

From algebra to abstract machine: a verified generic construction

Carlos Tomé Cortiñas

Department of Information and Computing Sciences
Utrecht University
The Netherlands
c.tomecortinas@students.uu.nl

Wouter Swierstra

Department of Information and Computing Sciences
Utrecht University
The Netherlands
w.s.swierstra@uu.nl

Abstract

Many functions over algebraic datatypes can be expressed in terms of a fold. Doing so, however, has one notable drawback: folds are not tail-recursive. As a result, a function defined in terms of a fold may raise a stack overflow when executed. This paper defines a datatype generic, tail-recursive higher-order function that is guaranteed to produce the same result as the fold. Doing so combines the compositional nature of folds and the performance benefits of a hand-written tail-recursive function in a single setting.

Keywords datatype generic programming, catamorphisms, dissection, dependent types, Agda, well-founded recursion

1 Introduction

Folds, or *catamorphisms*, are a pervasive programming pattern. Folds generalize many simple traversals over algebraic data types. Functions implemented by means of a fold are both compositional and structurally recursive. Consider, for instance, the following expression datatype, written in the programming language Agda [Norell 2007]:

```
data Expr : Set where
  Val  : ℕ → Expr
  Add  : Expr → Expr → Expr
```

We can write a simple evaluator, mapping expressions to natural numbers, as follows:

```
eval : Expr → ℕ
eval (Val n)      = n
eval (Add e1 e2) = eval e1 + eval e2
```

In the case for `Add e1 e2`, the `eval` function makes two recursive calls and sums their results. Such a function can be implemented using a fold, passing the addition and identity functions as the argument algebra.

```
fold : (ℕ → X) → (X → X → X) → Expr → X
fold φ1 φ2 (Val n)      = φ1 n
fold φ1 φ2 (Add e1 e2) = φ2 (fold φ1 φ2 e1) (fold φ1 φ2 e2)
eval : Expr → ℕ
eval = fold id _+_
```

Unfortunately, not everything in the garden is rosy. The operator `_+_` needs both of its parameters to be fully evaluated before it can reduce further. As a result, the size of the stack used during execution grows linearly with the size of the input, potentially leading to a stack overflow on large inputs.

To address this problem, we can manually rewrite the evaluator to be *tail-recursive*. Modern compilers typically map tail-recursive functions to machine code that runs in constant memory. To write such a tail-recursive function, we need to introduce an explicit stack storing both intermediate results and the subtrees that have not yet been evaluated.

```
data Stack : Set where
  Top  : Stack
  Left : Expr → Stack → Stack
  Right : ℕ → Stack → Stack
```

We can define a tail-recursive evaluation function by means of a pair of mutually recursive functions, `load` and `unload`. The `load` function traverses the expressions, pushing subtrees on the stack; the `unload` function unloads the stack, while accumulating a (partial) result.

```
mutual
load : Expr → Stack → ℕ
load (Val n)   stk = unload+ n stk
load (Add e1 e2) stk = load e1 (Left e2 stk)
unload+ : ℕ → Stack → ℕ
unload v Top      = v
unload v (Right v' stk) = unload+ (v' + v) stk
unload v (Left r stk)  = load r (Right v stk)
```

We can now define a tail-recursive version of `eval` by calling `load` with an initially empty stack:

```
tail-rec-eval : Expr → ℕ
tail-rec-eval e = load e Top
```

Implementing this tail-recursive evaluator comes at a price: Agda's termination checker flags the `load` and `unload` functions as potentially non-terminating by highlighting them orange. Indeed, in the very last clause of the `unload` function a recursive call is made to arguments that are not syntactically smaller. Furthermore, it is not clear at all that the tail-recursive evaluator produces the same result as our original one. It is precisely these issues that this paper tackles by making the following novel contributions:

- We give a verified proof of termination of `tail-rec-eval` using a carefully chosen *well-founded relation* (Sections 2 and 3). After redefining `tail-rec-eval` using this relation, we can prove the two evaluators equal in Agda.
- We generalize this relation and its corresponding proof of well-foundedness, inspired by McBride's work on *dissections* [McBride 2008], to show how to calculate an abstract machine from an algebra. To do so, we define a universe of algebraic data types and a generic fold operation (Section 4). Subsequently we show how to turn any structurally recursive function defined using a fold into its tail-recursive counterpart.
- Finally, we present how our proofs of termination and semantics preservation from our example are generalized to the generic fold (Sections 4.6 and 4.7).

Together these results give a verified function that computes a tail-recursive traversal from any algebra for any algebraic datatype. All the constructions and proofs presented in this paper have been implemented in and checked by Agda. The corresponding code is freely available online.¹

2 Termination and tail-recursion

Before tackling the generic case, we will present the termination and correctness proof for the tail-recursive evaluator presented in the introduction in some detail.

The problematic call for Agda's termination checker is the last clause of the `unload` function, that calls `load` on the expression stored on the top of the stack. From the definition of `load`, it is clear that we only ever push subtrees of the input on the stack. However, the termination checker has no reason to believe that the expression at the top of the stack is structurally smaller in any way. Indeed, if we were to redefine `load` as follows:

$$\text{load} (\text{Add } e_1 e_2) \text{ stk} = \text{load } e_1 (\text{Left } (f e_2) \text{ stk})$$

we might use some function $f : \text{Expr} \rightarrow \text{Expr}$ to push *arbitrary* expressions on the stack, potentially leading to non-termination.

The functions `load` and `unload` use the stack to store subtrees and partial results while folding the input expression. Thus, every node in the original tree is visited twice during the execution: first when the function `load` traverses the tree, until it finds the leftmost leaf; second when `unload` inspects the stack in searching of an unevaluated subtree. This process is depicted in Figure 1.

As there are finitely many nodes on a tree, the depicted traversal using `load` and `unload` must terminate – but how can we convince Agda's termination checker of this?

As a first approximation, we revise the definitions of `load` and `unload`. Rather than consuming the entire input in one




Figure 1. Traversing a tree with `load` and `unload`

go with a pair of mutually recursive functions, we rewrite them to compute one ‘step’ of the fold.

The function `unload` is defined by recursion over the stack as before, but with one crucial difference. Instead of always returning the final result, it may also² return a new configuration of our abstract machine, that is, a pair $\mathbb{N} \times \text{Stack}$:

$$\begin{aligned} \text{unload} &: \mathbb{N} \rightarrow \text{Stack} \rightarrow (\mathbb{N} \times \text{Stack}) \uplus \mathbb{N} \\ \text{unload } v \text{ Top} &= \text{inj}_2 v \\ \text{unload } v (\text{Right } v' \text{ stk}) &= \text{unload } (v' + v) \text{ stk} \\ \text{unload } v (\text{Left } r \text{ stk}) &= \text{load } r (\text{Right } v \text{ stk}) \end{aligned}$$

The other key difference arises in the definition of `load`:

$$\begin{aligned} \text{load} &: \text{Expr} \rightarrow \text{Stack} \rightarrow (\mathbb{N} \times \text{Stack}) \uplus \mathbb{N} \\ \text{load } (\text{Val } n) \text{ stk} &= \text{inj}_1 (n, \text{stk}) \\ \text{load } (\text{Add } e_1 e_2) \text{ stk} &= \text{load } e_1 (\text{Left } e_2 \text{ stk}) \end{aligned}$$

Rather than calling `unload` upon reaching a value, it returns the current stack and the value of the leftmost leaf. Even though the function never returns an `inj2`, its type is aligned with the type of `unload` so the definition of both functions resembles an abstract machine more closely.

Both these functions are now accepted by Agda's termination checker as they are clearly structurally recursive. We can use both these functions to define the following evaluator³:

$$\begin{aligned} \text{tail-rec-eval} &: \text{Expr} \rightarrow \mathbb{N} \\ \text{tail-rec-eval } e \text{ with load } e \text{ Top} \\ \dots \mid \text{inj}_1 (n, \text{stk}) &= \text{rec } (n, \text{stk}) \\ \text{where} \\ \text{rec} &: (\mathbb{N} \times \text{Stack}) \rightarrow \mathbb{N} \\ \text{rec } (n, \text{stk}) \text{ with unload } n \text{ stk} \\ \dots \mid \text{inj}_1 (n', \text{stk}') &= \text{rec } (n', \text{stk}') \\ \dots \mid \text{inj}_2 r &= r \end{aligned}$$

Here we use `load` to compute the initial configuration of our machine – that is, it finds the leftmost leaf in our initial expression and its associated stack. We proceed by repeatedly calling `unload` until it returns a value. This version of our evaluator, however, does not pass the termination checker. The new state, (n', stk') , is not structurally smaller than the initial state (n, stk) . If we work under the assumption that we

² \uplus is Agda's type of disjoint union.

³We ignore `load`'s impossible case, it can always be discharged with $\perp\text{-elim} : \forall \{X : \text{Set}\} \rightarrow \perp \rightarrow X$.

¹<https://github.com/carlostome/Dissection-thesis>

221 have a relation between the states $\mathbb{N} \times \text{Stack}$ that decreases
 222 after every call to `unload` and a proof that the relation is well-
 223 founded – we know this function will terminate eventually.
 224 We now define the following version of the tail-recursive
 225 evaluator:

```

226 tail-rec-eval : Expr → ℕ
227 tail-rec-eval e with load e Top
228 ... | inj1 (n, stk) = rec (n, stk) □1
229 where
230   rec : (c : ℕ × Stack) → Acc _<_ c → ℕ
231   rec (n, stk) (acc rs) with unload n stk
232   ... | inj1 (n', stk') = rec (n', stk') (rs □2)
233   ... | inj2 r = r
    
```

235 To complete this definition, we still need to define a suitable
 236 relation `_<_` between configurations of type $\mathbb{N} \times \text{Stack}$,
 237 prove the relation to be well-founded ($\square_1 : \text{Acc } _<_ (n, \text{stk})$)
 238 and show that the calls to `unload` produce ‘smaller’ states
 239 ($\square_2 : (n', \text{stk}') < (n, \text{stk})$). In the next section, we will
 240 define such a relation and prove it is well-founded.

243 3 Well-founded tree traversals

244 The type of configurations of our abstract machine can be
 245 seen as a variation of Huet’s *zippers* [1997]. The zipper asso-
 246 ciated with an expression $e : \text{Expr}$ is pair of a (sub)expression
 247 of e and its *context*. As demonstrated by McBride [2008], the
 248 zippers can be generalized further to *dissections*, where the
 249 values to the left and right of the current subtree may have
 250 different types. It is precisely this observation that we will ex-
 251 ploit when considering the generic tail-recursive traversals
 252 in the later sections; for now, however, we will only rely on
 253 the intuition that the configurations of our abstract machine,
 254 given by the type $\mathbb{N} \times \text{Stack}$, are an instance of *dissections*,
 255 corresponding to a partially evaluated expression:

```

257 Config : Set
258 Config = ℕ × Stack
    
```

260 These configurations, are more restrictive than dissections
 261 in general. In particular, the configurations presented in
 262 the previous section *only* ever denote a *leaf* in the input
 263 expression.

264 The tail-recursive evaluator, `tail-rec-eval` processes the
 265 leaves of the input expression in a left-to-right fashion. The
 266 leftmost leaf – that is the first leaf found after the initial
 267 call to `load` – is the greatest element; the rightmost leaf is
 268 the smallest. In our example expression from Section 1, we
 269 would number the leaves as follows:

270 This section aims to formalize the relation that orders
 271 elements of the `Config` type (that is, the configurations of
 272 the abstract machine) and prove it is *well-founded*. However,
 273 before doing so there are two central problems with our
 274 choice of `Config` datatype:




Figure 2. Numbered leaves of the tree

1. The `Config` datatype is too liberal. As we evaluate our input expression the configuration of our abstract machine changes constantly, but satisfies one important *invariant*: each configuration is a decomposition of the original input. Unless this invariant is captured, we will be hard pressed to prove the well-foundedness of any relation defined on configurations.
2. The choice of the `Stack` datatype, as a path from the leaf to the root is convenient to define the tail-recursive machine, but impractical when defining the coveted order relation. The top of a stack stores information about neighbouring nodes, but to compare two leaves we need *global* information about their positions relative to the root.

We will now address these limitations one by one. Firstly, by refining the type of `Config`, we will show how to capture the desired invariant (Section 3.1). Secondly, we explore a different representation of stacks, as paths from the root, that facilitates the definition of the desired order relation (Section 3.2). Finally we will define the relation over configurations, Section 3.3, and sketch the proof that it is well-founded.

3.1 Invariant preserving configurations

A value of type `Config` denotes a leaf in our input expression. In the previous example, the following `Config` corresponds to the third leaf:

As we observed previously, we would like to refine the type `Config` to capture the invariant that execution preserves: every `Config` denotes a unique leaf in our input expression, or equivalently, a state of the abstract machine that computes the fold. There is one problem still: the `Stack` datatype stores the values of the subtrees that have been evaluated, but does not store the subtrees themselves. In the example in Figure 3, when the traversal has reached the third leaf, all the subexpressions to its left have been evaluated.

In order to record the necessary information, we redefine the `Stack` type as follows:

```

data Stack+ : Set where
Left : Expr → Stack+ → Stack+
    
```

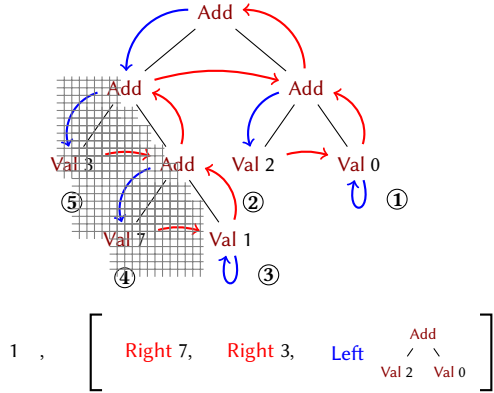


Figure 3. Example: Configuration of leaf number 3

```

Right : (n : ℕ) → (e : Expr) → eval e ≡ n → Stack+ → Stack+
Top   : Stack+

```

The **Right** constructor now not only stores the value n , but also records the subexpression e and the proof that e evaluates to n . Although we are modifying the definition of the **Stack** data type, we claim that the expression e and equality are not necessary at run-time, but only required for the proof of well-foundedness – a point we will return to in our discussion (Section 5). From now onwards, the type **Config** uses **Stack⁺** as its right component:

```
Config = ℕ × Stack+
```

The function **unload** was previously defined by induction over the stack (Section 2), thus, it needs to be modified to work over the new type of stacks, **Stack⁺**:

```

unload+ : (n : ℕ) → (e : Expr) → eval e ≡ n → Stack+
          → Config ⊔ ℕ
unload+ n e eq Top       = inj2 n
unload+ n e eq (Left e' stk) = load e' (Right n e eq stk)
unload+ n e eq (Right n' e' eq' stk)
= unload+ (n' + n) (Add e' e) (cong2 _+ _ e' eq' eq) stk

```

A value of type **Config** contains enough information to recover the input expression. This is analogous to the **plug** operation on zippers:

```

plugη : Expr → Stack+ → Expr
plugη e Top       = e
plugη e (Left t   stk) = plugη (Add e t) stk
plugη e (Right _ t _ stk) = plugη (Add t e) stk
plugCη : Config → Expr
plugCη (n, stk) = plugη (Val n) stk

```

Any two terms of type **Config** may still represent states of a fold over two entirely different expressions. As we aim to define an order relation comparing configurations during the fold of the input expression, we need to ensure that we only ever compare configurations within the same expression. We can *statically* enforce such requirement by defining a new wrapper data type over **Config** that records the original input expression:

```

data Configη (e : Expr) : Set where
  _-_- : (c : Config) → plugCη c ≡ e → Configη e

```

For a given expression $e : \text{Expr}$, any two terms of type $\text{Config}_\eta e$ are configurations of the same abstract machine during the tail-recursive fold over the expression e .

3.2 Up and down configurations

Next, we would like to formalize the left-to-right order on the configurations of our abstract machine. The **Stack** in the **Config** represents a path upwards, from the leaf to the root of the input expression. This is useful when navigating to neighbouring nodes, but makes it harder to compare the relative positions of two configurations. We now consider the value of **Config** corresponding to leaves with numbers 3 and 4 in our running example:




Figure 4. Comparison of configurations for leaves 3 and 4

The natural way to define the desired order relation is by induction over the **Stack**. However, there is a problem. The first element of both stacks does not provide us with sufficient information to decide which position is ‘smaller.’ The top of the stack only stores information about the location of the leaf with respect to its *parent* node. This kind of *local* information cannot be used to decide which one of the leaves is located in a position further to the right in the original input expression.

Instead, we would like to compare the *last* elements of both stacks. The common suffix of the stacks shows that both positions are in the left subtree of the root. Once these paths – read from right to left – diverge, we have found the exact node **Add** where one of the positions is in the left subtree and the other in the right.

When comparing two **Stacks**, we therefore want to consider them as paths *from the root*. Fortunately, this observation does not require us to change our definition of the **Stack** type; instead, we can define a variant of the **plug_η** function that interprets our contexts top-down rather than bottom-up:

```

plug⊥ : Expr → Stack+ → Expr
plug⊥ e Top       = e
plug⊥ e (Left t   stk) = Add (plug⊥ e stk) t
plug⊥ e (Right _ t _ stk) = Add t (plug⊥ e stk)
plugC⊥ : Config → Expr
plugC⊥ (n, stk) = plug⊥ (Val n) stk

```


We can convert freely between these two interpretations by reversing the stack. Furthermore, this conversion satisfies the plug_{\Downarrow} -to- plug_{\Uparrow} property, relating the two variants of plug :

```

444 convert : Config → Config
445 convert (n, s) = (n, reverse s)
446
447 plug⌊-to-plug⌈ : ∀ (c : Config)
448   → plugC⌊ c ≡ plugC⌈ (convert c)

```

As before, we can create a wrapper around Config that enforces that our Config denotes a leaf in the input expression e :

```

452 data Config⌊ (e : Expr) : Set where
453   _⌊_ : (c : Config) → plugC⌊ c ≡ e → Config⌊ e

```

As a corollary of the plug_{\Downarrow} -to- plug_{\Uparrow} property, we can define a pair of functions to switch between Config_{\Uparrow} and $\text{Config}_{\Downarrow}$:

```

456 Config⌊-to-Config⌈ : (e : Expr) → Config⌊ e → Config⌈ e
457 Config⌈-to-Config⌊ : (e : Expr) → Config⌈ e → Config⌊ e

```

3.3 Ordering configurations

Finally, we can define the ordering relation over values of type $\text{Config}_{\Downarrow}$. Even if the Config_{\Uparrow} is still used during execution of our tail-recursive evaluator, the $\text{Config}_{\Downarrow}$ type will be used to prove its termination.

The L_{\Downarrow} type defined below relates two configurations of type $\text{Config}_{\Downarrow} e$, that is, two states of the abstract machine evaluating the input expression e :

```

468 data L⌊ (e : Expr) → Config⌊ e → Config⌊ e → Set where
469   <-StepR : L r ⌋ ((t1, s1), ...) < ((t2, s2), ...)
470     → L Add l r ⌋ ((t1, Right l n eq s1), eq1) < ((t2, Right l n eq s2), eq2)
471   <-StepL : L l ⌋ ((t1, s1), ...) < ((t2, s2), ...)
472     → L Add l r ⌋ ((t1, Left r s1), eq1) < ((t2, Left r s2), eq2)
473   <-Base : (eq1 : Add e1 e2 ≡ Add e1 (plugC⌊ t1 s1))
474     → (eq2 : Add e1 e2 ≡ Add (plugC⌊ t2 s2) e2)
475     → L Add e1 e2 ⌋ ((t1, Right n e1 eq s1), eq1) < ((t2, Left e2 s2), eq2)

```

Despite the apparent complexity, the relation is straightforward. The constructors <-StepR and <-StepL cover the inductive cases, consuming the shared path from the root. When the paths diverge, the <-Base constructor states that the positions in the right subtree are ‘smaller than’ those in the left subtree.

Now we turn into showing that the relation is *well-founded*. We sketch the proof below:

```

484 <-WF : ∀ (e : Expr) → Well-founded (L e ⌋)
485 <-WF e x = acc (aux e x)
486 where
487   aux : ∀ (e : Expr) (x y : Config⌊ e)
488     → L e ⌋ y < x → Acc (L e ⌋) y
489   aux = ...

```

The proof follows the standard schema⁴ of most proofs of well-foundedness. It uses an auxiliary function, aux , that proves every configuration smaller than x is accessible.

⁴Most well-founded proofs in Agda standard library follow this pattern.

The proof proceeds initially by induction over our relation. The inductive cases, corresponding to the <-StepR and <-StepL constructors, recurse on the relation. In the base case, <-Base , we cannot recurse further on the relation. We then proceed by recursing over the original expression e ; without the type index, the subexpressions to the left e_1 and right e_2 are not syntactically related thus a recursive call is not possible. This step in the proof relies on only comparing configurations arising from traversing the same initial expression e .

3.4 A terminating and correct tail-recursive evaluator

We now have almost all the definitions in place to revise our tail-recursive fold, tail-rec-eval . However, we are missing one essential ingredient: we still need to show that the configuration decreases after a call to the unload^+ function.

Unfortunately, the function unload^+ and the relation that we have defined work on ‘different’ versions of the Stack : the relation compares stacks top-down; the unload^+ function manipulates stacks bottom-up. Furthermore, the function unload^+ as defined previously manipulates elements of the Config type directly, with no further type-level constraints relating these to the original input expression.

In the remainder of this section, we will reconcile these differences, complete the definition of our tail-recursive evaluator and finally prove its correctness.

Decreasing recursive calls To define our tail-recursive evaluator, we will begin by defining an auxiliary step function that performs a single step of computation. We will define the desired evaluator by iterating the step function, proving that it decreases in each iteration.

The step function calls unload^+ to produce a new configuration, if it exists. If the unload^+ function returns a natural number, $\text{inj}_2 v$, the entire input tree has been processed and the function terminates:

```

533 step : (e : Expr) → Config⌈ e → Config⌈ e ∪ ℕ
534 step e ((n, stk), eq)
535   with unload+ n (Val n) refl stk
536   ... | inj1 (n', stk') = inj1 ((n', stk'), ...)
537   ... | inj2 v         = inj2 v

```

We have omitted the second component of the result returned in the first branch, corresponding to a proof that $\text{plugC}_{\Uparrow} (n', \text{stk}') \equiv e$. The crucial lemma that we need to show to complete this proof, demonstrates that the unload^+ function respects our invariant:

```

544 unload+-plug⌈ :
545   ∀ (n : ℕ) (e : Expr) (eq : eval e ≡ x) (s : Stack+) (c : Config)
546   → unload+ n e eq s ≡ inj1 c
547   → ∀ (e' : Expr) → plug⌈ e s ≡ e' → plugC⌈ c ≡ e'

```

Finally, we can define the theorem stating that the step function always returns a smaller configuration:

```

551 step-< : ∀ (e : Expr) → (c c' : Configη e) → step e c ≡ inj1 c'
552       → ⊔ e ⊔ Configη-to-Config⊥ c' < Configη-to-Config⊥ c

```

Proving this statement directly is tedious, as there are many cases to cover and the expression e occurring in the types makes it difficult to identify and prove lemmas covering the individual cases. Therefore, we instead define another relation over non type-indexed configurations directly, and prove that there is an injection between both relations under suitable assumptions:

```

560 data _<_ : Config → Config → Set where
561   <-StepR : (t1 , s1) < (t2 , s2)
562         → (t1 , Right l n eq s1) < (t2 , Right l n eq s2)
563   <-StepL : (t1 , s1) < (t2 , s2)
564         → (t1 , Left r s1) < (t2 , Left r s2)
565   <-Base  : (e1 ≡ plugC⊥ t2 s2) → (e2 ≡ plugC⊥ t1 s1)
566         → (t1 , Right n e1 eq s1) < (t2 , Left e2 s2)
567   to : (e : Expr) (c1 c2 : Config)
568       → (eq1 : plugC⊥ c1 ≡ e) (eq2 : plugC⊥ c2 ≡ e)
569       → c1 < c2 → ⊔ e ⊔ (c1 , eq1) < (c2 , eq2)

```

Thus to complete the previous theorem, it is sufficient to show that the function `unload+` delivers a smaller `Config`:

```

572 unload+-< : ∀ (n : ℕ) (s : Stack+) (e : Expr) (s' : Stack+)
573           → unload+ n (Val n) refl s ≡ inj1 (t' , s)
574           → (t' , reverse s) < (n , reverse s)

```

The proof is done by induction over the stack supported; the complete proof requires some bookkeeping, covering around 200 lines of code, but is conceptually not complicated.

The function `tail-rec-eval` is now completed as follows⁵:

```

580 rec : (e : Expr) → (c : Configη e)
581     → Acc (⊔ e ⊔ ⊔) (Configη-to-Config⊥ c) → Configη e ⊔ ℕ
582 rec e c (acc rs) = with step e c | inspect (step e) c
583 ... | inj2 n | _ = inj2 n
584 ... | inj1 c' | [Is]
585     = rec e c' (rs (Configη-to-Config⊥ c) (step-< e c c' Is))
586 tail-rec-eval : Expr → ℕ
587 tail-rec-eval e with load e Top
588 ... | inj1 c = rec e (c , ...) (<-WF e c)

```

Agda's termination checker now accepts that the repeated calls to `rec` are on strictly smaller configurations.

3.5 Correctness

As we have indexed our configuration datatypes with the input expression, proving correctness of it is relatively straightforward. The type of the function `step` guarantees that the configuration returned points to a leaf in the input expression.

Proving the function `tail-rec-eval` correct amounts to show that the auxiliary function, `rec`, that is iterated until a value is produced, will behave the same as the original evaluator, `eval`. This is expressed by the following lemma, proven by induction over the accessibility predicate:

⁵`inspect` is an Agda idiom needed to remember that c' is the result of the call `step e c`.

```

606 rec-correct : ∀ (e : Expr) → (c : Configη e)
607             → (ac : Acc (⊔ e ⊔ ⊔) (Configη-to-Config⊥ c))
608             → eval e ≡ rec e c ac
609 rec-correct e c (acc rs)
610   with step e c | inspect (step e) c
611 ... | inj1 c' | [Is]
612     = rec-correct e c' (rs (Configη-to-Config⊥ c) (step-< e c c' Is))
613 ... | inj2 n | [Is] = step-correct n e eq c

```

At this point, we still need to prove the `step-correct` lemma that it is repeatedly applied. As the `step` function is defined as a wrapper around the `unload+` function, it suffices to prove the following property of `unload+`:

```

618 unload+-correct : ∀ (n : ℕ) (e : Expr) (eq : eval e ≡ n) (s : Stack+)
619                → ∀ (m : ℕ) → unload+ n e eq s ≡ inj2 m
620                → eval (plugη e s) ≡ m

```

This proof follows immediately by induction over s : `Stack+`.

The main correctness theorem now states that `eval` and `tail-rec-eval` are equal for all inputs:

```

624 correctness : ∀ (e : Expr) → eval e ≡ tail-rec-eval e
625 correctness e with load e Top
626 ... | inj1 c = rec-correct e (c , ...) (<-WF e c)
627 ... | inj2 _ = ⊥-elim ...

```

This finally completes the definition and verification of a tail-recursive evaluator.

4 A generic tail-recursive traversal

The previous section showed how to prove that our hand-written tail-recursive evaluation function was both terminating and equal to our original evaluator. In this section, we will show how we can generalize this construction to compute a tail-recursive equivalent of *any* function that can be written as a fold over a simple algebraic datatype. In particular, we generalize the following:

- The kind of datatypes, and their associated fold, that our tail-recursive evaluator supports, Section 4.1.
- The type of configurations of the abstract machine that computes the generic fold, Sections 4.2 and 4.3.
- The functions `load` and `unload` such that they work over our choice of generic representation, Section 4.4.
- The 'smaller than' relation to handle generic configurations, and its well-foundedness proof, Section 4.5.
- The tail-recursive evaluator, Section 4.6.
- The proof that the generic tail-recursive function is correct, Section 4.7.

Before we can define any such datatype generic constructions, however, we need to fix our universe of discourse.

4.1 The regular universe

In a dependently typed programming language such as Agda, we can represent a collection of types closed under certain operations as a *universe* [Altenkirch and McBride 2003; Martin-Löf 1984], that is, a data type U : `Set` describing the inhabitants of our universe together with its semantics,

661 $\text{el} : \mathbf{U} \rightarrow \text{Set}$, mapping each element of \mathbf{U} to its corre-
 662 sponding type. We have chosen the following universe of
 663 *regular* types [Morris et al. 2006; Noort et al. 2008]:

```
664 data Reg : Set1 where
665   0   : Reg
666   1   : Reg
667   I   : Reg
668   K   : (A : Set) → Reg
669   _⊕_ : (R Q : Reg) → Reg
670   _⊗_ : (R Q : Reg) → Reg
```

671 Types in this universe are formed from the empty type (0),
 672 unit type (1), and constant types (K A); the I constructor
 673 is used to refer to recursive subtrees. Finally, the universe
 674 is closed under both coproducts ($_⊕_$) and products ($_⊗_$).
 675 We could represent the *pattern* functor corresponding to the
 676 **Expr** type in this universe as follows:

```
677 exprF : Reg
678 exprF = K N ⊕ (I ⊗ I)
```

680 Note that as the constant functor K takes an arbitrary type A
 681 as its argument, the entire datatype lives in Set_1 . This could
 682 easily be remedied by stratifying this universe explicitly and
 683 parametrisng our development by a base universe.

684 We can interpret the inhabitants of **Reg** as a functor of
 685 type $\text{Set} \rightarrow \text{Set}$:

```
686 [ ] : Reg → Set → Set
687 [ 0 ] X = ⊥
688 [ 1 ] X = ⊤
689 [ I ] X = X
690 [ (K A) ] X = A
691 [ (R ⊕ Q) ] X = [ R ] X ⊔ [ Q ] X
692 [ (R ⊗ Q) ] X = [ R ] X × [ Q ] X
```

693 To show that this interpretation is indeed functorial, we
 694 define the following **fmap** operation:

```
695 fmap : (R : Reg) → (X → Y) → [ R ] X → [ R ] Y
696 fmap 0 f () = tt
697 fmap 1 f tt = tt
698 fmap I f x = f x
699 fmap (K A) f x = x
700 fmap (R ⊕ Q) f (inj1 x) = inj1 (fmap R f x)
701 fmap (R ⊕ Q) f (inj2 y) = inj2 (fmap Q f y)
702 fmap (R ⊗ Q) f (x, y) = fmap R f x, fmap Q f y
```

703 Finally, we can tie the recursive knot, taking the least fixpoint
 704 of the functor associated with the elements of our universe:

```
705 data μ (R : Reg) : Set where
706   In : [ R ] (μ R) → μ R
```

708 Next, we can define a *generic* fold, or *catamorphism*, to work
 709 on the inhabitants of the regular universe. For each code
 710 $R : \text{Reg}$, the **cata** R function takes an *algebra* of type
 711 $[R] X \rightarrow X$ as argument. This algebra assigns semantics
 712 to the ‘constructors’ of R. Folding over a tree of type μR
 713 corresponds to recursively folding over each subtree and
 714 assembling the results using the argument algebra:

```
716 cata : (R : Reg) → ([ R ] X → X) → μ R → X
717 cata R ψ (In r) = ψ (fmap R (cata R ψ) r)
```

718 Unfortunately, Agda’s termination checker does not accept
 719 this definition. The problem, once again, is that the recursive
 720 calls to **cata** are not made to structurally smaller trees, but
 721 rather **cata** is passed as an argument to the higher-order
 722 function **fmap**.

723 To address this, we fuse the **fmap** and **cata** functions into
 724 a single **map-fold** function [Norell 2008]:

```
725 map-fold : (R Q : Reg) → ([ Q ] X → X) → [ R ] (μ Q) → [ R ] X
726 map-fold 0 Q ψ () = tt
727 map-fold 1 Q ψ tt = tt
728 map-fold I Q ψ (In x) = ψ (map-fold Q Q ψ x)
729 map-fold (K A) Q ψ x = x
730 map-fold (R ⊕ Q) P ψ (inj1 x) = inj1 (map-fold R P ψ x)
731 map-fold (R ⊕ Q) P ψ (inj2 y) = inj2 (map-fold Q P ψ y)
732 map-fold (R ⊗ Q) P ψ (x, y) = map-fold R P ψ x, map-fold Q P ψ y
```

733 We can now define **cata** in terms of **map-fold** as follows:

```
734 cata : (R : Reg) ([ R ] X → X) → μ R → X
735 cata R ψ (In r) = map-fold R R ψ r
```

736 This definition is indeed accepted by Agda’s termination
 737 checker.

738 **Example** We can now revisit our example evaluator from
 739 the introduction. To define the evaluator using the generic
 740 **cata** function, we instantiate the catamorphism to work on
 741 the expressions and pass the desired algebra:

```
742 eval : μ exprF → N
743 eval = cata exprF φ
744   where φ : [ exprF ] N → N
745         φ (inj1 n) = n
746         φ (inj2 (n, n')) = n + n'
```

747 In the remainder of this paper, we will develop an alter-
 748 native traversal that maps any algebra to a tail-recursive
 749 function that is guaranteed to terminate and produce the
 750 same result as the corresponding call to **cata**.

751 4.2 Dissection

752 As we mentioned in the previous section, the configurations
 753 of our abstract machine from the introduction are instances
 754 of McBride’s dissections [2008]. We briefly recap this con-
 755 struction, showing how to calculate the type of abstract
 756 machine configurations for any type in our universe. The
 757 key definition, ∇ , computes a bifunctor for each element of
 758 our universe:

```
759 ∇ : (R : Reg) → (Set → Set → Set)
760 ∇ 0 XY = ⊥
761 ∇ 1 XY = ⊥
762 ∇ I XY = ⊤
763 ∇ (K A) XY = ⊥
764 ∇ (R ⊕ Q) XY = ∇ R X Y ⊔ ∇ Q X Y
765 ∇ (R ⊗ Q) XY = (∇ R X Y × [ Q ] Y)
766   ⊔ ([ R ] X × ∇ Q X Y)
```

This operation generalizes the zippers, by defining a bifunctor $\nabla R X Y$. You may find it useful to think of the special case, $\nabla R X (\mu R)$ as a configuration of an abstract machine traversing a tree of type μR to produce a result of type X . The last clause of the definition of ∇ is of particular interest: to *dissect* a product, we either *dissect* the left component pairing it with the second component interpreted over the second variable Y ; or we *dissect* the second component and pair it with the first interpreted over X .

A *dissection* is formally defined as the pair of the one-hole context and the missing value that can fill the context.

$$\begin{aligned} \mathcal{D} &: (R : \text{Reg}) \rightarrow (X Y : \text{Set}) \rightarrow \text{Set} \\ \mathcal{D} R X Y &= \nabla R X Y \times Y \end{aligned}$$

We can reconstruct Huet's zipper for generic trees of type μR by instantiating both X and Y to μR .

Given a *dissection*, we can define a *plug* operation that assembles the context and current value in focus to produce a value of type $\llbracket R \rrbracket Y$:

$$\begin{aligned} \text{plug} &: (R : \text{Reg}) \rightarrow (X \rightarrow Y) \rightarrow \mathcal{D} R X Y \rightarrow \llbracket R \rrbracket Y \\ \text{plug } \mathbb{0} &\quad \eta (\cdot), x \\ \text{plug } \mathbb{1} &\quad \eta (\cdot), x \\ \text{plug } \mathbb{I} &\quad \eta (tt), x = x \\ \text{plug } (K A) &\quad \eta (\cdot), x \\ \text{plug } (R \oplus Q) \eta (\text{inj}_1 r), x &= \text{inj}_1 (\text{plug } R \eta (r), x) \\ \text{plug } (R \oplus Q) \eta (\text{inj}_2 q), x &= \text{inj}_2 (\text{plug } Q \eta (q), x) \\ \text{plug } (R \otimes Q) \eta (\text{inj}_1 (dr), q), x &= (\text{plug } R \eta (dr), x), q \\ \text{plug } (R \otimes Q) \eta (\text{inj}_2 (r), dq), x &= (\text{fmap } R \eta r, \text{plug } Q \eta (dq), x) \end{aligned}$$

In the last clause of the definition, the *dissection* is over the right component of the pair leaving a value $r : \llbracket R \rrbracket X$ to the left. In that case, it is only possible to reconstruct a value of type $\llbracket R \rrbracket Y$, if we have a function η to recover Y s from X s.

In order to ease things later, we bundle a *dissection* together with the functor to which it *plugs* as a type-indexed type.

$$\begin{aligned} \text{data } \mathcal{D}_x (R : \text{Reg}) (X Y : \text{Set}) (\eta : X \rightarrow Y) (t_x : \llbracket R \rrbracket Y) : \text{Set} \text{ where} \\ _ _ : (d : \mathcal{D} R X Y) \rightarrow \text{plug } R \eta d \equiv t_x \rightarrow \mathcal{D}_x R X Y \eta t_x \end{aligned}$$

4.3 Generic configurations

While the *dissection* computes the bifunctor *underlying* our configurations, we still need to take a fixpoint of this bifunctor. Each configuration consists of a list of *dissections* and the current subtree in focus. To the left of the current subtree in focus, we store the partial results arising from the subtrees that we have already processed; on the right, we store the subtrees that still need to be visited.

As we did for the Stack^+ datatype from the introduction, we also choose to store the original subtrees that have been visited and their corresponding correctness proofs:

$$\begin{aligned} \text{record } \text{Computed} (R : \text{Reg}) (X : \text{Set}) (\psi : \llbracket R \rrbracket X \rightarrow X) : \text{Set} \text{ where} \\ \text{constructor } _ _ _ \\ \text{field} \\ \text{Tree} : \mu R \end{aligned}$$

$$\begin{aligned} \text{Value} &: X & 826 \\ \text{Proof} &: \text{cata } R \psi \text{ Tree} \equiv \text{Value} & 827 \end{aligned}$$

$$\begin{aligned} \text{Stack}^G &: (R : \text{Reg}) \rightarrow (X : \text{Set}) \rightarrow (\psi : \llbracket R \rrbracket X \rightarrow X) \rightarrow \text{Set} & 829 \\ \text{Stack}^G R X \psi &= \text{List } (\nabla R (\text{Computed } R X \psi) (\mu R)) & 830 \end{aligned}$$

A *stack* is a list of *dissections*. To the left we have the *Computed* results; to the right, we have the subtrees of type μR . Note that the Stack^G datatype is parametrised by the algebra ψ , as the *Proof* field of the *Computed* record refers to it.

As we saw in Section 3.5, we can define two different *plug* operations on these stacks:

$$\begin{aligned} \text{plug-}\mu_{\downarrow} &: (R : \text{Reg}) \rightarrow \{\psi : \llbracket R \rrbracket X \rightarrow X\} \\ &\quad \rightarrow \mu R \rightarrow \text{Stack}^G R X \psi \rightarrow \mu R & 838 \\ \text{plug-}\mu_{\downarrow} R t [] &= t & 839 \\ \text{plug-}\mu_{\downarrow} R t (h : hs) &= \text{In } (\text{plug } R \text{ Computed.Tree } h (\text{plug-}\mu_{\downarrow} R t hs)) & 840 \\ \text{plug-}\mu_{\uparrow} &: (R : \text{Reg}) \rightarrow \{\psi : \llbracket R \rrbracket X \rightarrow X\} \\ &\quad \rightarrow \mu R \rightarrow \text{Stack}^G R X \psi \rightarrow \mu R & 841 \\ \text{plug-}\mu_{\uparrow} R t [] &= t & 842 \\ \text{plug-}\mu_{\uparrow} R t (h : hs) &= \text{plug-}\mu_{\uparrow} R (\text{In } (\text{plug } R \text{ Computed.Tree } h t)) hs & 843 \end{aligned}$$

Both functions use the projection, *Computed.Tree*, as an argument to *plug* to extract the subtrees that have already been processed.

To define the configurations of our abstract machine, we are interested in *any* path through our initial input, but want to restrict ourselves to those paths that lead to a leaf. But what constitutes a leaf in this generic setting?

To describe leaves, we introduce the following predicate *NonRec*, stating when a tree of type $\llbracket R \rrbracket X$ does not refer to the variable X , that will be used to represent recursive subtrees:

$$\begin{aligned} \text{data } \text{NonRec} &: (R : \text{Reg}) \rightarrow \llbracket R \rrbracket X \rightarrow \text{Set} \text{ where} & 854 \\ \text{NonRec-}\mathbb{1} &: \text{NonRec } \mathbb{1} \text{ } tt & 855 \\ \text{NonRec-K} &: (B : \text{Set}) \rightarrow (b : B) \rightarrow \text{NonRec } (K B) b & 856 \\ \text{NonRec-}\oplus_1 &: (R Q : \text{Reg}) \rightarrow (r : \llbracket R \rrbracket X) \\ &\quad \rightarrow \text{NonRec } R r \rightarrow \text{NonRec } (R \oplus Q) (\text{inj}_1 r) & 857 \\ \text{NonRec-}\oplus_2 &: (R Q : \text{Reg}) \rightarrow (q : \llbracket Q \rrbracket X) \\ &\quad \rightarrow \text{NonRec } Q q \rightarrow \text{NonRec } (R \oplus Q) (\text{inj}_2 q) & 858 \\ \text{NonRec-}\otimes &: (R Q : \text{Reg}) \rightarrow (r : \llbracket R \rrbracket X) \rightarrow (q : \llbracket Q \rrbracket X) \\ &\quad \rightarrow \text{NonRec } R r \rightarrow \text{NonRec } Q q \rightarrow \text{NonRec } (R \otimes Q) (r, q) & 859 \end{aligned}$$

As an example, in the pattern functor for the *Expr* type, $K \mathbb{N} \oplus (\mathbb{I} \otimes \mathbb{I})$, terms built using the left injection are non-recursive:

$$\begin{aligned} \text{Val-NonRec} &: \forall (n : \mathbb{N}) \rightarrow \text{NonRec } (K \mathbb{N} \oplus (\mathbb{I} \otimes \mathbb{I})) (\text{inj}_1 n) & 860 \\ \text{Val-NonRec} &: n = \text{NonRec-}\oplus_1 (K \mathbb{N}) (\mathbb{I} \otimes \mathbb{I}) n (\text{NonRec-K } \mathbb{N} n) & 861 \end{aligned}$$

This corresponds to the idea that the constructor *Val* is a leaf in a tree of type *Expr*.

On the other hand, we cannot prove the predicate *NonRec* for terms using the right injection. The occurrences of recursive positions disallow us from framing the proof (The type *NonRec* does not have a constructor such as *NonRec-I* : $(x : X) \rightarrow \text{NonRec } \mathbb{I} x$).

This example also shows how ‘generic’ leaves can be recursive. As long as the recursion only happens in the functor layer (code \oplus) and not in the fixpoint level (code \mathbb{I}).

Crucially, any non-recursive subtree is independent of X – as is exhibited by the following coercion function:

```
coerce : (R : Reg) → (x : [R] X) → NonRec R x → [R] Y
```

Whose definition is not worth including as it follows directly by induction over the predicate.

We can now define the notion of leaf generically, as a substructure without recursive subtrees:

```
Leaf : Reg → Set → Set
Leaf R X = Σ ([R] X) (NonRec R)
```

Just as we saw previously, a configuration is now given by the current leaf in focus and the stack, given by a dissection, storing partial results and unprocessed subtrees:

```
ConfigG : (R : Reg) → (X : Set) → (ψ : [R] X → X) → Set
ConfigG R X ψ = Leaf R X × StackG R X ψ
```

Finally, we can recompute the original tree using a `plug` function as before:

```
plugC-μll : (R : Reg) {ψ : [R] X → X}
  → ConfigG R X ψ → μ R → Set
plugC-μll R ((l, isl), s) t = plug-μll R (ln (coerce l isl)) s t
```

Note that the `coerce` function is used to embed a leaf into a larger tree. A similar function can be defined for the ‘bottom-up’ zippers, that work on a reversed stack.

4.4 One step of a catamorphism

In order to write a tail-recursive catamorphism, we start by defining the generic operations that correspond to the functions `load` and `unload` given in the introduction (Section 2).

Load The function `loadG` traverses the input term to find its leftmost leaf. Any other subtrees the `loadG` function encounters are stored on the stack. Once the `loadG` function encounters a constructor without subtrees, it is has found the desired leaf.

We write `loadG` by appealing to an ancillary definition `first-cps`, that uses continuation-passing style to keep the definition tail-recursive and obviously structurally recursive. If we were to try to define `loadG` by recursion directly, we would need to find the leftmost subtree and recurse on it – but this subtree may not be obviously syntactically smaller.

The type of our `first-cps` function is daunting at first:

```
first-cps : (R Q : Reg) {ψ : [Q] X → X}
  → [R] (μ Q)
  → (∇ R (Computed Q X ψ) (μ Q) → (∇ Q (Computed Q X ψ) (μ Q)))
  → (Leaf R X → StackG Q X ψ → ConfigG Q X ψ ⊔ X)
  → StackG Q X ψ
  → ConfigG Q X ψ ⊔ X
```

The first two arguments are codes of type `Reg`. The code `Q` represents the datatype for which we are defining a traversal; the code `R` is the code on which we pattern match. In the

initial call to `first-cps` these two codes will be equal. As we define our function, we pattern match on `R`, recursing over the codes in (nested) pairs or sums – yet we still want to remember the original code for our data type, `Q`.

The next argument of type `[R] (μ Q)` is the data we aim to traverse. Note that the ‘outermost’ layer is of type `R`, but the recursive subtrees are of type `μ Q`. The next two arguments are two continuations: the first is used to gradually build the *dissection* of `R`; the second continues on another branch once one of the leaves have been reached. The last argument of type `StackG Q X ψ` is the current stack. The entire function will compute the initial configuration of our machine of type `ConfigG Q X ψ`⁶:

```
loadG : (R : Reg) {ψ : [R] X → X} → μ R
  → StackG R X ψ → ConfigG R X ψ ⊔ X
loadG R (ln t) s = first-cps R R t id (λ l → inj1 ∘ _.) s
```

We shall fill the definition of `first-cps` by cases. The clauses for the base cases are as expected. In $\mathbb{0}$ there is nothing to be done. The $\mathbb{1}$ and `K A` codes consist of applying the second continuation to the tree and the *stack*.

```
first-cps 0 Q () _
first-cps 1 Q x k f s = f(tt, NonRec-1) s
first-cps (K A) Q x k f s = f(x, NonRec-K A x) s
```

The recursive case, constructor `l`, corresponds to the occurrence of a subtree. The function `first-cps` is recursively called over that subtree with the stack incremented by a new element that corresponds to the *dissection* of the functor layer up to that point. The second continuation is replaced with the initial one.

```
first-cps l Q (ln x) k f s = first-cps Q Q x id (λ c → inj1 ∘ _.) (k tt = s)
```

In the coproduct, both cases are similar, just having to account for the use of different constructors in the continuations.

```
first-cps (R ⊕ Q) P (inj1 x) k f s = first-cps R P x (k ∘ inj1) cont s
  where cont (l, isl) = f((inj1 l), NonRec-⊕1 R Q l isl)
first-cps (R ⊕ Q) P (inj2 y) k f s = first-cps Q P y (k ∘ inj2) cont s
  where cont (l, isl) = f((inj1 l), NonRec-⊕2 R Q l isl)
```

The interesting clause is the one that deals with the product. First the function `first-cps` is recursively called on the left component of the pair trying to find a subtree to recurse over. However, it may be the case that there are no subtrees at all, thus it is passed as the first continuation a call to `first-cps` over the right component of the product. In case the continuation fails to find a subtree, it returns the leaf as it is.

```
first-cps (R ⊗ Q) P (r, q) k f s = first-cps R P r (k ∘ inj1 ∘ (, q)) cont s
  where cont (l, isl) = first-cps Q P q (k ∘ inj2 ∘ _.) (coerce l isl) cont'
  where cont' (l', isl') = f(l, l') (NonRec-⊗ R Q l l' isl isl')
```

⁶As in the introduction, we use a sum type \uplus to align its type with that of `unloadG`.

Unload Armed with load^G we turn our attention to unload^G . First of all, it is necessary to define an auxiliary function, **right**, that given a *dissection* and a value (of the type of the left variables), either finds a dissection $\mathcal{D} R X Y$ or it shows that there are no occurrences of the variable left. In the latter case, it returns the functor interpreted over $Y, \llbracket R \rrbracket Y$.

$\text{right} : (R : \text{Reg}) \rightarrow \nabla R X Y \rightarrow X \rightarrow \llbracket R \rrbracket X \uplus \mathcal{D} R X Y$

Its definition is simply by induction over the code R , with the special case of the product that needs to use another ancillary definition to look for the leftmost occurrence of the variable position within $\llbracket R \rrbracket X$.

The function unload^G is defined by induction over the *stack*. If the *stack* is empty the job is done and a final value is returned. In case the *stack* has at least one *dissection* in its head, the function **right** is called to check whether there are any more holes left. If there are none, a recursive call to unload^G is dispatched, otherwise, if there is still a subtree to be processed the function load^G is called.

```

unloadG : (R : Reg)
  → (ψ : ⌊ R ⌋ X → X)
  → (t : μ R) → (x : X) → cata R ψ t ≡ x
  → StackG R X ψ
  → ConfigG R X ψ ⊔ X

unloadG R ψ t x eq [] = inj2 x
unloadG R ψ t x eq (h : hs) with right R h (t, x, eq)
unloadG R ψ t x eq (h : hs) | inj1 r with compute R R r
... | (rx, rr), eq' = unloadG R ψ (ln rp) (ψ rx) (cong ψ eq') hs
unloadG R ψ t x eq (h : hs) | inj2 (dr, q) = loadG R q (dr : hs)

```

When the function **right** returns an inj_1 it means that there are not any subtrees left in the *dissection*. If we take a closer look, the type of the r in $\text{inj}_1 r$ is $\llbracket R \rrbracket (\text{Computed } R X \psi)$. The functor $\llbracket R \rrbracket$ is storing at its variable positions both values, subtrees and proofs.

However, what is needed for the recursive call is: first, the functor interpreted over values, $\llbracket R \rrbracket X$, in order to apply the algebra; second, the functor interpreted over subtrees, $\llbracket R \rrbracket (\mu R)$, to keep the original subtree where the value came from; Third, the proof that the value equals to applying a **cata** over the subtree. The function **compute** massages r to adapt the arguments for the recursive call to unload^G .

4.5 Relation over generic configurations

We can engineer a *well-founded* relation over elements of type $\text{Config}_\eta^G t$, for some concrete tree $t : \mu R$, by explicitly separating the functorial layer from the recursive layer induced by the fixed point. At the functor level, we impose the order over *dissections* of R , while at the fixed point level we define the order by induction over the *stacks*.

To reduce clutter in the definition, we give a non type-indexed relation over terms of type Config_η^G . We can later use the same technique as in Section 3.4 to recover a fully type-indexed relation over elements of type $\text{Config}_\eta^G t$ by requiring that the *zippers* respect the invariant, $\text{plugC-}\mu_\parallel c \equiv t$.

The relation is defined by induction over the Stack^G part of the *zippers* as follows.

```

data _<C_ : ConfigG R X ψ → ConfigG R X ψ → Set where
  Step : (t1, s1) <C (t2, s2) → (t1, h : s1) <C (t2, h : s2)
  Base : plugC-μ∥ R (t1, s1) ≡ e1 → plugC-μ∥ R (t2, s2) ≡ e1
        → (h1, e1) <∇ (h2, e2) → (t1, h1 : s1) <C (t2, h2 : s2)

```

This relation has two constructors:

- The **Step** constructor covers the inductive case. When the head of both *stacks* is the same, i.e., both Config_η^G s share the same prefix, it recurses directly on tail of both *stacks*.
- The constructor **Base** accounts for the case when the head of the *stacks* is different. This means that the paths given by the configuration denotes different subtrees of the same node. In that case, the relation we are defining relies on an auxiliary relation $\llcorner_ _ _ <\nabla_ _$ that orders *dissections* of type $\mathcal{D} R (\text{Computed } R X \psi) (\mu R)$.

We can define this relation on dissections directly, without having to consider the recursive nature of our datatypes. We define the required relation over dissections interpreted on any sets X and Y as follows:

```

data _\llcorner\_ \_ \_ <\nabla\_ \_ : (R : Reg) → D R X Y → D R X Y → Set where
  step-⊕1 : \llcorner R \llcorner (r, t1) <\nabla (r', t2)
            → \llcorner R ⊕ Q \llcorner (inj1 r, t1) <\nabla (inj1 r', t2)
  step-⊕2 : \llcorner Q \llcorner (q, t2) <\nabla (q', t2)
            → \llcorner R ⊕ Q \llcorner (inj2 q, t1) <\nabla (inj2 q', t2)
  step-⊗1 : \llcorner R \llcorner (dr, t1) <\nabla (dr', t2)
            → \llcorner R ⊗ Q \llcorner (inj1 (dr, q), t1) <\nabla (inj1 (dr', q), t2)
  step-⊗2 : \llcorner Q \llcorner (dq, t1) <\nabla (dq', t2)
            → \llcorner R ⊗ Q \llcorner (inj2 (r, dq), t1) <\nabla (inj2 (r, dq'), t2)
  base-⊗ : \llcorner R ⊗ Q \llcorner (inj2 (r, dq), t1) <\nabla (inj1 (dr, q), t2)

```

The idea is that we order the elements of a *dissection* in a left-to-right fashion. All the constructors except for **base-⊗** simply follow the structure of the dissection. To define the base case, **base-⊗**, recall that the *dissection* of the product of two functors, $R \otimes Q$, has two possible values. It is either a term of type $\nabla R X Y \times \llbracket Q \rrbracket Y$, such as $\text{inj}_1 (dr, q)$ or a term of type $\llbracket R \rrbracket X \times \nabla Q X Y$ like $\text{inj}_2 (r, dq)$. The former denotes a position in the left component of the pair while the latter denotes a position in the right component. The **base-⊗** constructor states that positions in right are smaller than those in the left.

This completes the order relation on configurations; we still need to prove our relation is *well-founded*. To prove this, we write a type-indexed version of each relation. The first relation, $_<C_$, has to be type-indexed by the tree of type μR to which both *zipper* recursively plug through $\text{plugC-}\mu_\parallel$. The auxiliary relation, $\llcorner_ _ _ <\nabla_ _$, needs to be type-indexed by the functor of type $\llbracket R \rrbracket X$ to which both *dissections* **plug**:

```

data _\llcorner\_ \_ \_ <\nabla\_ \_ {X Y : Set} {η : X → Y} : (R : Reg) → (tx : ⌊ R ⌋ Y)
  → Dx R X Y η tx → Dx R X Y η tx → Set where

```

1101 data $\llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} \{X : \text{Set}\} (R : \text{Reg}) \{ \psi : \llbracket R \rrbracket X \rightarrow X \} : (t : \mu R)$
 1102 $\rightarrow \text{Config}_{\llcorner}^G R X \psi t \rightarrow \text{Config}_{\llcorner}^G R X \psi t \rightarrow \text{Set}$ where

1103 The proof of *well-foundedness* of $\llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner}$ is a straight-
 1104 forward generalization of proof given for the example in
 1105 Section 3.3. The full proof of the following statement can
 1106 found in the accompanying code:

1107 $\llcorner_{\llcorner} \text{-WF} : (R : \text{Reg}) \rightarrow (t : \mu R) \rightarrow \text{Well-founded} (\llcorner R \llcorner \llcorner \llcorner \llcorner \llcorner_{\llcorner} \llcorner_{\llcorner})$
 1108

1109 4.6 A generic tail-recursive machine

1110 We are now ready to define a generic tail-recursive machine.
 1111 To do so we now assemble the generic machinery we have
 1112 defined so far. We follow the same outline as in Section 3.4.

1113 The first point is to build a wrapper around the function
 1114 unload^G that performs one step of the *catamorphism*. The
 1115 function step^G statically enforces that the input tree remains
 1116 the same both in its argument and in its result.

1117 $\text{step}^G : (R : \text{Reg}) \rightarrow (\psi : \llbracket R \rrbracket X \rightarrow X) \rightarrow (t : \mu R)$
 1118 $\rightarrow \text{Config}_{\llcorner}^G R X \psi t \rightarrow \text{Config}_{\llcorner}^G R X \psi t \uplus X$

1119 We omit the full definition. The function step^G performs
 1120 a call to unload^G , coercing the leaf of type $\llbracket R \rrbracket X$ in the
 1121 $\text{Config}_{\llcorner}^G$ argument to a generic tree of type $\llbracket R \rrbracket (\mu R)$.

1122 We show that unload^G preserves the invariant, by proving
 1123 the following lemma:

1124 $\text{unload-plug}_{\llcorner}^G : \forall (R : \text{Reg}) \{ \psi : \llbracket R \rrbracket X \rightarrow X \}$
 1125 $\rightarrow (t : \mu R) (x : X) (eq : \text{cata } R \psi t \equiv x) (s : \text{Stack}^G R X \psi)$
 1126 $\rightarrow (c : \text{Config}^G R X \psi)$
 1127 $\rightarrow \forall (e : \mu R) \rightarrow \text{plug-}\mu_{\llcorner} R t s \equiv e$
 1128 $\rightarrow \text{unload}^G R \psi t x \text{ eq } s \equiv \text{inj}_1 c \rightarrow \text{plug-}\mu_{\llcorner} R c \equiv e$

1129 Next, we show that applying the function step^G to a con-
 1130 figuration of the abstract machine produces a smaller config-
 1131 uration. As the function step^G is a wrapper over the unload^G
 1132 function, we only have to prove that the property holds for
 1133 unload^G .

1134 The unload^G function does two things. First, it calls the
 1135 function right to check whether the *dissection* has any more
 1136 recursive subtrees to the right that still have to be processed.
 1137 It then dispatches to either load^G , if there is, or recurses
 1138 otherwise. When there is a hole left, a new *dissection* is re-
 1139 turned by right . Thus showing that the new configuration
 1140 is smaller amounts to show that the *dissection* returned by
 1141 right is smaller by $\llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner}$. This amounts to proving the
 1142 following lemma:

1143 $\text{right-}\llcorner : \text{right } R \text{ dr } (t, y, eq) \equiv \text{inj}_2 (dr', t)$
 1144 $\rightarrow \llcorner R \llcorner ((dr', t)) \llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} \llcorner_{\llcorner} (dr, t)$

1145 We have simplified the type signature, leaving out the uni-
 1146 versally quantified variables and their types.

1147 Extending this result to show that the function unload^G
 1148 delivers a smaller value is straightforward. By induction over
 1149 the input stack we check if the traversal is done or not. If it
 1150 is not yet done, there is at least one dissection in the top of
 1151 the stack. The function right applied to that element returns
 1152 either a smaller dissection or a tree with all values on the

1153 leaves. If we obtain a new dissection, we use the $\text{right-}\llcorner$
 1154 lemma; in the latter case, we continue by induction over the
 1155 stack. In this fashion, we can prove the following statement
 1156 that our traversal decreases:

1157 $\text{step}^G \llcorner_{\llcorner} : (R : \text{Reg}) (\psi : \llbracket R \rrbracket X \rightarrow X) \rightarrow (t : \mu R)$
 1158 $\rightarrow (c_1 c_2 : \text{Config}_{\llcorner}^G R X \psi t)$
 1159 $\rightarrow \text{step}^G R \psi t c_1 \equiv \text{inj}_1 c_2 \rightarrow \llcorner R \llcorner \llcorner t \llcorner c_2 \llcorner_{\llcorner} \llcorner_{\llcorner} c_1$

1160 Finally, we can write the *tail-recursive machine*, tail-rec-cata ,
 1161 as the combination of an auxiliary recursor over the acces-
 1162 sibility predicate and a top-level function that initiates the
 1163 computation with suitable arguments:

1164 $\text{rec} : (R : \text{Reg}) (\psi : \llbracket R \rrbracket X \rightarrow X) (t : \mu R)$
 1165 $\rightarrow (c : \text{Config}_{\llcorner}^G R X \psi t)$
 1166 $\rightarrow \text{Acc} (\llcorner R \llcorner \llcorner t \llcorner_{\llcorner} \llcorner_{\llcorner}) (\text{Config}_{\llcorner}^G \text{-to-Config}_{\llcorner}^G c) \rightarrow X$
 1167 $\text{rec } R \psi t c (\text{acc } rs)$ with $\text{step}^G R \psi t c \mid \text{inspect} (\text{step}^G R \psi t) c$
 1168 $\dots \mid \text{inj}_1 z' \mid [Is] = \text{rec } R \psi t z' (rs z' (\text{step}^G \llcorner_{\llcorner} R \psi t c z' Is))$
 1169 $\dots \mid \text{inj}_2 x \mid [-] = x$
 1170 $\text{tail-rec-cata} : (R : \text{Reg}) \rightarrow (\psi : \llbracket R \rrbracket X \rightarrow X) \rightarrow \mu R \rightarrow X$
 1171 $\text{tail-rec-cata } R \psi x$ with $\text{load}^G R \psi x []$
 1172 $\dots \mid \text{inj}_1 c = \text{rec } R \psi (c, \dots) (\llcorner_{\llcorner} \text{-WF } R c)$

1173 4.7 Correctness, generically

1174 To prove our tail-recursive evaluator produces the same out-
 1175 put as the catamorphism is straight-forward. As we did in
 1176 the tail-rec-eval example (Section 3.5), we perform induction
 1177 over the accessibility predicate in the auxiliary recursor. In
 1178 the base case, when the function step^G returns a ground
 1179 value of type X , we have to show that such value is the result
 1180 of applying the *catamorphism* to the input. Recall that step^G
 1181 is a wrapper around unload^G , hence it suffices to prove the
 1182 following lemma:

1183 $\text{unload}^G \text{-correct} : \forall (R : \text{Reg}) (\psi : \llbracket R \rrbracket X \rightarrow X)$
 1184 $(t : \mu R) (x : X) (eq : \text{cata } R \psi t \equiv x)$
 1185 $(s : \text{Stack}^G R X \psi) (y : X)$
 1186 $\rightarrow \text{unload}^G R \psi t x \text{ eq } s \equiv \text{inj}_2 y$
 1187 $\rightarrow \forall (e : \mu R) \rightarrow \text{plug-}\mu_{\llcorner} R t s \equiv e \rightarrow \text{cata } R \psi e \equiv y$

1188 Our generic correctness result is an immediate consequence:

1189 $\text{correctness}^G : \forall (R : \text{Reg}) (\psi : \llbracket R \rrbracket X \rightarrow X) (t : \mu R)$
 1190 $\rightarrow \text{cata } R \psi t \equiv \text{tail-rec-cata } R \psi t$

1191 4.8 Example

1192 To conclude, we show how to generically implement the ex-
 1193 ample from the introduction (Section 1), and how the generic
 1194 construction gives us a *correct* tail-recursive machine for free.
 1195 First, we recap the *pattern* functor underlying the type Expr :

1196 $\text{exprF} : \text{Reg}$
 1197 $\text{exprF} = \mathbb{K} \mathbb{N} \oplus (I \otimes I)$

1198 The Expr type is then isomorphic to tying the knot over
 1199 exprF :

1200 $\text{Expr}^G : \text{Set}$
 1201 $\text{Expr}^G = \mu \text{exprF}$

The function `eval` is equivalent to instantiating the *catamorphism* with an appropriate algebra:

$$\begin{aligned} \psi &: \text{exprF } \mathbb{N} \rightarrow \mathbb{N} \\ \psi (\text{inj}_1 n) &= n \\ \psi (\text{inj}_2 (e_1, e_2)) &= e_1 + e_2 \\ \text{eval} &: \text{Expr}^G \rightarrow \mathbb{N} \\ \text{eval} &= \text{cata exprF } \psi \end{aligned}$$

Finally, a tail-recursive machine *equivalent* to the one we derived in Section 3.4, `tail-rec-eval`, is given by:

$$\begin{aligned} \text{tail-rec-eval}^G &: \text{Expr}^G \rightarrow \mathbb{N} \\ \text{tail-rec-eval}^G &= \text{tail-rec-cata exprF } \psi \end{aligned}$$

5 Discussion

There is a long tradition of calculating abstract machines from an evaluator, dating back as far as early work on the abstract machines for the evaluation of lambda calculus terms [Landin 1964]. In particular, Danvy [Ager et al. 2003; Danvy 2009] has published many examples showing how abstract machines arise from defunctionalizing an interpreter written in continuation-passing style. This work in turn, inspired McBride’s work on dissections [2008], that defines the key constructions on which this paper builds. McBride’s work, however, does not give a proof of termination or correctness.

The universe of regular types used in this paper is somewhat restricted: we cannot represent mutually recursive types [Yakushev et al. 2009], nested data types [Bird and Meertens 1998], indexed families [Dybjer 1994], or inductive-recursive types [Dybjer and Setzer 1999]. Fortunately, there is a long tradition of generic programming with universes in Agda, arguably dating back to Martin-Löf [1984]. It would be worthwhile exploring how to extend our construction to more general universes, such as the context-free types [Altenkirch et al. 2007], containers [Abbott et al. 2005; Altenkirch et al. 2015], or the ‘sigma-of-sigma’ universe [Chapman et al. 2010; Oury and Swierstra 2008]. Doing so would allow us to exploit dependent types further in the definition of our evaluators. For example, we might then define an interpreter for the well-typed lambda terms and derive a tail recursive evaluator automatically, rather than verifying the construction by hand [Swierstra 2012].

The termination proof we have given defines a well-founded relation and shows that this decreases during execution. There are other techniques for writing functions that are not obviously structurally recursive, such as the Bove-Capretta method [Bove and Capretta 2005], partiality monad [Danielsson 2012], or coinductive traces [Nakata and Uustalu 2009]. In contrast to the well-founded recursion used in this paper, however, these methods do not yield an evaluator that is directly executable, but instead defer the termination proof. Given that we can – and indeed have – shown termination of our tail-recursive abstract machines, the abstract machines are executable directly in Agda.

One drawback of our construction is that the stacks now not only store the value of evaluating previously visited subtrees, but also records the subtrees themselves. Clearly this is undesirable for an efficient implementation. It would be worth exploring if these subtrees may be made computationally irrelevant – as they are not needed during execution, but only used to show termination and correctness. One viable approach might be porting the development to Coq, where it is possible to make a clearer distinction between values used during execution and the propositions that may be erased.

References

- Michael Abbott, Thorsten Altenkirch, and Neil Ghani. 2005. Containers: constructing strictly positive types. *Theoretical Computer Science* 342, 1 (2005), 3–27.
- Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. 2003. A Functional Correspondence Between Evaluators and Abstract Machines. In *Proceedings of the 5th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDP '03)*. ACM, New York, NY, USA, 8–19. <https://doi.org/10.1145/888251.888254>
- Thorsten Altenkirch, Neil Ghani, Peter Hancock, Conor McBride, and Peter Morris. 2015. Indexed containers. *Journal of Functional Programming* 25 (2015).
- Thorsten Altenkirch and Conor McBride. 2003. Generic programming within dependently typed programming. In *Generic Programming*. Springer, 1–20.
- Thorsten Altenkirch, Conor McBride, and Peter Morris. 2007. Generic Programming with Dependent Types. In *Spring School on Datatype-Generic Programming*, Roland Backhouse, Jeremy Gibbons, Ralf Hinze, and Johan Jeuring (Eds.). LNCS, Vol. 4719. Springer-Verlag.
- Richard Bird and Lambert Meertens. 1998. Nested datatypes. In *International Conference on Mathematics of Program Construction*. Springer, 52–67.
- Ana Bove and Venanzio Capretta. 2005. Modelling general recursion in type theory. *Mathematical Structures in Computer Science* 15, 4 (2005), 671–708.
- James Chapman, Pierre-Évariste Dagand, Conor McBride, and Peter Morris. 2010. The Gentle Art of Levitation. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (ICFP '10)*. ACM, New York, NY, USA, 3–14. <https://doi.org/10.1145/1863543.1863547>
- Nils Anders Danielsson. 2012. Operational Semantics Using the Partiality Monad. In *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming (ICFP '12)*. ACM, New York, NY, USA, 127–138. <https://doi.org/10.1145/2364527.2364546>
- Olivier Danvy. 2009. *From Reduction-Based to Reduction-Free Normalization*. Springer Berlin Heidelberg, Berlin, Heidelberg, 66–164. https://doi.org/10.1007/978-3-642-04652-0_3
- Peter Dybjer. 1994. Inductive families. *Formal Aspects of Computing* 6, 4 (01 Jul 1994), 440–465. <https://doi.org/10.1007/BF01211308>
- Peter Dybjer and Anton Setzer. 1999. A Finite Axiomatization of Inductive-Recursive Definitions. In *Typed Lambda Calculi and Applications*, Jean-Yves Girard (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 129–146.
- Gérard Huet. 1997. The zipper. *Journal of functional programming* 7, 5 (1997), 549–554.
- Peter J Landin. 1964. The mechanical evaluation of expressions. *Comput. J.* 6, 4 (1964), 308–320.
- Per Martin-Löf. 1984. *Intuitionistic type theory*. Vol. 9. Bibliopolis Napoli.
- Conor McBride. 2008. Clowns to the Left of Me, Jokers to the Right (Pearl): Dissecting Data Structures. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '08)*. ACM, New York, NY, USA, 287–295. <https://doi.org/>

1321	10.1145/1328438.1328474	1376
1322	Peter Morris, Thorsten Altenkirch, and Conor McBride. 2006. Exploring the Regular Tree Types. In <i>Types for Proofs and Programs (TYPES 2004) (Lecture Notes in Computer Science)</i> , Christine Paulin-Mohring	1377
1323	Jean-Christophe Filliatre and Benjamin Werner (Eds.).	1378
1324	Keiko Nakata and Tarmo Uustalu. 2009. Trace-based coinductive operational semantics for while. In <i>International Conference on Theorem Proving in Higher Order Logics</i> . Springer, 375–390.	1379
1325	Thomas van Noort, Alexey Rodriguez, Stefan Holdermans, Johan Jeuring, and Bastiaan Heeren. 2008. A Lightweight Approach to Datatype-generic Rewriting. In <i>Proceedings of the ACM SIGPLAN Workshop on Generic Programming (WGP '08)</i> . ACM, New York, NY, USA, 13–24. https://doi.org/10.1145/1411318.1411321	1380
1326	Ulf Norell. 2007. <i>Towards a practical programming language based on dependent type theory</i> . Ph.D. Dissertation. Chalmers University of Technology.	1381
1327	Ulf Norell. 2008. Dependently typed programming in Agda. In <i>International School on Advanced Functional Programming</i> . Springer, 230–266.	1382
1328	Nicolas Oury and Wouter Swierstra. 2008. The Power of Pi. In <i>Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming (ICFP '08)</i> . ACM, New York, NY, USA, 39–50. https://doi.org/10.1145/1411204.1411213	1383
1329	Wouter Swierstra. 2012. From Mathematics to Abstract Machine: A formal derivation of an executable Krivine machine. In <i>Proceedings Fourth Workshop on Mathematically Structured Functional Programming, MSFP@ETAPS 2012, Tallinn, Estonia, 25 March 2012</i> . 163–177. https://doi.org/10.4204/EPTCS.76.10	1384
1330	Alexey Rodriguez Yakushev, Stefan Holdermans, Andres Löb, and Johan Jeuring. 2009. Generic Programming with Fixed Points for Mutually Recursive Datatypes. In <i>Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming (ICFP '09)</i> . ACM, New York, NY, USA, 233–244. https://doi.org/10.1145/1596550.1596585	1385
1331		1386
1332		1387
1333		1388
1334		1389
1335		1390
1336		1391
1337		1392
1338		1393
1339		1394
1340		1395
1341		1396
1342		1397
1343		1398
1344		1399
1345		1400
1346		1401
1347		1402
1348		1403
1349		1404
1350		1405
1351		1406
1352		1407
1353		1408
1354		1409
1355		1410
1356		1411
1357		1412
1358		1413
1359		1414
1360		1415
1361		1416
1362		1417
1363		1418
1364		1419
1365		1420
1366		1421
1367		1422
1368		1423
1369		1424
1370		1425
1371		1426
1372		1427
1373		1428
1374		1429
1375		1430