

# Inductive Families Need Not Store Their Indices

Edwin Brady<sup>1</sup>, Conor McBride<sup>1</sup> and James McKinna<sup>2</sup>

<sup>1</sup> Department of Computer Science, University of Durham

<sup>2</sup> School of Computer Science, University of St. Andrews

**Abstract.** We consider the problem of efficient representation of dependently typed data. In particular, we consider a language TT based on Dybjer’s notion of **inductive families** [11] and reanalyse their general form with a view to optimising the storage associated with their use. We introduce an execution language, ExTT, which allows the commenting out of computationally irrelevant subterms and show how to use properties of elimination rules to elide constructor arguments and tags in ExTT. We further show how some types can be collapsed entirely at run-time. Several examples are given, including a representation of the simply typed  $\lambda$ -calculus for which our analysis yields an 80% reduction in run-time storage requirements.

## 1 Introduction

Dependent type theory provides programmers with more than an integrated logic for reasoning about program correctness. It allows more precise types for programs and data in the first place, strengthening the typechecker’s language of guarantees. We have richer function types  $\forall x : S. T$  which adapt their return types to each argument; we also have richer data structures which do not just contain but *explain* data, exposing and enforcing their properties.

Moreover, we may reasonably expect more static detail about programs and data to yield better optimised dynamic behaviour. We need neither test what is guaranteed nor store what is determined by typechecking. Pollack’s implicit syntax [24] already supports the omission of much redundant information from concrete syntax for similar reasons.

This paper identifies some space optimisations which significantly reduce the storage overheads associated with **inductive families** in the sense of [11]. These are data-indexed collections of mutually recursive datatypes,  $D\vec{x}$ , available in systems such as those underlying LEGO [16], COQ [9], ALF [18] and also the language we use here — EPIGRAM [20]. A common example for illustrative purposes is **Vect**, the family of list types indexed by element type and length:

$$\text{data } \frac{A : \star \quad n : \mathbb{N}}{\text{Vect } A \ n : \star} \quad \text{where } \frac{}{\varepsilon : \text{Vect } A \ 0} \quad \frac{a : A \quad v : \text{Vect } A \ k}{a :: v : \text{Vect } A \ (s \ k)}$$

## 1.1 Programming with Inductive Families

Function types over inductive families can use specific indices to require and ensure properties of inputs and outputs — e.g., compatibility of length:

$$\text{let } \frac{u, v : \mathbf{Vect} \mathbb{N} n}{\mathbf{vAdd} \ u \ v : \mathbf{Vect} \mathbb{N} n} \quad \mathbf{vAdd} \ \varepsilon \quad \varepsilon \quad \mapsto \ \varepsilon \\ \mathbf{vAdd} \ (x :: u) \ (y :: v) \mapsto \ (x + y) :: \mathbf{vAdd} \ u \ v$$

The precise type prevents some bogus choices of output — we can only return  $\varepsilon$  on the first line, only a  $::$  on the second. The input possibilities become narrower too — adding  $(x :: u)$  to  $\varepsilon$ , or vice versa, is not even an issue.

By the same token, the potential for optimisation is clear. Once we know whether the first argument is  $\varepsilon$  or  $(x :: u)$ , we can *presuppose* the form of the second argument — we can ignore it in the  $\varepsilon$  case; in the  $::$  case, we can safely project out  $y$  and  $v$  without checking the constructor tag. Moreover, if we inspect  $n$ , implicitly passed to  $\mathbf{vAdd}$ , we need never check  $\mathbf{Vect}$  constructor tags at all.

We impose invariants on inductive families to improve reliability, but this paper seeks to exploit them for performance. Such optimisations are not available in conventional functional languages — there is no way that inspecting one argument can justify presuppositions about another. If we want to write vector addition using ordinary lists, we must not only consider how to handle length mismatch in our code, we must also effectively test for it at run-time.

## 1.2 Underlying Type Theory

Following [20], EPIGRAM programs elaborate to well typed terms in a type theory TT, based on Luo’s UTT [15] with inductive families [11] and equality as in [19]. Here is its syntax:

$t ::= \star_i$	(type of types)		$x$	(variable)	
	$\forall x : t. t$	(function space)		$\mathbf{D}$	(inductive family)
	$\lambda x : t. t$	(abstraction)		$\mathbf{c}$	(constructor)
	$t \ t$	(application)		$\mathbf{D-E}$	(elimination operator)

As usual, we may abbreviate the function space  $\forall x : S. T$  by  $S \rightarrow T$  if  $x$  is not free in  $T$ . There is an infinite hierarchy of predicative universes,  $\star_i : \star_{i+1}$ . We leave universe levels to the machine, as in [12].

Computation is by  $\beta$ -**reduction** for  $\lambda$ -abstractions and  $\iota$ -**reduction** for elimination operators. A **data** declaration typically elaborates to declarations of a family  $\mathbf{D} : \forall \vec{i} : \vec{I}. \star$ , constructors  $\mathbf{c}$ , and an elimination operator  $\mathbf{D-E}$  equipped with  $\iota$ -schemes. We write  $s \mapsto t$  if  $s$   $\beta$ - or  $\iota$ -reduces to  $t$ . In the usual way, every well typed TT term  $t$  computes to a **weak head-normal form**  $\text{WHNF}(t)$ .

A typical constructor has a type like this:<sup>1</sup>

$$\mathbf{c} : \forall \vec{a} : \vec{A}. \mathbf{D} \ \vec{r}_1 \rightarrow \dots \rightarrow \mathbf{D} \ \vec{r}_j \rightarrow \mathbf{D} \ \vec{s}$$

<sup>1</sup> To ease presentation, we keep the non-recursive arguments  $\vec{a}$  to the front and permit only first-order recursive arguments — neither restriction is crucial to this work.

For example, elaborating our `Vect` example, we acquire  $\text{Vect} : \forall A : \star. \forall n : \mathbb{N}. \star$ , and constructors:

$$\begin{aligned} \varepsilon &: \forall A : \star. \text{Vect } A \ 0 \\ :: &: \forall A : \star. \forall k : \mathbb{N}. \forall a : A. \forall v : \text{Vect } A \ k. \text{Vect } A \ (sk) \end{aligned}$$

Note that the variables left schematic in the `data` declaration have become explicitly quantified arguments. In naïve implementations these take up space — every  $\text{Vect } A \ n$  stores the sequence  $0, \dots, n-1$ , and  $n$  references to  $A$ . Even with perfect sharing, this is quite an overhead — the space implications for families with more complex invariants are quite drastic if this problem is left unchecked.

Basically, the elimination operator, **D-E** has a type of this form:

$$\left. \begin{aligned} \forall \vec{i} : \vec{I}. \forall x : D \ \vec{i}. & \quad \text{(indices, target)} \\ \forall P : \forall \vec{i} : \vec{I}. D \ \vec{i} \rightarrow \star. & \quad \text{(motive)} \\ \forall m_c : \forall \vec{a} : \vec{A}. \forall y_1 : D \ \vec{r}_1. \dots \quad \forall y_j : D \ \vec{r}_j. & \\ \quad P \ \vec{r}_1 \ y_1 \rightarrow \dots \rightarrow P \ \vec{r}_j \ y_j \rightarrow P \ \vec{s} \ (c \ \vec{a} \ \vec{y}). & \quad \text{(methods)} \\ \dots & \\ P \ \vec{i} \ x & \end{aligned} \right\}$$

The **target**, with given **indices**, explains what to eliminate; the **motive** explains what is to be achieved by the elimination; the **methods** explain how to achieve the motive for each canonical form the target can take, given appropriate inductive hypotheses. The associated  $\iota$ -rules for definitional equality have this form:

$$\Gamma \vdash \text{D-E } \vec{s} \ (c \ \vec{a} \ \vec{y}) \ P \ \vec{m} = m_c \ \vec{a} \ \vec{y} \ (\text{D-E } \vec{r}_1 \ y_1 \ P \ \vec{m}) \ \dots \ (\text{D-E } \vec{r}_j \ y_j \ P \ \vec{m})$$

When indices are used uniformly, such as the element type of `Vect`, we adapt the basic **D-E** slightly, abstracting these parameters once for all. This yields:

$$\begin{aligned} \text{Vect-E} &: \forall A : \star. \forall n : \mathbb{N}. \forall v : \text{Vect } A \ n. \\ &\quad \forall P : \forall n : \mathbb{N}. \forall v : \text{Vect } A \ k. \star. \\ &\quad \forall m_\varepsilon : P \ 0 \ (\varepsilon \ A). \\ &\quad \forall m_\cdot : \forall k : \mathbb{N}. \forall a : A. \forall v : \text{Vect } A \ k. (P \ k \ v) \rightarrow P \ (sk) \ (:\!:\ A \ k \ a \ v). \\ &\quad P \ n \ v \\ \text{Vect-E } A \ 0 &\quad (\varepsilon \ A) \quad P \ m_\varepsilon \ m_\cdot = m_\varepsilon \\ \text{Vect-E } A \ (s \ k) &\quad (\!:\!:\ A \ k \ a \ v) \ P \ m_\varepsilon \ m_\cdot = m_\cdot \ k \ a \ v \ (\text{Vect-E } A \ k \ v \ P \ m_\varepsilon \ m_\cdot) \end{aligned}$$

Implementing **Vect-E** appears to require non-linear matching — there are repeated arguments in both patterns, suggesting a run-time conversion check. In fact this is not needed — the repeated arguments coincide in any *well typed* application of **Vect-E**. We do not need to recheck the duplicate  $A$  or  $k$  in the patterns  $(\varepsilon \ A)$  or  $(\!:\!:\ A \ k \ a \ v)$ . So why store them?

In this paper we show how to streamline the implementation of  $\iota$ -rules so that unnecessary testing is avoided. We introduce extensions to the TT syntax for marking parts of terms to be ignored or removed. So equipped, we consider which constructor arguments can be ignored, and then play a similar game with constructor tags. Finally, we show how to eliminate some structures entirely and make a larger example smaller — the simply typed  $\lambda$ -calculus.

### 1.3 Related Work

Correctness preserving program transformations [10, 23] provide a basis for many optimisations in simply typed functional languages. In this paper we use substitution transformations to mark unused terms for deletion in a similar manner to Berardi’s pruning of simply typed  $\lambda$ -terms [4]. Program transformation techniques have also been applied to type theory; Magaud and Bertot [17] show an approach to changing data representation by transforming the constructors and elimination rule of a family and use this technique to change from unary natural numbers to a more efficient binary representation.

The COQ program extraction tool [22, 14] attempts to remove purely logical parts of proofs in order to produce executable programs. Our approach differs in that we do not separate the predicative and impredicative type universes and attempt to remove *all* terms which are unused.

Callaghan and Luo [7] use the well-typedness of elimination rules to avoid checking of repeated arguments, a technique which we apply and extend in this paper. Xi’s DML [26] also uses dependent types for optimisation, eliminating dead code [27] and array bounds checking [28].

## 2 Implementing Reduction Rules for Datatypes

The elimination operator **D-E** is the only means **TT** provides for inspecting data in the inductive family **D**. If we optimise **D-E**’s reduction behaviour, we optimise the programs which elaborate in terms of it. Moreover, if any data in the representation of **D**’s elements is not needed by **D-E**, then it is *never* needed at run-time. Let us look more closely at how  $\iota$ -rules are implemented.

### 2.1 Pattern Syntax and its Semantics

We implement  $\iota$ -rules **D-E**  $\vec{t}_i = e_i$  by pattern matching, marking with  $[\cdot]$  those parts of patterns  $p$  which *well typed* terms are *presupposed* to match. Unmarking these parts gives back a term,  $|p|$ .

$$p ::= x \quad (\text{pattern variable}) \quad | [t] \quad (\text{presupposed term}) \\ | c \vec{p} \quad (\text{constructor pattern}) \quad | [c] \vec{p} \quad (\text{presupposed-constructor pattern})$$

For each  $\iota$ -law, as above, we write a  $\iota$ -**scheme**, **D-E**  $\vec{p}_i \mapsto e_i$  with  $|\vec{p}_i| = \vec{t}_i$  and  $e_i$  a term over  $\vec{p}_i$ ’s **pattern variables**. The  $\iota$ -schemes are then compiled into an efficient case-expression [2]. However, our pattern syntax will facilitate the discussion without delving into those details.

The partial function **MATCH** tries to compute a **matching substitution** for a pattern and term (**MATCHES** lifts **MATCH** to sequences in the obvious way):

$$\begin{aligned} \text{MATCH}(x, t) &\Longrightarrow t/x \\ \text{MATCH}(c \vec{p}, t) &\Longrightarrow \text{MATCHES}(\vec{p}, \vec{t}) \quad \text{if } \text{WHNF}(t) \Longrightarrow c' \vec{t} \quad \text{and } c = c' \\ \text{MATCH}([t'], t) &\Longrightarrow \text{ID} \\ \text{MATCH}([c] \vec{p}, t) &\Longrightarrow \text{MATCHES}(\vec{p}, \vec{t}) \quad \text{if } \text{WHNF}(t) \Longrightarrow c' \vec{t} \\ \text{MATCHES}(\cdot, \cdot) &\Longrightarrow \text{ID} \\ \text{MATCHES}(p \vec{p}, t \vec{t}) &\Longrightarrow \text{MATCH}(p, t) \circ \text{MATCHES}(\vec{p}, \vec{t}) \end{aligned}$$

The first two lines of `MATCH` test constructors and bind pattern variables as usual in implementations of pattern matching from [21] onwards. The remaining two lines, however, presuppose the successful outcome of testing. To justify these presuppositions, we shall require that each  $\iota$ -scheme is  $\Gamma$ -**respectful** of well typed instances, i.e.

if  $\Gamma \vdash \mathbf{D-E} \vec{t} : T$  and  $\text{MATCHES}(\vec{p}_i, \vec{t}) \implies \sigma$  then  $\Gamma \vdash \mathbf{D-E} \sigma|\vec{p}_i| = \mathbf{D-E} \vec{t} : T$

A set of  $\iota$ -schemes,  $\mathbf{D-E} \vec{p}_i \mapsto e_i$  is  $\Gamma$ -**well-defined** if, for any  $\Gamma \vdash \mathbf{D-E} \vec{t} : T$  of the right arity, with a constructor-headed target, we have  $\text{MATCHES}(\vec{p}_i, \vec{t}) \implies \sigma$  for exactly one  $i$ . This yields  $\iota$ -reduction  $\mathbf{D-E} \vec{t} \mapsto \sigma e_i$ . A set of  $\iota$ -schemes which is  $\Gamma$ -respectful and  $\Gamma$ -well-defined for all  $\Gamma$  is said to **implement** the corresponding  $\iota$ -rules.

## 2.2 Standard Implementation

**Theorem.** For  $\mathbf{D} : \forall \vec{i} : \vec{I}. \star$ , with typical  $\mathbf{c} : \forall \vec{a} : \vec{A}. \mathbf{D} \vec{r}_1 \rightarrow \dots \rightarrow \mathbf{D} \vec{r}_j \rightarrow \mathbf{D} \vec{s}$ , this typical  $\iota$ -scheme implements the  $\iota$ -rules (the **standard** implementation):

$$\mathbf{D-E} [\vec{s}] (\mathbf{c} \vec{a} \vec{y}) P \vec{m} \mapsto m_{\mathbf{c}} \vec{a} \vec{y} (\mathbf{D-E} \vec{r}_1 y_1 P \vec{m}) \dots (\mathbf{D-E} \vec{r}_j y_j P \vec{m})$$

**Proof.** For any  $\Gamma$ , if  $\Gamma \vdash \mathbf{D-E} \vec{s}' (\mathbf{c} \vec{a}' \vec{y}') P' \vec{m}' : T$  then

$$\text{MATCHES}([\vec{s}] (\mathbf{c} \vec{a} \vec{y}) P \vec{m}, \vec{s}' (\mathbf{c} \vec{a}' \vec{y}') P' \vec{m}') \implies \sigma$$

but matching the other  $\iota$ -schemes fails, so these schemes are  $\Gamma$ -well-defined.

Moreover,  $\sigma$  is  $\vec{a}' / \vec{a} \circ \vec{y}' / \vec{y} \circ P' / P \circ \vec{m}' / \vec{m}$ . Typechecking,  $\mathbf{c} \vec{a}' \vec{y}' : \mathbf{D}(\vec{a}' / \vec{a} \circ \vec{y}' / \vec{y}) \vec{s} = \mathbf{D} \sigma \vec{s}$ . Hence  $\sigma \vec{s} = \vec{s}'$  as  $\mathbf{D-E} \vec{s}' (\mathbf{c} \vec{a}' \vec{y}')$  is well-typed. Hence our typical scheme is  $\Gamma$ -respectful.  $\square$

The standard implementation comments out the *indices* — just as well, because there is no guarantee that they generally take the constructor form which explicit matching requires. For example, `Vect-E` has standard implementation

$$\begin{aligned} \mathbf{Vect-E} [A] [0] \quad (\varepsilon A) \quad P m_{\varepsilon} m_{\cdot} \mapsto m_{\varepsilon} \\ \mathbf{Vect-E} [A] [s k] (\:: A k a v) P m_{\varepsilon} m_{\cdot} \mapsto m_{\cdot} k a v (\mathbf{Vect-E} A k v P m_{\varepsilon} m_{\cdot}) \end{aligned}$$

## 2.3 Alternative Implementations

Where the indices of a constructor's return type do happen to resemble constructor or variable patterns, we are free to consider alternative implementations of the corresponding  $\iota$ -schemes. We may certainly comment out a pattern variable from the target if we can recover it by matching an index. For example, this is also an implementation of `Vect-E`:

$$\begin{aligned} \mathbf{Vect-E} A 0 \quad (\varepsilon [A]) \quad P m_{\varepsilon} m_{\cdot} \mapsto m_{\varepsilon} \\ \mathbf{Vect-E} A (s k) (\:: [A] [k] a v) P m_{\varepsilon} m_{\cdot} \mapsto m_{\cdot} k a v (\mathbf{Vect-E} A k v P m_{\varepsilon} m_{\cdot}) \end{aligned}$$

But we can do better than that. There is no need to check the constructor tags on *both* the length and the target — one check will do. We may take either

$$\begin{aligned}
(\dagger) \text{Vect-E } A \ [0] \quad (\varepsilon \ [A]) \quad P \ m_\varepsilon \ m_\varepsilon &\mapsto m_\varepsilon \\
\text{Vect-E } A \ ([s] \ k) \ (\llbracket A \rrbracket [k] \ a \ v) \ P \ m_\varepsilon \ m_\varepsilon &\mapsto m_\varepsilon : k \ a \ v \ (\text{Vect-E } A \ k \ v \ P \ m_\varepsilon \ m_\varepsilon)
\end{aligned}$$

or, instead, privileging index length over vector contents

$$\begin{aligned}
(\ddagger) \text{Vect-E } A \ 0 \quad ([\varepsilon] \ [A]) \quad P \ m_\varepsilon \ m_\varepsilon &\mapsto m_\varepsilon \\
\text{Vect-E } A \ (s \ k) \ (\llbracket A \rrbracket [k] \ a \ v) \ P \ m_\varepsilon \ m_\varepsilon &\mapsto m_\varepsilon : k \ a \ v \ (\text{Vect-E } A \ k \ v \ P \ m_\varepsilon \ m_\varepsilon)
\end{aligned}$$

In the sequel, we show how to choose good alternative implementations for elimination operators by systematically exploiting the presence of constructor symbols in indices. This leads naturally to space optimisations, where we do not merely comment out unnecessary data from patterns — we delete them entirely from the representation of datatypes.

## 2.4 ExTT — an execution language for TT with deleted terms

We introduce ExTT, an execution language for terms in TT. ExTT extends TT’s syntax with **deleted** terms and patterns  $\{t\}$ , and also with deleted constructor patterns  $\{c\} \vec{p}$  corresponding to untagged tuples  $\{c\} \vec{t}$ . We extend the operational semantics thus:

$$\begin{aligned}
\text{MATCH}(\{t\}, \{t'\}) &\Longrightarrow \text{ID} \\
\text{MATCH}(\{c\} \vec{p}, t) &\Longrightarrow \text{MATCHES}(\vec{p}, \vec{t}) \ \text{if} \ \text{WHNF}(t) \Longrightarrow (\{c\} \vec{t})
\end{aligned}$$

We are careful to distinguish  $(\{c\} \vec{t})$ , which is represented by the empty tuple, from  $\{c\} \vec{t}$ , which is deleted altogether. The actual evaluation of terms in ExTT can be by any standard method, such as normalisation by evaluation [1, 5], compilation to G machine code [13] or program extraction [14].

The unmarking operation  $|\cdot|$  takes both patterns and terms in ExTT back to terms in TT by stripping out both  $[\cdot]$  and  $\{\cdot\}$  marks. Terms in ExTT arise only by optimisations from well typed TT terms hence ExTT needs no typing rules provided that these optimisations are safe.

We specify an **optimisation** by giving a substitution  $\llbracket \cdot \rrbracket$  from TT identifiers to ExTT terms, ID by default, together with the optimised ExTT  $\iota$ -schemes. For  $\iota$ -rules  $\Gamma \vdash \text{D-E } \vec{t}_i = e_i$ , these have form  $\text{D-E } \vec{p}_i \mapsto d_i$ , where  $|\vec{p}_i| = \vec{t}_i$ ,  $|d_i| = e_i$  and every undeleted free variable in  $d_i$  is a pattern variable in  $\vec{p}_i$ . For all  $\Gamma$ , these schemes must be  $\Gamma$ -well-defined in the obvious way, and  $\Gamma$ -respectful in that

$$\begin{aligned}
\text{if } \Gamma \vdash \text{D-E } \vec{t} : T \text{ and } \text{MATCHES}(\vec{p}_i, \llbracket \vec{t} \rrbracket) &\Longrightarrow \sigma \\
\text{then there exists a substitution } \tau \text{ such that } \Gamma \vdash \tau |\sigma(\text{D-E } \vec{p}_i)| &= \text{D-E } \vec{t} : T
\end{aligned}$$

The rôle of  $\tau$  is to instantiate the variables free in  $e_i$ , but deleted in  $d_i$ —these are not needed when executing ExTT terms, hence they need not be matched.

In the following sections, we establish several such optimisations.

## 3 Eliding Redundant Constructor Arguments

Recall the alternative implementation of Vect-E ( $\dagger$  above) which matches  $A$  and  $k$  in the indices rather than the target. When can we do this, in general?

Whenever  $c \vec{a}, c \vec{b} : D \vec{s}$  implies  $a_i = b_i$ , we say that the  $i$ th argument of  $c$  is **forceable**. eg., the  $A$  argument to  $\varepsilon$  is forceable since if  $\varepsilon a, \varepsilon b : \text{Vect } A \ 0$  then clearly  $a = b = A$ . For  $::$ ,  $A$  and  $k$  are forceable in the same way.

Constructor arguments which have been commented out owing to their repetition in a  $\iota$ -scheme are forceable. This is to be expected; such repeated arguments arise from the patterns describing constructor indices.

Consider a typical constructor, fully applied to variables,  $c \vec{a} \vec{y} : D \vec{s}$ . If we express  $\vec{s}$  as  $|\vec{p}|$  for any patterns  $\vec{p}$ , then any  $a_i$  appearing as a pattern variable in  $\vec{p}$  is forceable, by injectivity of constructors. We call these arguments **concretely forceable** since they can be retrieved in constant time by pattern matching on the indices.

To express  $\vec{s}$  as  $|\vec{p}|$ , we write a program to extract from a term a linear pattern with its variable set:

$$\begin{aligned}
\text{PAT } (V, x) &\Longrightarrow (x \cup V, x) \text{ \underline{if} } x \notin V \\
\text{PAT } (V, c \vec{t}) &\Longrightarrow (V', \text{LAZY}(c, \vec{p})) \text{ \underline{if} } \text{PATS } (V, \vec{t}) \Longrightarrow (V', \vec{p}) \\
\text{PAT } (V, t) &\Longrightarrow (V, [t]) \\
\text{PATS } (V, \cdot) &\Longrightarrow (V, \cdot) \\
\text{PATS } (V, t \vec{t}) &\Longrightarrow (V'', p \vec{p}) \\
&\quad \text{\underline{if} } \text{PAT } (V, t) \Longrightarrow (V', p) \text{ \underline{and} } \text{PATS } (V', \vec{t}) \Longrightarrow (V'', \vec{p}) \\
\text{LAZY } (c, [\vec{p}]) &\Longrightarrow [c \vec{p}] \\
\text{LAZY } (c, \vec{p}) &\Longrightarrow [c] \vec{p} \text{ \underline{otherwise}}
\end{aligned}$$

For our typical constructor  $c$ , we can extract the patterns which **D-E** will match by  $\text{PATS } (\emptyset, \vec{s}) \Longrightarrow (V, \vec{p})$ . If an argument  $a_i \in V$  then  $a_i$  is concretely forceable. It is instantiated by matching  $\vec{p}$ , hence we may presuppose it when we match the target, yielding the same result. Hence, we may then choose the alternative implementation:

$$\text{D-E } \vec{p} (c \vec{a}^{[V]} \vec{y}) P \vec{m} \mapsto m_c \cdots \quad \text{\underline{where} } \begin{array}{l} a^{[V]} \Longrightarrow [a] \text{ \underline{if} } a \in V \\ a^{[V]} \Longrightarrow a \text{ \underline{otherwise}} \end{array}$$

**Theorem.** The following is an optimisation (**forcing**):

$$\begin{aligned}
&\text{for } c : \forall \vec{a} : \vec{A}. D \vec{r}_1 \rightarrow \dots \rightarrow D \vec{r}_j \rightarrow D \vec{s} \quad \text{\underline{where} } \text{PATS } (\emptyset, \vec{s}) \Longrightarrow (V, \vec{p}) \\
&\text{take } [[c]] \Longrightarrow \lambda \vec{a}; \vec{y}. c \vec{a}^{\{V\}} \vec{y} \\
&\text{D-E } \vec{p} (c \vec{a}^{\{V\}} \vec{y}) P \vec{m} \mapsto m_c \vec{a} \vec{y} (\text{D-E } \vec{r}_1 y_1 P \vec{m}) \dots (\text{D-E } \vec{r}_j y_j P \vec{m}) \\
&\quad \text{\underline{where} } \begin{array}{l} a^{\{V\}} \Longrightarrow \{a\} \text{ \underline{if} } a \in V \\ a^{\{V\}} \Longrightarrow a \text{ \underline{otherwise}} \end{array}
\end{aligned}$$

**Proof.** Clearly,  $|\vec{p}| = \vec{s}$  and  $|c \vec{a}^{\{V\}} \vec{y}| = c \vec{a} \vec{y}$ , so if  $\Gamma \vdash \text{D-E } \vec{s}' (c \vec{a}' \vec{y}') P' \vec{m}' : T$  then, as before,  $\vec{s}' = (\vec{a}' / \vec{a} \circ \vec{y}' / \vec{y}) \vec{s}$ . Now,

$$\begin{aligned}
\text{MATCHES } (\vec{p}, (\vec{a}' / \vec{a} \circ \vec{y}' / \vec{y}) \vec{s}) &\Longrightarrow \circ_{a_i \in V} (a'_i / a_i) \\
\text{MATCHES } (c \vec{a}^{\{V\}} \vec{y}, c \vec{a}' \vec{y}') &\Longrightarrow \circ_{a_i \notin V} (a'_i / a_i) \circ \vec{y}' / \vec{y}
\end{aligned}$$

Hence any matching substitution  $\sigma$  for the left-hand side satisfies

$$\text{ID } |\sigma(\text{D-E } \vec{p} (c \vec{a}^{\{V\}} \vec{y}) P \vec{m})| = \text{D-E } \vec{s}' (c \vec{a}' \vec{y}') P' \vec{m}'$$

So these schemes are  $\Gamma$ -respectful. They are clearly  $\Gamma$ -well-defined, as they discriminate on the target's constructor.  $\square$

For our **Vect** example, forcing is given by:

$$\begin{aligned} \llbracket \varepsilon \rrbracket &\Longrightarrow \lambda A. \varepsilon \{A\} \\ \llbracket :: \rrbracket &\Longrightarrow \lambda A; k; a; v. :: \{A\} \{k\} a v \\ \mathbf{Vect-E} A \ [0] \quad (\varepsilon \{A\}) \quad P \ m_\varepsilon \ m_{::} &\mapsto m_\varepsilon \\ \mathbf{Vect-E} A \ ([s] \ k) \ ( :: \{A\} \{k\} a v) \ P \ m_\varepsilon \ m_{::} &\mapsto m_{::} \ k \ a \ v \ (\mathbf{Vect-E} \ A \ k \ v \ P \ m_\varepsilon \ m_{::}) \end{aligned}$$

In the implementation the deleted arguments really are removed from the now fully applied constructors. This is safe because these terms are only decomposed by **Vect-E** which does not expect the deleted arguments.

## 4 Eliding Redundant Constructor Tags

Recall the second alternative implementation of **Vect-E** (§) where case selection is by analysis of the length index rather than the target itself. For which types can we do case selection on an argument other than the target?

If  $c \vec{a}, c' \vec{b} : D \vec{s}$  implies  $c = c'$ , we say that the family **D** is **detaggable**. **Vect** is detaggable because the length index determines whether the constructor is  $\varepsilon$  (if the length index is 0) or  $::$  (if the length index is  $sk$ ).

For any set of  $\iota$ -schemes, if the index patterns are already mutually exclusive, we can decide which scheme applies without checking the target's constructor tag. The following program checks if two patterns are guaranteed to match disjoint sets of terms:

$$\begin{aligned} \text{DISJOINT}(c \vec{p}, c' \vec{q}) &\Longrightarrow \text{true} \ \underline{\text{if}} \ c \neq c' \\ \text{DISJOINT}(c \vec{p}, c \vec{q}) &\Longrightarrow \exists i. \text{DISJOINT}(p_i, q_i) \\ \text{DISJOINT}([c] \vec{p}, [c] \vec{q}) &\Longrightarrow \exists i. \text{DISJOINT}(p_i, q_i) \\ \text{DISJOINT}(p \ , \ q) &\Longrightarrow \text{false} \ \underline{\text{otherwise}} \end{aligned}$$

Of course if we are to match on the indices then we must actually examine their constructors, so the previous lazy definition of **PATS** is not sufficient. We compute the patterns we need for this optimisation with **EPATS** — the same as **PATS** but with **LAZY** replaced by **EAGER**:

$$\mathbf{EAGER}(c, \vec{p}) \Longrightarrow c \vec{p}$$

Given a family **D** with constructors  $c_i : \forall \vec{x}. \vec{X}_i. D \vec{s}_i$  where  $\mathbf{EPATS}(\emptyset, \vec{s}_i) \Longrightarrow (V_i, \vec{p}_i)$ , we say **D** is **concretely detaggable** if

$$\forall i \neq j. \exists k. \text{DISJOINT}(p_{ik}, p_{jk}) \Longrightarrow \text{true}$$

**Theorem.** We may optimise (**detag**) such a concretely detaggable **D** thus:

$$\begin{aligned} \llbracket c_i \rrbracket &\Longrightarrow \lambda \vec{x}. \{c_i\} \vec{x}^{(V)} \\ \mathbf{D-E} \ \vec{p}_i \ (\{c_i\} \ \vec{x}^{(V)}) \ P \ \vec{m} &\mapsto e_i \end{aligned}$$

**Proof.** These schemes are  $\Gamma$ -respectful for all  $\Gamma$  by the same argument as for

forcing—the switch to eager patterns does not affect the set of variables matched from the indices, nor the success of matching well-typed values. Deleting the constructor in the target can only improve the possibility of a match, but the disjointness condition directly ensures that the schemes remain  $\Gamma$ -well-defined.  $\square$

For our **Vect** example, detagging is given by:

$$\begin{aligned} \llbracket \varepsilon \rrbracket &\Longrightarrow \lambda A. \{\varepsilon\} \{A\} \\ \llbracket :: \rrbracket &\Longrightarrow \lambda A; k; a; v. \{\cdot\} \{A\} \{k\} a v \\ \mathbf{Vect-E} A \ 0 \quad (\{\varepsilon\} \{A\}) \quad P \ m_\varepsilon \ m_\cdot &\mapsto m_\varepsilon \\ \mathbf{Vect-E} A \ (s \ k) \ (\{\cdot\} \{A\} \{k\} a \ v) \ P \ m_\varepsilon \ m_\cdot &\mapsto m_\cdot \ k \ a \ v \ (\mathbf{Vect-E} \ A \ k \ v \ P \ m_\varepsilon \ m_\cdot) \end{aligned}$$

We achieve this space optimisation at the cost of using eager rather than lazy patterns. The number of constructor tests required increases by a constant (possibly zero!) factor and indices may sometimes be computed where they would previously be ignored. Clearly a real implementation would minimise the number of eager patterns required to make the distinction. An analysis of this space/time trade-off is beyond the scope of this paper, but for **Vect** it seems likely to be worthwhile since we have swapped one constructor test for another.

## 5 Run-Time Optimisation

In our **Vect-E** example, we have already deleted both  $\varepsilon$  and its argument. We might be tempted to go a step further, and comment out that entire target.

$$\mathbf{Vect-E} \ A \ 0 \ [\{\varepsilon\} \{A\}] \ P \ m_\varepsilon \ m_\cdot \mapsto m_\varepsilon$$

However, this  $\iota$ -scheme is not respectful and breaks subject reduction thus:

$$\begin{aligned} \dots ; x : \mathbf{Vect} \ A \ 0 \vdash \mathbf{Vect-E} \ A \ 0 \ x \ P \ m_\varepsilon \ m_\cdot &: P \ 0 \ x \\ &\mapsto m_\varepsilon : P \ 0 \ \varepsilon \end{aligned}$$

The pattern  $(\{\varepsilon\} \{A\})$  may not test tags or extract arguments, but it still only matches targets whose weak head-normal forms are constructor applications. The optimisations we have seen thus far are safe to use in any context, and we need to reduce under binders when performing the equality checks which ensure that **EPIGRAM** programs elaborate to well typed terms.

However, at run-time, we can employ a much more restricted notion of computation, reducing only in the *empty* context,  $\mathcal{E}$ . In this scenario, we can exploit the **adequacy** property of **TT** — if  $\mathcal{E} \vdash t : D \vec{s}$  then  $\text{WHNF}(t)$  is  $c \vec{t}$  for some  $\vec{t}$  — to gain further optimisations, not available in a general context.

In effect, we may employ weaker criteria for alternative implementations of elimination operators in run-time execution. We say that a **run-time optimisation** is given by a substitution and  $\iota$ -schemes in **ExTT** as before, except that these schemes need only be  $\mathcal{E}$ -respectful and  $\mathcal{E}$ -well-defined.

The adequacy property tells us that the target will always match a constructor pattern at run-time, hence we may safely presuppose a pattern from which

no information is gained, as suggested above. Moreover, by applying this observation inductively, we can sometimes extract another, more drastic optimisation from the guarantee of adequacy at run-time.

## 6 Collapsing Content-Free Families at Run-Time

Consider the less than or equal relation, declared and elaborated as follows:

$$\begin{array}{l} \text{data } \frac{x, y : \mathbb{N}}{x \leq y : \star} \quad \text{where } \frac{}{\text{leO} : 0 \leq y} \quad \frac{p : x \leq y}{\text{leS } p : sx \leq sy} \\ \leq : \mathbb{N} \rightarrow \mathbb{N} \rightarrow \star \\ \text{leO} : \forall y : \mathbb{N}. \leq 0 y \\ \text{leS} : \forall x, y : \mathbb{N}. \leq x y \rightarrow \leq (sx) (sy) \end{array}$$

The  $\leq$  family describes a property of its indices and stores no other data. It is not surprising therefore to find that much of its content can be deleted. Forcing and detagging yield:

$$\begin{array}{l} \llbracket \text{leO} \rrbracket \implies \lambda y. (\{\text{leO}\} \{y\}) \\ \llbracket \text{leS} \rrbracket \implies \lambda x; y; p. (\{\text{leS}\} \{x\} \{y\} p) \\ \leq\text{-E } 0 \quad y \quad (\{\text{leO}\} \{y\}) \quad P \quad m_{\text{leO}} \quad m_{\text{leS}} \mapsto m_{\text{leO}} y \\ \leq\text{-E } (sx) (sy) (\{\text{leS}\} \{x\} \{y\} p) \quad P \quad m_{\text{leO}} \quad m_{\text{leS}} \\ \mapsto m_{\text{leS}} x y p (\leq\text{-E } x y p \quad P \quad m_{\text{leO}} \quad m_{\text{leS}}) \end{array}$$

Now we are left with only one undeleted argument, the recursive  $p$  in  $\text{leS}$ . This argument serves two purposes — firstly it is the target of the recursive call and secondly it is passed to the method  $m_{\text{leS}}$ . We might think that  $p$  can also be elided — ultimately it can only be examined by  $\leq\text{-E}$  which, by induction, can be shown never to examine it. In a partial evaluation setting, however, where we may reduce under binders, we must at least check that the target is canonical for reduction to be possible. If not, we run the risk of reducing a proof of something which cannot be constructed, such as  $5 \leq 4$ !

At run-time, on the other hand, we never need to check that  $p$  is canonical because the adequacy property tells us that it must be. Hence, at run-time, we no longer need to store the recursive argument — the entire family collapses:

$$\begin{array}{l} \llbracket \text{leO} \rrbracket \implies \lambda y. (\{\text{leO}\} y) \\ \llbracket \text{leS} \rrbracket \implies \lambda x; y; p. (\{\text{leS}\} x y p) \\ \llbracket \leq\text{-E} \rrbracket \implies \lambda x; y; p; P; m_{\text{leO}}; m_{\text{leS}}. \leq\text{-E } x y \{p\} P m_{\text{leO}} m_{\text{leS}} \end{array}$$

For which families can we do this run-time optimisation? If  $x, y : D \vec{s}$  implies  $x = y$  we say that the family  $D$  is **collapsible**.  $\leq$  is collapsible because any value in  $x \leq y$  is determined entirely by the indices  $x$  and  $y$ .

$$\begin{array}{l} \leq\text{-E } 0 \quad y \quad \{\text{leO}\} y \quad P \quad m_{\text{leO}} \quad m_{\text{leS}} \mapsto m_{\text{leO}} y \\ \leq\text{-E } (sx) (sy) \{\text{leS}\} x y p \quad P \quad m_{\text{leO}} \quad m_{\text{leS}} \\ \mapsto m_{\text{leS}} x y (\{p\}) (\leq\text{-E } x y \{p\} P m_{\text{leO}} m_{\text{leS}}) \end{array}$$

We say a family is **concretely collapsible** if it is detaggable and for each constructor  $c : \forall \vec{a} : \vec{A}. D \vec{r}_1 \rightarrow \dots \rightarrow D \vec{r}_j \rightarrow D \vec{s}, \text{EPATS}(\emptyset, \vec{s}) \Longrightarrow (\vec{a}, \vec{p})$ . That is, the constructor tag and all the non-recursive arguments are *cheaply* recoverable from the indices.

**Theorem.** We may optimise a concretely collapsible family *at run-time*:

$$\begin{aligned} \text{D-E } \vec{p} \{c \vec{a} \vec{y}\} P \vec{m} & \\ \mapsto m_c \vec{a} (\{y_1\}) \dots (\{y_n\}) (D\text{-E } \vec{r}_1 \{y_1\} P \vec{m}) \dots (D\text{-E } \vec{r}_n \{y_n\} P \vec{m}) & \\ \llbracket c \rrbracket \Longrightarrow \lambda \vec{a}; \vec{y}. (\{c \vec{a} \vec{y}\}) & \\ \llbracket D\text{-E} \rrbracket \Longrightarrow \lambda \vec{i}; x; p; \vec{m}. D\text{-E } \vec{i} \{x\} P \vec{m} & \end{aligned}$$

**Proof.** These schemes are  $\mathcal{E}$ -well-defined by the same argument as for detagging. They are  $\mathcal{E}$ -respectful because the only possible left-hand sides have the form  $\mathcal{E} \vdash D\text{-E } \vec{s}' (c \vec{a}' \vec{y}') P' \vec{m}'$ , hence, by disjointness, the only possible match, even with the target deleted, is with the scheme for  $c$ , with matching substitution  $\sigma = \vec{a}' / \vec{a} \circ P' / P \circ \vec{m}' / \vec{m}$ , binding all the undeleted free variables on the right-hand side because  $\text{EPATS}(\emptyset, \vec{s}) \Longrightarrow (\vec{a}, \vec{p})$ . Taking  $\tau = \vec{y}' / \vec{y}$ , we see that

$$\mathcal{E} \vdash \tau | \sigma (D\text{-E } \vec{p} \{c \vec{a} \vec{y}\} P \vec{m}) | = D\text{-E } \vec{s}' (c \vec{a}' \vec{y}') P' \vec{m}'$$

hence these schemes are  $\mathcal{E}$ -respectful.  $\square$

## 7 Examples

### 7.1 The Finite Sets

The finite sets, indexed over a natural number  $n$ , are a family of types with  $n$  elements. Effectively, they are a representation of bounded numbers and are declared as follows:

$$\text{data } \frac{n : \mathbb{N}}{\text{Fin } n : \star} \quad \text{where } \frac{}{\text{f0} : \text{Fin } sn} \quad \frac{i : \text{Fin } n}{\text{fs } i : \text{Fin } sn}$$

The forcing optimisation elides the indices from the elaborated constructors:

$$\begin{aligned} \llbracket \text{f0} \rrbracket &\Longrightarrow \lambda n. \text{f0 } \{n\} \\ \llbracket \text{fs} \rrbracket &\Longrightarrow \lambda n; i. \text{fs } \{n\} i \end{aligned}$$

After stripping the forceable arguments, the shape of the resulting type matches that of  $\mathbb{N}$  — that is, the base constructor takes no arguments and the step constructor takes a single recursive argument. In principle, any optimisations which apply to  $\mathbb{N}$  such as Magaud and Bertot's binary representation [17] should also apply to  $\text{Fin}$ . We hope to recover Xi's efficient treatment of bounded numbers in this way [28] and perhaps extend it to other forms of validation.

## 7.2 Comparison of Natural Numbers

The Compare family from [20] represents the result of comparing two numbers, storing which is the greater and by how much:

$$\begin{array}{c} \text{data} \quad \frac{m, n : \mathbb{N}}{\text{Compare } m \ n : \star} \quad \text{where} \quad \frac{}{\text{lt } y : \text{Compare } x \ (x + (s \ y))} \\ \frac{}{\text{eq} : \text{Compare } x \ x} \\ \frac{}{\text{gt } x : \text{Compare } (y + (s \ x)) \ y} \end{array}$$

Compare is an example of a family which is collapsible, but not concretely collapsible. Clearly there is only one possible element of `Compare m n` for each  $m$  and  $n$ , and given this element we can extract their difference in constant time. If we were to collapse Compare we would replace this simple inspection by the recomputation of the difference each time the same value was used. We restrict concretely collapsible families to those where the recomputation of values is cheap. Nonetheless, by forcing, Compare need only store which index is larger and by how much:

$$\begin{array}{l} \llbracket \text{lt} \rrbracket \implies \lambda x; y. \text{lt } \{x\} \ y \\ \llbracket \text{eq} \rrbracket \implies \lambda x. \text{eq } \{x\} \\ \llbracket \text{gt} \rrbracket \implies \lambda x; y. \text{gt } x \ \{y\} \end{array}$$

## 7.3 Accessibility Predicates

In [6], Bove and Capretta use special-purpose accessibility predicates to prove termination of general recursive functions. For example, **quicksort** terminates on the `nil`, and it terminates on `cons x xs` if it terminates on **filter** ( $< x$ ) `xs` and **filter** ( $\geq x$ ) `xs`. This is expressed by the `qsAcc` predicate below:

$$\begin{array}{c} \text{data} \quad \frac{l : \text{List } \mathbb{N}}{\text{qsAcc } l : \star} \\ \text{where} \quad \frac{}{\text{qsNil} : \text{qsAcc } \text{nil}} \\ \frac{\text{qsl} : \text{qsAcc } (\text{filter } (< \ x) \ xs) \quad \text{qsr} : \text{qsAcc } (\text{filter } (\geq \ x) \ xs)}{\text{qsCons } \text{qsl } \text{qsr} : \text{qsAcc } (\text{cons } x \ xs)} \end{array}$$

**quicksort** itself is defined by induction over `qsAcc`, so a naïve implementation would need to store the proofs. However, `qsAcc` is concretely collapsible:

$$\begin{array}{l} \llbracket \text{qsNil} \rrbracket \implies \{\text{qsNil}\} \\ \llbracket \text{qsCons} \rrbracket \implies \lambda x; xs; qsl; qsr. \{\text{qsCons } x \ xs \ qsl \ qsr\} \end{array}$$

Collapsing replaces computation over `qsAcc` by computation over its indices, restoring the intended operational semantics of the original program! These accessibility predicates are concretely collapsible because their indices are constructed from the constructor patterns of programs.

## 7.4 The Simply Typed $\lambda$ -calculus

We define the simply typed  $\lambda$ -calculus in a similar fashion to [20], making extensive use of inductive families to specify invariants on the data structures. We begin with  $\text{STy}$ , representing simple monomorphic types:

$$\text{data } \overline{\text{STy}} : \star \quad \text{where } \overline{\iota} : \text{STy} \quad \frac{s, t : \text{STy}}{s \Rightarrow t : \text{STy}}$$

We represent contexts by  $\text{Vect}$ s of types,  $\text{Ctx} = \text{Vect STy}$ . The explicit length allows a safe de Bruijn representation of variables, via the  $\text{Fin}$  family, hence our untyped terms,  $\text{Expr}$ , are at least well scoped—the length is forceable for each constructor:

$$\text{data } \overline{\text{Expr } n} : \star$$

$$\text{where } \frac{i : \text{Fin } n}{\text{eVar } i : \text{Expr } n} \quad \frac{S : \text{STy} \quad t : \text{Expr } sn}{\text{eLam } S t : \text{Expr } n} \quad \frac{f, s : \text{Expr } n}{\text{eApp } f s : \text{Expr } n}$$

The  $\text{Var}$  relation gives types to variables.  $\text{Var } G i T$  states that the  $i$ th member of the context  $G$  has type  $T$ . Clearly  $\text{Var}$  is concretely collapsible.

$$\text{data } \frac{G : \text{Ctx } n \quad i : \text{Fin } n \quad T : \text{STy}}{\text{Var } G i T : \star}$$

$$\text{where } \frac{}{\text{stop} : \text{Var } (S::G) \text{f0 } S} \quad \frac{v : \text{Var } G i T}{\text{pop } v : \text{Var } (S::G) (\text{fs } i) T}$$

Finally, we have the well typed terms, indexed over contexts, the original raw terms and types. This gives us a particularly safe representation — no typechecker can return the wrong well typed term. This indexing also enables us to synchronise terms safely with value environments during evaluation in the style of Augustsson and Carlsson [3].

$$\text{data } \frac{G : \text{Ctx } n \quad e : \text{Expr } n \quad T : \text{STy}}{\text{Term } G e T : \star}$$

$$\text{where } \frac{v : \text{Var } G i T}{\text{var } v : \text{Term } G (\text{eVar } i) T} \quad \frac{b : \text{Term } (S::G) e T}{\text{lam } b : \text{Term } G (\text{eLam } S e) (S \Rightarrow T)}$$

$$\frac{f : \text{Term } G fe (S \Rightarrow T) \quad a : \text{Term } G ae S}{\text{app } f a : \text{Term } G (\text{eApp } fe ae) T}$$

$\text{Term}$  seems to involve a horrifying amount of duplication. Fortunately, many of the arguments are forceable and thanks to the indexing over raw terms,  $\text{Term}$  is detaggable. After optimisation, this is all that remains:

$$\begin{aligned} \llbracket \text{var} \rrbracket &\Longrightarrow \lambda n; G; i; T; v. \{\text{var}\} \{n\} \{G\} \{i\} \{T\} \{v\} \\ \llbracket \text{lam} \rrbracket &\Longrightarrow \lambda n; G; S; e; T; b. \{\text{lam}\} \{n\} \{G\} \{S\} \{e\} \{T\} b \\ \llbracket \text{app} \rrbracket &\Longrightarrow \lambda n; G; fe; S; T; f; ae; a. \{\text{app}\} \{n\} \{G\} \{fe\} S \{T\} f \{ae\} a \end{aligned}$$

The only non-recursive arguments which survive are the domain types of applications. Typechecking thus consists of ensuring that these can be determined.

## 8 Conclusions and Further Work

The ideas presented here have been tested in a prototype implementation. Execution in this system is by extraction to a Haskell coding of `ExTT` values without the deleted subterms. We have used the GHC profiling tools [25] to assess the space usage of programs.

Our experiments show a significant reduction in space requirements over a naïve implementation particularly where there is extensive indexing. For vector operations, a 10-20% saving in memory usage is typical (depending on the length of the vector), but for the typechecker, a saving of over 80% has been observed.

Although remarkably straightforward, these optimisations only present themselves when one takes dependently typed programming seriously. The forcing optimisation largely overcomes the space penalty of adopting dependent types, but detagging derives new dynamic benefit from previously unavailable static information. Collapsing, too, has significant consequences, deleting accessibility arguments and all the equational reasoning from run-time code, not because we deem them to be proof-irrelevant, but because they actually are.

We suspect that these optimisations are the first of many. For example, as we erase forceable indices, it is worth identifying operations which affect nothing else, such as `weaken` :  $\text{Fin } n \rightarrow \text{Fin } (sn)$ , which embeds a value in a higher indexed set — this is effectively the identity function. This optimisation applies wherever functions exist only to manage invariants.

We might also consider the low level implementation of high level types, such as the natural numbers. By replacing `N-E` with an appropriate elimination rule for unbounded binary numbers [17] we can achieve a significant speed-up. Any other data structure with the same shape *after optimisation*, eg. `Fin`, can be treated similarly. In a practical implementation, such optimisations are essential for comparable performance to its conventional counterparts.

Optimisation of a new language with a new type system naturally presents new problems and new opportunities. While we can never hope to produce a completely optimal program in all cases, this research leads us to believe that the presence of much more static information can only give us greater scope for optimisation in both time and space.

## References

1. Thorsten Altenkirch, Martin Hofmann, and Thomas Streicher. Categorical reconstruction of a reduction free normalization proof, 1996.
2. Lennart Augustsson. Compiling pattern matching. In Jean-Pierre Jouannaud, editor, *Functional Programming Languages and Computer Architecture*, pages 368–381. Springer-Verlag, September 1985.
3. Lennart Augustsson and Magnus Carlsson. An exercise in dependent types: A well-typed interpreter. <http://www.cs.chalmers.se/~augustss/cayenne/>, 1999.
4. Stefano Berardi. Pruning simply typed lambda terms. *Journal of Logic and Computation*, 6(5):663–681, 1996.

5. U. Berger and H. Schwichtenberg. An inverse of the evaluation functional for typed  $\lambda$ -calculus. In R. Vemuri, editor, *Proceedings of the Sixth Annual IEEE Symposium on Logic in Computer Science*, pages 203–211. IEEE Computer Society Press, 1991.
6. Ana Bove and Venanzio Capretta. Modelling general recursion in type theory, September 2002.
7. Paul Callaghan and Zhaohui Luo. Implementation techniques for inductive types in plastic. In Bengt Nordström Thierry Coquand, Peter Dybjer and Jan Smith, editors, *Types for Proofs and Programs*, volume 1956 of *LNCS*, pages 94–113. Springer-Verlag, 1999.
8. Paul Callaghan, Zhaohui Luo, James McKinna, and Robert Pollack, editors. *Types for Proofs and Programs*, volume 2277 of *LNCS*. Springer-Verlag, 2000.
9. Coq Development Team. The Coq proof assistant — reference manual, 2001.
10. André Luís de Medeiros Santos. *Compilation By Transformation In Non-Strict Functional Languages*. PhD thesis, University of Glasgow, 1995.
11. Peter Dybjer. Inductive families. *Formal Aspects Of Computing*, 6:440–465, 1994.
12. Robert Harper and Randy Pollack. Type checking with universes. *Theoretical Computer Science*, 89(1):107–136, 1991.
13. Thomas Johnsson. Efficient compilation of lazy evaluation, 1984.
14. Pierre Letouzey. A new extraction for Coq. In Herman Geuvers and Freek Wiedijk, editors, *Types for proofs and programs*, volume 2646 of *LNCS*. Springer-Verlag, 2002.
15. Zhaohui Luo. *Computation and Reasoning – A Type Theory for Computer Science*. International Series of Monographs on Computer Science. OUP, 1994.
16. Zhaohui Luo and Robert Pollack. LEGO proof development system: User's manual. Technical report, LFCS, University of Edinburgh, 1992.
17. Nicolas Magaud and Yves Bertot. Changing data structures in type theory: A study of natural numbers. In Callaghan et al. [8], pages 181–196.
18. Lena Magnusson. *The implementation of ALF – A Proof Editor based on Martin-Löf's Monomorphic Type Theory with Explicit Substitutions*. PhD thesis, Chalmers University of Technology, Göteborg, 1994.
19. Conor McBride. *Dependently Typed Functional Programs and their proofs*. PhD thesis, University of Edinburgh, May 2000.
20. Conor McBride and James McKinna. The view from the left. *Journal of Functional Programming*, 14(1), 2004.
21. Fred McBride. *Computer Aided Manipulation of Symbols*. PhD thesis, Queen's University of Belfast, 1970.
22. Christine Paulin-Mohring. *Extraction de programmes dans le Calcul des Constructions*. PhD thesis, Paris 7, 1989.
23. Simon L Peyton Jones and André L M Santos. A transformation-based optimiser for Haskell. *Science of Computer Programming*, 32:3–47, 1998.
24. Randy Pollack. Implicit syntax. Technical report, LFCS, University of Edinburgh, 1992.
25. Patrick M. Sansom and Simon L. Peyton Jones. Time and space profiling for non-strict, higher order functional languages, 1995.
26. Hongwei Xi. *Dependent Types in Practical Programming*. PhD thesis, Department of Mathematical Sciences, Carnegie Mellon University, December 1998.
27. Hongwei Xi. Dead code elimination through dependent types, 1999.
28. Hongwei Xi and Frank Pfenning. Eliminating array bound checking through dependent types, 1998.