

Programming with Ornaments

HSIANG-SHANG KO

National Institute of Informatics, Japan

JEREMY GIBBONS

University of Oxford, UK

Abstract

Dependently typed programming advocates the use of various indexed versions of the same shape of data, but the formal relationship among these structurally similar datatypes usually needs to be established manually and tediously. Ornaments have been proposed as a formal mechanism to manage the relationships between such datatype variants. In this paper, we conduct a case study under an ornament framework; the case study concerns programming binomial heaps and their operations — including insertion and minimum extraction — by viewing them as lifted versions of binary numbers and numeric operations. We show how current dependently typed programming technology can lead to a clean treatment of the binomial heap constraints when implementing heap operations. We also identify some gaps between the current technology and an ideal dependently typed programming language that we would wish to have for our development.

1 Introduction

Dependently typed programming is characterised by the use of more informative types, in particular inductive families (Dybjer, 1994) — or indexed datatypes — to guarantee program correctness by construction. Having various versions of the same kind of data, however, causes a management problem. For example, there are various kinds of list datatype, such as vectors (with embedded length constraints) and ordered lists; but formally these list datatypes and their operations are unrelated, and whenever a new variant of lists is invented, similar operations for the new variant must be developed from scratch, along with conversions to and from existing list-like datatypes. This work is tedious, and can make the programmer hesitant to employ informative types so as to avoid the hassle of managing them.

To address the variant-management problem, McBride (2011) proposed *ornaments* as a framework for relating structurally similar datatypes. An ornament specifies a particular way of relating two datatype definitions intensionally, and stating that one of them is more informative than the other; from an ornament, we can compute a forgetful function converting the more informative datatype to the less informative one. McBride also proposed *algebraic ornaments*, which provide a systematic way to synthesise new datatypes by adding indices to existing ones. McBride's work spawned a number of subsequent developments:

The ornament-induced forgetful function was extended by Ko & Gibbons (2013a) to an isomorphism for promoting less informative datatypes to more informative ones; at the heart of the construction of the promotion isomorphism is a *parallel composition* operation on ornaments that, in general, can be used to combine multiple constraints on a datatype and synthesise a new variant of the datatype with all those constraints embedded at once. Dagand & McBride (2014) then proposed *functional ornaments*, generalising ornaments to work also for functions. Algebraic ornamentation has also been generalised to a relational setting, and shown to work with relational calculation techniques (Ko & Gibbons, 2013b).

In this paper, we set out to do an experiment to evaluate how effective the current ornament framework is in assisting the development of dependently typed programs, by working out a small yet complete library of binomial heaps in AGDA (Norell, 2007; Norell, 2009; Bove & Dybjer, 2009). Our experiment is tied to AGDA to some extent, as we employ AGDA-specific features like interactive programming support, pattern synonyms, and instance arguments (Devriese & Piessens, 2011), but we believe that the experiment has broader applicability since AGDA is representative of many dependently typed languages. The case study is based on Okasaki’s (1999) idea of numerical representations, lifting operations on binary numbers such as increment and decrement to operations on binomial heaps such as insertion and extraction — this lifting is well supported by the ornament framework. Interestingly, the numerically developed extraction operation leads us to an algorithm for minimum extraction that is conceptually simpler than the standard one.

Before presenting the case study in Section 5, we will first give a tour of related constructions used in this paper, including index-first inductive families and their encodings in Section 2, ornaments for relating index-first inductive families in Section 3, and an extensible function-promoting mechanism that works in conjunction with ornaments in Section 4. We gloss over many technical details about the theory of ornaments and about alternative universes for them, since we intend this paper to provide a more operational/pragmatic understanding of ornaments without touching their theoretical side; more comprehensive discussions of such details can be found in the first author’s DPhil thesis (Ko, 2014). Throughout the paper, there are special *Gap* paragraphs that suggest what still needs to be developed in the ornament framework or for dependently typed programming languages in general so as to give better support for dependently typed programming; Section 6 will give a summary of these gaps, and, along with other remarks, concludes this paper. Most of the code presented in this paper is collected in two supplementary AGDA source files, which have been checked with AGDA version 2.5.1.1 and standard library version 0.12.

2 Descriptions: encoding index-first inductive families

We first look at *index-first inductive families* (Chapman *et al.*, 2010), which are a variant of inductive families (Dybjer, 1994) that can naturally yield efficient representations. This ability will turn out to be important for the ornamental constructions in Section 3.3 to be useful in practice.

2.1 Index-first inductive families

A good starting point for understanding index-first inductive families is a comparison between the usual inductive families (as supported by AGDA) and inductively defined type-computing functions. Consider vectors (length-indexed lists). In AGDA there are two approaches to defining vectors, one as an inductive family:

```
data Vec (A : Set) : Nat → Set where
  []   : Vec A zero
  _::_ : A → {n : Nat} → Vec A n → Vec A (suc n)
```

and the other as a type-computing function:

```
Vec : (A : Set) → Nat → Set
Vec A zero   = ⊤
Vec A (suc n) = A × Vec A n
```

In the first definition, choice of constructor determines the index in the type of the constructed vector, whereas in the second definition, the index determines which constructors are valid (i.e., whether we can construct an empty vector of type \top or a nonempty one of type $A \times \text{Vec } A \ n$). The two approaches have their own pros and cons:

- The inductive family comes with an induction principle, which, roughly speaking, means that in AGDA the programmer can perform dependent pattern matching on vectors, making some subsequent programs like vector append as natural as their list counterparts. Without any optimisations (like those proposed by Brady *et al.* (2004)), however, the inductive family leads to an inefficient representation of vectors, having to store the constructor choices (whether a vector starts with nil or cons) and also the length of the tail in every cons constructor, despite the fact that the constructor choices and the tail lengths are determined by the indices.
- In contrast, the type-computing function yields an (almost) optimal representation of vectors — a vector of length n is an n -tuple (ending with tt) and does not need to store constructor choices and tail lengths. Most interestingly, the optimisation of representation is achieved by directly expressing how the structure of a vector depends on its length, i.e., the index in its type (instead of relying on any compiler optimisations): If the index is zero, the vector must be empty; if the index is $\text{suc } n$, the vector must have a head element and a tail indexed by n . The approach has two drawbacks, though: This definition of vectors is dependent on the length, so whenever we want to analyse a vector we need to analyse its length first, making vector programs more verbose to write. In other words, the vectors do not come with their own induction principle, and must be analysed by induction on the natural number index. More seriously, this approach works only for situations where the type indices grow strictly as the data grow; otherwise the recursive definition is not well-founded.

Index-first inductive families can be regarded as reconciling type-computing functions with the usual inductive families, allowing the index of a type to influence what information can be provided to construct an element of that type, and at the same time removing the

restriction that the definition should be inductive on the index. Dagand & McBride (2012; 2014) proposed a new syntax for datatype declarations mimicking type-computing functions, with which we can (choose to) do case analysis on the index and decide what constructor(s) are available for each case. For example, they would define vectors as an index-first inductive family by

```
data Vec [A : Set] (n : Nat) : Set where
  VecA n      ⇐ Nat-case n
  VecA 0      ∋ nil
  VecA (suc m) ∋ cons (a : A) (vs : VecA m)
```

This syntax strongly depends on EPIGRAM’s pattern matching mechanism (McBride & McKinna, 2004). In this paper we informally adopt a more AGDA-like variation. For example, we describe index-first vectors in the following syntax:

```
indexfirst data Vec (A : Set) : Nat → Set where
  Vec A zero   ∋ []
  Vec A (suc n) ∋ _::_ A (Vec A n)
```

We start the definition with the keyword **indexfirst** to mark explicitly that this is not valid AGDA. The definition says that a $\text{Vec } A \text{ zero}$ can be constructed only by $[]$, while a $\text{Vec } A \text{ (suc } n)$ can be constructed only by $_{_}::__{_}$, which takes two arguments of types A and $\text{Vec } A \text{ } n$. In effect, this is the same as the type-computing function, and does not really show the full power of index-first inductive families. A more compelling example is the type family of λ -terms indexed with the number of free variables:

```
indexfirst data Term : Nat → Set where
  Term n ∋ var (Fin n)
        | lam (Term (suc n))
        | app (Term n) (Term n)
```

This definition is valid even though n never decreases. In contrast, the equivalent type-computing function

```
Term : Nat → Set
Term n = Fin n ⊔ Term (suc n) ⊔ (Term n × Term n)
```

is obviously non-terminating, and would be rejected by AGDA.

2.2 Universe construction

We have not really explained what index-first inductive families are; below we will do so by constructing a model of index-first inductive families with the help of a *universe* (Martin-Löf, 1984). This makes it possible to carry out (part of) our experiment with index-first datatypes without extending AGDA, but ultimately we would hope that such datatypes can be natively supported.

The standard way of constructing inductive families is to take the least fixed point of base functors of type $(I \rightarrow \text{Set}) \rightarrow I \rightarrow \text{Set}$, where $I : \text{Set}$ is the index set. For example, the base functor we use for index-first vectors (parametrised by the element type) is

$$\begin{aligned}
\mathit{Vec}F &: \mathit{Set} \rightarrow (\mathit{Nat} \rightarrow \mathit{Set}) \rightarrow \mathit{Nat} \rightarrow \mathit{Set} \\
\mathit{Vec}F A X \mathit{zero} &= \top \\
\mathit{Vec}F A X (\mathit{suc } n) &= A \times X n
\end{aligned}$$

This closely resembles the type-computing function except that the recursive reference to $\mathit{Vec } A$ is replaced by an invocation of the extra argument X . $\mathit{Vec } A$ is then defined as the least fixed point of $\mathit{Vec}F A$, implying that $\mathit{Vec } A$ and $\mathit{Vec}F A (\mathit{Vec } A)$ are isomorphic. In particular, $\mathit{Vec } A \mathit{zero}$ is isomorphic to $\mathit{Vec}F A (\mathit{Vec } A) \mathit{zero} = \top$, and $\mathit{Vec } A (\mathit{suc } n)$ is isomorphic to $\mathit{Vec}F A (\mathit{Vec } A) (\mathit{suc } n) = A \times \mathit{Vec } A n$, exactly corresponding to how we defined index-first vectors previously.

Not all functions of type $(I \rightarrow \mathit{Set}) \rightarrow I \rightarrow \mathit{Set}$ are valid base functors, and we will constrain ourselves by fixing the ways in which functions of type $(I \rightarrow \mathit{Set}) \rightarrow I \rightarrow \mathit{Set}$ can be constructed. This is where universe construction can help: By defining a datatype whose elements encode ways to construct valid base functors, and later quantifying over this datatype, we will be able to carry out our subsequent constructions for only the encoded base functors. For index-first inductive families, it turns out that we can encode just the “response” part of base functors: Using the base functor for vectors as an example, after swapping the index and recursive reference arguments,

$$\begin{aligned}
\mathit{Vec}F' &: \mathit{Set} \rightarrow \mathit{Nat} \rightarrow (\mathit{Nat} \rightarrow \mathit{Set}) \rightarrow \mathit{Set} \\
\mathit{Vec}F' A \mathit{zero} &= \lambda X \mapsto \top \\
\mathit{Vec}F' A (\mathit{suc } n) &= \lambda X \mapsto A \times X n
\end{aligned}$$

the parts on the right-hand side of the equations are referred to as “responses”, suggesting an interactive interpretation of base functors — $\mathit{Vec}F' A$ receives and analyses an index before it responds with what information is required to construct an element whose type has the given index. We will shortly define a universe RDesc whose elements — called *response descriptions* — encode a range of such responses:

$$\begin{aligned}
\mathit{RDesc} &: \mathit{Set} \rightarrow \mathit{Set}_1 \\
\llbracket - \rrbracket &: \{I : \mathit{Set}\} \rightarrow \mathit{RDesc } I \rightarrow (I \rightarrow \mathit{Set}) \rightarrow \mathit{Set}
\end{aligned}$$

Base functors will then be encoded by

$$\begin{aligned}
\mathit{Desc} &: \mathit{Set} \rightarrow \mathit{Set}_1 \\
\mathit{Desc } I &= I \rightarrow \mathit{RDesc } I \\
\mathbb{F} &: \{I : \mathit{Set}\} \rightarrow \mathit{Desc } I \rightarrow (I \rightarrow \mathit{Set}) \rightarrow (I \rightarrow \mathit{Set}) \\
\mathbb{F} D X i &= \llbracket D i \rrbracket X
\end{aligned}$$

and we can take the least fixed point of the encoded base functors by

$$\begin{aligned}
&\mathbf{data } \mu \{I : \mathit{Set}\} (D : \mathit{Desc } I) : I \rightarrow \mathit{Set} \mathbf{where} \\
&\quad \mathit{con} : \mathbb{F} D (\mu D) \Rightarrow \mu D
\end{aligned}$$

where $X \Rightarrow Y$ abbreviates $\{i : I\} \rightarrow X i \rightarrow Y i$.

The datatype of response descriptions and its decoding function determine the range of base functors we can describe. For this paper, we focus on inductive families that are finitely branching, captured by the following definition:

data RDesc ($I : \text{Set}$) : Set_1 **where**
 $v : (is : \text{List } I) \rightarrow \text{RDesc } I$
 $\sigma : (S : \text{Set}) (D : S \rightarrow \text{RDesc } I) \rightarrow \text{RDesc } I$
 $\llbracket _ \rrbracket : \{I : \text{Set}\} \rightarrow \text{RDesc } I \rightarrow (I \rightarrow \text{Set}) \rightarrow \text{Set}$
 $\llbracket v \ is \ \rrbracket X = \mathbb{P} \ is \ X$ -- see below
 $\llbracket \sigma \ S \ D \rrbracket X = \Sigma [s \in S] \llbracket D \ s \rrbracket X$

The operator \mathbb{P} computes the product of a finite number of types in a type family, whose indices are given in a list:

$\mathbb{P} : \{I : \text{Set}\} \rightarrow \text{List } I \rightarrow (I \rightarrow \text{Set}) \rightarrow \text{Set}$
 $\mathbb{P} [] \quad X = \top$
 $\mathbb{P} (i :: is) X = X \ i \times \mathbb{P} \ is \ X$

Thus, in a response, given $X : I \rightarrow \text{Set}$, we are allowed to form dependent sums (by σ) and the product of a finite number of types in X (via v , suggesting variable positions in the base functor). We will informally refer to the index part of a σ as a *field* of the datatype. For convenience, we define an abbreviation

syntax $\sigma \ S \ (\lambda s \mapsto D) = \sigma [s \in S] \ D$

so as to use σ as a binder.

2.3 Encoding datatypes as descriptions

Let us now look at several kinds of datatype and how they can be encoded as descriptions.

Non-indexed datatypes. A non-indexed datatype can be seen as an inductive family trivially indexed by \top , and hence can be encoded as an inhabitant of $\text{Desc } \top$. For example, the datatype of natural numbers

indexfirst data Nat : Set **where**
 Nat \ni zero
 | suc Nat

can be encoded as follows:

data ListTag : Set **where**
 'nil : ListTag
 'cons : ListTag
 NatD : Desc \top
 NatD tt = σ ListTag λ { 'nil $\mapsto v []$
 ; 'cons $\mapsto v (tt :: [])$ }

The index request is necessarily tt, and we respond with a field of type ListTag representing the constructor choices. (Any two-element set can be used to represent the constructor choices; the reason that we choose the names ListTag, 'nil, and 'cons will be clear when we relate natural numbers with lists in Section 3.) A pattern-matching lambda function follows,

which computes the trailing responses to the two possible values ‘nil and ‘cons for the field: If the field receives ‘nil, then we are constructing zero, which takes no recursive values, so we write $v []$ to end this branch; if the ListTag field receives ‘cons, then we are constructing a successor, which takes a recursive value at index tt, so we write $v (tt :: [])$.

To see that we have indeed defined the type of natural numbers, we can try to construct an inhabitant of $\mu NatD \text{ tt}$ interactively:

$$\begin{aligned} n &: \mu NatD \text{ tt} \\ n &= \{\mu NatD \text{ tt}\}_0 \end{aligned}$$

The shaded box is called an *interaction point*, or informally a *hole* in the program. Shown in the hole is the goal type, for which we have to construct an appropriate value. The subscript number is used to disambiguate when there are multiple holes.

To construct a μ -value, a natural — indeed, a forced — choice is to use the only constructor con:

$$\begin{aligned} n &: \mu NatD \text{ tt} \\ n &= \text{con } \{\mathbb{F} NatD (\mu NatD \text{ tt})\}_1 \end{aligned}$$

The goal type computes to $[\text{NatD tt}] (\mu NatD)$, which, if we compute under lambda, normalises to

$$\begin{aligned} \Sigma \text{ListTag } \lambda \{ & \text{'nil} \mapsto \top \\ & ; \text{'cons} \mapsto \mu NatD \text{ tt} \times \top \} \end{aligned}$$

Goal 1 thus expects a (dependent) pair whose first component has type ListTag, indicating whether this con node is a zero node or a successor node. Here we choose to construct a successor by supplying ‘cons, which determines the type of the second component:

$$\begin{aligned} n &: \mu NatD \text{ tt} \\ n &= \text{con } (\text{'cons}, \{\mu NatD \text{ tt} \times \top\}_2) \end{aligned}$$

Now we should construct an inner inhabitant of $\mu NatD \text{ tt}$ (paired with a tt) by going through the above process again. We can choose to supply ‘cons several times, but eventually we should choose to supply ‘nil to end the construction:

$$\begin{aligned} n &: \mu NatD \text{ tt} \\ n &= \text{con } (\text{'cons}, \text{con } (\text{'cons}, \text{con } (\text{'cons}, \text{con } (\text{'nil}, \{\top\}_3), \text{tt}), \text{tt}), \text{tt}) \end{aligned}$$

Filling tt into Goal 3 ends the construction. It should be evident now that inhabitants of $\mu NatD \text{ tt}$ indeed represent natural numbers, with $\text{con } (\text{'nil}, \text{tt}) : \mu NatD \text{ tt}$ being zero and $\lambda m \mapsto \text{con } (\text{'cons}, m, \text{tt}) : \mu NatD \text{ tt} \rightarrow \mu NatD \text{ tt}$ being the successor function.

Parametrised datatypes. Parametrised datatypes can simply be encoded as functions mapping parameters to descriptions. For example, the datatype of lists is parametrised by the element type:

$$\begin{aligned} \text{indexfirst data List } (A : \text{Set}) : \text{Set where} \\ \text{List } A \ni & [] \\ & | _::_ A (\text{List } A) \end{aligned}$$

This can be encoded as a function mapping element types to descriptions:

$$\begin{aligned} \text{ListD} &: \text{Set} \rightarrow \text{Desc } \top \\ \text{ListD } A \text{ tt} &= \sigma \text{ ListTag } \lambda \{ \text{'nil} \mapsto v [] \\ &\quad ; \text{'cons} \mapsto \sigma[- \in A] v (\text{tt} :: []) \} \end{aligned}$$

$\text{ListD } A$ is the same as NatD except that, in the ‘cons case, we use σ to insert a field of type A for storing an element.

Indexed datatypes. When a datatype is nontrivially indexed, we can perform arbitrary computation on the index request to produce suitable responses, but for most situations, the kind of computation we need is just pattern matching. For example, the datatype of vectors is encoded as

$$\begin{aligned} \text{VecD} &: \text{Set} \rightarrow \text{Desc } \text{Nat} \\ \text{VecD } A \text{ zero} &= v [] \\ \text{VecD } A (\text{suc } n) &= \sigma[- \in A] v (n :: []) \end{aligned}$$

which is directly comparable to the index-first base functor VecF' .

There are more sophisticated scenarios requiring extra information to be supplied in the descriptions. For example, consider the following variant of finite sets:

indexfirst data $\text{Fin}' : \text{Nat} \rightarrow \text{Set}$ **where**
 $\text{Fin}' n \ni \text{zero}$
 $\text{Fin}' (\text{suc } n) \ni \text{suc } (\text{Fin}' n)$

The inhabitants of $\text{Fin}' n$ are natural numbers less than or equal to n : Regardless of what n is, zero is always an available constructor; if we know that n is a successor, then we can also choose to construct an inhabitant by suc . One way to encode this kind of datatype is to distinguish the responses that are always available from those that are available subject to the result of pattern matching, and the distinction can be encoded as an implicit constructor tag (not appearing explicitly in the datatype declaration):

data $\text{ConMenu} : \text{Set}$ **where**
 $\text{'always} : \text{ConMenu}$
 $\text{'indexed} : \text{ConMenu}$
 $\text{Fin}'D : \text{Desc } \text{Nat}$
 $\text{Fin}'D n = \sigma \text{ ConMenu } \lambda \{ \text{'always} \mapsto v [] \\ \quad ; \text{'indexed} \mapsto \text{case } n \text{ of } \lambda \{ \text{zero} \mapsto \Sigma \perp \lambda () \\ \quad ; (\text{suc } m) \mapsto v (m :: []) \} \}$

Note that the empty response $\Sigma \perp \lambda ()$ produced in the case where n is pattern-matched with zero is also implicitly inserted to make the function total.

Gap 1 (elaboration of datatype declarations to descriptions)

Dagand & McBride (2012) presented an elaboration mechanism for encoding index-first datatype declarations to descriptions, but that mechanism needs to be further extended to deal with more sophisticated cases like always-available constructors and incomplete pattern

matching on the index. To circumvent the gap in this paper, we will present an index-first datatype declaration side by side with its encoding as a description, leaving the formal connection between them unspecified.

2.4 Functions on index-first datatypes

There is no problem defining functions on the encoded datatypes, except that it has to be done with the raw representation. This is easy when the datatypes are not defined by pattern matching on their index. For example, list append can be defined by

$$\begin{aligned} _+_ &: \mu (ListDA) \text{tt} \rightarrow \mu (ListDA) \text{tt} \rightarrow \mu (ListDA) \text{tt} \\ \text{con } ('nil, \text{tt}) &+_ = \text{ys} \\ \text{con } ('cons, x, xs, \text{tt}) &+_ = \text{con } ('cons, x, xs + \text{ys}, \text{tt}) \end{aligned}$$

To improve readability, we can define the following abbreviations, exploiting AGDA's ability to define "pattern synonyms":

$$\begin{aligned} \text{List} &: \text{Set} \rightarrow \text{Set} \\ \text{List } A &= \mu (ListDA) \text{tt} \\ \mathbf{pattern} \ [] &= \text{con } ('nil, \text{tt}) \\ \mathbf{pattern} \ _::_ a as &= \text{con } ('cons, a, as, \text{tt}) \end{aligned}$$

List append can then be rewritten in the usual form:

$$\begin{aligned} _+_ &: \text{List } A \rightarrow \text{List } A \rightarrow \text{List } A \\ [] &+_ = \text{ys} \\ (x :: xs) &+_ = x :: (xs + \text{ys}) \end{aligned}$$

The situation gets trickier when the datatypes are defined by pattern matching on their index. Consider vector append, for example:

$$\begin{aligned} _+_ &: \{m n : \text{Nat}\} \rightarrow \mu (VecDA) m \rightarrow \mu (VecDA) n \rightarrow \mu (VecDA) (m + n) \\ \text{con } xs + \text{ys} &= \{ \}_0 \end{aligned}$$

We cannot further decompose xs because its type gets stuck at $[\text{VecDA } m] (\mu (VecDA))$. In other words, before m is known to have a more specific form, we cannot determine which vector constructor is applicable here. We thus need to perform pattern matching on m before we can proceed:

$$\begin{aligned} _+_ &: \{m n : \text{Nat}\} \rightarrow \mu (VecDA) m \rightarrow \mu (VecDA) n \rightarrow \mu (VecDA) (m + n) \\ _+_ \{m = \text{zero}\} &(\text{con } \text{tt}) \text{ys} = \text{ys} \\ _+_ \{m = \text{suc } m\} &(\text{con } (x, xs, \text{tt})) \text{ys} = \text{con } (x, xs + \text{ys}, \text{tt}) \end{aligned}$$

This is obviously inferior to what we can already do with AGDA's native datatypes: The append program for the vectors defined as a native AGDA datatype is exactly the same as the list append program, using patterns corresponding to the declared constructors, whereas the above definition has to deal with indices first.

Gap 2 (elaboration of pattern matching on index-first data)

```

data Inv {A B : Set} (f : A → B) : B → Set where
  ok : (x : A) → Inv f (f x)
und : {A B : Set} {f : A → B} {y : B} → Inv f y → A
und (ok x) = x

data  $\mathbb{E}$  {I J : Set} (e : J → I) : List J → List I → Set where
  [] :  $\mathbb{E}$  e [] []
  -::- : {j : J} {i : I} (eq : e j ≡ i)
        {js : List J} {is : List I} (eqs :  $\mathbb{E}$  e js is) →  $\mathbb{E}$  e (j :: js) (i :: is)

data ROrn {I J : Set} (e : J → I) : RDesc I → RDesc J → Set1 where
  v : {js : List J} {is : List I} (eqs :  $\mathbb{E}$  e js is) → ROrn e (v is) (v js)
  σ : (S : Set) {D : S → RDesc I} {E : S → RDesc J}
      (O : (s : S) → ROrn e (D s) (E s)) → ROrn e (σ S D) (σ S E)
  Δ : (T : Set) {D : RDesc I} {E : T → RDesc J}
      (O : (t : T) → ROrn e D (E t)) → ROrn e D (σ T E)
  ∇ : {S : Set} (s : S) {D : S → RDesc I} {E : RDesc J}
      (O : ROrn e (D s) E) → ROrn e (σ S D) E

```

Fig. 1. Definitions for response ornaments

Type-theoretically, pattern matching can be explained in terms of the elimination rule, i.e., the induction principle, as done by McBride & McKinna (2004), Goguen *et al.* (2006), and Cockx *et al.* (2014). There is a generic induction principle for index-first datatypes (see, e.g., Section 5.1 of Chapman *et al.* (2010)), but that induction principle does not directly correspond to “proper” pattern matching (against declared constructors, rather than having to analyse the index first), resulting in a gap to be bridged. It is possible to simulate in AGDA proper pattern matching for the encoded index-first datatypes to some extent by, e.g., employing McBride & McKinna’s view idiom, but the syntax is still far from transparent. In Section 5.5 we will experiment with writing more complex programs; to employ a natural pattern matching notation (and also get more readable type information at interaction points), these programs are actually implemented on AGDA’s native versions of the datatypes in the experimental code.

3 Ornaments: relating structurally similar datatypes

Informally, ornaments relate one datatype with another that has the same recursive structure but is more informative. Here “more informative” refers to two things: having more fields and more refined indices. For example, there is an ornament from natural numbers to lists, as the latter has an extra field associated with the cons constructor; and there is also an ornament from lists to vectors, as the latter partitions the type of lists into a family of types indexed by length. On the other hand, there is no ornament from Peano-style natural numbers to binary trees, because they have different recursive structures, one being linear and the other branching.

More formally, ornaments are typed as

$$\text{Orn} : \{I J : \text{Set}\} (e : J \rightarrow I) (D : \text{Desc } I) (E : \text{Desc } J) \rightarrow \text{Set}_1$$

An ornament of type $\text{Orn } e D E$ is a relation from a less informative description $D : \text{Desc } I$ to a more informative description $E : \text{Desc } J$. Logically prior to forming the ornaments, a refining relation on the index sets I and J is first specified, in the form of a function e of type $J \rightarrow I$ (which appears in the type of the ornaments), so every J -index is related to exactly one I -index. Since D and E are collections of response descriptions, an ornament from D to E is actually a J -indexed collection of refining relations from $D (e j)$ to $E j$. For technical (and partly historical) reasons, we use an auxiliary datatype Inv (defined in Figure 1) to index this collection, such that $\text{Inv } e i$ contains those elements of J that are mapped to i by e :

$$\begin{aligned} \text{Orn} & : \{I J : \text{Set}\} (e : J \rightarrow I) (D : \text{Desc } I) (E : \text{Desc } J) \rightarrow \text{Set}_1 \\ \text{Orn } \{I\} e D E & = \{i : I\} (j : \text{Inv } e i) \rightarrow \text{ROrn } e (D i) (E (\text{und } j)) \end{aligned}$$

The datatype ROrn of *response ornaments*, which is where we place the actual definition of how two datatypes are related, is shown in Figure 1. (This figure is intended to provide type information for the constructors. In Section 3.1 below we will explain the constructors by going through concrete examples and showing how to use the constructors to relate datatypes, when the type information will become helpful.) One construction derived from an ornament is a forgetful function that turns an inhabitant of the more informative datatype to one of the less informative datatype, retaining the recursive structure:

$$\begin{aligned} \text{ornForget} & : \{I J : \text{Set}\} \{e : J \rightarrow I\} \{D : \text{Desc } I\} \{E : \text{Desc } J\} \\ & (O : \text{Orn } e D E) \rightarrow \mu E \Rightarrow (\mu D \circ e) \end{aligned}$$

For example, the forgetful function derived from the ornament from natural numbers to lists is the *length* function, and the forgetful function derived from the ornament from lists to vectors computes the underlying list of a vector. Note that the type of *ornForget* gives a reason for the index-refining relation to be specified as a function e : For any index j , an inhabitant of type $\mu E j$ can be transformed into $\mu D (e j)$.

3.1 Understanding ornaments by examples

The most interesting feature of ornaments is their intensional (syntactic) nature: An ornament can be seen as a field-by-field comparison between two datatypes, akin to a “diff” between two descriptions of types $\text{Desc } I$ and $\text{Desc } J$. When writing an ornament, imagine that we are consuming the two descriptions being compared, relating them field by field and index by index. Below we will go through some examples to bring out this intuition, and also explain how the constructors of ROrn are intended to be used.

Our first example is an ornament from natural numbers to lists whose elements are of a given type A . To specify the type of the ornament, first we need to determine the refining relation on the index sets; in this case there is only one choice, namely $! = \lambda _ \mapsto \text{tt}$, because naturals are trivially indexed.

$$\begin{aligned} \text{NatD-ListD} & : (A : \text{Set}) \rightarrow \text{Orn } ! \text{NatD } (\text{ListD } A) \\ \text{NatD-ListD } A & = \{\text{Orn } ! \text{NatD } (\text{ListD } A)\}_0 \end{aligned}$$

An ornament is a collection of response ornaments indexed by the more informative indices; in this case, again, there is only one such index, namely tt .

$$\begin{aligned} \text{NatD-ListD} & : (A : \text{Set}) \rightarrow \text{Orn} ! \text{NatD} (\text{ListD } A) \\ \text{NatD-ListD } A \text{ (ok tt)} & = \{ \text{ROrn} ! (\text{NatD tt}) (\text{ListD } A \text{ tt}) \}_1 \end{aligned}$$

Now we should relate the response descriptions NatD tt and $\text{ListD } A \text{ tt}$, shown as indices in the goal type. Both response descriptions (shown in Section 2.3) start with a common field of type List Tag . To indicate that the field is a common one in both response descriptions, we use the σ constructor of ROrn :

$$\begin{aligned} \text{NatD-ListD} & : (A : \text{Set}) \rightarrow \text{Orn} ! \text{NatD} (\text{ListD } A) \\ \text{NatD-ListD } A \text{ (ok tt)} & = \sigma \text{ ListTag} \{ (t : \text{ListTag}) \rightarrow \text{ROrn} ! (D' t) (E' t) \}_2 \\ & \quad \text{-- where } D' = \lambda \{ 'nil \mapsto v []; 'cons \mapsto v (\text{tt} :: []) \} \\ & \quad \text{-- and } E' = \lambda \{ 'nil \mapsto v []; 'cons \mapsto \sigma[- \in A] v (\text{tt} :: []) \} \end{aligned}$$

To make $D' t$ and $E' t$ reduce, we pattern-match on t , splitting the goal into two:

$$\begin{aligned} \text{NatD-ListD} & : (A : \text{Set}) \rightarrow \text{Orn} ! \text{NatD} (\text{ListD } A) \\ \text{NatD-ListD } A \text{ (ok tt)} & = \\ & \quad \sigma \text{ ListTag} \lambda \{ 'nil \mapsto \{ \text{ROrn} ! (v []) (v []) \}_3 \\ & \quad \quad ; 'cons \mapsto \{ \text{ROrn} ! (v (\text{tt} :: [])) (\sigma[- \in A] v (\text{tt} :: [])) \}_4 \} \end{aligned}$$

Let us look at Goal 4 first. There is a field of type A (for storing list elements) only on the more informative side and not on the less informative side; to indicate this, we use the Δ constructor:

$$\begin{aligned} \text{NatD-ListD} & : (A : \text{Set}) \rightarrow \text{Orn} ! \text{NatD} (\text{ListD } A) \\ \text{NatD-ListD } A \text{ (ok tt)} & = \\ & \quad \sigma \text{ ListTag} \lambda \{ 'nil \mapsto \{ \text{ROrn} ! (v []) (v []) \}_3 \\ & \quad \quad ; 'cons \mapsto \Delta[- \in A] \{ \text{ROrn} ! (v (\text{tt} :: [])) (v (\text{tt} :: [])) \}_5 \} \end{aligned}$$

Here again we regard Δ as a binder and write $\Delta[t \in T] O t$ for $\Delta T (\lambda t \mapsto O t)$. Now both goals require us to relate variable positions, which is done by the v constructor of ROrn :

$$\begin{aligned} \text{NatD-ListD} & : (A : \text{Set}) \rightarrow \text{Orn} ! \text{NatD} (\text{ListD } A) \\ \text{NatD-ListD } A \text{ (ok tt)} & = \\ & \quad \sigma \text{ ListTag} \lambda \{ 'nil \mapsto v \{ \mathbb{E} ! [] [] \}_6 \\ & \quad \quad ; 'cons \mapsto \Delta[- \in A] v \{ \mathbb{E} ! (\text{tt} :: []) (\text{tt} :: []) \}_7 \} \end{aligned}$$

The datatype $\mathbb{E} e js$ is just a more structured representation of the equality $map e js \equiv is$, guaranteeing that the recursive structure is the same (i.e., the two index lists have the same length) and that the corresponding pairs of indices respect the refining relation e . Here the lists in each goal are indeed of the same length, and the indices in Goal 7 are trivial, so both goals can be automatically discharged.

$$\begin{aligned} \text{NatD-ListD} & : (A : \text{Set}) \rightarrow \text{Orn} ! \text{NatD} (\text{ListD } A) \\ \text{NatD-ListD } A \text{ (ok tt)} & = \\ & \quad \sigma \text{ ListTag} \lambda \{ 'nil \mapsto v [] \\ & \quad \quad ; 'cons \mapsto \Delta[- \in A] v (\text{refl} :: []) \} \end{aligned}$$

The forgetful function $\text{ornForget} (\text{NatD-ListD } A)$ discards the A -typed element associated with each cons node (because the field is marked using Δ as being exclusive to the description

of lists) and keeps everything else, including the recursive structure (a series of ‘cons nodes ending with a ‘nil node). That is, it computes the length of a list.

Our second example is an ornament from lists to vectors, starting from the following ornament type:

$$\begin{aligned} \text{ListD-VecD} &: (A : \text{Set}) \rightarrow \text{Orn} ! (\text{ListD } A) (\text{VecD } A) \\ \text{ListD-VecD } A \text{ (ok } n) &= \{ \text{ROrn} ! (\text{ListD } A \text{ tt}) (\text{VecD } A \text{ } n) \}_0 \end{aligned}$$

An ornament is a collection of response ornaments, one for each index of the more informative datatype — in this case, vectors, which are indexed by natural numbers.

To make $\text{VecD } A \text{ } n$ reduce, we pattern-match on n :

$$\begin{aligned} \text{ListD-VecD} &: (A : \text{Set}) \rightarrow \text{Orn} ! (\text{ListD } A) (\text{VecD } A) \\ \text{ListD-VecD } A \text{ (ok zero)} &= \{ \text{ROrn} ! (\text{ListD } A \text{ tt}) (\text{v } []) \}_1 \\ \text{ListD-VecD } A \text{ (ok (suc } n)) &= \{ \text{ROrn} ! (\text{ListD } A \text{ tt}) (\sigma[_ \in A] \text{ v } (n :: [])) \}_2 \end{aligned}$$

This time there is a field of type `ListTag` that appears only on the less informative side, namely in $\text{ListD } A \text{ tt}$. The value for this field, however, should be fixed for both goals: For Goal 1 it should be ‘nil, and for Goal 2 it should be ‘cons (otherwise the goals will be unachievable). To indicate so, we use the ∇ constructor:

$$\begin{aligned} \text{ListD-VecD} &: (A : \text{Set}) \rightarrow \text{Orn} ! (\text{ListD } A) (\text{VecD } A) \\ \text{ListD-VecD } A \text{ (ok zero)} &= \nabla[\text{'nil}] \{ \text{ROrn} ! (\text{v } []) (\text{v } []) \}_3 \\ \text{ListD-VecD } A \text{ (ok (suc } n)) &= \nabla[\text{'cons}] \{ \text{ROrn} ! (\sigma[_ \in A] \text{ v } (\text{tt} :: [])) \\ &\quad (\sigma[_ \in A] \text{ v } (n :: [])) \}_4 \end{aligned}$$

Even though ∇ is not a binder, we write $\nabla[s] O$ for $\nabla s O$ to avoid the parentheses around O when O is a complex expression. Goal 3 is easily discharged using v . As for Goal 4, we use σ to relate the element field on both sides, and then use v to end the ornament.

$$\begin{aligned} \text{ListD-VecD} &: (A : \text{Set}) \rightarrow \text{Orn} ! (\text{ListD } A) (\text{VecD } A) \\ \text{ListD-VecD } A \text{ (ok zero)} &= \nabla[\text{'nil}] \text{ v } [] \\ \text{ListD-VecD } A \text{ (ok (suc } n)) &= \nabla[\text{'cons}] \sigma[_ \in A] \text{ v } (\text{refl} :: []) \end{aligned}$$

The forgetful function $\text{ornForget } (\text{ListD-VecD } A)$ erases the index, reinstalls the constructor tags ‘nil and ‘cons, and keeps everything else. That is, it computes the underlying list of a vector.

3.2 Ornamental descriptions

Often we want to define a more informative datatype that is an ornamentation of an existing datatype. *Ornamental descriptions* are provided for conveniently writing both the description and the ornament in a single definition. If an ornament is a “diff” between two descriptions, then an ornamental description is a “patch” specified using an existing description as a template, and this patch can be “applied” to produce a more informative description. The definitions are shown in Figure 2. Note that `OrnDesc` is indexed by only one template description. Historically, ornamental descriptions were in fact McBride’s (2011) original formulation of ornaments, but we needed to formulate ornaments as binary relations between

$$\begin{aligned}
& \text{OrnDesc} : \{I : \text{Set}\} (J : \text{Set}) (e : J \rightarrow I) (D : \text{Desc } I) \rightarrow \text{Set}_1 \\
& \text{OrnDesc } \{I\} J e D = \{i : I\} (j : \text{Inv } e i) \rightarrow \text{ROrnDesc } J e (D i) \\
& \mathbf{data} \text{ROrnDesc } \{I : \text{Set}\} (J : \text{Set}) (e : J \rightarrow I) : \text{RDesc } I \rightarrow \text{Set}_1 \mathbf{where} \\
& \quad v : \{is : \text{List } I\} (js : \mathbb{P} is (\text{Inv } e)) \rightarrow \text{ROrnDesc } J e (v is) \\
& \quad \sigma : (S : \text{Set}) \{D : S \rightarrow \text{RDesc } I\} \\
& \quad \quad (O : (s : S) \rightarrow \text{ROrnDesc } J e (D s)) \rightarrow \text{ROrnDesc } J e (\sigma S D) \\
& \quad \Delta : (T : \text{Set}) \{D : \text{RDesc } I\} \\
& \quad \quad (O : T \rightarrow \text{ROrnDesc } J e D) \rightarrow \text{ROrnDesc } J e D \\
& \quad \nabla : \{S : \text{Set}\} (s : S) \{D : S \rightarrow \text{RDesc } I\} \\
& \quad \quad (O : \text{ROrnDesc } J e (D s)) \rightarrow \text{ROrnDesc } J e (\sigma S D)
\end{aligned}$$

Fig. 2. Definitions for ornamental descriptions

descriptions (Ko & Gibbons, 2013a) in order to define parallel composition (Section 3.3) as a categorical pullback.

Suppose we wish to write an ornamental description of ordered lists indexed by a lower bound, using the description of lists as a template. Also suppose that the list elements are of type Val , on which there is an ordering relation $_{\leq} : Val \rightarrow Val \rightarrow \text{Set}$. Eventually, this ornamental description should give us the following datatype:

$$\begin{aligned}
& \mathbf{indexfirst} \mathbf{data} \text{OrdList} : Val \rightarrow \text{Set} \mathbf{where} \\
& \quad \text{OrdList } b \ni \text{nil} \\
& \quad \quad | \text{cons } (x : Val) (b \leq x) (\text{OrdList } x)
\end{aligned}$$

OrnDesc , like Orn , is an indexed collection of response ornamental descriptions of type ROrnDesc . We start by assuming that the lower bound is b :

$$\begin{aligned}
& \text{OrdListOD} : \text{OrnDesc } Val ! (\text{ListD } Val) \\
& \text{OrdListOD } (\text{ok } b) = \{ \text{ROrnDesc } Val ! (\text{ListD } Val \text{ tt}) \}_0
\end{aligned}$$

The first field of $\text{ListD } Val \text{ tt}$ is ListTag , which we should keep by using the σ constructor of ROrnDesc :

$$\begin{aligned}
& \text{OrdListOD} : \text{OrnDesc } Val ! (\text{ListD } Val) \\
& \text{OrdListOD } (\text{ok } b) = \sigma \text{ListTag} \\
& \quad \{ (t : \text{ListTag}) \rightarrow \\
& \quad \quad \text{ROrnDesc } Val ! ((\lambda \{ \text{'nil} \mapsto v [] ; \text{'cons} \mapsto \sigma[_ \in Val] v (\text{tt} :: []) \}) t) \}_1
\end{aligned}$$

Again we should pattern-match on t to reveal the two cases for the two list constructors:

$$\begin{aligned}
& \text{OrdListOD} : \text{OrnDesc } Val ! (\text{ListD } Val) \\
& \text{OrdListOD } (\text{ok } b) = \\
& \quad \sigma \text{ListTag } \lambda \{ \text{'nil} \mapsto \{ \text{ROrnDesc } Val ! (v []) \}_2 \\
& \quad \quad ; \text{'cons} \mapsto \{ \text{ROrnDesc } Val ! (\sigma[_ \in Val] v (\text{tt} :: [])) \}_3 \}
\end{aligned}$$

Goal 2 can be met using the v constructor of ROrnDesc . Goal 3, which corresponds to the refinement of the cons constructor, is more interesting. We first use σ to keep the element field:

$$\begin{aligned}
& \text{OrdListOD} : \text{OrnDesc } Val ! (\text{ListD } Val) \\
& \text{OrdListOD } (\text{ok } b) =
\end{aligned}$$

$$\sigma \text{ListTag } \lambda \{ \text{'nil} \mapsto v \text{tt} \\ ; \text{'cons} \mapsto \sigma[x \in \text{Val}] \{ \text{ROrnDesc Val ! (v (tt :: [])) \}_4 \}$$

We should not close Goal 4 just yet, because we should require that the element x is bounded below by b . This is done by inserting a new field of type $b \leq x$ by using the Δ constructor of ROrnDesc :

$$\begin{aligned} \text{OrdListOD} &: \text{OrnDesc Val ! (ListD Val)} \\ \text{OrdListOD (ok } b) &= \\ &\sigma \text{ListTag } \lambda \{ \text{'nil} \mapsto v \text{tt} \\ & ; \text{'cons} \mapsto \sigma[x \in \text{Val}] \Delta[- \in (b \leq x)] \\ & \quad \{ \text{ROrnDesc Val ! (v (tt :: [])) \}_5 \} \end{aligned}$$

Now we can refine the index for the tail to x , meaning that the tail should be bounded below by x :

$$\begin{aligned} \text{OrdListOD} &: \text{OrnDesc Val ! (ListD Val)} \\ \text{OrdListOD (ok } b) &= \\ &\sigma \text{ListTag } \lambda \{ \text{'nil} \mapsto v \text{tt} \\ & ; \text{'cons} \mapsto \sigma[x \in \text{Val}] \Delta[- \in (b \leq x)] v(\text{ok } x, \text{tt}) \} \end{aligned}$$

Note that the argument to v is a tuple containing exactly as many indices as those in the index list in the template description, keeping the recursive branching structure by construction.

We can interpret an ornamental description using $[-]$ to extract a description of type Desc Val of the new datatype:

$$\begin{aligned} [\text{OrdListOD}] b &= \\ &\sigma \text{ListTag } \lambda \{ \text{'nil} \mapsto v [] \\ & ; \text{'cons} \mapsto \sigma[x \in \text{Val}] \sigma[- \in (b \leq x)] v(x :: []) \} \end{aligned}$$

Alternatively, we can interpret the ornamental description using $[-]$ to extract an ornament of type $\text{Orn ! (ListD Val) [OrdListOD]}$ from the template description to the new description:

$$\begin{aligned} [\text{OrdListOD}] (\text{ok } b) &= \\ &\sigma \text{ListTag } \lambda \{ \text{'nil} \mapsto v [] \\ & ; \text{'cons} \mapsto \sigma[x \in \text{Val}] \Delta[- \in (b \leq x)] v(\text{refl} :: []) \} \end{aligned}$$

Gap 3 (interactive support for constructing of ornaments and ornamental descriptions)

Both Dagand & McBride (2014) and Williams *et al.* (2014) propose concrete syntax for defining ornamental descriptions, and Williams *et al.* also devise an “ornament from” mechanism for defining ornaments (between two existing datatypes, rather than creating a new datatype on top of an existing one) in terms of forgetful functions, which should conform to some syntax restrictions. We have shown that the construction of ornaments (and ornamental descriptions) can be performed in an interactive manner, so perhaps ornaments can also be constructed in a special tactic language, in which the programmer can go through the interactive construction process by giving instructions linking two existing datatype declarations, and get valid ornaments by construction.

3.3 Ornamental promotion isomorphisms

The syntactic nature of ornaments allows us to compute from them things other than the forgetful function. Indeed, the forgetful function can be seen as part of an isomorphism connecting the two datatypes related by an ornament: From an ornament $O : \text{Orn } e D E$ we can compute an *optimised promotion predicate* $\text{OptP } O : \{i : I\} (j : \text{Inv } e i) \rightarrow \mu D i \rightarrow \text{Set}$ that captures the necessary information for augmenting an inhabitant of the less informative datatype to an inhabitant of the more informative datatype. More precisely, for every $j : \text{Inv } e i$ there is an isomorphism

$$\mu E (\text{und } j) \cong \Sigma (\mu D i) (\text{OptP } O j)$$

which says that an inhabitant of the more informative datatype is interconvertible with an inhabitant of the less informative datatype paired with a proof that it satisfies the optimised promotion predicate. A representative example is the ornament from lists to ordered lists. In this case, the optimised promotion predicate computed from the ornament is

indexfirst data $\text{Ordered} : \text{Val} \rightarrow \text{List Val} \rightarrow \text{Set}$ **where**

$\text{Ordered } b [] \ni \text{nil}$

$\text{Ordered } b (x :: xs) \ni \text{cons } (leq : b \leq x) (\text{ord} : \text{Ordered } x xs)$

An inhabitant of $\text{Ordered } b xs$ is exactly a series of inequality proofs that chains together the lower bound b and all the elements of xs in order. Given such a series of inequality proofs, we can promote xs to an $\text{OrdList } b$; conversely, any $\text{OrdList } b$ can be separated into a list xs and a series of inequality proofs of type $\text{Ordered } b xs$. That is, for every b there is an isomorphism

$$\text{OrdList } b \cong \Sigma (\text{List Val}) (\text{Ordered } b)$$

We call this the *ornamental promotion isomorphism*.

We could construct the optimised promotion predicate and the isomorphism directly, as Dagand & McBride (2014) do with “reornaments”, but in fact there is a more general construction — *parallel composition* of ornaments. If we think of ornaments as diffs, then parallel composition is analogous to three-way merging: Given two ornaments $O : \text{Orn } e D E$ and $P : \text{Orn } f D F$ with the same less informative end D , we can compute an ornamental description $O \otimes P$ relative to D that incorporates all modifications to D recorded in O and in P . For example, we can compose in parallel the ornaments from lists to ordered lists and to vectors to get ordered vectors:

indexfirst data $\text{OrdVec} : \text{Val} \rightarrow \text{Nat} \rightarrow \text{Set}$ **where**

$\text{OrdVec } b \text{ zero} \ni \text{nil}$

$\text{OrdVec } b (\text{suc } n) \ni \text{cons } (x : \text{Val}) (leq : b \leq x) (xs : \text{OrdVec } x n)$

With parallel composition, the optimised promotion predicate for an ornament $O : \text{Orn } e D E$ can easily be defined as the parallel composition of O and the *singleton* ornamentation for D . Here we give only a quick and intuitive explanation. (For a more detailed account, see Section 3.3.1 of Ko (2014).) The singleton ornamentation for D gives the singleton datatype (Sheard & Linger, 2007) indexed by μD , and each type in the family has exactly one inhabitant. For example, the singleton datatype for lists is

indexfirst data ListS ($A : \text{Set}$) : List $A \rightarrow \text{Set}$ **where**
 ListS $A [] \ni \text{nil}$
 ListS $A (x :: xs) \ni \text{cons (ListS } A \text{ } xs)$

This datatype can be obtained from the following ornamental description:

$ListSOD : (A : \text{Set}) \rightarrow \text{OrnDesc (List } A) ! (ListD A)$
 $ListSOD A (\text{ok } [] \quad) = \nabla[\text{'nil}] \vee \text{tt}$
 $ListSOD A (\text{ok } (x :: xs)) = \nabla[\text{'cons}] \nabla[x] \vee (\text{ok } xs, \text{tt})$

Note that all the fields are deleted. What ornamental description do we get if we compose in parallel this ornament with the one from lists to ordered lists? The extra field for ordered lists storing inequality proofs will be inserted, but the fields for lists storing constructor choices and elements will be deleted. Thus the resulting datatype Ordered is left with only the extra field for ordered lists, and inhabitants of Ordered store only the inequality proofs.

We do not give the construction of the ornamental isomorphism in this paper, but instead refer the reader to Chapter 4 of Ko (2014), where parallel composition is shown to be a categorical pullback. The ornamental isomorphism can then be constructed from the pullback property, independently of the choice of universe encoding. It is worth mentioning that an arbitrary choice of encoding may make it impossible to establish the pullback property: Indeed, the encoding used in Ko & Gibbons (2013a) does not satisfy the pullback property because there is a product constructor $_*_* : \text{RDesc } I \rightarrow \text{RDesc } I \rightarrow \text{RDesc } I$ that turns out to be problematic; we thus switched to the version presented by Ko (2014) in this paper, in particular enforcing that all the recursive positions are grouped in a list, so that the pullback property can be proved.

4 Promotions and upgrades: lifting ornaments for function types

Ornaments give rise to constructions for promoting an element of a basic datatype to a structurally similar element of a more informative datatype. Having such constructions only for datatypes, however, is not enough — we need to promote operations on these datatypes too. Lifting of ornaments for function types has been considered by Ko & Gibbons (2013a) and more systematically developed by Dagand & McBride (2014). Here we propose another systematic lifting mechanism, which is less complicated and more easily extensible than Dagand & McBride’s approach.

4.1 Promotions: axiomatising the ornamental relationship

Since the ornamental constructions can be summarised as the promotion isomorphisms — for instance, $\text{OrdList } b \cong \Sigma (\text{List } Val) (\text{Ordered } b)$ — it is natural to axiomatise promotability as the existence of such isomorphisms. It turns out that we require only a part of the inverse properties to make the function-promoting mechanism work. Hence we dismantle the promotion isomorphism and include only the necessary inverse property in the axiomatisation:

record Promotion ($X Y : \text{Set}$) : Set_1 **where**
field

Predicate : $X \rightarrow \text{Set}$

forget : $Y \rightarrow X$

complement : $(y : Y) \rightarrow \text{Predicate } (\text{forget } y)$

promote : $(x : X) \rightarrow \text{Predicate } x \rightarrow Y$

coherence : $(x : X) (p : \text{Predicate } x) \rightarrow \text{forget } (\text{promote } x p) \equiv x$

In prose, when we say that $X : \text{Set}$ can be promoted to $Y : \text{Set}$, we mean that there is a *promotion predicate* $P : X \rightarrow \text{Set}$ such that Y and $\Sigma X P$ are interconvertible by *forget* Δ *complement* : $Y \rightarrow \Sigma X P$ (where $_{\Delta}$ is defined by $(f \Delta g) x = (f x, g x)$) and *uncurry promote* : $\Sigma X P \rightarrow Y$. We might think of every element $y : Y$ as consisting of a “core” element *forget* $y : X$ extended with some additional information of type P (*forget* y). Consequently, when we construct an element $y : Y$ by promoting $x : X$, we expect that its core, *forget* y , is exactly x — this is the *coherence* property that we impose on *promote* and *forget*. From every ornamental promotion isomorphism we can construct a promotion for which *Predicate* is the optimised promotion predicate, *forget* is the forgetful function, and the remaining three components are drawn from the appropriate parts of the isomorphism. For example, the promotion induced by the ornament from lists to ordered lists with lower bound b has type Promotion (List Val) (OrdList b); it uses Ordered b as its *Predicate* component and the forgetful function of the ornament as its *forget* component.

While promotions are defined to be an abstraction of the ornamental relationship, we do not strive to strengthen the definition so as to rule out other unintended instantiations. What matters to us is that the function-promoting mechanism works as expected when the basic promotions used are induced by ornaments. To see the generality of promotions, note that promotions can be seen as a variant of *bidirectional transformations*, a subject that has been extensively explored. (See, e.g., Czarnecki *et al.* (2009) for a comprehensive survey.) In a promotion from X to Y , Y is the source type and X is the view type, with *forget* and *promote* being respectively the forward (*get*) and backward (*put*) transformations. The promotion predicate corresponds to the idea of *view complement*, which represents the information that is present in the source but not in the view. (In fact we chose the name *complement* with this analogy in mind.) Finally, the *coherence* property is exactly the *PutGet* property, one of the two inverse-like properties that constitute *well-behavedness* of bidirectional transformations.

4.2 Towards promotions for function types

Let us first explain the idea of function promotion with a small example. Since there is an ornament from natural numbers to lists, we can induce a promotion $p : \text{Promotion Nat (List } A)$ (for some fixed $A : \text{Set}$), which we will use throughout this example. The promotion predicate of p is *Vec* A : To promote a natural number n to a list, we need to supply n elements, i.e., a vector of type *Vec* $A n$. With respect to this promotion p , we can promote the following function

$$\begin{aligned} \text{double} &: \text{Nat} \rightarrow \text{Nat} \\ \text{double zero} &= \text{zero} \\ \text{double (suc } n) &= \text{suc (suc (double } n)) \end{aligned}$$

to a “similarly behaving” function $g : \text{List } A \rightarrow \text{List } A$. By “similarly behaving” we mean that g and double map similar arguments to similar results. Since, as specified by p , we regard natural numbers and lists as similar when they have the same recursive structure, g should process the recursive structure (i.e., nil and cons nodes) of its input in the same way as double , i.e.,

$$(xs : \text{List } A) \rightarrow \text{double} (\text{length } xs) \equiv \text{length} (g \text{ } xs)$$

or as a commutative diagram:

$$\begin{array}{ccc} \text{List } A & \xrightarrow{g} & \text{List } A \\ \text{length} \downarrow & & \downarrow \text{length} \\ \text{Nat} & \xrightarrow{\text{double}} & \text{Nat} \end{array}$$

Dagand & McBride (2014) called this the *coherence property*. (Note that it is stated in terms of length , i.e., the *forget* function of p .)

By analogy with the data promotion mechanism offered by ornaments, rather than explicitly writing a coherent function g , we can instead show that double satisfies a promotion predicate. We already know that g should produce lists whose lengths are computed by double ; the missing information for constructing g , then, is how g processes the list elements while respecting the length constraint. The natural promotion predicate to use in this case is thus

$$\lambda f \mapsto (n : \text{Nat}) \rightarrow \text{Vec } A \ n \rightarrow \text{Vec } A \ (f \ n)$$

(Note, again, that this is defined in terms of the promotion predicate of p , i.e., $\text{Vec } A$.) Given a promotion proof for double , say

$$\begin{aligned} \text{duplicate}' &: (n : \text{Nat}) \rightarrow \text{Vec } A \ n \rightarrow \text{Vec } A \ (\text{double } n) \\ \text{duplicate}' \text{ zero} & \quad [] \quad = \quad [] \\ \text{duplicate}' (\text{suc } n) (x :: xs) &= x :: x :: \text{duplicate}' \ n \ xs \end{aligned}$$

we can synthesise a function $\text{duplicate} : \text{List } A \rightarrow \text{List } A$ by

$$\begin{aligned} \text{duplicate} &: \text{List } A \rightarrow \text{List } A \\ \text{duplicate} &= \text{uncurry promote} \circ (\text{double} * \text{duplicate}' _) \circ (\text{forget } \Delta \text{ complement}) \\ &\textbf{where open Promotion } p \end{aligned}$$

(where $(f * g) (x, y) = (f \ x, g \ y)$). That is, we decompose the input list into the underlying natural number and a vector of elements, process the two parts separately with double and $\text{duplicate}'$, and finally combine the results back to a list.

The sketch above suggests that we can construct promotions on function types from smaller promotions. The actual story has a few twists, though, as we will show below.

4.2.1 A failed attempt

A first attempt might be to write a promotion combinator of type

$$\text{Promotion } X Y \rightarrow \text{Promotion } Z W \rightarrow \text{Promotion } (X \rightarrow Z) (Y \rightarrow W)$$

We can argue that this combinator does not make sense, however: Consider the ornament-induced promotion from Nat to List Nat . Having the above combinator would mean that we can construct a promotion from $\text{Nat} \rightarrow \text{Nat}$ to $\text{List Nat} \rightarrow \text{List Nat}$. We have seen in the *double-duplicate* example that the natural choice of the promotion predicate is

$$\lambda f \mapsto (n : \text{Nat}) \rightarrow \text{Vec Nat } n \rightarrow \text{Vec Nat } (f n)$$

The problem arises when we try to define *forget* and *complement*. Being able to define these two functions means that every function of type $\text{List Nat} \rightarrow \text{List Nat}$ can be decomposed into an underlying function of type $\text{Nat} \rightarrow \text{Nat}$ acting exclusively on the recursive structure and a complement function on the list elements. This is clearly not the case, though, as we can easily find functions which manipulate the recursive structure in a way that depends on the list elements — for example, we can sum the elements of the input list and then produce a list of zeros whose length is the sum (and does not depend on the length of the input list at all).

We should note that the above argument is not rigorous, since it does not really show that a combinator of the given type cannot exist — we could have chosen other promotion predicates and/or constructed “strange” yet passable *forget* and *complement* functions. (Note that there is no law in the definition of promotions that put restrictions on *complement*.) The argument merely shows that a combinator of the type does not match well with our intention. Also note that taking variance into account (changing the type of the first argument to the combinator to $\text{Promotion } Y X$) does not help: In the case of *double* and *duplicate*, for example, we want to uniformly promote Nat to List Nat in the type of *double*, regardless of variance.

4.2.2 Promoting to evidently coherent functions only

The attempt in Section 4.2.1 fails because not every function of type $\text{List Nat} \rightarrow \text{List Nat}$ behaves like a function of type $\text{Nat} \rightarrow \text{Nat}$. One way out is to restrict the promotion to only those functions on lists that do behave like functions on natural numbers. In general, we could aim for an alternative combinator of the following type:

$$(p : \text{Promotion } X Y) (q : \text{Promotion } Z W) \rightarrow \\ \text{Promotion } (X \rightarrow Z) (\Sigma[f \in Y \rightarrow W] \Sigma[g \in X \rightarrow Z] \text{FCoherent } p q f g)$$

where

$$\text{FCoherent} : \text{Promotion } X Y \rightarrow \text{Promotion } Z W \rightarrow (Y \rightarrow W) \rightarrow (X \rightarrow Z) \rightarrow \text{Set} \\ \text{FCoherent } p q f g = (y : Y) \rightarrow \text{Promotion}.\text{forget } q (f y) \equiv g (\text{Promotion}.\text{forget } p y)$$

That is, we promote $X \rightarrow Z$ to only the subset of $Y \rightarrow W$ that are evidently coherent with some function of type $X \rightarrow Z$. This combinator is easy to define, and works reasonably well

in simple cases: The *promote* function now produces a triple consisting of the promoted function, the input core function, and a coherence proof for the two functions. After a promotion, we do a projection to get either the promoted function or the coherence proof.

This combinator does not work so well for curried functions, however. For example, composing three promotions respectively from basic types X_0 , X_1 , and X_2 to more informative types Y_0 , Y_1 , and Y_2 using the combinator in the right order, we get a promotion from $X_0 \rightarrow X_1 \rightarrow X_2$ to

$$\begin{aligned} & \Sigma[f \in Y_0 \rightarrow \Sigma[f' \in Y_1 \rightarrow Y_2] \Sigma[g' \in X_1 \rightarrow X_2] \text{FCoherent } _ _ f' g'] \\ & \Sigma[g \in X_0 \rightarrow X_1 \rightarrow X_2] \text{FCoherent } _ _ f g \end{aligned}$$

with which it is inconvenient to extract even just the promoted functions, let alone the coherence proofs.

In general, this trick of “promoting to an evidently coherent subset” could be used to develop a variety of promotion combinators, but these combinators do not work well when they are used to construct promotions between more complex types. For a different example, sometimes we want to promote X to $(i : I) \rightarrow Y i$ provided that X can be promoted to $Y i$ for every $i : I$. This is useful when, for example, promoting $\text{List } A \rightarrow \text{List } A$ to $(n : \text{Nat}) \rightarrow \text{Vec } A n \rightarrow \text{Vec } A n$, where the promoted function type has an additional argument n for indexing the vectors. Again, there is no forgetful function from $(i : I) \rightarrow Y i$ to X , and we must restrict the promoted type to $\Sigma[f \in (i : I) \rightarrow Y i] \Sigma[x \in X] ((i : I) \rightarrow \text{forget } (f i) \equiv x)$. This combinator, like the one for function types above, can greatly complicate the extraction of promoted functions and coherence proofs. We will develop a better solution next, in which this trick will still play a role though.

4.3 Upgrades: lifting promotions for function types

(Part of) the problem with promotions for function types in Section 4.2.1 was that we cannot define a sensible *forget*. In typical use cases, however, after we *promote* a basic function, there is no need to apply *forget* to the promoted function to retrieve the basic function — we already have the basic function in the first place. This suggests that we use a slightly different formulation of promotion for function types, which does not support *forget*:

```
record Upgrade (X Y : Set) : Set1 where
  field
    Predicate   : X → Set
    Coherence   : X → Y → Set
    complement  : (x : X) (y : Y) → Coherence x y → Predicate x
    promote     : (x : X) → Predicate x → Y
    coherence   : (x : X) (p : Predicate x) → Coherence x (promote x p)
```

An upgrade does not require the existence of a *forget* function, but includes a *Coherence* property which is used to reformulate the types of *complement* and *coherence*. The relationship between promotions and upgrades is most clearly shown by the following combinator *toUpgrade*, which says that promotions are special cases of upgrades, namely those whose *Coherence* property is stated in terms of *forget*:

```

toUpgrade : {X Y : Set} → Promotion X Y → Upgrade X Y
toUpgrade p = record
  { Predicate   = Promotion.Predicate p
  ; Coherence   = λ x y ↦ Promotion.forget p y ≡ x
  ; complement  = λ {..} y refl ↦ Promotion.complement p y }
  ; promote    = Promotion.promote p
  ; coherence   = Promotion.coherence p }

```

We can then formulate in general terms the promoting construction in the *double-duplicate* example:

```

_→_ : {X Y Z W : Set} →
  Promotion X Y → Upgrade Z W → Upgrade (X → Z) (Y → W)
_→_ {X} {Y} p u = record
  { Predicate   =
    λ f ↦ (x : X) → Promotion.Predicate p x → Upgrade.Predicate u (f x)
  ; Coherence   =
    λ f g ↦ (y : Y) → Upgrade.Coherence u (f (Promotion.forget p y)) (g y)
  ; complement  = ...
  ; promote    =
    λ f com ↦ uncurry (Upgrade.promote u) ∘ (f * com _) ∘
      (Promotion.forget p Δ Promotion.complement p)
  ; coherence   = ... }

```

In particular, *promote* does exactly what we did in the *double-duplicate* example: Given $f : X \rightarrow Z$ and its complement function com , we construct a function of type $Y \rightarrow W$ by separating the argument of type Y into its core of type X and a complement, processing them respectively using f and com , and invoking *promote* of the upgrade u to combine the processed pair. The *double-duplicate* example can now be handled by the upgrade

```

Nat-List A → toUpgrade (Nat-List A) : Upgrade (Nat → Nat) (List A → List A)

```

where $Nat-List : (A : Set) \rightarrow Promotion Nat (List A)$ encapsulates the promotion isomorphism induced by the ornament from natural numbers to lists. Note that the computed coherence property looks exactly as we want it to, and that computation works nicely for curried functions too, solving one problem mentioned in Section 4.2.2.

The promotion–upgrade mechanism is easily extensible, since the combinators are shallowly embedded (cf. Dagand & McBride’s (2014) approach, which uses deep embedding). For example, the other problem mentioned in Section 4.2.2 is additional arguments in promoted function types. We can extend the mechanism to deal with additional arguments by writing a new combinator:

```

new : (I : Set) {X : Set} {Y : I → Set}
  (u : (i : I) → Upgrade X (Y i)) → Upgrade X ((i : I) → Y i)
new I u = record
  { Predicate   = λ x ↦ (i : I) → Upgrade.Predicate (u i) x
  ; Coherence   = λ x y ↦ (i : I) → Upgrade.Coherence (u i) x (y i)
  ; complement  = λ x y coh i ↦ Upgrade.complement (u i) x (y i) (coh i)

```

$$\begin{aligned} & ; \textit{promote} &= \lambda x \textit{com} i \mapsto \textit{Upgrade.promote} (u i) x (\textit{com} i) \\ & ; \textit{coherence} &= \lambda x \textit{com} i \mapsto \textit{Upgrade.coherence} (u i) x (\textit{com} i) \} \\ \textit{syntax new } I & (\lambda i \mapsto u) = \forall^+[i \in I] u \end{aligned}$$

The promotion mentioned in Section 4.2.2 can then be handled by the upgrade

$$\begin{aligned} & \forall^+[n \in \text{Nat}] (\textit{List-Vec} A n \rightarrow \textit{toUpgrade} (\textit{List-Vec} A n)) : \\ & \textit{Upgrade} (\textit{List} A \rightarrow \textit{List} A) ((n : \text{Nat}) \rightarrow \textit{Vec} A n \rightarrow \textit{Vec} A n) \end{aligned}$$

where $\textit{List-Vec} : (A : \text{Set}) (n : \text{Nat}) \rightarrow \text{Promotion} (\textit{List} A) (\textit{Vec} A n)$ is induced by the ornament from lists to vectors. The upgrade itself looks like a dependent function type, thanks to the mixfix binder $\forall^+[n \in \text{Nat}]$.

Higher-order functions. It might seem that we lose the ability to promote higher-order functions, as $_ \rightarrow _$ requires a promotion as its first argument, which does not seem to be able to relate function types. But in fact we can deal with higher-order functions surprisingly straightforwardly, by injecting upgrades back into promotions:

$$\begin{aligned} \textit{fromUpgrade} & : \{X Y : \text{Set}\} (u : \textit{Upgrade} X Y) \rightarrow \\ & \text{Promotion } X (\Sigma[y \in Y] \Sigma[x \in X] \textit{Upgrade.Coherence} u x y) \\ \textit{fromUpgrade } u & = \mathbf{record} \\ & \{ \textit{Predicate} &= \textit{Upgrade.Predicate} u \\ & ; \textit{forget} &= \textit{proj}_1 \circ \textit{proj}_2 \\ & ; \textit{complement} &= \lambda \{ (y, x, \textit{coh}) \mapsto \textit{Upgrade.complement} u x y \textit{coh} \} \\ & ; \textit{promote} &= \lambda x \textit{com} \mapsto \textit{Upgrade.promote} u x \textit{com} , x , \\ & & \textit{Upgrade.coherence} u x \textit{com} \\ & ; \textit{coherence} &= \lambda _ _ \mapsto \textit{refl} \} \end{aligned}$$

This is a general formulation of the trick of “promoting to evidently coherent subsets” used in Section 4.2.2.

Here we give a nontrivial but not perfect example of promoting a higher-order function. Consider a version of $\textit{concatMap}$ specialised to lists whose elements are themselves lists:

$$\begin{aligned} \textit{concatMap} & : (\textit{List} A \rightarrow \textit{List} B) \rightarrow \textit{List} (\textit{List} A) \rightarrow \textit{List} B \\ \textit{concatMap } f & = \textit{concat} \circ \textit{map } f \end{aligned}$$

If we know that the function f preserves a particular length, and that all lists in the input list have that length, then we can determine the length of the output list in advance. Formally:

$$\begin{aligned} & (n : \text{Nat}) (f : \textit{List} A \rightarrow \textit{List} B) \rightarrow \\ & ((xs : \textit{List} A) \rightarrow \textit{length} xs \equiv n \rightarrow \textit{length} (f xs) \equiv n) \rightarrow \\ & (m : \text{Nat}) (xss : \textit{List} (\textit{List} A)) \rightarrow \\ & \textit{length} xss \equiv m \times \textit{All} (\lambda xs \rightarrow \textit{length} xs \equiv n) xss \rightarrow \\ & \textit{length} (\textit{concatMap } f xss) \equiv m * n \end{aligned}$$

This is exactly the type of promotion proof computed by the upgrade

$$\begin{aligned} & \forall^+[n \in \text{Nat}] (\textit{fromUpgrade} (\textit{List-Vec} A n \rightarrow \textit{List-Vec} A n) \rightarrow \\ & \forall^+[m \in \text{Nat}] (\textit{List}^2\textit{-Vec}^2 A m n \rightarrow \textit{toUpgrade} (\textit{List-Vec} B (m * n)))) : \end{aligned}$$




Fig. 3. *Left*: a binomial heap of size 3 consisting of two binomial trees storing elements 2, 5, and 8. *Right*: the result of inserting 11 into the heap. (Note that the digits of the underlying binary numbers are ordered with the least significant digit first.)

$$\begin{aligned}
& \text{Upgrade } ((\text{List } A \rightarrow \text{List } B) \rightarrow \text{List } (\text{List } A) \rightarrow \text{List } B) \\
& ((n : \text{Nat}) \rightarrow (\Sigma [g \in \text{Vec } A \ n \rightarrow \text{Vec } A \ n] \Sigma [f \in \text{List } A \rightarrow \text{List } A] \\
& \quad ((xs : \text{Vec } A \ n) \rightarrow \text{toList } (g \ xs) \equiv f \ (\text{toList } \ xs))) \rightarrow \\
& (m : \text{Nat}) \rightarrow \text{Vec } (\text{Vec } A \ n) \ m \rightarrow \text{Vec } B \ (m * n)
\end{aligned}$$

whose type shows the promoted function type, which works on vectors. (We omit the definition of $\text{List}^2\text{-Vec}^2 : (A : \text{Set}) (m \ n : \text{Nat}) \rightarrow \text{Promotion } (\text{List } (\text{List } A)) (\text{Vec } (\text{Vec } A \ n) \ m)$, which can be constructed either in an ad hoc way or somehow compositionally.) A promoted function, say concatMapV , will satisfy the coherence property

$$\begin{aligned}
& (n : \text{Nat}) (g : \text{Vec } A \ n \rightarrow \text{Vec } B \ n) (f : \text{List } A \rightarrow \text{List } B) \\
& \quad (c : (xs : \text{Vec } A \ n) \rightarrow \text{toList } (g \ xs) \equiv f \ (\text{toList } \ xs)) \rightarrow \\
& (m : \text{Nat}) (xss : \text{Vec } (\text{Vec } A \ n) \ m) \rightarrow \\
& \quad \text{toList } (\text{concatMapV } (g, f, c) \ xss) \equiv \text{concatMap } f \ (\text{map } \text{toList } (\text{toList } \ xss))
\end{aligned}$$

which is also computed, and proved, by the upgrade.

The (g, f, c) -triple argument to concatMapV looks rather bulky — why can it not just be g ? This has to do with the semantic way in which we promote functions: Note that a function of type $\text{Vec } A \ n \rightarrow \text{Vec } A \ n$ (which should be distinguished from $\{n : \text{Nat}\} \rightarrow \text{Vec } A \ n \rightarrow \text{Vec } A \ n$) is less general than a function of type $\text{List } A \rightarrow \text{List } A$. Since concatMapV computes by calling concatMap , it must somehow supply an extension of g — i.e., a function $f : \text{List } A \rightarrow \text{List } A$ such that $\text{toList } (g \ xs) \equiv f \ (\text{toList } \ xs)$ for all $xs : \text{Vec } A \ n$ — to concatMap ; here concatMapV simply requires such an extension as an additional argument. (This might not be as problematic as it seems: If we consistently use function promotion, that is, functional arguments to promoted functions are themselves promoted functions, then we can simply bundle the promoted functions with the basic functions and coherence proofs and pass the triples around.)

5 Case study: binomial heaps

Having gone through the previous sections on index-first inductive families, ornaments, and function promotion, we are now ready for our case study on implementing binomial heaps and their operations by relating them to binary numbers and numeric operations.

Every schoolchild is familiar with the idea of *positional number systems*, in which a number is represented as a list of digits. Each digit is associated with a weight, and the interpretation of the list is the weighted sum of the digits. For example, the weights used for binary numbers are powers of 2. Some container data structures and associated operations strongly resemble positional representations of natural numbers and associated operations. For example, a *binomial heap* (illustrated in Figure 3) can be thought of as a binary number in which every 1-digit stores a *binomial tree* — the actual vehicle for storing elements — whose size is exactly the weight of that digit. The number of elements stored in a binomial heap is therefore exactly the value of the underlying binary number. Inserting a new element into a binomial heap is analogous to incrementing a binary number, with carries corresponding to combining smaller binomial trees into larger ones. Okasaki (1999) thus proposed to design container data structures by analogy with positional representations of natural numbers, and called such data structures *numerical representations*. Using an ornament, it is easy to express the relationship between a numerically represented container datatype (e.g., binomial heaps) and its underlying numeric datatype (e.g., binary numbers). But the ability to express the relationship alone is not too surprising. What is more interesting is that the ornament can give rise to upgrades such that

- the coherence properties of the upgrades semantically characterise the resemblance between container operations and corresponding numeric operations, and
- the promotion predicates give the precise types of the container operations that guarantee such resemblance.

In this section, we show how ornaments induce insertion and extraction operations on binomial heaps, given increment and decrement operations on binary numbers. The induced extraction operation is not quite what is usually wanted, though, since the element it extracts is not necessarily the minimum element. We also show how to build a derived operation to extract the minimum element, by identifying a more refined datatype of binomial heaps, and we show how ornaments induce the conversions to and from this more refined datatype.

5.1 Binomial trees

The basic building blocks of binomial heaps are *binomial trees*, in which elements are stored. We assume that elements are drawn from a set Val on which there is a decidable ordering

$$_ \leq _ : Val \rightarrow Val \rightarrow Set$$

with comparison function

$$_ \leq ? _ : (x\ y : Val) \rightarrow (x \leq y) \uplus (y \leq x)$$




Fig. 4. *Left*: inductive definition of binomial trees. *Right*: decomposition of binomial trees of ranks 1 to 3.

Binomial trees are defined inductively on their *rank*, which is a natural number (see Figure 4):

- A binomial tree of rank 0 is a single node storing an element of type *Val*, and
- a binomial tree of rank $\text{suc } r$ consists of two binomial trees of rank r , with one attached under the other's root node.

From this definition we can readily deduce that a binomial tree of rank r has 2^r elements. To actually define binomial trees as a datatype, however, an alternative view is more helpful: A binomial tree of rank r is a root node and a forest of binomial trees of ranks 0 to $r - 1$. (Figure 4 shows how binomial trees of ranks 1 to 3 can be decomposed according to this view.) We could thus define a datatype $\text{BTree} : \text{Nat} \rightarrow \text{Set}$ of binomial trees — which is indexed with the rank — as follows: For any rank $r : \text{Nat}$, the type $\text{BTree } r$ has a field of type *Val*, namely the root node, and r recursive positions indexed from $r - 1$ down to 0. For our experiment, it is not necessary to encode the BTree datatype as a description, since we do not need any ornamental constructions involving BTree ; but BTree turns out to be an interesting example, so we present the encoding anyway:

$\text{BTreeD} : \text{Desc Nat}$

$\text{BTreeD } r = \sigma[_ \in \text{Val}] \vee (\text{descend } r)$

indexfirst data $\text{BTree} : \text{Nat} \rightarrow \text{Set}$ **where**

$\text{BTree } r \ni \text{node Val } (\mathbb{P} (\text{descend } r) \text{ BTree})$

This is an interesting example since we are exploiting the full computational power of Desc , computing the list of recursive indices from the index request r by $\text{descend } r$, which yields a list from $r - 1$ down to 0:

$\text{descend} : \text{Nat} \rightarrow \text{List Nat}$

$\text{descend zero} = []$

$\text{descend } (\text{suc } n) = n :: \text{descend } n$

For use in min-heaps, however, we should also ensure that elements in binomial trees are in *heap order*, i.e., the root of any binomial tree (including subtrees) is the minimum element in the tree. As demonstrated earlier with ordered lists (Section 3.2), ordering invariants can be easily encoded with inductive families. Here we go for a slightly more space-efficient

variant: Instead of a lower bound, we index the BTree datatype by a given root stored elsewhere. A more appropriate name for the revised datatype is thus BForest, since an inhabitant of the datatype now defines only a *forest* of binomial trees that are considered attached under the root specified as the index (which is not stored in the inhabitant itself):

indexfirst data BForest : Nat → Val → Set **where**

BForest $r\ x \ni \text{forest } (\mathbb{P} (\text{descend } r) (\lambda i \rightarrow \Sigma^1 [t \in \Sigma \text{Val } (\text{BForest } i)]\ x \leq \text{proj}_1\ t))$

(Σ^1 is a special dependent pair type former, which will be introduced below.) A binomial tree, then, is just a pair of its root and the forest of subtrees. We thus redefine BTree as a dependent pair type:

BTree : Nat → Set

BTree $r = \Sigma \text{Val } (\text{BForest } r)$

For clarity, we give alternative names to the two projection functions:

$\text{root} : \{r : \text{Nat}\} \rightarrow \text{BTree } r \rightarrow \text{Val}$

$\text{root} = \text{proj}_1$

$\text{children} : \{r : \text{Nat}\} (t : \text{BTree } r) \rightarrow \text{BForest } r (\text{root } t)$

$\text{children} = \text{proj}_2$

As an aside, the encoding of the BForest datatype as a description requires a bit of twist, but can still be done fairly easily. We can even define BForest as an ornamentation of the BTree datatype from the previous paragraph, but (again) this is not essential since we will exclusively use the BForest and BTree datatypes defined in this paragraph from now on, making no use of the relationship between BForest and the first version of BTree.

The most important operation on binomial trees is combining two smaller binomial trees of the same rank into a larger one, corresponding to carries in positional arithmetic. Given two binomial trees of the same rank r , one can be *attached* under the root of the other, forming a single binomial tree of rank $\text{suc } r$, corresponding exactly to the inductive definition of binomial trees. To establish heap ordering, we should also require a proof that that the root of the first tree is no less than the root of the second tree, and put this proof into the new tree. Here we adopt (a variation of) the technique (McBride, 2014) of using AGDA's *instance arguments* (Devriese & Piessens, 2011) to pass around such proofs implicitly. We define a special dependent pair type whose second component can be suppressed and automatically taken care of by AGDA's instance searching mechanism:

record $\Sigma^1 (A : \text{Set}) (B : A \rightarrow \text{Set}) : \text{Set} **where**$

constructor $\langle _ \rangle$

field

$\text{fst} : A$

$\{\{\text{snd}\}\} : B \text{fst}$

syntax $\Sigma^1 A (\lambda x \rightarrow B) = \Sigma^1 [x \in A] B$

Typically, we deconstruct a Σ^1 -typed element by matching it against the constructor $\langle _ \rangle$ without naming (and thus explicitly using) its second component, which will nevertheless be considered as an available instance during instance searching. On the other hand, when

we construct a new Σ^1 -typed element using the constructor $\langle _ \rangle$, supplying only its first component, its second component will be automatically determined to be the unique available instance having the right type. (Typechecking fails unless there is exactly one available instance having the right type.) This special dependent pair type was used to define BForest, and allows ordering proofs in elements of BForest to be handled implicitly by instance searching. We can now define *attach*, which takes two trees of the same rank and a proof of ordering of their roots, by

$$\begin{aligned} \text{attach} &: \{r : \text{Nat}\} (t u : \text{BTree } r) \{ _ : \text{root } u \leq \text{root } t \} \rightarrow \text{BTree } (\text{suc } r) \\ \text{attach } t (y, \text{forest } us) &= y, \text{forest } (\langle t \rangle, us) \end{aligned}$$

Using *attach* requires that we know how the roots of two trees are ordered in advance. To write a more general combining function, we first define the following proof-relevant version of **if_then_else_**:

$$\begin{aligned} \mathbf{if_then_else_} &: \{A B C : \text{Set}\} \rightarrow A \uplus B \rightarrow (\{ _ : A \} \rightarrow C) \rightarrow (\{ _ : B \} \rightarrow C) \rightarrow C \\ \mathbf{if } \text{inj}_1 a \mathbf{ then } t \mathbf{ else } u &= t \{ \{ a \} \} \\ \mathbf{if } \text{inj}_2 b \mathbf{ then } t \mathbf{ else } u &= u \{ \{ b \} \} \end{aligned}$$

and then we can compare the roots of two binomial trees to decide which will be attached under the other:

$$\begin{aligned} \text{link} &: \{r : \text{Nat}\} \rightarrow \text{BTree } r \rightarrow \text{BTree } r \rightarrow \text{BTree } (\text{suc } r) \\ \text{link } t u &= \mathbf{if } \text{root } t \leq_? \text{root } u \mathbf{ then } \text{attach } u t \mathbf{ else } \text{attach } t u \end{aligned}$$

Instance searching manages all ordering proofs behind the scenes, and makes it impossible to get the two branches the wrong way around, as remarked by McBride (2014) for his development.

5.2 From binary numbers to binomial heaps

The datatype `Bin` : Set of binary numbers is just a specialised datatype of lists of binary digits:

```
indexfirst data Bin : Set where
  Bin  $\ni$  nil
      | zero Bin
      | one Bin
```

This datatype can be encoded as follows:

```
data BinTag : Set where
  'nil : BinTag
  'zero : BinTag
  'one : BinTag

BinD : Desc  $\top$ 
BinD tt =  $\sigma$  BinTag  $\lambda$  { 'nil  $\mapsto$  v []
                          ; 'zero  $\mapsto$  v (tt :: [])
                          ; 'one  $\mapsto$  v (tt :: []) }
```

The intended interpretation of binary numbers is given by

$$\begin{aligned} \text{toNat} &: \text{Bin} \rightarrow \text{Nat} \\ \text{toNat nil} &= 0 \\ \text{toNat (zero } b) &= 0 + 2 \times \text{toNat } b \\ \text{toNat (one } b) &= 1 + 2 \times \text{toNat } b \end{aligned}$$

That is, binary numbers of type `Bin` are written least significant bit first, and the i -th digit (counting from zero) has weight 2^i . We refer to the position of a digit as its *rank*, i.e., the i -th digit is said to have rank i .

As stated in the opening of Section 5, binomial heaps are binary numbers whose 1-digits are decorated with binomial trees of matching rank, and they can be expressed straightforwardly as an ornamentation of binary numbers. To ensure that the binomial trees in binomial heaps have the correct rank, the datatype `BHeap` : `Nat` \rightarrow `Set` is indexed with the starting rank: If a binomial heap of type `BHeap` r is nonempty (i.e., not `nil`), then its first digit has rank r (and stores a binomial tree of rank r when the digit is one), and the rest of the heap is indexed with `suc` r :

$$\begin{aligned} \text{indexfirst data } \text{BHeap} &: \text{Nat} \rightarrow \text{Set} \text{ where} \\ \text{BHeap } r &\ni \text{ nil} \\ &| \text{ zero (BHeap (suc } r)) \\ &| \text{ one (BTree } r) \text{ (BHeap (suc } r)) \end{aligned}$$

This datatype can be encoded as

$$\begin{aligned} \text{BHeapOD} &: \text{OrnDesc Nat ! BinD} \\ \text{BHeapOD (ok } r) &= \sigma \text{ BinTag } \lambda \{ \text{'nil} \mapsto v \text{ tt} \\ & \quad ; \text{'zero} \mapsto v \text{ (ok (suc } r) , \text{ tt)} \\ & \quad ; \text{'one} \mapsto \Delta[t \in \text{BTree } r] v \text{ (ok (suc } r) , \text{ tt)} \} \end{aligned}$$

In applications, we would use binomial heaps of type `BHeap zero`, since this type encompasses binomial heaps of all sizes.

5.3 Increment and insertion

The increment operation on binary numbers is defined by

$$\begin{aligned} \text{incr} &: \text{Bin} \rightarrow \text{Bin} \\ \text{incr nil} &= \text{one nil} \\ \text{incr (zero } b) &= \text{one } b \\ \text{incr (one } b) &= \text{zero (incr } b) \end{aligned}$$

The corresponding operation on binomial heaps is insertion of a binomial tree into a binomial heap (of matching rank), whose direct implementation is

$$\begin{aligned} &\text{-- Later this function will be synthesised using function promotion} \\ \text{insT} &: \{r : \text{Nat}\} \rightarrow \text{BTree } r \rightarrow \text{BHeap } r \rightarrow \text{BHeap } r \\ \text{insT } t \text{ nil} &= \text{one } t \text{ nil} \end{aligned}$$

$$\begin{aligned} \mathit{insT} \ t \ (\mathit{zero} \ h) &= \mathit{one} \ t \ h \\ \mathit{insT} \ t \ (\mathit{one} \ u \ h) &= \mathit{zero} \ (\mathit{insT} \ (\mathit{link} \ t \ u) \ h) \end{aligned}$$

Conceptually, incr puts a 1-digit into the least significant position of a binary number, triggering a series of carries, i.e., summing 1-digits of smaller ranks into 1-digits of larger ranks; insT follows the pattern of incr , but since 1-digits now have to store a binomial tree of matching rank, insT takes an additional binomial tree as input and links binomial trees of smaller ranks into binomial trees of larger ranks whenever carries occur. Having defined insT , inserting a single element into a binomial heap of type $\mathit{BHeap} \ \mathit{zero}$ is then inserting, by insT , a rank-0 binomial tree (i.e., a single node) storing the element into the heap.

$$\begin{aligned} \mathit{insV} &: \mathit{Val} \rightarrow \mathit{BHeap} \ \mathit{zero} \rightarrow \mathit{BHeap} \ \mathit{zero} \\ \mathit{insV} \ x &= \mathit{insT} \ (x, \mathit{forest} \ \mathit{tt}) \end{aligned}$$

It is apparent that the program structure of insT strongly resembles that of incr — they manipulate the list-of-binary-digits structure in the same way. But can we characterise the resemblance semantically? It turns out that the coherence property of the following upgrade from the type of incr to that of insT is an appropriate answer:

$$\begin{aligned} \mathit{upg} &: \mathit{Upgrade} \left(\begin{array}{c} \mathit{Bin} \quad \rightarrow \mathit{Bin} \\ \{\mathit{r} : \mathit{Nat}\} \rightarrow \mathit{BTree} \ \mathit{r} \rightarrow \mathit{BHeap} \ \mathit{r} \rightarrow \mathit{BHeap} \ \mathit{r} \end{array} \right) \\ \mathit{upg} &= \forall^+ [\mathit{r} \in \mathit{Nat}] \ \forall^+ [_ \in \mathit{BTree} \ \mathit{r}] \ \mathit{Bin-BHeap} \ \mathit{r} \rightarrow \mathit{toUpgrade} \ (\mathit{Bin-BHeap} \ \mathit{r}) \end{aligned}$$

The upgrade upg says that, compared to the type of incr , the type of insT has two new arguments — the implicit argument $\mathit{r} : \mathit{Nat}$ (declared using $\forall^+ [\mathit{r} \in \mathit{Nat}]$), which is the same as $\forall^+ [\mathit{r} \in \mathit{Nat}]$ except that the new argument in the upgraded function type is implicit) and the explicit argument of type $\mathit{BTree} \ \mathit{r}$ — and that the two occurrences of $\mathit{BHeap} \ \mathit{r}$ in the type of insT refine the corresponding occurrences of Bin in the type of incr using the promotion

$$\begin{aligned} &\text{-- Induced from the ornament } [\mathit{BHeapOD}] \\ \mathit{Bin-BHeap} &: (\mathit{r} : \mathit{Nat}) \rightarrow \mathit{Promotion} \ \mathit{Bin} \ (\mathit{BHeap} \ \mathit{r}) \end{aligned}$$

induced by the ornament $[\mathit{BHeapOD}]$ (ok r) from Bin to $\mathit{BHeap} \ \mathit{r}$. The upgrade computes a coherence property, which instantiates for incr and insT to

$$\begin{aligned} &\text{-- Computed by } \mathit{Upgrade.Coherence} \ \mathit{upg} \ \mathit{incr} \ \mathit{insT} \\ \{\mathit{r} : \mathit{Nat}\} \ (t : \mathit{BTree} \ \mathit{r}) \ (b : \mathit{Bin}) \ (h : \mathit{BHeap} \ \mathit{r}) &\rightarrow \\ (\mathit{toBin} \ h \equiv b) &\rightarrow (\mathit{toBin} \ (\mathit{insT} \ t \ h) \equiv \mathit{incr} \ b) \end{aligned}$$

where toBin extracts the underlying binary number of a binomial heap:

$$\begin{aligned} \mathit{toBin} &: \{\mathit{r} : \mathit{Nat}\} \rightarrow \mathit{BHeap} \ \mathit{r} \rightarrow \mathit{Bin} \\ \mathit{toBin} &= \mathit{ornForget} \ [\mathit{BHeapOD}] \end{aligned}$$

That is, given a binomial heap $h : \mathit{BHeap} \ \mathit{r}$ whose underlying binary number is $b : \mathit{Bin}$, after inserting a binomial tree into h by insT , the underlying binary number of the result is $\mathit{incr} \ b$. This says exactly that insT manipulates the underlying binary number in the same way as incr .

We have seen that the coherence property computed by upg is appropriate for characterising the resemblance of incr and insT ; proving that it holds for incr and insT is a

separate matter, though. We can, however, avoid doing the implementation of insertion and the coherence proof separately: instead of implementing $insT$ directly, we can implement insertion with a more precise type in the first place such that, from this more precisely typed version, we can derive automatically a function $insT$ that satisfies the coherence property. This process is fully supported by the mechanism of upgrades. Specifically, the more precise type for insertion is given by the promotion predicate of upg (applied to $incr$), the more precisely typed version of insertion acts as a promotion proof of $incr$ (with respect to upg), and the promotion gives us $insT$, accompanied by a proof that $insT$ is coherent with $incr$.

Let $BHeap^U$ be the optimised promotion predicate for the ornament from Bin to $BHeap\ r$:

```

indexfirst data BHeapU : Nat → Bin → Set where
  BHeapU r nil      ⊃ nil
  BHeapU r (zero b) ⊃ zero (h : BHeapU (suc r) b)
  BHeapU r (one b)  ⊃ one (t : BTree r) (h : BHeapU (suc r) b)
BHeapU : Nat → Bin → Set
BHeapU r b = OptP [BHeapOD] (ok r) b

```

Here a more helpful interpretation is that $BHeap^U$ is a datatype of binomial heaps additionally indexed with the underlying binary number. The type of promotion proofs for $incr$ then expands to

```

-- Computed by Upgrade.Predicate upg incr
{r : Nat} → BTree r → (b : Bin) → BHeapU r b → BHeapU r (incr b)

```

A function of this type is explicitly required to transform the underlying binary number structure of its input in the same way as $incr$. Insertion can now be implemented as

```

insTU : {r : Nat} → BTree r → (b : Bin) → BHeapU r b → BHeapU r (incr b)
insTU t nil      nil      = one t nil
insTU t (zero b) (zero h) = one t h
insTU t (one b)  (one u h) = zero (insTU (link t u) b h)

```

which is very similar to the original $insT$. It is interesting to note that all the constructor choices for binomial heaps in $insT^U$ are actually completely determined by the types. This fact is easier to observe if we desugar $insT^U$ to its raw representation:

```

insTU : {r : Nat} → BTree r → (b : Bin) → BHeapU r b → BHeapU r (incr b)
insTU t (con ('nil , tt)) (con      tt ) = con (t , con tt      , tt)
insTU t (con ('zero , b , tt)) (con (  h , tt)) = con (t , h      , tt)
insTU t (con ('one , b , tt)) (con (u , h , tt)) = con ( insTU (link t u) b h , tt)

```

in which no constructor tags for binomial heaps are present. This means that the types would determine which constructors to use when programming $insT^U$, establishing the coherence property by construction. Indeed, when programming $insT^U$ in AGDA interactively, most of $insT^U$ can be directly generated by AGDA (upon instruction by the programmer). Finally, since $insT^U$ is a promotion proof for $incr$, we can invoke the promoting operation of upg and get $insT$:

```

insT : {r : Nat} → BTree r → BHeap r → BHeap r
insT = Upgrade.promote upg incr insTU

```

which is automatically coherent with *incr*:

$$\begin{aligned} \text{incr-insT-coherence} &: \{r : \text{Nat}\} (t : \text{BTree } r) \rightarrow \text{toBin } (\text{insT } t \ h) \equiv \text{incr } (\text{toBin } h) \\ \text{incr-insT-coherence} &= \text{Upgrade.coherence } \text{upg } \text{incr } \text{insT}^U \end{aligned}$$

To sum up: We define *Bin*, *incr*, and then *BHeap* as an ornamentation of *Bin*, describe in *upg* how the type of *insT* is an upgraded version of the type of *incr*, and implement *insT*^U, whose type is supplied by *upg*. We can then derive *insT*, the coherence property of *insT* with respect to *incr*, and its proof, all automatically using *upg*. Compared to Okasaki’s implementation (1999), besides rank-indexing, which elegantly transfers the management of rank-related invariants to the type system, the only extra work is the straightforward marking of the differences between *Bin* and *BHeap* (in *BHeapOD*) and between the type of *incr* and that of *insT* (in *upg*). The reward is large in comparison: We get a coherence property that precisely characterises the structural behaviour of insertion with respect to increment, and an enriched function type that guides the implementation of insertion such that the coherence property is satisfied by construction.

Similar to lifting increment to insertion, addition on binary numbers can also be lifted to merging of two heaps. The lifting process is the same, though, and is thus omitted from this paper.

5.4 Decrement and extraction

Binomial heaps would be useless if we could only insert but not *extract* elements. Thinking in terms of numerical representations, element extraction on binomial heaps corresponds to *decrement* on binary numbers. Decrement is slightly more complicated than increment since it is inherently a partial operation, applicable to non-zero numbers only. (Defining the result of decrementing zero as zero counts as a workaround at best, as it destroys the natural property that increment is left inverse to decrement.) We will explore the simplest approach, weakening the result type to *Maybe Bin*:

$$\begin{aligned} \text{decr} &: \text{Bin} \rightarrow \text{Maybe Bin} \\ \text{decr nil} &= \text{nothing} \\ \text{decr } (\text{zero } b) &= \text{mapMaybe one } (\text{decr } b) \\ \text{decr } (\text{one } b) &= \text{just } (\text{zero } b) \end{aligned}$$

where *mapMaybe* : {A B : Set} → (A → B) → Maybe A → Maybe B is the usual functorial map for *Maybe*.

Following the route taken in Section 5.3, we will promote *decr* : *Bin* → *Maybe Bin* to a function

$$\text{extract} : \{r : \text{Nat}\} \rightarrow \text{BHeap } r \rightarrow \text{Maybe } (\text{BTree } r \times \text{BHeap } r)$$

which divides a nonempty binomial heap of given rank *r* into a tree of rank *r* and a smaller heap of the same rank; if the heap has a rank-*r* zero digit, then the tree has to be “borrowed” from a higher-rank one digit.

The first step is to write an upgrade relating the types of *decr* and *extract*. The language of promotions and upgrades presented in Section 4 cannot deal with the sub-upgrade from


```

 $\_+ \times \_ :$  (X : Set) {Y Z : Set}  $\rightarrow$  Promotion Y Z  $\rightarrow$  Promotion Y (X  $\times$  Z)
X  $\_+ \times$  p = record
  { Predicate =  $\lambda$  y  $\rightarrow$  X  $\times$  Promotion.Predicate p y
  ; forget    = Promotion.forget p  $\circ$  proj2
  ; complement =  $\lambda$  { (x, z)  $\mapsto$  x, Promotion.complement p z }
  ; promote   =  $\lambda$  { y (x, com)  $\mapsto$  x, Promotion.promote p y com }
  ; coherence =  $\lambda$  { y (x, com)  $\mapsto$  Promotion.coherence p y com }
  }
data Maybe' {A : Set} (X : A  $\rightarrow$  Set) : Maybe A  $\rightarrow$  Set where
  nothing : Maybe' X nothing
  just     : {a : A}  $\rightarrow$  X a  $\rightarrow$  Maybe' X (just a)
MaybeP : {A B : Set}  $\rightarrow$  Promotion A B  $\rightarrow$  Promotion (Maybe A) (Maybe B)
MaybeP p = record
  { Predicate = Maybe' (Promotion.Predicate p)
  ; forget    = mapMaybe (Promotion.forget p)
  ; complement =  $\lambda$  { nothing  $\mapsto$  nothing
                    ; (just b)  $\mapsto$  just (Promotion.complement p b) }
  ; promote   =  $\lambda$  { nothing _  $\mapsto$  nothing
                    ; (just a) (just x)  $\mapsto$  just (Promotion.promote p a x) }
  ; coherence =  $\lambda$  { nothing _  $\mapsto$  refl
                    ; (just a) (just x)  $\mapsto$  cong just (Promotion.coherence p a x) }

```

Fig. 5. Two more promotion combinators for upgrading decrement

Maybe Bin to Maybe (BTree $r \times$ BHeap r), but we can easily extend the language with two combinators

```

 $\_+ \times \_ :$  (X : Set) {Y Z : Set}  $\rightarrow$  Promotion Y Z  $\rightarrow$  Promotion Y (X  $\times$  Z)
MaybeP : {A B : Set}  $\rightarrow$  Promotion A B  $\rightarrow$  Promotion (Maybe A) (Maybe B)

```

to support that, as shown in Figure 5. (The most important information in Figure 5 is how the *Predicate* and *forget* fields are defined for the two combinators; the rest of the code is presented simply to show that it is easy to complete the definitions.) These may not be the most generic combinators that we can define, but they show that the framework can be easily extended. The upgrade we write is then

```

upg : Upgrade ( Bin  $\rightarrow$  Maybe Bin )
         ({r : Nat}  $\rightarrow$  BHeap r  $\rightarrow$  Maybe (BTree r  $\times$  BHeap r))
upg =  $\forall^+$ [[r  $\in$  Nat]] (Bin-BHeap r  $\rightarrow$ 
                       toUpgrade (MaybeP (BTree r  $\_+ \times$  Bin-BHeap r)))

```

The promotion predicate for *decr* with respect to this upgrade, or rather, a more informative type for extraction, is

```

-- Computed by Upgrade.Predicate upg decr
{r : Nat} {b : Bin}  $\rightarrow$ 
BHeapU r b  $\rightarrow$  Maybe' ( $\lambda$  b'  $\mapsto$  BTree r  $\times$  BHeapU r b') (decr b)

```

where the result is nothing when *decr* $b =$ nothing, or just p for some pair $p :$ BTree $r \times$ BHeap^U $r b'$ when *decr* $b =$ just b' . (See Figure 5 for the definition of Maybe'.) A function of this type is guaranteed to produce a tree and a heap (wrapped in just as a pair) if and only

if *decr* executes successfully on the underlying number of the input heap, and when that is the case, the underlying number of the resulting heap will be the result of *decr*.

$$\begin{aligned}
\text{extract}^U &: \{r : \text{Nat}\} (b : \text{Bin}) \rightarrow \\
&\quad \text{BHeap}^U r b \rightarrow \text{Maybe}' (\lambda b' \mapsto \text{BTree } r \times \text{BHeap}^U r b') (\text{decr } b) \\
\text{extract}^U \text{ nil} &\quad \text{nil} = \text{nothing} \\
\text{extract}^U (\text{zero } b) &(\text{zero } h) = \text{mapMaybe}' (\lambda \{((x, t \triangleleft ts), h) \mapsto (x, \text{forest } ts), \\
&\quad \text{one } t h\}) \\
&\quad (\text{extract}^U b h) \\
\text{extract}^U (\text{one } b) &(\text{one } t h) = \text{just } (t, \text{zero } h)
\end{aligned}$$

where we introduce a pattern synonym **pattern** $_ \triangleleft _ t \text{ ts} = \text{forest } ((t), \text{ts})$ for separating the leftmost tree from a forest, and *mapMaybe'* is a carefully indexed version of functorial map on *Maybe'*:

$$\begin{aligned}
\text{mapMaybe}' &: \{A B : \text{Set}\} \{C : A \rightarrow \text{Set}\} \{D : B \rightarrow \text{Set}\} \{f : A \rightarrow B\} \rightarrow \\
&\quad (\{a : A\} \rightarrow C a \rightarrow D (f a)) \rightarrow \\
&\quad \{ma : \text{Maybe } A\} \rightarrow \text{Maybe}' C ma \rightarrow \text{Maybe}' D (\text{mapMaybe } f ma) \\
\text{mapMaybe}' g &\text{ nothing} = \text{nothing} \\
\text{mapMaybe}' g &(\text{just } c) = \text{just } (g c)
\end{aligned}$$

We now get a coherent-by-construction definition of *extract* via the upgrade:

$$\begin{aligned}
\text{extract} &: \{r : \text{Nat}\} \rightarrow \text{BHeap } r \rightarrow \text{Maybe } (\text{BTree } r \times \text{BHeap } r) \\
\text{extract} &= \text{Upgrade.promote upg decr extract}^U
\end{aligned}$$

where coherence is witnessed by

$$\begin{aligned}
&\text{decr-extract-coherence} : \\
&\quad \{r : \text{Nat}\} (h : \text{BHeap } r) \rightarrow \text{mapMaybe } (\text{toBin} \circ \text{proj}_2) (\text{extract } h) \equiv \text{decr } (\text{toBin } h) \\
&\text{decr-extract-coherence} = \text{Upgrade.coherence upg decr extract}^U
\end{aligned}$$

5.5 Minimum extraction

The type of the *extract* operation guarantees that it extracts from a heap of rank *r* a subtree of rank *r*, but places no constraints on which particular subtree of that rank. (In fact, the implementation picks the subtree containing the root out of the “first” tree present in the heap, that is, the one with lowest rank; but that fact is not captured in the type.) A more useful operation on min-heaps is to extract the minimum element of the heap. Because of the heap ordering property, this minimum element will be the root of some tree, but not necessarily of the first tree. But we can preprocess the heap to ensure that the minimum element is the root of the first tree; then *extract* will indeed return a subtree containing the minimum element as its root.

We say that a heap is *normalised* if the minimum element is the root of the first tree. A normalised heap is of course a heap; we express the normalisation property as a refined datatype BHeap^N of heaps, carrying an extra boolean index stating whether the heap is normalised. The relationship between ordinary heaps BHeap and this refinement is an

ornamentation. Actually, we do the ornamentation in two steps: first to a datatype BHeap^M that indexes a heap by its minimum element, and second to BHeap^N to state whether the root of the first tree has this minimum value. We can normalise a heap by swapping the first tree with an equally sized subtree of the tree with the minimum root. Finally, we get for free the forgetful mapping back from BHeap^N via BHeap^M to BHeap , which allows us to apply *extract*.

Compare this with the standard algorithm for extracting the minimum element from a binomial heap, namely (i) to find the tree with minimum root, (ii) separate it from the heap, (iii) extract its root, and (iv) merge its children back into the heap. This would have no connection with our *extract* function, though; since we have this operation at our disposal, it is appealing to be able to reuse it. With the standard algorithm, on the other hand, step (i) has no direct analogue in the numerical representation, so would need to be defined from first principles; and steps (ii) and (iv) amount to implementing numerical operations to set a specified 1-digit (say, the k -th digit) to 0 and then to add back $2^k - 1$. Our algorithm is arguably conceptually simpler, as well as benefitting from the leverage provided by ornaments.

In the interests of brevity, we will provide less detail from now on — as foreshadowed on page 10, we present this section as a kind of thought experiment, and rather than defining the ornaments and deriving from them the refined datatypes and the forgetful functions, we will directly present the results that would have arisen.

We start by refining BHeap to be indexed by its minimum element. A heap of type $\text{BHeap}^M r mv$ is essentially a $\text{BHeap } r$; but the additional index mv is nothing iff the heap is empty, and otherwise $mv = \text{just } v$ where v is the minimum element of the heap. The nil and zero constructors are analogous to those for BHeap . But the one constructor of BHeap is refined here into three constructors: *min* indicates that the minimum element is the root of the first tree, and there are more trees to follow; *sin* indicates that the minimum element is again the root of the first tree, but that there are no more trees to follow; and *one* indicates that the minimum element is the root of a higher-ranked tree, not the first tree. Note that in the *min* and *sin* cases, the element v is not duplicated: Instead of storing the original tree, they store only its forest of children.

```

indexfirst data BHeapM : Nat → Maybe Val → Set where
  BHeapM r nothing ⊃ nil
  BHeapM r (just v) ⊃ min (BForest r v) {v' : Val} {{v ≤ v'}}
                    (BHeapM (suc r) (just v'))
                    | sin (t : BForest r v) (BHeapM (suc r) nothing)
                    | one (t : BTree r) {v ≤ root t} (BHeapM (suc r) (just v))
  BHeapM r mv ⊃ zero (BHeapM (suc r) mv)

```

The function *toBHeap^M* lifts an ordinary BHeap to a more informative BHeap^M :

```

toBHeapM : {r : Nat} → BHeap r → Σ (Maybe Val) (BHeapM r)
toBHeapM nil = ⟨ nil ⟩
toBHeapM (zero h) = let ⟨ h' ⟩ = toBHeapM h in ⟨ zero h' ⟩
toBHeapM (one t h) with toBHeapM h

```

$$\begin{aligned} \text{toBHeap}^M (\text{one } t \ h) \mid \text{nothing } , h' &= \langle \langle \text{sin } (\text{children } t) \ h' \rangle \rangle \\ \text{toBHeap}^M (\text{one } t \ h) \mid \text{just } v \ , h' &= \mathbf{if} \ \text{root } t \leq_? v \ \mathbf{then} \ \langle \langle \text{min } (\text{children } t) \ h' \rangle \rangle \\ &\quad \mathbf{else} \ \langle \langle \text{one } t \ h' \rangle \rangle \end{aligned}$$

We use the pattern synonym **pattern** $\langle \langle - \rangle \rangle b = -, b$ to hide the first components, which are indices in the types of the heaps and are inferable. Note that the dependent pair representation $\Sigma (\text{Maybe } Val) (\text{BHeap}^M r)$ provides constant-time access to the minimum element, by first projection from the dependent pair. The function toBHeap^M does some real work, refining the original one constructor into three, and cannot be expected to be automatically derivable; on the other hand, if we had expressed BHeap^M as an ornamentation of BHeap , we would have been able to derive for free the following forgetful function:

```
-- Derivable from an ornament from BHeap to BHeapM
fromBHeapM : {r : Nat} {mv : Maybe Val} → BHeapM r mv → BHeap r
fromBHeapM nil          = nil
fromBHeapM (min ts h) = one ⟨ ts ⟩ (fromBHeapM h)
fromBHeapM (sin ts h) = one ⟨ ts ⟩ (fromBHeapM h)
fromBHeapM (one t h)  = one t (fromBHeapM h)
fromBHeapM (zero h)   = zero (fromBHeapM h)
```

The reader may have noticed that the sin constructor of BHeap^M takes a trailing heap which is necessarily empty because of the index of its type. Could we not just get rid of this seemingly redundant argument? This, in fact, points out a limitation of the ornament framework (by design): Ornamentation does not change the recursive structure. The sin constructor is a refinement of the one constructor of BHeap , and has to take the same number of recursive arguments as the latter. In a sense, this problem is not one inherent to the ornament framework, but originates from our initial choice of allowing binomial heaps to be represented with redundant zeros; that is, if we started with a non-redundant representation, then sin would not have to take that argument. On the other hand, the ornament framework is indeed limited, and cannot help us to refine a redundant representation into a non-redundant one.

The next step is to introduce a further ornamentation so we can tell from the index in the type whether a binomial heap is normalised or not — in fact, it is normalised iff the first non-zero constructor is not one:

```
indexfirst data BHeapN : Nat → Maybe Val → Bool → Set where
  BHeapN r nothing true  ⊃ nil
  BHeapN r (just v) true  ⊃ min (BForest r v) {v' : Val} {{v ≤ v'}}
                                {norm : Bool} (BHeapN (suc r) (just v') norm)
                                | sin (t : BForest r v) {norm : Bool}
                                (BHeapN (suc r) nothing norm)
  BHeapN r (just v) false ⊃ one (t : BTree r) {{v ≤ root t}}
                                {norm : Bool} (BHeapN (suc r) (just v) norm)
  BHeapN r mv          norm ⊃ zero (BHeapN (suc r) mv norm)
```

As before, if we had defined BHeap^N as an ornamentation of BHeap^M , then we would have got for free the forgetful function:

```

-- Derivable from an ornament from BHeapM to BHeapN
fromBHeapN : {r : Nat} {mv : Maybe Val} {norm : Bool} →
  BHeapN r mv norm → BHeapM r mv
fromBHeapN nil = nil
fromBHeapN (min ts h) = min ts (fromBHeapN h)
fromBHeapN (sin ts h) = sin ts (fromBHeapN h)
fromBHeapN (one t h) = one t (fromBHeapN h)
fromBHeapN (zero h) = zero (fromBHeapN h)

```

In addition, the relationship between BHeap^M and BHeap^N (at least in a non-indexed-first variant on the boolean index) turns out to be an *algebraic ornamentation* (McBride, 2011; Ko & Gibbons, 2013b), so we could get not only the forgetful function from BHeap^N to BHeap^M but also its converse:

```

-- Derivable from an algebraic ornament from BHeapM to BHeapN
toBHeapN : {r : Nat} {mv : Maybe Val} → BHeapM r mv → Σ Bool (BHeapN r mv)
toBHeapN nil = ⟨ nil ⟩
toBHeapN (min t h) = let ⟨ h' ⟩ = toBHeapN h in ⟨ min t h' ⟩
toBHeapN (sin t h) = let ⟨ h' ⟩ = toBHeapN h in ⟨ sin t h' ⟩
toBHeapN (one t h) = let ⟨ h' ⟩ = toBHeapN h in ⟨ one t h' ⟩
toBHeapN (zero h) = let ⟨ h' ⟩ = toBHeapN h in ⟨ zero h' ⟩

```

Gap 4 (representation optimisation for algebraic ornamentation)

None of the existing formulations of general algebraic ornamentation (McBride, 2011; Dagand & McBride, 2014; Ko & Gibbons, 2013b) takes representation optimisation into account — they all work by inserting some equality constraints as fields into descriptions, and cannot further classify constructors in an index-first manner, which happened for BHeap^N. (Dagand & McBride (2014)’s reornaments, while being algebraic and achieving representation optimisation, are specialised only to the ornamental algebras.) Ko & Gibbons (2013b) noted that, in order to fully exploit the optimising power of index-first inductive families, we might need to switch to *coalgebraic ornamentation* (for inductive datatypes), which remains to be investigated.

Now, normalisation takes a binomial heap with information about ordering among the roots of the binomial trees and turns it into a normalised binomial heap. Note that working on BHeap^M instead of BHeap means that all ordering information is reified through different constructors, so we can simply use pattern matching to determine whether we are looking at a binomial tree whose root is the minimum, instead of having to perform comparisons and deal with ordering proofs on the fly.

```

normalise : {r : Nat} {mv : Maybe Val} → BHeapM r mv → BHeapN r mv true
normalise nil = nil
normalise (min t h) = let ⟨ h' ⟩ = toBHeapN h in min t h'
normalise (sin t h) = let ⟨ h' ⟩ = toBHeapN h in sin t h'
normalise (one t h) with normalise-aux h
normalise (one t h) | us , f with f t

```

$$\begin{aligned} \text{normalise } (\text{one } t \ h) \mid us, _ \mid \langle \langle h' \rangle \rangle &= \text{min } us \ h' \\ \text{normalise } (\text{zero } h) &= \text{zero } (\text{normalise } h) \end{aligned}$$

We introduce another pattern synonym **pattern** $\langle _ \rangle b = b, _$ to hide inferable second components. The only complicated case in *normalise* is for the one constructor, where the root of the first tree is not the minimum; this case is handled by an auxiliary function *normalise-aux*. This function takes the trailing heap, extracts the forest under the minimum root (which is determined by the index), and also returns the rest of the heap but with a hole into which the first tree can be inserted. The size and ordering constraints are carefully encoded in the type of the function and kept track of by AGDA; with some effort we can hide all this extra information to some extent, and present the program almost as if it were simply typed.

$$\begin{aligned} \text{normalise-aux} : & \\ \{r : \text{Nat}\} \{v : \text{Val}\} \rightarrow & \text{BHeap}^M (\text{suc } r) (\text{just } v) \rightarrow \\ & \text{BForest } r \ v \times \\ & ((u : \text{BTree } r) \{ _ : v \leq \text{root } u \} \rightarrow \\ & \Sigma (\text{Val} \times \text{Bool}) \lambda \{ (v', \text{norm}) \rightarrow \Sigma^1 [_ \in \text{BHeap}^N (\text{suc } r) (\text{just } v') \text{norm}] \ v \leq v' \}) \\ \text{normalise-aux } (\text{min } (t \triangleleft ts) \{v'\} \ h) &= \\ \text{forest } ts, \lambda u \mapsto & \text{let } \langle u' \rangle = \text{link}' \ t \ u \\ & \langle \langle h' \rangle \rangle = \text{toBHeap}^N \ h \\ & \text{in if } \text{root } u' \leq_? v' \text{ then } \langle \langle \text{min } (\text{children } u') \ h' \rangle \rangle \\ & \text{else } \langle \langle \text{one } u' \ h' \rangle \rangle \\ \text{normalise-aux } (\text{sin } (t \triangleleft ts) \ h) &= \\ \text{forest } ts, \lambda u \mapsto & \text{let } \langle u' \rangle = \text{link}' \ t \ u \\ & \langle \langle h' \rangle \rangle = \text{toBHeap}^N \ h \\ & \text{in } \langle \langle \text{sin } (\text{children } u') \ h' \rangle \rangle \\ \text{normalise-aux } (\text{one } t \ h) \text{ with } & \text{normalise-aux } h \\ \text{normalise-aux } (\text{one } t \ h) \mid (t' \triangleleft ts), f &= \\ \text{forest } ts, \lambda u \mapsto & \text{let } \langle u' \rangle = \text{link}' \ t' \ u \\ & ((v', _), \langle h' \rangle) = f \ u' \\ & \text{in if } \text{root } t \leq_? v' \text{ then } \langle \langle \text{min } (\text{children } t) \ h' \rangle \rangle \\ & \text{else } \langle \langle \text{one } t \ h' \rangle \rangle \\ \text{normalise-aux } (\text{zero } h) \text{ with } & \text{normalise-aux } h \\ \text{normalise-aux } (\text{zero } h) \mid (t \triangleleft ts), f &= \\ \text{forest } ts, \lambda u \mapsto & \text{let } \langle u' \rangle = \text{link}' \ t \ u \\ & \langle \langle h' \rangle \rangle = f \ u' \\ & \text{in } \langle \langle \text{zero } h' \rangle \rangle \end{aligned}$$

Here, *link'* is a more informative version of the function *link* from Section 5.1, not only linking together two trees of some rank r to a single tree of rank $\text{suc } r$, but also combining lower bounds for the two input roots into a lower bound for the output root:

$$\begin{aligned} \text{link}' : \{r : \text{Nat}\} (t' : \text{BTree } r) \{v : \text{Val}\} \{ _ : v \leq \text{root } t \} \{ _ : v \leq \text{root } t' \} \rightarrow \\ \Sigma^1 [u \in \text{BTree } (\text{suc } r)] \ v \leq \text{root } u \end{aligned}$$

When used in the context of *normalise* (*one t h*), the type of *normalise-aux* declares that it takes h and returns a pair, whose first component is the children of the minimum-rooted

tree, and whose second component is a heap with a hole able to accommodate t . We will not go through the implementation, but the rank and ordering constraints implicit in the uses of the $\langle _ \rangle$ constructor and the $\langle\langle _ \rangle\rangle$ pattern synonym mean that it would be hard to make mistakes; the only constraints that are not currently expressed are that the output is a permutation of the input, for which some kind of linear typing seems appropriate, as noted by McBride (2014).

Gap 5 (hiding inferable terms)

The most important result that function *normalise-aux* returns is a BHeap^N , but this is buried inside several levels of nested tuples, which provide type indices and ordering proofs. Similar things can be said for *link'* and *toBHeap^N*. To prevent the program from being cluttered by these indices and proofs, we have used the $\langle _ \rangle$ constructor and the $\langle\langle _ \rangle\rangle$ pattern synonym to make them less noticeable typographically. This points to the need for some mechanism for eliding such ancillary data, and a nice implementation of dependent intersection types (Kopylov, 2003) could be helpful here. Bernardy & Moulin’s work on “type theory in colour” (2013) might also be helpful, but it does not solve our problem directly, because in our index-first setting the indices are used in a computationally relevant way.

Finally we can implement *extract-min* by turning a BHeap into a BHeap^M (i.e., establishing ordering among the roots), normalising it to a BHeap^N , and erasing it to a BHeap , whose first tree contains the minimum as its root, and then applying *extract*.

$$\begin{aligned} \text{extract-min} &: \{r : \text{Nat}\} \rightarrow \text{BHeap } r \rightarrow \text{Maybe } (\text{BTree } r \times \text{BHeap } r) \\ \text{extract-min} &= \text{extract} \circ \text{fromBHeap}^M \circ \text{fromBHeap}^N \circ \text{normalise} \circ \text{proj}_2 \circ \text{toBHeap}^M \end{aligned}$$

We remark that, up to now, we have not actually proved that the value extracted by *extract-min* is indeed the minimum. That would require us to start from basic definitions such as “the minimum element in a heap”; the whole proof does not have much to do with ornamentation, so we omit it here (but include it in the supplementary code). But it is worth noting that the sophisticated use of indexed datatypes does facilitate the proof — the indices carry enough information for most properties to be established by straightforward induction on them (i.e., simple recursive programs defined by pattern matching).

6 Discussion

This paper stems from a part of Chapter 3 of the first author’s DPhil dissertation (Ko, 2014), which in turn is partly based on both authors’ article in *Progress in Informatics* (Ko & Gibbons, 2013a). The definition of the optimised promotion predicate in terms of parallel composition first appeared in Ko & Gibbons (2013a), and the mechanism of promotions and upgrades is inspired by Dagand & McBride (2014) and developed in Ko (2014) as a more straightforward and flexible alternative. In this paper we make a slight simplification of the function promotion mechanism, and present one way to deal with higher-order functions. Ko (2014) uses insertion into binomial heaps as an example; that example is fully developed

in this paper, by also lifting decrement to extraction and then to minimum extraction via adding a normalising step.

Williams *et al.* (2014) have also conducted a case study on ornaments, lifting OCaml library functions semi-automatically to more informative datatypes. Their case study is larger, but mainly deals with only non-dependent types. The real power of ornaments lies in its ability to relate inductive families, and we believe that the case study about binomial heaps — which tries to push programming with inductive families to an extreme — has helped to point out how ornaments can help in a full dependently typed setting, and what conveniences are still lacking to support programming with ornaments.

We have seen that, in the case study on binomial heaps, the ornament framework is capable of generating all sorts of types and operations that help the programmer to move between more informative and less informative variants of datatypes, namely Bin (binary numbers), BHeap (binomial heaps indexed with starting rank), BHeap^U (BHeap indexed additionally with the underlying binary number), BHeap^M (BHeap indexed additionally with the minimum element), and BHeap^N (BHeap^M indexed additionally with a normalisation flag). The ornament framework also effectively guides the lifting of increment and decrement to insertion and extraction, and establishes coherence by construction. Finally, by carefully designing the datatypes BHeap^M and BHeap^N and employing McBride (2014)’s instance argument technique for managing ordering witnesses, we are able to express the heap normalisation algorithm (for minimum extraction) in a syntactically simple program which is nevertheless guaranteed by its type to maintain the rank and ordering constraints.

During our development, we have also encountered some inconveniences, which we identify as “gaps”. Of the five gaps we identify (on pages 8, 10, 15, 37, and 39), Gaps 1, 2, and 3 are about higher-level presentation of index-first inductive families and ornaments, and Gap 5 is about hiding inferable terms in programs — these all call for better language support, so that we could read and write dependently typed programs more easily. The remaining Gap 4 concerns coalgebraic ornamentation, which calls for further investigation into programming methodology. We look forward to advances that can fill these gaps.

Acknowledgements

Most of this work was done while the first author was in Oxford, supported by a Clarendon Scholarship and by the UK EPSRC-funded project *Reusability and Dependent Types* (EP/G034516/1). The first author then relocated to Japan and is supported by the Japan Society for the Promotion of Science (JSPS) Grant-in-Aid for Scientific Research (A) No. 25240009. We are indebted to Conor McBride for inspiration, not least in introducing us to ornaments and index-first datatypes. We would also like to thank Ralf Hinze and Venanzio Capretta for their insightful comments on the binomial heaps development during the first author’s DPhil viva. Thanks are also due to the anonymous referees for their thorough reviewing and invaluable comments.

References

- Bernardy, Jean-Philippe, & Moulin, Guilhem. (2013). Type theory in color. *Pages 61–72 of: International Conference on Functional Programming*. ICFP'13. ACM.
- Bove, Ana, & Dybjer, Peter. (2009). Dependent types at work. *Pages 57–99 of: Language Engineering and Rigorous Software Development*. Lecture Notes in Computer Science, vol. 5520. Springer-Verlag.
- Brady, Edwin, McBride, Conor, & McKinna, James. (2004). Inductive families need not store their indices. *Pages 115–129 of: Types for Proofs and Programs*. Lecture Notes in Computer Science, vol. 3085. Springer-Verlag.
- Chapman, James, Dagand, Pierre-Évariste, McBride, Conor, & Morris, Peter. (2010). The gentle art of levitation. *Pages 3–14 of: International Conference on Functional Programming*. ICFP'10. ACM.
- Cockx, Jesper, Devriese, Dominique, & Piessens, Frank. (2014). Pattern matching without K. *Pages 257–268 of: International Conference on Functional Programming*. ICFP'14. ACM.
- Czarnecki, Krzysztof, Foster, J. Nathan, Hu, Zhenjiang, Lämmel, Ralf, Schürr, Andy, & Terwilliger, James F. (2009). Bidirectional transformations: A cross-discipline perspective. *Pages 260–283 of: International Conference on Model Transformation*. Lecture Notes in Computer Science, vol. 5563. Springer-Verlag.
- Dagand, Pierre-Évariste, & McBride, Conor. (2012). *Elaborating inductive definitions*. arXiv:1210.6390.
- Dagand, Pierre-Évariste, & McBride, Conor. (2014). Transporting functions across ornaments. *Journal of functional programming*, **24**(2–3), 316–383.
- Devriese, Dominique, & Piessens, Frank. (2011). On the bright side of type classes: Instance arguments in Agda. *Pages 143–155 of: International Conference on Functional Programming*. ICFP'11. ACM.
- Dybjer, Peter. (1994). Inductive families. *Formal aspects of computing*, **6**(4), 440–465.
- Goguen, Healfdene, McBride, Conor, & McKinna, James. (2006). Eliminating dependent pattern matching. *Pages 521–540 of: Algebra, Meaning, and Computation*. Lecture Notes in Computer Science, vol. 4060. Springer-Verlag.
- Ko, Hsiang-Shang. (2014). *Analysis and synthesis of inductive families*. DPhil thesis, University of Oxford.
- Ko, Hsiang-Shang, & Gibbons, Jeremy. (2013a). Modularising inductive families. *Progress in informatics*, **10**, 65–88.
- Ko, Hsiang-Shang, & Gibbons, Jeremy. (2013b). Relational algebraic ornaments. *Pages 37–48 of: Dependently Typed Programming*. DTP'13. ACM.
- Kopylov, Alexei. (2003). Dependent intersection: A new way of defining records in type theory. *Pages 86–95 of: Logic in Computer Science*. LICS'03. IEEE.
- Martin-Löf, Per. (1984). *Intuitionistic type theory*. Bibliopolis, Napoli.
- McBride, Conor. (2011). *Ornamental algebras, algebraic ornaments*.
- McBride, Conor. (2014). How to keep your neighbours in order. *Pages 297–309 of: International Conference on Functional Programming*. ICFP'14. ACM.
- McBride, Conor, & McKinna, James. (2004). The view from the left. *Journal of functional programming*, **14**(1), 69–111.
- Norell, Ulf. (2007). *Towards a practical programming language based on dependent type theory*. Ph.D. thesis, Chalmers University of Technology.
- Norell, Ulf. (2009). Dependently typed programming in Agda. *Pages 230–266 of: Advanced Functional Programming*. Lecture Notes in Computer Science, vol. 5832. Springer-Verlag.
- Okasaki, Chris. (1999). *Purely functional data structures*. Cambridge University Press.

- Sheard, Tim, & Linger, Nathan. (2007). Programming in Ω mega. *Pages 158–227 of: Central European Functional Programming School*. Lecture Notes in Computer Science, vol. 5161. Springer-Verlag.
- Williams, Thomas, Dagand, Pierre-Évariste, & Rémy, Didier. (2014). Ornaments in practice. *Pages 15–24 of: Workshop on Generic Programming*. WGP'14. ACM.