

Seamless, Correct, and Generic Programming over Serialised Data

GUILLAUME ALLAIS, University of Strathclyde, UK

In typed functional languages, one can typically only manipulate data in a type-safe manner if it first has been deserialised into an in-memory tree represented as a graph of nodes-as-structs and subterms-as-pointers.

We demonstrate how we can use QTT as implemented in Idris 2 to define a small universe of serialised datatypes, and provide generic programs allowing users to process values stored contiguously in buffers.

Our approach allows implementors to prove the full functional correctness by construction of the IO functions processing the data stored in the buffer.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.


Additional Key Words and Phrases: functional programming, correct-by-construction, idris, serialisation

ACM Reference Format:

Guillaume Allais. 2018. Seamless, Correct, and Generic Programming over Serialised Data. *J. ACM* 37, 4, Article 111 (August 2018), 31 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

In (typed) functional language we are used to manipulating structured data by pattern-matching on it. We include an illustrative example below.



```
data Tree
  = Leaf
  | Node Tree Bits8 Tree

sum : Tree -> Nat
sum t = case t of
  Leaf => 0
  Node l b r =>
    let m = sum l
        n = sum r
    in (m + cast b + n)
```

On the left, an example of a binary tree storing bytes in its nodes and nothing at its leaves. On the right, a small Idris 2 snippet defining the corresponding inductive type and declaring a function summing up all of the nodes' contents. It proceeds by pattern-matching: if the tree is a leaf then we immediately return 0, otherwise we start by summing up the left and right subtrees, cast the byte to a natural number and add everything up. Simply by virtue of being accepted by the typechecker,

Author's address: [Guillaume Allais](mailto:guillaume.allais@ens-lyon.org), guillaume.allais@ens-lyon.org, University of Strathclyde, 16 Richmond Street, Glasgow, Scotland, UK, G1 1XQ.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

0004-5411/2018/8-ART111 \$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

we know that this function is covering (all the possible patterns have been handled) and total (all the recursive calls are performed on smaller trees).

At runtime, the tree will quite probably be represented by constructors-as-structs and substructures-as-pointers: each constructor will be a struct with a tag indicating which constructor is represented and subsequent fields will store the constructors' arguments. Each argument will either be a value (e.g. a byte) or a pointer to either a boxed value or a substructure. If we were to directly write a function processing a value in this encoding, proving that a dispatch over a tag is covering, and that the pointer-chasing is terminating relies on global invariants tying the encoding to the inductive type. Crucially, the functional language allows us to ignore all of these details and program at a higher level of abstraction where we can benefit from strong guarantees.

Unfortunately not all data comes structured as inductive values abstracting over a constructors-as-structs and substructures-as-pointers runtime representation. Data that is stored in a file or received over the network is typically represented in a contiguous format.

We include below a textual representation of the above tree using node and leaf constructors and highlighting the data in red.

```
(node (node (node leaf 1 leaf) 5 leaf) 10 (node leaf 20 leaf))
```

This looks almost exactly like the list of bytes we get when using a naïve serialisation format based on a left-to-right in-order traversal of this tree. In the encoding below, leaves are represented by the byte 00, and nodes by the byte 01 (each byte is represented by two hexadecimal characters, we have additionally once again highlighted the bytes corresponding to data stored in the nodes):

```
(node (node leaf 1 leaf) 5 leaf)
01 01 01 00 01 00 05 00 0a 01 00 14 00
(node leaf 1 leaf)
```

The idiomatic way to process such data in a functional language is to first deserialise it as an inductive type and then call the `sum` function we defined above. If we were using a lower-level language however, we could directly process the serialised data without the need to fully deserialise it. Even a naïve port of `sum` to C can indeed work directly over buffers:

```
1 int sumAt (uint8_t buf[], int *ptr) {
2   uint8_t tag = buf[*ptr]; (*ptr)++;
3   switch (tag) {
4     case 0: return 0;
5     case 1:
6       int m = sumAt(buf, ptr);
7       uint8_t b = buf[*ptr]; (*ptr)++;
8       int n = sumAt(buf, ptr);
9       return (m + (int) b + n);
10    default: exit(-1); }}
```

This function takes a buffer of bytes, and a pointer currently indicating the start of a tree and returns the corresponding sum. We start (line 2) by reading the byte the pointer is referencing and immediately move the pointer past it. This is the tag indicating which constructor is at the root of the tree and so we inspect it (line 3). If the tag is 0 (line 4), the tree is a leaf and so we return 0 as the sum. If the tag is 1 (line 5), then the tree starts with a node and the rest of the buffer contains first the left subtree, then the byte stored in the node, and finally the right subtree. We start by summing the left subtree (line 6), after which the pointer has been moved past its end and is now pointing at the byte stored in the node. We can therefore dereference the byte and move the pointer past it (line 7), compute the sum over the right subtree (line 8), and finally add up all the components, not

99 forgetting to cast the byte to an int (line 9). If the tag is anything other than 0 or 1 (line 10) then
 100 the buffer does not contain a valid tree and so we immediately exit with an error code.

101 As we can readily see, this program directly performs pointer arithmetic, explicitly mentions
 102 buffer reads, and relies on undocumented global invariants such as the structure of the data stored
 103 in the buffer, or the fact that the pointer is being moved along and points directly past the end of a
 104 subtree once `sumAt` has finished computing its sum.

105 Our goal with this work is to completely hide all of these dangerous aspects and offer the user the
 106 ability to program over serialised data just as seamlessly and correctly as if they were processing
 107 inductive values. We will see that Quantitative Type Theory (QTT) [Atkey 2018; McBride 2016] as
 108 implemented in Idris 2 [Brady 2021] empowers us to do just that purely in library code.

1.1 Seamless Programming over Serialised Data

111 Forgetting about correctness for now, this can be summed up by the the following code snippet in
 112 which we compute the sum of the bytes stored in our type of binary trees.

```

113
114     sum : Pointer.Mu Tree _ -> IO Nat
115     sum ptr = case !(view ptr) of
116       "Leaf" # _ => pure Z
117       "Node" # l # b # r =>
118         do m <- sum l
119           n <- sum r
120           pure (m + cast b + n)
  
```

121 We reserve for later our detailed explanations of the concepts used in this snippet (`Pointer.Mu`
 122 in Section 5, `view` in Section 7.4). For now, it is enough to understand that the function is an `IO`
 123 process inspecting a buffer that contains a tree stored in serialised format and computing the same
 124 sum as the pure function seen in the previous section. In both cases, if we uncover a leaf ("`Leaf`" #
 125 _) then we return zero, and if we uncover a node ("`Node`" # `l` # `b` # `r`) with a left branch `l`, a stored
 126 byte `b`, and a right branch `r`, then we recursively compute the sums for the left and right subtrees,
 127 cast the byte to a natural number and add everything up. Crucially, the two functions look eerily
 128 similar, and the one operating on serialised data does not explicitly perform error-prone pointer
 129 arithmetic, or low-level buffer reads. This is the first way in which our approach shines.

130 One major difference between the two functions is that we can easily prove some of the pure
 131 function's properties by a structural induction on its input whereas we cannot prove anything
 132 about the `IO` process without first explicitly postulating the `IO` monad's properties. Our second
 133 contribution tackles this issue.

1.2 Correct Programming over Serialised Data

136 We will see that we can refine that second definition to obtain a correct-by-construction version of
 137 `sum`, with almost exactly the same code.

```

138
139     sum : Pointer.Mu Tree t ->
140         IO (Singleton (Data.sum t))
141     sum ptr = case !(view ptr) of
142       "Leaf" # _ => pure [| Z |]
143       "Node" # l # b # r =>
144         do m <- sum l
145           n <- sum r
146           pure [| [| m + [| cast b |] |] + n |]
  
```

148 In the above snippet, we can see that the `Pointer.Mu` is indexed by a phantom parameter: a runtime
 149 irrelevant `t` which has type `(Data.Mu Tree)`. And so the return type can mention the result of the
 150 pure computation `(Data.sum t)`. `Singleton` is, as its name suggests, a singleton type (cf. Section 6)
 151 i.e. the natural number we compute is now proven to be equal to the one computed by the pure `sum`
 152 function. The implementation itself only differs in that we had to use idiom brackets [McBride and
 153 Paterson 2008], something we will explain in Section 6.2.

154 In other words, our approach also allows us to prove the functional correctness of the `IO` proce-
 155 dures processing trees stored in serialised format in a buffer. This is our second main contribution.

156 1.3 Generic Programming over Serialised Data

158 Last but not least, as Altenkirch and McBride demonstrated [Altenkirch and McBride 2002]: “With
 159 dependently (sic) types, generic programming is just programming; it is not necessary to write a
 160 new compiler each time a useful universe presents itself.”

161 In this paper we carve out a universe of inductive types that can be uniformly serialised and
 162 obtain all of our results by generic programming. In practice this means that we are not limited to
 163 the type of binary trees with bytes stored in the nodes we used in the examples above. We will for
 164 instance be able to implement a generic and correct-by-construction definition of `fold` operating
 165 on data stored in a buffer whose type declaration can be seen below (we will explain how it is
 166 defined in Section 7.5).

```
167 fold : {cs : Data nm} -> (alg : Alg cs a) ->
168 forall t. Pointer.Mu cs t ->
169 IO (Singleton (Data.fold alg t))
```

171 This data-genericity is our third contribution.

172 1.4 Plan

174 In summary, we are going to define a library for the seamless, correct, and generic manipulation of
 175 algebraic types in serialised format.

176 Section 2 introduces the language of descriptions capturing the subset of inductively defined
 177 types that our work can handle. It differs slightly from usual presentations in that it ensures the
 178 types can be serialised and tracks crucial invariants towards that goal. Section 3 gives a standard
 179 meaning to these data descriptions as strictly positive endofunctors whose fixpoints give us the
 180 expected inductive types. We will use this standard meaning in the specification layer of our work.
 181 Section 4 explores the serialisation format we have picked for these trees: a depth-first, left-to-right
 182 infix traversal of the trees, with additional information stored to allow for the random access of any
 183 subtree. Section 5 defines the type of pointers to trees stored in a buffer and shows how we can use
 184 such pointers to write the corresponding tree to a file. Section 6 introduces the terminology of *views*
 185 and *singleton* types that is crucial to the art of programming in a correct-by-construction manner.
 186 Section 7 defines IO primitives that operate on serialised trees stored in an underlying buffer. They
 187 encapsulate all the unsafe low-level operations and offer a high-level interface that allows users to
 188 implement correct-by-construction procedures. Section 8 defines a set of serialisation combinators
 189 that allows users to implement correct-by-construction procedures writing values into a buffer.
 190 Section 9 discusses some preliminary performance results for the library.

191 2 OUR UNIVERSE OF DESCRIPTIONS

193 We first need to pin down the domain of our discourse. To talk generically about an entire class of
 194 datatypes without needing to modify the host language we have decided to perform a universe
 195 construction [Benke et al. 2003; Löh and Magalhães 2011; Morris 2007]. That is to say that we are

196

going to introduce an inductive type defining a set of codes together with an interpretation of these codes as bona fide host-language types. We will then be able to program generically over the universe of datatypes by performing induction on the type of codes [Pfeifer and Rueß 1999].

The universe we define is in the tradition of a sums-of-products vision of inductive types [Jansson and Jeuring 1997] where the data description records additional information about the static and dynamic size of the data being stored. In our setting, constructors are essentially arbitrarily nested tuples of values of type unit, bytes, and recursive substructures. A datatype is given by listing a choice of constructors.

2.1 Descriptions

We start with these constructor descriptions; they are represented internally by an inductive family `Desc` declared below.

```
data Desc : (rightmost : Bool) ->
           (static : Nat) -> (offsets : Nat) ->
           Type
```

This family has three indices corresponding to three crucial invariants being tracked. First, an index telling us whether the current description is being used in the `rightmost` branch of the overall constructor description. Second, the `statically` known size of the described data in the number of bytes it occupies. Third, the number of `offsets` that need to be stored to compensate for subterms not having a statically known size. The reader should think of `rightmost` as an ‘input’ index whereas `static` and `offsets` are ‘output’ indices.

Next we define the family proper by giving its four constructors.

```
data Desc where
  None : Desc r 0 0
  Byte : Desc r 1 0
  Prod : {sl, sr, ol, or : Nat} ->
         Desc False sl ol -> Desc r sr or ->
         Desc r (sl + sr) (ol + or)
  Rec : Desc r 0 (ifThenElse r 0 1)
```

Each constructor can be used anywhere in a description so their return `rightmost` index can be an arbitrary boolean.

`None` is the description of values of type unit. The static size of these values is zero as no data is stored in a value of type unit. Similarly, they do not require an offset to be stored as we statically know their size.

`Byte` is the description of bytes. Their static size is precisely one byte, and they do not require an offset to be stored either.

`Prod` gives us the ability to pair two descriptions together. Its static size and the number of offsets are the respective sums of the static sizes and numbers of offsets of each subdescription. The description of the left element of the pair will never be in the `rightmost` branch of the overall constructors description and so its index is `False` while the description of the right element of the pair is in the `rightmost` branch precisely whenever the whole pair is; hence the propagation of the `r` arbitrary value from the return index into the description of the right component.

Last but not least, `Rec` is a position for a subtree. We cannot know its size in bytes statically and so we decide to store an offset unless we are in the `rightmost` branch of the overall description. Indeed, there are no additional constructor arguments behind the `rightmost` one and so we have no reason to skip past the subterm. Consequently we do not bother recording an offset for it.

246 2.2 Constructors

247 We represent a constructor as a record packing together a name for the constructor, the description
 248 of its arguments (which is, by virtue of being used at the toplevel, in rightmost position), and the
 249 values of the `static` and `offsets` invariants. The two invariants are stored as implicit fields because
 250 their value is easily reconstructed by Idris 2 using unification and so users do not need to spell
 251 them out explicitly.

```
252 record Constructor (nm : Type) where
253   constructor (::)
254   name : nm
255   {static : Nat}
256   {offsets : Nat}
257   description : Desc True static offsets
258
```

259 Note that we used `::` as the name of the constructor for records of type `Constructor`. This
 260 allows us to define constructors by forming an expression reminiscent of Haskell's type declarations:
 261 `name :: type`. Returning to our running example, this gives us the following encodings for leaves
 262 that do not store anything and nodes that contain a left branch, a byte, and a right branch.

```
263
264 Leaf : Constructor String           Node : Constructor String
265 Leaf = "Leaf" :: None              Node = "Node" :: Prod Rec (Prod Byte Rec)
266
```

267 2.3 Datatypes

268 A datatype description is given by a number of constructors together with a vector (also known as
 269 a length-indexed list) associating a description to each of these constructors.

```
270
271 record Data (nm : Type) where                                     Tree : Data String
272   constructor MkData                                           Tree = MkData [Leaf, Node]
273   {consNumber : Nat}
274   constructors : Vect consNumber (Constructor nm)
275
```

276 We can then encode our running example as a simple `Data` declaration: a binary tree whose node
 277 stores bytes is described by the choice of either a `Leaf` or `Node`, as defined above.

278 Now that we have a language that allows us to give a description of our inductive types, we are
 279 going to give these descriptions a meaning as trees.

280 3 MEANING AS TREES

281 We now see descriptions as functors and, correspondingly, datatypes as the initial objects of the
 282 associated functor-algebras. This is a standard construction derived from Malcolm's work [Malcolm
 283 1990], itself building on Hagino's categorically-inspired definition of a lambda calculus with a
 284 generic notion of datatypes [Hagino 1987].

285 In our work these trees will be used primarily to allow users to give a precise specification of the
 286 IO procedures they actually want to write in order to process values stored in buffers. We expect
 287 these inductive trees and the associated generic programs consuming them to be mostly used at
 288 the \emptyset modality i.e. to be erased during compilation.

290 3.1 Descs as Functors

291 We define the meaning of descriptions as strictly positive endofunctors on `Type` by induction on
 292 said descriptions. `Meaning` gives us the action of the functors on objects.

293
294

```

295
296 Meaning : Desc r s n -> Type -> Type           record Tuple (a, b : Type) where
297 Meaning None x = ()                           constructor (#)
298 Meaning Byte x = Bits8                        fst : a
299 Meaning Rec x = x                             snd : b
300 Meaning (Prod d e) x
301     = Tuple (Meaning d x) (Meaning e x)

```

Both `None` and `Byte` are interpreted by constant functors (respectively the one returning the unit type, and the one returning the type of bytes). `Rec` is the identity functor. Finally `(Prod d e)` is interpreted as the pairing of the interpretation of `d` and `e` respectively. We use our own definition of pairing rather than the standard library's because it gives us better syntactic sugar.

This gives us the action of descriptions on types, let us now see their action on morphisms. We once again proceed by induction on the description.

```

308 fmap : (d : Desc r s o) -> (a -> b) -> Meaning d a -> Meaning d b
309 fmap None f v = v
310 fmap Byte f v = v
311 fmap (Prod d e) f (v # w) = (fmap d f v # fmap e f w)
312 fmap Rec f v = f v

```

All cases but the one for `Rec` are structural. Verifying that these definitions respect the functor laws is left as an exercise for the reader.

3.2 Data as Trees

Given a datatype description `cs`, our first goal is to define what it means to pick a constructor. The `Index` record is a thin wrapper around a finite natural number known to be smaller than the number of constructors this type provides.

```

321 record Index (cs : Data nm) where
322   constructor MkIndex
323   getIndex : Fin (consNumber cs)

```

We use this type rather than `Fin` directly because it plays well with inference. In the following code snippet, implementing a function returning the description corresponding to a given index, we use this to our advantage: the `cs` argument can be left implicit because it already shows up in the type of the `Index` and can thus be reconstructed by unification.

```

328 description : {cs : Data nm} -> (k : Index cs) ->
329     let cons = index (getIndex k) (constructors cs) in
330     Desc True (static cons) (offsets cons)
331 description {cs} k
332   = description (index (getIndex k) (constructors cs))

```

This type of indices also allows us to provide users with syntactic sugar enabling them to use the constructors' names directly rather than confusing numeric indices. The following function runs a decision procedure `isConstructor` at the type level in order to turn any raw string `str` into the corresponding `Index`.

```

337 fromString : {cs : Data String} -> (str : String) ->
338     {auto 0 _ : IsJust (isConstructor str cs)} ->
339     Index cs
340 fromString {cs} str with (isConstructor str cs)
341   _ | Just k = MkIndex k

```

344 If the name is valid then `isConstructor` will return a valid `Index` and Idris 2 will be able to
 345 automatically fill-in the implicit proof. If the name is not valid then Idris 2 will not find the index
 346 and will raise a compile time error. We include a successful example on the left and a failing test on
 347 the right hand side (**failing** blocks are only accepted in Idris 2 if their body leads to an error).

```
348
349 indexLeaf : Index Tree           failing
350 indexLeaf = "Leaf"              notIndexCons : Index Tree
                                    notIndexCons = "Cons"
```

352 Once equipped with the ability to pick constructors, we can define the type of algebras for the
 353 functor described by a `Data` description. For each possible constructor, we demand an algebra for
 354 the functor corresponding to the meaning of the constructor's description.

```
355 Alg : Data nm -> Type -> Type
356 Alg cs x = (k : Index cs) -> Meaning (description k) x -> x
357
```

358 We can then introduce the fixpoint of data descriptions as the initial algebra, defined as the
 359 following inductive type.

```
360 data Mu : Data nm -> Type where
361   (#) : Alg cs (assert_total (Mu cs))
```

362 Note that here we are forced to use `assert_total` to convince Idris 2 to accept the definition.
 363 Indeed, unlike Agda, Idris 2 does not (yet!) track whether a function's arguments are used in a
 364 strictly positive manner. Consequently the positivity checker is unable to see that `Meaning` uses its
 365 second argument in a strictly positive manner and that this is therefore a legal definition.

366 Now that we can build trees as fixpoints of the meaning of descriptions, we can define convenient
 367 aliases for the `Tree` constructors. Note that the leftmost `(#)` use in each definition corresponds to
 368 the `Mu` constructor while later ones are `Tuple` constructors. Idris 2's type-directed disambiguation
 369 of constructors allows us to use this uniform notation for all of these pairing notions.

```
371 leaf : Mu Tree           node : Mu Tree -> Bits8 -> Mu Tree -> Mu Tree
372 leaf = "Leaf" # ()      node l b r = "Node" # l # b # r
373
```

374 This enables us to define our running example as an inductive value:

```
375 example : Mu Tree
376 example = node (node (node leaf 1 leaf) 5 leaf) 10 (node leaf 20 leaf)
377
```

378 3.3 Generic Fold

379 `Mu` gives us the initial fixpoint for these algebras i.e. given any other algebra over a type `a`, from
 380 a term of type `(Mu cs)`, we can compute an `a`. We define the generic `fold` function over inductive
 381 values as follows:

```
382 fold : {cs : Data nm} -> Alg cs a -> Mu cs -> a
383 fold alg (k # t) = alg k (assert_total $ fmap _ (fold alg) t)
384
```

385 We first match on the term's top constructor, use `fmap` (defined in Section 3.1) to recursively
 386 apply the fold to all the node's subterms and finally apply the algebra to the result.

387 Here we only use `assert_total` because Idris 2 does not see that `fmap` only applies its argument
 388 to strict subterms. This limitation could easily be bypassed by mutually defining an inlined and
 389 specialised version of `(fmap _ (fold alg))`, as we demonstrate in Appendix A. In an ideal type
 390 theory these supercompilation steps, whose sole purpose is to satisfy the totality checker, would be
 391 automatically performed by the compiler [Mendel-Gleason 2012].

392

Further generic programming can yield other useful programs e.g. a generic proof that tree equality is decidable or a generic definition of zippers [Löh and Magalhães 2011].

4 SERIALISED REPRESENTATION

Before we can give a meaning to descriptions as pointers into a buffer we need to decide on a serialisation format. The format we have opted for is split in two parts: a header containing data that can be used to check that a user’s claim that a given file contains a serialised tree of a given type is correct, followed by the actual representation of the tree.

For instance, the following binary snippet is a hex dump of a file containing the serialised representation of a binary tree belonging to the type we have been using as our running example. The raw data is semantically highlighted: 8-bytes-long `offsets`, a `type` description of the stored data, some `nodes` of the tree and the `data` stored in the nodes.

```

87654321  00 11 22 33 44 55 66 77 88 99 AA BB CC DD EE FF
00000000: 07 00 00 00 00 00 00 00 00 02 00 02 03 02 01 03 01
00000010: 17 00 00 00 00 00 00 00 00 01 0c 00 00 00 00 00 00
00000020: 00 01 01 00 00 00 00 00 00 00 00 00 01 00 05 00 0a
00000030: 01 01 00 00 00 00 00 00 00 00 00 14 00

```

More specifically, this block is the encoding of the `example` given in the previous section and, knowing that a `leaf` is represented here by `00` and a `node` is represented by `01` the careful reader can check (modulo ignoring the type description and offsets for now) that the data is stored in a depth-first, left-to-right traversal of the tree (i.e. we get exactly the bit pattern we saw in the naïve encoding presented in Section 1).

4.1 Header

In our example, the header is as follows:

```
07 00 00 00 00 00 00 00 00 02 00 02 03 02 01 03
```

The header consists of an offset allowing us to jump past it in case we do not care to inspect it, followed by a binary representation of the `Data` description of the value stored in the buffer. This can be useful in a big project where different components produce and consume such serialised values: if we change the format in one place but forget to update it in another, we want the program to gracefully fail to load the file using an unexpected format. We detail in Section 10.2.1 how dependent type providers can help structure a software project to prevent such issues.

The encoding of a data description starts with a byte giving us the number of constructors, followed by these constructors’ respective descriptions serialised one after the other. `None` is represented by `00`, `Byte` is represented by `01`, `(Prod d e)` is represented by `02` followed by the representation of `d` and then that of `e`, and `Rec` is represented by `03`.

Looking once more at the header in the running example, the `Data` description is indeed 7 bytes long like the offset states. The `Data` description starts with `02` meaning that the type has two constructors. The first one is `00` i.e. `None` (this is the encoding of the type of `Leaf`), and the second one is `02 03 02 01 03` i.e. `(Prod Rec (Prod Byte Rec))` (that is to say the encoding of the type of `Node`). According to the header, this file does contain a `Tree`.

4.2 Tree Serialisation

Our main focus in the definition of this format is that we should be able to process any of a node’s subtrees without having to first traverse the subtrees that come before it. This will allow us to, for instance, implement a function looking up the value stored in the rightmost node in our running

example type of binary trees in time linear in the depth of the tree rather than exponential. To this end each node needs to store an offset measuring the size of the subtrees that are to the left of any relevant information.

If a given tag is associated to a description of type (Desc True s o) then the representation in memory of the associated node will look something like the following.

tag	o offsets	tree ₁ ... byte ₁ ... tree _k ... byte _s tree _{o+1}
0	1	1 + 8 * o 8 * o + s + $\sum_{i=1}^o o_i$

On the first line we have a description of the data layout and on the second line we have the offset of various positions in the block with respect to the tag's address.

For the data layout, we start with the tag then we have o offsets, and finally we have a block contiguously storing an interleaving of subtrees and s bytes dictated by the description. In this example the rightmost value in the description is a subtree and so even though we have o offsets, we actually have $(o + 1)$ subtrees stored.

The offsets of the tag with respect to its own address is 0. The tag occupies one byte and so the offset of the block of offsets is 1. Each offset occupies 8 bytes and so the constructor's arguments are stored at offset $(1 + 8 * o)$. Finally each value's offset can be computed by adding up the offset of the start of the block of constructor arguments, the offsets corresponding to all of the subtrees that come before it, and the number of bytes stored before it; in the case of the last byte that gives $1 + 8 * o + \sum_{i=1}^o o_i + s - 1$ hence the formula included in the diagram.

Going back to our running example, this translates to the following respective data layouts and offsets for a leaf and a node.

Leaf	Node				
00	01	offset	left subtree	byte	right subtree
0	0	1	9	9 + o ₁	10 + o ₁

Now that we understand the format we want, we ought to be able to implement pointers and the functions manipulating them.

5 MEANING AS POINTERS INTO A BUFFER

Now that we know the serialisation format, we can give a meaning to constructor and data descriptions as pointers into a buffer. For reasons that will become apparent in Section 7.5 when we start programming over serialised data in a correct-by-construction manner, our types of 'pointers' will be parameterised not only by the description of the type of the data stored but also by a runtime-irrelevant inductive value of that type. For now, it is enough to think of these indices as a lightweight version of the 'points to' assertions used in separation logic [Reynolds 2002] when reasoning about imperative programs. We expand on this analogy in Appendix C where we also discuss the connection with the combinators defined in Section 7.

5.1 Tracking Buffer Positions

We start with the definition of the counterpart to Mu for serialised values.

```

491 record Mu (cs : Data nm) (t : Data.Mu cs) where
492   constructor MkMu
493   muBuffer : Buffer
494   muPosition : Int
495   muSize : Int

```

A tree sitting in a buffer is represented by a record packing the buffer, the position at which the tree's root node is stored, and the size of the tree. Note that according to our serialisation format the size is not stored in the file but using the size of the buffer, the stored offsets, and the size of the static data we will always be able to compute a value corresponding to it.

```

500 record Meaning (d : Desc r s o) (cs : Data nm)
501   (t : Data.Meaning d (Data.Mu cs)) where
502   constructor MkMeaning
503   subterms : Vect o Int
504   meaningBuffer : Buffer
505   meaningPosition : Int
506   meaningSize : Int

```

The counterpart to a `Meaning` stores additional information. For a description of type `(Desc r s o)` on top of the buffer, the position at which the root of the meaning resides, and the size of the layer we additionally have a vector of `o` offsets that allow us to efficiently access any value we want.

5.2 Writing a Tree to a File

Once we have a pointer to a tree in a buffer, we can easily write it to a file be it for safekeeping or sending over the network.

```

515 writeToFile : {cs : Data nm} -> FilePath ->
516   forall t. Pointer.Mu cs t -> IO ()
517 writeToFile fp (MkMu buf pos size) = do
518   desc <- getInt buf 0
519   let start = 8 + desc
520       bufSize = 8 + desc + size
521       buf <- if pos == start then pure buf else do
522         Just newbuf <- newBuffer bufSize
523         | Nothing => failWith "{__LOC__} Couldn't allocate buffer"
524         copyData buf 0 start newbuf 0
525         copyData buf start size newbuf start
526         pure buf
527   Right () <- writeBufferToFile fp buf bufSize
528   | Left (err, _) => failWith (show err)
529   pure ()

```

We first start by reading the size of the header stored in the buffer. This allows us to compute both the `start` of the data block as well as the size of the buffer (`bufSize`) that will contain the header followed by the tree we want to write to a file. We then check whether the position of the pointer is exactly the beginning of the data block. If it is then we are pointing to the whole tree and the current buffer can be written to a file as is. Otherwise we are pointing to a subtree and need to separate it from its surrounding context first. To do so we allocate a new buffer of the right size and use the standard library's `copyData` primitive to copy the raw bytes corresponding to the header first, and the tree of interest second. We can then write the buffer we have picked to a file and happily succeed.

Now that we have pointers and can save the tree they are standing for, we are only missing the ability to look at the content they are pointing to. But first we need to introduce some basic tools to be able to talk precisely about this stored content.

6 INTERLUDE: VIEWS AND SINGLETONS

The precise indexing of pointers by a runtime-irrelevant copy of the value they are pointing to means that inspecting the buffer's content should not only return runtime information but also refine the index to reflect that information at the type-level. As a consequence, the functions we are going to define in the following subsections are views.

6.1 Views

A view in the sense of Wadler [Wadler 1987], and subsequently refined by McBride and McKinna [McBride and McKinna 2004] for a type T is a type family V indexed by T together with a function which maps values t of type T to values of type $V t$. By inspecting the $V t$ values we can learn something about the t input. The prototypical example is perhaps the 'snoc' ('cons' backwards) view of right-nested lists as if they were left-nested. We present the `Snoc` family below.

```
data Snoc : List a -> Type where
  Lin : Snoc []
  (:<) : (init : List a) -> (last : a) -> Snoc (init ++ [last])
```

By matching on a value of type `(Snoc xs)` we get to learn either that `xs` is empty (`Lin`, nil backwards) or that it has an initial segment `init` and a last element `last` (`init :< last`). The function `unsnoc` demonstrates that we can always *view* a `List` in a `Snoc`-manner.

```
unsnoc : (xs : List a) -> Snoc xs
unsnoc [] = Lin
unsnoc (x :: xs@_) with (unsnoc xs)
  _ | [<] = [] :< x
  _ | init :< last = (x :: init) :< last
```

Here we defined `Snoc` as an inductive family but it can sometimes be convenient to define the family recursively instead. In which case the `Singleton` inductive family can help us connect runtime values to their runtime-irrelevant type-level counterparts.

6.2 The Singleton type

The `Singleton` family has a single constructor which takes an argument `x` of type `a`, its return type is indexed precisely by this `x`.

```
data Singleton : {0 a : Type} -> (x : a) -> Type where
  MkSingleton : (x : a) -> Singleton x
```

More concretely this means that a value of type `(Singleton t)` has to be a runtime relevant copy of the term `t`. Note that Idris 2 performs an optimisation similar to Haskell's newtype unwrapping: every data type that has a single non-recursive constructor with only one non-erased argument is unwrapped during compilation. This means that at runtime the `Singleton` / `MkSingleton` indirections will have disappeared.

We can define some convenient combinators to manipulate singletons. We reuse the naming conventions typical of applicative functors which will allow us to rely on Idris 2's automatic desugaring of *idiom brackets* [McBride and Paterson 2008] into expressions using these combinators.

```

589 pure : (x : a) -> Singleton x
590 pure = MkSingleton

```

591 First `pure` is a simple alias for `MkSingleton`, it turns a runtime-relevant value `x` into a singleton
 592 for this value.

```

593 (<$>) : (f : a -> b) -> Singleton t -> Singleton (f t)
594 f <$> MkSingleton t = MkSingleton (f t)

```

595
 596 Next, we can ‘map’ a function under a `Singleton` layer: given a pure function `f` and a runtime
 597 copy of `t` we can get a runtime copy of `(f t)`.

```

598 (<*>) : Singleton f -> Singleton t -> Singleton (f t)
599 MkSingleton f <*> MkSingleton t = MkSingleton (f t)

```

600 Finally, we can apply a runtime copy of a function `f` to a runtime copy of an argument `t` to get a
 601 runtime copy of the result `(f t)`.

602 As we mentioned earlier, Idris 2 automatically desugars idiom brackets using these combinators.
 603 That is to say that `[| x |]` will be elaborated to `(pure x)` while `[| f t1 ··· tn |]` will become `(f <$> t1`
 604 `<*> ··· <*> tn)`. This lets us apply `Singleton`-wrapped values almost as seamlessly as pure values.

605 We are now equipped with the appropriate notions and definitions to look at a buffer’s content.
 606

607 7 INSPECTING A BUFFER’S CONTENT

608 We can now describe the combinators allowing our users to inspect the value they have a pointer
 609 for. We are going to define the most basic of building blocks (`poke` and `out`), combine them to derive
 610 useful higher-level combinators (`layer` and `view`), and ultimately use these to implement a generic
 611 correct-by-construction version of `fold` operating over trees stored in a buffer (cf. Section 7.5).

612 Readers may be uneasy about the unsafe implementations of the basic building blocks but we
 613 argue that it is a necessary evil by drawing an extended analogy to separation logic in Appendix C.
 614

615 7.1 Poking the Buffer

616 Our most basic operation consists in poking the buffer to unfold the description by exactly one
 617 step. The type of the function is as follows: provided a pointer for a meaning `t`, we return an `IO`
 618 process computing the one step unfolding of the meaning.
 619

```

620 poke : {θ cs : Data nm} -> {d : Desc r s o} ->
621       forall t. Pointer.Meaning d cs t ->
622       IO (Poke d cs t)

```

623 The result type of this operation is defined by case-analysis on the description. In order to
 624 keep the notations user-friendly, we mutually define a recursive function `Poke` interpreting the
 625 straightforward type constructors and an inductive family `Poke’` with interesting return indices.
 626

```

627 Poke : (d : Desc r s o) -> (cs : Data nm) ->
628       Data.Meaning d (Data.Mu cs) -> Type
629 Poke None _ t = ()
630 Poke Byte cs t = Singleton t
631 Poke Rec cs t = Pointer.Mu cs t
632 Poke d@(Prod _ _) cs t = Poke’ d cs t

```

633 Poking a buffer containing `None` will return a value of the unit type as no information whatsoever
 634 is stored there.

635 If we access a `Byte` then we expect that inspecting the buffer will yield a runtime-relevant copy
 636 of the type-level byte we have for reference. Hence the use of `Singleton`.
 637

638 If the description is `Rec` this means we have a substructure. In this case we simply demand a
639 pointer to it.

640 Last but not least, if we access a `Prod` of two descriptions then the type-level term better be a pair
641 and we better be able to obtain a `Pointer.Meaning` for each of the sub-meanings. Because Idris 2
642 does not currently support definitional eta equality for records, it will be more ergonomic for users
643 if we introduce `Poke'` rather than yielding a `Tuple` of values. By matching on `Poke'` at the value
644 level, they will see the pair at the type level also reduced to a constructor-headed tuple.

```
645 data Poke' : (d : Desc r s o) -> (cs : Data nm) ->
646           Data.Meaning d (Data.Mu cs) -> Type where
647   (#) : Pointer.Meaning d cs t ->
648         Pointer.Meaning e cs u ->
649         Poke' (Prod d e) cs (t # u)
650
```

651 The implementation of this operation proceeds by case analysis on the description. As we are
652 going to see shortly, it is necessarily somewhat unsafe as we claim to be able to connect a type-level
653 value to whatever it is that we read from the buffer. Let us go through each case one-by-one.

```
654 poke {d = None} el = pure ()
655
```

656 If the description is `None` we do not need to fetch any information from the buffer and can
657 immediately return `()`.

```
658
659 poke {d = Byte} el = do
660   bs <- getBits8 (meaningBuffer el) (meaningPosition el)
661   pure (unsafeMkSingleton bs)
662
```

663 If the description is `Byte` then we read a byte at the determined position. The only way we can
664 connect this value we just read to the type index is to use the unsafe combinator `unsafeMkSingleton`
665 to manufacture a value of type `(Singleton t)` instead of the value of type `(Singleton bs)` we would
666 expect from wrapping `bs` in the `MkSingleton` constructor. As we explain in Appendix C.2.1, in
667 separation logic this would correspond to declaring an axiom about the `poke` language construct.

```
668 poke {d = Prod {s1, ol} d e} {t} (MkMeaning sub buf pos size) = do
669   let (subl, subr) = splitAt ol sub
670       let sizel = sum subl + cast s1
671           let left = MkMeaning subl buf pos sizel
672               let posr = pos + sizel
673                   let right = MkMeaning subr buf posr (size - sizel)
674                       pure (rewrite etaTuple t in left # right)
675
```

676 If the description is the product of two sub-descriptions then we want to compute the `Pointer.Meaning`
677 corresponding to each of them. We start by splitting the vector of offsets to distribute them between
678 the left and right subtrees.

679 We can readily build the pointer for the `left` subdescription: it takes the left offsets, the buffer,
680 and has the same starting position as the whole description of the product as the submeanings are
681 stored one after the other. Its size (`sizel`) is the sum of the space reserved by all of the left offsets
682 (`sum subl`) as well as the static size occupied by the rest of the content (`s1`).

683 We then compute the starting position of the right subdescription: we need to move past the whole
684 of the left subdescription, that is to say that the starting position is the sum of the starting position
685 for the whole product and `sizel`. The size of the right subdescription is then easily computed by
686 subtracting `sizel` from the overall `size` of the paired subdescriptions.

687 We can finally use the lemma `etaTuple` saying that a tuple is equal to the pairing of its respective
 688 projections in order to turn `t` into `(fst t # snd t)` which lets us use the `Poke'` constructor (`#`) to
 689 return our pair of pointers.

```
690 poke {d = Rec} (MkMeaning _ buf pos size) = pure (MkMu buf pos size)
```

692 Lastly, when we reach a `Rec` description, we can discard the vector of offsets and return a
 693 `Pointer.Mu` with the same buffer, starting position and size as our input pointer.

695 7.2 Extracting one layer

696 By repeatedly poking the buffer, we can unfold a full layer. The result of this operation is defined
 697 by induction on the description. It is identical to the definition of `Poke` except for the `Prod` case:
 698 here, instead of being content with a pointer for each of the subdescriptions, we demand a full
 699 layer for them too.

```
700
701 Layer : (d : Desc r s o) -> (cs : Data nm) ->
702         Data.Meaning d (Data.Mu cs) -> Type
703 Layer None _ _ = ()
704 Layer Byte _ t = Singleton t
705 Layer Rec cs t = Pointer.Mu cs t
706 Layer d@(Prod _ _) cs t = Layer' d cs t
707
708 data Layer' : (d : Desc r s o) -> (cs : Data nm) ->
709         Data.Meaning d (Data.Mu cs) -> Type where
710   (#) : Layer d cs t -> Layer e cs u -> Layer' (Prod d e) cs (t # u)
```

711 This function can easily be implemented by induction on the description and repeatedly calling
 712 `poke` to expose the values one by one.

```
713
714 layer : {0 cs : Data nm} -> {d : Desc r s o} ->
715         forall t. Pointer.Meaning d cs t -> IO (Layer d cs t)
716 layer el = poke el >>= go d where
717
718   go : forall r, s, o. (d : Desc r s o) ->
719         forall t. Poke d cs t -> IO (Layer d cs t)
720   go None p = pure ()
721   go Byte p = pure p
722   go (Prod d e) (p # q) = [| layer p # layer q |]
723   go Rec p = pure p
```

725 7.3 Exposing the top constructor

726 Now that we can deserialise an entire layer of `Meaning`, the only thing we are missing to be able
 727 to generically manipulate trees is the ability to expose the top constructor of a tree stored at a
 728 `Pointer.Mu` position. Remembering the data layout detailed in Section 4.2, this will amount to
 729 inspecting the tag used by the node and then deserialising the offsets stored immediately after it.

730 The `Out` family describes the typed point of view: to get your hands on the index of a tree's
 731 constructor means obtaining an `Index`, and a `Pointer.Meaning` to the constructor's arguments
 732 (remember that these high-level 'pointers' store a vector of offsets). The family's index (`k # t`)
 733 ensures that the structure of the runtime irrelevant tree is adequately described by the index (`k`)
 734 and the `Data.Meaning (t)` the `Pointer.Meaning` is for.

735

```

736 data Out : (cs : Data nm) -> (t : Data.Mu cs) -> Type where
737   (#) : (k : Index cs) ->
738     forall t. Pointer.Meaning (description k) cs t ->
739     Out cs (k # t)

```

The type of the `out` function is as expected: given a pointer to a tree `t` of type `cs` we can get a value of type `(Out cs t)`. That is to say, we can get a view allowing us to reveal what the index of the tree's head constructor is.

```

743 out : {cs : Data nm} -> forall t. Pointer.Mu cs t ->
744     IO (Out cs t)

```

The implementation is fairly straightforward except for another unsafe step meant to reconcile the information we read in the buffer with the runtime-irrelevant tree index.

```

748 out {t} mu = do
749   tag <- getBits8 (muBuffer mu) (muPosition mu)
750   let Just k = MkIndex <$> natToFin (cast tag) (consNumber cs)
751       | _ => failWith "Invalid representation"
752   let 0 sub = unfoldAs k t
753       val <- (k #) <$> getConstructor k {t = sub.fst}
754           (rewrite sym sub.snd in mu)
755   pure (rewrite sub.snd in val)

```

We start by reading the tag `k` corresponding to the constructor choice: we obtain a byte by calling `getBits8`, cast it to a natural number and then make sure that it is in the range `[0 .. consNumber cs]` using `natToFin`. We then use the unsafe `unfoldAs` postulate to step the type-level `t` to something of the form `(k # val)`.

```

760 %unsafe
761 0 unfoldAs :
762   (k : Index cs) -> (0 t : Data.Mu cs) ->
763   (val : Data.Meaning (description k) (Data.Mu cs)
764   ** t === (k # val))

```

The declaration of `unfoldAs` is marked as runtime irrelevant because it cannot possibly be implemented (`t` is runtime irrelevant and so cannot be inspected) and so its output should not be relied upon in runtime-relevant computations. Its type states that there exists a `Meaning` called `val` such that `t` is equal to `(k # val)`.

Now that we know the head constructor we want to deserialize and that we have the ability to step the runtime irrelevant tree to match the actual content of the buffer, we can use `getConstructor` to build such a value.

```

773 getConstructor : (k : Index cs) ->
774   forall t. Pointer.Mu cs (k # t) ->
775   IO (Pointer.Meaning (description k) cs t)
776 getConstructor (MkIndex k) mu
777 = let offs : Nat; offs = offsets (index k $ constructors cs) in
778   getOffsets (muBuffer mu) (1 + muPosition mu) offs
779 $ let size = muSize mu - 1 - cast (8 * offs) in
780   \ subterms, pos => MkMeaning subterms (muBuffer mu) pos size

```

To get a constructor, we start by getting the vector of offsets stored immediately after the tag. We then compute the size of the remaining `Meaning` description: it is the size of the overall tree, minus 1 (for the tag) and 8 times the number of offsets (because each offset is stored as an 8 bytes

784

number). We can then use the record constructor `MkMeaning` to pack together the vector of offsets, the buffer, the position past the offsets and the size we just computed.

```

785 number). We can then use the record constructor MkMeaning to pack together the vector of offsets,
786 the buffer, the position past the offsets and the size we just computed.
787
788 getOffsets : Buffer -> (pos : Int) ->
789   (n : Nat) ->
790   forall t. (Vect n Int -> Int -> Pointer.Meaning d cs t) ->
791   IO (Pointer.Meaning d cs t)
792 getOffsets buf pos 0 k = pure (k [] pos)
793 getOffsets buf pos (S n) k = do
794   off <- getInt buf pos
795   getOffsets buf (8 + pos) n (k . (off ::))

```

The implementation of `getOffsets` is straightforward: given a continuation that expect `n` offsets as well as the position past the last of these offsets, we read the 8-bytes-long offsets one by one and pass them to the continuation, making sure that we move the current position accordingly before every recursive call.

7.4 Offering a convenient View

We can combine `out` and `layer` to obtain the `view` function we used in our introductory examples in Section 1.1. A `(View cs t)` value gives us access to the `(Index cs)` of `t`'s top constructor together with the corresponding `Layer` of deserialised values and pointers to subtrees.

```

805 data View : {cs : Data nm} -> (t : Data.Mu cs) -> Type where
806   (#) : (k : Index cs) ->
807     forall t. Layer (description k) cs t ->
808     View cs (k # t)

```

The implementation of `view` is unsurprising: we use `out` to expose the top constructor index and a `Pointer.Meaning` to the constructor's payload. We then use `layer` to extract the full `Layer` of deserialised values that the pointer references.

```

813 view : {cs : Data nm} ->
814   forall t. Pointer.Mu cs t ->
815   IO (View cs t)
816
817 view ptr = do k # el <- out ptr
818   vs <- layer el
819   pure (k # vs)

```

It is worth noting that although a `view` may be convenient to consume, a performance-minded user may decide to directly use the `out` and `poke` combinators to avoid deserialising values that they do not need. We present a case study in Appendix B comparing the access patterns of two implementations of the function fetching the byte stored in a tree's rightmost node depending on whether we use `view` or the lower level `poke` combinator.

By repeatedly calling `view`, we can define the correct-by-construction generic deserialisation function that turns a pointer to a tree into a runtime value equal to this tree.

```

827 deserialise : {cs : Data nm} -> forall t.
828   Pointer.Mu cs t -> IO (Singleton t)
829

```

We can measure the benefits of our approach by comparing the runtime of a function directly operating on buffers to its pure counterpart composed with a deserialisation step. For functions like `rightmost` that only explore a very small part of the full tree, the gains are spectacular: the

834 process operating on buffers is exponentially faster than its counterpart which needs to deserialise
835 the entire tree first (cf. Section 9).

836

837 7.5 Generic Fold

838 The implementation of the generic `fold` over a tree stored in a buffer is going to have the same
839 structure as the generic fold over inductive values: first match on the top constructor, then use `fmap`
840 to apply the fold to all the substructures and, finally, apply the algebra to the result. We start by
841 implementing the buffer-based counterpart to `fmap`. Let us go through the details of its type first.

842

```
843 fmap : (d : Desc r s o) ->
844       (f : Data.Mu cs -> b) ->
845       (forall t. Pointer.Mu cs t -> IO (Singleton (f t))) ->
846       forall t. Pointer.Meaning d cs t ->
847       IO (Singleton (Data.fmap d f t))
```

848 The first two arguments to `fmap` are similar to its pure counterpart: a description `d` and a (here
849 runtime-irrelevant) function `f` to map over a `Meaning`. Next we take a function which is the buffer-
850 aware counterpart to `f`: given any runtime-irrelevant term `t` and a pointer to it in a buffer, it returns
851 an `IO` process computing the value `(f t)`. Finally, we take a runtime-irrelevant meaning `t` as well
852 as a pointer to its representation in a buffer and compute an `IO` process which will return a value
853 equal to `(Data.fmap d f t)`.

854 We can now look at the definition of `fmap`.

```
855 fmap d f act ptr = poke ptr >>= go d where
856
857 go : (d : Desc{}) -> forall t. Poke d cs t ->
858     IO (Singleton (Data.fmap d f t))
859 go None {t} v = pure byIrrelevance
860 go Byte v = pure v
861 go (Prod d e) (v # w)
862   = do fv <- fmap d f act v
863       fw <- fmap e f act w
864       pure [| fv # fw |]
865 go Rec v = act v
```

866 We poke the buffer to reveal the value the `Pointer.Meaning` named `ptr` is pointing at and then
867 dispatch over the description `d` using the `go` auxiliary function.

868 If the description is `None` we use `byIrrelevance` which happily builds any `(Singleton t)` provided
869 that `t`'s type is proof irrelevant.

870 If the description is `Byte`, the value is left untouched and so we can simply return it immediately.

871 If we have a `Prod` of two descriptions, we recursively apply `fmap` to each of them and pair the
872 results back.

873 Finally, if we have a `Rec` we apply the function operating on buffers that we know performs the
874 same computation as `f`.

875 We can now combine `out` and `fmap` to compute the correct-by-construction `fold`: provided an
876 algebra for a datatype `cs` and a pointer to a tree of type `cs` stored in a buffer, we return an `IO` process
877 computing the fold.

```
878 fold : {cs : Data nm} -> (alg : Alg cs a) ->
879     forall t. Pointer.Mu cs t ->
880     IO (Singleton (Data.fold alg t))
```

881

882

We first use `out` to reveal the constructor choice in the tree's top node, we then recursively apply (`fold alg`) to all the substructures by calling `fmap`, and we conclude by applying the algebra to this result.

```

886 fold alg ptr
887   = do k # t <- out ptr
888       rec <- assert_total (fmap _ _ (fold alg) t)
889       pure (alg k <$> rec)

```

We once again (cf. Section 3.3) had to use `assert_total` because it is not obvious to Idris 2 that `fmap` only uses its argument on subterms. This could have also been avoided by mutually defining `fold` and a specialised version of (`fmap (fold alg)`) at the cost of code duplication and obfuscation. We once again include such a definition in Appendix A.

8 SERIALISING DATA

So far all of our example programs involved taking an inductive value apart and computing a return value in the host language. But we may instead want to compute another value in serialised form. We include below one such example: a `map` function which takes a function `f` acting on bytes and applies it to all of the ones stored in the nodes of our type of `Trees`.

```

902 map : (f : Bits8 -> Bits8) ->
903       (ptr : Pointer.Mu Tree t) ->
904       Serialising Tree (Data.map f t)
905 map f ptr = case !(view ptr) of
906   "Leaf" # () => "Leaf" # ()
907   "Node" # l # b # r => "Node" # map f l # [| f b |] # map f r

```

It calls the `view` we just defined to observe whether the tree is a leaf or a node. If it's a leaf, it returns a leaf. If it's a node, it returns a node where the `map` has been recursively applied to the left and right subtrees while the function `f` has been applied to the byte `b`.

In this section we are going to spell out how we can define high-level constructs allowing users to write these correct-by-construction serialisers.

8.1 The Type of Serialisation Processes

A serialisation process for a tree `t` that belongs to the datatype `cs` is a function that takes a buffer and a starting position and returns an `IO` process that serialises the term in the buffer at that position and computes the position of the first byte past the serialised tree.

```

919 record Serialising (cs : Data nm) (t : Data.Mu cs) where
920   constructor MkSerialising
921   runSerialising : Buffer -> Int -> IO Int

```

We do not expect users to define such processes by hand and in fact prevent them from doing so by not exporting the `MkSerialising` constructor. Instead, we provide high-level, invariant-respecting combinators to safely construct such serialisation processes.

8.2 Building Serialisation Processes

Our main combinator is (`#`): by providing a node's constructor index and a way to serialise all of the node's subtrees, we obtain a serialisation process for said node. We will give a detailed explanation of `All` below.

```

932 (#) : {cs : Data nm} -> (k : Index cs) ->
933     {0 t : Meaning (description k) (Data.Mu cs)} ->
934     All (description k) (Serialising cs) t ->
935     Serialising cs (k # t)

```

936 The keen reader may refer to the accompanying code to see the implementation. Informally (cf.
937 Section 4.2 for the description of the format): first we write the tag corresponding to the choice
938 of constructor, then we leave some space for the offsets, in the meantime we write all of the
939 constructor's arguments and collect the offsets associated to each subtree while doing so, and
940 finally we fill in the space we had left blank with the offsets we have thus collected.

941 The `All` quantifier performs the pointwise lifting of a predicate over the functor described by a
942 `Desc`. It is defined by induction over the description.

```

943 All : (d : Desc r s o) -> (p : x -> Type) -> Meaning d x -> Type
944 All None p t = ()
945 All Byte p t = Singleton t
946 All Rec p t = p t
947 All d@(Prod _ _) p t = All' d p t
948

```

949 If the description is `None` then there is nothing to apply the predicate to and so we return the unit
950 type. If the description is `Byte` we only demand that we have a runtime copy of the byte so that we
951 may write it inside a buffer. This is done using the `Singleton` family discussed in Section 6.2. If the
952 description is `Rec` then we demand that the predicate holds. Finally, if the description is a the `Prod`
953 of two subdescriptions, we once again use an auxiliary family purely for ergonomics. It is defined
954 mutually with `All` and does the expected structural operation.

```

955 data All' : (d : Desc r s o) -> (p : x -> Type) ->
956     Meaning d x -> Type where
957     (#) : All d p t -> All e p u -> All' (Prod d e) p (t # u)
958

```

959 It should now be clear that `(All (description k) (Serialising cs))` indeed corresponds to having
960 already defined a serialisation process for each subtree.

961 This very general combinator should be enough to define all the serialisers we may ever want.
962 By repeatedly pattern-matching on the input tree and using `(#)`, we can for instance define the
963 correct-by-construction generic serialisation function.

```

964 serialise : {cs : Data nm} -> (t : Data.Mu cs) -> Serialising cs t
965

```

966 We nonetheless include other combinators purely for performance reasons.

968 8.3 Copying Entire Trees

969 We introduce a `copy` combinator for trees that we want to serialise as-is and have a pointer for.
970 Equipped with this combinator, we are able to easily write e.g. the `swap` function which takes a
971 binary tree apart and swaps its left and right branches (if the tree is non-empty).

```

972 swap : Pointer.Mu Tree t -> Serialising Tree (Data.swap t)
973 swap ptr = case !(view ptr) of
974     "Leaf" # () => leaf
975     "Node" # l # b # r => node (copy r) b (copy l)
976

```

977 We could define this `copy` combinator at a high level either by composing `deserialise` and
978 `serialise`, or by interleaving calls to `view` and `(#)`. This would however lead to a slow implementa-
979 tion that needs to traverse the entire tree in order to simply copy it.

980

981 Instead, we implement `copy` by using the `copyData` primitive for `Buffers` present in Idris 2's
 982 standard library. This primitive allows us to grab a slice of the source buffer corresponding to the
 983 tree and to copy the raw bytes directly into the target buffer.

```
984 copy : Pointer.Mu cs t -> Serialising cs t
985 copy ptr = MkSerialising $ \ buf, pos => do
986   let size = muSize ptr
987       copyData (muBuffer ptr) (muPosition ptr) size buf pos
988       pure (pos + size)
```

989 This is the one combinator that crucially relies on our format only using offsets and not absolute
 990 addresses and on the accuracy of the size information we have been keeping in `Pointer.Mu` and
 991 `Pointer.Meaning`. As we can see in Section 9, this is spectacularly faster than a deep copying process
 992 traversing the tree.
 993

994 8.4 Executing a Serialisation Action

995 Now that we can describe actions serialising a value to a buffer, the last basic building block we are
 996 still missing is a function actually performing such actions. This is provided by the `execSerialising`
 997 function declared below.
 998

```
999 execSerialising : {cs : Data nm} -> {0 t : Data.Mu cs} ->
1000   Serialising cs t -> IO (Pointer.Mu cs t)
```

1001 By executing a `(Serialising cs t)`, we obtain an `IO` process returning a pointer to the tree `t`
 1002 stored in a buffer. We can then either compute further with this tree (e.g. by calling `sum` on it), or
 1003 write it to a file for safekeeping using the function `writeToFile` introduced in Section 5.2.
 1004

1005 8.5 Evaluation Order

1006 The careful reader may have noticed that we can and do run arbitrary `IO` operations when building
 1007 a value of type `Serialising` (cf. the `map` example in Section 8 where we perform a call to `view` to
 1008 inspect the input's shape).
 1009

1010 This is possible thanks to Idris 2 elaborating `do`-blocks using whichever appropriate bind operator
 1011 is in scope. In particular, we have defined the following one to use when building a serialisation
 1012 process:

```
1013 (>>=) : IO a -> (a -> Serialising cs t) -> Serialising cs t
1014 io >>= f = MkSerialising $ \buf, start =>
1015   do x <- io
1016       runSerialising (f x) buf start
```

1017 By using this bind we can temporarily pause writing to the buffer to make arbitrary `IO` requests
 1018 to the outside world. In particular, this allows us to interleave reading from the original buffer and
 1019 writing into the target one thus having a much better memory footprint than if we were to first use
 1020 the `IO` monad to build in one go the whole serialisation process for a given tree and then execute it.
 1021

1022 9 BENCHMARKS

1023 Now that we have the ability to read, write, and program directly over trees stored in a buffer we
 1024 can run some experiments to see whether this allows us to gain anything over the purely functional
 1025 programming style.
 1026

1027 For all of these tests we generate a full tree of a given depth and compare the time it takes to run
 1028 the composition of deserialising the tree and applying the pure function to the time it takes to run
 1029


its pointer-based counterpart. Each test is run 20 times in a row, and the duration averaged. We manually run `chezscheme`'s garbage collector before the start of each time measurement.

All of our plots use a logarithmic y axis because the runtime of the deserialisation-based function is necessarily exponential in the depth of the full tree.


The `sum` function explores the entirety of the tree and as such the difference between the deserialisation-based and the pointer-based functions is minimal.

Measure memory footprint?

The `rightmost` function only explores the rightmost branch of the tree and we correspondingly see an exponential speedup for the pointer-based function which is able to efficiently skip past every left subtree.



The deep `copy` is unsurprisingly also exponential in the depth of the tree being copied whereas the version based on the `copyData` primitive for buffers is vastly faster.



10 CONCLUSION

We have seen how, using a universe of descriptions indexed by their static and dynamic sizes, we can define a precise language of values serialised in a buffer. This allowed us to develop a library to manipulate such trees in a seamless, correct, and generic manner either using low-level combinators like `poke` or high-level programs like a data-polymorphic `fold`. We then provided users with convenient tools to write serialisation processes thus allowing them to compositionally build correct-by-construction values stored in buffers.

10.1 Related Work

This work sits at the intersection of many domains: data-generic programming, the efficient runtime representation of functional data, programming over serialised values, and the design of serialisation formats. Correspondingly, a lot of related work is worth discussing. In many cases the advantage of our approach is precisely that it is at the intersection of all of these strands of research.

10.1.1 Data-Generic Programming. There is a long tradition of data-generic programming [Gibbons 2006] and we will mostly focus here on the approach based on the careful reification of a precise universe of discourse as an inductive family in a host type theory, and the definition of generic programs by induction over this family.

One early such instance is Pfeifer and Rueß' 'polytypic proof construction' [Pfeifer and Rueß 1999] meant to replace unsafe meta-programs deriving recursors (be they built-in support, or user-written tactics).

In his PhD thesis, Morris [Morris 2007] declares various universes for strictly positive types and families and defines by generic programming further types (the type of one-hole contexts), modalities (the universal and existential predicate lifting over the functors he considers), and functions (map, boolean equality).

Löh and Magalhães [Löh and Magalhães 2011] define a more expressive universe over indexed functors that is closed under composition and fixpoints. They also detail how to define additional generic construction such as a proof of decidable equality, various recursion schemes, and zippers. This work, quite similar to our own in its presentation, offers a natural candidate universe for us to use to extend our library.

10.1.2 Efficient Runtime Representation of Inductive Values. Although not dealing explicitly with programming over serialised data, Monnier's work [Monnier 2019] with its focus on performance and in particular on the layout of inductive values at runtime, partially motivated our endeavour. Provided that we find a way to get the specialisation and partial evaluation of the generically defined views, we ought to be able to achieve –purely in user code– Monnier's vision of a representation where n-ary tuples have constant-time access to each of their component.

10.1.3 Working on Serialised Data. LoCal [Vollmer et al. 2019] is the work that originally motivated the design of this library. We have demonstrated that generic programming within a dependently typed language can yield the sort of benefits other language can only achieve by inventing entirely new intermediate languages and compilation schemes.

LoCal was improved upon with a re-thinking of the serialisation scheme making the approach compatible with parallel programming [Koparkar et al. 2021]. This impressive improvement is a natural candidate for future work on our part: the authors demonstrate it is possible to reap the benefits of both programming over serialised data and dividing up the work over multiple processors with almost no additional cost in the case of a purely sequential execution.

10.1.4 Serialisation Formats. The PADS project [Mandelbaum et al. 2007] aims to let users quickly, correctly, and compositionally describe existing formats they have no control over. As they reminds us, ad-hoc serialisation formats abound be it in networking, logging, billing, or as output of measurement equipments in e.g. gene sequencing or molecular biology. Our current project is not offering this kind of versatility as we have decided to focus on a specific serialisation format with strong guarantees about the efficient access to subtrees. But our approach to defining correct-by-construction components could be leveraged in that setting too and bring users strong guarantees about the traversals they write.

1128 ASN.1 [Larmouth 1999] gives users the ability to define a high-level specification of the exchange
1129 format (the ‘abstract syntax’) to be used in communications without the need to concern themselves
1130 with the actual encoding as bit patterns (the ‘transfer syntax’). This separation between specification
1131 and implementation means that parsing and encoding can be defined once and for all by generic
1132 programming (here, a compiler turning specifications into code in the user’s host language of
1133 choice). The main difference is once again our ability to program in a correct-by-construction
1134 manner over the values thus represented.

1135 Yallop’s automatic derivation of serialisers using an OCaml preprocessor [Yallop 2007] highlights
1136 the importance of empowering domain experts to take advantage of the specifics of the problem
1137 they are solving to minimise the size of the encoded data. By detecting sharing using a custom
1138 equality function respecting α -equivalence instead of the default one, he was able to serialise large
1139 lambda terms using only a quarter of the bytes OCaml’s standard library marshaller.

1140

1141 10.2 Limitations and Future Work

1142 Although our design is already proven to be functional by two implementations in Idris 2 and Agda
1143 respectively, we can always do better. In this section we are going to see what benefits future work
1144 could bring across the whole project.

1145

1146 *10.2.1 A More Robust Library.* For sake of ease of presentation we have not dealt with issues
1147 necessitating buffer resizing: in Section 8, we defined `execSerialising` by allocating a fixed size
1148 buffer and not worrying whether the whole content would fit. A real library would need to adopt a
1149 more robust approach akin to the one used in the implementation of Idris 2’s own serialisation
1150 code: whenever we are about to write a byte to the buffer, we make sure there is either enough
1151 space left or we grow it.

1152 In our library, the data types descriptions currently need to be defined as values in the host
1153 language. This opens up the opportunity for bugs if, say, we write a server in Idris 2 and a client in
1154 Agda and accidentally use two slightly different descriptions in the projects. This could be solved at
1155 the language level by equipping our dependently typed languages with type providers like Idris 1
1156 had [Christiansen 2013]. This way the format could be loaded at compile time from the same file
1157 thus ensuring all the components are referring to the exact same specification.

1158

1159 *10.2.2 A More Efficient Library.* Looking at the code generated by Idris 2, we notice that our generic
1160 programs are not specialised and partially evaluated even when the types they are working on are
1161 statically known. Refactoring the library to use a continuation-passing-style approach does help
1162 the compiler generate slightly more specialised code but the results are in our opinion not good
1163 enough to justify forcing users to program in this more cumbersome style. A possible alternative
1164 would be to present users with macros rather than generic programs so that the partial evaluation
1165 would be guaranteed to happen at typechecking time. This however makes the process of defining
1166 the generic programs much more error prone. A more principled approach would be to extend
1167 Idris 2 with a proper treatment of staging e.g. by using a two-level type theory as suggested by
1168 Kovács [Kovács 2022].

1169 Our serialisation format has been designed to avoid pointer-chasing and thus ensures entire
1170 subtrees can be easily copied by using the raw bytes. Correspondingly it currently does not support
1171 sharing. This could however be a crucial feature for trees with a lot of duplicated nodes and we
1172 would like to allow users to, using the same interface, easily pick between different serialisation
1173 formats so that the library ends up using the one that suits their application best. To this end, we
1174 could take inspiration from Yallop’s definition of preprocessors generating serialisers [Yallop 2007].
1175 It maintains an object map containing the already serialised nodes and uses it to maximally detect
1176 sharing and maintain it both when serialising and deserialising.

1176

1177 Our current approach allows us to define a correct-by-construction `sum` operating directly on
1178 serialised data but it does not eliminate the call stack used in the naïve functional implementation.
1179 Converting a fold to a tail recursive function in a generic manner is a well studied problem
1180 and the existing solutions [McBride 2008; Tomé Cortiñas and Swierstra 2018] should be fairly
1181 straightforward, if time-consuming, to port to our setting.

1182 *10.2.3 A More Expressive Universe of Descriptions.* We have used a minimal universe to demonstrate
1183 our approach but a practical application would require the ability to store more than just raw bytes.
1184 An easy extension is to add support for all of the numeric types of known size that Idris 2 offers
1185 (`Bits{8, 16, 32, 64}`), `Int{8, 16, 32, 64}`), for `Bool` as well as a unbounded data such as `Nat`, or `String`
1186 as long as an extra offset is provided for each value.

1187 The storage of values smaller than a byte (here `Bool`) naturally raises the question of bit packing:
1188 why store eight booleans as eight bytes when they could fit in a single one? Our recent work [Allais
1189 2023] on the efficient runtime representation of inductive families as values of Idris 2's primitive
1190 types points us in the direction of a solution.

1191 A natural next candidate is a universe allowing the definition of parametrised types [Löh and
1192 Magalhães 2011]: we should be able to implement functions over arbitrary (`List a`) values stored
1193 in a buffer, provided that we know that `a` is serialisable. This was already an explicit need in
1194 ASN.1 [Larmouth 1999], reflecting that protocols often leave 'holes' where the content of the
1195 protocol's higher layer is to be inserted.

1196 Next, we will want to consider a universe of indexed data: we can currently natively model
1197 algebraic datatypes such as lists or trees, we can use the host language to compute the description
1198 of vectors by induction on their length, but we cannot model arbitrary type families [Dybjer 1994]
1199 e.g. correct-by-construction red-black trees.

1200 Last but not least we may want to have a universe of descriptions closed under least fixpoints [Mor-
1201 ris 2007] in order to represent rose trees for instance.

1202 *10.2.4 A More Expressive Library.* Using McBride's generalisation of one hole contexts [McBride
1203 2008] we ought to be able to give a more precise type to the combinator (`#`) used to build serialisation
1204 processes. When defining the serialisation of a given subtree, we ought to have access to pointers to
1205 the result of serialising any subtree to the left of it. In particular this would make building complete
1206 binary trees a lot faster by allowing us to rely on `copyData` for duplicating branches rather than
1207 running the computation twice.

1208 Last, but not least we currently do not support in-place updates to the data stored in a buffer. This
1209 could however be beneficial for functions like `map`. It remains to be seen whether we can somehow
1210 leverage Idris 2's linear quantity annotation to provide users with serialised value that can be safely
1211 updated in place. This would turn our ongoing metaphor involving Hoare triples [Hoare 1969],
1212 heap pointers, and separation logic [Reynolds 2002] into a bona fide shallow embedding. Poulsen,
1213 Rouvoet, Tolmach, Krebbers, and Visser's pioneering work [Poulsen et al. 2018; Rouvoet 2021]
1214 on definitional interpreters for languages with references and the use of a shallowly embedded
1215 separation logic to minimise bookkeeping give us a clear set of techniques to adapt to our setting.
1216

1217 ACKNOWLEDGMENTS

1218 We would like to thank Wouter Swierstra for his helpful comments on a draft of this paper.

1219 This research was partially funded by the Engineering and Physical Sciences Research Council
1220 (grant number EP/T007265/1).
1221

1222
1223
1224
1225

REFERENCES

- 1226
1227 Guillaume Allais. 2023. Builtin Types Viewed as Inductive Families. In *Programming Languages and Systems - 32nd European*
1228 *Symposium on Programming, ESOP 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software,*
1229 *ETAPS 2023, Paris, France, April 22-27, 2023, Proceedings (Lecture Notes in Computer Science, Vol. 13990)*, Thomas Wies
1230 (Ed.). Springer, 113–139. https://doi.org/10.1007/978-3-031-30044-8_5
- 1231 Thorsten Altenkirch and Conor McBride. 2002. Generic Programming within Dependently Typed Programming. In *Generic*
1232 *Programming, IFIP TC2/WG2.1 Working Conference on Generic Programming, July 11-12, 2002, Dagstuhl, Germany (IFIP*
1233 *Conference Proceedings, Vol. 243)*, Jeremy Gibbons and Johan Jeuring (Eds.). Kluwer, 1–20.
- 1234 Robert Atkey. 2018. Syntax and Semantics of Quantitative Type Theory. In *Proceedings of the 33rd Annual ACM/IEEE*
1235 *Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, Anuj Dawar and Erich Grädel (Eds.).
1236 ACM, 56–65. <https://doi.org/10.1145/3209108.3209189>
- 1237 Marcin Benke, Peter Dybjer, and Patrik Jansson. 2003. Universes for Generic Programs and Proofs in Dependent Type
1238 Theory. *Nordic J. of Computing* 10, 4 (Dec. 2003), 265–289. <http://dl.acm.org/citation.cfm?id=985799.985801>
- 1239 Edwin C. Brady. 2021. Idris 2: Quantitative Type Theory in Practice. In *35th European Conference on Object-Oriented*
1240 *Programming, ECOOP 2021, July 11-17, 2021, Aarhus, Denmark (Virtual Conference) (LIPIcs, Vol. 194)*, Anders Møller and
1241 Manu Sridharan (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 9:1–9:26. [https://doi.org/10.4230/LIPIcs.](https://doi.org/10.4230/LIPIcs.ECOOP.2021.9)
1242 [ECOOP.2021.9](https://doi.org/10.4230/LIPIcs.ECOOP.2021.9)
- 1243 David Raymond Christiansen. 2013. Dependent type providers. In *Proceedings of the 9th ACM SIGPLAN workshop on Generic*
1244 *programming, WGP 2013, Boston, Massachusetts, USA, September 28, 2013*, Jacques Carette and Jeremiah Willcock (Eds.).
1245 ACM, 25–34. <https://doi.org/10.1145/2502488.2502495>
- 1246 Peter Dybjer. 1994. Inductive Families. *Formal Aspects Comput.* 6, 4 (1994), 440–465. <https://doi.org/10.1007/BF01211308>
- 1247 Jeremy Gibbons. 2006. Datatype-Generic Programming. In *Datatype-Generic Programming - International Spring School,*
1248 *SSDGP 2006, Nottingham, UK, April 24-27, 2006, Revised Lectures (Lecture Notes in Computer Science, Vol. 4719)*, Roland Carl
1249 Backhouse, Jeremy Gibbons, Ralf Hinze, and Johan Jeuring (Eds.). Springer, 1–71. [https://doi.org/10.1007/978-3-540-](https://doi.org/10.1007/978-3-540-76786-2_1)
1250 [76786-2_1](https://doi.org/10.1007/978-3-540-76786-2_1)
- 1251 Tatsuya Hagino. 1987. A Typed Lambda Calculus with Categorical Type Constructors. In *Category Theory and Computer*
1252 *Science, Edinburgh, UK, September 7-9, 1987, Proceedings (Lecture Notes in Computer Science, Vol. 283)*, David H. Pitt, Axel
1253 Poigné, and David E. Rydeheard (Eds.). Springer, 140–157. https://doi.org/10.1007/3-540-18508-9_24
- 1254 C. A. R. Hoare. 1969. An Axiomatic Basis for Computer Programming. *Commun. ACM* 12, 10 (1969), 576–580. <https://doi.org/10.1145/363235.363259>
- 1255 Paul Hudak. 1996. Building Domain-Specific Embedded Languages. *ACM Comput. Surv.* 28, 4es (1996), 196. <https://doi.org/10.1145/242224.242477>
- 1256 Patrik Jansson and Johan Jeuring. 1997. Polyp - A Polytypic Programming Language. In *Conference Record of POPL '97: The*
1257 *24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at the Symposium,*
1258 *Paris, France, 15-17 January 1997*, Peter Lee, Fritz Henglein, and Neil D. Jones (Eds.). ACM Press, 470–482. <https://doi.org/10.1145/263699.263763>
- 1259 Chaitanya Koparkar, Mike Rainey, Michael Vollmer, Milind Kulkarni, and Ryan R. Newton. 2021. Efficient tree-traversals:
1260 reconciling parallelism and dense data representations. *Proc. ACM Program. Lang.* 5, ICFP (2021), 1–29. <https://doi.org/10.1145/3473596>
- 1261 András Kovács. 2022. Staged compilation with two-level type theory. *Proc. ACM Program. Lang.* 6, ICFP (2022), 540–569.
1262 <https://doi.org/10.1145/3547641>
- 1263 John Larmouth. 1999. *ASN.1 Complete*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- 1264 Andres Löh and José Pedro Magalhães. 2011. Generic programming with indexed functors. In *Proceedings of the seventh*
1265 *ACM SIGPLAN workshop on Generic programming, WGP@ICFP 2011, Tokyo, Japan, September 19-21, 2011*, Jaakko Järvi
1266 and Shin-Cheng Mu (Eds.). ACM, 1–12. <https://doi.org/10.1145/2036918.2036920>
- 1267 Grant Malcolm. 1990. Data Structures and Program Transformation. *Sci. Comput. Program.* 14, 2-3 (1990), 255–279.
1268 [https://doi.org/10.1016/0167-6423\(90\)90023-7](https://doi.org/10.1016/0167-6423(90)90023-7)
- 1269 Yitzhak Mandelbaum, Kathleen Fisher, David Walker, Mary F. Fernández, and Artem Gleyzer. 2007. PADS/ML: a functional
1270 data description language. In *Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
1271 *Languages, POPL 2007, Nice, France, January 17-19, 2007*, Martin Hofmann and Matthias Felleisen (Eds.). ACM, 77–83.
1272 <https://doi.org/10.1145/1190216.1190231>
- 1273 Conor McBride. 2008. Clowns to the left of me, jokers to the right (pearl): dissecting data structures. In *Proceedings of the 35th*
1274 *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA,*
January 7-12, 2008, George C. Necula and Philip Wadler (Eds.). ACM, 287–295. <https://doi.org/10.1145/1328438.1328474>
- 1275 Conor McBride. 2016. I Got Plenty o' Nuttin'. In *A List of Successes That Can Change the World - Essays Dedicated to Philip*
1276 *Wadler on the Occasion of His 60th Birthday (Lecture Notes in Computer Science, Vol. 9600)*, Sam Lindley, Conor McBride,
1277 Philip W. Trinder, and Donald Sannella (Eds.). Springer, 207–233. https://doi.org/10.1007/978-3-319-30936-1_12

- 1275 Conor McBride and James McKinna. 2004. The view from the left. *J. Funct. Program.* 14, 1 (2004), 69–111. <https://doi.org/10.1017/S0956796803004829>
- 1276
- 1277 Conor McBride and Ross Paterson. 2008. Applicative programming with effects. *J. Funct. Program.* 18, 1 (2008), 1–13. <https://doi.org/10.1017/S0956796807006326>
- 1278
- 1279 Gavin Mendel-Gleason. 2012. *Types and verification for infinite state systems*. Ph. D. Dissertation. Dublin City University.
- 1280 Stefan Monnier. 2019. Inductive types deconstructed: the calculus of united constructions. In *Proceedings of the 4th ACM SIGPLAN International Workshop on Type-Driven Development, TyDe@ICFP 2019, Berlin, Germany, August 18, 2019*, David Darais and Jeremy Gibbons (Eds.). ACM, 52–63. <https://doi.org/10.1145/3331554.3342607>
- 1281
- 1282 Peter W. J. Morris. 2007. *Constructing Universes for Generic Programming*. Ph. D. Dissertation. University of Nottingham, UK. <https://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.519405>
- 1283
- 1284 Holger Pfeifer and Harald Rueß. 1999. Polytypic Proof Construction. In *Theorem Proving in Higher Order Logics, 12th International Conference, TPHOLS'99, Nice, France, September, 1999, Proceedings (Lecture Notes in Computer Science, Vol. 1690)*, Yves Bertot, Gilles Dowek, André Hirschowitz, Christine Paulin-Mohring, and Laurent Théry (Eds.). Springer, 55–72. https://doi.org/10.1007/3-540-48256-3_5
- 1285
- 1286 Casper Bach Poulsen, Arjen Rouvoet, Andrew Tolmach, Robbert Krebbers, and Eelco Visser. 2018. Intrinsically-typed definitional interpreters for imperative languages. *Proc. ACM Program. Lang.* 2, POPL (2018), 16:1–16:34. <https://doi.org/10.1145/3158104>
- 1287
- 1288
- 1289 John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*. IEEE Computer Society, 55–74. <https://doi.org/10.1109/LICS.2002.1029817>
- 1290
- 1291
- 1292 Arjen Rouvoet. 2021. *Correct by Construction Language Implementations*. Ph. D. Dissertation. Delft University of Technology, Netherlands. <https://doi.org/10.4233/uuid:f0312839-3444-41ee-9313-b07b21b59c11>
- 1293
- 1294 Carlos Tomé Cortiñas and Wouter Swierstra. 2018. From algebra to abstract machine: a verified generic construction. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Type-Driven Development, TyDe@ICFP 2018, St. Louis, MO, USA, September 27, 2018*, Richard A. Eisenberg and Niki Vazou (Eds.). ACM, 78–90. <https://doi.org/10.1145/3240719.3241787>
- 1295
- 1296 Michael Vollmer, Chaitanya Koparkar, Mike Rainey, Laith Sakka, Milind Kulkarni, and Ryan R. Newton. 2019. LoCal: a language for programs operating on serialized data. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, Kathryn S. McKinley and Kathleen Fisher (Eds.). ACM, 48–62. <https://doi.org/10.1145/3314221.3314631>
- 1297
- 1298
- 1299 Philip Wadler. 1987. Views: A Way for Pattern Matching to Cohabit with Data Abstraction. In *Conference Record of the Fourteenth Annual ACM Symposium on Principles of Programming Languages, Munich, Germany, January 21-23, 1987*. ACM Press, 307–313. <https://doi.org/10.1145/41625.41653>
- 1300
- 1301
- 1302 Jeremy Yallop. 2007. Practical generic programming in OCaml. In *Proceedings of the ACM Workshop on ML, 2007, Freiburg, Germany, October 5, 2007*, Claudio V. Russo and Derek Dreyer (Eds.). ACM, 83–94. <https://doi.org/10.1145/1292535.1292548>
- 1303
- 1304
- 1305
- 1306
- 1307
- 1308
- 1309
- 1310
- 1311
- 1312
- 1313
- 1314
- 1315
- 1316
- 1317
- 1318
- 1319
- 1320
- 1321
- 1322
- 1323

1324 A SAFE IMPLEMENTATIONS OF FOLD

1325 We include below the alternative definitions of `fold` (respectively processing inductive data and
1326 data stored in a buffer) which are seen as total by Idris 2. Each of them is mutually defined with
1327 what is essentially the supercompilation of $(\lambda d \Rightarrow \text{fmap } d (\text{fold } \text{alg}))$.

```
1328 parameters {cs : Data nm} (alg : Alg cs a)
1329
1330 fold : Data.Mu cs -> a
1331 fmapFold : (d : Desc{}) ->
1332           Data.Meaning d (Data.Mu cs) -> Data.Meaning d a
1333
1334 fold (k # t) = alg k (fmapFold (description k) t)
1335
1336 fmapFold None t = t
1337 fmapFold Byte t = t
1338 fmapFold (Prod d e) (s # t)
1339   = (fmapFold d s # fmapFold e t)
1340 fmapFold Rec t = fold t
1341
1342 parameters {cs : Data nm} (alg : Alg cs a)
1343
1344 fold : Pointer.Mu cs t -> IO (Singleton (fold alg t))
1345 fmapFold : (d : Desc{}) -> forall t. Pointer.Meaning d cs t ->
1346           IO (Singleton (fmapFold alg d t))
1347
1348 fold ptr
1349   = do k # t <- out ptr
1350       rec <- fmapFold (description k) t
1351       pure (alg k <$> rec)
1352
1353 fmapFold d ptr = poke ptr >>= go d where
1354
1355 go : (d : Desc{}) -> forall t. Poke d cs t ->
1356     IO (Singleton (fmapFold alg d t))
1357 go None {t} v = rewrite etaUnit t in pure (pure ())
1358 go Byte v = pure v
1359 go (Prod d e) (v # w)
1360   = do v <- fmapFold d v
1361       w <- fmapFold e w
1362       pure [| v # w |]
1363 go Rec v = fold v
```

1364 B ACCESS PATTERNS: VIEWING VS. POKING

1365 In this example we implement `rightmost`, the function walking down the rightmost branch of our
1366 type of binary trees and returning the content of its rightmost node (if it exists).

1367 The first implementation is the most straightforward: use `view` to obtain the top constructor
1368 as well as an entire layer of deserialised values and pointers to substructures and inspect the
1369 constructor. If we have a leaf then there is no byte to return. If we have a node then call `rightmost`
1370 recursively and inspect the result: if we got `Nothing` back we are at the rightmost node and can
1371 return the current byte, otherwise simply propagate the result.

1372

```

1373 rightmost : Pointer.Mu Tree t -> IO (Maybe Bits8)
1374 rightmost ptr = case !(view ptr) of
1375   "Leaf" # _ => pure Nothing
1376   "Node" # _ # b # r => do
1377     mval <- rightmost r
1378     case mval of
1379       Just _ => pure mval
1380       Nothing => pure (Just (getSingleton b))

```

In the alternative implementation we use `out` to expose the top constructor and then, in the node case, call `poke` multiple times to get our hands on the pointer to the right subtree. We inspect the result of recursively calling `rightmost` on this subtree and only deserialise the byte contained in the current node if the result we get back is `Nothing`.

```

1385 rightmost : Pointer.Mu Tree t -> IO (Maybe Bits8)
1386 rightmost ptr = case !(out ptr) of
1387   "Leaf" # _ => pure Nothing
1388   "Node" # e1 => do
1389     (_ # br) <- poke e1
1390     (b # r) <- poke br
1391     mval <- rightmost !(poke r)
1392     case mval of
1393       Just _ => pure mval
1394       Nothing => do
1395         b <- poke b
1396         pure (Just (getSingleton b))

```

This will give rise to two different access patterns: the first function will have deserialised all of the bytes stored in the nodes along the tree's rightmost path whereas the second will only have deserialised the rightmost byte. Admittedly deserialising a byte is not extremely expensive but in a more realistic example we could have for instance been storing arbitrarily large values in these nodes. In that case it may be worth trading convenience for making sure we are not doing any unnecessary work.

1403 C ANALOGY TO SEPARATION LOGIC

1405 Some readers may feel uneasy about the fact that parts of our library are implemented using Idris 2
 1406 escape hatches. This section justifies this practice by drawing an analogy to separation logic and
 1407 highlighting that this practice corresponds to giving an axiomatisation of the runtime behaviour of
 1408 our library's core functions.

1410 C.1 Interlude: Separation Logic

1411 A Hoare triple [Hoare 1969] of the form

$$1412 \{ P \} e \{ v. Q \}$$

1415 states that under the precondition P , and binding the result of evaluating the expression e as v , we
 1416 can prove that Q holds. One of the basic predicates of separation logic [Reynolds 2002] is a 'points
 1417 to' assertion ($\ell \mapsto t$) stating that the label ℓ points to a memory location containing t .

1418 A separation logic proof system then typically consists in defining a language and providing
 1419 axioms characterising the behaviour of each language construct. The simplest example involving
 1420 memory is perhaps a language with pointers to bytes and a single `deref` construct dereferencing a

1422 pointer. We can then give the following axiom

$$1423 \quad \{ \ell \mapsto bs \} \text{deref } \ell \{ v. bs = v * \ell \mapsto bs \}$$

1425 to characterise deref by stating that the value it returns is precisely the one the pointer is ref-
1426 erencing, and that the pointer is still valid and still referencing the same value after it has been
1427 dereferenced.

1428 The axioms can be combined to prove statements about more complex programs such as the
1429 following silly one for instance. Here we state that if we dereference the pointer a first time, discard
1430 the result and then dereference it once more then we end up in the same situation as if we had
1431 dereferenced it only once.

$$1432 \quad \{ \ell \mapsto bs \} \text{deref } \ell; \text{deref } \ell \{ v. bs = v * \ell \mapsto bs \}$$

1433 Note that in all of these rules bs is only present in the specification layer. deref itself cannot
1434 possibly return bs directly, it needs to actually perform an effectful operation that will read the
1435 memory cell's content.

1439 C.2 Characterising Our Library

1440 We are going to explain that we can see our library as a small embedded Domain Specific Language
1441 (eDSL) [Hudak 1996] that has `poke` and `out` as sole language constructs. Our main departure
1442 from separation logic is that we want to program in a correct-by-construction fashion and so
1443 the types of `poke` and `out` have to be just as informative as the axioms we would postulate in
1444 separation logic. This dual status of the basic building blocks being both executable programs *and*
1445 an axiomatic specification of their respective behaviour is precisely why their implementations in
1446 Idris 2 necessarily uses unsafe features.

1447 We are going to write $\ell \xrightarrow{\llbracket d \rrbracket (\mu cs)} t$ for the assumption that we own a pointer ℓ of type
1448 (`Pointer.Meaning d cs t`), and $\ell \xrightarrow{\mu cs} t$ for the assumption that we own a pointer ℓ of type (`Pointer.Mu`
1449 `cs t`). In case we do not care about the type of the pointer at hand (e.g. because it can be easily
1450 inferred from the context), we will simply write $\ell \mapsto t$.

1451 **C.2.1 Axioms for `poke`.** Thinking in terms of Hoare triples, if we have a pointer ℓ to a term t known
1452 to be a single byte then (`poke` ℓ) will return a byte bs and allow us to observe that t is equal to that
1453 byte.

$$1454 \quad \{ \ell \xrightarrow{\llbracket \text{Byte} \rrbracket (\mu cs)} t \} \text{poke } \ell \{ bs. t = bs * \ell \mapsto t \}$$

1455 Similarly, if the pointer ℓ is for a pair then (`poke` ℓ) will reveal that the term t can be taken apart
1456 into the pairing of two terms t_1 and t_2 and return a pointer for each of these components.

$$1457 \quad \{ \ell \xrightarrow{\llbracket \text{Prod } d_1 d_2 \rrbracket (\mu cs)} t \}$$

1458 `poke` ℓ

$$1459 \quad \{ (\ell_1, \ell_2). \exists t_1. \exists t_2. t = (t_1 \# t_2) * \ell \mapsto t * \ell_1 \xrightarrow{\llbracket d_1 \rrbracket (\mu cs)} t_1 * \ell_2 \xrightarrow{\llbracket d_2 \rrbracket (\mu cs)} t_2 \}$$

1460 Last but not least, poking a pointer with the `Rec` description will return another pointer for the
1461 same value but at a different type.

$$1462 \quad \{ \ell \xrightarrow{\llbracket \text{Rec} \rrbracket (\mu cs)} t \} \text{poke } \ell \{ \ell_1. \ell_1 \xrightarrow{\mu cs} t * \ell \mapsto t \}$$

1471 *C.2.2 Example of a Derived Rule for layer.* Given that `layer` is defined in terms of `poke`, we do not
 1472 need to postulate any axioms to characterise it and can instead prove lemmas. We will skip the
 1473 proofs here but give an example of a derived rule. Using the description (`Prod Rec (Prod Byte Rec)`)
 1474 of the arguments to a node in our running example of binary trees, `layer`'s behaviour would be
 1475 characterised by the following statement.

$$1476 \quad \{ \ell \xrightarrow{\llbracket \text{Prod Rec (Prod Byte Rec)} \rrbracket (\mu cs)} t \}$$

$$1477 \quad \text{layer } \ell$$

$$1478 \quad \{ (\ell_1, bs, \ell_2). \exists t_1. \exists t_2. t = (t_1 \# bs \# t_2) * \ell_1 \xrightarrow{\mu cs} t_1 * \ell_2 \xrightarrow{\mu cs} t_2 * \ell \leftrightarrow t \}$$

1481 It states that provided a pointer to such a meaning, calling `layer` would return a triple of a pointer
 1482 ℓ_1 for the left subtree, the byte bs stored in the node, and a pointer ℓ_2 for the right subtree.

1483 *C.2.3 Axiom for out.* The only other construct for our small DSL is the function `out`. Things are a
 1484 lot simpler here as the return type is not defined by induction on the description. As a consequence
 1485 we only need the following axiom.

$$1486 \quad \{ \ell \xrightarrow{\mu cs} t \} \text{out } \ell \{ (k, \ell_1). \exists t_1. t = (k \# t_1) * \ell \leftrightarrow t * \ell_1 \xrightarrow{\llbracket cs_k \rrbracket (\mu cs)} t_1 \}$$

1487 It states that under the condition that ℓ points to t , `(out ℓ)` returns a pair of an index and a pointer
 1488 to the meaning of the description associated to that index by cs , and allows us to learn that t is
 1489 constructed using that index and that meaning.

1490 By combining `out` and `layer` we could once more define a derived rule and prove e.g. that every
 1491 tree can be taken apart as either a leaf or a node with a pointer to a left subtree, a byte, and a
 1492 pointer to a right subtree i.e. what `view` does in our library.

1493 Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519