

A Typical Synergy

Dynamic Types and Generalised Algebraic Datatypes

Thomas van Noort, Peter Achten, and Rinus Plasmeijer

Institute for Computing and Information Sciences
Radboud University Nijmegen
P.O. Box 9010, 6500 GL Nijmegen, The Netherlands
{thomas,p.achten,rinus}@cs.ru.nl

Abstract. We present a typical synergy between dynamic types (dynamics) and generalised algebraic datatypes (GADTs). The former provides a clean approach to integrating dynamic typing in a statically typed language. It allows values to be wrapped together with their type in a uniform package, deferring type unification until run time using a pattern match annotated with the desired type. The latter allows for the explicit specification of constructor types, as to enforce their structural validity. In contrast to ADTs, GADTs are heterogeneous structures since each constructor type is implicitly universally quantified. Unfortunately, pattern matching only enforces structural validity and does not provide instantiation information on polymorphic types. Consequently, functions that manipulate such values, such as a type-safe update function, are cumbersome due to boilerplate type representation administration. In this paper we focus on improving such functions by providing a new GADT annotation via a natural synergy with dynamics. We formally define the semantics of the annotation and touch on novel other applications of this technique such as type dispatching and enforcing type equality invariants on GADT values.

1 Introduction

In this paper we discuss a typical synergy between two concepts: dynamic types (dynamics) and generalised algebraic datatypes (GADTs).

Types play an important role in strongly typed functional programming languages such as Clean and Haskell. Using static type checking, erroneous behaviour at run time is prevented. Moreover, more efficient code can be generated using the knowledge provided by the types at compile time. However, in dynamic systems that deal with user input, some types will only be known at run time. Using dynamics, monomorphic [1] and polymorphic [2] values can be wrapped together with their type in a black box. The dynamic is unwrapped by pattern matching on the required type in a function definition, instead of specifying the type explicitly in its signature. This approach defers part of the type checking process until run time, exactly when the final required type information is made available. Fortunately, this does not take place at the cost of the advantage of

static typing since the type system guarantees that when pattern matching succeeds, the unwrapped dynamic can be used safely as dictated by the specified type. Of course, pattern matching can fail and cause a run-time error, but this is not different from conventional pattern matching.

Algebraic datatypes (ADTs) in functional languages allow us to inductively define structures. Unfortunately, it does not allow us to enforce structural validity at compile time. With the arrival of generalised algebraic datatypes (GADTs) [6,10,14], this restriction is relieved by allowing constructors to explicitly dictate their types. On the one hand, this prevents us from constructing ill-structured (i.e., ill-typed) values, and on the other hand this ensures structural validity once a GADT value is pattern matched. In contrast to ADTs, GADTs are heterogeneous structures since each constructor occurrence is implicitly universally quantified. Pattern matching such a value only introduces information regarding the structure of the constructor types, leaving the type variables polymorphic. However, more information on their instantiation is often required, typically in functions that manipulate such values. Conventional approaches to this problem are cumbersome, due to boilerplate type representation administration.

The main contribution of this paper is to define a type-safe update function on GADT values via an annotation, achieved by a natural synergy between dynamics and GADTs.

Overview

This paper is organised as follows. First, we elaborate on both dynamics and GADTs (Section 2). We motivate the need for the synergy by defining an update function on λ -terms (Section 3). Then, we formally define a semantics for the new annotation via a synergy between dynamics and GADTs (Section 4). We conclude with related work (Section 5) and a discussion on future work and other applications of this technique (Section 6). In this paper we use Clean's dynamics and Haskell's GADTs. For the sake of presentation, our examples use Haskell syntax, augmented with Clean's notation for dynamics.

2 Preliminaries

We start by introducing dynamics (Section 2.1) and GADTs (Section 2.2).

2.1 Dynamic Types

The advantage of statically typed languages is that types are verified at compile time, preventing erroneous behaviour at run time due to ill-typed values. However, static typing sometimes does not suffice since a type might only be known at run time. Using dynamic types, values are wrapped in a black box, not exposing the type of the contents to the outside world. But unlike existential types [9], both the value and its type are unwrapped by pattern matching the black box, thereby obtaining a value of the matched content type.

In Clean, the keyword **dynamic** provides the mechanism to wrap values together with their type in a dynamic [13], obtaining a value of type *Dynamic*:

```
wrapInt :: Int → Dynamic
wrapInt x = dynamic x
```

Unwrapping the integer value is achieved by pattern matching on the dynamic value using the `::` annotation, thereby providing a required type:

```
unwrapInt :: Dynamic → Int
unwrapInt (x :: Int) = x
unwrapInt (x :: String) = stringToInt x
unwrapInt _ = 0
```

The first arm of the function pattern matches on a value x of type *Int* in the dynamic. If this is the case, the value is returned unchanged. However, the value in the dynamic is possibly a string and has to be converted to an integer first. Due to run-time type unification, the dynamic pattern match can fail in case the wrapped value is not of the type *Int* or *String*. It is our responsibility to provide a catch-all arm which either returns a default value or a run-time error message.

Instead of defining a function for each value type that is turned into a dynamic, we define a single function:

```
wrap :: TC α ⇒ α → Dynamic
wrap x = dynamic x
```

Since this function is polymorphic in the argument type, we require the context to provide the type code (i.e., the value representation of the type) of α which is stored together with the value x , using Clean's built-in *TC* class constraint. A type code also contains the definition of the type it describes, because dynamics can be (de)serialised across modules and verifying name equivalence in a dynamic pattern match does not suffice. Consequently, *TC* instances are only available for nonabstract types.

Unlike Haskell, Clean supports type-dependent dynamics [11], which allows us to use pattern variables in the type of a dynamic pattern match:

```
unwrap :: TC α ⇒ Dynamic → α
unwrap (x :: α^) = x
unwrap _ = error "unwrap: incorrect type"
```

We require x to be of type α and refer to the same variable in the result type of *unwrap* using the [^] annotation. This causes both types to be coerced automatically at run time. Therefore, a type code is required for α such that it can be compared with the type code obtained from the dynamic pattern match. The context in which this function is used determines which type code is provided.

Pattern variables can also be used to enforce type equality, for example, to define function application of dynamics:

$$\begin{aligned}
\text{apply} &:: TC \alpha \Rightarrow \text{Dynamic} \rightarrow \text{Dynamic} \rightarrow \text{Maybe } \alpha \\
\text{apply } (f &:: \beta \rightarrow \alpha^{\wedge}) (x &:: \beta) &= \text{Just } (f \ x) \\
\text{apply } _ & \quad \quad \quad _ &= \text{Nothing}
\end{aligned}$$

The dynamic pattern matches in the first arm share the same scope. Therefore, they only succeed once the argument type of the function matches the type of the argument. Because the result type of the function in the first dynamic pattern match refers to α in the result type of *apply*, a type code is required for this type. As an example, consider the following expressions:

$$\begin{aligned}
\text{apply } (\mathbf{dynamic} \text{ fst}) (\mathbf{dynamic} (1, "2")) &\rightsquigarrow \text{Just } 1 \\
\text{apply } (\mathbf{dynamic} \text{ fst}) (\mathbf{dynamic} 1) &\rightsquigarrow \text{Nothing}
\end{aligned}$$

While the first expression succeeds, the second expression fails since the argument is not a pair. Finally, dynamics preserve lazy behaviour of functional programs:

$$\text{apply } (\mathbf{dynamic} \text{ fst}) (\mathbf{dynamic} (1, \perp)) \rightsquigarrow \text{Just } 1$$

Although the value \perp is part of the tuple that is wrapped in a dynamic, it is not evaluated when (un)wrapped.

2.2 Generalised Algebraic Datatypes

Algebraic datatypes are an oft-used abstraction in functional languages since they provide an inductive approach to defining complex structures by enumerating the alternatives of a type and the associated fields. For example, in Haskell, an ADT representing λ -terms could be defined as follows:

```

data Lam = Undef
        | Const Value
        | App Lam Lam

```

The *Undef* constructor has no fields, while the *Const* constructor has a single field for a value. The *App* constructor has two fields, which both can be any term. The values are enumerated by another ADT:

```

data Value = VInt Int
           | VFun (Value → Value)

```

Next, we define an evaluation function:

```

eval :: Lam → Value
eval Undef      = ⊥
eval (Const x) = x
eval (App f x) = case eval f of
                VFun f → f (eval x)
                _      → error "eval: not a function"

```

The arms for *Undef* and *Const* are straightforward. However, since nothing prevents us from constructing ill-typed terms, the arm for *App* has to ensure that its first field actually evaluates to a function.

With the arrival of generalised abstract datatypes, we are able to enforce structural validity by providing an explicit type signature to each constructor. Consequently, a GADT imposes a heterogeneous structure since all constructors are implicitly universally quantified. We illustrate the use of GADTs by defining the *Lam* type again, this time describing typed λ -terms:

```
data Lam :: * -> * where
  Undef :: Lam  $\alpha$ 
  Const ::  $\alpha$  -> Lam  $\alpha$ 
  App   :: Lam ( $\beta$  ->  $\alpha$ ) -> Lam  $\beta$  -> Lam  $\alpha$ 
```

The *Lam* type is parameterised by the result type of the term once it is evaluated. With each constructor, we explicitly specify its result type. The *Undef* constructor represents an undefined value. Since its result type α is free and not bound by any fields, it can be unified with any other type. The *Const* constructor lifts any value to the *Lam* type. The *App* constructor is more explicit about the types of its two field. The argument type of the function term must match the type of the argument term. Then, its result type is the result type of the function term. The explicit constructor types prevent us from constructing ill-typed terms. Consider the following examples:

$$\perp 1 \quad \equiv \quad \text{App Undef (Const 1)}$$

Since the return type of the *Undef* constructor can be anything, it is instantiated to a function as *App* requires, thereby returning a value of type *Lam* α . When we provide a term that does not return a function, the term becomes ill typed:

$$0 1 \quad \equiv \quad \text{App (Const 0) (Const 1)}$$

A more useful example actually applies a function. Consider the absolute value of an integer:

$$\text{abs } 1 \quad \equiv \quad \text{App (Const abs) (Const 1)}$$

This term is well typed and returns a value of type *Lam Int*.

Type information described in the type of the constructors is also employed when the constructors are pattern matched in a function definition. Since only well-typed terms can be constructed, we can now safely and concisely define the evaluation function:

```
eval :: Lam  $\alpha$  ->  $\alpha$ 
eval Undef   =  $\perp$ 
eval (Const x) = x
eval (App f x) = eval f (eval x)
```

The result type of the function depends on the term that is evaluated. Each constructor dictates the type of its fields as well as the result type. For example, evaluating the first field of the *App* constructors returns a function, which can safely be applied to its evaluated second field. However, be aware that the exact types of these fields are not known since we are dealing with a heterogeneous structure.

3 Motivation

In this section, we motivate the need for the typical synergy between dynamics and GADTs in the context of update functions on GADTs (Section 3.1). Next, we discuss why the conventional approach is not suited to this problem (Section 3.2) and how the synergy elegantly improves on these issues using a new GADT annotation (Section 3.3).

3.1 Setting the Scene

We use as the running example the definition from Section 2.2 that represents λ -terms. Our goal is to define an update function that takes such a term, and updates a field of a constructor at a specified position with a new value. Then, the desired type of the update function becomes:

$$update :: Lam\ \alpha \rightarrow Path \rightarrow \beta \rightarrow Lam\ \alpha$$

The argument and result type of the function are the same since we only consider updates that do not affect the top-level type of the term. However, an update can change the structure. The path depicts the location of the update in the heterogeneous structure:

$$\mathbf{type}\ Path = [Int]$$

The path is represented as a list of integers. The length of the list indicates the recursive level (where the empty list is the root) of the target and each value the field (where 0 is the first field) that must be considered. Since the path possibly dictates an update anywhere in the heterogeneous structure, the type of the new value is unrestricted. Hence, the challenge we face lies in only allowing type-safe updates.

3.2 Conventional Approach

The conventional approach to this problem makes extensive use of equality types [3,5]. By comparing the value representations of the type of the old and new value, a proof of type equality can be obtained to ensure only type-safe updates.

First, we modify our original *Lam* definition from Section 2.2 to the following:

```
data LamR :: * → * where
  UndefR :: LamR α
  ConstR :: RepOf α → LamR α
  AppR   :: RepOf (LamR (β → α)) → RepOf (LamR β) → LamR α
```

The difference is that the types of the constructor fields now include a type representation:

```
type RepOf α = (α, Rep α)
```

The *Rep* type enumerates the possible types, including the integer type, the function type, and the *Lam_R* type:

```
data Rep :: * → * where
  RInt   :: Rep Int
  RFun   :: Rep α → Rep β → Rep (α → β)
  RLamR :: Rep α → Rep (LamR α)
```

The *Rep* type is only a witness of a type, for example, the type *Lam_R (Int → Int)* is witnessed by the value *RLam_R (RFun RInt RInt)*. For the sake of brevity, this representation type only reflects monomorphic types. Given such witnesses, we are able to construct the actual proof that the types of such *Rep* values are the same. Such a proof is constructed by the following GADT:

```
data Equal :: * → * → * where
  Refl :: Equal α α
```

The *Equal* type consists of a single constructor *Refl*, one that proves that both of the type arguments are the same. Then, we define a type equality function that performs a point-wise comparison of type representations, using Haskell's **do** notation:

```
eqR :: Rep α → Rep β → Maybe (Equal α β)
eqR RInt      RInt      = Just Refl
eqR (RFun x1 x2) (RFun y1 y2) = do Refl ← eqR x1 y1
                                     Refl ← eqR x2 y2
                                     return Refl
eqR (RLamR x)   (RLamR y) = do Refl ← eqR x y
                                     return Refl
eqR -           -           = Nothing
```

Given two *Rep* values, this function either returns *Just Refl* if the type representations are the same, thereby implicitly indicating that the types α and β are the same as well, or *Nothing*. In the arms for *RFun* and *RLam_R* we have to explicitly pattern match the result of the recursion as to obtain its type equality proof. Finally, we define a catch-all arm which returns *Nothing* for *Rep* values that are not equal.

Then, using the modified *Lam* definition and a type representation added to the new value, we are finally able to define our update function:

$$\begin{aligned}
& \text{update}_R :: \text{Lam}_R \alpha \rightarrow \text{Path} \rightarrow \text{RepOf } \beta \rightarrow \text{Lam}_R \alpha \\
& \text{update}_R \text{ Undef}_R \quad [] \quad - \quad = \text{Undef}_R \\
& \text{update}_R (\text{Const}_R (x, rx)) [0] \quad (y, ry) = \mathbf{case} \text{ eq}_R \text{ rx ry of} \\
& \qquad \qquad \qquad \text{Just Refl} \rightarrow \text{Const}_R (y, ry) \\
& \qquad \qquad \qquad \text{Nothing} \rightarrow \text{Const}_R (x, rx) \\
& \text{update}_R (\text{App}_R (f, rf) x) [0] \quad (y, ry) = \mathbf{case} \text{ eq}_R \text{ rf ry of} \\
& \qquad \qquad \qquad \text{Just Refl} \rightarrow \text{App}_R (y, ry) x \\
& \qquad \qquad \qquad \text{Nothing} \rightarrow \text{App}_R (f, rf) x \\
& \text{update}_R (\text{App}_R f (x, rx)) [1] \quad (y, ry) = \mathbf{case} \text{ eq}_R \text{ rx ry of} \\
& \qquad \qquad \qquad \text{Just Refl} \rightarrow \text{App}_R f (y, ry) \\
& \qquad \qquad \qquad \text{Nothing} \rightarrow \text{App}_R f (x, rx) \\
& \text{update}_R (\text{App}_R (f, rf) x) (0 : p) y \quad = \text{App}_R (\text{update}_R f p y, rf) x \\
& \text{update}_R (\text{App}_R f (x, rx)) (1 : p) y \quad = \text{App}_R f (\text{update}_R x p y, rx) \\
& \text{update}_R x \quad - \quad - \quad = x
\end{aligned}$$

In the arm for *Undef_R* there is nothing left to do, we only have to make sure that the path is fully consumed. The *Const_R* is the first interesting case, since we have to verify that the types match by testing the equality of the *Rep* values. Once these values are the same, we provide a proof that α and β are equal types by pattern matching on the *Refl* constructor. Then, in the arms for *App_R* we use the same approach and either replace its first or second field, or dispatch on the head of the path and continue to recurse in either of its fields. Finally, a catch-all arm is included to return the original term once the provided path is incorrect. Whenever the function is applied, all the type representations need to be provided explicitly:

$$\begin{aligned}
& \text{update}_R (\text{Const}_R (\text{abs}, \text{RFun } \text{RInt } \text{RInt})) [0] (\text{neg}, \text{RFun } \text{RInt } \text{RInt}) \\
& \quad \rightsquigarrow \\
& \text{Const}_R (\text{neg}, \text{RFun } \text{RInt } \text{RInt})
\end{aligned}$$

Although this approach guarantees type-safe updates, it is not a very elegant definition. First of all, the invasive inclusion of *Rep* values in the datatype clutters the update function with type equality witnesses and manual proofs. Moreover, the types of the values that are updated have to be known beforehand since these are enumerated in the *Rep* type and traversed in the type equality function. Above all, this approach does not scale up to more complex structures and update functions.

3.3 The Synergy

The conventional approach requires us to carry around type representations that are used to convince the type checker of type equality. When we look back at Section 2.1, we notice that this is actually what Clean's *TC* type class provides. We propose to adapt the original *Lam* definition from Section 2.2 again:

data $Lam_T :: \star \rightarrow \star$ **where**

$$\begin{aligned} Undef_T &:: Lam_T \alpha \\ Const_T &:: TC \alpha \Rightarrow \alpha \rightarrow Lam_T \alpha \\ App_T &:: (TC \beta, TC \alpha) \Rightarrow Lam_T (\beta \rightarrow \alpha) \rightarrow Lam_T \beta \rightarrow Lam_T \alpha \end{aligned}$$

Instead of including Rep values, we include TC class constraints with the constructors that can be updated. Then, we define the update function using the new $::^G$ annotation on the field of a GADT constructor:

$$\begin{aligned} update_T &:: TC \beta \Rightarrow Lam_T \alpha \rightarrow Path \rightarrow \beta \rightarrow Lam_T \alpha \\ update_T Undef_T & \quad [] \quad - = Undef_T \\ update_T (Const_T (x ::^G \beta^\wedge)) & [0] \quad y = Const_T y \\ update_T (App_T (f ::^G \beta^\wedge) x) & [0] \quad y = App_T y x \\ update_T (App_T f (x ::^G \beta^\wedge)) & [1] \quad y = App_T f y \\ update_T (App_T f x) & \quad (0 : p) y = App_T (update_T f p y) x \\ update_T (App_T f x) & \quad (1 : p) y = App_T f (update_T x p y) \\ update_T x & \quad - \quad - = x \end{aligned}$$

Let us take a look at the differences between this update function and the conventional definition $update_R$ from Section 3.2. First of all, this function operates on the Lam_T type that is decorated with TC constraints, and its type contains a TC constraint to obtain a type code for the new value of type β . Although the update function was intended to be polymorphic at first, this constraint only forbids abstract types to occur as new values, as discussed earlier in Section 2.1. Another difference is that the function is no longer cluttered with verbose type equality witnesses and manual proofs. Instead, the fields of the constructors are annotated using the $::^G$ annotation, accessing the instantiated polymorphic type information. For example, in the arm for $Const_T$, the annotation denotes that x is of type β , or even more specific, the type of the new value as determined by the context in which this function is used. Note that the catch-all arm now also takes care of any failing tests for type equality. Comparing the use of this update function to the conventional approach emphasises the elegance of our approach:

$$update_T (Const_T abs) [0] neg \rightsquigarrow Const_T neg$$

Instead of explicitly providing type representations and equality proofs, it is now the context that implicitly determines which fields are eligible for an update.

4 Semantics

In this section we present the formal semantics of the synergy. We formally define a core functional language and the GADT annotation extension (Section 4.1). Then, we describe the idea behind the translation from the extended language to the core language by means of an example (Section 4.2), followed by a formal approach (Section 4.3).

(program)	$\pi ::= \bar{\delta} \bar{\phi}$
(datatype declaration)	$\delta ::= \mathbf{type} \ T \ \bar{\alpha} = \tau$ $\quad \ \mathbf{data} \ T \ \bar{\alpha} = \overline{C \ \tau}$ $\quad \ \mathbf{data} \ T :: \kappa \ \mathbf{where} \ \overline{C} :: \sigma$
(qualified type)	$\sigma ::= \overline{TC} \ \alpha \Rightarrow \tau$
(base type)	$\tau ::= \alpha \mid Int \mid T$ $\quad \ \tau_1 \ \tau_2 \mid \tau_1 \rightarrow \tau_2$ $\quad \ \mathit{Dynamic}$
(annotation type)	$\omega ::= \alpha \mid \alpha^\wedge \mid Int \mid T$ $\quad \ \omega_1 \ \omega_2 \mid \omega_1 \rightarrow \omega_2$ $\quad \ \mathit{Dynamic}$
(kind)	$\kappa ::= \star \mid \kappa_1 \rightarrow \kappa_2$
(function declaration)	$\phi ::= \mathbf{fix} \ f :: \sigma = \epsilon$
(expression)	$\epsilon ::= \perp \mid i \mid x \mid C$ $\quad \ \epsilon_1 \ \epsilon_2 \mid \lambda x \rightarrow \epsilon \mid \mathbf{case} \ \epsilon_s \ \mathbf{of} \ \overline{\rho} \rightarrow \bar{\epsilon}$ $\quad \ \mathbf{dynamic} \ e :: \omega$
(nested pattern)	$\rho ::= \varrho \mid C \ \bar{\rho}$
(base pattern)	$\varrho ::= _ \mid i \mid x$ $\quad \ x :: \omega$

Fig. 1. The core language FC

4.1 Formal Language

The functional core language FC, which forms the basis of our semantics, is depicted in Fig. 1. It is a common subset of Clean and Haskell, extended with Clean’s dynamics and Haskell’s GADTs. An FC program consists of zero or more datatype declarations and function declarations. A datatype is either a type synonym, an ADT, or a GADT. A type comes in three flavours: a qualified type, a base type, and an annotation type. A qualified type only includes the TC constraint, as to facilitate dynamics, where we write τ as a shorthand for the qualified type $\cdot \Rightarrow \tau$ with no constraints. Second, a base type comprises the polymorphic types. Very much like a base type, we define a separate annotation type, but one that also allows the use of the $^\wedge$ annotation. A named function is defined by its type and body. Amongst the well-known expressions, our language supports the case construct to pattern match values, typically the arguments of a function, and dynamic values. In the language of patterns we distinguish a nested pattern from a base pattern, as to prepare for the language extension. Finally, we do not explicitly include lists and tuples of arbitrary arity in the language of expressions, patterns and types, since these can easily be realized through predefined ADTs. We do not provide operational semantics and typing for the core language since these have been studied in-depth elsewhere [2,4,6].

For example, the Lam_T type from Section 3.3 translates to the following:

data $Lam_T^\circ :: \star \rightarrow \star$ **where**
 $Undef_T^\circ :: Lam_T^\circ \alpha$
 $Const_T^\circ :: TC \alpha \Rightarrow TypeOf \alpha \rightarrow Lam_T^\circ \alpha$
 $App_T^\circ :: (TC \beta, TC \alpha) \Rightarrow TypeOf (Lam_T (\beta \rightarrow \alpha)) \rightarrow TypeOf (Lam_T \beta) \rightarrow Lam_T^\circ \alpha$

The extended definition, as well as its constructors, is given a new name. Since all fields of the constructors are annotated in the update function from Section 3.3, all fields of the $Const_T^\circ$ and App_T° constructor now contain a typed value. Note that in order to only have to translate patterns instead of complete functions later on, the addition of type information is nonrecursive:

type $TypeOf \alpha = (\alpha, Dynamic)$

A typed value is simply the original value paired with its type stored in a dynamic. As a $Dynamic$ can contain a value of any type, not necessarily the type α , we use the following function to obtain correctness by construction:

fix $typeOf :: TC \alpha \Rightarrow \alpha \rightarrow TypeOf \alpha =$
 $\lambda x \rightarrow (x, \mathbf{dynamic} \perp :: \alpha^\wedge)$

Since we only need the type of a value, it suffices to wrap \perp instead of an actual value. As described in Section 2.1, the \wedge annotation refers to context-dependent type information. The context in which $typeOf$ is used determines the type that is stored in the dynamic. Note that constructors can contain GADT values, like App_T° , which requires such types to be stored in a dynamic. Unfortunately, type code facilities are yet to be defined for GADTs. Although GADTs greatly complicate the type inference process [10,12], we hypothesize that storing such values in dynamics is not different from ADT values since it does not affect the unification of type codes that describe a GADT.

Then, the conversion from the original to the extended definition injects the type information in constant time using the function $typeOf$:

fix $toLam_T^\circ :: Lam_T \alpha \rightarrow Lam_T^\circ \alpha =$
 $\lambda x \rightarrow \mathbf{case} \ x \ \mathbf{of}$
 $Undef_T \rightarrow Undef_T^\circ$
 $Const_T \ x \rightarrow Const_T^\circ (typeOf \ x)$
 $App_T \ f \ x \rightarrow App_T^\circ (typeOf \ f) (typeOf \ x)$

The conversion only renames the $Undef_T$ constructor since it has no fields. The fields of the $Const_T$ and App_T constructor are extended with their types. As the function $typeOf$ dictates, this requires a type code for the field types. The translation relies critically on this assumption, which is enforced by only considering a FC^+ program well typed if and only if each constructor has TC constraints on every type variable occurring in its annotated fields. Fortunately, as mentioned before in Section 2.1, the TC constraint is easily discharged for

$$\boxed{\llbracket \pi_{\mathcal{FC}^+} \rrbracket \equiv \pi_{\mathcal{FC}}}$$

$$\frac{\overline{\llbracket \delta \rrbracket} \equiv \overline{\delta'} \phi^\circ \quad \overline{\llbracket \phi \rrbracket} \equiv \overline{\phi'}}{\overline{\llbracket \delta \rrbracket \phi} \equiv \overline{\delta'} \phi^\circ \overline{\phi'}} \text{ (T-PROG)}$$

Fig. 3. Translation of programs

any nonabstract type, which only forbids the use of the GADT annotation in combination with abstract types.

Finally, we define the translation of the actual $::^{\mathcal{G}}$ annotation, accessing the inserted type information. For example, the \mathcal{FC}^+ function $update_T$, as defined in Section 4.1, is translated to \mathcal{FC} :

```

fix  $update_T :: TC \beta \Rightarrow Lam_T \alpha \rightarrow Path \rightarrow \beta \rightarrow Lam_T \alpha =$ 
 $\lambda x \rightarrow \lambda p \rightarrow \lambda y \rightarrow$  case  $(x, p)$  of
  ...
   $(x, (0 : [])) \rightarrow$  case  $toLam_T^\circ x$  of
     $Const_T^\circ(x, \_ :: \beta^\wedge) \rightarrow Const_T y$ 
     $App_T^\circ(f, \_ :: \beta^\wedge)(x, \_) \rightarrow App_T y x$ 
     $\_ \rightarrow x$ 
   $(x, (1 : [])) \rightarrow$  case  $toLam_T^\circ x$  of
     $App_T^\circ(f, \_)(x, \_ :: \beta^\wedge) \rightarrow App_T f y$ 
     $\_ \rightarrow x$ 
  ...

```

The conversion from the original to the extended GADT is applied to the scrutinee of the case expression. This provides type information in the pattern match, allowing it to interact naturally like a conventional dynamic, in this case with the type of the function using the \wedge annotation. Note that since the conversion function is specific to a program, and not to each case expression, the fields that do not use the annotation must discard the inserted type information as in the case for App_T° .

4.3 Formal Translation

We continue by defining the formal translation from the extended language \mathcal{FC}^+ to the core language \mathcal{FC} . We conjecture that the translation is sound, every well-typed \mathcal{FC}^+ program is translated to a well-typed \mathcal{FC} program.

Let us begin by translating programs, as depicted by Rule T-PROG in Fig. 3. A program in the \mathcal{FC}^+ language is translated to the \mathcal{FC} language by translating both the datatype declarations and the function declarations.

In Fig. 4 we define the translation of datatype declarations. Type synonyms and ADTs are left unchanged, as defined by Rules T-DATA-TSYN and T-DATA-ADT. We distinguish GADTs by using the metafunction $annotated(T)$ to test if it is

$$\boxed{\llbracket \delta_{\text{FC}+} \rrbracket \equiv \overline{\delta}_{\text{FC}}; \phi_{\text{FC}}}$$

$$\frac{}{\llbracket \text{type } T \overline{\alpha} = \tau \rrbracket \equiv \text{type } T \overline{\alpha} = \tau; \cdot} \text{ (T-DATA-TSYN)}$$

$$\frac{}{\llbracket \text{data } T \overline{\alpha} = \overline{C} \tau \rrbracket \equiv \text{data } T \overline{\alpha} = \overline{C} \tau; \cdot} \text{ (T-DATA-ADT)}$$

$$\frac{\neg \text{annotated}(T)}{\llbracket \text{data } T :: \kappa \text{ where } \overline{C} :: \sigma \rrbracket \equiv \text{data } T :: \kappa \text{ where } \overline{C} :: \sigma; \cdot} \text{ (T-DATA-GADT-1)}$$

$$\begin{array}{l}
\overline{\llbracket \sigma \rrbracket}_C \equiv \overline{\sigma^\circ}; \overline{x}; \overline{\epsilon^\circ} \\
\text{arity}(T) \equiv \overline{\alpha} \\
\text{annotated}(T) \\
\delta^\circ \equiv \text{data } T^\circ :: \kappa \text{ where } \overline{C}^\circ :: \sigma^\circ \\
\phi^\circ \equiv \text{fix } \text{to} T^\circ :: T \overline{\alpha} \rightarrow T^\circ \overline{\alpha} = \\
\lambda x \rightarrow \text{case } x \text{ of} \\
\frac{}{C \overline{x} \rightarrow C^\circ \overline{\epsilon^\circ}}
\end{array}$$

$$\frac{}{\llbracket \text{data } T :: \kappa \text{ where } \overline{C} :: \sigma \rrbracket \equiv \text{data } T :: \kappa \text{ where } \overline{C} :: \sigma \delta^\circ; \phi^\circ} \text{ (T-DATA-GADT-2)}$$

Fig. 4. Translation of datatypes

$$\boxed{\llbracket \sigma_{\text{FC}+} \rrbracket_C \equiv \sigma_{\text{FC}}; \overline{x}; \overline{\epsilon}_{\text{FC}}}$$

$$\frac{\llbracket \tau \rrbracket_{C;0} \equiv \tau^\circ; \overline{x}; \overline{\epsilon^\circ}}{\llbracket TC \overline{\alpha} \Rightarrow \tau \rrbracket \equiv TC \overline{\alpha} \Rightarrow \tau^\circ; \overline{x}; \overline{\epsilon^\circ}} \text{ (T-QTYPE)}$$

Fig. 5. Translation of qualified types

pattern matched somewhere in the program using the GADT annotation (e.g., $\text{annotated}(Lam^\circ) \equiv \text{True}$). If not, the original definition is returned without any modifications, as defined by Rule T-DATA-GADT-1. However, an annotated GADT requires some effort. In Rule T-DATA-GADT-2, the translation results in the original definition, an extended definition δ° , and a conversion function ϕ° . By translating the types of the constructors parameterized by the respective constructor name, we obtain extended types together with corresponding pattern variables and expressions that extend these variables. The former is used to define the constructor types of the extended definition, the latter two to define the corresponding conversion function. The metafunction $\text{arity}(T)$ provides zero or more fresh type variables, determined by the arity of the type T (e.g., $\text{arity}(Lam_T) \equiv \alpha$).

The translation of qualified types, parameterised by a constructor name, is shown in Fig. 5. A qualified type propagates the translation to its base type, adding a parameter which represents the index of the constructor field type under translation.

In Fig. 6 we define the parameterised translation of such types, resulting in an extended type, pattern variables, and expressions that extends these

$$\boxed{\llbracket \tau_{\text{FC}^+} \rrbracket_{C;n} \equiv \tau_{\text{FC}}; \bar{x}; \bar{c}_{\text{FC}}}$$

$$\frac{}{\llbracket T \rrbracket_{C;n} \equiv T; \cdot; \cdot} \text{ (T-TYPE-DATA)} \qquad \frac{}{\llbracket \tau_1 \tau_2 \rrbracket_{C;n} \equiv \tau_1 \tau_2; \cdot; \cdot} \text{ (T-TYPE-APP)}$$

$$\frac{\neg \text{annotated}(C, n) \quad \llbracket \tau_2 \rrbracket_{C;n+1} \equiv \tau_2^\circ; \bar{x}_2; \bar{c}_2^\circ}{\llbracket \tau_1 \rightarrow \tau_2 \rrbracket_{C;n} \equiv \tau_1 \rightarrow \tau_2^\circ; x_1 \bar{x}_2; x_1 \bar{c}_2^\circ} \text{ (T-TYPE-FUN-1)}$$

$$\frac{\text{annotated}(C, n) \quad \llbracket \tau_2 \rrbracket_{C;n+1} \equiv \tau_2^\circ; \bar{x}_2; \bar{c}_2^\circ}{\llbracket \tau_1 \rightarrow \tau_2 \rrbracket_{C;n} \equiv \text{TypeOf } \tau_1 \rightarrow \tau_2^\circ; x_1 \bar{x}_2; (\text{typeOf } x_1) \bar{c}_2^\circ} \text{ (T-TYPE-FUN-2)}$$

Fig. 6. Translation of base types

$$\boxed{\llbracket \phi_{\text{FC}^+} \rrbracket \equiv \phi_{\text{FC}}}$$

$$\frac{\llbracket \epsilon \rrbracket \equiv \epsilon'}{\llbracket \text{fix } f :: \sigma = \epsilon \rrbracket \equiv \text{fix } f :: \sigma = \epsilon'} \text{ (T-FUN)}$$

Fig. 7. Translation of functions

variables. Since we are only interested in the fields of a constructor type, and the type of an empty constructor is either a type constructor or a type application, Rules T-TYPE-DATA and T-TYPE-APP result in an unchanged type and no pattern variables or expressions. The function type is the interesting case. If a constructor field is not annotated in the program, as shown in Rule T-TYPE-FUN-1, it is returned unchanged together with a fresh pattern variable and expression that corresponds to the identity. Otherwise, the translation in Rule T-TYPE-FUN-2 extends the type of the constructor field with additional type information and ensures that the fresh pattern variable is extended as well. In both cases we recurse in the translation by incrementing the second parameter to denote the next constructor field.

Next, we define the translation of functions by Rule T-FUN in Fig. 7. A function is translated by translating its body expression, which localises the conversion and thus does not change the type of a function.

The translation of expressions is shown in Fig. 8. The basic building blocks of expressions: bottom, integers, identifiers, and constructors, are left unchanged, as can be seen in Rules T-EXP-BOT, T-EXP-INT, T-EXP-ID, and T-EXP-CON respectively. Translation of an application is defined by Rule T-EXP-APP as translating both of its expressions. Rule T-EXP-ABS defines the translation of an abstraction by translating the body expression. For case expressions, we define two separate rules, testing if one of its patterns uses the GADT annotation. If not, it suffices to only translate the scrutinee and the expression of each pattern, as defined by Rule T-EXP-CASE-1. Otherwise, Rule T-EXP-CASE-2 defines that the conversion function must be applied

$$\boxed{[\epsilon_{\text{FC}^+}] \equiv \epsilon_{\text{FC}}}$$

$$\frac{}{[\perp] \equiv \perp} \text{ (T-EXP-BOT)} \qquad \frac{}{[i] \equiv i} \text{ (T-EXP-INT)}$$

$$\frac{}{[x] \equiv x} \text{ (T-EXP-ID)} \qquad \frac{}{[C] \equiv C} \text{ (T-EXP-CON)}$$

$$\frac{[\epsilon_1] \equiv \epsilon'_1 \quad [\epsilon_2] \equiv \epsilon'_2}{[\epsilon_1 \ \epsilon_2] \equiv \epsilon'_1 \ \epsilon'_2} \text{ (T-EXP-APP)} \qquad \frac{[\epsilon] \equiv \epsilon'}{[\lambda x \rightarrow \epsilon] \equiv \lambda x \rightarrow \epsilon'} \text{ (T-EXP-ABS)}$$

$$\frac{x ::^{\mathcal{G}} \alpha \notin \bar{\theta} \quad \frac{[\epsilon_s] \equiv \epsilon'_s \quad \bar{\theta} \equiv \bar{\rho} \quad [\epsilon] \equiv \epsilon'}{\text{case } \epsilon_s \text{ of } \bar{\theta} \rightarrow \epsilon] \equiv \text{case } \epsilon'_s \text{ of } \bar{\rho} \rightarrow \epsilon'} \text{ (T-EXP-CASE-1)}$$

$$\frac{x ::^{\mathcal{G}} \alpha \in \bar{\theta} \quad \text{btype}(\epsilon_s) \equiv T \quad \frac{[\epsilon_s] \equiv \epsilon'_s \quad [\bar{\theta}] \equiv \bar{\rho} \quad [\epsilon] \equiv \epsilon'}{\text{case } \epsilon_s \text{ of } \bar{\theta} \rightarrow \epsilon] \equiv \text{case } \text{to} T^{\circ} \epsilon'_s \text{ of } \bar{\rho} \rightarrow \epsilon'} \text{ (T-EXP-CASE-2)}$$

$$\frac{[\epsilon] \equiv \epsilon'}{[\text{dynamic } \epsilon :: \omega] \equiv \text{dynamic } \epsilon' :: \omega} \text{ (T-EXP-DYN)}$$

Fig. 8. Translation of expressions

$$\boxed{[\theta_{\text{FC}^+}] \equiv \rho_{\text{FC}}}$$

$$\frac{}{[\varrho] \equiv \varrho} \text{ (T-PAT-BASE)} \qquad \frac{[\bar{\theta}]_{C;\text{index}} \equiv \bar{\rho}}{[C \ \bar{\theta}] \equiv C^{\circ} \bar{\rho}} \text{ (T-PAT-CON)}$$

Fig. 9. Translation of patterns

to the translated scrutinee. The name of this function is determined by the metafunction $\text{btype}(\epsilon_s)$ which determines the base name of the type of the scrutinee ϵ_s (e.g., $\text{btype}(\text{Const}_T \ 1) \equiv \text{Lam}_T$). Furthermore, each pattern is translated so that the actual use of the annotation is translated. As we will see in a moment, the translation of patterns takes care of renaming the constructors, which is required since the scrutinee is converted to the extended type. Finally, Rule T-EXP-DYN defines the translation of a dynamic, simply translating its expression.

Patterns possibly provide access to the inserted type information, their translation is shown in Fig. 9. A base pattern is left untouched, as depicted in Rule T-PAT-BASE. In Rule T-PAT-CON, the constructor in a constructor pattern is renamed and its fields are all translated, parameterised by the name of the original constructor and a metavalue index that provides the index of each constructor field.

$$\boxed{\llbracket \vartheta_{\text{FC}^+} \rrbracket_{C;n} \equiv \rho_{\text{FC}}}$$

$$\frac{\neg \text{annotated}(C, n)}{\llbracket \rho \rrbracket_{C;n} \equiv \rho} \text{ (T-FPAT-PAT-1)} \qquad \frac{\text{annotated}(C, n)}{\llbracket \rho \rrbracket_{C;n} \equiv (\rho, -)} \text{ (T-FPAT-PAT-2)}$$

$$\frac{}{\llbracket x ::^{\mathcal{G}} \omega \rrbracket_{C;n} \equiv (x, - :: \omega)} \text{ (T-FPAT-ANN)}$$

Fig. 10. Translation of field patterns

In Fig. 10 we conclude the translation from FC^+ to FC by defining the translation of field patterns, being the language extension itself. Since the conversion function that inserts type information is specific to a program, we have to verify if the current field pattern is annotated somewhere in the program. Rule T-FPAT-PAT-1 states that if a field is never annotated, then it does not need to be translated. Otherwise, the additional information is discarded, as defined by Rule T-FPAT-PAT-2. The core of the translation is captured by Rule T-FPAT-ANN. A GADT annotation is erased by translating it to a dynamic type annotation, yielding a pair that matches the original value and the type that is stored in the dynamic.

5 Related Work

The foundations of structured programming on GADTs [7] provide an elegant approach to defining algebras on GADTs. While such algebras provide an abstraction mechanism to define an update function, explicit type representations and equality types [3,5] are still required. In Section 3.2, we discussed the disadvantages of such an approach. In our work, type representations and type equality proofs are implicitly provided by dynamics, which significantly improves the elegance of the function definitions.

Another approach to heterogeneous structures reflects the structure of a value directly in its type [8]. For example, the type of a heterogeneous list is basically a structure of nested tuples. Then, functions are defined on such structures using the type class mechanism, dispatching on the type structure. To enforce type-safe updates, yet another type class is defined to reflect type equality. Consequently, this approach results in rather verbose definitions since all action takes place on the level of type classes. Since the structure of the types are available, direct manipulation enables type-changing functions. Looking at the type of the update function in Section 3.3, our approach seems to forbid any type-changing updates. However, subterms can be replaced by arbitrary complex terms, thereby changing the underlying type structure.

6 Conclusion

We have presented the typical synergy between dynamics and GADTs to elegantly define functions that manipulate GADTs, requiring instantiation information on polymorphic types. Our approach comprises a new GADT annotation and improves upon boilerplate type representation administration in conventional approaches, since the functions are no longer cluttered with type equality witnesses and manual proofs. Also, by using dynamics, the need to maintain a closed enumeration of the used types is eliminated. Above all, our approach scales up to more complex structures and functions due to its simplicity. We have shown that the language extension is easily translated to a functional core that supports both dynamics and GADTs.

One of the major limitations in our approach is that the use of type codes limits the use of the GADT annotation to non-abstract types. It remains as future work to define type codes for such types, as well as investigating if dynamics can be implemented without type codes as class constraints. This would improve our approach considerably since it will not require us to decorate GADTs beforehand with type code constraints. Also, we plan to verify our hypothesis that storing GADTs in dynamics is no different from conventional ADTs.

Despite these limitations, the translation to dynamics provides novel opportunities, such as type dispatching and enforcing type equality invariants on GADTs. These opportunities require a more intricate translation than described in this paper, since this class of functions projects values instead of manipulating the values as such.

Acknowledgements

The authors would like to thank the anonymous reviewers for their helpful comments and suggestions. This work has been funded by the Technology Foundation STW through its project on “Demand Driven Workflow Systems” (07729).

References

1. Abadi, M., Cardelli, L., Pierce, B., Plotkin, G.: Dynamic typing in a statically typed language. *ACM Transactions on Programming Languages and Systems* 13(2), 237–268 (1991)
2. Abadi, M., Cardelli, L., Pierce, B., Rémy, D., Taylor, R.: Dynamic typing in polymorphic languages. *Journal of Functional Programming* 5(1), 81–110 (1994)
3. Baars, A., Swierstra, D.: Typing dynamic typing. In: Peyton Jones, S. (ed.) *Proceedings of the 7th International Conference on Functional Programming, ICFP 2002, Pittsburgh, PA, USA*, pp. 157–166. ACM, New York (2002)
4. Cartwright, R., Donahue, J.: The semantics of lazy (and industrious) evaluation. In: *Proceedings of the 2nd Symposium on LISP and Functional Programming, LFP 1982, Pittsburgh, PA, USA*, pp. 253–264. ACM, New York (1982)
5. Cheney, J., Hinze, R.: A lightweight implementation of generics and dynamics. In: Chakravarty, M. (ed.) *Proceedings of the 6th Haskell Workshop, Haskell '02, Pittsburgh, PA, USA*, pp. 90–104. ACM, New York (2002)

6. Cheney, J., Hinze, R.: First-class phantom types. Technical Report TR2003-1901, Cornell University (2003)
7. Johann, P., Ghani, N.: Foundations for structured programming with GADTs. In: Necula, G., Wadler, P. (eds.) Proceedings of the 35th Symposium on Principles of Programming Languages, POPL 2008, San Francisco, CA, USA, pp. 297–308. ACM, New York (2008)
8. Kiselyov, O., Lämmel, R., Schupke, K.: Strongly typed heterogeneous collections. In: Nilsson, H. (ed.) Proceedings of the 8th Haskell Workshop, Haskell 2004, Snowbird, UT, USA, pp. 96–107. ACM, New York (2004)
9. Läufer, K., Odersky, M.: Polymorphic type inference and abstract data types. ACM Transactions on Programming Languages and Systems 16(5), 1411–1430 (1994)
10. Peyton Jones, S.L., Vytiniotis, D., Weirich, S., Washburn, G.: Simple unification-based type inference for GADTs. In: Lawall, J. (ed.) Proceedings of the 11th International Conference on Functional Programming, ICFP 2006, Portland, Oregon, USA, pp. 50–61. ACM, New York (2006)
11. Pil, M.: Dynamic types and type dependent functions. In: Koopman, P., Clack, C. (eds.) IFL 1999. LNCS, vol. 1868, pp. 169–185. Springer, Heidelberg (2000)
12. Schrijvers, T., Peyton Jones, S., Sulzmann, M., Vytiniotis, D.: Complete and decidable type inference for GADTs. In: Hutton, G., Tolmach, A. (eds.) Proceedings of the 14th International Conference on Functional Programming, ICFP 2009, Edinburgh, Scotland, pp. 341–352. ACM, New York (2009)
13. Vervoort, M., Plasmeijer, R.: Lazy dynamic input/output in the lazy functional language Clean. In: Trinder, P., Michaelson, G.J., Peña, R. (eds.) IFL 2003. LNCS, vol. 3145, pp. 101–117. Springer, Heidelberg (2004)
14. Xi, H., Chen, C., Chen, G.: Guarded recursive datatype constructors. In: Morrisett, G., Aiken, A. (eds.) Proceedings of the 30th Symposium on Principles of Programming Languages, POPL 2003, New Orleans, LA, USA, pp. 224–235. ACM, New York (2003)