Custom Representations of Inductive Families

(Extended Abstract, Student Research Paper)

Constantine Theocharis[0009−0001−0198−2750] and
Edwin Brady[0000−0002−9734−367X]

University of St Andrews, UK
{kt81,ecb10}@st-andrews.ac.uk

Abstract. Inductive families provide a convenient way of programming with dependent types. Yet, when it comes to compilation, their default linked-tree runtime representations, as well as the need to convert between different indexed views of the same data when programming with dependent types, can lead to unsatisfactory runtime performance. In this paper, we aim to introduce a language with dependent types, and inductive families with custom representations. Representations are a version of Wadler's views [13], refined to inductive families like in Epigram [11]. However, representations come with compilation guarantees: a represented inductive family will not leave any runtime traces behind, without having to rely on automated optimisations such as deforestation [14]. This way, we can build a library of convenient inductive families based on a minimal set of primitives, whose re-indexing and conversion functions are erased at compile-time. In addition, we show how we can express inductive data optimisation techniques, such as representing Nat-like types as GMP-style big integers, without special casing in the compiler. With dependent types, reasoning about data representations is also possible; for example, we get computationally irrelevant isomorphisms between the original and represented data.

Keywords: Dependent types · Memory representation · Inductive families

1 Introduction

Inductive families are a generalisation of inductive data types found in some programming languages with dependent types. Every inductive definition is equipped with an eliminator that captures the notion of mathematical induction over the data, and in particular, enables structural recursion over the data. This is a powerful tool for programming as well as theorem proving. However, this abstraction has a cost when it comes to compilation: the runtime representation of inductive types is a linked tree structure. This representation is not always the most efficient for all operations, and often forces users to rely on more efficient machine primitives to achieve desirable performance, at the cost of structural recursion and dependent pattern matching. This is the first problem we aim to address in this paper.

Despite advances in the erasure of irrelevant indices in inductive families [4] and the use of theories with irrelevant fragments [2,12], there is still a need to convert between different indexed views of the same data. For example, the function to convert from List T to Vec T n by forgetting the length index n is *not* erased by any current language with dependent types, unless vectors are defined as a refinement of lists with an erased length field (which hinders dependent pattern matching due to the presence of non-structural witnesses), or a Church encoding is used in a Curry-style context [8] (which restricts the flexibility of data representation). This is the second problem we aim to address in this paper.

Wadler's views [13] provide a way to abstract over inductive interfaces, so that different views of the same data can be defined and converted between seamlessly. In the context of inductive families, views have been used in Epigram [11] that utilise the index refinement machinery of dependent pattern matching to avoid certain proof obligations with eliminator-like constructs. While views exhibit a nice way to transport across a bijection between the original data and the viewed data, they do not utilise this bijection to erase the view from the program. Despite deforestation handling this erasure to some extent, it is not guaranteed to erase all traces of the view from the program, and the optimisation might be difficult to predict.

In this paper, we propose an extension λ_{REP} to a core language with dependent types and inductive families λ_{IND}, which allows programmers to define custom, correct-by-construction data representations. This is done through user-defined translations of the constructors and eliminators of an inductive type to a concrete implementation, which form a bijective view of the original data called a 'representation'. Representations are defined internally to the language, and require coherence properties that ensure a representation is faithful to its the original inductive family. In the final version of the paper, we plan to contribute the following:

- A dependent type system with inductive families λ_{IND}, and its extension by representations λ_{REP}.
- A formulation of common optimisations such as the 'Nat-hack', and similarly for other inductive types, as representations.
- A demonstration of zero-cost data reuse when reindexing by using representations.
- A translation from λ_{REP} to λ_{IND} that erases all inductive types with representations from the program.
- An implementation of this system and accompanying examples in Superfluid, a programming language with inductive types and dependent pattern matching.

2 A tour of data representations

A common optimisation done by programming languages with dependent types such as Idris 2 and Lean is to represent natural numbers as GMP-style [1] big

integers. The definition of natural numbers looks like

$$\textbf{data}\ \mathsf{Nat}\ \left\{\begin{array}{c}\mathsf{0} : \mathsf{Nat} \\ \mathsf{1+} : \mathsf{Nat} \to \mathsf{Nat}\end{array}\right\} \tag{1}$$

and generates a Peano-style induction principle $\mathsf{elim}_{\mathsf{Nat}}$ of type[1]

$$(P : \mathsf{Nat} \to \mathcal{U}) \to P\ \mathsf{0} \to ((n : \mathsf{Nat}) \to \overline{P\ n} \to P\ (\mathsf{1+}\ n)) \to (s : \mathsf{Nat}) \to P\ s\,.$$

Without further intervention, the Nat type is represented in unary form, where each digit becomes an empty heap cell at runtime. This is inefficient for a lot of the basic operations on natural numbers, especially since computers are particularly well-equipped to deal with numbers natively, so many real-world implementations will treat Nat specially, swapping the default inductive type representation with one based on GMP integers. This is done by performing the replacements

$$|\mathsf{0}| = \mathsf{0} \tag{2}$$
$$|\mathsf{1+}| = \mathsf{1}\ + \tag{3}$$
$$|\mathsf{elim}_{\mathsf{Nat}}\ P\ m_{\mathsf{0}}\ m_{\mathsf{1+}}\ s| = \mathsf{ubig\text{-}elim}\ |s|\ |m_{\mathsf{0}}|\ |m_{\mathsf{1+}}| \tag{4}$$

where $|\cdot|$ denotes a source translation into a compilation target language with primitives $\mathsf{ubig}\text{-}*$.[2]

In addition to the constructors and eliminators, the compiler might also define translations for commonly used definitions which have a more efficient counterpart in the target, such as recursively-defined addition, multiplication, etc. The recursively-defined functions are well-suited to proofs and reasoning, while the GMP primitives are more efficient for computation.

The issue with this approach is that it only works for the data types which the compiler can recognise as special. Particularly in the presence of dependent types, other data types might end up being equivalent to Nat or another 'nicely-representable' type, but in a non-trivial way that the compiler cannot recognise. Hence, one of our goals is to extend this optimisation to work for any data type. To achieve this this, our framework requires that representations are fully typed in a way that ensures the behaviour of the representation of a data type matches the behaviour of the data type itself.

[1] Recursive parameters like $\overline{P\ n}$ are lazy, which makes the eliminator more efficient when they are not used.

[2] Idris 2 will in fact look for any 'Nat-like' types and apply this optimisation. A Nat-like type is any type with two constructors, one with arity zero and the other with arity one. A similar optimisation is also done with list-like and boolean-like types because they have a canonical representation in the target runtime, Chez Scheme.

2.1 The well-typed Nat-hack

A representation definition looks like

$$\textbf{repr}\ \mathsf{Nat}\ \textbf{as}\ \mathsf{UBig}\ \left\{\begin{array}{rl} \mathsf{0} & \textbf{as}\ \mathsf{0} \\ \mathsf{1+}\ n & \textbf{as}\ \mathsf{1}+n \\ \mathsf{elim_{Nat}} & \textbf{as}\ \mathsf{ubig\text{-}elim} \\ & \textbf{by}\ \mathsf{ubig\text{-}elim\text{-}zero\text{-}id}, \\ & \quad\ \mathsf{ubig\text{-}elim\text{-}add\text{-}one\text{-}id} \end{array}\right\}$$

Nat is represented as the type UBig of GMP-style unlimited-size unsigned integers, with translations for the constructors 0 and 1+, and the eliminator $\mathsf{elim_{Nat}}$. Additionally, the eliminator satisfies the expected computation rules of the Nat eliminator, which are postulated as propositional equalities. This representation is valid in a signature containing the primitives

$$\begin{gathered} \mathsf{0}, \mathsf{1} : \mathsf{UBig} \qquad +, \times : \mathsf{UBig} \to \mathsf{UBig} \to \mathsf{UBig} \\ \mathsf{ubig\text{-}elim} : (P : \mathsf{UBig} \to \mathcal{U}) \to P\ \mathsf{0} \to ((n : \mathsf{UBig}) \to \overline{P\ n} \to P\ (\mathsf{1}+n)) \\ \to (s : \mathsf{UBig}) \to P\ s \end{gathered}$$

and propositional equalities

$$\mathsf{ubig\text{-}elim\text{-}zero\text{-}id} :_{\forall Pbr}\ \mathsf{ubig\text{-}elim}\ P\ b\ r\ \mathsf{0} = b$$
$$\mathsf{ubig\text{-}elim\text{-}add\text{-}one\text{-}id} :_{\forall Pbrn}\ \mathsf{ubig\text{-}elim}\ P\ b\ r\ (\mathsf{1}+n) = r\ n\ (\lambda\_.\ \mathsf{ubig\text{-}elim}\ P\ b\ r\ n)\,.$$

Representations can also be defined for functions on Nat, such as addition, multiplication, and other numeric operations, in terms of UBig primitives.

$$\textbf{repr}\ \mathsf{add}\ \textbf{as}\ +\ \textbf{by}\ +\text{-}\mathsf{id} \qquad \textbf{repr}\ \mathsf{mul}\ \textbf{as}\ \times\ \textbf{by}\ \times\text{-}\mathsf{id}$$

These will be replaced during a translation process back to λ_{IND}, like rewriting rules [6], given that we have the appropriate lemmas to justify them in the signature.

This will effectively erase the Nat type from the compiled program, replacing all occurrences with the UBig type and its primitives. In a way, the hard work is done by the postulates above; we expect that the underlying implementation of UBig indeed satisfies them, which is a separate concern from the correctness of the representation itself. However, postulates are only needed when the representation target is a primitive; the next examples use defined types as targets, so that the coherence of the target eliminator follows from the coherence of other eliminators used in its implementation.

2.2 Vectors are just certain lists

In addition to representing inductive types as primitives, we can use representations to share the underlying data when converting between indexed views of

the same data. For example, we can define a representation of Vec in terms of List, so that the conversion from one to the other is 'compiled away'. We can do this by representing the indexed type as a refinement of the unindexed type by an appropriate relation. For the case of Vec, we know intuitively that

$$\mathsf{Vec}\ T\ n \simeq \{l : \mathsf{List}\ T \mid \mathsf{length}\ l = n\} := \mathsf{List'}\ T\ n$$

so we can start by choosing $\mathsf{List'}\ T\ n$ as the representation of $\mathsf{Vec}\ T\ n$.[3] We are then tasked with providing terms that correspond to the constructors of Vec but for List'. These can be defined as

$$\begin{aligned} &\mathsf{nil} : \mathsf{List'}\ T\ 0 && \mathsf{cons} : T \to \mathsf{List'}\ T\ n \to \mathsf{List'}\ T\ (1{+}\ n) \\ &\mathsf{nil} = (\mathsf{nil}, \mathsf{refl}) && \mathsf{cons}\ x\ (xs, p) = (\mathsf{cons}\ x\ xs, \mathsf{cong}\ (1{+})\ p) \end{aligned}$$

Next we need to define the eliminator for List', which should have the form

$$\begin{aligned} &\mathsf{elim\text{-}List'} : (E : (n : \mathsf{Nat}) \to \mathsf{List'}\ T\ n \to \mathsf{Type}) \\ &\quad \to E\ 0\ \mathsf{nil} \\ &\quad \to ((x : T) \to (n : \mathsf{Nat}) \to (xs : \mathsf{List'}\ T\ n) \to \overline{E\ n\ xs} \to E\ (1{+}\ n)\ (\mathsf{cons}\ x\ xs)) \\ &\quad \to (n : \mathsf{Nat}) \to (v : \mathsf{List'}\ T\ n) \to E\ n\ v \end{aligned}$$

Dependent pattern matching does a lot of the heavy lifting by refining the length index and equality proof by matching on the underlying list. However we still need to substitute the lemma $\mathsf{cong}\ (1{+})\ (1{+}\text{-}\mathsf{inj}\ p) = p$ in the recursive case.

$$\begin{aligned} &\mathsf{elim\text{-}List'}\ P\ b\ r\ 0\ (\mathsf{nil}, \mathsf{refl}) = b \\ &\mathsf{elim\text{-}List'}\ P\ b\ r\ (1{+}\ m)\ (\mathsf{cons}\ x\ xs, e) = \mathsf{subst}\ (\lambda p.\ P\ (1{+}\ m)\ (\mathsf{cons}\ x\ xs, p)) \\ &\qquad\qquad (1{+}\text{-}\mathsf{cong\text{-}id}\ e)\ (r\ x\ (xs, 1{+}\text{-}\mathsf{inj}\ e) \\ &\qquad\qquad (\lambda\_.\ \mathsf{elim\text{-}List'}\ P\ b\ r\ m\ (xs, 1{+}\text{-}\mathsf{inj}\ e))) \end{aligned}$$

Finally, we need to prove that the eliminator satisfies the expected computation rules propositionally. These are

$$\mathsf{elim\text{-}List'\text{-}nil\text{-}id} : \mathsf{elim\text{-}List'}\ P\ b\ r\ 0\ (\mathsf{nil}, \mathsf{refl}) = b$$

$$\begin{aligned} \mathsf{elim\text{-}List'\text{-}cons\text{-}id} : \mathsf{elim\text{-}List'}\ &P\ b\ r\ (1{+}\ m)\ (\mathsf{cons}\ x\ xs, \mathsf{cong}\ (1{+})\ p) \\ &= r\ x\ (xs, p)\ (\lambda\_.\ \mathsf{elim\text{-}List'}\ P\ b\ r\ m\ (xs, p)) \end{aligned}$$

which we leave as an exercise, though they are evident from the definition of elim-List'. This completes the definition of the representation of Vec as List', which would be written as

$$\textbf{repr}\ \mathsf{Vec}\ T\ n\ \textbf{as}\ \mathsf{List'}\ T\ n \left\{ \begin{array}{r} \mathsf{nil}\ \textbf{as}\ \mathsf{nil} \\ \mathsf{cons}\ \textbf{as}\ \mathsf{cons} \\ \mathsf{elim}_{\mathsf{Vec}}\ \textbf{as}\ \mathsf{elim\text{-}List'} \\ \textbf{by}\ \mathsf{elim\text{-}List'\text{-}nil\text{-}id}, \\ \mathsf{elim\text{-}List'\text{-}cons\text{-}id} \end{array} \right\}$$

[3] We will take the subset $\{x : A \mid P\ x\}$ to mean a Σ-type $(x : A) \times P\ x$ where the right component is irrelevant and erased at runtime.

Now the hard work is done; Every time we are working with a $v : \mathsf{Vec}\ T\ n$, its form will be (l, p) at runtime, where l is the underlying list and p is the proof that the length of l is n. Under the assumption that the Σ-type's right component is irrelevant and erased at runtime, every vector is simply a list at runtime, where the length proof has been erased. In the full paper we will show how this erasure is achieved in practice in Superfluid using Quantitative Type Theory [2].

We can utilise this representation to convert between Vec and List at zero runtime cost, by using the **repr** and **unrepr** operators of the language (defined in section 3). Specifically, we can define the functions

$$\begin{aligned}
&\mathsf{forget\text{-}length} : \mathsf{Vec}\ T\ n \to \mathsf{List}\ T\\
&\mathsf{forget\text{-}length}\ v = \mathbf{let}\ (l, _) = \mathbf{repr}\ v\ \mathbf{in}\ l\\
\\
&\mathsf{recall\text{-}length} : (l : \mathsf{List}\ T) \to \mathsf{Vec}\ T\ (\mathsf{length}\ l)\\
&\mathsf{recall\text{-}length}\ l = \mathbf{unrepr}\ (l, \mathsf{refl})
\end{aligned}$$

and it holds by reflexivity that forget-length is a left inverse of recall-length.

2.3 General reindexing

The idea from the previous example can be generalised to any data type. In general, suppose that we have two inductive families

$$\mathsf{F} : \Xi \to \mathcal{U} \qquad \mathsf{G} : \Xi \to X\ \xi \to \mathcal{U}$$

for some index family $X : \Xi \to \mathcal{U}$. If we hope to represent G as some refinement of F then we must be able to provide a way to compute G's extra indices X from F, like we computed Vec's extra Nat index from List with length in the previous example. This means that we need to provide a function

$$\mathsf{comp} :_{\forall \xi}\ \mathsf{F}\ \xi \to X\ \xi$$

which can then be used to form the family

$$\mathsf{F}^{\mathsf{comp}}\ \xi\ x := \{f : \mathsf{F}\ \xi \mid \mathsf{comp}\ f = x\}.$$

If G is 'equivalent' to the algebraic ornament of F by the algebra defining comp (given by an isomorphism between the underlying polynomial functors), then it is also equivalent to the Σ-type above. The 'recomputation lemma' of algebraic ornaments [7] then arises from its projections. Our system allows us to *set* the representation of G as $\mathsf{F}^{\mathsf{comp}}$, so that the forgetful map from G to F is the identity at runtime.

2.4 Zero-copy deserialisation

The machinery of representations can be used to implement zero-copy deserialisation of data formats into inductive types. For example, consider the following

record for a player in a game:

$$\textbf{data}\ \mathsf{Player}\ \left\{\begin{array}{l}\mathsf{player} : (\mathsf{position} : \mathsf{Position}) \\ \quad \to (\mathsf{direction} : \mathsf{Direction}) \\ \quad \to (\mathsf{items} : \mathsf{Fin}\ \mathsf{MAX\_INVENTORY}) \\ \quad \to (\mathsf{inventory} : \mathsf{Inventory}\ \mathsf{items}) \to \mathsf{Player}\end{array}\right\}$$

We can use the Fin type to maintain the invariant that the inventory has a maximum size. Additionally, we can index the Inventory type by the number of items it contains, which might be defined similarly to Vec:

$$\textbf{data}\ \mathsf{Inventory}\ (n : \mathsf{Nat})\ \left\{\begin{array}{l}\mathsf{empty} : \mathsf{Inventory}\ 0 \\ \quad \mathsf{add} : \mathsf{Item} \to \mathsf{Inventory}\ n \to \mathsf{Inventory}\ (1+\ n)\end{array}\right\}$$

We can use the full power of inductive families to model the domain of our problem in the way that is most convenient for us. If we were writing this in a lower-level language, we might choose to use the serialised format directly when manipulating the data, relying on the appropriate pointer arithmetic to access the fields of the serialised data, to avoid copying overhead. Representations allow us to do this while still being able to work with the high-level inductive type.

We can define a representation for Player as a pair of a byte buffer and a proof that the byte buffer contents correspond to a player record. Similarly, we can define a representation for Inventory as a pair of a byte buffer and a proof that the byte buffer contents correspond to an inventory record of a certain size. The projection $\mathsf{inventory} : (p : \mathsf{Player}) \to \mathsf{Inventory}\ p.\mathsf{items}$ is compiled into some code to slice into the inventory part of the player's byte buffer. We assume that the standard library already represents Fin in the same way as Nat, so that reading the items field is a constant-time operation (we do not need to build a unary numeral). We can thus define the representation of Player as

$$\textbf{repr}\ \mathsf{Player}\ \textbf{as}\ \{\mathsf{Buf} \mid \mathsf{IsPlayer}\}\ \left\{\begin{array}{r}\mathsf{player}\ \textbf{as}\ \mathsf{buf\text{-}is\text{-}player} \\ \mathsf{elim}_{\mathsf{Player}}\ \textbf{as}\ \mathsf{elim\text{-}buf\text{-}is\text{-}player} \\ \textbf{by}\ \mathsf{elim\text{-}buf\text{-}is\text{-}player\text{-}id}\end{array}\right\}$$

with an appropriate definition of IsPlayer which refines a byte buffer. We will provide the full details of this construction in the final paper.

2.5 Transitivity

Representations are transitive, so in the previous example, the 'terminal' representation of Vec also depends on the representation of List. It is possible to define a custom representation for List itself, for example a heap-backed array or a finger tree, and Vec would inherit this representation. However it will still be the case that $\textbf{Repr}\ (\mathsf{Vec}\ T\ n) \equiv \mathsf{List}\ T$, which means the **repr**/**Repr** operators only look at the immediate representation of a term, not its terminal representation. Regardless, we can construct predicates that find types which satisfy a

certain 'eventual' representation. For example, given a Buf type of byte buffers, we can consider the set of all types which are eventually represented as a Buf:

$$\textbf{data}\ \mathsf{ReprBuf}\ (T : \mathcal{U})\ \left\{\begin{array}{c}\mathsf{buf} : \mathsf{ReprBuf}\ \mathsf{Buf} \\ \mathsf{from} : \mathsf{ReprBuf}\ (\textbf{Repr}\ T) \to \mathsf{ReprBuf}\ T \\ \mathsf{refined} : \mathsf{ReprBuf}\ T \to \mathsf{ReprBuf}\ \{t : T \mid P\ t\}\end{array}\right\}$$

Every such type comes with a projection function to the Buf type

$$\begin{aligned}&\mathsf{as\text{-}buf} :_{\forall T.\mathsf{ReprBuf}\ T}\ T \to \mathsf{Buf}\\ &\mathsf{as\text{-}buf}\ \mathsf{buf}\ x = x\\ &\mathsf{as\text{-}buf}\ (\mathsf{from}\ t)\ x = \mathsf{as\text{-}buf}\ t\ (\textbf{repr}\ x)\\ &\mathsf{as\text{-}buf}\ (\mathsf{refined}\ t)\ (x, \_) = \mathsf{as\text{-}buf}\ t\ x\end{aligned}$$

which eventually computes to the identity function after applying **repr** the appropriate amount of times. Upon compilation, every type is converted to its terminal representation, and all **repr** calls are erased, so the as-buf function is effectively the identity function at runtime.[4]

3 A type system for data representations

This section provides an overview of the language λ_{REP}, which has dependent types and representations for inductive families and global function definitions. We start with a core language with inductive families λ_{IND}, that is extended with data representations to form λ_{REP}. All of the examples in the paper are written in a surface language that elaborates to λ_{REP}.

The core languge λ_{IND}, is a dependent type theory with Π and a Coquand-style universe hierarchy $\mathcal{U}_i$ [10, 2.1], extended with strictly positive inductive families and global definitions similarly to [9]. We follow a similar approach to [5] by packaging named inductive constructions and global function definitions into a signature Σ, and indexing contexts by signatures.

3.1 Extending λ_{IND} with representations

We extend the language λ_{IND} to form λ_{REP}, which allows users to define custom representations for inductive types and global functions. First, we add a type former

$$\textbf{Repr} : \text{Ty}\ (\Sigma \mid \Gamma) \to \text{Ty}\ (\Sigma \mid \Gamma) \tag{5}$$

along with two new terms in the syntax, forming an isomorphism

$$\textbf{repr} : \text{Tm}\ (\Sigma \mid \Gamma)\ T \simeq \text{Tm}\ (\Sigma \mid \Gamma)\ (\textbf{Repr}\ T) : \textbf{unrepr}\,. \tag{6}$$

[4] We do not guarrantee that an invocation of as-buf will be entirely erased, but rather that any invocation will eventually produce the identity function without having to perform a case analysis on its T subject.

which holds definitionally and preserves Π and universes. The type **Repr** T is the defined representation of the type T. The term **repr** takes a term of type T to its representation of type **Repr** T, and **unrepr** undoes the effect of **repr**, treating a represented term as an inhabitant of its original type. These new constructs satisfy certain equalities, some of which are given in fig. 1. New valid signature items are introduced, corresponding to representation definitions for components of inductive families and global function definitions.

$$\frac{\textsc{Repr-Ctor-Id}\quad \textbf{repr}\ \mathsf{c}\ \Pi\ \textbf{as}\ \kappa \in \Sigma}{\Sigma \mid \Gamma \vdash \textbf{repr}\ (\mathsf{c}\ \delta\ \pi) \equiv \kappa[\delta, \pi] : A[\delta, \xi[\pi]]} \qquad \frac{\textsc{Repr-Data-Id}\quad \textbf{repr}\ \mathsf{D}\ \Delta\ \Xi\ \textbf{as}\ A \in \Sigma}{\Sigma \mid \Gamma \vdash \textbf{Repr}\ (\mathsf{D}\ \delta\ \psi) \equiv A[\delta, \psi]\ \text{type}}$$

Fig. 1. Definitional equalities for **Repr** and **repr** relating to data types and constructors with defined representations. Similar equalities hold for representations of global function definitions and eliminators, albeit propositionally.

We state some basic lemmas below. The proof are left to the full version of the paper, along with the formalisation of computational irrelevance.

Lemma 1. *The term formers* ***repr*** *and* ***unrepr*** *are injective, i.e.*

$$\frac{\Sigma \mid \Gamma \vdash \boldsymbol{repr}\ t \equiv \boldsymbol{repr}\ t' : \boldsymbol{Repr}\ T}{\Sigma \mid \Gamma \vdash t \equiv t' : T} \qquad \frac{\Sigma \mid \Gamma \vdash \boldsymbol{unrepr}\ t \equiv \boldsymbol{unrepr}\ t' : T}{\Sigma \mid \Gamma \vdash t \equiv t' : \boldsymbol{Repr}\ T}$$

Lemma 2. *The type former* ***Repr*** *is injective up to internal isomorphism, i.e.*

$$\frac{\Sigma \mid \Gamma \vdash \boldsymbol{Repr}\ T \equiv \boldsymbol{Repr}\ T'\ \textit{type}}{\Sigma \mid \Gamma \vdash p : T \simeq T'}$$

Moreover, this isomorphism is computationally irrelevant.

4 Translating from λ_{REP} to λ_{IND}

We can define a translation step $\mathcal{R}$ from λ_{REP} to $\lambda_{\text{IND}}^{\text{Ext}}$, meant to be applied during the compilation process. More specifically, the translation target is the extensional flavour of λ_{IND} by adding the equality reflection rule. We do this by translating well-formed contexts, substitutions, types, and terms in a mutual manner such that definitional equality is preserved. $\mathcal{R}$ preserves the structure of λ_{REP}, but maps constructs to their terminal representations. Eliminator coherence rules are preserved by reflecting the propositional coherence rules provided by the defined representations. We will prove some desired properties of $\mathcal{R}$ [3] such as typing and computational soundness, and preservation of consistency. The final program can then be converted into a simply-typed language which erases irrelevant data. We can recover a program in λ_{IND} by translating extensional typing derivations to intensional proofs [15].

5 Implementation

Superfluid is a programming language with dependent types, $\mathcal{U} : \mathcal{U}$, quantities, inductive families and dependent pattern matching. Its compiler is written in Haskell and the compilation target is JavaScript. Dependent pattern matching in Superfluid is elaborated to a core language with internal eliminators. The $\mathcal{R}$ transformation is then applied to the core program, which erases all inductive constructs with defined representations. This is finally translated to JavaScript, erasing all irrelevant data. As a result, we are able to represent Nat as JavaScript's BigInt, and List T/SnocList T/Vec T n as JavaScript's arrays with the appropriate index refinement, such that we can convert between them without any runtime overhead.

References

1. The GNU MP bignum library. `https://gmplib.org/`, accessed: 2024-12-8
2. Atkey, R.: Syntax and semantics of quantitative type theory. In: LICS 2018. pp. 56–65. LICS '18, Association for Computing Machinery, New York, NY, USA (Jul 2018)
3. Boulier, S., Pédrot, P.M., Tabareau, N.: The next 700 syntactical models of type theory. In: CPP 2017. pp. 182–194. CPP 2017, Association for Computing Machinery, New York, NY, USA (Jan 2017)
4. Brady, E., McBride, C., McKinna, J.: Inductive families need not store their indices. In: Types for Proofs and Programs. pp. 115–129. Springer Berlin Heidelberg (2004)
5. Cockx, J., Abel, A.: Elaborating dependent (co)pattern matching. Proc. ACM Program. Lang. **2**(ICFP), 1–30 (Jul 2018)
6. Cockx, J., Tabareau, N., Winterhalter, T.: The taming of the rew: a type theory with computational assumptions. Proc. ACM Program. Lang. **5**(POPL), 1–29 (Jan 2021)
7. Dagand, P.E., McBride, C.: A categorical treatment of ornaments. arXiv [cs.PL] (Dec 2012)
8. Diehl, L., Firsov, D., Stump, A.: Generic zero-cost reuse for dependent types. Proc. ACM Program. Lang. **2**(ICFP), 1–30 (Jul 2018)
9. Goguen, H., McBride, C., McKinna, J.: Eliminating dependent pattern matching. In: Algebra, Meaning, and Computation, pp. 521–540. Lecture notes in computer science, Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
10. Gratzer, D., Kavvos, G.A., Nuyts, A., Birkedal, L.: Multimodal dependent type theory. arXiv [cs.LO] (Nov 2020)
11. Mcbride, C., Mckinna, J.: The view from the left. JFP **14**(1), 69–111 (Jan 2004)
12. Moon, B., Eades, III, H., Orchard, D.: Graded modal dependent type theory. In: Programming Languages and Systems. pp. 462–490. Springer (2021)
13. Wadler, P.: Views: a way for pattern matching to cohabit with data abstraction. In: Proceedings of the 14th ACM SIGACT-SIGPLAN symposium on Principles of programming languages. pp. 307–313. POPL '87, Association for Computing Machinery, New York, NY, USA (Oct 1987)
14. Wadler, P.: Deforestation: transforming programs to eliminate trees. Theor. Comput. Sci. **73**(2), 231–248 (Jun 1990)
15. Winterhalter, T., Sozeau, M., Tabareau, N.: Eliminating reflection from type theory. In: CPP 2019. ACM, New York, NY, USA (Jan 2019)