



Kiss, C., Pickering, M., & Wu, N. (2018). Generic Deriving of Generic Traversals. *Proceedings of the ACM on Programming Languages*, 2(ICFP), [85]. <https://doi.org/10.1145/3236780>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1145/3236780](https://doi.org/10.1145/3236780)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via ACM at <https://dl.acm.org/citation.cfm?doid=3243631.3236780> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available: <http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>



Generic Deriving of Generic Traversals

CSONGOR KISS, Imperial College London, United Kingdom

MATTHEW PICKERING, University of Bristol, United Kingdom

NICOLAS WU, University of Bristol, United Kingdom

Functional programmers have an established tradition of using traversals as a design pattern to work with recursive data structures. The technique is so prolific that a whole host of libraries have been designed to help in the task of automatically providing traversals by analysing the generic structure of data types. More recently, lenses have entered the functional scene and have proved themselves to be a simple and versatile mechanism for working with product types. They make it easy to focus on the salient parts of a data structure in a composable and reusable manner.

This paper uses the combination of lenses and traversals to give rise to a library with unprecedented expressivity and flexibility for querying and modifying complex data structures. Furthermore, since lenses and traversals are based on the generic shape of data, this information is used to generate code that is as efficient as hand-optimised versions. The technique leverages the structure of data to produce generic abstractions that are then eliminated by the standard workhorses of modern functional compilers: inlining and specialisation.

CCS Concepts: • **Software and its engineering** → **Data types and structures**; *Functional languages*; *Reusability*;

Additional Key Words and Phrases: generic programming, traversals, lenses, extensibility

ACM Reference Format:

Csongor Kiss, Matthew Pickering, and Nicolas Wu. 2018. Generic Deriving of Generic Traversals. *Proc. ACM Program. Lang.* 2, ICFP, Article 85 (September 2018), 30 pages. <https://doi.org/10.1145/3236780>

1 INTRODUCTION

Traversals are a ubiquitous way of querying and manipulating data. They provide a reliable interface for working with data types in a structured and predictable manner. An appropriate suite of traversals is a valuable tool that eases the task of constructing programs that interact with diverse data. Unfortunately, writing traversals quickly becomes tedious work that requires continuous curation as code evolves over time. Naturally, our desire is to have our traversals provided for us.

Our goal is to identify a declarative family of useful traversals that is expressive enough for a wide range of practical programming tasks. Furthermore, we want to completely remove the burden of writing these traversals by automatically deriving them whenever possible. Not only that, but we also want to generate code that performs as well as hand-written code.

The most famous existing solution, Scrap Your Boilerplate (SYB) [Lämmel and Peyton Jones 2003], treats this problem by performing run-time type tests to decide which part of the tree to traverse. This leads to a flexible interface at the expense of performance: the approach is famously

Authors' addresses: Csongor Kiss, Department of Computing, Imperial College London, United Kingdom, csongor.kiss14@imperial.ac.uk; Matthew Pickering, Department of Computer Science, University of Bristol, United Kingdom, matthew.pickering@bristol.ac.uk; Nicolas Wu, Department of Computer Science, University of Bristol, United Kingdom, nicolas.wu@bristol.ac.uk.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2018 Copyright held by the owner/author(s).

2475-1421/2018/9-ART85

<https://doi.org/10.1145/3236780>

slow. In this paper we present `generic-lens`, a Haskell library that provides a suite of traversals that is both faster and richer than SYB. In essence, we believe that it is time to scrap your SYB.

We are not the first to optimise SYB, and like our predecessors [Adams et al. 2015; Yallop 2017], we make use of the fact that much of the information required for traversals is *statically* known, thereby avoiding *dynamic* checks at run-time. Our innovation is to work with types and use these to infer generated code generically in a suitable form for an automatic evaluator to optimise effectively. By leveraging the static information that is provided by the generic structure of data, we are able to produce much better generated code. The generic abstraction is eliminated.

To have a taste of the `generic-lens` library, consider a data type of weighted trees. There are two type parameters, which correspond to the type of elements and weights in the tree:

```
data WTree a w = Leaf a
  | Fork (WTree a w) (WTree a w)
  | WithWeight (WTree a w) w
```

Suppose the object is to gather all the weights in the tree. The `generic-lens` library provides a traversal for this data type called `param` which takes a type-level integer as an argument. This allows a traversal to specify which parameter it wants to focus on. Counting from the right the parameter `@0` indicates focus on the 0th parameter, which is `w`, the weights in the tree.

```
weights :: WTree Int Int → [Int]
weights = toListOf (param @0)
```

The `toListOf` combinator takes a traversal and turns it into a fold which summarises what it is focusing on. Applying `weights` to a tree will then correctly return a list of all the weights, even though the node values are also integers. Thankfully, if an incorrect index is used, the `generic-lens` library produces a bespoke compile-time error to help identify the mistake.

Contributions. The primary contribution of this paper is a demonstration of how guilt-free generic programming using existing language features can be achieved. More specifically, our contributions are the following:

- (1) We specify a high-level interface for describing a family of useful lenses, prisms, and traversals in a type-directed manner.
- (2) We introduce a technique that allows generic traversals over multiple type parameters.
- (3) We outline the implementation of `generic-lens`, a library that implements this interface using generics.
- (4) We provide benchmarks which demonstrate that `generic-lens` is as fast as hand-written code. We also discuss the optimisations which we require a compiler to perform.

The remainder of this paper is structured as follows. Section 2 motivates the use of the generic lenses and traversals. Section 3 describes the interface that our library `generic-lens` supplies. Section 4 gives the background necessary for the implementation of our library. We then move onto implementing generic traversals that are directed by types in Section 5, by parameters in Section 6, and by constraints in Section 7. The performance of `generic-lens` is considered in Section 8 and we evaluate it with benchmarks in Section 9. Finally, we discuss related work in Section 10, and conclude in Section 11.

2 TYPE-DIRECTED QUERIES

Suppose you are running a biscuit distribution company. You have customers who place orders for biscuits which you need to keep track of and process. In addition, you allow customers to prioritise their biscuit orders, which are then distributed from an entirely separate distribution facility.

Lenses and Traversals

view :: *Lens* *s t a b* → *s* → *a*

update :: *Lens* *s t a b* → *b* → *s* → *t*

modify :: *Lens* *s t a b* → (*a* → *b*) → *s* → *t*

toListOf :: *Traversal* *s s a a* → *s* → [*a*]

over :: *Traversal* *s t a b* → (*a* → *b*) → *s* → *t*

traverseOf :: *Traversal* *s t a b* → (∀*g*.*Applicative* *g* ⇒ (*a* → *g b*) → *s* → *g t*)

(◦) :: *o*₁, *o*₂ ∈ {*Lens*, *Traversal*} ⇒ *o*₁ *s t c d* → *o*₂ *c d a b* → (*o*₁ ∨ *o*₂) *s t a b*

Generic Lenses

field :: ∀*name* *s t a b*. *HasField* *name* *s t a b* ⇒ *Lens* *s t a b*

typed :: ∀*a s*. *HasType* *a s* ⇒ *Lens* *s s a a*

position :: ∀*pos* *s t a b*. *HasPosition* *pos* *s t a b* ⇒ *Lens* *s t a b*

super :: ∀*sup sub*. *Subtype* *sup sub* ⇒ *Lens* *sub sub sup sup*

Generic Traversals

types_ :: ∀*a s*. *HasTypes* *s a* ⇒ *Traversal* *s s a a*

param :: ∀*pos* *s t a b*. *HasParam* *pos* *s t a b* ⇒ *Traversal* *s t a b*

constraints :: ∀*c s t*. *HasConstraints* *c s t* ⇒
(∀*g*.*Applicative* *g* ⇒ (∀*a b c*. *a b* ⇒ *a* → *g b*) → *s* → *g t*)

Fig. 1. The generic lens and traversal interface to generic-lens

To this end, you implement *Item*, a data type to represent a type of biscuit, and *Invoice*, a data type to represent a single order. It is parameterised by the type of priority that is assigned to the orders. Finally, *Orders*, a top-level data structure which contains the normal and priority queues. The priority queue has an augmented priority field that keeps track of the priority level.

data *Item* = *Item* { *name* :: *String*, *cost* :: *Cost* }

newtype *Cost* = *Cost* *Double*

data *Invoice* *p* = *Invoice* { *item* :: *Item*, *name* :: *String*, *number* :: *Int*, *priority* :: *p* }

data *Orders* = *Orders* [*Invoice* *Int*] [*Invoice* (*Int*, *Double*)]

Here is how a Bourbon Biscuit is represented, along with the invoices of some eager biscuit eaters.

bourbon :: *Item*

bourbon = *Item* "Bourbon" (*Cost* 100)

orders = *Orders* [*Invoice* *bourbon* "Earl" 1 0, *Invoice* *bourbon* "Johnny" 2 2]
[*Invoice* *bourbon* "George" 2 (0, 3)]

Generic queries from *generic-lens* will now be used to interrogate specific aspects of this data structure. First a specification by example will be provided before an in-depth explanation in the next section. The interface of the *generic-lens* library is summarised in Figure 1.

Starting from the simplest example, *field* derives a *Lens* which focuses on a named field in a data type. The lens *field* @"name" focuses on the "name" field of an *Item*. This makes use of visible type application [Eisenberg et al. 2016] to supply the static argument "name" to *field*. Once this field has been focused it can be viewed or updated.

```

> ghci> view (field @"name") bourbon
> "Bourbon"
> ghci> update (field @"cost") (Cost 110) bourbon
> Item "Bourbon" (Cost 110)

```

Why is the in-built record selector not used? For it is not *compositional*. Lenses can be composed together using the composition operator \circ in order to inspect nested fields. For example, finding the name of an item is achieved by composing the two *field* lenses like so:

```

nameOfItem :: Invoice p → String
nameOfItem = view (field @"item" ∘ field @"name")

```

These lenses are read left-to-right: a lens that finds the field called "item" is applied, followed by a lens that finds the field "name".

This is all well and good if only nested products are presented but no good at all for modifying many parts of a data structure at once. As a special thank you to your customers, you wish to decrease the cost of all invoices. This is achieved with a *types* traversal which creates a *Traversal* that focuses on every part of a data structure with a specific type.

```

thankYou :: Orders → Orders
thankYou = over (types @Cost) (λ(Cost c) → Cost (c × 0.85))

```

Later you realise that you only really want to thank your priority customers. In order to do this the focus must first be restricted to the priority queue before the previous incantation is reused. The *position* lens selects the *k*th field of a data type by its position in the data declaration.

```

thankYouPriority :: Orders → Orders
thankYouPriority = over (position @2 ∘ types @Cost) (λ(Cost c) → Cost (c × 0.85))

```

The shows that a lens and a traversal have been composed to get a traversal. A lens is a special case of a traversal that also allows a value to be extracted by focusing on one item.

Finally, like any good business, you give your customers the choice and opportunity to upgrade their standard orders to premium orders. In order to do so, an *Invoice Int* must be modified into an *Invoice (Int, Double)*. The *param* traversal is used in order to modify the 0th type parameter from the right of *Invoice* from an *Int* to a *(Int, Double)*.

```

upgrade :: Double → Invoice Int → Invoice (Int, Double)
upgrade bribe invoice = over (param @0) (λi → (i, bribe)) invoice

```

The above example highlights how a traversal can change the type of its argument. Traditionally type changing is difficult to implement in generic traversal frameworks such as SYB.

At the end of the year, your auditors want to see a summary of all the items you have sold this year. They don't care whether they were priority orders or not: you just need to extract all the *Items* you have sold. The *types* traversal can be used to focus on all *Items* in the tree to be extracted:

```

audit :: Orders → [Item]
audit = toListOf (types @Item)

```

The *toListOf* combinator summarises a *Traversal* by returning all the parts it focuses on.

These examples have demonstrated how it is possible to concisely traverse, modify, inspect, and analyse the biscuit pipeline. This was made possible by the use of lenses and traversals that are generically derived from the data types involved. The next section describes the generic combinators that were used in these examples.

3 INTERFACE

This section discuss ways of identifying certain parts of algebraic data types using a type-directed approach. These can be classified into the following three categories, based on the underlying structure of a data type.

Lenses : Patterns applicable to data types made from *products*

Prisms : Patterns applicable to data types made from *sums*

Traversals : Patterns applicable to data types made from *sums of products*

These abstractions are known together as *optics*. This paper will concentrate on lenses and traversals. Prisms follow the same principles so are discussed only briefly.

3.1 Lenses

A lens focuses on one part of a product. The focus can then be viewed and updated whilst the rest of the structure remains unchanged. A lens $l :: \text{Lens } s \ t \ a \ b$ can be read as saying that l is a lens whose source is of type s , its focus is on a value of type a which when changed to a value of type b replaces the value of type a in s and produces a product of type t .

A lens can be used in three primitive ways.

$\text{view} :: \text{Lens } s \ t \ a \ b \rightarrow s \rightarrow a$

$\text{update} :: \text{Lens } s \ t \ a \ b \rightarrow b \rightarrow s \rightarrow t$

$\text{modify} :: \text{Lens } s \ t \ a \ b \rightarrow (a \rightarrow b) \rightarrow s \rightarrow t$

Lenses with this interface are already well established [Pickering et al. 2017]. The *view* operation extracts a component from its context. The *update* operation updates a structure. The *modify* function is included as a means of efficiently viewing and updating a structure in a single step.

A contribution of this paper is to derive a number of generic lenses. The first consideration is the ways in which access to different parts of a product data type can be specified. To achieve this various *Lenses* that focus on precisely one part of a product are derived.

3.1.1 By name. For a data type with named fields it is possible to specify the lens that focuses on a field with a certain name. As each field must have a unique name, this provides a way of specifying a unique field in a larger product. A combinator named *field* is defined that provides this lens for all suitable types.

$\text{field} :: \text{HasField } \text{name } s \ t \ a \ b \Rightarrow \text{Lens } s \ t \ a \ b$

HasField name s t a b instances are derived generically. The constraint means that the type s has a field called *name* of type a , and if the field is changed from a to b , a structure of type t is obtained. To illustrate this, consider a change of the *cost* field of an *Item*:

```
> ghci> modify (field @"cost") (\(Cost c) -> (Cost (c + 5))) bourbon
> Item "Bourbon" (Cost 105)
```

In this case, the type of the derived instance is the following:

$\text{field @"cost" @Item} :: \text{Lens } \text{Item } \text{Item } \text{Cost } \text{Cost}$

That is, the type of the *Cost* field can not be changed within *Item* – of course, as it is fixed to be *Cost* in *Item*'s definition. Contrast this with *Invoice p*: it is parameterised by a type variable. *Invoice a* can be changed to *Invoice b*, as long as inner a can be changed into b .

$\text{field @"priority" @(Invoice _)} :: \text{Lens } (\text{Invoice } a) (\text{Invoice } b) a b$

Accordingly, this version of the lens readily allows type-changing manipulations to be carried out:

```
> ghci> modify (field @"priority") (\i → (i, 0)) (Invoice bourbon "Johnny" 2 2)
> Invoice bourbon "Johnny" 2 (2, 0)
```

A lens focuses on exactly one part of a structure, meaning in this case that it must contain the field of interest. But what happens if a field that does not exist is requested? A type error is thrown!

```
> ghci> view (field @"weight") bourbon
> error: * The type Item does not contain a field named weight.
```

Thus the existence of a field can be *statically* determined. Notice that the error message generated by the library is informative and hides away the underlying complexities. The approach not only provides a pleasant user experience, but also obviates the need for any dynamic checks.

3.1.2 By type. Often it is burdensome to access the fields by name, as it can change over time. In many cases it does not matter how exactly the subpart can be located, as long as it is uniquely identified by its type. The *typed* lens focuses on the unique type in a product.

$$\text{typed} :: \text{HasType } a \ s \Rightarrow \text{Lens } s \ s \ a \ a$$

For example, *Item* has one *Cost* field so the *typed* lens can be used to update and modify it.

```
> ghci> update (typed @Cost) (Cost 200) bourbon
> Item "Bourbon" (Cost 200)
```

Often in practice, the type of interest can even be inferred from the context, and does not need to be specified explicitly. For example:

```
> ghci> modify typed ("Chocolate "++) bourbon
> Item "Chocolate Bourbon" (Cost 100)
```

It is clear that appending text requires a *String*, and therefore *typed* knows which field to select. As expected, requesting a lens for a type not contained in the product yields a type error.

The *typed* lens is monomorphic as it is complicated to specify precisely when it is safe to change the type. For instance, the target could be “lost” if the field’s type is changed to something already present in the structure, as this type would no longer be uniquely identifiable. These complications do not arise for the *field* and *position* lenses.

3.1.3 By position. Not all product types have named fields. For example, consider *Orders*: it contains no named fields, yet it may still be desirable to restrict attention to either the first or second order queue. In this case fields can be referred to positionally.

$$\text{position} :: \text{HasPosition } pos \ s \ t \ a \ b \Rightarrow \text{Lens } s \ t \ a \ b$$

Indexed from 1, the *position* lens focuses on the *k*th field in a product. *HasPosition* instances are derived generically for all product types.

```
> ghci> view (position @2) orders
> [Invoice bourbon "George" 2 (0, 3)]
```

Trying to access an “out of bounds” element results in a type error:

```
> ghci> view (position @3) orders
> error: * The type Orders does not contain a field at position 3
```

This lens works equally well for any product data type including in-built types such as tuples.

3.1.4 *By structure.* Finally, the *super* lens generalises the *field* lens to focus on a collection of fields rather than just one. The *Subtype sup sub* constraint holds if the data type *sub* contains all the fields labels (with the same types) as *sup* contains.

$$\text{super} :: \text{Subtype sup sub} \Rightarrow \text{Lens sub sub sup sup}$$

Consider a new data type *WeighedItem* which adds a new weight field to *Item* so the postage cost for orders can be calculated. The *super* lens will be used to extract a value of type *Item* from *WeighedItem*. As such, *WeighedItem* is a subtype of *Item* as it contains all the fields which *Item* contains. Thus *super @Item @WeighedItem :: Lens WeighedItem WeighedItem Item Item*.

```
newtype Weight = Weight Double
data WeighedItem = WItem { name :: String, cost :: Cost, weight :: Weight }
```

Now *super* can be used to modify several fields at once.

```
> ghci> view (super @Item) (WItem "Bourbon" (Cost 2000) (Weight 0.03))
> Item "Bourbon" (Cost 2000)
> ghci> update (super @Item) bourbon (WItem "Bourbon+" (Cost 500) (Weight 0.03))
> WItem "Bourbon" (Cost 100) (Weight 0.03)
```

This kind of lens is particularly useful in data processing pipelines where additional steps add computed fields to a data type.

3.2 Prisms

Prisms are the dual to lenses: while a lens focuses on one part of a product, a prism focuses on one part of a sum. As such, the focused value might not be present. Prisms can be used in the other way; they can construct the sum by injecting in the focused part. Prisms for data types which are made from sums can be derived. A prism $p :: \text{Prism } s \ t \ a \ b$ consumes values of type *s* and, supposing an *a* can be turned into a *b*, it produces values of type *t*.

```
match :: Prism s t a b → s → Either t a
build  :: Prism s t a b → b → t
```

Since prisms behave similarly to lenses, they are described only briefly in this section to give an intuition of their use. The remainder of the paper focuses on lenses and traversals.

Consider a simple sum type *D* that makes use of the sum of constructors:

```
data D = DInt Int | DPair Bool String
```

As with lenses, three different ways of deriving prisms for sum types are provided.

3.2.1 *By name.* The *_Ctor* prism selects a constructor by its name.

$$_Ctor :: \text{AsConstructor name } s \ t \ a \ b \Rightarrow \text{Prism } s \ t \ a \ b$$

```
> ghci> match (_Ctor @"DInt") (DInt 1)
> Right 1
```

3.2.2 *By type.* The `_Typed` prism selects a constructor by the type inside the constructor. Constructors that contain multiple values are viewed as a tuple of those values.

```
_Typed :: AsType s a ⇒ Prism s s a a
> ghci> build _Typed (False, "wurbble") :: D
> DPair False "wurbble"
```

3.2.3 *By structure.* The `_Sub` prism allows a substructure to be injected into a superstructure.

```
_Sub :: AsSubtype sub sup ⇒ Prism sup sup sub sub
```

A sum `Sub` is a subtype of another sum `Sup` if a value of `Sub` can be given (modulo naming of constructors) whenever a value of `Sup` is expected. Consider the data type `E`, a supertype of `D`:

```
data E = EInt Int | EPair Bool String | EChar Char
```

The parameter `_Sub @D` is used to pattern match on values of `E` as if they were `D` (in this case a failure as `D` has no corresponding `Char` constructor):

```
> ghci> match (_Sub @D) (EChar 'a')
> Left (EChar 'a')
```

Or in the other direction, build values of `E` from `D`:

```
> ghci> build _Sub (DInt 10) :: E
> EInt 10
```

The combination of prisms and lenses make for an extremely powerful and versatile querying language when combined with traversals, which are discussed next.

3.3 Traversals

For algebraic data types (i.e. those constructed using a combination of sums and products), *traversals* are derived. A traversal written `Traversal s t a b` walks over a value of type `s`, modifying all `as` into `bs`, resulting in a value of type `t`. For example, imagine a traversal `tree :: Traversal (Tree a) (Tree b) a b` that focuses on all the elements in a tree. While the most general combinator is `traverseOf`, the specialisations `over` and `toListOf` which modify and summarise will usually be used instead.

```
over :: Traversal s t a b → (a → b) → s → t
toListOf :: Traversal s s a a → s → [a]
traverseOf :: Traversal s t a b → (∀g.Applicative g ⇒ (a → g b) → s → g t)
```

The different traversals that can be generically derived are now described.

3.3.1 *By type.* The `types` function allows all values of a given type in a data type to be traversed.

```
types :: HasTypes s a ⇒ Traversal s s a a
```

Recalling an example introduced in the previous section, `types @Cost` generates a traversal that considers all values of type `Cost` wherever they are located in a structure. This can be used to uniformly modify all the costs in a data structure.

```
costInc :: HasTypes t Cost ⇒ t → t
costInc = over (types @Cost) (λ(Cost c) → Cost (c + 5))
```

By using the `types` combinator, it is not necessary to spell out the recursion over `Orders`. Furthermore, the function `costInc` is polymorphic and will work for any data structure containing costs. For similar reasons to the `typed` lens, the `types` traversal is monomorphic and can not change types.

However, there is a danger lurking in the shadows: when using *types*, a type that is too general must be avoided. Consider the running example again where only the priorities of a normal invoice are to be modified. A first attempt might be:

$$\begin{aligned} \text{modifyPriority} &:: (\text{Int} \rightarrow \text{Int}) \rightarrow \text{Invoice Int} \rightarrow \text{Invoice Int} \\ \text{modifyPriority} &= \text{over } (\text{types } @\text{Int}) \end{aligned}$$

This will have unexpected consequences as there are other values of type *Int* in the *Invoices*, namely the order number. The modification function will also unfortunately update all of those as well. The tree contains many *Ints* used in different ways. For this reason, the *types* combinator should be used with care. The programmer must maintain good type discipline to avoid semantically different types being traversed together. This is because the way *modifyPriority* was specified did not quite reflect what was actually *meant*. The intention is to update only the *Ints* that are in the priority positions. A type-based query is insufficient here because it cannot distinguish between uses of *Int*. This problem did not exist for the lens version, because that requires the type to appear exactly once, avoiding such clashes.

3.3.2 By parameter. In the previous example, the intention was to select only the values in the priority fields, or in other words, those that correspond to the *p* type parameter. To solve this the library provides traversals that are defined over a specific type parameter. Positional indexing is used to refer to the type parameter of interest.

$$\text{param} :: \text{HasParam } \text{pos } s \ t \ a \ b \Rightarrow \text{Traversal } s \ t \ a \ b$$

Numbering starts from the outside, meaning that the last parameter has the index 0. Trying to access an out of bounds type parameter results in a type error. Using *param*, the *modifyPriority* function can be revised:

$$\begin{aligned} \text{treeIncParam} &:: \text{HasParam } 0 \ s \ s \ \text{Int} \ \text{Int} \Rightarrow s \rightarrow s \\ \text{treeIncParam} &= \text{over } (\text{param } @0) (+1) \end{aligned}$$

This revised definition now properly distinguishes between the different *Ints* in the *Invoices*.

3.3.3 By constraint. The most general type of traversal is the *constrained* traversal. A constrained traversal focuses on all positions in a data type. It does this by requiring that the types in all positions satisfy a constraint, and then uniformly applies a function in terms of this constraint to all fields.

A constrained traversal thus has the following type:

$$\text{constraints} :: \text{HasConstraints } c \ s \ t \Rightarrow \text{Applicative } g \Rightarrow (\forall a \ b. c \ a \ b \Rightarrow a \rightarrow g \ b) \rightarrow s \rightarrow g \ t$$

The user can instantiate the traversal to any type class of their choosing, thereby specifying the traversal strategy. The traversing function has to be one that only has knowledge of what information is available in the class *c*. Via the ad-hoc overloading mechanism of type classes, the function is instantiated to the version specified for each field in *s*.

There are many choices to which *c* could be instantiated. In fact, it is the most general traversal and subsumes the two other traversals that have been discussed. Users might also decide to instantiate *c* to a constraint based on *Data* or *Generic* in order to specify dynamically how each field is processed. The *constraints* traversal just provides a framework whilst the constraint determines how to deal precisely with each subpart.

3.4 Composition

The final ingredient is an overloaded composition operator \circ which can be used to compose together any combination of lenses, prisms and traversals.

The type of this operator can be thought of abstractly as

$$(\circ) :: o_1, o_2 \in \{Lens, Traversal\} \Rightarrow o_1 \text{ s t c d} \rightarrow o_2 \text{ c d a b} \rightarrow (o_1 \vee o_2) \text{ s t a b}$$

The join operation is specified by defining a *Traversal* to be above a *Lens*. The composition operator is presented in this way as in the full generality there are more components (such as prisms) to the hierarchy [Pickering et al. 2017]. The intuition is that a lens is a special case of a traversal where there is exactly one focused element. Being more restrictive allows lenses to support the additional operations of viewing that traversals do not support.

Summary. This section has described the various widgets that allow values to be traversed, modified, and inspected using generic lenses and traversals. These operations form an interface for the generic-lens library, which is summarised in Figure 1.

4 BACKGROUND: LENSES, TRAVERSALS, AND GENERICS

This section begins to describe an efficient implementation of the interface found in Figure 1, while introducing the necessary background that is the foundation for generic traversals.

4.1 Lenses and Traversals

The concrete representations of *Lens* and *Traversal* and the associated operators are described first, before explaining the implementation of the different derived lenses and traversals.

Lenses. The representation of lenses used is the van Laarhoven representation [van Laarhoven 2009]. A van Laarhoven lens is a function of the following type:

$$\text{type Lens s t a b} = \forall f. \text{Functor } f \Rightarrow (a \rightarrow f \text{ b}) \rightarrow (s \rightarrow f \text{ t})$$

The *view* and *update* functions are implemented as required by the interface by suitably instantiating f to the *Const* and *Identity* functor respectively [McBride and Paterson 2008].

$$\begin{aligned} \text{view} &:: \text{Lens s t a b} \rightarrow s \rightarrow a \\ \text{view } l &= \text{getConst} \cdot l \text{ Const} \end{aligned}$$

$$\begin{aligned} \text{update} &:: \text{Lens s t a b} \rightarrow b \rightarrow s \rightarrow t \\ \text{update } l \text{ b} &= \text{runIdentity} \cdot l (\text{const } (\text{Identity } b)) \end{aligned}$$

Traversals. The van Laarhoven representation is also convenient as the type is similar to that of traversals. The type *Traversal s t a b* is implemented with functions of the following type:

$$\text{type Traversal s t a b} = \forall f. \text{Applicative } f \Rightarrow (a \rightarrow f \text{ b}) \rightarrow (s \rightarrow f \text{ t})$$

Again, the interface is implemented by instantiating the applicative f to *Identity* and *Const*, which provides the correct specialisation to implement the functions.

$$\begin{aligned} \text{over} &:: \text{Traversal s t a b} \rightarrow (a \rightarrow b) \rightarrow s \rightarrow t \\ \text{over } t \text{ f} &= \text{runIdentity} \cdot t (\text{Identity} \cdot f) \\ \text{toListOf} &:: \text{Traversal s s a a} \rightarrow s \rightarrow [a] \\ \text{toListOf } t &= \text{getConst} \cdot t (\text{Const} \cdot \text{singleton}) \end{aligned}$$

$$\begin{aligned} \text{traverseOf} &:: \text{Traversal s t a b} \rightarrow (\forall g. \text{Applicative } g \Rightarrow (a \rightarrow g \text{ b}) \rightarrow s \rightarrow g \text{ t}) \\ \text{traverseOf} &= \text{id} \end{aligned}$$

$$\begin{aligned} \text{singleton} &:: a \rightarrow [a] \\ \text{singleton } x &= [x] \end{aligned}$$

Additional justification for this representation is not provided in this paper as it has been extensively studied elsewhere [Bird et al. 2013; Jaskelioff and O'Connor 2015; O'Connor 2011]. In any case, the choice is not crucial to this work. Instead, the same techniques could be applied to the profunctor [Pickering et al. 2017] and other encodings.

Composition. With this representation, \circ simply becomes composition:

$$\begin{aligned} (\circ) &:: \text{Lens } s \ t \ c \ d \rightarrow \text{Lens } c \ d \ a \ b \rightarrow \text{Lens } s \ t \ a \ b \\ (\circ) &= (\cdot) \end{aligned}$$

In order to make lenses and traversals compose, a lens composed with a traversal must result in a traversal. This is easy to observe as the type of a *Traversal* is more constrained than that of a *Lens* because *Functor* is a superclass of *Applicative*.

4.2 Generic Programming

Datatype-generic programming allows data types to be decomposed into their constituent parts, which are shown below:

$$\begin{array}{lll} \mathbf{data} \ f \ :+ : g & = L \ f \ | \ R \ g & \mathbf{data} \ V & \mathbf{newtype} \ M \ (m :: Meta) \ a = M \ a \\ \mathbf{data} \ f \ : \times : g & = f \ : \times : g & \mathbf{data} \ U = U & \mathbf{data} \ Meta = MetaData \ Symbol \\ \mathbf{newtype} \ K \ a = K \ a & & & | MetaCons \ Symbol \\ & & & | MetaSel \ (Maybe \ Symbol) \end{array}$$

This is a sum-of-products representation similar to that proposed by Hinze [2000]. Algebraic data types can be uniformly viewed in this way: choice between constructor variants is encoded as (potentially nested) binary sums ($:+$). A single field of type a inside a constructor is stored as $K \ a$; multiple fields are collected in (potentially nested) binary products ($: \times$). Datatypes with no constructors are represented by V , and constructors with no fields by U .

Additional metadata (name of the datatype, names of constructors, and (optional) names of fields) can be attached to the nodes via M . The meta constructor M makes use of *datatype promotion* [Yorgey et al. 2012], which allows $Meta$'s constructors to be used in a type context. In general, promotion allows data types like $Meta$ and $Bool$ to be used as kinds.

The isomorphism between concrete types and their sum-of-products view is witnessed by an instance of the *Generic* type class:

```
class Generic a where
  type family Rep a :: ★
  from :: a → Rep a
  to :: Rep a → a
```

As the type of the generic view is different for each type, *Generic* associates the concrete type and their representation type via the *Rep* type family [Chakravarty et al. 2005]. Writing these instances is laborious, but straightforward. The Glasgow Haskell Compiler (GHC) provides built-in support for deriving these instances [Magalhães et al. 2010]. In practice, this requires a *deriving Generic* clause to be appended to the data definitions: this is omitted in the presentation.

Now consider the following definition of linked lists:

```
data List a = Empty | Cons a (List a)
```

The generic view of $List \ Int$ has the type

$$\begin{aligned}
 \text{Rep } (\text{List } \text{Int}) &\equiv M (\text{MetaData } \text{"List"}) (M (\text{MetaCons } \text{"Empty"}) U \\
 &\quad \text{:+}: M (\text{MetaCons } \text{"Cons"}) (M (\text{MetaSel } \text{'Nothing}) (K \text{ Int}) \\
 &\quad \text{:x}: M (\text{MetaSel } \text{'Nothing}) (K (\text{List } \text{Int})))
 \end{aligned}$$

Reflecting the algebraic structure of *List Int* to the type-level in this way allows the shape and metadata of the type to be statically introspected. Using this information allows safe and optimal transformations to be derived without having to write boilerplate code.

5 GENERIC TRAVERSALS WITH TYPES

Now that the basic parts of the interface have been implemented, the next task is deriving interesting traversals. Traversals are given special attention since the principles for deriving lenses are the same.

The principle of the implementation is simple. In order to generate an optic for a specific data type, it is converted to its generic representation using *from* :: *Generic a* \Rightarrow *a* \rightarrow *Rep a*. The type family *Rep* turns a type into the type of its generic representation. The type of its generic representation directly corresponds to the structure of the generic representation.

In order to implement a function that consumes the generic representation, a type class must be defined for the desired function to be implemented, followed by a type class instance for each clause of the function. Thus, a separate case will be defined to deal with empty types, products, sums and so on.

There are two complexities that must be considered in the implementation. First, is establishing whether or not access to a specific value at the leaves of a data structure is desired. The second complexity in the implementation is maintaining good type-inference behaviour. To give a sense of how these complexities are dealt with, the implementation of the *types* traversal is first described, followed by an outline of how type inference works.

Implementing types. The first focus is on how traversals are implemented using this machinery. As a reminder, the *types* traversal is indexed by a type, and it provides access to all subparts of a structure that have the specified type. This section implements a naive first attempt at *types* with deficiencies that will be addressed later on.

First, a type class *HasTypes_ s t a b* is created, which represents that *s* contains some (zero or more) values of type *a*, and changing these *as* into *bs* results in a structure of type *t*. The sole member of this type class is the *types_* combinator.

```
class HasTypes_ s t a b where
  types_ :: Traversal s t a b
```

HasTypes_ is the abstract interface that is going to be instantiated by induction over the generic view. Matching on individual cases is done using the auxiliary type class *GHasTypes s t a b*. Then each type whose generic representation admits a *GHasTypes* instance has a *HasTypes_* instance itself derivable via the isomorphism.

```
class GHasTypes s t a b where
  gtypes :: Traversal s t a b
instance (GHasTypes (Rep s) (Rep t) a b, Generic s, Generic t)
  => HasTypes_ s t a b where
  types_ = isoRep · gtypes
```

Where iso_{Rep} is a lens that views a value as its (isomorphic) generic representation by using the methods from the *Generic* type class. Changes made on the generic representations are reflected on the original type, including changes of types.

$$iso_{Rep} :: (Generic\ s, Generic\ t) \Rightarrow Lens\ s\ t\ (Rep\ s)\ (Rep\ t)$$

$$iso_{Rep}\ f = fmap\ to\ \cdot f\ \cdot from$$

The above instance is defined *for all* types, meaning that all *types* queries require s and t to have a *Generic* instance, as per the constraints. However, certain types do not admit a *Generic* instance, namely the non-algebraic primitive types. For these, overlapping instances are defined that are picked instead of the general one above.

instance *HasTypes_ Char Char a b where*
types_ _ = pure

The code above shows the instance for *Char*, and is similar for other primitive types such as *Double*, *Float*, *Int*, and *Integer*. Given that these types are not actually containers, they can not possibly contain any interesting values, thus their traversal is defined as the no-op *pure*.

Additionally, the iso_K and iso_M lenses are defined, which focus on the value inside the K node and the generic structure wrapped by metadata nodes respectively.

$$iso_K :: Lens\ (K\ a)\ (K\ b)\ a\ b$$

$$iso_M :: Lens\ (M\ m\ s)\ (M\ m\ t)\ s\ t$$

$$iso_K\ f\ s = K\ \$\ f\ (unK\ s)$$

$$iso_M\ f\ s = M\ \$\ f\ (unM\ s)$$

The generic cases are now dealt with one-by-one. As the *types* traversal is oblivious to metadata such as constructor names, *gtypes* simply skips over these.

instance *GHasTypes s t a b \Rightarrow GHasTypes (M m s) (M m t) a b where*
gtypes = iso_M \cdot gtypes

Next, sums are handled, which correspond to the constructors of a datatype. Since the aim is to discover every node in the structure, the function *gtypes* is recursively called on whichever case alternative is present. The constraints *GHasTypes* $l_1\ l_2\ a\ b$ and *GHasTypes* $r_1\ r_2\ a\ b$ ensure that both cases can indeed be traversed.

instance (*GHasTypes* $l_1\ l_2\ a\ b$, *GHasTypes* $r_1\ r_2\ a\ b$)
 \Rightarrow *GHasTypes* ($l_1\ \text{:+}\ r_1$) ($l_2\ \text{:+}\ r_2$) $a\ b$ **where**
gtypes $f\ (L\ l) = L\ \$\ gtypes\ f\ l$
gtypes $f\ (R\ r) = R\ \$\ gtypes\ f\ r$

Products are treated similarly: both left and right trees are traversed, looking for *as*.

instance (*GHasTypes* $l_1\ l_2\ a\ b$, *GHasTypes* $r_1\ r_2\ a\ b$)
 \Rightarrow *GHasTypes* ($l_1\ \text{:}\times\ r_1$) ($l_2\ \text{:}\times\ r_2$) $a\ b$ **where**
gtypes $f\ (L\ \text{:}\times\ r) = (\text{:}\times\text{:})\ \$\ gtypes\ f\ l\ \langle*\rangle\ gtypes\ f\ r$

Now the interesting case – that is, when a field of type a is encountered. In this case the search can be stopped with focus on this leaf node.

instance *GHasTypes (K a) (K b) a b where*
gtypes = iso_K

What if the leaf is not actually the desired type, but something else? In the case of a *shallow* traversal, this is where the traversal would stop. Given that these traversals are *deep*, the search continues further to see if this leaf contains any more values of type a , by recursively invoking a *HasTypes_* constraint, and the corresponding *types_* traversal, now for the leaf. Note that this instance is

overlapped by the previous one, as it is strictly more general than the previous case, and it is picked when the above does not match.

```
instance HasTypes_ s t a b  $\Rightarrow$  GHasTypes (K s) (K t) a b where
  gtypes = isoK · types_
```

When the leaves are primitives this process stops, since their *HasTypes_* instance is *pure*.

Two cases remain: *U*, when the field contains no value at all (isomorphic to unit), and *V*, which corresponds to types with no constructors. Both of these are just skipped.

```
instance GHasTypes U U a b where
  gtypes _ = pure
instance GHasTypes V V a b where
  gtypes _ = pure
```

All cases have now been covered, but care is required; the two overlapping instances for the *K* cases can lead to surprising results when the query is changing types. Consider the type *IntPair*:

```
data IntPair a = IntPair Int a
```

As expected, updating the *Ints* in an *IntPair Int* updates both the monomorphic value, and the field corresponding to the type variable:

```
> ghci> over types_ ((+10) :: Int  $\rightarrow$  Int) (IntPair 1 (2 :: Int)) :: IntPair Int
> IntPair 11 12
```

However, when a function that changes the types is mapped, the monomorphic *Int* is left alone:

```
> ghci> over types_ (show :: Int  $\rightarrow$  String) (IntPair 1 (2 :: Int)) :: IntPair String
> IntPair 1 "2"
```

While technically the correct behaviour, it is rather confusing. Therefore this combinator is restricted to only allow monomorphic updates.

```
type HasTypes s a = HasTypes_ s s a a
types ::  $\forall a s.$  HasTypes s a  $\Rightarrow$  Traversal s s a a
types = types_
```

Another thing to note is the abundance of explicit type annotations in the above examples. The promise of type inference is that the type of functions can be inferred from their use without providing type signatures.

However, here the type really is an *input* to the function, as it determines the nature of the traversal. The type abstraction $\forall a$ can be instantiated using visible type applications [Eisenberg et al. 2016] to provide hints to the compiler. This allows the much more concise form of application:

```
> ghci> over (types @Int) (+10) (IntPair 1 (2 :: Int))
> IntPair 11 12
```

Note that the order in which type variables are quantified for *types* has been carefully chosen: it is much more common to provide the targeted type than the structure's type.

5.1 More Efficient Traversals

Let us now evaluate our traversal so far. Consider the following simple datatype:

```
data T = MkT Int String (Maybe Bool)
```

Suppose that the task is to collect all the *Int* values in a given *T* into a list, as follows:

```
> ghci> toListOf (types @Int) (MkT 10 "a long string" (Just True))
> [10]
```

Clearly, the only interesting field is the first one. At runtime, there is no need to inspect either the *String* field or the *Maybe Bool* field. However, our naive implementation does inspect both. Worse, it traverses the whole string, character by character! In the *Maybe Bool* case, the compiler’s inliner comes to the rescue: by inlining the traversal’s definition sufficiently many times, it is able to tell that only the *pure* function is ever called, thus the whole field can be skipped.

The bigger problem is the string as the *String* type in Haskell is defined as a linked list of characters, the generated traversal is recursive. As such, there is no hope that the inliner could ever work out that this traversal is fruitless. For large types that contain many recursive subparts, the performance penalty is significant.

Interesting types. In order to avoid this penalty at runtime, the subparts that need to be traversed must be identified at compile-time; we call these subparts “interesting”. A type is interesting if it immediately contains the queried type or it contains other interesting types. Crucially, a mutually recursive group that does not contain the queried type need not be traversed at runtime – our predicate aims to filter out precisely these cases.

This section proceeds by defining the “interesting” predicate inductively on the type of the generic structure. To express this type-level computation, the *Interesting* closed type family is used. Closed type families [Eisenberg et al. 2014] comprise an ordered set of potentially overlapping type equations.

The first two arguments are the generic structure and the queried type. The third argument keeps a list of already seen types. This is to break loops in case of (mutually) recursive types.

```

type family Interesting (rep :: ★) (a :: ★) (seen :: [★]) :: Bool where
  Interesting (l :+: r)    t seen = Interesting l t seen ∨ Interesting r t seen
  Interesting (l :×: r)    t seen = Interesting l t seen ∨ Interesting r t seen
  Interesting (K t)       t seen = 'True
  Interesting (K Char)   _ _    = 'False
  ...
  Interesting (K Word)  _ _    = 'False
  Interesting (K r)     t seen = InterestingUnless (Elem r seen) (Rep r) t (r' : seen)
  Interesting (M _ f)   t seen = Interesting f t seen
  Interesting _         _ _    = 'False

type family InterestingUnless (s :: Bool) f (a :: ★) (seen :: [★]) :: Bool where
  InterestingUnless 'True _ _ _ = 'False
  InterestingUnless 'False f a seen = Interesting f a seen

```

In addition to overlapping equations, two properties of closed type families are used:

- (1) Pattern matching on members of the open type universe ★ (in the *rep* argument). This is not essential in our work since an inductive universe for the generic constructors could have been defined.
- (2) Non-linear patterns: the pattern *Interesting (K t) t seen* matches when the type of the field matches the query.

Most of the cases are self-explanatory. For *:+:* and *:×:* nodes, a requirement is that either branch is interesting. Fields are interesting when their type matches the query. Otherwise, for primitive types such as *Char* the search stops. The rest of the fields are inspected further, as they might contain the query. Their inspection is done by recursively invoking the *Interesting* predicate on their representation type. *Elem r seen* is a predicate that returns true when the field *r* is already in the *seen* set. This recursive branch could be written as

$$\text{Interesting } (K \ r) \ t \ \text{seen} = \text{If } (\text{Elem } r \ \text{seen}) \ \text{'False } (\text{Interesting } (\text{Rep } r) \ t \ (r' : \text{seen}))$$

However, type families are eagerly evaluated [Vytiniotis et al. 2011], so both branches of *If* are evaluated. This would be disastrous, as without the *seen* predicate, the recursive branch would diverge for mutually recursive groups. Instead, *InterestingUnless* is implemented to encode the conditional and ensure that the recursive case is only evaluated for fields not visited already.

The implementation of *types* is now refined by using *Interesting* to eliminate unnecessary runtime traversals. An auxiliary class, *HasTypesOpt*, is introduced, which is indexed by a boolean flag: whether to inspect its argument or not.

```
class HasTypesOpt (i :: Bool) s t a b where
  typesOpt :: Traversal s t a b
instance (Generic s, Generic t, GHasTypes (Rep s) (Rep t) a b)
  => HasTypesOpt 'True s t a b where
  typesOpt = isoRep · gtypes
instance HasTypesOpt 'False s s a b where
  typesOpt _ = pure
```

The default *HasTypes_* instance is revised by dispatching to the helper *HasTypesOpt* with the flag set to the *Interesting* predicate applied to the structure.

```
instance (Generic s, Generic t, HasTypesOpt (Interesting (Rep s) a [s]) s t a b)
  => HasTypes_ s t a b where
  types_ = typesOpt @(Interesting (Rep s) a [s])
```

Now depending on the result of computing *Interesting* for this type, the traversal will either stop or carry on recursively if more vales should be traversed. This saves unnecessary work from being performed as was the case in the example of traversing *T*.

The approach outlined above works for many polymorphic recursive data types, but has its limitations. Consider the type of perfectly balanced trees:

```
data Perfect a = Single a | Balanced (Perfect (a, a))
```

This can be considered to be divergent polymorphic recursion in the sense that the type changes at every level, and that no finite fix point can be found. To tackle such cases it is worth considering augmenting the predicate with a number that acts as a depth bound.

6 GENERIC TRAVERSALS WITH PARAMETERS

The implementation of the *param* traversal which focuses on all values corresponding to a type parameter is discussed in this section. The motivation is to derive a traversal which is able to change the type of the elements in a container. Recall that for *types* this was disallowed as the behaviour of the traversal is not easy to specify. On the other hand, *param* doesn't have this problem as only the necessary positions are changed in order to change the type of a parameter.

Implementing *param* poses a number of new challenges:

Locating parameters The traversal must distinguish between values that correspond to the queried type parameter, and values that were monomorphically defined (or correspond to other parameters). Even (and especially) when the query is not type-changing, it must not confuse a monomorphic *Int* with a type parameter instantiated to *Int*.

Multiple parameters It is important that *param* works for any number of type parameters, as opposed to the special case where there is just one.

Type inference When specifying a type-changing traversal, it is important to infer how the type of the structure will change.

Our solution to these problems uses only the *Generic* class and a set of type-level algorithms. Previous approaches [Magalhães et al. 2010] extended the generic representation to allow working with a single type parameter, but our approach is more flexible and uses existing machinery.

6.1 Locating Parameters

We first tackle the problems of locating parameters and dealing with multiple parameters simultaneously. The type family *Rep* in the *Generic* class takes as argument a type with kind \star . This means that it must be called on type constructors that are fully saturated.

Recall the *Invoice a* type from our biscuit factory. When in the normal queue case it is instantiated to *Invoice Int*, *Rep* is unable to distinguish between the *priority* and *number* fields by their types, as both are *Ints*, since $Rep (Invoice Int) \equiv K Item : \times : K String : \times : K Int : \times : K Int$. In order to implement *param* the two must be somehow distinct.

To solve this, each type parameter is tagged with a unique index, corresponding to the parameter's position in the original type. The position of each type parameter can then be tracked in the generic representation. The *Param* newtype wraps a value of type *a*. It is indexed by a type-level natural representing which parameter it corresponds to.

newtype $Param (i :: Nat) a = Param \{ unParam :: a \}$

Then, given any concrete instantiation of a type, its type parameters are iterated through and wrapped in *Param* constructors with increasing indices. The *Index* type family does that.

type family $Index (t :: k) (i :: Nat) :: k$ **where**
 $Index (t a) i = Index t (i + 1) (Param i a)$
 $Index t _ = t$

This conversion allows us to track the parameters:

$Index (Invoice Int) 0 \equiv Invoice (Param 0 Int)$
 $Index (Either Int String) 0 \equiv Either (Param 1 Int) (Param 0 String)$

Note that numbering starts at the last parameter, as it is the outermost one. With the new indexing in place, determining the origin of types in the generic representation is no longer a problem. If the type is wrapped in a *Param* constructor, it was a type parameter, otherwise it was an ordinary field.

$Rep (Index (Invoice Int) 0) \equiv K Item : \times : K String : \times : K Int : \times : K (Param 0 Int)$

Only one problem remains: the functions *to* and *from* operate on *Rep a*. $to :: Rep a \rightarrow a$ – how is this turned into $Rep (Index a 0) \rightarrow a$? $Rep (Index a 0)$ extends *Rep a* by wrapping certain fields in the *Param* newtype. Newtype wrappers have no runtime representation, which means that $Rep (Index a 0)$ and *Rep a* are representationally equal: they are the same at runtime. This means that they can be safely coerced [Breitner et al. 2014].

The $Generic_N$ class can be thought of as an extension of the *Generic* class, whose to_N and $from_N$ functions take care of the coercions, by requiring that *Rep a* and $Rep_N a$ are indeed coercible.

<p>class $(Coercible (Rep a) (Rep_N a),$ $Generic a) \Rightarrow Generic_N (a :: \star)$ where type $Rep_N a :: \star$ $to_N :: Rep_N a \rightarrow a$ $from_N :: a \rightarrow Rep_N a$</p>	<p>instance $(Coercible (Rep a) (Rep_N a),$ $Generic a) \Rightarrow Generic_N a$ where type $Rep_N a = Rep (Index a 0)$ $to_N = coerce (to @a)$ $from_N = coerce (from @a)$</p>
--	---

To reflect, the first two difficulties have been taken care of. By using the *Index* type family to label each of the parameter positions, it is possible to keep track of which fields arise from parameters and which ones do not. This also works for any number of parameters.

Once this information has been identified, notice that the problem of traversing the *i*th parameter can be formulated as a typed traversal of *Param i a* over the *Indexed* view of the structure. The implementation of the *param* traversal is thus:

```
instance (GenericN s, GenericN t, GHasTypes (RepN s) (RepN t) (Param i a) (Param i b))
  ⇒ HasParam i s t a b where
  param = isoRepN · gtypes · paramIso @i
```

Here, *paramIso* is the lens that focuses on the values by forgetting the *Param i* wrapper.

```
paramIso :: ∀ i a b. Lens (Param i a) (Param i b) a b
paramIso f p = Param ($) f (unParam p)
```

The function *iso_{Rep_N}* is analogous to *iso_{Rep}*, but for the new *Rep_N* representation.

```
isoRepN :: (GenericN s, GenericN t) ⇒ Lens s t (RepN s) (RepN t)
isoRepN f s = toN ($) f (fromN s)
```

Now we turn to the problem of improving type inference for our type changing traversals.

6.2 Type Inference

The four parameters of *Traversal s t a b* have interesting connections, and even from partial information the rest can be inferred. More formally, the *HasParam i s t a b* class is defined with a single function, *param :: Traversal s t a b*, which describes the traversal of the *i*th type parameter, *a*, of *s*. *t* is the result of changing the *i*th parameter of *s* to *b*.

```
class HasParam (i :: Nat) s t a b | i s b → t, i s → a, i t a → s, i t → b where
  param :: Traversal s t a b
```

In order to resolve which instance of *HasParam* to use, the types of all five type parameters must be known. The user is expected to provide *i* by using type applications but the other four can be inferred in different situations.

The four functional dependencies [Sulzmann et al. 2007] each specify what type information can be inferred if some of the types are known. They act as a specification as to what relationship must hold between the type variables in each instance.

i s b → t The source type and modification function are known. Then the target type *t* can be uniquely determined. This ensures that if *i* is provided as a type argument and the traversal is fully applied, then the result type can be inferred.

i s → a From a position and the source type only, the type of the parameter at that position can be uniquely determined.

i t a → s The result type and modification function are known. Then the source type *s* can be uniquely determined. This dependency helps traversals to be composed together where there would otherwise be ambiguous type variables in the middle of the composition.

i t → b From a position and the target type only, the type of the parameter at that position can be uniquely determined.

Without these functional dependencies it would be very difficult to use these optics without explicitly writing type signatures.

Now, the instance for *HasParam* which was defined above must be modified in order to implement the stated functional dependencies. The way that we prove to the compiler that each functional dependency holds is by defining a type family which witnesses each assertion.

There are two kind of dependencies: $i\ s \rightarrow a$ and $i\ t \rightarrow b$ both *get* the parameter at index i , while $i\ s\ b \rightarrow t$ and $i\ t\ a \rightarrow s$ both *set* it. In order to assign an operational meaning to these functional dependencies, two type families are defined that express the getting and the setting relations respectively. First, *GetParam* peels off the parameters of its argument one by one until it reaches the i th.

```
type family GetParam (t :: k) (i :: Nat) :: ★ where
  GetParam (t a) 0 = a
  GetParam (t _) i = GetParam t (i - 1)
```

Similarly, *PutParam* digs into its argument to find and the i th parameter.

```
type family PutParam (t :: k) (i :: Nat) (b :: ★) :: k where
  PutParam (t _) 0 b = t b
  PutParam (t a) i b = (PutParam t (i - 1) b) a
```

Notice that both *GetParam* and *PutParam* operate on poly-kinded arguments, even though we only intend to call them on types of kind ★. However, as both functions peel off the arguments, intermediate recursive calls operate on higher-kinded types.

$$\text{GetParam (Either Int String :: ★) 1} \equiv \text{GetParam (Either Int :: ★} \rightarrow \text{★) 0} \equiv \text{Int}$$

GetParam and *PutParam* highlight another important aspect of closed type families: they are not parametric, as arguments that have polymorphic kinds can be matched on. They can also decompose application forms, as in the $t\ a$ pattern [Weirich et al. 2011].

Now we have a method of proving these dependencies. The proofs are supplied as instance constraints which allow the compiler to conclude the validity of the functional dependencies. This leaves the final definition for *param*.

```
instance (a ~ GetParam s i, b ~ GetParam t i, t ~ PutParam s i b, s ~ PutParam t i a,
  GenericN s, GenericN t, GHasTypes (RepN s) (RepN t) (Param i a) (Param i b))
  => HasParam i s t a b where
  param = isoRepN · gtypes · paramIso @i
```

As an example, consider the *Poly a b* type, which is a list that alternates between elements of type a and type b (note the polymorphic recursion in the tail).

```
data Poly a b = PNil | PCons a (Poly b a)
```

With *param*, a traversal can be specified that updates the *Strings* that correspond to the a parameter:

```
> ghci> over (param@1) length (PCons "wafer" (PCons "oreo" (PCons "nice" PNil)))
> PCons 5 (PCons "oreo" (PCons 4 PNil))
```

7 GENERIC TRAVERSALS WITH CLASS

Inspecting the inductive definition of *GHasTypes s t a b* in Section 5, notice that all the inductive cases do is merely “forward the focus” to their children. The first time any decision is made is at K : whether to stop, or keep going via the mutually recursive *HasTypes_* class. Defining traversals that employ a different operation on fields would require writing a very similar inductive definition for each traversal, only differing at the last case: the fields.

Instead, we define an *extensible* generic traversal that is parameterised over a type class [Boling-broke 2011] that provides the action applied to the fields. The requirement is that all fields have an instance of this class, giving the name *constrained traversal*.

```
type TraversalC (c :: ★ → ★ → Constraint) s t
  = ∀f. Applicative f ⇒ (∀ a b. c a b ⇒ a → f b) → s → f t
```

Instead of specifying up-front the type of the focus a , we say that we target every field of every type, as expressed by the *rank-2* [Peyton Jones et al. 2007] quantification of the variables a and b in the first argument. Different instantiations of c can relate a and b in different ways. *HasConstraints* classifies types that can be traversed in this way, and *GHasConstraints* provides a concrete definition by induction over the generic structure.

```
class HasConstraints (c :: ★ → ★ → Constraint) s t where
  constraints :: TraversalC c s t

class GHasConstraints (c :: ★ → ★ → Constraint) s t where
  gconstraints :: TraversalC c s t
```

The nodes \times , $+$, U , V and M are treated analogously to *HasTypes₋*. We target our focus at the values, as specified by the action c .

```
instance c a b ⇒ GHasConstraints c (K a) (K b) where
  gconstraints = isoK
```

Here, iso_K is instantiated to the constrained traversal

```
isoK :: c a b ⇒ TraversalC c a b
isoK :: ∀f. (Functor f, c a b) ⇒ (∀ a1 b1. c a1 b1 ⇒ a1 → f b1) → K a → f (K b)
```

Since its function argument can be applied to *any* a_1 and b_1 , it is certainly applicable to a and b (as the $c a b$ instance is given).

To show that this traversal is indeed the most general, we allude briefly to an implementation of *HasTypes₋* in terms of *HasConstraints*. Note that compared to *HasTypes₋ s t a b*, the type parameters of *HasTypesC a b s t* are swapped. This is because the traversal will be constrained by *HasTypesC a b* – intuitively, we require that each field be traversable with an $a \rightarrow b$ action.

```
class HasTypesC a b s t where
  typesC :: Traversal s t a b
```

The decision at the leaf nodes can be encoded via two corresponding instances. The first instance describes what to do when the target of the focus is a .

```
instance HasTypesC a b a b where
  typesC f s = f s
```

Here note that this instance allows the field transformation to select the queried types

```
isoK :: (∀ a1 b1. HasTypes- a b a1 b1 ⇒ a1 → f b1) → K a → f (K b)
```

When $a \sim a_1$ and $b \sim b_1$, instance resolution picks the above instance, applying the transformation. Otherwise, the more general instance is selected, which guides the recursion:

```
instance (Generic s, Generic t, HasConstraints (HasTypesC a b) s t)
  ⇒ HasTypesC a b s t where
  typesC f = constraints @(HasTypesC a b) (typesC f)
```

We omit here the definition for primitives, which can be defined analogously to *HasTypes₋*.

If *HasConstraints* is indeed the most general traversal, then why not use it to define *HasTypes_?* The answer is of a practical nature: the additional burden on the constraint solver slows down compilation times, and the optimiser misses inlining and specialisation opportunities more easily.

8 PERFORMANCE

When working generically we must always ask whether the abstraction comes at the cost of performance. In this case, it is pleasing that our use of generics is optimised away by the compiler. There are four crucial reasons why we can be confident that GHC will produce efficient code.

Evidence generation. By using a type-directed approach, the call hierarchy is known at compile time and this information can be used to unroll the definitions. This unrolling is achieved during *evidence generation*.

Specialisation. Functions using our methods will have constrained types but this overhead is eliminated via *specialisation*.

Inlining. We define our operations such that the composition operator is not recursive and can hence be readily *inlined*.

Internal representation. Finally, we choose an internal representation of our optics such that they expose the optimisation opportunities to the compiler.

In this section we describe the optimisations which we rely on to produce efficient code. We explain each of these techniques in turn. Our running example in this section is the *incList* function which maps over a list of trees and increments the *Ints* inside the tree.

```
data Tree a = Leaf a | Branch (Tree a) (Tree a)
incList :: [Tree Int] → [Tree Int]
incList [] = []
incList (x : xs) = over (types @Int) (+1) x : incList xs
```

8.1 Evidence Generation

During compilation, type class constraints are desugared into arguments to the function [Wadler and Blott 1989]. The argument is known as a dictionary and contains a field for each method of a type class. Type class methods are then desugared as lookup functions into this dictionary.

The definition of *incList* uses *types* so the constraint solver must generate an evidence of *HasTypes (Tree Int) Int*, it does so by creating an appropriate dictionary.

The instance for *HasTypes s a* has constraints *Generic s*, and *GHasTypes (Rep s) a*. We focus on *HasTypes* and *GHasTypes*, treating the dictionary for *Generic (Tree Int)* implicitly. Thus the produced dictionaries are *HasTypesDict* and *GHasTypesDict*, corresponding to the appropriate classes.

```
data HasTypesDict s a = HasTypesDict { types :: Traversal s s a }
data GHasTypesDict s a = GHasTypesDict { gtypes :: Traversal s s a }
```

The necessary evidence generated for *GHasTypes (Rep (Tree Int)) Int* comes by providing the dictionaries for this type. The simplified representation for *Tree Int* without metadata nodes is:

$$\text{Rep (Tree Int)} \equiv K \text{ Int} \text{ :+} : (K \text{ (Tree Int)} \text{ :}\times : K \text{ (Tree Int)})$$

By working through this structure methodically, we arrive at the following dictionary definitions:

```
hasTypesDictTreeInt :: HasTypesDict (Tree Int) Int
hasTypesDictTreeInt = HasTypesDict { types = isoRep · gtypes ghasTypesDictTreeInt }
ghasTypesDictTreeInt :: GHasTypesDict (Rep (Tree Int)) Int
```

$$\begin{aligned}
ghasTypesDict_{TreeInt} &= GHasTypesDict \{ gtypes = \lambda f \ l r1 \rightarrow \mathbf{case} \ l r1 \ \mathbf{of} \\
&\quad L \ l \rightarrow L \ \langle \$ \rangle \ gtypes \ ghasTypesDict_{KInt} \ f \ l \\
&\quad R \ r \rightarrow R \ \langle \$ \rangle \ gtypes \ ghasTypesDict_{:\times} \ f \ r \} \\
ghasTypesDict_{KInt} &:: GHasTypesDict \ (K \ Int) \ Int \\
ghasTypesDict_{KInt} &= GHasTypesDict \ { gtypes = iso_K \} \\
ghasTypesDict_{KTreeInt} &:: GHasTypesDict \ (K \ (Tree \ Int)) \ Int \\
ghasTypesDict_{KTreeInt} &= GHasTypesDict \ { gtypes = iso_K \cdot types \ ghasTypesDict_{TreeInt} \} \\
ghasTypesDict_{:\times} &:: GHasTypesDict \ (K \ (Tree \ Int) \ :\times: \ K \ (Tree \ Int)) \ Int \\
ghasTypesDict_{:\times} &= GHasTypesDict \ { gtypes = \lambda f \ (l \ :\times: \ r) \rightarrow \ (: \times: \) \ \langle \$ \rangle \\
&\quad gtypes \ ghasTypesDict_{KTreeInt} \ f \ l \ \langle * \rangle \ gtypes \ ghasTypesDict_{KTreeInt} \ f \ r \}
\end{aligned}$$

Evidence is first generated by using the instance for $+$, before recursing into both branches and finding evidence for $:\times:$ and the $K \ Int$ nodes. As such, we have a dictionary for each type constructor. The constraint solver will terminate as it will observe that the $ghasTypesDict_{TreeInt}$ dictionary can be used when trying to solve the recursive case. Thus, these dictionaries form a mutually recursive group. The dictionaries generated are straightforward transcriptions of the instances, with instance constraints solved and β -reduced. The definition of $ghasTypesDict_{TreeInt}$ is still not as efficient as it could be, and we discuss how it can be further improved with inlining in Section 8.2.1.

We see that the process of generating evidence also unrolls definitions. If we had instead defined $types$ as a function over a normal data type without any type direction, it would be self-recursive and hence not able to be eliminated in the same manner. This process is safe as $types$ are finite and statically known at compile time. Without additional language pragmas, the restrictions on instance contexts guarantee that the constraint solving process terminates.

8.2 Inlining

Once the structure is in place, there is still indirection present which can be removed. The first step of doing this is inlining. Inlining is the process of replacing a function's name by its definition. It is the most crucial optimisation in the compiler's pipeline as it enables all other optimisations to occur. We already saw how the compiler generates naive verbose code which is simplified when inlined. This is in general true for all programs.

However, whilst always safe in a pure language like Haskell, we must still be careful about when we inline. If we inline too little then we miss optimisation opportunities. If we inline too much then the size of our program becomes very large and takes a long time to compile.

The compiler contains a set of balanced heuristics to decide whether to inline a definition [Peyton Jones and Marlow 2002]. These include factors such as: the syntactic size of a function, as a measure to stop a lot of code duplication; whether a function is recursive, recursive functions are never inlined; whether a function is applied to known arguments, there is a good chance that the body will scrutinise the arguments and perform more simplification and so on.

There are also manners in which the user can influence these automatic decisions. One in particular is the use of `INLINE` pragmas which can be used to mark definitions as very desirable to `INLINE`. In our use cases, marking some instance methods as `INLINE` was necessary to unstick the optimiser and enable it to perform much more simplification.

In addition, the optimiser will also evaluate programs by β -reducing, evaluating case expressions with a known scrutinee and perform commuting conversions. For a full account of the simple core transformations which the simplifier performs in order to generate simpler code, one should consult [Peyton Jones and Santos 1998].

8.2.1 *Optimising Dictionaries.* We recall that our generated dictionaries are mutually recursive. This isn't surprising, as we expect *gtypes* to be recursive in general if we are trying to traverse a recursive data structure. Mutually recursive blocks of functions must be treated with care, as repeatedly inlining them causes the inliner to diverge. Each mutually recursive group is thus appointed a loop-breaker function, which is never inlined, but other definitions can be freely inlined into each other in order to create a single self-recursive definition. After the dictionaries are inlined into each other, we end up with the following evidence which has the correct unrolled shape we were looking for.

$$\begin{aligned}
ghasTypesDict'_{TreeInt} &:: GHasTypesDict (Rep (Tree Int)) Int \\
ghasTypesDict'_{TreeInt} &= GHasTypesDict \{ gtypes = \lambda f \{ l r1 \rightarrow \mathbf{case} \{ l r1 \} \mathbf{of} \\
&\quad L \ l \rightarrow L \langle \$ \rangle \ iso_K \ f \ l \\
&\quad R \ b \rightarrow R \langle \$ \rangle (\lambda f \ (l : \times : r) \rightarrow (: \times :)) \\
&\quad \quad \langle \$ \rangle (\iso_K \cdot \iso_{Rep} \cdot gtypes \ ghasTypesDict'_{TreeInt}) \ f \ l \\
&\quad \quad \langle * \rangle (\iso_K \cdot \iso_{Rep} \cdot gtypes \ ghasTypesDict'_{TreeInt}) \ f \ r \} \ f \ b \}
\end{aligned}$$

In this case, $ghasTypesDict'_{TreeInt}$ acts as the loop-breaker.

8.3 Specialisation

As we have seen, the evidence generation procedure and inlining are sufficient on their own to eliminate much of the generic overhead of a statically known parameter as long as we call the class method directly. However, we use class methods inside bigger functions and when we do they give rise to class constraints. When these larger functions are called, the dictionary must be solved and the required evidenced passed to the function.

For instance, we might want to write the more general type signature for *inclList* to be parametric over the choice of data structure contained in the list as long as it contains integers. We will call this generalised version *inclListGen*. If we call *inclListGen* and instantiate *s* to be *Tree Int* then we should expect that the definition would be identical to *inclList*.

$$\begin{aligned}
inclListGen &:: HasTypes \ s \ Int \Rightarrow [s] \rightarrow [s] \\
inclListGen \ [] &= [] \\
inclListGen \ (x : xs) &= \mathbf{over} \ (types @ Int) \ (+1) \ x : inclListGen \ xs
\end{aligned}$$

This problem is not trivial. When *inclListGen* is called, the evidence witnessing the constraint *HasTypes* will be passed to it. In order to eliminate this dictionary, it needs to be pushed inwards to the call of *types*. Since *inclListGen* is recursive, it cannot be inlined. Instead, we rely on *specialisation*.

The specialiser looks for calls to overloaded functions called at a known type. It then creates a new type-specialised definition which does not take a dictionary argument and a rewrite rule which rewrites the old version to the new version.

Suppose that we know that the value of *s* is *Tree Int*, and that the evidence dictionary for *HasTypes* is called *treeIntHasTypes*. The naive desugaring of calling $inclListGen @ (Tree Int) \ xs$ is:

$$inclListGen \ treeIntHasTypes \ xs$$

The specialiser then observes this call to *inclListGen* takes a dictionary argument and creates a specialised version $inclListGen_{TreeInt}$ with the following definition:

$$\begin{aligned}
inclListGen_{TreeInt} &:: [Tree Int] \rightarrow [Tree Int] \\
inclListGen_{TreeInt} \ xs &= (\lambda hasTypesDict \ xs \rightarrow \mathbf{case} \ xs \ \mathbf{of} \\
&\quad [] \rightarrow []
\end{aligned}$$

$$(x : xs) \rightarrow \text{over } (\text{types } \text{hasTypesDict}) (+1) x : \text{inclListGen } \text{hasTypesDict } xs \\ \text{treeIntHasTypes } xs$$

The right-hand side of the definition is the same as the right-hand side of *inclListGen* applied to *treeIntHasTypes*. Then, an additional rewrite rule is generated which replaces the overloaded call with the specialised definition.

$$\{-\# \text{ RULES "specinclListGen" forall } xs . \text{inclListGen } \text{treeIntHasTypes } xs \\ = \text{inclListGen_TreeInt } \#\}$$

That's the whole process. After β -reduction, the dictionary selector *types* is now adjacent to its dictionary and hence *types* can be inlined, and the correct method from *treeIntHasTypes* can be selected. Notice that in the definition of *inclListGen_{TreeInt}* there is still an overloaded call to *inclListGen*, this will be rewritten when the rewrite rule is applied and then *inclListGen_{TreeInt}* will become self-recursive. After these two steps, all occurrences of *treeIntHasTypes* and thus the overloading overhead is eliminated.

Once again, specialisation is an *enabling* transformation. Later optimisation passes will perform more complicated rearranging with the express goal of improving our code.

8.4 Internal Representation

After this unrolled pipeline of functions is created, the question remains how this can become the same as hand-written definitions later in the compilation process. How precisely do inlining and β -reduction lead to good code? How and why depends on the internal representation of lenses and traversals we choose in the library.

8.4.1 Lenses. In the case of lenses, the inliner does a sufficient job of combining the composition of lenses into a single lens without further intervention. The lens composition operator is not recursive and hence is readily inlined which leads to much further simplification.

$$\mathbf{data} \text{ Lens}_1 s t a b = \text{Lens}_1 (s \rightarrow a) (b \rightarrow s \rightarrow t) \\ (\circ) :: \text{Lens}_1 s t c d \rightarrow \text{Lens}_1 c d a b \rightarrow \text{Lens}_1 s t a b \\ (\text{Lens}_1 \text{ get}_1 \text{ set}_1) \circ (\text{Lens}_1 \text{ get}_2 \text{ set}_2) = \text{Lens}_1 (\text{get}_2 \cdot \text{get}_1) \\ (\lambda b s \rightarrow \text{set}_1 (\text{set}_2 b (\text{get}_1 s)) s)$$

In fact, the naive encoding given above for lenses does not produce the best results. Whilst it does collapse a sequence of compositions appropriately, the type ensures that in order to implement a modification operation, we must perform a *get* followed by a *set* and hence deconstruct *s* twice. We can get around this problem by using the existential encoding which means that we can directly implement an updating function by only deconstructing the source once.

$$\mathbf{data} \text{ Lens}_2 s t a b = \forall c. \text{Lens}_2 (s \rightarrow (a, c)) ((b, c) \rightarrow t)$$

Intuitively, the *get* function separates *s* into the the part we are focusing on of type *a* and its complement *c*. In turn, the *set* function recombines a value of type *b* with the complement.

$$(\bullet) :: \text{Lens}_2 s t c d \rightarrow \text{Lens}_2 c d a b \rightarrow \text{Lens}_2 s t a b \\ (\text{Lens}_2 \text{ get}_1 \text{ set}_1) \bullet (\text{Lens}_2 \text{ get}_2 \text{ set}_2) = \text{Lens}_2 \text{ get } \text{set } \mathbf{where} \\ \text{get } s = \mathbf{let} (c, \text{com}_1) = \text{get}_1 s; (a, \text{com}_2) = \text{get}_2 c \mathbf{in} (a, (\text{com}_1, \text{com}_2)) \\ \text{set } (b, (\text{com}_1, \text{com}_2)) = \text{set}_1 ((\text{set}_2 (b, \text{com}_2)), \text{com}_1) \\ \text{modify} :: \text{Lens}_2 s t a b \rightarrow (a \rightarrow b) \rightarrow (s \rightarrow t) \\ \text{modify } (\text{Lens}_2 \text{ get } \text{set}) f s = \mathbf{let} (a, c) = \text{get } s \mathbf{in} \text{set } ((f a), c)$$

Using this definition, chained modifications can be fused into a single function. Thus, in our implementation, the lenses we use are of this latter encoding. Once we have fused them together, we turn them into whichever encoding that we want the library to produce. By default, it is the van Laarhoven encoding as found in the lens library [Kmett 2018].

8.4.2 Traversals. Optimising traversals in the same manner is slightly trickier as we must find an encoding of a traversal which does not require a recursive composition operator. In order to do this, we use a van Laarhoven style representation. The composition operator for these traversals is the function composition operator. However, this is not sufficient, the downside of using this composition operator is that it does not perform normalisation as happened with lenses. It is necessary to appeal to the *Applicative* laws in order to rearrange and normalise these compositions. The following technique is due to Eric Mertens and can be found implemented in the lens library.

A van Laarhoven *Traversal* is a function with the following type.

$$\mathbf{type} \text{ Traversal } s \ t \ a \ b = \forall g. \text{Applicative } g \Rightarrow (a \rightarrow g \ b) \rightarrow s \rightarrow g \ t$$

The result type of these functions is a value constructed using *Applicative* operators. *Applicative* expressions have a normal form of a single *pure* followed by a sequence of left-associated applications using the combinator $\langle * \rangle$ [McBride and Paterson 2008]. In order to rewrite this normal form, we must re-associate all uses of $\langle * \rangle$ and then fuse together all uses of *pure*.

This first step is achieved by instantiating g to be *Curried*.

$$\mathbf{data} \text{ Curried } f \ a = \text{Curried } \{ \text{runCurried} :: \forall r. f \ (a \rightarrow r) \rightarrow f \ r \}$$

$$\mathbf{instance} \text{ Functor } f \Rightarrow \text{Functor } (\text{Curried } f) \ \mathbf{where}$$

$$\text{fmap } f \ (\text{Curried } v) = \text{Curried } (\lambda \text{far} \rightarrow v \ (\text{fmap } (\cdot) \ \text{far}))$$

$$\mathbf{instance} \text{ Functor } f \Rightarrow \text{Applicative } (\text{Curried } f) \ \mathbf{where}$$

$$\text{pure } a = \text{Curried } (\lambda \text{far} \rightarrow \text{fmap } (\$a) \ \text{far})$$

$$\text{Curried } m \ \langle * \rangle \ \text{Curried } n = \text{Curried } (m \cdot m \cdot \text{fmap } (\cdot))$$

It is the definition of $\langle * \rangle$ which performs the reassociation. Notice that the *Applicative* instance for *Curried* delegates all calls to *pure* to the underlying functor. We will fuse those together with an additional layer termed *Yoneda* which intercepts all the calls to *fmap* and fuses them together.

$$\mathbf{data} \text{ Yoneda } f \ a = \text{Yoneda } \{ \text{runYoneda} :: \forall r. (a \rightarrow r) \rightarrow f \ r \}$$

$$\mathbf{instance} \text{ Functor } (\text{Yoneda } f) \ \mathbf{where}$$

$$\text{fmap } f \ (\text{Yoneda } v) = \text{Yoneda } (\lambda k \rightarrow v \ (k \cdot f))$$

$$\mathbf{instance} \text{ Applicative } f \Rightarrow \text{Applicative } (\text{Yoneda } f) \ \mathbf{where}$$

$$\text{pure } a = \text{Yoneda } (\lambda f \rightarrow \text{pure } (f \ a))$$

$$\text{Yoneda } m \ \langle * \rangle \ \text{Yoneda } n = \text{Yoneda } (\lambda f \rightarrow m \ (f \cdot) \ \langle * \rangle \ n \ \text{id})$$

This time, we notice that *Yoneda* just delegates the definitions of the *Applicative*. Putting this together, we instantiate g to be *Curried (Yoneda g)* and then use *lowerCurriedYoneda* in order to return to a simple type parameterised by an *Applicative* constraint.

$$\text{liftCurriedYoneda} :: \text{Applicative } g \Rightarrow g \ a \rightarrow \text{Curried } (\text{Yoneda } g) \ a$$

$$\text{lowerCurriedYoneda} :: \text{Applicative } g \Rightarrow \text{Curried } (\text{Yoneda } g) \ a \rightarrow g \ a$$

This process performs the reassociating and fusion that we desired. However, in practice, it is difficult to be sure that the compiler will remove this overhead. On the other hand, it does not affect performance in common use cases such as modifying or summarising. This is because when

g is instantiated to a known *Applicative*, the non-recursive *Applicative* methods can be inlined. We usually instantiate to g to either *Const* or *Identity* which are completely eliminated.

Using similar techniques to traversals, van Laarhoven or profunctor representation of lenses and prisms could also be optimised, but these simple minded techniques are the most reliable and very effective in generating good programs without impacting compile times significantly.

8.5 Inspection Testing

In reality, compiler optimisations such as inlining are guided by heuristics, thus the exact shape of the optimised program is difficult to predict. To aid in reasoning about program transformations, [Breitner \[2018\]](#) developed a technique called *inspection testing*. Properties of the generated code can be verified by inspecting the intermediate representation after optimisations and comparing it against the hand-optimised version. The test suite of `generic-lens` uses `inspection-testing` to guarantee that the abstraction overhead is eliminated by the optimiser.

9 BENCHMARKS

We compare the performance of `generic-lens` with hand-written code as well as five other generic programming libraries which derive traversals for data types. There are no other libraries which derive lenses or prisms in a similar way so we could not compare this aspect of the library.

- (gl) `generic-lens` The library which we describe in this paper.
- (gp) `geniplate-0.7.6` [[Augustsson 2018](#)] A library which provides a similar interface to the `uniplate` library below but uses Template Haskell in order to generate traversals.
- (up) `uniplate-1.6.12` [[Mitchell and Runciman 2007](#)] A library which provides an interface for traversing data. A traversal for a data type which has a `Data` instance is derived by using `Data.Generics.Uniplate.Data`.
- (lens) `lens-4.16` [[Kmett 2018](#)] This library provides a reimplementaion of the `uniplate` interface to generate van Laarhoven style traversals rather than `uniplate` traversals.
- (syb) `syb-0.7` [[Lämmel and Peyton Jones 2003](#)] One of the first generics libraries using the `Data` type class to dynamically decide which nodes to traverse.
- (ol) `one-liner-1.0` [[Visscher and Xia 2018](#)] A library implementing profunctor style generic traversals using generics in a similar style to `generic-lens`. It generates the most general constrained traversal which we instantiate suitably to turn it into an ordinary traversal.

We implement a collection of benchmarks which modify, update, and summarise data types of three different sizes. `Tree` is a simple data type representing binary trees with 2 constructors. `Logic` is a deep embedding of propositional logic with 6 constructors. `HsModule` is a large data type representing a Haskell syntax tree with many constructors.

Our benchmarks have been compiled with `-fexpose-all-unfoldings` and enable a later specialisation pass. The former ensures more predictable cross-module inlining. The latter is a more aggressive change which in particular helps the `HsMod` benchmark by performing a specialisation pass towards the end of the compilation.

We show three results in [Figure 2](#). The y-axis is a log scale where we normalise against the hand-written code. The number above each column indicates the time relative to the hand-written code. For example, a value of 2 indicates that the benchmark took twice as much time as the hand-written definition. We include one modification benchmark for each different data type to indicate the relative performance of the libraries.

One should notice how the performance of our `generic-lens` library, labelled (gl), is comparable to the hand written examples and `geniplate` (gp) which uses Template Haskell to analyse the data definitions and to produce the optimal code. On the other hand, SYB (syb) is consistently very slow.

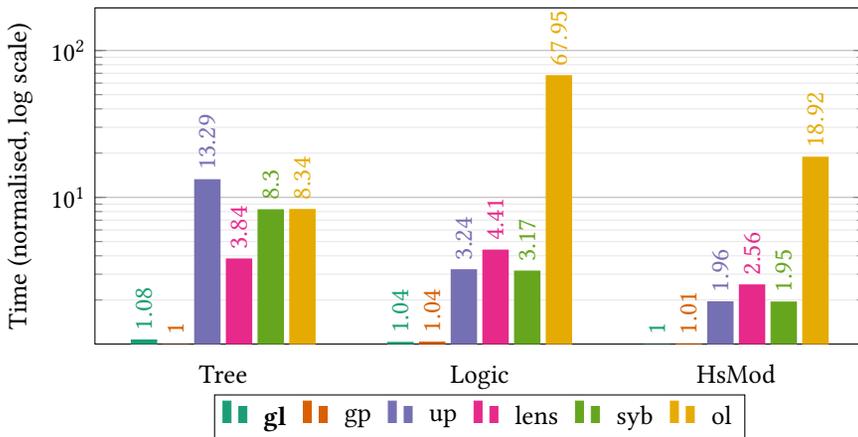


Fig. 2. Benchmarks showing execution time normalised as a factor of hand-written code

one-liner (ol) performs an order of magnitude worse than SYB in these benchmarks. Experiments at higher inlining thresholds indicate that the performance can be comparable with our library.

We also have a large suite of other benchmarks which we have used to validate our approach. These include effectful traversals, summarising and traversing dense and sparse structures. We observe that in all these cases, the performance is very close to the hand-written definitions.

10 RELATED WORK

The scrap your boilerplate work pioneered the generation of code suitable for generic traversals [Lämmel and Peyton Jones 2003; Lämmel and Peyton Jones 2004]. Our work can be seen as the latest refinement in a long line of techniques. We believe that we are the first to provide an efficient embedding in Haskell using existing language features which does not rely on Template Haskell [Adams and DuBuisson 2012]. Furthermore, the type-changing variants which are enabled by our use of generics is novel. Magalhães [2014] considered multiple parameters, but has a different solution to ours and is less flexible.

Our method is to use the generic structure of data, which was prominently investigated by Hinze [2000], and has led to an implementation in GHC that allows generic type classes to be derived [Magalhães et al. 2010]. *GHC.Generics* was intended to be a low-level way for library authors to implement generic programming libraries. Later de Vries and Löh [2014] implemented *generics-sop*, a high-level interface. In theory, we could have based our implementation on *generics-sop* as it is also shallow and non-recursive like *GHC.Generics* so we would expect similar performance characteristics. However, many of the library functions are currently implemented using recursion. In principle, these could be rewritten in terms of type classes so that any "recursive" calls are type-directed in the style of our definitions. In private communication, the authors have indicated that they are intending to rework *generics-sop* in this manner. It would be worthwhile to attempt to implement our interface using *generics-sop* once their work is completed. Visscher and Xia [2018] have implemented the low-level one-liner library, which provides constrained profunctor traversals. As we saw in the benchmarks, we predict their performance could be comparable to ours with suitable compiler hints. In our interface we already provide several different traversal schemes rather than just the constrained variant. We also consider problems such as removing redundant traversals and type inference.

Lenses have quickly been adopted into the Haskell ecosystem since [van Laarhoven 2009] and have been described from first principles in [Pickering et al. 2017]. There are many different libraries implementing lenses in the Haskell ecosystem but by far the most prominent and well used is `lens` [Kmett 2018]. The interface of `generic-lens` is compatible with that of `lens`. The `lens` library can also derive lenses using Template Haskell. We find this unsatisfactory as users must decide up-front which lenses they want for their data types. As a result, users usually invent named fields with a specific naming strategy and derive lenses with names based from them. They can't derive the other queries we specify, in part because defining all such lenses up front would incur namespace pollution. It is more desirable to have the flexibility we provide in order to specify precisely at use-sites the mode of inspection.

Yallop [2017] discusses how structured multi-staged programming techniques can be used to improve SYB. The work observes the benefits which can be gained from a simple binding time analysis. The insight is that when applied, the types a traversal must deal with are already known. Once the target of the traversal is known, a specialised version for that specific structure is created. Further, if the traversal function is also statically known, it can be partially evaluated at each node to eliminate the dynamic type checks. The constrained traversal can be seen as an implementation of the same idea in the current work. By statically knowing the type, we generate a specialised traversal which targets all fields but has the right unrolled structure. By instantiating the constraint (which we do statically) we can then select which parts of the data type to target.

They further develop this approach in order to analyse additional local transformations which can be applied in order to improve the generated code. In section 8 we explained how compiler optimisations performed by GHC amount to achieving the same thing. As one particular example, they observe the need for `let`-insertion to deal with recursive definitions. Their solution is to use a fixpoint combinator which supports memoisation. They then insert `let` bindings for each recursive call. For us, this memoisation is performed in the same manner by the constraint solver.

Their work is valuable as it carefully analyses the optimisations needed to create optimal code. Our work is complementary as we observe that with an automatic partial evaluator (GHC's optimiser), it is sufficient to instruct a simple unrolling before passing the code to be automatically optimised once it is no longer recursive. Like ours, Yallop's implementation also hinges on an observation that implicit information can be treated as static, allowing much of the overhead to be eliminated. [Magalhães 2012] also considered the behaviour of the optimiser on similar generic programs.

Adams et al. [2014] analysed the poor performance of SYB and performed code optimisations expressed in HERMIT [Farmer 2015] to get better performance, in their extended work they show how this can be done with an adapted version of GHC [Adams et al. 2015]. A large problem with their approach is that they rely on being able to symbolically evaluate type casts. The implementation of *Typeable* [Peyton Jones et al. 2016] implements type equality by comparing fingerprints using primitive operations. There is no hope for a compiler to evaluate these operations without additional guidance. Our solution does not generate any domain specific idiosyncrasies which we need to modify the optimiser in order to eliminate.

11 CONCLUSION

This paper has presented `generic-lens`, a library which has advanced the state-of-the-art for traversing and manipulating data structures in two key ways: first, the library introduces a suite of new traversals that can be easily combined with lenses for extremely flexible and powerful queries. Second, the code that is generated competes with hand-optimised definitions. By embracing the principled approach of generic programming, our library offers these benefits while remaining robust and possible. For these reasons, we believe that `generic-lens` promises to be a successful library that will be established as an essential tool for the working functional programmer.

REFERENCES

- Michael D Adams and Thomas M DuBuisson. 2012. Template your boilerplate: Using Template Haskell for efficient generic programming. In *ACM SIGPLAN Notices*, Vol. 47. ACM, 13–24.
- Michael D. Adams, Andrew Farmer, and José Pedro Magalhães. 2014. Optimizing SYB is Easy!. In *Proceedings of the ACM SIGPLAN 2014 Workshop on Partial Evaluation and Program Manipulation (PEPM '14)*. ACM, New York, NY, USA, 71–82. <https://doi.org/10.1145/2543728.2543730>
- Michael D. Adams, Andrew Farmer, and José Pedro Magalhães. 2015. Optimizing SYB Traversals is Easy! *Sci. Comput. Program.* 112, P2 (Nov. 2015), 170–193. <https://doi.org/10.1016/j.scico.2015.09.003>
- Lennart Augustsson. 2018. geniplate-mirror-0.7.6 library. <http://hackage.haskell.org/package/geniplate-mirror-0.7.6>
- Richard Bird, Jeremy Gibbons, Stefan Mehner, Janis Voigtlaender, and Tom Schrijvers. 2013. Understanding Idiomatic Traversals Backwards and Forwards. In *Haskell Symposium*. <http://www.comlab.ox.ac.uk/jeremy.gibbons/publications/uitbaf.pdf>
- Maximilian C. Bolingbroke. 2011. Constraint Kinds for GHC. <http://blog.omega-prime.co.uk/2011/09/10/constraint-kinds-for-ghc/>
- Joachim Breitner. 2018. Beyond correct and fast: Inspection Testing. *CoRR* abs/1803.07130 (2018). arXiv:1803.07130 <http://arxiv.org/abs/1803.07130>
- Joachim Breitner, Richard A. Eisenberg, Simon Peyton Jones, and Stephanie Weirich. 2014. Safe Zero-cost Coercions for Haskell. *SIGPLAN Not.* 49, 9 (Aug. 2014), 189–202. <https://doi.org/10.1145/2692915.2628141>
- Manuel M. T. Chakravarty, Gabriele Keller, Simon Peyton Jones, and Simon Marlow. 2005. Associated types with class. In *POPL '05: Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM Press, 1–13. <https://doi.org/10.1145/1040305.1040306>
- Edsko de Vries and Andres Löb. 2014. True sums of products. In *Proceedings of the 10th ACM SIGPLAN workshop on Generic programming*. ACM, 83–94.
- Richard A. Eisenberg, Dimitrios Vytiniotis, Simon Peyton Jones, and Stephanie Weirich. 2014. Closed Type Families with Overlapping Equations. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)*. ACM, New York, NY, USA, 671–683. <https://doi.org/10.1145/2535838.2535856>
- Richard A. Eisenberg, Stephanie Weirich, and Hamidhasan G. Ahmed. 2016. Visible Type Application. In *Proceedings of the 25th European Symposium on Programming Languages and Systems - Volume 9632*. Springer-Verlag New York, Inc., New York, NY, USA, 229–254. https://doi.org/10.1007/978-3-662-49498-1_10
- Andrew Farmer. 2015. *HERMIT: Mechanized Reasoning during Compilation in the Glasgow Haskell Compiler*. Ph.D. Dissertation. University of Kansas, USA. <http://hdl.handle.net/1808/19416>
- Ralf Hinze. 2000. A New Approach to Generic Functional Programming. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '00)*. ACM, New York, NY, USA, 119–132. <https://doi.org/10.1145/325694.325709>
- Mauro Jaskieloff and Russell O'Connor. 2015. A Representation Theorem for Second-Order Functionals. *Journal of Functional Programming* 25, e13 (2015). <https://doi.org/10.1017/S0956796815000088>
- Edward Kmett. 2018. lens-4.16 library. <https://hackage.haskell.org/package/lens-4.16>
- Ralf Lämmel and Simon Peyton Jones. 2003. Scrap Your Boilerplate: A Practical Design Pattern for Generic Programming. In *Types in Languages Design and Implementation*. ACM Press, New York, NY, USA, 26–37. <https://doi.org/10.1145/640136.604179>
- Ralf Lämmel and Simon Peyton Jones. 2004. Scrap More Boilerplate: Reflection, Zips, and Generalised Casts. In *Proceedings of the Ninth ACM SIGPLAN International Conference on Functional Programming (ICFP '04)*. ACM, New York, NY, USA, 244–255. <https://doi.org/10.1145/1016850.1016883>
- José Pedro Magalhães, Atze Dijkstra, Johan Jeuring, and Andres Löb. 2010. A Generic Deriving Mechanism for Haskell. In *Proceedings of the Third ACM Haskell Symposium on Haskell (Haskell '10)*. ACM, New York, NY, USA, 37–48. <https://doi.org/10.1145/1863523.1863529>
- José Pedro Magalhães. 2012. Optimisation of generic programs through inlining. In *Symposium on Implementation and Application of Functional Languages*. Springer, 104–121.
- José Pedro Magalhães. 2014. Generic Programming with Multiple Parameters. In *Functional and Logic Programming*, Michael Codish and Eijiro Sumii (Eds.). Springer International Publishing, Cham, 136–151.
- Conor McBride and Ross Paterson. 2008. Applicative programming with effects. *Journal of functional programming* 18, 1 (2008), 1–13.
- Neil Mitchell and Colin Runciman. 2007. Uniform boilerplate and list processing. In *Proceedings of the ACM SIGPLAN workshop on Haskell*. ACM, 49–60.
- Russell O'Connor. 2011. Functor is to Lens as Applicative is to Biplate: Introducing Multiplate. *CoRR* abs/1103.2841 (2011). <https://doi.org/arXiv:1103.2841> Presented at WGP 2011.

- Simon Peyton Jones and Simon Marlow. 2002. Secrets of the glasgow haskell compiler inliner. *Journal of Functional Programming* 12, 4-5 (2002), 393–434.
- Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. 2007. Practical Type Inference for Arbitrary-rank Types. *J. Funct. Program.* 17, 1 (Jan. 2007), 1–82. <https://doi.org/10.1017/S0956796806006034>
- Simon Peyton Jones, Stephanie Weirich, Richard A Eisenberg, and Dimitrios Vytiniotis. 2016. A reflection on types. In *A List of Successes That Can Change the World*. Springer, 292–317.
- Simon L Peyton Jones and AndréL M Santos. 1998. A transformation-based optimiser for Haskell. *Science of computer programming* 32, 1-3 (1998), 3–47.
- Matthew Pickering, Jeremy Gibbons, and Nicolas Wu. 2017. Profunctor Optics: Modular Data Accessors. *Programming Journal* 1, 2 (2017), 7. <https://doi.org/10.22152/programming-journal.org/2017/1/7>
- Martin Sulzmann, Gregory J. Duck, Simon Peyton Jones, and Peter J. Stuckey. 2007. Understanding Functional Dependencies via Constraint Handling Rules. *J. Funct. Program.* 17, 1 (Jan. 2007), 83–129. <https://doi.org/10.1017/S0956796806006137>
- Twan van Laarhoven. 2009. CPS-Based Functional References. (July 2009). <http://www.twanvl.nl/blog/haskell/cps-functional-references>
- Sjoerd Visscher and Li-yao Xia. 2018. one-liner-1.0 library. <http://hackage.haskell.org/package/one-liner-1.0>
- Dimitrios Vytiniotis, Simon Peyton Jones, Tom Schrijvers, and Martin Sulzmann. 2011. Outsidein(x) Modular Type Inference with Local Assumptions. *J. Funct. Program.* 21, 4-5 (Sept. 2011), 333–412. <https://doi.org/10.1017/S0956796811000098>
- Phillip Wadler and Stephen Blott. 1989. How to Make Ad-hoc Polymorphism Less Ad Hoc. In *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '89)*. ACM, New York, NY, USA, 60–76. <https://doi.org/10.1145/75277.75283>
- Stephanie Weirich, Dimitrios Vytiniotis, Simon Peyton Jones, and Steve Zdancewic. 2011. Generative Type Abstraction and Type-level Computation. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '11)*. ACM, New York, NY, USA, 227–240. <https://doi.org/10.1145/1926385.1926411>
- Jeremy Yallop. 2017. Staged Generic Programming. *Proc. ACM Program. Lang.* 1, ICFP, Article 29 (Aug. 2017), 29 pages. <https://doi.org/10.1145/3110273>
- Brent A. Yorgey, Stephanie Weirich, Julien Cretin, Simon Peyton Jones, Dimitrios Vytiniotis, and José Pedro Magalhães. 2012. Giving Haskell a Promotion. In *Proceedings of the 8th ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI '12)*. ACM, New York, NY, USA, 53–66. <https://doi.org/10.1145/2103786.2103795>