

Axiom Groups in FC

Richard Eisenberg

August 20, 2012

1 Introduction

Currently, the consistency of an FC context Γ requires that no two axioms (created from Haskell type families) overlap. We wish to lift that restriction by defining a new form of axiom that encapsulates an ordered list of potentially overlapping axioms. Because an old-style single axiom can be seen as a degenerate case of an axiom group, this replaces the old notion of an axiom.

2 Motivation

The primary motivation for overlapping axioms is to be able to define Boolean equality at the type level without needing a quadratic number of instance declarations. We would like the following to work:

```
type family (a :: k) ==: (b :: k) :: Bool
type instance where
  a ==: a = True
  a ==: b = False
```

The concrete syntax is something Dimitris, Simon, and I cooked up and is subject to change.

3 Old axioms

To make a meaningful comparison, it is helpful to review the current state (as proposed in the *Down with kinds* paper) of axioms.

The *old* form of an axiom binding is

$$C:\forall \Delta. (F \bar{\tau} \sim \sigma)$$

The variables in the telescope Δ are the free type, kind, and coercion variables taken from the $\bar{\tau}$. Different axioms arising from the same type family might have Δ s of different sizes. If we imagine the code above to be rewritten with separate instances, the first instance would have a Δ of $\kappa:\star, \alpha:\kappa$ and the second would have a Δ of $\kappa:\star, \alpha:\kappa, \beta:\kappa$. The different arities are acceptable because the type inference engine figures out which axiom to apply and creates the FC term with the appropriate axiom and the appropriate number of applications. Upon application to an appropriate Θ such that $\Gamma \vdash_{\text{tc}} \Delta \rightsquigarrow \Theta$, the axiom becomes a coercion according to the following rule:

$$\frac{C:\forall \Delta. (F \bar{\tau} \sim \sigma) \in \Gamma \quad \Gamma \vdash_{\text{tc}} \Delta \rightsquigarrow \Theta}{\Gamma \vdash_{\text{co}} C \Theta : \Theta_1(F \bar{\tau}) \sim \Theta_2(\sigma)} \quad \text{CT_VARAX}$$

The type of the coercion $C \Theta$ is the declared type of C with the appropriate substitutions taken out of Θ .

(Θ is a mapping from type variables to triples (τ_1, τ_2, γ) where $\Gamma \vdash_{\text{co}} \gamma : \tau_1 \sim \tau_2$ (modulo internal dependencies). $\Theta_1(\cdot)$ substitutes the corresponding τ_1 and $\Theta_2(\cdot)$ substitutes the corresponding τ_2 .)

4 New axiom groups

The *new* form is

$$C:F \overline{\{\Delta\} \bar{\tau}^n \rightsquigarrow \sigma^m}$$

This type contains a list of m equations. The name of the type family, F , must be the same for each equation and thus is included outside the list. Each equation lists its free type, kind, and coercion variables in Δ . Then comes the list of type patterns $\bar{\tau}$, the left-hand side of one of the equations, and σ , the right-hand side of that equation. We require each equation to have the same number, n , of type patterns. (The value of n does *not* depend on which equation we are looking at.)

This axiom C will be applied to a list of n coercions. One might think that this should be applied to n types, or perhaps n types-and-coercions, but each parameter is a coercion, due to the implementation of coercion optimization within GHC. The idea is that the left-hand sides of the coercions' types will match up with the type patterns $\bar{\tau}$. In fact, even with the new proposed syntax in *Down with kinds*, one can think of the application to the telescoped coercion Θ to be an application to a list of coercions stored in the Θ .

The intuition here is that we wish to search through the list of equations, in order, looking for an equation that matches. When we find one, we can type $C \bar{\gamma}^n$ according to the match. Here is the well-formedness and coercion typing rules:

$$\frac{\begin{array}{l} C \# \Gamma \\ \text{for each } i \leq m \\ \Gamma, \Delta_i \vdash_{\text{ty}} F \bar{\tau}_i^n \sim \sigma_i : \star \end{array}}{\vdash_{\text{wf}} \Gamma, C:F \overline{\{\Delta\} \bar{\tau}^n \rightsquigarrow \sigma^m}} \quad \text{GWF_AXIOMGROUP}$$

$$\frac{\begin{array}{l} C:F \overline{\{\Delta\} \bar{\tau}^n \rightsquigarrow \sigma^m} \in \Gamma \\ \Gamma \vdash_{\text{co}} \gamma : \tau_1 \sim \tau_2 \\ \Gamma \vdash_{\text{ty}} F \bar{\tau}_1^n : \kappa \\ (i, \Omega) =_{\Gamma} \text{select}_1 \bar{\gamma}^n \overline{\{\Delta\} \bar{\tau}^n}^m \end{array}}{\Gamma \vdash_{\text{co}} C \bar{\gamma}^n : \Omega_L(F \bar{\tau}_i) \sim \Omega_R(\sigma_i)} \quad \text{CT_AXIOMGROUP}$$

The *select* function (defined below) does the work of choosing which equation matches, returning the number of the equation i and a substitution Ω . This Ω is a mapping from type and coercion variables to coercions (which will be built from the $\bar{\gamma}^n$ passed to C). The substitution $\Omega_L(\cdot)$ substitutes variables with the left-hand side of the type of the mapped coercion; $\Omega_R(\cdot)$ substitutes with the right-hand side. (If Ω contains $\alpha \mapsto \gamma$ and $\Gamma \vdash_{\text{co}} \gamma : \tau_1 \sim \tau_2$ (modulo internal dependencies), $\Omega_L(\alpha) = \tau_1$ and $\Omega_R(\alpha) = \tau_2$.)

The 1 subscript to the *select* operation says to start the search at the first equation.

Here is the definition of *select*:

$$\boxed{(i, \Omega) =_{\Gamma} \text{select}_j \bar{\gamma}^n \overline{\{\Delta\} \bar{\tau}^n}^m} \quad \text{The select algorithm}$$

$$\frac{\begin{array}{l} j \leq m \\ i = j \\ \Omega =_{\Gamma} \text{match} \bar{\gamma}^n \{\Delta_i\} \bar{\tau}_i^n \\ \forall \Omega'. \forall (r < i). \text{fail} =_{\Gamma} \text{match} \overline{\Omega'(\bar{\gamma})}^n \{\Delta_r\} \bar{\tau}_r^n \end{array}}{(i, \Omega) =_{\Gamma} \text{select}_j \bar{\gamma}^n \overline{\{\Delta\} \bar{\tau}^n}^m} \quad \text{SELECT_MATCH}$$

$$\frac{\begin{array}{l} j' = j + 1 \\ i = j \\ \text{fail} =_{\Gamma} \text{match} \bar{\gamma}^n \{\Delta_i\} \bar{\tau}_i^n \\ (r, \Omega) =_{\Gamma} \text{select}_{j'} \bar{\gamma}^n \overline{\{\Delta\} \bar{\tau}^n}^m \end{array}}{(r, \Omega) =_{\Gamma} \text{select}_j \bar{\gamma}^n \overline{\{\Delta\} \bar{\tau}^n}^m} \quad \text{SELECT_NOMATCH}$$

Unsurprisingly, *select* iterates through the list of equations and returns the result from the one that matches, delegating the match operation to *match*. (The various indexing operations are just to get things to work with ott. A more formal presentation would fix that up.)

The only interesting piece to this judgment is the last premise to the first rule. It asserts that, for all possible substitutions Ω' , all previous equations will not match. At first, this seems redundant, because we start matching at equation #1 and stop at the first match. However, it is important to note that the domain of Ω' is *not* the domain of Δ . This premise is stating that for any possible assignment of type variables in $\bar{\gamma}^n$ (*not* the type variables in $\bar{\tau}$), the match will fail. Another way of saying this is that the left-hand types in the types of $\bar{\gamma}^n$ cannot unify with the $\bar{\tau}$. Say a coercion is $\langle \alpha \rangle$ for some type variable $\alpha \in \text{dom } \Gamma$. According to *match* (defined below), $\langle \alpha \rangle$ will not match with **Int**. However, maybe α will be substituted with **Int** at some point. Because of this possibility, we must make sure that before committing to a match, no possible match could have worked no matter what happens to the variables in the $\bar{\gamma}^n$. In this case, the overall coercion will be ill-typed, because there is no unambiguous match.

Here is the match relation:

$$\boxed{\Omega =_{\Gamma} \text{match } \bar{\gamma}^n \{ \Delta \} \bar{\tau}^n} \quad \text{Coercion list matching}$$

$$\frac{\frac{\frac{\Gamma \vdash_{\text{co}} \gamma : \sigma_1 \sim \sigma_2^n}{\Omega_L(\tau) = \sigma_1^n}}{\text{dom } \Omega = \text{dom } \Delta}}{\Omega =_{\Gamma} \text{match } \bar{\gamma}^n \{ \Delta \} \bar{\tau}^n} \quad \text{MATCHI_MATCH}$$

For a Ω such that the domain of Ω matches the free variables in the $\bar{\tau}$, check if the substitution returns the desired types, and succeed if so. Unfortunately, this judgment gives us no way of finding Ω . So, we use the algorithmic version presented below, with all its auxiliary judgments:

$$\boxed{\Omega =_{\Gamma} \text{match } \bar{\gamma}^n \{ \Delta \} \bar{\tau}^n} \quad \text{Coercion list matching (algorithmic)}$$

$$\frac{n = 0}{\epsilon =_{\Gamma} \text{match } \bar{\gamma}^n \{ \Delta \} \bar{\tau}^n} \quad \text{MATCH_NIL}$$

$$\frac{\begin{array}{l} \bar{\gamma}^n = \gamma, \bar{\gamma}'^m \\ \bar{\tau}^n = \tau, \bar{\tau}'^m \\ n = m + 1 \\ \Delta' \text{ is the smallest prefix of } \Delta \text{ s.t. } \Gamma, \Delta' \vdash_{\text{ty}} \tau : \kappa \\ \Omega_1 =_{\Gamma} \text{match } \bar{\gamma}'^m \{ \Delta \setminus \Delta' \} \bar{\Omega}_{2L}(\bar{\tau}')^m \\ \Gamma \vdash \gamma \overset{\Delta'}{\leftrightarrow} \tau : \Omega_2 \\ \Omega = \Omega_1, \Omega_2 \end{array}}{\Omega =_{\Gamma} \text{match } \bar{\gamma}^n \{ \Delta \} \bar{\tau}^n} \quad \text{MATCH_CONS}$$

The “smallest prefix” condition is there to select out the prefix of Δ that contains all of the free variables in τ .

$$\boxed{\Gamma \vdash \gamma \overset{\Delta}{\leftrightarrow} \tau : \Omega} \quad \text{Coercion / type pattern match}$$

$$\frac{\Gamma \vdash_{\text{co}} \gamma : \tau \sim \sigma}{\Gamma \vdash \gamma \overset{\epsilon}{\leftrightarrow} \tau : \epsilon} \quad \text{MATCH1_EXACT}$$

$$\frac{\Gamma \vdash_{\text{co}} \gamma : \phi}{\Gamma \vdash \gamma \overset{\beta:\kappa}{\leftrightarrow} \beta : \beta \mapsto \gamma} \quad \text{MATCH1_VAR}$$

$$\begin{array}{c}
\Gamma \vdash_{\text{co}} \gamma : \sigma \overline{\sigma_1}^n \sim \sigma \overline{\sigma_2}^n \\
\Delta_0 \text{ is the smallest prefix of } \Delta \text{ s.t. } \Gamma, \Delta_0 \vdash_{\text{ty}} \tau : \kappa \\
\Gamma \vdash \langle \sigma \rangle \overset{\Delta_0}{\rightsquigarrow} \tau : \Omega_0 \\
\text{for each } \tau'_i \in \overline{\tau'}^n \\
\quad i = r + 1 \\
\quad \Delta_i \text{ is the smallest prefix of } \Delta \setminus \Delta_r \text{ s.t. } \Gamma, \Delta_i \vdash_{\text{ty}} \Omega_{rL}(\tau'_i) : \kappa_i \\
\quad \Gamma \vdash \mathbf{nth}^i \gamma \overset{\Delta_i}{\rightsquigarrow} \Omega_{rL}(\tau'_i) : \Omega'_i \\
\quad \Omega_i = \Omega_r, \Omega'_i \\
\Omega = \Omega_n \\
\tau \text{ is not a type application} \\
\hline
\Gamma \vdash \gamma \overset{\Delta}{\rightsquigarrow} \tau \overline{\tau'}^n : \Omega \qquad \text{MATCH1_APP}
\end{array}$$

In the last rule, we go left-to-right, carefully building up a substitution Ω as we go. Once again, some care must be taken to get the Δ right so that the following lemma holds:

Lemma 1 (Substitution domain). *If $\Gamma \vdash \gamma \overset{\Delta}{\rightsquigarrow} \tau : \Omega$ then the domain of Ω is the domain of Δ .*

Acknowledgments

Thanks to José Pedro Magalhães for starting this design process. Thanks also to Dimitrios Vytiniotis and Simon Peyton Jones for helping me start this project.