,1,0,0,1)
```

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

It turns out that we can use the same idea to define a general context that applies to arbitrary sequences of addition, multiplication, and inverses.

5.6 Modulo fields context

Using the insight from continued fractions, we can define general *field contexts* that support not only addition and multiplication but also additive and multiplicative inverses. We define the context F as:

$$F ::= \Box \mid a + F \mid m * F \mid F^{-1}$$

It turns out that we use the same convergent representation as for continued fractions, where we keep four numbers:

$$\frac{x * h_1 + h_2}{x * k_1 + k_2}$$

and define the fold operation as:

$$(\!|\Box|\!) = \frac{x*1+0}{x*0+1} \qquad (\!|m * F|\!) = \frac{x*m+0}{x*0+1} \bullet (\!|F|\!)$$

$$(\!|a + F|\!) = \frac{x*1+a}{x*0+1} \bullet (\!|F|\!) \qquad (\!|F^{-1}|\!) = \frac{x*0+1}{x*1+0} \bullet (\!|F|\!)$$

We can apply a field context to an expression by substituting the expression for x. Similarly, we can compose two field contexts by substituting the second context into the first context and simplifying the expression. Our context is defined as:

$$\begin{array}{lll} (fctx) & \mathsf{ctx}\, F & = (\!|F|\!) \\ (fcomp) & \frac{x*h_1+h_2}{x*k_1+k_2} \bullet \frac{y*h_1'+h_2'}{y*k_1'+k_2'} & = \frac{y*(h_1'*h_1+k_1'*h_2)+(h_2'*h_1+k_2'*h_2)}{y*(h_1'*k_1+k_1'*k_2)+(h_2'*k_1+k_2'*k_2)} \\ (fapp) & \mathsf{app}\, \frac{x*h_1+h_2}{x*k_1+k_2}\, e & = \frac{e*h_1+h_2}{e*k_1+k_2} \end{array}$$

and the context laws hold.

5.7 An advanced expression evaluator

Using the field contexts, we can extend our expression evaluator to support arbitrary field operations. Our implementation arises directly from the obvious expression evaluator which folds the expression into a rational number:

```
fun eval(e : expr) : rat
  match e
    Add(e1,e2) -> eval(e1) + eval(e2)
    Mul(e1,e2) -> eval(e1) * eval(e2)
    Neg(e1)    -> from-int(-1) * eval(e1)
    Inv(e1)    -> from-int(1) / eval(e1)
    Lit(n)     -> from-int(n)
```

We can directly define the field context as a datatype, where we define `empty()`, `add(a)`, `mul(a)`, and `inv()` to correspond to the fold operations:

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

```
type fctx
  Fctx( h1 : rat, h2 : rat, k1 : rat, k2 : rat )

fun fapp(f : fctx, r : rat) : rat
  (f.h1*r + f.h2) / (f.k1*r + f.k2)

fun add(a : rat) : fctx
  Fctx( from-int(1), a, from-int(0), from-int(1) )

...
```

Then we use the TRMC algorithm with the composition of defunctionalized contexts and field contexts to obtain a tail-recursive version that uses a field context for the field operations and a defunctionalized context for the recursive calls that leave an expression to be evaluated:

```
type accum
  Hole
  AddL(f : fctx, k : accum, e : expr)
  MulL(f : fctx, k : accum, e : expr)

fun eval(e : expr, acc : accum, f : fctx)
  match e
    Add(e1,e2) -> eval(e1, AddL(f, acc, e2), empty())
    Mul(e1,e2) -> eval(e1, MulL(f, acc, e2), empty())
    Neg(e1)    -> eval(e1, acc, comp(f, mul(-1)))
    Inv(e1)    -> eval(e1, acc, comp(f, inv()))
    Lit(n)     -> app(acc, fapp(f, from-int(n)))

fun app(acc : accum, r : rat)
  match acc
    Hole        -> r
    AddL(f,k,e) -> eval(e, k, comp(f, add(r)))
    MulL(f,k,e) -> eval(e, k, comp(f, mul(r)))
```

The final derived program is actually quite sophisticated and fully tail-recursive. We believe that deriving this algorithm manually would be nontrivial. Moreover, it only allocates a small amount of memory while descending the left-spine. In contrast, a simple application of defunctionalized contexts without field contexts would require us to allocate a constructor even in the `Neg` and `Inv` cases, which would be less efficient.

6 Modulo constructor contexts

As shown in the introduction, the most interesting instantiation is of course the modulo *cons* transformation on constructor contexts. We can define a constant constructor context K as:

$$\mathsf{K} \quad ::= \quad \Box \mid C^k\, v_1 \ldots \mathsf{K} \ldots v_k$$

We define the ($\star$) condition in the TRMC translation to restrict the context E to K contexts only. A possible way to define the contexts is to directly use K as a runtime context:

$$\begin{array}{lll} (\mathit{kctx}) & \mathsf{ctx}\,\mathsf{K} & = \mathsf{K} \\ (\mathit{kcomp}) & \mathsf{K}_1 \bullet \mathsf{K}_2 & = \mathsf{K}_1[\mathsf{K}_2] \\ (\mathit{kapp}) & \mathsf{app}\,\mathsf{K}\,e & = \mathsf{K}[e] \end{array}$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Similar to general evaluation contexts (Section 4.1), the context laws hold trivially for such definition (Appendix B.6 in the supplement) – and just as with general evaluation contexts, the *map* function translates to:

$$
\begin{array}{l}
map'\,xs\,f\,k \;=\; \mathsf{match}\,xs\,\{ \\
\quad Nil \to \mathsf{app}\,k\,Nil \\
\quad Cons\,x\,xx \to \mathsf{let}\,y = f\,x\,\mathsf{in}\,map'\,xx\,f\,(k \bullet (\mathsf{ctx}\,(Cons\,y\,\Box)))\,\}
\end{array}
$$

Even though this is a valid instantiation, it does not yet imply that this can be efficient. In particular, the composition $\mathsf{K}_1[\mathsf{K}_2]$ could a fresh context every time and it may be difficult to implement such substitution efficiently at runtime as it needs to copy K_1 along the path to the hole. What we are looking for instead is an *in-place updating* instantiation that can compose in constant time.

6.1 Minamide

Minamide (1998) presents a "hole calculus" that can directly express our contexts in a functional way but also allows an efficient in-place updating implementation. Using the hole calculus as our target calculus, we can instantiate the translation function using Minamide's system.

We define the context type as a "hole function" $(\hat{\lambda}x.\,e)$, where $ctx\,\alpha \equiv hfun\,\alpha\,\alpha.$ and instantiate the context operations to use the primitives as given by Minamide (1998):

$$
\begin{array}{lll}
(hctx) & \mathsf{ctx}\,\mathsf{K} & = \hat{\lambda}x.\,\mathsf{K}[x] \\
(hcomp) & k_1 \bullet k_2 & = hcomp\,k_1\,k_2 \\
(happ) & \mathsf{app}\,k\,e & = happ\,k\,e
\end{array}
$$

Satisfyingly, our primitives turn out to map directly to the hole calculus primitives. The reduction rules for *happ* and *hcomp* specialized to our calculus are (Minamide, 1998, fig. 5):

$$
\begin{array}{lll}
(happly) & happ\,(\hat{\lambda}x.\,\mathsf{K})\,v & \longrightarrow \mathsf{K}[x{:=}v] \\
(hcompose) & hcomp\,(\hat{\lambda}x.\,\mathsf{K}_1)\,(\hat{\lambda}y.\,\mathsf{K}_2) & \longrightarrow \hat{\lambda}y.\,\mathsf{K}_1[x{:=}\mathsf{K}_2]
\end{array}
$$

This means that for any context k, we have $k \cong \hat{\lambda}x.\,\mathsf{K}[x]$ **(1)**. We can now show that our context laws are satisfied for this system, with composition:

$$
\begin{array}{lll}
 & \mathsf{app}\,(k_1 \bullet k_2)\,e & \\
= & \mathsf{app}\,(hcomp\,k_1\,k_2)\,e & \{\ (hcomp)\ \} \\
= & happ\,(hcomp\,k_1\,k_2)\,e & \{\ (happ)\ \} \\
\cong & happ\,(hcomp\,(\hat{\lambda}x.\,\mathsf{K}_1[x])\,(\hat{\lambda}y.\,\mathsf{K}_2[y]))\,e & \{\ (1),\ \mathbf{2}\ \} \\
\cong & happ\,(\hat{\lambda}y.\,\mathsf{K}_1[x][x{:=}\mathsf{K}_2[y]])\,e & \{\ (hcomp)\ \} \\
\cong & (\mathsf{K}_1[x][x{:=}\mathsf{K}_2[y]])[y{:=}e] & \{\ (happly)\ \} \\
= & \mathsf{K}_1[\mathsf{K}_2[e]] & \{\ contexts\ \} \\
\cong & \mathsf{K}_1[happ\,(\hat{\lambda}y.\,\mathsf{K}_2[y])\,e] & \{\ (happly)\ \} \\
\cong & happ\,(\hat{\lambda}x.\,\mathsf{K}_1[x])\,(happ\,(\hat{\lambda}y.\,\mathsf{K}_2[y])\,e) & \{\ (happly)\ \} \\
\cong & happ\,k_1\,(happ\,k_2\,e) & \{\ (1),\ (2)\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k_2\,e) & \{\ (happ)\ \}
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

and application:

$$\begin{array}{lll} & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{K})\,e & \\ = & \mathsf{app}\,(\hat{\lambda}x.\,\mathsf{K}[x])\,e & \{\ (hctx)\ \} \\ = & \mathit{happ}\,(\hat{\lambda}x.\,\mathsf{K}[x])\,e & \{\ (happ)\ \} \\ \cong & \mathsf{K}[x][x{:=}e] & \{\ (happly)\ \} \\ = & \mathsf{K}[e] & \{\ contexts\ \} \end{array}$$

The hole calculus is restricted by a linear-type discipline where the contexts $ctx\,\alpha \equiv hfun\,\alpha\,\alpha$ have a linear type. This is what enables an efficient in-place update implementation while still having a pure functional interface. For our needs, we need to check separately that the translation ensures that all uses of a context k are indeed linear. Type judgments in Minamide's system (Minamide, 1998, fig. 4) are denoted as $\Gamma\,;\,H \vdash_{\mathrm{M}} e : \tau$ where Γ is the normal type environment, and H for linear bindings containing at most one linear value. The type environment Γ can itself contain linear values with a linear type (like *hfun*) but only pass those linearly to a single premise. The environment restricted to nonlinear values is denoted as $\Gamma|_{\mathsf{N}}$. We can now show that our translation can indeed be typed under the linear type discipline:

Theorem 3. (*TRMC uses contexts linearly*)
If $\Gamma|_{\mathsf{N}}\,;\,\varnothing \vdash_{\mathrm{M}} \mathsf{fun} f = \lambda x_1 \ldots x_n.\, e : \tau_1 \to \ldots \to \tau_n \to \tau$ and k fresh
then $\Gamma|_{\mathsf{N}}, f\,;\,\varnothing \vdash_{\mathrm{M}} \mathsf{fun} f' = \lambda x_1 \ldots x_n.\, \lambda k.\, [\![e]\!]_{f,k} : \tau_1 \to \ldots \to \tau_n \to ((\tau,\,\tau)\,hfun) \to \tau$.

To show this, we need a variant of the general replacement lemma (Hindley & Seldin, 1986, Lemma 11.18; Wright & Felleisen, 1994, Lemma 4.2) to reason about linear substitution in an evaluation context:

Lemma 1. (*Linear replacement*)
If $\Gamma|_{\mathsf{N}}\,;\,\varnothing \vdash_{\mathrm{M}} \mathsf{K}[e] : \tau$ for a constructor context K then there is a sub-deduction $\Gamma|_{\mathsf{N}}\,;\,\varnothing \vdash_{\mathrm{M}} e : \tau'$ at the hole and $\Gamma|_{\mathsf{N}}\,;\,x : \tau' \vdash_{\mathrm{M}} \mathsf{K}[x] : \tau$.

Interestingly, this lemma requires constructor contexts and we would not be able to derive the Lemma for general contexts as the linear type environment is not propagated through applications. The proofs can be found in Appendix B.7 in the supplement, which also contains the full type rules adapted to our calculus.

6.2 In-place update

The instantiation with Minamide's system is using fast in-place updates and proven sound, but it is still a bit unsatisfactory as *how* such in-place mutation is done (or why this is safe) is only described informally. In Minamide's system, a suggested implementation for a context is as a tuple $\langle \mathsf{K}, x@i \rangle$ where K is (a pointer to) a context and $x@i$ is the address of the hole as the *ith* field of object x (in K). The empty tuple $\langle\rangle$ is used for an empty context ($\Box$). Composition and application directly update the hole pointed to by $x@i$ by overwriting the hole with the child context or value.

In contrast, Bour *et al*. (2021) show a TRMC translation for OCaml that uses *destination passing* style which makes it more explicit how the in-place update of the hole works. In particular, the general construct $x.i := v$ overwrites the *ith* field of any object x with v.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

To gain more insight of why in-place update is possible and correct, we are going to use the explicit heap semantics of Perceus (Xie *et al.*, 2021; Lorenzen & Leijen, 2022). In such semantics, the heap is explicit and all objects are explicitly reference counted. Using the Perceus derivation rules, we can soundly translate our current calculus to the Perceus target calculus where the reference counting instructions (dup and drop) are derived automatically by the derivation rules (Xie *et al.*, 2021, fig. 5). The Perceus heap semantics reduces the derived expressions using reduction steps of the form $H \mid e_1 \longmapsto_r H' \mid e_2$, which reduces a heap H and an expression e to a new heap H' and expression e_2 (Xie *et al.*, 2021, fig. 7). The heap H maps objects x with a reference count $n \geqslant 1$ to values, denoted as $x \mapsto^n v$. In this system, we can express in-place updates directly, and it turns out we can even *calculate* the in-place updating reduction rules for comp and app from the context laws. Before we do that though, we first need to establish some terminology and look carefully at what "in-place update" actually means.

6.2.1 *The essence of in-place update*

Let's consider a generic copy function, $(x.i\,\mathsf{as}\,y)$, that changes the *ith* field of an object x to y, for any generic constructor C:

$$x.i\,\mathsf{as}\,y \;=\; \mathsf{match}\,x\,\{\,C^k\,x_1 \ldots x_i \ldots x_k \to C^k\,x_1 \ldots y \ldots x_k\,\}$$

When we apply the Perceus algorithm (Xie *et al.*, 2021), we need to insert a single drop:

$$x.i\,\mathsf{as}\,y \;=\; \mathsf{match}\,x\,\{\,C^k\,x_1 \ldots x_i \ldots x_k \to \mathsf{drop}\,x_i;\; C^k\,x_1 \ldots y \ldots x_k\,\}$$

In the special case that x is unique at runtime (i.e., the reference count of x is 1), we can now derive the following:

$$\begin{array}{lll}
 & H,\, x \mapsto^1 C^k\,x_1 \ldots x_i \ldots x_k \mid x.i\,\mathsf{as}\,y & \{\; x \notin H,\; \mathbf{1} \;\} \\
= & H,\, x \mapsto^1 C^k\,x_1 \ldots x_i \ldots x_k \mid \mathsf{match}\,x\,\{ & \\
 & C^k\,x_1 \ldots x_i \ldots x_k \to \mathsf{drop}\,x_i;\; C^k\,x_1 \ldots y \ldots x_k\,\} & \{\; \mathit{def}\,.\;\} \\
\longrightarrow_{\mathsf{r}} & H,\, x \mapsto^1 C\,\overline{x_j} \mid & \\
 & \mathsf{dup}(\overline{x_j});\; \mathsf{drop}(x);\; \mathsf{drop}(x_i);\; C^k\,x_1 \ldots y \ldots x_k & \{\; (\mathit{match}_r)\;\} \\
\longrightarrow^*_{\mathsf{r}} & H',\, x \mapsto^1 C\,\overline{x_j} \mid \mathsf{drop}(x);\; \mathsf{drop}(x_i);\; C^k\,x_1 \ldots y \ldots x_k & \{\; (\mathit{dup}_r), H'\;\mathit{has}\;\overline{x_j}\;\mathit{dup}'d, \mathbf{2}\;\} \\
\longrightarrow_{\mathsf{r}} & H' \mid \mathsf{drop}(\overline{x_j});\; \mathsf{drop}(x_i);\; C^k\,x_1 \ldots y \ldots x_k & \{\; (\mathit{drop}_r)\;\} \\
\longrightarrow_{\mathsf{r}} & H \mid \mathsf{drop}(x_i);\; C^k\,x_1 \ldots y \ldots x_k & \{\; \mathit{cancel}\;H'\;\mathit{dup}\;\overline{x_j}\;(2)\;\} \\
\cong & H \mid \mathsf{let}\,z = C^k\,x_1 \ldots y \ldots x_k\;\mathsf{in}\;\mathsf{drop}(x_i);\; z & \{\; \mathit{drop}\;\mathit{commutes}\;\} \\
\longrightarrow_{\mathsf{r}} & H,\, z \mapsto^1 C\,x_1 \ldots y \ldots x_k \mid \mathsf{drop}(x_i);\; z & \{\; (\mathit{con}_r),\; \mathsf{fresh}\,z,\; \mathbf{3}\;\} \\
= & H,\, x \mapsto^1 C\,x_1 \ldots y \ldots x_k \mid \mathsf{drop}(x_i);\; x & \{\; \alpha\;\mathit{rename}\;(1),(3)\;\}
\end{array}$$

And this is the essence of in-place mutation: when an object is unique, an in-place update corresponds to allocating a fresh copy, discarding the original (due to the uniqueness of x), and α-renaming to reuse the original "address".

We will write $(x.i := z)$ for $(x.i\,\mathsf{as}\,z)$ in the special case of updating a field in a unique constructor, where we can derive the following reduction rule:

$$(\mathit{assign}) \quad H,\, x \mapsto^1 C\ldots x_i \ldots \mid x.i := y \quad \longrightarrow^*_{\mathsf{r}} \quad H,\, x \mapsto^1 C\ldots y \ldots \mid \mathsf{drop}\,x_i;\; x$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

and in the case the field is a $\Box$, we can further refine this to:

(*assignn*) $H, x \mapsto^1 C \ldots \Box_i \ldots \mid x.i := y \quad \longrightarrow_{\mathsf{r}}^* \quad H, x \mapsto^1 C \ldots y \ldots \mid x$

For convenience, we will from now on use the notation $C \ldots x_i \ldots$, and $C \ldots \Box_i \ldots$ to denote the *ith* field in a constructor if there is no ambiguity.

6.2.2 Linear chains

We need a bit more generality to express hole updates in contexts. In particular, we will see that all objects along the path from the top of the context to the hole are unique by construction. We call such unique path a *linear chain*, denoted as $[H]_x^n$:

$[H]_x^n = [x \mapsto^n v_0, x_1 \mapsto^1 v_1, \ldots, x_m \mapsto^1 v_m]_x^n \quad (m \geqslant 0)$

where for all $x_i \in (\mathsf{dom}(H) - \{x\})$, we have $x_i \in \mathsf{fv}(v_{i-1})$ (and therefore for all $y \in \mathsf{dom}(H)$ we have $\mathsf{reachable}(H, x)$). Since the objects in H besides x are all unique and not reachable otherwise, we also say that x dominates H. When the dominator is also unique, we call it a *unique linear chain* (of the form $[H]_x^1$). We can define linear chains inductively as well since a single object always forms a linear chain:

(*linearone*) $x \mapsto^n v \; = [x \mapsto^n v]_x^n$

and we can always extend with a unique linear chain:

(*linearcons*) $x \mapsto^n \ldots z \ldots, [H]_z^1 \; = [x \mapsto^n \ldots z \ldots, H]_x^n$

Using (*linearcons*) we can derive that we can append a unique linear chain as well:

(*linearapp*) $[H_1, y \mapsto^1 \ldots z \ldots]_x^n, [H_2]_z^1 \; = [H_1, y \mapsto^1 \ldots z \ldots, H_2]_x^n$

6.2.3 Contexts as a linear chain

To simplify the proofs, we assume in this subsection that all fields in K contexts are variables:

$\mathsf{K} ::= \Box \mid C\, x_1 \ldots \mathsf{K} \ldots x_n$

since we can always arrange any K to have this form by let-binding the values v. It turns out that a constructor context then always evaluates to a unique linear chain:

Lemma 2. (*Contexts evaluate to unique linear chains*)
For any K, we have $H \mid \mathsf{K}[C \ldots \Box_i \ldots] \longrightarrow_{\mathsf{r}}^* H, [H', y \mapsto^1 C \ldots \Box_i \ldots]_x^1 \mid x$.

We can show this by induction on the shape of K (Appendix B.8 in the supplement).

6.2.4 Calculating the fold

Following Minamide's approach, we are going to denote our contexts as a tuple $\langle x, y@i \rangle$ where x is (a pointer to) a constructor context and $y@i$ is the address of the hole as the *ith* field of object y. We define $\mathsf{ctx}\,\mathsf{K} = (\!|\mathsf{K}|\!)$. For an empty context, we use an empty tuple ($(\!|\Box|\!) = \langle\rangle$), but otherwise we can specify the fold as:

(*foldspec*) $H \mid (\!|\mathsf{K}[C \ldots \Box_i \ldots]|\!) \;\cong\; H \mid \mathsf{let}\, x = \mathsf{K}[C \ldots \Box_i \ldots]\, \mathsf{in}\, \langle x, [x]@i \rangle$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

where we use the notation $[x]$ do denote the last object of the linear chain formed by K (Lemma 2). We can now calculate the definition of $(\!|_|\!)$ from its specification (see Appendix B.9 in the supplement), where we get following definition for $(\!|_|\!)$:

$$\begin{array}{ll}(\!|\Box|\!) & = \langle\rangle \\ (\!|C\ldots\Box_i\ldots|\!) & = \mathsf{let}\, x = C\ldots\Box_i\ldots \mathsf{in}\, \langle x, x@i\rangle \\ (\!|C\ldots\mathsf{K}\ldots|\!) & = \mathsf{let}\, \langle z, x@i\rangle = (\!|\mathsf{K}|\!)\, \mathsf{in}\, \langle C\ldots z\ldots, x@i\rangle \quad (\mathsf{K}\neq\Box)\end{array}$$

This builds up the context using let bindings, while propagating the address of the hole. As before, the intention is that the compiler expands the fold statically. For example, the *map* function translates to:

$$\begin{array}{l} map'\, xs\, f\, k = \mathsf{match}\, xs\, \{ \\ \quad Nil \to \mathsf{app}\, k\, Nil \\ \quad Cons\, x\, xx \to \mathsf{let}\, y = f\, x\, \mathsf{in}\, map'\, xx\, f\, (k \bullet (\mathsf{let}\, z = Cons\, y\, \Box\, \mathsf{in}\, \langle z, z@2\rangle))\, \}\end{array}$$

where $z@2$ correctly denotes the address of the hole field in the context.

6.2.5 *Updating a context*

Before we can define in-place application, we need an in-place substitution operation $\mathsf{subst}\, \langle x, y@i\rangle\, z$ that substitutes z at the hole (at $y@i$) in the context x. Note that in our representation of a context as a tuple $\langle x, y@i\rangle$ we treat $y@i$ purely as an address and do not reference count y as such. The y part is a "weak" pointer, and we cannot use it directly without also having an "real" reference. This means that if we want to define an in-place substitution, we cannot define it directly as $y.i := z$ (since we have no real reference to y). Instead, we are going to calculate an in-place updating substitution from its specification:

$$(subspec)\, H, [H', y \mapsto^1 C\ldots\Box_i\ldots]^1_x \mid \mathsf{subst}\, \langle x, y@i\rangle\, z \;\cong\; H, [H', y \mapsto^1 C\ldots z\ldots]^1_x \mid x$$

We do this by induction of the shape of the linear chain. For the singleton case, we have:

$$\begin{array}{lll} & H, [y \mapsto^1 C\ldots\Box_i\ldots]^1_y \mid \mathsf{subst}\, \langle y, y@i\rangle\, z & \\ = & H, [y \mapsto^1 C\ldots\Box_i\ldots]^1_y \mid y.i := z & \{\ define,\ (we\ have\ a\ y\ reference!)\ \} \\ \longrightarrow & H, [y \mapsto^1 C\ldots z\ldots]^1_y \mid y & \{\ (assignn)\ \}\end{array}$$

and for the extension we have:

$$\begin{array}{lll} & H, [x \mapsto^1 C\ldots x'_j\ldots, [H', y \mapsto^1 C\ldots\Box_i\ldots]^1_{x'}]^1_x \mid \mathsf{subst}\, \langle x, y@i\rangle\, z & \\ = & H, [x \mapsto^1 C\ldots x'_j\ldots, [H', y \mapsto^1 C\ldots\Box_i\ldots]^1_{x'}]^1_x & \\ & \mid \mathsf{dup}\, x';\ x.j := \Box;\ x.j := \mathsf{subst}\, \langle x', y@i\rangle\, z & \{\ define\ \} \\ \longrightarrow^* & H, [x \mapsto^1 C\ldots\Box_j\ldots, [H', y \mapsto^1 C\ldots\Box_i\ldots]^1_{x'}]^1_x & \\ & \mid x.j := \mathsf{subst}\, \langle x', y\rangle\, z & \{\ (dup_r),\ (assign)\ \} \\ \cong & H, [x \mapsto^1 C\ldots\Box_j\ldots, [H', y \mapsto^1 C\ldots z\ldots]^1_{x'}]^1_x \mid x.j := x' & \{\ induction\ hyp.\ \} \\ \longrightarrow & H, [x \mapsto^1 C\ldots x'_j\ldots, [H', y \mapsto^1 C\ldots z\ldots]^1_{x'}]^1_x \mid x & \{\ (assignn)\ \}\end{array}$$

This leads to the following inductive definition of subst:

$$\begin{array}{l} H \mid \mathsf{subst}\, \langle x, x@i\rangle\, z \;\; = H \mid x.i := z \\ H \mid \mathsf{subst}\, \langle x, y@i\rangle\, z \;\; = H \mid \mathsf{dup}\, x';\ x.j := \Box;\ x.j := \mathsf{subst}\, \langle x', y@i\rangle\, z \\ \quad \mathsf{where}\, x \neq y \wedge [x \mapsto^1 C\ldots x'_j\ldots, [H']^1_{x'}]^1_x \in H\end{array}$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

That is, to update the last element of the chain in-place, we need traverse down while separating the links such that when we reach the final element it has a unique reference count and can be updated in-place. We then traverse back up fixing up all the links again. Of course, we would not actually use this implementation in practice – the derivation here just shows that the substitution specification is sound, and we can thus implement the (*subspec*) reduction by instead using the tuple address $y@i$ directly to update the hole in-place. In essence, due to the uniqueness of the elements in the chain, the y is uniquely reachable through x, and thus it is safe to use it directly in this case.

6.2.6 Calculating application and composition

With the specification for fold and in-place substitution, we can use the context laws to calculate the in-place updating version of application and composition. Starting with application, we can calculate (for $\mathsf{K} \neq \Box$):

$$
\begin{array}{lll}
 & H \mid \mathsf{app}\,(\mathsf{ctx}\,\mathsf{K})\,e & \\
= & H \mid \mathsf{app}\,(\!|\mathsf{K}|\!)\,e & \{\ \mathit{def}.\ \} \\
\cong & H \mid \mathsf{app}\,(\mathsf{let}\,x = \mathsf{K}[\Box]\,\mathsf{in}\,\langle x, [x]@i\rangle)\,e & \{\ \mathit{fold\ specification},\ \mathsf{K} \neq \Box\ \} \\
\cong & H,\ [H', y \mapsto^1 C\ldots\Box_i\ldots]^1_x \mid \mathsf{app}\,\langle x, [x]@i\rangle\ e & \{\ \mathit{lemma}\,2,\ \mathbf{1}\ \} \\
= & H,\ [H', y \mapsto^1 C\ldots\Box_i\ldots]^1_x \mid \mathsf{app}\,\langle x, y@i\rangle\ e & \{\ \mathit{def}.\ \} \\
\cong & H,\ z \mapsto^1 v,\ [H', y \mapsto^1 C\ldots\Box_i\ldots]^1_x \mid \mathsf{app}\,\langle x, y@i\rangle\ z & \{\ e\ \mathit{is\ terminating}\ \mathbf{2}\ \} \\
= & H,\ z \mapsto^1 v,\ [H', y \mapsto^1 C\ldots\Box_i\ldots]^1_x \mid \mathsf{subst}\,\langle x, y@i\rangle\ z & \{\ \mathit{define}\ \} \\
\cong & H,\ z \mapsto^1 v,\ [H', y \mapsto^1 C\ldots z\ldots]^1_x \mid x & \{\ (\mathit{subspec})\ \} \\
\cong & H,\ z \mapsto^1 v \mid \mathsf{K}[z] & \{\ \mathit{lemma}\,2,\ (1)\ \} \\
\cong & H \mid \mathsf{K}[e] & \{\ (2)\ \}
\end{array}
$$

And thus we define application directly in terms of in-place substitution as:

$$(\mathit{uapp})\ H \mid \mathsf{app}\,\langle x, y@i\rangle\ z \quad \longrightarrow_{\mathsf{r}} \quad H \mid \mathsf{subst}\,\langle x, y@i\rangle\ z$$

We arrived exactly at the "obvious" implementation where the hole inside a unique context is updated in-place in constant time. This also corresponds to the informal implementation given in Section 2.3. For composition, it turns out we can define it in terms of applications:

$$(\mathit{ucomp})\ H \mid \langle x_1, y_1@i\rangle \bullet \langle x_2, y_2@j\rangle \quad \longrightarrow_{\mathsf{r}} \quad H \mid \langle \mathsf{app}\,\langle x_1, y_1@i\rangle\ x_2,\ y_2@j\rangle$$

where the derivation is in Appendix B.10 in the supplement. Again we arrived at the efficient translation where the hole in the first unique context is updated in-place (and in constant time) with a pointer to the second context. The full rules for application and composition are (with the derivations for the empty contexts in Appendix B.10 in the supplement):

$$
\begin{array}{lllll}
(\mathit{uapph}) & H \mid \mathsf{app}\,\langle\rangle\ x & \longrightarrow_{\mathsf{r}} & H \mid x \\
(\mathit{uapp}) & H \mid \mathsf{app}\,\langle x, y@i\rangle\ z & \longrightarrow_{\mathsf{r}} & H \mid \mathsf{subst}\,\langle x, y@i\rangle\ z \\
(\mathit{ucomp}) & H \mid \langle x_1, y_1@i\rangle \bullet \langle x_2, y_2@j\rangle & \longrightarrow_{\mathsf{r}} & H \mid \langle \mathsf{app}\,\langle x_1, y_1@i\rangle\ x_2,\ y_2@j\rangle \\
(\mathit{ucompl}) & H \mid \langle\rangle \bullet \langle x_2, y_2@j\rangle & \longrightarrow_{\mathsf{r}} & H \mid \langle x_2, y_2@j\rangle \\
(\mathit{ucompr}) & H \mid \langle x_1, y_1@i\rangle \bullet \langle\rangle & \longrightarrow_{\mathsf{r}} & H \mid \langle x_1, y_1@i\rangle
\end{array}
$$

Note that (*ucompr*) is not really needed since by construction our translation never generates empty contexts for the second argument. The rules also correspond to the informal implementation given in Section 2.3 where `Id` was used to represent the empty tuple.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

With these definitions, we still need to show that we can be *efficient* and that we never get *stuck*. For efficiency, we need to show that a context $\langle x, y@i \rangle$ is always a linear chain so we don't have to check that at runtime in (*subspec*). This follows by construction since any initial context ctx K is a linear chain (Lemma 2), and any composition as well (*ucomp*). Second, the reference count of the dominator should always be 1 or otherwise (*subspec*) may not apply – that is, contexts should be used linearly. This follows indirectly from Lemma 4 where we show that our translation adheres to Minamide's linear-type discipline. A more direct approach would be to show that Perceus never derives a dup operation for a context k in our translation. However, we refrain from doing so here, as it turns out that with general algebraic effect handlers, the linearity of a context may no longer be guaranteed!

7 Modulo first-class constructor contexts

In Koka, constructor contexts are first-class values in the language (Lorenzen *et al.*, 2024). A constructor context can be used more than once, as for example in the expression `val c = ctx Cons(1,_) in (c ++. [2] , c ++. [3])`, where the context `c` is shared and it evaluates correctly to `([1,2],[1,3])`. This abstraction can safely encapsulate the limited form of mutation necessary to implement a Minamide tuple, while still having a purely functional interface that does not rely on linear types.

As we will show in Section 7.2, a `ctx` K expression is compiled such that each constructor context has at runtime a representation of its linear chain (the *context path*). The Koka compiler compiles a context like `ctx Node(Node(Node(Leaf,1,Leaf),2,_),5,Leaf)` internally into a Minamide tuple:

```
val x = Node₃(Node(Leaf,1,Leaf),2,hole) in Ctx(Node₁(x,5,Leaf), x@3)
```

where each constructor along the context path is annotated with a child index (1 and 3) leading from the root down to the hole.

When we compose or apply a context, we have to determine whether the context is shared. If the contexts happen to be used linearly, then all operations execute in constant time, just as in Minamide's approach. But if the context is shared, we will have to copy it along the context path. This gives us a full functional semantics and any subsequent substitutions on the same context work correctly (but will take linear time in the length of the context path).

The ability to copy contexts if necessary is useful for programmers and we discuss some examples in Section 8. But perhaps surprisingly, it is also important for the TRMC transformation itself as we show next.

7.1 Nonlinear control

A long-standing issue in a TRMc transformation is that it is unsound in the presence of non-local control operations like *call*/*cc*, *shift*/*reset* (Danvy & Filinski, 1990; Sitaram & Felleisen, 1990; Shan, 2007), or in general with algebraic effect handlers (Plotkin & Power, 2003; Plotkin & Pretnar, 2009), whenever a continuation or handler resumption can be invoked more than once. Note that if only single-shot continuations or resumptions

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

are allowed (as in OCaml Dolan *et al.*, 2015) for example), the control flow is still always linear and the TRMc transformation still sound. Since the Koka language relies foundationally on general effect handlers (Leijen, 2017, 2021; Xie & Leijen, 2021), we need to tackle this problem. Algebraic effect handlers extend the syntax with a handle expression, $\mathsf{handle}\, h\, e$, and operations, op, that are handled by a handler h. There are two more reduction rules (Leijen, 2014):

$$\begin{array}{lll} (\mathit{return}) & \mathsf{handle}\, h\, v & \longrightarrow \quad v \\ (\mathit{handle}) & \mathsf{handle}\, h\, \mathsf{E}[\mathsf{op}\, v] & \longrightarrow \quad e[x{:=}v,\ \mathit{resume}{:=}\lambda y.\, \mathsf{handle}\, h\, \mathsf{E}[y]] \\ & & \quad \mathsf{where}\, (\mathit{op} \mapsto \lambda x.\, \lambda \mathit{resume}.\, e) \in h \ \wedge \mathit{op} \notin \mathsf{E} \end{array}$$

That is, when an operation is invoked it yields all the way up to the innermost handler for that operation and continues from there with the operation clause. Besides the operation argument, it also receives a resumption *resume* that allows the operation to return to the original call site with a result y. The culprit here is that the resumption captures the delimited evaluation context E in a lambda expression, and this can violate linearity assumptions. In particular, if we regard a TRMC context k as a linear value (as in Minamide), then such k may be in the context E of the (*handle*) rule and captured in a nonlinear lambda. Whenever the operation clause calls the resumption more than once, any captured linear values may be used more than once!

A nice example in practice of this occurs in the well-known Knapsack problem as described by Wu *et al.* (2014) where they use multiple resumptions to implement a non-determinism handler:

```
effect nondet
  ctl flip() : bool   // a control operation that may resume more than once
  ctl fail() : a      // or not at all

fun select( xs : list<a> ) : nondet a    // pick an element from a list
  match xs
    Nil       -> fail()
    Cons(x,xx) -> if flip() then x else select(xx)

fun knapsack(w : int, vs : list<int> ) : <nondet,div> list<int>
  if w < 0 then fail()
  elif w == 0 then []
  else val v = select(vs) in Cons(v, knapsack(w - v, vs))
```

An effectful function in Koka has three arguments where the type `a -> e b` denotes a function from type `a` to `b` with a potential (side) effect `e`. The `select` function picks an element from a list using the operations of the `nondet` effect. The `knapsack` function picks items from a list of item weights `vs` that together do not exceed the capacity `w` (of the knapsack). Since it calls `select` it has the `nondet` effect and additionally it has a divergence effect. We can now provide an effect handler that systematically explores all solutions using multiple resumptions:

```
val solutions = handler
  return(x)  [x]
  ctl fail() []
  ctl flip() resume(True) ++ resume(False)

fun test() : div list<list<int>>
  with solutions
  knapsack(3,[3,2,1])
```

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

That is, the `solutions` handler implements the `flip` function by resuming twice and appending the results. Even though `knapsack` returns a single solution as a list, the `test` function returns a list of all possible solution lists (as `[[3],[2,1],[1,2],[1,1,1]]`). The `knapsack` function is in the modulo *cons* fragment and gets translated to a tail-recursive version by our translation into:

```
fun knapsack'(w :int, vs :list<int>, k :ctx<list<int>> ) : <nondet,div> list<int>
  if w < 0 then k ++. fail() elif w == 0 then k ++. []
  else val v = select(vs)
       knapsack'(w - v, vs, k ++ Cons(v, _))
```

Instead of having a runtime that captures evaluation contexts E directly, Koka usually uses an explicit monadic transformation to translate effectful computations into pure lambda calculus. The effect handling is then implemented explicitly using a generic multi-prompt control monad `eff` (Xie & Leijen, 2020, 2021). This transforms our `knapsack` function into:

```
fun knapsack'(w:int,vs:list<int>,k:ctx<list<int>>) : eff<<nondet,div>,list<int>>
  if w < 0 then ... elif w == 0 then Pure( k ++. [] )
  else match select(vs)
         Pure(v)    -> knapsack'(w - v, vs, k ++ Cons(v,_))
         Yield(yld) -> Yield( yield-extend(yld,
                              fn(v) knapsack'(w - v, vs, k ++ Cons(v,_))))
```

Every computation in the effect monad either returns with a result (`Pure`) or is yielding up to a handler (`Yield`). As described by Xie & Leijen (2021), the Koka compiler backend implements this monad as a primitive and can generate efficient C code without needing to allocate closures in the fast non-yielding path.

In our example, we inlined the monadic bind operation where the result `select(vs)` is explicitly matched. We see that in the `Yield` case, the continuation expression (namely `fn(v)knapsack'(w - v, vs, k ++ Cons(v,_))`) is now explicitly captured under a lambda expression – including the supposedly linear context `k`! This is how we can end up at runtime with a context that is shared (with a reference count > 1) and where the rule (*ucomp*) should not be applied.

7.2 Dynamic copying via reference counting

Our context composition is defined in terms of context application, which in turn relies on the in-place substitution (Section 6.2.5):

$$(subspec)\; H, [H', y \mapsto^1 C \ldots \Box_i \ldots]^1_x \mid \mathsf{subst}\, \langle x, y@i \rangle\, z \;\cong\; H, [H', y \mapsto^1 C \ldots z \ldots]^1_x \mid x$$

This is the operation that eventually fails if the runtime context x is not unique. In Section 6.2.5, the substitution operation was calculated to recursively visit the full linear chain of the context. This suggests a solution for any non-unique context: we can actually traverse the context at runtime and create a fresh copy instead.

It is not immediately clear though how to implement such operation at runtime: the linear chains up to now are just a proof technique and we cannot actually visit the elements of the chain at runtime as we do not know which field in a chain element points to the next element. What we need to do is to explicitly annotate each constructor C^k (of arity k) in a context also with an index i corresponding to the field that points to the next element, as C^k_i. It turns out, we can actually do this efficiently while constructing the context – and we

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

can do it systematically just by modifying our fold function to keep track of this *context path* at construction:

$$
\begin{aligned}
(\!|\Box|\!) &= \langle\rangle \\
(\!|C \ldots \Box_i \ldots|\!) &= \mathsf{let}\, x = C_i \ldots \Box_i \ldots \mathsf{in}\, \langle x,\, x@i \rangle \\
(\!|C \ldots \mathsf{K}_i \ldots|\!) &= \mathsf{let}\, \langle z, x@j \rangle = (\!|\mathsf{K}|\!)\, \mathsf{in}\, \langle C_i \ldots z \ldots,\, x@j \rangle \quad (\mathsf{K} \neq \Box)
\end{aligned}
$$

With such indices present at runtime, we can define non-unique substitution as:

$$(\mathit{subapp}) \quad H, [H']_x^{n+1} \mid \mathsf{subst}\, \langle x, y@i \rangle\, z \quad \cong \quad H, [H']_x^{n+1} \mid \mathsf{append}\, x\, z$$

where append follows the context path at runtime copying each element as we go and eventually appending z at the hole:

$$
\begin{aligned}
H, x \mapsto^n C_i \ldots \Box_i \ldots \mid \mathsf{append}\, x\, z &\longrightarrow_r H, x \mapsto^n C_i \ldots \Box_i \ldots \mid x.i\, \mathsf{as}\, z \\
H, x \mapsto^n C_i \ldots y_i \ldots \mid \mathsf{append}\, x\, z &\longrightarrow_r H, x \mapsto^n C_i \ldots y_i \ldots \mid \mathsf{dup}\, y_i;\; x.i\, \mathsf{as}\, (\mathsf{append}\, y_i\, z)
\end{aligned}
$$

We can show the context laws still hold for these definitions (see Appendix B.11 in the supplement). The append operation in particular can be implemented efficiently at runtime using a fast loop that updates the previous element at each iteration (essentially using manual TRMC!). In the Koka runtime system, it happens to be the case that there is already an 8-bit field index in the header of each object which is used for stackless freeing. We can thus use that field for context paths since if a context is freed it is fine to discard the context path anyways. The runtime cost of the hybrid technique is mostly due to an extra uniqueness check needed when doing context composition to see if we can safely substitute in-place (see Section 7.3). As we see in the benchmark section, this turns out to be quite fast in practice. Moreover, the Koka compiler uses static-type information when possible to avoid this check if a function is guaranteed to be used only with a linear effect type.

7.3 Efficient code generation

As an example of the code generation of our TRMC scheme, we consider the `map` function from our benchmarks in Section 9. The `map` function is specialized by the compiler for the increment function, and after the TRMC transformation we have:

```
fun map_trmc'( xs : list<int32>, k : ctx<list<int32>> ) : list<int32>
  match xs
    Nil -> k ++. Nil
    Cons(x,xx) ->
      val y = x+1
      val c = ctx Cons(y, _)
      map_trmc'(xx, k ++ c)

fun map_trmc( xs : list<int32> ) : list<int32>
  map_trmc'(xs, ctx _)
```

Here the empty context `ctx _` is the Minamide tuple as a value type containing the final result and hole address, something like `Ctx(invalid, null)`. For efficiency, we represent the empty tuple with a `null` address for the hole. The single-cell context `ctx Cons(y, _)` is represented by the minamide tuple `val c = Cons₂(y, □) in Ctx(c, c@2)`. Eventually, the `Ctx` value type is passed in registers (`x19` and `x21`), and the generated code for arm64 becomes:

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

```
map_trmc':
    ...                          ; setup
    mov x21, x2                  ; x21 is the hole address of the tuple
    mov x19, x1                  ; x19 the final result part of the tuple
    cmp x0, #5                   ; is it Nil?
    b.ne  LBB3_5                 ; if not, goto to Cons branch
    ...
  LBB3_5:                        ; Cons branch
    mov x20, x3                  ; set up loop variables in registers
    mov x23, #x100000000         ; used for fast int32 arithmetic
    mov w24, #x020202            ; Cons header: fields=2,ctx path index=2,tag=2,rc=0
    mov w25, #1
  LBB3_6:                        ; tail call entry
    ldp x26, x22, [x0, #8]       ; load pair: x = x26 and xx = x22
    ldr w8, [x0, #4]             ; load ref count in w8
    cbnz  w8, LBB3_10            ; if not unique, goto slower copying path
  LBB3_7:
    add x8, x23, x26, lsl #31; increment x from/to a boxed int32 representation
    asr x8, x8, #31
    orr x8, x8, #0x1
    stp x24, x8, [x0]            ; store pair in-place: header and the incremented x
    mov x8, x0
    str x25, [x8, #16]!          ; set tail to invalid (1) for now (not really needed)
    cbz x21, LBB3_16             ; if this an empty tuple (hole==NULL), goto slow path
    str x0, [x21]                ; else store our Cons result into the current hole
  LBB3_9:
    mov x0, x22                  ; continue with the tail (x22)
    mov x21, x8                  ; and set x21 to the new hole
    cmp x22, #5                  ; is it a Nil?
    b.ne  LBB3_6                 ; if not, make a tail call
    b LBB3_2                     ; otherwise return
    ...

  map_trmc:
    mov x3, x1                   ; set up the empty Minamide tuple
    mov w1, #1                   ; final result is invalid for now (1)
    mov x2, #0                   ; with the initial hole==NULL
    b map_trmc'                  ; and jump
```

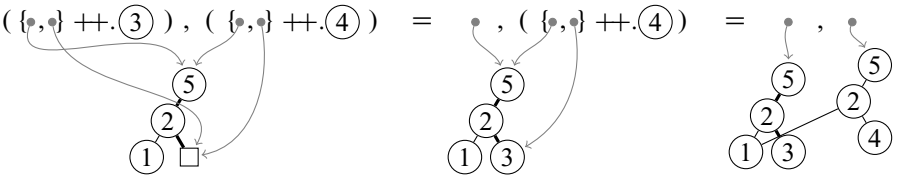
Note in particular how the header for the `Cons` node in the context is set as `mov w24, #x020202` where, from left-to-right, we initialize the tag (`0x02`), the context path field (`0x02`), and the total number of fields (also `0x02`). As such, maintaining context paths comes for free since it is done as part of header initialization. Also we see the reuse of Perceus reference counting (Xie *et al*., 2021; Lorenzen & Leijen, 2022) in action, where the `Cons` node that is matched (in `x0`) is reused for the context `Cons` node (also in `x0`). Since the effect inferred for the specialized map function is `total`, the check for uniqueness of the context is removed as it the context is guaranteed to be used lineraly.

7.4 Dynamic copying without reference counting

Lorenzen *et al*. (2024) show that it is possible to support first-class constructor contexts even in languages without precise reference counts. Their proposed implementation (also suggested by Gabriel Scherer) uses a special distinguished value for a runtime hole □ that is never used by any other object. A substitution now first checks the value at the hole:

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

if it is a □ value, the hole is substituted for the first time and we just overwrite the hole in-place (in constant time). However, any subsequent substitution on the same context will find some object instead of □. At this point, we first dynamically copy the context path (in linear time) and then update the copy in-place.



The illustration above (due to Lorenzen *et al*., 2024) shows a more complex example of a shared tree context that is applied to two separate nodes. The runtime context path is denoted here by bold edges. The intermediate state is interesting as it is both a valid tree, but also a part of the tree is shared with the remaining context, where the hole points to a regular node now. When that context is applied, only the context path (node 5 and 2) is copied first where all other nodes stay shared (in this case, only node 1).

However, it turns out that this simple approach is not sound without further restrictions. For general first-class contexts, the second context can be arbitrary (instead of always a constant `ctx` in the TRMC case), the context composition operation `c1 ++ c2` needs an extra check in order to avoid creating cycles: we check if `c2` has an already overwritten hole or if the hole in `c2` is at the same address as in `c1`. In either case, `c2` is copied along the context path.

Figure 3 shows a partial implemention in C code of how one can implement constructor contexts in a runtime for languages without precise reference counting. We assume that `HOLE` is the distinguished value for unfilled holes (□). When we compose two contexts, we need to ensure we can handle shared contexts as well where we copy a context along the context path if needed (using `ctx_copy`).

In the application and composition functions, the check `(A)` sees if the hole in `c1` is already overwritten (where `*c1.hole != HOLE`). In that case, we copy `c1` along the context path as shown in Section 7.2 to maintain referential transparency.

However, in the composition operation we also need to do a similar check for `c2` as well in order to avoid cycles: the second check `(B)` checks if `c2` has an already overwritten hole, but also if the hole in `c2` is the same as in `c1`. In either case, `c2` is copied along the context path. Effectively, both checks ensure that the new context that is returned always ends with a single fresh `HOLE`. Let's consider some examples of shared contexts. A basic example is a simple shared context, as in:

```
val c = ctx Cons(1,_) in (c ++. [2], c ++. [3])
```

which evaluates to `([1,2],[1,3])`. Here, during the second application, check `(A)` ensures the shared context `c` is copied such that the list `[1,2]` stays unaffected.

A more tricky example is composing a context with itself:

```
val c = ctx Cons(1,_) in (c ++ c) ++. [2]
```

which evaluates to `[1,1,2]`. The check `(B)` here copies the appended `c` (since `c1.hole == c2.hole`). In this example, the potential for a cycle is immediate, but generally it can be obscured with a shared context inside another. Consider:

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

```
struct ctx_t {              // a Minamide context
  heap_block_t*   root;
  heap_block_t**  hole;
};

struct ctx_t ctx_copy( struct ctx_t c ) {
  struct ctx_t d = { .root = c.root, .hole = c.hole };
  if( c.root == NULL ) return d;
  heap_block_t** prev = &(c.root);
  heap_block_t** next = &(d.root);

  while( prev != c.hole ) {
    *next = heap_block_copy( *prev );
    prev = (*prev)->children + ((*prev)->ctx_path);
    next = (*next)->children + ((*next)->ctx_path);
  }
  d.hole = next;
  return d;
}

// (++.) : cctx<a,b> -> b -> a
heap_block_t*  ctx_apply( struct ctx_t c1, heap_block_t* x )
{
  // is c1 an empty context?
  if (c1.root == NULL) return x;

  // copy c1 ?
  struct ctx_t  d1 = (*c1.hole != HOLE ?  ctx_copy(c1) : c1);       // (A)

  *d1.hole = x;
  return d1.root;
}

// (++) : cctx<a,b> -> cctx<b,c> -> cctx<a,c>
struct ctx_t  ctx_compose( struct ctx_t c1, struct ctx_t c2 )
{
  // is c1 or c2 an empty context?
  if (c1.root == NULL) return c2;
  if (c2.root == NULL) return c1;

  // copy c1 ?
  struct ctx_t  d1 = (*c1.hole != HOLE  ?  ctx_copy(c1) : c1 );     // (A)

  // copy c2 ? (needed to avoid cycles)
  struct ctx_t  d2 = ((*c2.hole != HOLE || c1.hole == c2.hole)
                    ? ctx_copy(c2) : c2 );                        // (B)

  *d1.hole = d2.root;
  d1.hole  = d2.hole;
  return d1;
}
```

Fig. 3. Implementing constructor composition and application in the runtime system (for languages without precise reference counts).

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

```
val c1 = ctx Cons(1,_)
val c2 = ctx Cons(2,_)
val c3 = ctx Cons(3,_)
val c  = c1 ++ c2 ++ c3 in  (c ++ c2) ++. [4]
```

which evaluates to `[1,2,3,2,4]`. The check `(B)` again copies the appended `c2` in `c ++ c2` (since `*c2.hole != HOLE`).

Note that the (B) check in composition is sufficient to avoid cycles. In order to create a cycle in the context path, either `c1` must be in the context path of `c2` (I), or the `c2` in the context path of `c1` (II). For case (I), if `c1` is at the end of `c2`, then their holes are at the same address where `c1.hole == c2.hole`. Otherwise, if `c1` is not at the end, then `*c1.hole != HOLE` and we have copied `c1` already due to check (A). For case (II) the argument is similar: if `c2` is at the end of `c1` we again have `c1.hole == c2.hole`, and otherwise `*c2.hole != HOLE`.

The implementation using precise reference counting is not very different from the one without reference counting. The main difference is in the checks (A) and (B), which become:

```
// copy c1 ?
struct ctx_t  d1 = (!is_unique(c1.root) ?  ctx_copy(c1) : c1 );   // (A)

// copy c2 ? (needed to maintain ctx paths where each node beside the root is
unique)
struct ctx_t  d2 = (!is_unique(c2.root) ?  ctx_copy(c2) : c2 );   // (B)
```

This is the implementation that is used in the Koka runtime system. The (B) check here is required to maintain the invariant that context paths always form *unique chains* (Section 7.2). From this property, it follows directly that no cycles can occur in the context path.

7.5 Runtime behavior

Interestingly, the two implementations, with or without precise reference counting, do differ in their runtime performance characteristics, which are dual to each other in terms of space and time.

7.5.1 Time

The implementation without reference counting only copies on demand when the hole is already filled, whereas our earlier implementation with reference counts copies whenever the context is found to be not unique upon filling the hole. The latter can be a problem if the context is later discarded without being used. Consider the knapsack program, which in its last iteration may call itself on a one-element list `[x]` with `x = w`. For this special case, the code reduces to:

```
fun knapsack'(x : int, k : ctx<list<int>>) : <nondet,div> list<int>
  val v = if flip() then x else fail() in k ++. Cons(v, [])
```

This computation is run twice, where the first run successfully returns `k ++. (Cons(x, []))` but the second run fails. The reference counting-based implementation has to copy `k` in the first run, since its reference count is not one (due to `k` being captured for the second run). In contrast, assuming that the hole in `k` is not yet filled, the

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

new implementation can simply fill the hole of `k` with `Cons(x, [])` in the first run without copying. Since `k` is discarded in the second run, no copying is needed at all. We will come back to this point in Section 9, where we see that the reference counting implementation in Koka does not perform well in a backtracking search, presumably due to this issue.

7.5.2 Space

The implementation without precise reference counts can use more space though than the one based on reference counting. This can occur when a context accidentally holds on to values that have been written into its hole. Consider an earlier state of the knapsack program, where it may process a list `vs = Cons(v, vv)` with `v > w`. Then we can simplify the code to:

```
if flip()
  then knapsack(w - v, Cons(v, vv), k ++ Cons(v, _))
  else val v' = select(vv) in knapsack(w - v', Cons(v, vv), k ++ Cons(v, _))
```

Following the `flip()`, we first try to use `v` as our element. But since `v > w`, this computation fails and we backtrack. However, our new algorithm may have written `Cons(v, _)` into the hole of `k`. This value is now garbage, but this may not be obvious to a garbage collector or reference counting scheme, since `k` is still live. Only when backtracking to the second run do we copy `k` and discard the old value.

In contrast, the implementation based on reference counting would have copied (and discarded) `k` in the first run already. Unlike the new implementation, it is *garbage-free* (Xie *et al*., 2021) and guarantees that no space is used for values that are no longer needed. For this reason, we prefer the implementation via reference counting in Koka, using the other implementation for GC-based languages.

8 Programming with first-class constructor contexts

First-class constructor contexts turn out to be a powerful feature, and they allow us to write many programs by hand that would be hard to generate automatically from a general TRMC transformation. In this section, we explore some of these programs, all of which can be written in Koka.

8.1 Modulo cons products

The `partition` function calls a predicate on each element of a list and appends it to one of two piles depending on the result:

```
fun partition(p : a -> bool, xs : list<a>) : (list<a, list<a>)
  match xs
    Nil -> (Nil, Nil)
    Cons(x, xx) ->
      val (yes, no) = partition(p, xx)
      if p(x)
        then (Cons(x, yes), no)
        else (yes, Cons(x, no))
```

The recursive call to `partition` is followed by a pattern match on the resulting tuple, an if-statement and finally the constructor application. This does not fit the TRMc transformation

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

directly, but it also might not seem too different – and indeed this function was suggested as fruitful target for an expanded TRMC translation both by Bour *et al*. (2021) and the conference version of this paper.

However, in order to make this function tail-recursive, the `p(x)` call would have to be moved *before* the recursive call. That can be done by a compiler if `p` is pure, but what if `p` may perform side effects? Thus, even an extended TRMc transformation could only apply if the user first rewrote their code to:

```
fun partition(p : a -> bool, xs : list<a>) : (list<a, list<a>)
  match xs
    Nil -> (Nil, Nil)
    Cons(x, xx) ->
      val ok = p(x)
      val (yes, no) = partition(p, xx)
      if ok
        then (Cons(x, yes), no)
        else (yes, Cons(x, no))
```

The conference version of this paper describes a transformation that recognizes that the pattern match on the returned tuple is mirrored in the creation of a new tuple and looks for constructor contexts inside the created tuple.

However, it may not be worth implementing such specific transformation as we can easily rewrite it manually using two explicit first-class constructor contexts for `yes` and `no`:

```
fun partition(p, xs, yes, no)
  match xs
    Nil -> (yes ++. Nil, no ++. Nil)
    Cons(x, xx) ->
      if p(x)
        then partition(p, xx, yes ++ ctx Cons(x, _), no)
        else partition(p, xx, yes, no ++ ctx Cons(x, _))
```

The resulting code is clearer than the version with an explicit `ok` variable, and not just more efficient but arguably even clearer than the original version. For this reason, we now recommend that programmers use first-class constructor contexts directly for examples like this.

8.2 Difference lists

Another example of future work described by Bour *et al*. (2021) is the `flatten` function. This function calls itself recursively and passes the result to the `append` function on lists:

```
fun append(xs : list<a>, ys : list<a>) : list<a>
  match xs
    Nil -> ys
    Cons(x, xx) -> Cons(x, append(xx, ys))

fun flatten(xss : list<list<a>>) : list<a>
  match xss
    Nil -> Nil
    Cons(xs, xss) -> append(xs, flatten(xss))
```

While `append` is tail-recursive modulo cons, `flatten` is not. However, `append` is just a sequence of constructor applications ending in the second argument, and we can easily

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

rewrite it using a first-class constructor context returned from `append` (i.e., a *difference* list):

```
fun append(acc : ctx<list<a>>, xs : list<a>) : ctx<list<a>>
  match xs
    Nil -> acc
    Cons(x, xs) -> append(acc ++ ctx Cons(x, _), xs)

fun flatten-acc(acc : ctx<list<a>>, xss : list<list<a>>) : list<a>
  match xss
    Nil -> acc ++. Nil
    Cons(xs, xss) -> flatten-acc(append(acc, xs), xss)

fun flatten( xss : list<list<a>> ) : list<a>
  flatten-acc(ctx _, xss)
```

8.3 Composing constructor contexts

Another example which illustrates the usefulness of first-class conntexts that can be stored in data structures is the composition of constructor contexts with defunctionalized evaluation contexts. While constructor contexts naturally apply to the map over a list, they do not apply directly to a map over trees:

```
type tree<a>
  Leaf
  Bin(l : tree<a>, a : a, r : tree<a>)

fun tmap(t, f)
  match t
    Bin(l, x, r) -> Bin(tmap(l, f), f(x), tmap(r, f))
    Leaf -> Leaf
```

Here, the first recursive call to `tmap` is not in a constructor context and thus the TRMc transformation alone is not enough to make this tail-recursive. However, instead of resorting to full defunctionalized evaluation contexts, we can use them only for descending into the left child and keep using constructor contexts to descend into the right branch:

```
type accum<a,b>
  Hole
  Accum(acc : accum<a,b>, top : ctx<tree<b>>, x : a, r : tree<a>)

fun tmap-acc(t, f, acc, top)
  match t
    Bin(l, x, r)    -> tmap-acc(l, f, Accum(acc, top, x, r), ctx _)
    Leaf            -> tmap-app(f, top ++. Leaf, acc)

fun tmap-app(f, l, acc)
  match acc
    Hole -> l
    Accum(acc, top, x, r) -> tmap-acc(r, f, acc, top ++ ctx Bin(l, f(x), _))
```

This function immediately follows from the technique described in Section 5.3. It extends the `acc` accumulator whenever it goes into the left subtree and extends the `top` accumulator whenever it goes into the right subtree. While a version using only defunctionalized evaluation contexts corresponds to pointer reversal (Schorr & Waite, 1967), this version reverses only the pointers going to the right child, but leaves the pointers to the left child intact.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

8.4 Polymorphic recursion

In this paper, we have limited ourselves to recursive functions where each recursive call has the same return type. However, there are some functions where the recursive call might have a different return type due to polymorphic recursion. For example, Okasaki (1999) presents the following random access list:

```
type seq<a>
  Empty
  Zero(       s : seq<(a, a)> )
  One ( x : a, s : seq<(a, a)> )

fun cons(x : a, s : seq<a>) : seq<a>
  match s
    Empty      -> One(x, Empty)
    Zero(ps)   -> One(x, ps)
    One(y, ps) -> Zero(cons((x, y), ps))
```

Here the recursive call instantiates `a` with `(a,a)`, and the hole in `Zero(□)` has type `seq <(a,a)>`. It turns out that for polymorphically recursive code, performing the translation can lead to code that is not typeable is System F. This issue is well known for defunctionalized evaluation contexts, where GADTs are required to regain typability (Pottier & Gauthier, 2004). Analogously, we give two type parameters to first-class constructor contexts `cctx<a,b>` where `a` corresponds to the type of the root and `b` to the type of the hole. Our primitive operations have the general types:

```
alias ctx<a> = cctx<a,a>

fun (++)( c1 : cctx<a,b>, c2 : cctx<b,c> ) : cctx<a,c>
fun (++.)( c : cctx<a,b>, x : b) : a
```

It turns out that this encapsulates the necessary type information to type the result of the translation for polymorphic recursion. Even though Koka has an intermediate core representation based on System F, the application and composition functions are primitives and Koka transforms the above function without problems. Our `cons` function is translated to:

```
fun cons(x : a, s : seq<a>, acc : cctx<seq<b>, seq<a>>) : seq<b>
  match s
    Empty      -> acc ++. One(x, Empty)
    Zero(ps)   -> acc ++. One(x, ps)
    One(y, ps) -> cons((x, y), ps, acc ++ ctx Zero(_))
```

8.5 Multiple holes

A limitation of our current approach is that we do not support multiple holes in a constructor context. For example, consider the following function which builds a perfectly balanced binary tree of height `n`:

```
fun tree(n : int) : tree<int>
  if n <= 0 then Leaf else
    val t = tree(n-1)
    Bin(t, n, t)
```

This function cannot directly be made tail-recursive with our current approach. The issue is that the value `t` occurs twice in the constructor. This means that we would need our context

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

to have two holes instead of one. But the compact representation of our context paths makes it impossible to fork the path, which means that we can only support copying for single-hole contexts. When constructor contexts are only used linearly and never copied, it is possible to support multiple holes directly (Bagrel, 2024).

8.6 TRMC as a source-to-source transformation

Since the introduction of first-class constructor contexts to the language, the TRMc transformation has become source-to-source. This means that there is no longer a direct reason for compiler writers to support TRMc directly as long as first-class contexts are supported, since the latter allow programmers to recreate the effects of the TRMc transformation manually.

Naturally, this raises the question of which parts of the transformation a compiler should handle automatically. In Koka, we currently support the TRMc transformation to ensure that the naive version of `map` is automatically optimized. However, we do not implement all extensions that are possible. For example, we have not implemented the extension proposed in Section 7.1 of the conference version of this paper since first-class constructor contexts make it easily possible to derive it manually. Furthermore, we do not implement an automatic translation for other kinds of contexts like the ones proposed in Section 4.

One choice for future compiler writers would be to automatically apply exactly those contexts for which benchmarks indicate that they always improve performance. This seems to be the case for constructor contexts and semiring contexts (which are also supported by GCC, see Section 4.4). However, traditional CPS and defunctionalized contexts tend to decrease performance. It can be valuable to use contexts that decrease performance since they still protect against stack overflows, but this should only be done for those functions that actually present a risk of causing a stack overflow.

Another option would be to let users make the choice on which functions should be transformed and with which context. Possible contexts might either be baked into the language or even declared directly by users. In Section 5, we describe how several contexts can be composed so that the compiler can achieve tail-recursive functions where some of the calls are eliminated using more performant contexts than other calls.

Finally, another sensible option would be to perform no TRMC pass in the compiler at all and instead leave the transformation completely to programmers. We believe that the implementations using an explicit accumulator presented in this section are often similar in clarity to direct-style implementations and it would not be infeasible to simply perform the TRMC transformation by hand where required. In return, the language enjoys simpler semantics and there is a smaller implementation burden in the compiler.

9 Benchmarks

The Koka compiler has a full implementation the TRMC algorithm as described in this paper for constructor contexts (since v2.0.3, Aug 2020). We measure the impact of TRMC relative to other variants on various tests: the standard map function over a list (*map*), mapping over a balanced binary tree (*tmap*), balanced insertion in a red-black tree (*rbtree*),

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

and finally the *knapsack* problem as shown in Section 7. Each test program scales the repetitions to process the same number of total elements (100 000 000) for each test size.

The *map* test repeatedly maps the increment function over a shared list of numbers from 1 to N and sums the result list. This means that the `map` function repeatedly copies the original list and Perceus cannot apply reuse here (Lorenzen & Leijen, 2022). For example, the test for the standard (and TRMC) `map` function in Koka is written as:

```
fun map-std( xs : list<a>, f : a -> b ) : list<b>
  match xs
    Cons(x,xx) -> Cons(f(x),xx.map-std(f))
    Nil        -> Nil

fun test(n : int)
  val xs = list(1,n)
  val x  = fold-int(0, 100_000_000/max(n,1), 0) fn(i,acc)
             acc + xs.map-std(fn(x) x + 1).sum
  println("total: " ++ x.show)
```

For each test, we measured five different variants:

- *trmc*: the TRMC version which is exactly like the standard (*std*) version.
- *std*: the standard non-tail-recursive version. This is the same source as the *trmc* version but compiled with the `-fno-trmc` flag.
- *acc*: this is the accumulator style definition where the accumulated result list- or tree-visitor is reversed in the end.
- *acc (no reuse)*: this is the accumulator style version but with Perceus reuse disabled for the accumulator. The performance of this variant may be more indicative for systems without reuse. Accumulator reuse is important as it allows the accumulated result to be reversed "in place".
- *cps*: the CPS style version with an explicit continuation function. This allocates a closure for every element that eventually allocates the result element for the final result. Perceus does not reuse the memory underlying closures.

The benchmark results are shown in Figure 4. For the *map* function, we see that our TRMC translation is always faster than the alternatives for any size list. For a tree map (*tmap*), this is also the case, except for one-element trees where the standard *tmap* is slightly faster (6%). However, when we consider a slightly more realistic example of balanced insertion into a tree, TRMC is again as fast or faster in all cases. The *rbtree* benchmark is interesting as during traversal down to the insertion point, there two recursive cases where TRMC applies, but also two recursive cases where TRMC does not apply. Here, we see that it still helps to apply TRMC where possible as looping is apparently faster than a recursive call in this benchmark.

Finally, *knapsack* implements the example from Section 7 with a backtracking effect. Unfortunately, the TRMC variant, which uses the *hybrid* approach to copy the context on demand, is less fast than the alternatives. It is not *that* much slower though – about 25% at worst. The reason for this is that there is less sharing. For the accumulator version, at each choice point the current accumulated result is shared between each choice, building a tree of choices. At the end, many of these choices are just discarded (as the knapsack is too full), and only for valid solutions a result list is constructed (as a copy). However, for the *hybrid trmc* approach, we copy the context on demand *at each choice point*, and when we reach a point where the knapsack is too full the entire result is discarded, keeping only

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

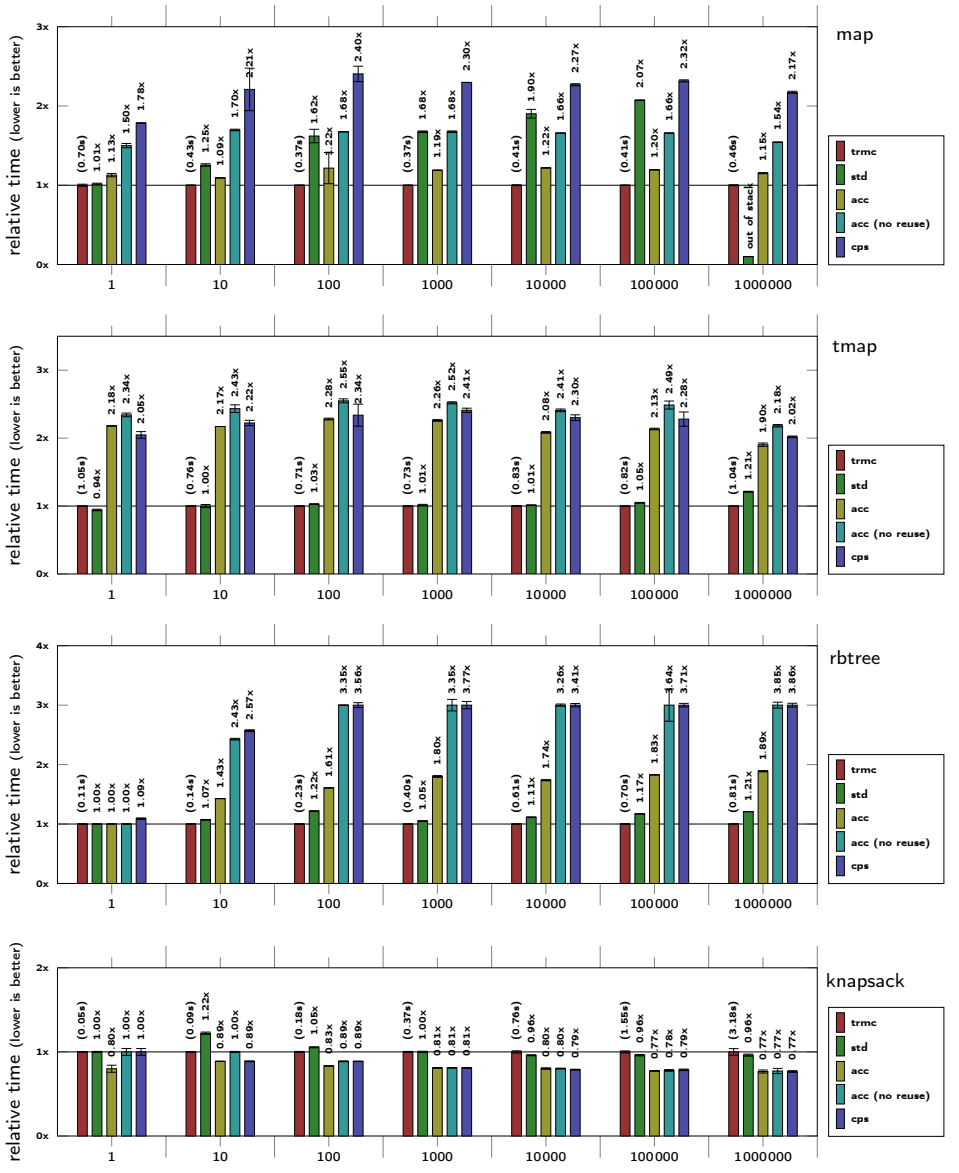


Fig. 4. Benchmarks on Ubuntu 20.04 (AMD 5950x), Koka v2.4.1-dev. The benchmarks are map over a list (*map*), map over a tree (*tmap*), balanced red-black tree insertion (*rbtree*), and the *knapsack* program that uses nonlinear control flow. Each workload is scaled to process the same number of total elements (usually 100 000 000). The tested variants are TRMC (*trmc*), the standard non-tail-recursive style (*std*), accumulator style (*acc*), accumulator style without Perceus reuse (*acc (no reuse)*), and finally CPS style (*cps*).

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

valid solutions. As such, the *trmc* variant copies more than the other approaches depending on how many of the generated solutions are eventually kept. Still, in Koka we prefer the hybrid approach to avoid code duplication.

10 Related work

Tail recursion modulo *cons* was a known technique in the LISP community as early as the 1970s. Risch (1973) describes the TRMc transformation in the context of REMREC system which also implemented the modulo associative operators instantiation described in Section 4.4. A more precise description of the TRMc transformation was given by Friedman & Wise (1975).

More recently, Bour *et al*. (2021) describe an implementation for OCaml which also explores various language design issues with TRMc. The implementation is based on *destination passing* style where the result is always directly written into the destination hole. This entails generating an initial unrolling of each function. For example, the `map` function is translated (in pseudo code) as:

```
fun map( xs, f )
  match xs
    Nil -> Nil
    Cons(x,xx) ->
      val y = f(x)
      val dst = Cons(y,□)
      map_dps( xx, f, dst@2 )
      dst
```

```
fun map_dps( xs, f, dst@i ) : ()
  match xs
    Nil -> dst.i := Nil
    Cons(x,xx) ->
      val y = f(x)
      val dst' = Cons(y,□)
      dst.i := dst'
      map_dps( xx, f, dst'@2 )
```

This can potentially be more efficient since there is only one extra argument for the destination address (instead of our representation as a Minamide tuple of the final result with the hole address), but it comes at the price of duplicating code. Note that the `map_dps` function returns just a unit value and is only called for its side effect. As such it seems quite different from our general TRMC based on context composition and application. However, the destination passing style may still be reconciled with our approach: with a Minamide tuple the first iteration always uses an "empty" tuple, while every subsequent iteration has a tuple with the fixed final result as its first element, where only the hole address (i.e., the destination) changes. Destination passing style uses this observation to specialize for each case, doing one unrolling for the first iteration (with the empty tuple), and then iterating with only the second hole address as the destination.

The algorithm rules by Bour *et al*. (2021) directly generate a destination passing style program. For example, the core translation rule for a constructor with a hole is:

$$\frac{n' = |I| + 1 \quad d'.n' \leftarrow \Box[U] \rightsquigarrow_{dps} \mathbb{T}[d_l.n_l \leftarrow \mathsf{K}_l]^l}{d.n \leftarrow \mathsf{K}[C((e_i)^{i\in I}, \Box, (e_j)^j)][U] \rightsquigarrow_{dps} \begin{array}{l} \mathsf{let}\, d' = C((e_i)^{i\in I}, \mathit{Hole}, (e_j)^j)\, \mathsf{in} \\ d.n \leftarrow \mathsf{K}[d']; \\ \mathbb{T}[d_l.n_l \leftarrow \mathsf{K}_l]^l \end{array}} \text{[DPS-REIFY]}$$

Here a single rule does various transformations that we treat as orthogonal, such as folding, extraction, instantiation of composition, and the actual TRMc transformation. Allain *et al*. (2025) expand on the design and give a formal proof of correctness using separation logic.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

In logic languages, difference lists (Clark & Tärnlund, 1977) can be used to encode a form of TRMc: difference lists are usually presented as a pair (L, X) where X is a logic variable which is the last element of the list L. With in-place update of the unification variable X, one can thus append to L in constant time – quite similar to our constructor contexts. This is also done in the experimental Ozma backend of the Scala language (Doeraene & Van Roy, 2013). Engels (2022) describes an implemention of TRMC for the Elm language that can also tail-optimize calls to the right of a list append by keeping the last cell of the right-appended list as a context. Pottier & Protzenko (2013) implement a type system inspired by separation logic, which allows the user to implement a safe version of in place updating TRMc through a mutable intermediate datatype. Laziness works similar to TRMc for the functions we consider: recursive calls guarded by a constructor are thunked and incremental forcing can happen without using the stack. The listless machine (Wadler, 1984) is an elegant model for this behavior.

Hughes (1986) considers the function `reverse` and shows how the fast version can be derived from the naive version by defining a new representation of lists as a composition of partially applied append functions (which are sometimes also called difference lists). His function `rep(xs)` (defined as `fn(ys) xs ++ ys`) creates such abstract list and is equal to our `ctx` when instantiated to append functions and list contexts (Section 4.1). Similarly, his `abs(f)` function (defined as `f []`) corresponds to our `app` k `[]` in that case, and finally, the correctness condition would correspond to our (*appctx*) law. The idea of calculating programs from a specification has a long history, and we refer the reader to early work by Bird (1984), Wand (1980), and Meertens (1986), and more recent work by Gibbons (2022) and Hutton (2021).

Defunctionalization (Reynolds, 1972; Danvy & Nielsen, 2001) has often been used to eliminate all higher-order calls and obtain a first-order version of a program. Wand & Friedman (1978) describes a defunctionalization algorithm in the context of LISP. Minamide *et al*. (1996) introduce special primitives `pack` and `open` (that correspond roughly to our `ctx` and `app`) and describe a type system for correct usage. Bell *et al*. (1997) and Tolmach & Oliva (1998) perform the conversion automatically at compile-time. Danvy & Nielsen (2001) propose to apply defunctionalization only to the closures of self-recursive calls, which should produce equal results as our approach in Section 4.3. However, they do not give an algorithm for this and the technique has so far mainly been used manually (Danvy & Goldberg, 2002; Gibbons, 2022).

An early implementation of TRMc in a typed language was in the OPAL compiler (Didrich *et al*., 1994). Similar to Bour *et al*. (2021), they also used destination passing style compilation with an extra destination argument where the final result is written to. Like Koka and Lean, OPAL also managed memory using reference counting and could reuse matched constructors (Schulte & Grieskamp, 1992). Reuse combines well with TRMc and in recent work Lorenzen & Leijen (2022) show how this can be used to speed up balanced insertion into red-black trees using the *functional but in-place* (FBIP) technique. Sobel & Friedman (1998) propose to reuse the closures of a CPS-transformed program for newly allocated constructors and show that this approach succeeds for all anamorphisms. However, reuse based on dynamic reference counts can improve upon this by for example also reusing the original data for the accumulator (and generalize to nonlinear control).

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

We are using the linearity of the Perceus heap semantics (Xie *et al*., 2021; Lorenzen & Leijen, 2022) to reason about linear chains and the essence of in-place updates. In our case, these linear chains are used to reason about the shape of a separate part of the heap. This suggest that separation logic (Reynolds, 2002) could also be used effectively for such proofs. For example, Moine *et al*. (2023) use separation logic to reason about space usage under garbage collection.

11 Conclusion and future work

In this paper, we explored tail recursion modulo *context* and tried to bring the general principles out of the shadows of specific algorithms and into the light of equational reasoning. We have a full implementation of the modulo *cons* instantiation and look forward to explore future extensions to other instantiations as described in this paper.

Acknowledgments

We thank the anonymous reviewers of POPL 2023 and JFP for their helpful feedback. Gabriel Scherer and Jeremy Gibbons provided feedback on earlier drafts of this paper.

Conflicts of Interest

None.

References

Abelson, H., Kent Dybvig, R., Haynes, C. T., Rozas, G. J., Adams, N. I., Friedman, D. P., Kohlbecker, E., et al. (1998) Revised 5 report on the algorithmic language scheme. In *Higher-Order and Symbolic Computation*, vol. 11. Springer, pp. 7–105.

Allain, C., Bour, F., Clément, B., Pottier, F. & Scherer, G. (2025) Tail modulo cons, OCaml, and relational separation logic. In *Proceedings of the ACM on Programming Languages*, number POPL. New York, NY, USA: ACM.

Appel, A. W. (1991) *Compiling with Continuations*. Cambridge University Press. doi:10.1017/CBO9780511609619.

Appel, A. W. & McAllester, D. (2001) An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst. (TOPLAS)* **23**(5), 657–683. ACM New York, NY, USA.

Appel, A. W., Mellies, P.-A., Richards, C. D. & Vouillon, J. (2007) A very modal model of a modern, major, general type system. In *Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pp. 109–122.

Bagrel, T. (2024) Destination-passing style programming: A haskell implementation. *Journeées Francophones Des Langages Applicatifs (JFLA)*. Saint-Jacut-de-la-Mer, France. https://inria.hal.science/hal-04406360/document. hal-04406360.

Bell, J. M., Bellegarde, F. & Hook, J. (1997) Type-Driven Defunctionalization, ICFP'97. New York, NY, USA: Association for Computing Machinery, pp. 25–37. doi:10.1145/258948.258953.

Bird, R. S. (1984) The promotion and accumulation strategies in transformational programming. *ACM Trans. Program. Lang. Syst.* **6**(4), 487–504. doi:10.1145/1780.1781.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Bloch, J. (2008) *Effective Java (2nd Edition) (The Java Series)*, 2nd ed. USA: Prentice Hall PTR. doi:10.5555/1377533.

Bour, F., Clément, B. & Scherer, G. (2021) Tail modulo cons. *Journeées Francophones Des Langages Applicatifs (JFLA)*, April. Saint Médard d'Excideuil. https://hal.inria.fr/hal-03146495/document. hal-03146495.

Clark, K. L. & Tärnlund, S.-k. (1977) A first order theory of data and programs. In *IFIP Congress*, pp. 939–944.

Danvy, O. & Filinski, A. (1990) Abstracting control. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming*. LFP'90. Nice, France, pp. 151–160. doi:10.1145/91556.91622.

Danvy, O. & Goldberg, M. (2002) There and back again. In *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming*. ICFP'02. Pittsburgh, PA, USA, pp. 230–234. doi:10.1145/581478.581500.

Danvy, O. & Nielsen, L. R. (2001) Defunctionalization at work. In *Proceedings of the 3rd ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*. PPDP'01. Florence, Italy, pp. 162–174. doi:10.1145/773184.773202.

Didrich, K., Fett, A., Gerke, C., Grieskamp, W. & Pepper, P. (1994) OPAL: design and implementation of an algebraic programming language. In *Programming Languages and System Architectures*. Springer, pp. 228–244. doi:10.1007/3-540-57840-4_34.

Doeraene, S. & Van Roy, P. (2013) A new concurrency model for scala based on a declarative dataflow core. In *Proceedings of the 4th Workshop on Scala*. SCALA'13. New York, NY, USA: Association for Computing Machinery. doi:10.1145/2489837.2489841.

Dolan, S., White, L., Sivaramakrishnan, K. C., Yallop, J. & Madhavapeddy, A. (2015) Effective concurrency through algebraic effects. In *OCaml Workshop*.

Doligez, D. & Leroy, X. (1993) A concurrent, generational garbage collector for a multithreaded implementation of ML. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL'93. New York, NY, USA: Association for Computing Machinery, pp. 113–123. doi:10.1145/158511.158611.

Dreyer, D., Ahmed, A. & Birkedal, L. (2009) Logical step-indexed logical relations. In *2009 24th Annual IEEE Symposium on Logic in Computer Science*. IEEE, pp. 71–80.

Dvořák, Z. (2004) Declarative world inspiration. In *GCC Developers' Summit*, 25. See also https://github.com/gcc-mirror/gcc/blob/master/gcc/tree-tailcall.cc.

Engels, J. (2022) Tail Recursion, but modulo Cons. Accessed June 06, 2022. https://jfmengels.net/modulo-cons/.

Flanagan, C., Sabry, A., Duba, B. F. & Felleisen, M. (1993) The essence of compiling with continuations. In *Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*. PLDI'93. Albuquerque, New Mexico, USA, pp. 237–247. doi:10.1145/155090.155113.

Friedman, D. P. & Wise, D. S. (1975) *Unwinding Stylized Recursion into Iterations*. 19. Computer Science Department Indiana University, Bloomingdale, Indiana. https://legacy.cs.indiana.edu/ftp/techreports/TR19.pdf.

Gibbons, J. (2022) Continuation-passing style, defunctionalization, accumulations, and associativity. *Art Sci. Eng. Program.* **6**, Article 7.

Harper, R. (2012) 15-150 Equational Reasoning Guide. https://www.cs.cmu.edu/~15150/previous-semesters/2012-spring/resources/handouts/equational.pdf. Notes for the 15-150 CMU Functional Programming course.

Hindley, J. & Seldin, J. (1986) *Introduction to Combinators and Lambda-Calculus.* Cambridge University Press. doi:10.1017/CBO9780511809835.

Huet, G. P. (1997) The zipper. *J. Funct. Program.* **7**(5), 549–554. doi:10.1017/S0956796897002864.

Hughes, R. J. M. (1986) A novel representation of lists and its application to the function "reverse." *Inf. Process. Lett.* **22**(3), 141–144. Elsevier. doi:10.1016/0020-0190(86)90059-1.

Hutton, G. (2021) It's as Easy as 1, 2, 3. https://www.cs.nott.ac.uk/~pszgmh/123.pdf. Unpublished draft.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Krogh-Jespersen, M., Svendsen, K. & Birkedal, L. (2017) A relational model of types-and-effects in higher-order concurrent separation logic. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages*, pp. 218–231.

Leijen, D. (2014) Koka: programming with row polymorphic effect types. In *MSFP'14, 5th Workshop on Mathematically Structured Functional Programming*. doi:10.4204/EPTCS.153.8.

Leijen, D. (2017) Type directed compilation of row-typed algebraic effects. In *Proc. of the 44th ACM SIGPLAN Symp. on Principles of Programming Languages (POPL'17)*, Paris, France, pp. 486–499. doi:10.1145/3009837.3009872.

Leijen, D. (2021) The Koka Language. `https://koka-lang.github.io`.

Leijen, D. & Lorenzen, A. (2023) Tail recursion modulo context: An equational approach. *Proc. ACM Program. Lang.* **7**(POPL), 1152–1181. ACM New York, NY, USA. doi:10.1145/3571233.

Lorenzen, A. & Leijen, D. (2022) Reference counting with frame limited reuse. Proc. ACM Program. Lang. **6**. ICFP. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3547634.

Lorenzen, A., Leijen, D. & Swierstra, W. (2023) FP^2: Fully in-place functional programming. *Proc. ACM Program. Lang.* **7**(ICFP), 275–304. ACM New York, NY, USA. doi:10.1145/3607840.

Lorenzen, A., Leijen, D., Swierstra, W. & Lindley, S. (2024) The functional essence of imperative binary search trees. In *Proceedings of the ACM on Programming Languages*, number PLDI. ACM New York, NY, USA.

Maurer, L., Downen, P., Ariola, Z. M. & Jones, S. P. (2017) Compiling without continuations. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pp. 482–494.

Meertens, L. (1986) Algorithmics: Towards programming as a mathematical activity. In *Mathematics and Computer Science*, pp. 289–334.

Minamide, Y. (1998) A functional representation of data structures with a hole. In *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL'98. San Diego, California, USA, pp. 75–84. doi:10.1145/268946.268953.

Minamide, Y., Morrisett, G. & Harper, R. (1996) Typed closure conversion. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL'96. New York, NY, USA: Association for Computing Machinery, pp. 271–283. doi:10.1145/237721.237791.

Moine, A., Charguéraud, A. & Pottier,F. (2023) A high-level separation logic for heap space under garbage collection (extended version). In *Proceedings of the 50th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. POPL'23, pp. 1–27.

Okasaki, C. (1999) *Purely Functional Data Structures*. New York: Colombia University.

Plotkin, G. D. & Power, J. (2003) Algebraic operations and generic effects. *Appl. Categor. Struct.* **11**(1), 69–94. doi:10.1023/A:1023064908962.

Plotkin, G. D. & Pretnar, M. (2009) Handlers of algebraic effects. In *18th European Symposium on Programming Languages and Systems*. ESOP'09. York, UK, pp. 80–94. doi:10.1007/978-3-642-00590-9_7.

Pottier, F. & Gauthier, N. (2004) Polymorphic Typed Defunctionalization, POPL'04. New York, NY, USA: Association for Computing Machinery, pp. 89–98. doi:10.1145/964001.964009.

Pottier, F. & Protzenko, J. (2013) Programming with permissions in mezzo. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming*. ICFP'13. Boston, Massachusetts, USA: ACM, pp. 173–184. doi:10.1145/2500365.2500598.

Reinking, A., Xie, N., de Moura, L. & Leijen, D. (2021) Perceus: Garbage free reference counting with reuse. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. PLDI 2021. New York, NY, USA: ACM, pp. 96–111. doi:10.1145/3453483.3454032.

Reynolds, J. C. (1972) Definitional interpreters for higher-order programming languages. In *Proceedings of the ACM Annual Conference - Volume 2*. Boston, Massachusetts, USA: ACM, pp. 717–740. doi:10.1145/800194.805852.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

Reynolds, J. C. (2002) Separation logic: A logic for shared mutable data structures. In *Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science*. LICS'02. USA: IEEE Computer Society, pp. 55–74.

Risch, T. (1973) *REMREC - A Program for Automatic Recursion Removal*. Inst. för Informationsbehandling, Uppsala Universitet. https://user.it.uu.se/~torer/publ/remrec.pdf.

Schorr, H. & Waite, W. M. (1967) An efficient machine-independent procedure for garbage collection in various list structures. *Commun. ACM* **10**(8), 501–506. ACM New York, NY, USA.

Schulte, W. & Grieskamp, W. (1992) Generating efficient portable code for a strict applicative language. In *Declarative Programming, Sasbachwalden 1991*. Springer, pp. 239–252. doi:10.1007/978-1-4471-3794-8_16.

Shan, C.-c. (2007) A static simulation of dynamic delimited control. *Higher-Order Symb. Comput.* **20**(4), 371–401. doi:10.1007/s10990-007-9010-4.

Sitaram, D. & Felleisen, M. (1990) Control delimiters and their hierarchies. *LISP Symb. Comput.* **3**(1), 67–99. doi:10.1007/BF01806126.

Sobel, J. & Friedman, D. P. (1998) Recycling continuations. In *Proc. of the Third ACM SIGPLAN Int. Conf. on Functional Programming*. ICFP'98, Baltimore, Maryland, USA, pp. 251–260. doi:10.1145/289423.289452.

Timany, A., Krebbers, R., Dreyer, D. & Birkedal, L. (2024) A logical approach to type soundness. *J. ACM* **71**(6), 1–75. ACM New York, NY.

Tolmach, A. & Oliva, D. P. (1998) From ML to Ada: Strongly-typed language interoperability via source translation. *J. Funct. Program.* **8**(4), 367–412. Cambridge University Press: . doi:10.1017/S0956796898003086.

Turon, A., Dreyer, D. & Birkedal, L. (2013) Unifying refinement and hoare-style reasoning in a logic for higher-order concurrency. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming*, pp. 377–390.

Ullrich, S. & de Moura, L. (2019) Counting immutable beans – reference counting optimized for purely functional programming. In *Proceedings of the 31st Symposium on Implementation and Application of Functional Languages (IFL'19)*, Singapore.

Wadler, P. (1984) Listlessness is better than laziness: Lazy evaluation and garbage collection at compile-time. In *Proceedings of the 1984 ACM Symposium on LISP and Functional Programming*. LFP'84. New York, NY, USA: Association for Computing Machinery, pp. 45–52. doi:10.1145/800055.802020.

Wadler, P. (1987) The Concatenate Vanishes. *Note, University of Glasgow*.

Wand, M. (1980) Continuation-based program transformation strategies. *J. ACM* **27**(1), 164–180. doi:10.1145/322169.322183.

Wand, M. & Friedman, D. P. (1978) Compiling lambda-expressions using continuations and factorizations. *Comput. Lang.* **3**(4), 241–263. Pergamon Press, Inc., USA. doi:10.1016/0096-0551(78)90042-5.

Wright, A. K. & Felleisen, M. (1994) A syntactic approach to type soundness. *Inf. Comput.* **115**(1), 38–94. doi:10.1006/inco.1994.1093.

Wu, N., Schrijvers, T. & Hinze, R. (2014) Effect handlers in scope. In *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell*, Haskell'14. Göthenburg, Sweden, pp. 1–12. doi:10.1145/2633357.2633358.

Xie, N. & Leijen, D. (2020) Effect handlers in haskell, evidently. In *Proceedings of the 13th ACM SIGPLAN International Symposium on Haskell*. Haskell 2020. New York, NY, USA: Association for Computing Machinery, pp. 95–108. doi:10.1145/3406088.3409022.

Xie, N. & Leijen, D. (2021) Generalized evidence passing for effect handlers: Efficient compilation of effect handlers to C. In *Proc. ACM Program. Lang.*, vol. 5. ICFP. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3473576.

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

A Further Benchmarks

Figure A1 shows benchmark results of the `map` benchmark. This time we included the results for OCaml 4.14.0 which has support for TRMc (Bour *et al.*, 2021) using the `[@tail_mod_cons]` attribute. For example, the TRMc `map` function is expressed as:

```
let[@tail_mod_cons] rec map_trmc xs f =
  match xs with
  | [] -> []
  | x :: xx -> let y = f x in y :: map_trmc xx f
```

Comparing across systems is always difficult since there are many different aspects, in particular the different memory management of both systems where Koka uses Perceus style reference counting (Xie et al., 2021) and OCaml uses generational garbage collection, with a copying collector for the minor generation, and a mark-sweep collector for the major heap (Doligez & Leroy, 1993).

The results at least indicate that our approach, using Minamide-style tuples of the final result object and a hole address, is competitive with the OCaml approach based on direct destination passing style. For our translation, the *trmc* translation is always as fast or faster as the alternatives, but unfortunately this is not the case in OCaml (yet) where it requires larger lists to become faster then the standard recursion.

OCaml is also faster for lists of size 10 where *std* is about 25% faster than Koka's *trmc*. We believe this is in particular due to memory management. For the micro benchmark, such small lists always fit in the minor heap with very fast bump allocation. Since in the benchmark the result is always immediately discarded no live data need to be traced in the minor heap for GC – perfect! In contrast, Koka uses regular malloc/free with reference counting with the associated overheads. However, once the workload increases with larger lists, the overhead of garbage collection and copying to the major heap becomes larger, and in such situation Koka becomes (significantly) faster. Also, the time to process the 100

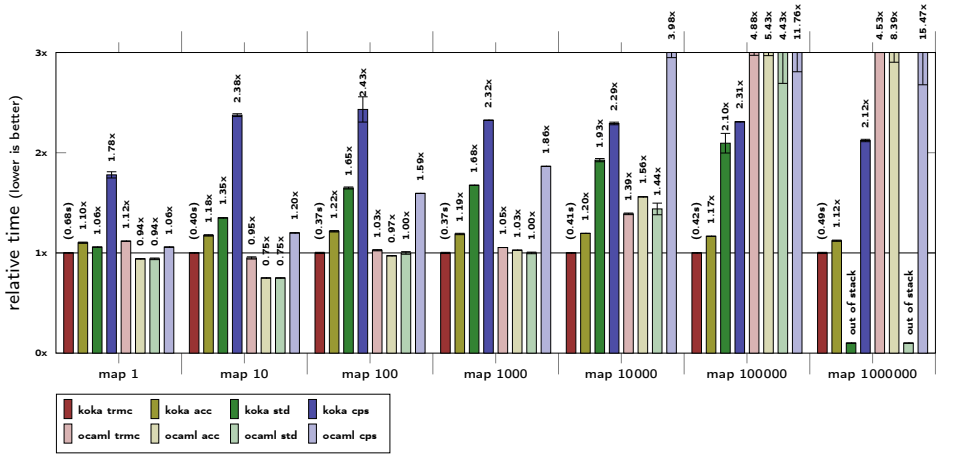


Fig. A1. Benchmarks on Ubuntu 20.04 (AMD 5950x), Koka v2.4.1-dev, OCaml 4.14.0. The benchmark repeatedly maps the increment function over a list of a given size and sums the result list. Each workload is scaled to process the same number of total elements (100 000 000). The tested variants of `map` are TRMC (trmc), accumulator style (acc), the standard non-tail-recursive style (std), and finally CPS style (cps).

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

M elements stays relatively stable for Koka (around 0.45 s) no matter the sizes of the lists, while with GC we see that processing on larger lists takes much longer.

B Proofs

B.1 The TRMC Algorithm is Sound

In Section 3.2, we used equational reasoning to derive the TRMC algorithm which makes it sound by construction. However, since we use an assumed specification (a), we generally need to make sure we only use this inductively on smaller terms (Hutton, 2021; Gibbons, 2022). However, we use the specification on the original term in the directly tail-recursive case, where we translate a tail-recursive expression $f\ e$ to its corresponding tail-recursive call in the translation $f'\ e\ k$. Intuitively, this is correct since we map each original recursive call to a recursive call in the translation, but technically it is not inductive.

It is a bit beyond the scope of this paper, we can show formally that such recursive reasoning step is actually sound in general. In particular, we can use a slightly more powerful notion of equality. We base our development on the step-indexed logical relation of Appel & McAllester (2001).

Definition 1.
For closed terms e_1 and e_2, we write $F \vDash e_1 \leqslant_i e_2$ if for all $0 < j < i$ we have $e_1 \longmapsto^j v$ implies $e_2 \longmapsto^* v$ under the global environment of top-level functions F. We write $F \vDash e_1 \cong_i e_2$ if $F \vDash e_1 \leqslant_i e_2$ and $F \vDash e_2 \leqslant_i e_1$.

The above definition is similar to the $e \leqslant_k f : \tau$ relation of Appel & McAllester (2001) (Section 3), but simplified in two ways: we do not consider a typing relation $: \tau$ and we require that the values produced by e_1 and e_2 are the same instead of merely being equivalent for $i - j$ steps. The latter choice implies that this notion of equality is not congruent under lambdas and we cannot rewrite under lambdas (which is the case for our development).

We use the subscript i instead of k to avoid confusion with our notation for continuations. We omit the environment F where it is clear from the context.

Definition 2.
For open terms e_1 and e_2 with free variables Γ, we write $F \vDash e_1 \leqslant_i e_2$ if for all substitutions σ mapping Γ to values, we have $F \vDash \sigma(e_1) \leqslant_i \sigma(e_2)$ and similarly for $F \vDash e_1 \cong_i e_2$. We write $F \vDash e_1 \leqslant e_2$ if $F \vDash e_1 \leqslant_i e_2$ for all i and similarly for $F \vDash e_1 \cong e_2$.

This definition is similar to the $\Gamma \vDash_k e \leqslant f : \tau$ relation but simplified to use the same substitution σ for both e_1 and e_2 instead of substitutions σ_1, σ_2 whose values are equivalent for i steps.

The main induction principle we need is the following:

$$\frac{f \notin \mathsf{E},\ e'\ \ \forall e'',\ (F, f\,x = e'' \vDash \mathsf{E}[f\,x] \leqslant e') \Rightarrow (F, f\,x = e'' \vDash \mathsf{E}[e] \leqslant e')}{F, f\,x = e \vDash \mathsf{E}[f\,x] \leqslant e'}\ [\textsc{unfolding}]$$

Informally, this lemma states that we can show that $\mathsf{E}[f\,x] \leqslant e'$ by unfolding f exactly one step and then rewrite using the inequality we aim to prove. The key insight that makes this

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

sound is that we quantify over all possible implementations of f. The "free theorem" of f then implies that our assumption can only access the inequality and not any other properties of f.

Lemma 3. (*Unfolding Lemma*)
Let f be a function, E be an evaluation context and e, e' be expressions where E, e' do not mention f. If for any implementation of f with $\mathsf{E}[f\,x] \leqslant e'$ we have $\mathsf{E}[e] \leqslant e'$, then for $f\,x = e$, we have $\mathsf{E}[f\,x] \leqslant e'$.

Proof. Construct a sequence of functions f_i as follows: $f_0\,x = f_0\,x$ and $f_{i+1}\,x = e[f_i/f]$.

- We show that $f_j\,x \leqslant f_{j+1}\,x$ for all j by induction on j. Case $j = 0$: Obvious, since $f_0\,x$ diverges. Case $j = j' + 1$: By induction on e. Case $e = f\,v$: Then $f_{j'}\,v \leqslant f_j\,v \leqslant f_{j+1}\,v$. All other cases follow from the inductive hypothesis since $f_{j'}\,x$ and $f_j\,x$ only differ in the recursive calls.
- We show that $\mathsf{E}[f_i\,x] \leqslant e'$ for all i. We have $f_0\,x \leqslant e'$ since f_0 diverges on all inputs. We have $\mathsf{E}[f_0\,x] \leqslant e'$, since the hole in an E context is always evaluated. Assume that $\mathsf{E}[f_i\,x] \leqslant e'$. Then for $f = f_i$, the assumption gives us $\mathsf{E}[e] \leqslant e'$ and thus $\mathsf{E}[f_{i+1}\,x] \leqslant e'$.

To finish the proof, we thus only have to show that $\mathsf{E}[f\,x] \leqslant \mathsf{E}[f_j\,x]$ for some j. Assume that $\mathsf{E}[f\,x]$ converges in i steps to v. Then $f\,x$ converges in $j \leqslant i$ steps to some value w. We show that $f_j\,x$ also converges to w by induction on j, which implies that $\mathsf{E}[f_j\,x]$ converges to v.

- Case $j = 0$: Since $f\,x$ has to take at least one step, it does not yield a value.
- Case $j = j' + 1$: By the induction hypothesis, we have $f\,x \leqslant_{j'} f_{j'}\,x$. We show that for $f\,x = e$, we have $f\,x \leqslant_j f_j\,x$ by induction on e. Case $e = f\,v$: Then $f\,v \leqslant_{j'} f_{j'}\,v \leqslant f_j\,v$. All other cases follow from the inductive hypothesis since $f\,x$ and $f_j\,x$ only differ in the recursive calls.

Our proof of the unfolding lemma is in the style of Scott Induction, where we create a chain of approximations f_i to the least fixpoint f. An alternative approach would be to define contextual equality of terms using weakest preconditions (Turon *et al.*, 2013; Krogh-Jespersen *et al.*, 2017; Timany *et al.*, 2024, Section 8) and to replace the approximation functions by Löb induction (Appel *et al.*, 2007; Dreyer *et al.*, 2009) in the weakest-precondition relation.

Proof of Theorem 1:

Let $f\,x = e_f$ and $f'\,x\,k = [\![e_f]\!]_{f,k}$, then $app\,k\,(f\,x) \cong f'\,x\,k$.

Proof. We focus on the (*tail*) case of the translation, while (*tlet*), (*tmatch*), (*tapp*), and (*base*) can be handled as in the text of the paper.

We show $f'\,x\,k = [\![e_f]\!]_{f,k} \leqslant app\,k\,(f\,x)$ by the unfolding lemma with $\mathsf{E} = \Box$, $e' = app\,k\,(f\,x)$. That is, we assume that $f'\,x\,k \leqslant app\,k\,(f\,x)$ and show that this implies that $[\![e_f]\!]_{f,k} \leqslant app\,k\,(f\,x)$:

$$
\begin{array}{lll}
 & [\![\mathsf{E}[f\ e_0]]\!]_{f,k} & \\
= & f'\ e_0\ (k \bullet \mathsf{ctx}\ \mathsf{E}) & \{\ (\mathit{tail})\ \} \\
\cong & \mathsf{let}\ x = e_0\ \mathsf{in}\ f'\ x\ (k \bullet \mathsf{ctx}\ \mathsf{E}) & \{\ (\mathit{letfloat})\ \} \\
\leqslant & \mathsf{let}\ x = e_0\ \mathsf{in}\ \mathsf{app}\ (k \bullet \mathsf{ctx}\ \mathsf{E})\ (f\ x) & \{\ \mathit{assumption} : f'\ x\ k \leqslant \mathsf{app}\ k\ (f\ x)\ \} \\
\cong & \mathsf{let}\ x = e_0\ \mathsf{in}\ \mathsf{app}\ k\ (\mathsf{app}\ (\mathsf{ctx}\ \mathsf{E})\ (f\ x)) & \{\ (\mathit{appcomp})\ \} \\
\cong & \mathsf{let}\ x = e_0\ \mathsf{in}\ \mathsf{app}\ k\ \mathsf{E}[f\ x] & \{\ (\mathit{appctx})\ \} \\
\cong & \mathsf{app}\ k\ \mathsf{E}[f\ e_0] & \{\ (\mathit{letfloat})\ \}
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

We show $app\,k\,(f\,x) = \mathsf{app}\,k\,(e_f) \leqslant f'\,x\,k$ by the unfolding lemma with $\mathsf{E} = app\,k\,\Box$, $e' = f'\,x\,k$. That is, we assume that $app\,k\,(f\,x) \leqslant f'\,x\,k$ and show that this implies that $\mathsf{app}\,k\,(e_f) \leqslant f'\,x\,k$:

$$
\begin{array}{lll}
 & \mathsf{app}\,k\,\mathsf{E}[f\,e_0] & \\
\cong & \mathsf{app}\,k\,(\mathsf{app}\,(\mathsf{ctx}\,\mathsf{E})\,(f\,e_0)) & \{\ (appctx)\ \} \\
\cong & \mathsf{app}\,(k \bullet \mathsf{ctx}\,\mathsf{E})\,(f\,e_0) & \{\ (appcomp)\ \} \\
\leqslant & f'\,e_0\,(k \bullet \mathsf{ctx}\,\mathsf{E}) & \{\ assumption : app\,k\,(f\,x) \leqslant f'\,x\,k\ \} \\
= & [\![\mathsf{E}[f\,e_0]]\!]_{f,k} & \{\ (tail)\ \}
\end{array}
$$

□

B.2 Context Laws for Defunctionalized Contexts

$$
\begin{array}{lll}
 & \mathsf{app}\,(k_1 \bullet k_2)\,e & \\
= & \mathsf{app}\,(k_1 \bullet Hole)\,e & \{\ assumption\ \} \\
= & \mathsf{app}\,k_1\,e & \{\ def \bullet\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,Hole\,e) & \{\ def\ app\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k_2\,e) & \{\ def\ k_2\ \}
\end{array}
$$

and case $k_2 = \mathsf{A}_i\,x_1 \ldots x_m\,k_3$

$$
\begin{array}{lll}
 & \mathsf{app}\,(k_1 \bullet k_2)\,e & \\
= & \mathsf{app}\,(k_1 \bullet \mathsf{A}_i\,x_1 \ldots x_m\,k_3)\,e & \{\ assumption\ \} \\
= & \mathsf{app}\,(\mathsf{A}_i\,x_1 \ldots x_m\,(k_1 \bullet k_3))\,e & \{\ def \circ\ \} \\
= & [\![\,\mathsf{E}_i[e \mid x_1, \ldots, x_m]\,]\!]_{f,k} & \{\ def\ app,\ k = k_1 \bullet k_3\ \} \\
= & \mathsf{app}\,(k_1 \bullet k_3)\,(\mathsf{E}_i[e \mid x_1, \ldots, x_m]) & \{\ spec\,(b)\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k_3\,(\mathsf{E}_i[e \mid x_1, \ldots, x_m])) & \{\ inductive\ hypothesis\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,(\mathsf{A}_i\,x_1 \ldots x_m\,k_3)\,e) & \{\ def\ app\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k_2\,e) & \{\ def\ app\ \}
\end{array}
$$

For application we have:

$$
\begin{array}{lll}
 & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{E}_i)\,e & \\
= & \mathsf{app}\,(\mathsf{A}_i\,x_1, \ldots x_m\,Hole)\,e & \{\ def\ ctx\ \} \\
= & [\![\,\mathsf{E}_i[e \mid x_1, \ldots, x_m]\,]\!]_{f,k} & \{\ def\ app,\ k = Hole\ \} \\
= & \mathsf{app}\,Hole\,(\mathsf{E}_i[e \mid x_1, \ldots x_m]) & \{\ spec\,(b)\ \} \\
= & \mathsf{E}_i[e \mid x_1, \ldots x_m] & \{\ def\ app\ \} \\
= & \mathsf{E}_i[e] &
\end{array}
$$

□

B.3 Context Laws for Right-biased-contexts

$$
\begin{array}{lll}
 & \mathsf{app}\,(k_1 \bullet k_2)\,e & \\
= & \mathsf{app}\,(k_2 \odot k_1)\,e & \{\ (rcomp)\ \} \\
= & e \odot (k_2 \odot k_1) & \{\ (rapp)\ \} \\
= & (e \odot k_2) \odot k_1 & \{\ assoc.\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k_2\,e) & \{\ (rapp)\ \}
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

and for context application we have:

$$
\begin{array}{lll}
 & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{A})\,e & \\
= & \mathsf{app}\,(\!|\mathsf{A}|\!)\,e & \{\ (rctx)\ \} \\
= & e \odot (\!|\mathsf{A}|\!) & \{\ (rapp)\ \}
\end{array}
$$

We proceed by induction over A.

Case $\mathsf{A} = \Box$:

$$
\begin{array}{lll}
= & e \odot (\!|\Box|\!) & \\
= & e \odot \mathit{unit} & \{\ \mathit{fold}\ \} \\
= & e & \{\ \mathit{unit}\ \} \\
= & \Box[e] & \{\ \Box\ \}
\end{array}
$$

and the case $\mathsf{A} = \mathsf{A}' \odot v$:

$$
\begin{array}{lll}
= & e \odot (\!|\mathsf{A}' \odot v|\!) & \\
= & e \odot ((\!|\mathsf{A}'|\!) \odot v) & \{\ \mathit{fold}\ \} \\
= & (e \odot (\!|\mathsf{A}'|\!)) \odot v & \{\ \mathit{assoc.}\ \} \\
= & \mathsf{A}'[e] \odot v & \{\ \mathit{induction\ hyp.}\ \} \\
= & \mathsf{A}[e] & \{\ \mathsf{A}\ \mathit{context}\ \}
\end{array}
$$

B.4 General Monoid Contexts

$$
\begin{array}{lll}
 & \mathsf{app}\,((l_1, r_1) \bullet (l_2, r_2))\,e & \\
= & \mathsf{app}\,(l_1 \odot l_2, r_2 \odot r_1)\,e & \{\ (acomp)\ \} \\
= & (l_1 \odot l_2) \odot e \odot (r_2 \odot r_1) & \{\ (aapp)\ \} \\
= & (l_1 \odot (l_2 \odot e \odot r_2) \odot r_1) & \{\ \mathit{assoc.}\ \} \\
= & \mathsf{app}\,(l_1, r_1)\,(\mathsf{app}\,(l_2, r_2)\,e) & \{\ (aapp)\ \}
\end{array}
$$

and

$$
\begin{array}{lll}
 & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{A})\,e & \\
= & \mathsf{app}\,(\!|\mathsf{A}|\!)\,e & \{\ (actx)\ \} \\
= & l \odot e \odot r & \{\ (aapp), \mathit{for}\,(l, r) = (\!|\mathsf{A}|\!)\ \}
\end{array}
$$

We proceed by induction over A: case $\mathsf{A} = \Box$:

$$
\begin{array}{lll}
= & l \odot e \odot r & \{\ \mathit{for}\,(l, r) = (\!|\Box|\!)\ \} \\
= & \mathit{unit} \odot e \odot \mathit{unit} & \{\ \mathit{fold}\ \} \\
= & e & \{\ \mathit{unit}\ \} \\
= & \Box[e] & \{\ \Box\ \}
\end{array}
$$

and $\mathsf{A} = v \odot \mathsf{A}'$:

$$
\begin{array}{lll}
= & l \odot e \odot r & \{\ \mathit{for}\,(l, r) = (\!|v \odot \mathsf{A}'|\!)\ \} \\
= & (v \odot l) \odot e \odot r & \{\ \mathit{fold}, \mathit{for}\,(l, r) = (\!|\mathsf{A}'|\!)\ \} \\
= & v \odot (l \odot e \odot r) & \{\ \mathit{assoc.}, \mathit{for}\,(l, r) = (\!|\mathsf{A}'|\!)\ \} \\
= & v \odot \mathsf{A}'[e] & \{\ \mathit{induction\ hyp.}, \mathit{for}\,(l, r) = (\!|\mathsf{A}'|\!)\ \} \\
= & \mathsf{A}[e] & \{\ \mathsf{A}\ \mathit{context}\ \}
\end{array}
$$

and $\mathsf{A} = \mathsf{A}' \odot v$:

$$
\begin{array}{lll}
= & l \odot e \odot r & \{\ \mathit{for}\,(l, r) = (\!|\mathsf{A}' \odot v|\!)\ \} \\
= & l \odot e \odot (r \odot v) & \{\ \mathit{fold}, \mathit{for}\,(l, r) = (\!|\mathsf{A}'|\!)\ \} \\
= & (l \odot e \odot r) \odot v & \{\ \mathit{assoc.}, \mathit{for}\,(l, r) = (\!|\mathsf{A}'|\!)\ \} \\
= & \mathsf{A}'[e] \odot v & \{\ \mathit{induction\ hyp.}, \mathit{for}\,(l, r) = (\!|\mathsf{A}'|\!)\ \} \\
= & \mathsf{A}[e] & \{\ \mathsf{A}\ \mathit{context}\ \}
\end{array}
$$

□

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

B.5 Context Laws for Exponent Contexts

We prove the composition law by induction on k_2:

$$
\begin{array}{lll}
 & \mathsf{app}\,(k_1 \bullet k_2)\,e & \\
= & \mathsf{app}\,(k_1 + k_2)\,e & \\
= & \mathsf{app}\,k_1\,e & \{\ \mathit{case}\,k_2 = 0\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,0\,e) & \{\ (\mathit{xapp})\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k_2\,e) & \{\ k_2 = 0\ \}
\end{array}
$$

and

$$
\begin{array}{lll}
 & \mathsf{app}\,(k_1 \bullet k_2)\,e & \\
= & \mathsf{app}\,(k_1 + (k' + 1))\,e & \{\ \mathit{case}\,k_2 = k' + 1\ \} \\
= & \mathsf{app}\,((k_1 + k') + 1)\,e & \{\ \mathit{assoc.}\ \} \\
= & \mathsf{app}\,(k_1 + k')\,(g\,e) & \{\ (\mathit{xapp})\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k'\,(g\,e)) & \{\ \mathit{inductive\ hyp.}\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,(k' + 1)\,e) & \{\ (\mathit{xapp})\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k_2\,e) & \{\ k_2 = k' + 1\ \}
\end{array}
$$

Appliction can be derived as:

$$
\begin{array}{lll}
 & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{A})\,e & \\
= & \mathsf{app}\,(\!|\mathsf{A}|\!)\,e & \{\ (\mathit{xctx})\ \}
\end{array}
$$

We proceed by induction over A: case $\mathsf{A} = \Box$:

$$
\begin{array}{lll}
= & \mathsf{app}\,(\!|\Box|\!)\,e & \\
= & \mathsf{app}\,0\,e & \{\ \mathit{fold}\ \} \\
= & e & \{\ (\mathit{xapp})\ \} \\
= & \Box[e] & \{\ \Box\ \}
\end{array}
$$

and $\mathsf{A} = g\,\mathsf{A}'$:

$$
\begin{array}{lll}
= & \mathsf{app}\,(\!|g\,\mathsf{A}'|\!)\,e & \\
= & \mathsf{app}\,((\!|\mathsf{A}'|\!) + 1)\,e & \{\ \mathit{fold}\ \} \\
= & \mathsf{app}\,(\!|\mathsf{A}'|\!)\,(g\,e) & \{\ (\mathit{xapp})\ \} \\
= & \mathsf{A}'[g\,e] & \{\ \mathit{induction\ hyp.}\ \} \\
= & \mathsf{A}[e] & \{\ \mathsf{A}\ \mathit{context}\ \}
\end{array}
$$

□

B.6 Constructor Contexts

Composition:

$$
\begin{array}{lll}
 & \mathsf{app}\,(k_1 \bullet k_2)\,e & \\
= & \mathsf{app}\,(k_1[k_2])\,e & \{\ (\mathit{kcomp})\ \} \\
= & (k_1[k_2])[e] & \{\ (\mathit{kapp})\ \} \\
= & k_1[k_2[e]] & \{\ \mathit{contexts}\ \} \\
= & k_1[\mathsf{app}\,k_2\,e] & \{\ (\mathit{kapp})\ \} \\
= & \mathsf{app}\,k_1\,(\mathsf{app}\,k_2\,e) & \{\ (\mathit{kapp})\ \}
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

$$\frac{x : \tau \in \Gamma}{\Gamma\,;\, \varnothing \vdash_{\mathrm{M}} x : \tau}\ [\textsc{var}] \qquad \frac{\Gamma|_{\mathsf{N}} \uplus \{x : \tau_1\}\,;\, \varnothing \vdash_{\mathrm{M}} M : \tau_2}{\Gamma\,;\, \varnothing \vdash_{\mathrm{M}} \lambda x : \tau_1 . M : \tau_1 \to \tau_2}\ [\textsc{abs}]$$

$$\frac{}{\Gamma\,;\, x : \tau \vdash_{\mathrm{M}} x : \tau}\ [\textsc{hle}] \qquad \frac{\Gamma\,;\, x : \tau_1 \vdash_{\mathrm{M}} M : \tau_2}{\Gamma\,;\, \varnothing \vdash_{\mathrm{M}} \hat{\lambda} x : \tau_1 . M : (\tau_1, \tau_2)\, \mathit{hfun}}\ [\textsc{hfun}]$$

$$\frac{\Gamma_1\,;\, \varnothing \vdash_{\mathrm{M}} M_1 : \tau_1 \to \tau_2 \quad \Gamma_2\,;\, \varnothing \vdash_{\mathrm{M}} M_2 : \tau_1}{\Gamma_1 \uplus \Gamma_2\,;\, \varnothing \vdash_{\mathrm{M}} M_1\, M_2 : \tau_2}\ [\textsc{app}]$$

$$\frac{\Gamma_1\,;\, \varnothing \vdash_{\mathrm{M}} M_1 : (\tau_1, \tau_2)\, \mathit{hfun} \quad \Gamma_2\,;\, H \vdash_{\mathrm{M}} M_2 : \tau_1}{\Gamma_1 \uplus \Gamma_2\,;\, H \vdash_{\mathrm{M}} \mathit{happ}\, M_1\, M_2 : \tau_2}\ [\textsc{happ}]$$

$$\frac{\Gamma_i\,;\, H_i \vdash_{\mathrm{M}} M_i : \tau_i \quad C^k : \tau_1 \to \ldots \to \tau_k \to \tau}{\uplus_i \Gamma_i\,;\, \oplus_i H_i \vdash_{\mathrm{M}} C^k\, M_1 \ldots M_k : \tau}\ [\textsc{cons}] \qquad \frac{C^k : \tau \in \Gamma}{\Gamma \vdash_{\mathrm{M}} C^k : \tau}\ [\textsc{con}]$$

$$\frac{\Gamma_1\,;\, \varnothing \vdash_{\mathrm{M}} M : \tau_1 \quad \Gamma_2\,;\, \varnothing \vdash_{\mathrm{PAT}} p_i : \tau_1 \mapsto M_i : \tau_2}{\Gamma_1 \uplus \Gamma_2\,;\, \varnothing \vdash_{\mathrm{M}} \mathsf{match}\, M\, \{\, \overline{p_i \mapsto M_i}\, \} : \tau_2}\ [\textsc{match}] \qquad \frac{f : \tau \in \Gamma}{\Gamma \vdash_{\mathrm{M}} f : \tau}\ [\textsc{fun}]$$

$$\frac{\Gamma, f : \tau\,;\, \varnothing \vdash_{\mathrm{M}} \lambda x . e : \tau}{\Gamma\,;\, \varnothing \vdash_{\mathrm{M}} \mathsf{fun} f = \lambda x . e : \tau}\ [\textsc{fundecl}] \qquad \frac{\Gamma\,;\, \varnothing \vdash_{\mathrm{M}} C^k : \tau_1 \to \ldots \to \tau_k \to \tau \quad \Gamma, x_1 : \tau_1, \ldots, x_k : \tau_k\,;\, \varnothing \vdash M_i : \tau'}{\Gamma\,;\, \varnothing \vdash_{\mathrm{PAT}} C^k\, x_1 \ldots x_k : \tau \mapsto e_i : \tau'}\ [\textsc{pat}]$$

$$\frac{\Gamma_1\,;\, \varnothing \vdash_{\mathrm{M}} M_1 : \tau_1 \quad \Gamma_2, x : \tau_1\,;\, \varnothing \vdash_{\mathrm{M}} M_2 : \tau_1}{\Gamma_1 \uplus \Gamma_2\,;\, \varnothing \vdash_{\mathrm{M}} \mathsf{let}\, x = M_1\, \mathsf{in}\, M_2 : \tau_2}\ [\textsc{let}]$$

Fig. B1. Minamide's type system adapted to our language

and application:

$$\begin{array}{lll} & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{K})\,e & \\ = & \mathsf{app}\,\mathsf{K}\,e & \{\ (\mathit{kctx})\,\} \\ = & \mathsf{K}[e] & \{\ (\mathit{kapp})\ \} \end{array}$$

□

B.7 Constructor Contexts and Minamide

The hole calculus is restricted by a linear-type discipline where the contexts $\mathit{ctx}\,\alpha \equiv \mathit{hfun}\,\alpha\,\alpha$ have a linear type. This is what enables an efficient in-place update implementation while still having a pure functional interface. For our needs, we need to check separately that the translation ensures that all uses of a context k are indeed linear. Type judgments in Minamide's system (Minamide, 1998, fig. 4) are denoted as $\Gamma\,;\, H \vdash_{\mathrm{M}} e : \tau$ where Γ is the normal type environment and H is the linear one containing at most one linear value. The type environment Γ can still contain linear values with a linear-type but

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

only pass those to one of the premises. The environment restricted to nonlinear values is denoted at $\Gamma|_{\mathsf{N}}$. We can now show that our translation can be typed in Minamide's system:

Lemma 4. (*TRMC uses contexts linearly*)
If $\Gamma|_{\mathsf{N}}\,;\,\varnothing \vdash_{\mathrm{M}} \mathsf{fun} f = \lambda xs.\, e : \tau_1 \to \ldots \to \tau_n \to \tau$ and k fresh
then $\Gamma|_{\mathsf{N}}, f\,;\,\varnothing \vdash_{\mathrm{M}} \mathsf{fun} f' = \lambda xs.\, \lambda k.\, [\![e]\!]_{f,k} : \tau_1 \to \ldots \to \tau_n \to ((\tau,\, \tau)\, hfun) \to \tau$.

To show this, we need a variant of the general replacement lemma (Hindley & Seldin, 1986, Lemma 11.18; Wright & Felleisen, 1994, Lemma 4.2) to reason about linear substitution in an evaluation context:

Lemma 5. (*Linear replacement*)
If $\Gamma|_{\mathsf{N}}\,;\,\varnothing \vdash_{\mathrm{M}} \mathsf{K}[e] : \tau$ for a constructor context K, then there is a sub-deduction $\Gamma|_{\mathsf{N}}\,;\,\varnothing \vdash_{\mathrm{M}} e : \tau'$ at the hole and $\Gamma|_{\mathsf{N}}\,;\,x : \tau' \vdash_{\mathrm{M}} \mathsf{K}[x] : \tau$.

Proof. By induction over the constructor context K.

Case $\Box$.

$\Gamma|_{\mathsf{N}}\;;\varnothing \vdash_{\mathrm{M}} \Box[e] : \tau$ { *assumption* }
$\Gamma|_{\mathsf{N}}\;;\varnothing \vdash_{\mathrm{M}} e : \tau$ { *subject reduction* }
$\Gamma|_{\mathsf{N}}\;;x : \tau \vdash_{\mathrm{M}} x : \tau$ { [*hle*] }
$\Gamma|_{\mathsf{N}}\;;x : \tau \vdash_{\mathrm{M}} \Box[x] : \tau'$ { *definition* }
$\Gamma|_{\mathsf{N}}\;;x : \tau \vdash_{\mathrm{M}} \mathsf{E}[x] : \tau'$ { *definition* }

Case $C^k\, w_1 \ldots \mathsf{K}' \ldots w_k$.

$\Gamma|_{\mathsf{N}}\;;\varnothing \vdash_{\mathrm{M}} C^k\, w_1 \ldots \mathsf{K}'[e] \ldots w_k : \tau$ { *assumption* }
$\Gamma|_{\mathsf{N}}\;;\varnothing \vdash_{\mathrm{M}} w_i : \tau_i \quad \text{for } i \neq j$ { [*cons*] *and nonlinearity* }
$\Gamma|_{\mathsf{N}}\;;\varnothing \vdash_{\mathrm{M}} \mathsf{K}'[e] : \tau_j$ { [*cons*] }
$\Gamma|_{\mathsf{N}}\;;x : \tau' \vdash_{\mathrm{M}} \mathsf{K}'[x] : \tau_j$ { *inductive hypothesis* }
$\Gamma|_{\mathsf{N}}\;;x : \tau' \vdash_{\mathrm{M}} C^k\, w_1 \ldots \mathsf{K}'[x] \ldots w_k : \tau$ { [*cons*] }

Again we see that our maximal context is an evaluation context as we would not be able to derive the Lemma for contexts under lambda's for example (as the linear-type environment is not propagated under lambda's).

Proof. (*Of Theorem 4*) By the FUNDECL and ABS rules we obtain:

$\Gamma_1 \;= \Gamma|_{\mathsf{N}}, f : \tau_1 \to \ldots \to \tau_n \to \tau,\, x_1 : \tau_1,\, \ldots,\, x_n : \tau_n$
$\Gamma_1 \;;\varnothing \vdash_{\mathrm{M}} e : \tau$ { *inductive property* }

By the FUNDECL and ABS rules, we need to derive:

$$\Gamma_2 \;= \Gamma|_{\mathsf{N}},\, f : \tau_1 \to \ldots \to \tau_n \to \tau,\; f' : \tau_1 \to \ldots \to \tau_n \to ((\tau,\, \tau)\, hfun) \to \tau,\, x_1 : \tau_1,\, \ldots,\, x_n : \tau_n$$

$$\Gamma_2,\quad k : ((\tau,\, \tau)\, hfun)\,;\, \varnothing \vdash_{\mathrm{M}} [\![e]\!]_{f,k} : \tau$$

In particular, we have $\Gamma_1 \subseteq \Gamma_2$. We proceed by induction over the translation function while maintaining the inductive property.

Case (*base*).

$$[\![e]\!]_{f,k} \;= \mathsf{app}\, k\, e = happ\, k\, e$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

$k \quad : (\tau, \tau)\, hfun\,;\, \varnothing \vdash_{\mathrm{M}} k : (\tau, \tau)\, hfun \quad \{\, [hle] \,\}$
$\Gamma_1 \quad ;\, \varnothing \vdash_{\mathrm{M}} e : \tau \quad \{\, assumption \,\}$
$\Gamma_2 \quad ;\, \varnothing \vdash_{\mathrm{M}} e : \tau \quad \{\, weakening \,\}$
$\Gamma_2, \quad k : (\tau, \tau)\, hfun\,;\, \varnothing \vdash_{\mathrm{M}} happ\, k\, e \quad \{\, [happ] \,\}$

Case (*tail*), $e = \mathsf{K}[f\, e_1 \ldots e_n]$.

$[\![e]\!]_{f,k} \quad = f'\, e_1 \ldots e_n\, (k \bullet \mathsf{ctx}\, \mathsf{K}) = f'\, e_1 \ldots e_n\, (hcomp\, k\, (\hat{\lambda} x.\, \mathsf{K}[x]))$

$\Gamma_1 \quad ;\, \varnothing \vdash_{\mathrm{M}} \mathsf{K}[f\, e_1 \ldots e_n] : \tau \quad \{\, assumption \,\}$
$\Gamma_2 \quad ;\, \varnothing \vdash_{\mathrm{M}} \mathsf{K}[f\, e_1 \ldots e_n] : \tau \quad \{\, weakening \,\}$
$\Gamma_2 \quad ;\, x : \tau' \vdash_{\mathrm{M}} \mathsf{K}[x] : \tau \quad \{\, linear\ replacement\ with\ nonlinearity\ of\ \Gamma_2 \,\}$
$\Gamma_2 \quad ;\, \varnothing \vdash_{\mathrm{M}} \hat{\lambda} x.\, \mathsf{K}[x] : (\tau, \tau)\, hfun \quad \{\, [hfun] \,\}$
$\Gamma_2, \quad k : (\tau, \tau)\, hfun\,;\, \varnothing \vdash_{\mathrm{M}}$
$\quad hcomp\, k\, (\hat{\lambda} x.\, \mathsf{K}[x]) : (\tau, \tau)\, hfun \quad \{\, hcomp, [happ], [hfun] \,\}$
$\Gamma_2 \quad ;\, \varnothing \vdash_{\mathrm{M}} f\, e_1 \ldots e_n : \tau' \quad \{\, linear\ replacement\ with\ nonlinearity\ of\ \Gamma_2 \,\}$
$\Gamma_2 \quad ;\, \varnothing \vdash_{\mathrm{M}} e_i : \tau_i \quad \{\, [app] \,\}$
$\Gamma_2, \quad k : (\tau, \tau)\, hfun\,;\, \varnothing \vdash_{\mathrm{M}} f'\, e_1 \ldots e_n$
$\quad (hcomp\, k\, (\hat{\lambda} x.\, \mathsf{K}[x])) \quad \{\, [app] \,\}$

Case (*let*), $e = \mathsf{let}\, x = e_1\, \mathsf{in}\, e_2$.

$[\![e]\!]_{f,k} \quad = \mathsf{let}\, x = e_1\, \mathsf{in}\, [\![e_2]\!]_{f,k}$

$\Gamma_1 \quad ;\, \varnothing \vdash_{\mathrm{M}} \mathsf{let}\, x = e_1\, \mathsf{in}\, e_2 : \tau \quad \{\, assumption \,\}$
$\Gamma_1 \quad ;\, \varnothing \vdash_{\mathrm{M}} e_1 : \tau_1 \quad \{\, [let] \,\}$
$\Gamma_2 \quad ;\, \varnothing \vdash_{\mathrm{M}} e_1 : \tau_1 \quad \{\, weakening \,\}$
$\Gamma_1, \quad x : \tau_1\,;\, \varnothing \vdash_{\mathrm{M}} e_2 : \tau \quad \{\, [let] \,\}$
$\Gamma_2, \quad k : (\tau, \tau)\, hfun, x : \tau_1\,;\, \varnothing \vdash_{\mathrm{M}} [\![e_2]\!]_{f,k} : \tau \quad \{\, inductive\ hypothesis \,\}$
$\Gamma_2, \quad k : (\tau, \tau)\, hfun, ;\, \varnothing \vdash_{\mathrm{M}} \mathsf{let}\, x = e_1\, \mathsf{in}\, [\![e_2]\!]_{f,k} : \tau \quad \{\, [let] \,\}$

Case (*match*), $e = \mathsf{match}\, e_1\, \{p_i \mapsto e_i\}$.

$[\![e]\!]_{f,k} \quad = \mathsf{match}\, e_1\, \{p_i \mapsto [\![e_i]\!]_{f,k}\}$

$\Gamma_1 \quad ;\, \varnothing \vdash_{\mathrm{M}} \mathsf{match}\, e_1\, \{p_i \mapsto e_i\} : \tau \quad \{\, assumption \,\}$
$\Gamma_1 \quad ;\, \varnothing \vdash_{\mathrm{M}} e_1 : \tau' \quad \{\, [match] \,\}$
$\Gamma_2 \quad ;\, \varnothing \vdash_{\mathrm{M}} e_1 : \tau' \quad \{\, weakening \,\}$
$\Gamma_1 \quad ;\, \varnothing \vdash_{\mathrm{PAT}} p_i \mapsto e_i : \tau \quad \{\, [match] \,\}$
$\Gamma_1 \quad ;\, \varnothing \vdash_{\mathrm{M}} C^k : \tau_1 \to \ldots \to \tau_k \to \tau' \quad \{\, [pat] \,\}$
$\Gamma_1, \quad x_1 : \tau_1, \ldots, x_k : \tau_k\,;\, \varnothing \vdash_{\mathrm{M}} e_i : \tau \quad \{\, [pat] \,\}$
$\Gamma_2, \quad k : (\tau, \tau)\, hfun, x_1 : \tau_1, \ldots, x_k : \tau_k\,;\, \varnothing \vdash_{\mathrm{M}} [\![e_i]\!]_{f,k} : \tau \quad \{\, inductive\ hypothesis \,\}$
$\Gamma_2, \quad k : (\tau, \tau)\, hfun\,;\, \varnothing \vdash_{\mathrm{PAT}} p_i \mapsto [\![e_i]\!]_{f,k} : \tau \quad \{\, [pat] \,\}$
$\Gamma_2, \quad k : (\tau, \tau)\, hfun\,;\, \varnothing \vdash_{\mathrm{M}} \mathsf{match}\, e_1\, \{p_i \mapsto [\![e_i]\!]_{f,k}\} : \tau \quad \{\, [match] \,\}$

□

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

B.8 Contexts Form Linear Chains

Proof. (*Of Lemma 2*) By induction on the shape of K:

Case $C \ldots \Box_i \ldots$:

$$
\begin{array}{lll}
 & H \mid C \ldots \Box_i \ldots & \\
\longrightarrow_{\mathsf{r}}^{*} & H,\, x \mapsto^1 C \ldots \Box_i \ldots \mid x & \{\ (con_r)\ \} \\
= & H,\, [x \mapsto^1 C \ldots \Box_i \ldots]^1_x \mid x & \{\ \textit{linear chain}\ \}
\end{array}
$$

Case $C \ldots \mathsf{K}'[C' \ldots \Box_i \ldots] \ldots$

$$
\begin{array}{lll}
 & H \mid (\!| C \ldots \mathsf{K}'[C' \ldots \Box_i \ldots] \ldots |\!) & \\
\longrightarrow_{\mathsf{r}}^{*} & H,\, [H', y \mapsto^1 C' \ldots \Box_i \ldots]^1_{x'} \mid \langle C \ldots x' \ldots \rangle & \{\ \textit{induction hyp.}\ \} \\
\longrightarrow_{\mathsf{r}} & H,\, x \mapsto^1 C \ldots x' \ldots,\, [H',\, y \mapsto^1 C' \ldots \Box_i \ldots]^1_{x'} \mid x & \{\ (con_r)\ \} \\
= & H,\, [x \mapsto^1 C \ldots x' \ldots,\, H',\, y \mapsto^1 C' \ldots \Box_i \ldots]^1_x \mid x & \{\ \textit{linear chain}\ \}
\end{array}
$$

□

B.9 Deriving Constructor Context Fold

Given the specification:

$$(\textit{foldspec}) \quad H \mid (\!| \mathsf{K}[C \ldots \Box_i \ldots] |\!) \cong H \mid \mathsf{let}\, x = \mathsf{K}[C \ldots \Box_i \ldots]\, \mathsf{in}\, \langle x, [x]@i \rangle$$

we can calculate the fold using induction over the shape of K. In the case that $\mathsf{K} = \Box$, we derive:

$$
\begin{array}{lll}
 & H \mid (\!| C \ldots \Box_i \ldots |\!) & \\
\cong & H \mid \mathsf{let}\, x = C \ldots \Box_i \ldots \mathsf{in}\, \langle x, [x]@i \rangle & \{\ \textit{specification}\ \} \\
\cong & H, x \mapsto^1 C \ldots \Box_i \ldots \mid \langle x, [x]@i \rangle & \{\ (let_r),\ (con_r),\ \mathbf{1}\ \} \\
= & H, [x \mapsto^1 C \ldots \Box_i \ldots]^1_x \mid \langle x, [x]@i \rangle & \{\ \textit{linear chain}\ \} \\
= & H, [x \mapsto^1 C \ldots \Box_i \ldots]^1_x \mid \langle x, x@i \rangle & \{\ \textit{def}.\ \} \\
\cong & H \mid \mathsf{let}\, x = C \ldots \Box_i \ldots \mathsf{in}\, \langle x, x@i \rangle & \{\ (let_r),\ (con_r),\ \mathbf{1}\ \}
\end{array}
$$

and otherwise, K has the form $C' \ldots \mathsf{K}' \ldots$ where $(\!| \mathsf{K}'[C \ldots \Box_i \ldots] |\!) = \mathsf{let}\, x = \mathsf{K}'[C \ldots \Box_i \ldots]\, \mathsf{in}\, \langle x, [x]@i \rangle$ (by induction):

$$
\begin{array}{lll}
 & H \mid (\!| C' \ldots \mathsf{K}'[C \ldots \Box_i \ldots] \ldots |\!) & \\
\cong & H \mid \mathsf{let}\, x = C' \ldots \mathsf{K}'[C \ldots \Box_i \ldots] \ldots \mathsf{in}\, \langle x, [x]@i \rangle & \{\ \textit{specification}\ \} \\
\cong & H \mid \mathsf{let}\, z = \mathsf{K}'[C \ldots \Box_i \ldots]\, \mathsf{in}\, \mathsf{let}\, x = C \ldots z \ldots \mathsf{in}\, \langle x, [x]@i \rangle & \{\ (let_r)\ \} \\
\cong & H \mid \mathsf{let}\, \langle z, [z]@i \rangle = (\!| \mathsf{K}'[C \ldots \Box_i \ldots] |\!)\, \mathsf{in} & \\
 & \mathsf{let}\, x = C \ldots z \ldots \mathsf{in}\, \langle x, [x]@i \rangle & \{\ \textit{calculate}\ \} \\
\cong & H,\, [H', y \mapsto^1 C \ldots \Box_i \ldots]^1_z,\, x \mapsto^1 C \ldots z \ldots \mid \langle x, [x.i] \rangle & \{\ (let_r),\ \textit{lemma 2},\ \mathbf{1}\ \} \\
= & H,\, [x \mapsto^1 C \ldots z \ldots,\, [H', y \mapsto^1 C \ldots \Box_i \ldots]^1_z]^1_x,\ \mid \langle x, [x]@i \rangle & \{\ \textit{linear chain}\ \} \\
= & H,\, [x \mapsto^1 C \ldots z \ldots,\, [H', y \mapsto^1 C \ldots \Box_i \ldots]^1_z]^1_x,\ \mid \langle x, y@i \rangle & \{\ \textit{def}.\ \} \\
\cong & H \mid \mathsf{let}\, \langle z, y@i \rangle = (\!| \mathsf{K}'[C \ldots \Box_i \ldots] |\!)\, \mathsf{in}\, \langle C \ldots z \ldots, y@i \rangle & \{\ (let_r),\ (con_r)\,(1)\ \}
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

B.10 Deriving Constructor Context Composition

We can calculate for a $\mathsf{K}_1, \mathsf{K}_2 \neq \Box$:

$$
\begin{array}{lll}
 & H \mid \mathsf{app}\,(\mathsf{ctx}\,\mathsf{K}_1 \bullet \mathsf{ctx}\,\mathsf{K}_2)\,e & \\
\cong & H \mid \mathsf{app}\,(\mathsf{let}\,x_1 = \mathsf{K}_1[\Box]\,\mathsf{in}\,\langle x_1, [x_1]@i\rangle) \bullet (\mathsf{ctx}\,\mathsf{K}_2)\,e & \{\textit{fold specification},\ \mathsf{K}_1 \neq \Box\} \\
\cong & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots \Box_i \ldots]^1_{x_1} \mid \mathsf{app}\,(\langle x_1, [x_1]@i\rangle \bullet \mathsf{ctx}\,\mathsf{K}_2)\,e & \{\ \textit{lemma}\ 2,\ \mathbf{1}\ \} \\
\cong & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots \Box_i \ldots]^1_{x_1},\ [H_2, y_2 \mapsto^1 C_2 \ldots \Box_j \ldots]^1_{x_2} & \\
 & \quad \mid \mathsf{app}\,(\langle x_1, [x_1]@i\rangle \bullet \langle x_2, [x_2]@j\rangle)\,e & \{\textit{fold specification},\ \mathsf{K}_2 \neq \Box,\ \textit{lemma}\ 2,\ \mathbf{2}\} \\
= & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots \Box_i \ldots]^1_{x_1},\ [H_2, y_2 \mapsto^1 C_2 \ldots \Box_j \ldots]^1_{x_2} & \\
 & \quad \mid \mathsf{app}\,(\langle x_1, y_1@i\rangle \bullet \langle x_2, y_2@j\rangle)\,e & \{\ \textit{def}.\ \} \\
= & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots \Box_i \ldots]^1_{x_1},\ [H_2, y_2 \mapsto^1 C_2 \ldots \Box_j \ldots]^1_{x_2} & \\
 & \quad \mid \mathsf{app}\,\langle \mathsf{app}\,\langle x_1, y_1@i\rangle,\ y_2@j\rangle\,e & \{\ \textit{calculate}\ \} \\
= & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots x_2 \ldots]^1_{x_1},\ [H_2, y_2 \mapsto^1 C_2 \ldots \Box_j \ldots]^1_{x_2} & \\
 & \quad \mid \mathsf{app}\,\langle x_1, y_2@j\rangle)\,e & \{\ (\textit{uapp})\ \} \\
= & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots x_2 \ldots]^1_{x_1},\ [H_2, y_2 \mapsto^1 C_2 \ldots \Box_j \ldots]^1_{x_2}, & \\
 & \quad z \mapsto^1 v \mid \mathsf{app}\,\langle x_1, y_2@j\rangle\,z & \{\ e\ \textit{terminating},\ \mathbf{3}\ \} \\
\cong & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots \Box_i \ldots]^1_{x_1},\ [H_2, y_2 \mapsto^1 C_2 \ldots z \ldots]^1_{x_2}, & \\
 & \quad z \mapsto^1 v \mid \mathsf{app}\,\langle x_1, y_1@i\rangle\,x_2 & \{\ (\textit{app})\ \} \\
\cong & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots \Box_i \ldots]^1_{x_1},\ [H_2, y_2 \mapsto^1 C_2 \ldots \Box_j \ldots]^1_{x_2}, & \\
 & \quad z \mapsto^1 v \mid \mathsf{app}\,\langle x_1, y_1@i\rangle\,(\mathsf{app}\,\langle x_2, y_2@j\rangle\,z) & \{\ (\textit{app})\ \} \\
\cong & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots \Box_i \ldots]^1_{x_1},\ [H_2, y_2 \mapsto^1 C_2 \ldots z \ldots]^1_{x_2} & \\
 & \quad \mid \mathsf{app}\,\langle x_1, y_1@i\rangle\,(\mathsf{app}\,\langle x_2, y_2.j\rangle\,e & \{\ (3)\ \} \\
\cong & H,\ [H_1, y_1 \mapsto^1 C_1 \ldots \Box_i \ldots]^1_{x_1} \mid \mathsf{app}\,\langle x_1, y_1@i\rangle\,(\mathsf{app}\,\mathsf{ctx}\,\mathsf{K}_2)\,e & \{\ (2)\ \} \\
\cong & H \mid \mathsf{app}\,(\mathsf{ctx}\,\mathsf{K}_1)\,(\mathsf{app}\,(\mathsf{ctx}\,\mathsf{K}_2)\,e) & \{\ (1)\ \}
\end{array}
$$

and thus define composition as:

$$(\textit{ucomp})\ H \mid \langle x_1, y_1@i\rangle \bullet \langle x_2, y_2@j\rangle \quad \longrightarrow_{\mathsf{r}} \quad H \mid \langle \mathsf{app}\,\langle x_1, y_1@i\rangle\,x_2,\ y_2@j\rangle$$

In case the context is empty, we can calculate immediately:

$$
\begin{array}{lll}
 & H \mid \mathsf{app}\,(\mathsf{ctx}\,\Box)\,e & \\
= & H \mid \mathsf{app}\,(\!|\Box|\!)\,e & \{\ \textit{def}.\ \} \\
\cong & H \mid \mathsf{app}\,\langle\rangle\,e & \{\ \textit{fold specification}\ \} \\
\cong & H \mid e & \{\ \textit{calculate}\ \} \\
= & H \mid \Box[e] & \{\ \textit{contex}\ \}
\end{array}
$$

For the empty contexts, we can calculate for application:

$$
\begin{array}{lll}
 & \mathsf{app}\,(\mathsf{ctx}\,\Box \bullet \mathsf{ctx}\,\mathsf{K}_2)\,e & \\
= & \mathsf{app}\,(\!|\Box|\!) \bullet \mathsf{ctx}\,\mathsf{K}_2)\,e & \{\ \textit{def}.\ \} \\
\cong & \mathsf{app}\,(\langle\rangle \bullet \mathsf{ctx}\,\mathsf{K}_2)\,e & \{\ \textit{fold specification}\ \} \\
\cong & \mathsf{app}\,(\mathsf{ctx}\,\mathsf{K}_2)\,e & \{\ \textit{calculate}\ \} \\
\cong & \mathsf{K}[e] & \{\ (\textit{appctx})\ \} \\
\cong & \Box[\mathsf{K}[e]] & \{\ \textit{contexts}\ \}
\end{array}
$$

and similarly for $\mathsf{K}_2 = \Box$ (but note that in our translation we never have $k \bullet \mathsf{ctx}\,\Box$).

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press

B.11 Soundness of the Hybrid Approach

We need to show the context laws still hold for the hybrid approach.

At runtime, a context K is always a linear chain resulting from the fold or composition. We write $H \mid \hat{\mathsf{K}}$ for a non-empty context $[H', y \mapsto^m C \ldots \Box_i \ldots]^n_x \mid \langle x, y@i \rangle$ if we have $H_0 \mid (\!|\mathsf{K}|\!) \cong H_0,\ [H', y \mapsto^m C \ldots \Box_i \ldots]^1_x \mid \langle x, y@i \rangle$.

Application:

$$
\begin{array}{lll}
 & H \mid \mathsf{app}\,\hat{\mathsf{K}}\,e & \\
= & H,\ [H', y \mapsto^m C \ldots \Box_i \ldots]^{n+1}_x \mid \mathsf{app}\,\langle x, y@i \rangle\,e & \{\ (\mathsf{A}),\ \mathbf{1}\ \} \\
\cong & H,\ z \mapsto^1 v,\ [H', y \mapsto^m C_i \ldots \Box_i \ldots]^{n+1}_x \mid \mathsf{app}\,\langle x, y@i \rangle\,z & \{\ e\ \textit{is terminating}\ \mathbf{2}\ \} \\
= & H,\ z \mapsto^1 v,\ [H', y \mapsto^m C_i \ldots \Box_i \ldots]^{n+1}_x \mid \mathsf{append}\,x\,z & \{\ \textit{calculate}\ \}
\end{array}
$$

Now proceed by induction on H'. $H' = \Box$:

$$
\begin{array}{lll}
 & H,\ z \mapsto^1 v,\ [y \mapsto^{n+1} C_i \ldots \Box_i \ldots]^{n+1}_y \mid \mathsf{append}\,y\,z & \{\ \textit{singleton}\ \} \\
\cong & H,\ z \mapsto^1 v,\ [y \mapsto^{n+1} C_i \ldots \Box_i \ldots]^{n+1}_y \mid y.i\ \mathsf{as}\ z & \{\ \textit{calculate}\ \} \\
\cong & H,\ z \mapsto^1 v,\ [y \mapsto^n C_i \ldots \Box_i \ldots]^n_y,\ [x' \mapsto^1 C_i \ldots z \ldots]^1_{x'} \mid x' & \{\ (\textit{as})\ \} \\
\cong & H,\ z \mapsto^1 v,\ [y \mapsto^n C_i \ldots \Box_i \ldots]^n_y \mid \hat{\mathsf{K}}[z] & \{\ (1)\ \} \\
\cong & H,\ [y \mapsto^n C_i \ldots \Box_i \ldots]^n_y \mid \hat{\mathsf{K}}[e] & \{\ (2)\ \}
\end{array}
$$

and

$$
\begin{array}{lll}
 & H,\ z \mapsto^1 v,\ [x \mapsto^{n+1} C' \ldots y_i \ldots, [H_1]^1_y]^{n+1}_x & \\
 & \quad \mid \mathsf{append}\,x\,z & \\
\cong & H,\ z \mapsto^1 v,\ [x \mapsto^{n+1} C' \ldots y_i \ldots, [H_1]^1_y]^{n+1}_x & \\
 & \quad \mid \mathsf{dup}\,y_i;\ x.i\ \mathsf{as}\,(\mathsf{append}\,y_i\,z) & \{\ (\textit{append})\ \} \\
\cong & H,\ z \mapsto^1 v,\ [x \mapsto^{n+1} C' \ldots y_i \ldots, [H_1]^2_y]^{n+1}_x & \\
 & \quad \mid x.i\ \mathsf{as}\,(\mathsf{append}\,y\,z) & \\
\cong & H,\ z \mapsto^1 v,\ [x \mapsto^{n+1} C' \ldots y_i \ldots, [H_1]^1_y]^{n+1}_x,\ [H_2]^1_{y'} & \\
 & \quad \mid x.i\ \mathsf{as}\ y' & \{\ \textit{induction hyp.}\ \} \\
\cong & H,\ z \mapsto^1 v,\ [x \mapsto^n C' \ldots y_i \ldots, [H_1]^1_y]^n_x, & \\
 & [x' \mapsto^1 C' \ldots y'_i \ldots, [H'']^1_{y'}]^1_{x'} \mid x' & \{\ (\textit{as})\ \} \\
\cong & H,\ z \mapsto^1 v,\ [x \mapsto^n C' \ldots y_i \ldots, [H_1]^1_y]^n_x & \\
 & \quad \mid C' \ldots \hat{\mathsf{K}}'[z] \ldots & \\
\cong & H,\ z \mapsto^1 v,\ [x \mapsto^n C' \ldots y_i \ldots, [H_1]^1_y]^n_x & \\
 & \quad \mid \hat{\mathsf{K}}[z] & \\
\cong & H,\ [x \mapsto^n C' \ldots y_i \ldots, [H_1]^1_y]^n_x & \\
 & \quad \mid \hat{\mathsf{K}}[e] & \{\ (2)\ \}
\end{array}
$$

https://doi.org/10.1017/S0956796825100117 Published online by Cambridge University Press