

Type Invariants for Haskell

Tom Schrijvers *

Katholieke Universiteit Leuven, Belgium
tom.schrijvers@cs.kuleuven.be

Louis-Julien Guillemette

Université de Montréal, Canada
guillelj@iro.umontreal.ca

Stefan Monnier

Université de Montréal, Canada
monnier@iro.umontreal.ca

Abstract

Haskell’s multi-parameter type classes, together with functional dependencies, allow the specification of complex type-level operations, and the recent introduction of open type families in GHC makes such type-level programming even more accessible and flexible. But type-level code is special in that its correctness is crucial to the safety of the program; so except in those cases simple enough for the type checker to see trivially that the code is correct (or harmless), type-level programs need to come with their specification and correctness proof.

In this article, we propose an extension to Haskell that allows the specification of invariants for type classes and open type families, together with accompanying evidence that those invariants hold. To accommodate the open nature of type classes and type families, the evidence itself needs to be open and every subcase of the proof can be provided independently from the others.

Categories and Subject Descriptors D.3.2 [Programming Languages]: Language Classifications—Functional Languages; F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs—Type Structure

General Terms Algorithms, Languages

Keywords Haskell, type checking, type functions, type families

1. Introduction

Haskell’s multi-parameter type classes (Duggan and Ophel 2002), together with functional dependencies (Jones 2000), allow the specification of complex type-level operations, and the recent introduction of open type families in GHC makes such type-level programming even more accessible and flexible. But type-level code is special in that its correctness is crucial to the safety of the program; In surprisingly many cases the type checker’s limited reasoning power is actually sufficient to verify completely automatically that the type-level code is correct. But in the general case, the programmer needs to help the type checker by instructing it to exploit some particular properties of the types manipulated.

As it happens, GHC’s type system is already sufficiently powerful that in most if not all cases, the code can be rewritten in such a way that the type checker can check the correctness of the type

annotations, even for fairly complex cases such as when proving that a compiler is type preserving (Guillemette and Monnier 2008). But such rewrites are unsatisfactory because they tend to either incur substantial run-time cost, or force wide-reaching changes which break the modularity of the code, or both.

So we basically want our type language to include a proof assistant. One alternative is to retrofit a system such as the Calculus of Inductive Constructions (Paulin-Mohring 1993) or Twelf (Pfenning and Schürmann 1999) into GHC’s type language, but since GHC’s type system is already fairly complex and powerful we decided instead to try to limit ourselves to a small extension of its type language, so as to make good use of the existing machinery. One of the side benefits is that the system arguably integrates better with the rest of Haskell and should hopefully be more palatable for Haskell coders.

Concretely, we propose an extension of GHC’s type system which allows the programmer to specify on the one hand some type invariants, and on the other the corresponding proof. GHC already supports some forms of type invariants: a type class declaration can specify that all instances of this class should also be instances of some other classes, and for multi-parameter classes it can also specify functional dependencies between the parameters. Our type invariants subsume those two particular cases, except for the fact that they require more annotations: being more general, not only does the programmer need to provide explicit proofs of those invariants, but she additionally needs to explicitly specify when the invariant is used. To this end, each invariant provides a corresponding coercion function.

Another difference with the functional dependencies and the class context invariants, is that our type invariants can be specified separately from any class declaration, and can also apply to a combination of several classes or type families. This makes it possible to retroactively specify that all `Num` instances are also instances of `Additive`, without having to change the `Num` class.

The specific contributions of this paper are:

- We clearly motivate the need for type invariants with examples (Section 2).
- We present type invariants for Haskell, and explain and illustrate the choices made in our approach (Section 3).
- The formalization of our approach (Section 4) respects as much as possible the open and modular nature of Haskell’s type classes and type families.
- We provide an external proof language that infers proof steps and greatly simplifies writing proofs (Section 5.1).
- As experimental evaluation, we have written and checked a number of invariants and their proofs (Section 5.2).
- A prototype implementation in Haskell is available at <http://www.cs.kuleuven.be/~toms/Haskell/>.

* Post-doctoral researcher of the Fund for Scientific Research - Flanders.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

At the end of this paper, we discuss related work (Section 6) and conclude (Section 7).

2. Motivation and Examples

There are many examples where type families would be much more useful if they would support additional invariants beyond the basic axioms, i.e. the type family instances.

2.1 Indexed List Processing

Parity Indexing Ki Yung Ahn proposed the following list GADT on the Haskell mailing list, which is *parity*-indexed, i.e. indexed by whether the length is odd or even:

```
data Even
data Odd

data List a l where
  Nil :: List a Even
  Cons :: a -> List a l -> List a (Flip l)
```

For instance, `Cons True Nil :: List Bool Odd` and `Cons False (Cons True Nil) :: List Bool Even`. The type family `Flip` expresses what happens to the parity of the list length if we add an element:

```
type family Flip l

type instance Flip Even = Odd
type instance Flip Odd = Even
```

While some functions using these lists are easy to type check, e.g.,

```
headList :: List a l -> a
headList (Cons x xs) = x
```

it is unfortunately not possible to consider some other basic functions as well-typed:

```
tailList :: List a l -> List a (Flip l)
tailList (Cons x xs) = xs
```

From the GADT pattern match we know that `Flip k ~ l`, where `k` is the parity of `xs`. From the signature, we can also see that `Flip l ~ k`. Once we eliminate `k` from these equations, we end up with the need to prove the following constraint:

$$\forall l. \text{Flip (Flip l)} \sim l. \quad (1)$$

While we may think that this equation readily holds for any `l`, this is not so. Due to the openness of the type family `Flip`, one may add at any time an additional type instance, e.g.,

```
type instance Flip Char = Even
```

such that `Flip (Flip Char) = Odd` and Equation (1) does not hold for `Char`.

There are various solutions to the problem, such as cluttering either the type signature of `tailList` or the `List` constructors themselves with the required equality. Yet, this defers the equality to elsewhere in the program. It would be much neater if Equation (1) could simply be enforced on all instances of the type family, and that subsequently we could simply use Equation (1) wherever needed.

In (Ahn and Sheard 2008) Ki Yung Ahn decided to use an inelegant encoding of the indexing types that avoids type families.

Length Indexing Length-indexed lists are a classical example:

```
data Z
data S n

data List a l where
```

```
Nil :: List a Z
Cons :: a -> List a n -> List a (S n)
```

where we need a type family `Add` for expressing the signature of `append`:

```
append :: List a k -> List a l -> List a (Add k l)
append Nil ys = ys
append (Cons x xs) ys = Cons x (append xs ys)
```

```
type family Add m n
type instance Add Z n = n
type instance Add (S m) n = S (Add m n)
```

The same type family would also serve us for the signature of `merge`:

```
merge :: List a k -> List a l -> List a (Add k l)
merge Nil ys = ys
merge (Cons x xs) ys = Cons x (merge ys xs)
```

In the second clause, the GADT pattern match exposes the equality `k ~ S k'`, where `k'` is the length of `xs`. Now the type checker expects the type `Add k' l` for the expression `(merge ys xs)`. However, it infers, based on the signature of `merge`, the type `Add l k'`. Hence, for this code to type check, commutativity of `Add` must hold:

$$\forall k. \forall l. \text{Add k l} \sim \text{Add l k} \quad (2)$$

However, for the same reason as before, we cannot rely on this invariant to hold.

Invariants as term-level functions Without support for type invariants, it is still possible to implement the invariants at the term level, so as to make the above code examples type-check. We first need to reify coercions at the term level using a GADT which witnesses the equivalence of two types, as well as reify the types (such as the length annotations) using singleton types:

```
data Eqv s t where
  Eqv :: (s ~ t) => Eqv s t
```

```
data Nat n where
  Nz :: Nat Z
  Ns :: Nat s -> Nat (S s)
```

An invariant such as the commutativity of addition is then implemented as a function that analyzes the term-level representatives and constructs the proof:

```
comm :: Nat i -> Nat j -> Eqv (Add i j) (Add j i)
comm Nz Nz = Eqv
comm Nz (Ns j) = case addZ j of Eqv -> Eqv
comm (Ns i) Nz = case addZ i of Eqv -> Eqv
comm (Ns i) (Ns j) =
  case comm i j of
    Eqv -> case (comm (Ns i) j, comm i (Ns j)) of
      (Eqv, Eqv) -> Eqv
```

where `addZ` similarly implements an auxiliary invariant stating that:

$$\forall n. \text{Add n Z} \sim n. \quad (3)$$

To apply the invariant and have the `merge` function type-check, the length of the lists must be computed separately:

```
merge (Cons x xs) ys =
  case comm (length xs) (length ys) of
    Eqv -> Cons x (merge ys xs)
```

```
length :: List a l -> Nat l
length = ...
```

Of course, implementing invariants in this way has a cost at run-time, and since Haskell does not enforce that such a “proof” covers all the cases and will not loop indefinitely, we would need to rely on an external verifier to check the totality of that function.

2.2 Type Preserving Compilation

In Guillemette and Monnier’s type preserving compiler (2008), an abstract syntax tree for System F is type-indexed with the source-level types. Type functions are used to relate the source-level types before and after various program transformations, i.e. continuation-passing-style (CPS) transformation and type. One of the invariants that comes up in this context is the commutativity of CPS transformation and type substitution:

$$\forall t, x, s. \text{CPS} (\text{Subst } t \ x \ s) \sim \text{Subst} (\text{CPS } t) \ x \ (\text{CPS } s)$$

where $\text{CPS } t$ reflects the type of an expression of source type t after CPS transformation, and $\text{Subst } t \ x \ s$ is the type after substitution of the type variable x with type s .

3. Outline of Our Approach

In this section, we outline our approach, what works and what doesn’t, and a number of complications that arise.

3.1 Domain Restriction

It is quite clear early on that invariants are not intended to be instantiated at just any types. For instance, we do not wish the invariant (2) to be instantiated to $\text{Add } Z \ \text{Char} \sim \text{Add } \text{Char} \ Z$, as there is not even an instance of Add that tells us what to do with $\text{Add } \text{Char} \ Z$.

An extreme case is a type family EMPTY without any instance at all. Since there is no instance, we could readily assume that the two following invariants hold:

$$\begin{aligned} \forall x. \text{EMPTY } x \sim \text{Int} \\ \forall x. \text{EMPTY } x \sim \text{Bool} \end{aligned} \quad (4)$$

In the absence of any instance, it is vacuously true that all instances of EMPTY satisfy the above two invariants. Yet, there is no mechanism to restrict EMPTY from being used with types for which there is no instance. So things can go very wrong if we do, as we can then derive $\text{Int} \sim \text{EMPTY } \text{Char} \sim \text{Bool}$. Type soundness is at stake!

We need a mechanism to restrict the use of invariants to *appropriate* types only. Perhaps the most obvious approach is to use a richer kinding system than the one Haskell offers where all proper types are of the same kind $*$. A richer kind system could allow the user to classify types into distinct kinds. For instance, analogous to Haskell data type definitions, we could specify that type-level natural numbers are of kind Nat :

```
datakind Nat = Z | S Nat
```

Then we could restrict the use of type variables in invariants to the appropriate kind:

$$\forall k : \text{Nat}. \forall l : \text{Nat}. \text{Add } k \ l \sim \text{Add } l \ k \quad (5)$$

The above `datakind` has two disadvantages: Firstly, it provides a closed definition of kinds whereas type classes and type families are open. Secondly and more importantly, it requires extending Haskell with another type system feature.

Instead we propose a simpler solution that uses an existing Haskell feature for classifying types and that is open, namely type classes. Instead of a `datakind Nat`, we propose to use a type class `Nat`:

```
class Nat n
instance Nat Z
instance Nat n => Nat (S n)
```

Note that the type class does not need to have any methods; we are only interested in its type-level aspects. Using this type class, we can add the same style of constraint context to the invariant with which we are already familiar in function signatures.

$$\forall k. \forall l. (\text{Nat } k, \text{Nat } l) \Rightarrow \text{Add } k \ l \sim \text{Add } l \ k \quad (6)$$

3.2 Using Invariants

Due to their generality, invariants are unfortunately inappropriate for automatic use by the type checker algorithm. Firstly, they cannot be treated as type family instances, also called top-level equations in (Schrijvers et al. 2008), because the left-hand side is not necessarily in the proper *constructor form* to ensure either termination or confluence. Secondly, they cannot be treated as local equations, as provided in a function signature, because they contain schema (i.e. universally quantified) variables. In other words, and unsurprisingly, the completion algorithm of type families cannot be used to transform the invariants into a terminating and confluent rewrite system.

Because of the above issues, we propose that the programmer explicitly indicates when an invariant should be used and at what types. For this purpose the concrete syntax for invariant definitions introduces a name for the invariant. This then introduces a coercion function of the same name, which can be used by the programmer to apply the invariant where it is needed.

Example 1. The definition of the commutativity invariant in concrete syntax is:

```
type invariant add_comm =
  (Nat x, Nat y) => Add x y ~ Add y x
```

In return for this specification, we get a coercion function which has the signature:

```
add_comm :: (Nat x, Nat y)
=> (Add x y ~ Add y x => a)
-> a
```

It is used as follows in the merge example:

```
merge (Cons x xs) ys =
  Cons x (add_comm (merge ys xs))
```

□

3.3 Proving Invariants

While type family instances are really the programmer-provided axioms of type families, invariants should be derivable from those axioms. A invariant which is not derivable is not a proper invariant, and may introduce unsoundness. Hence, to make sure an invariant actually holds, we require a proof.

Because of the proof obligation for invariants, we also call them *lemmas* throughout the rest of the paper. However, we must remember that invariants differ from lemmas on the account of openness. Whereas lemmas usually prove an existing property, invariants also impose a property on future instances.

Case-based proofs Because of the open and modular nature of type classes and type families, we expect our proofs to be open and modular as well. Modularity of a proof means that it should be based on case analysis where each case can be written separately. Openness means that when new instances are added, then also new cases of the proof can be added. Of course, in a (closed) program, the type checker should verify whether all cases are covered.

As we have already introduced type classes to restrict the instantiation of type variables in invariants, we will also use the type class instances to determine our proof cases.

Example 2. For the commutativity invariant we have the following four proof cases based on the cartesian product of the `Nat` type class instances:

```

proofcase add_comm Z      Z      = ...
proofcase add_comm Z      (S m) = ...
proofcase add_comm (S n) Z      = ...
proofcase add_comm (S n) (S m) = ...

```

□

Coercions as proofs Each proof case must provide evidence that a particular instance of the lemma holds. The consequent of the lemma being a type equivalence, we need to prove the two types equal. System F_C (Sulzmann et al. 2007a), the intermediate language of GHC, provides us with exactly such evidence, called type equality coercions. Hence, we simply reuse these, and so the body of a proof case consists of a coercion.

Example 3. The full first proof case of the commutativity axiom looks as follows:

```
proofcase add_comm Z Z = refl (Add Z Z)
```

This case concerns the instance of the axiom with two Zs:

$$\text{Add } Z \ Z \sim \text{Add } Z \ Z \quad (7)$$

This equation holds because the left- and right-hand sides are identical. The `refl` coercion constructor is evidence for that, a type being equal to itself. □

Example 4. Recall the parity invariant:

```

type invariant parity p =
  Parity p => Flip (Flip p) ~ p

```

where

```

class Parity p

instance Parity Odd
instance Parity Even

```

The proof case for `Odd` looks like:

```

proofcase parity Odd =
  trans (app (refl Flip) flip_odd)
        flip_even

```

The two `Flip` type family instances, denoted respectively `flip_odd` and `flip_even`, are used here as axioms. The proof constructor `app` composes the two proofs for `Flip ~ Flip` and `Flip Odd ~ Even` into a proof for `Flip (Flip Odd) ~ Flip Even`. Finally, the `trans` constructor composes the latter proof with `Flip Even ~ Odd` into the overall proof `Flip (Flip Odd) ~ Odd`. □

3.4 Invariant Coercions

Many proofs over inductively defined types, such as the natural numbers, require induction themselves. For instance, in order to show the commutativity invariant for the second case, Z and $S \ m$, we rely on the invariant to hold for m and Z . Hence, we must extend the coercion language of System F_C with induction steps, which we represent by coercions that simply refer to some invariant. With these lemma coercions we can build both inductive proofs and proofs involving auxiliary lemmas.

Example 5. The above mentioned proof case of the commutativity lemma is:

```

proofcase add_comm (S n) Z =
  trans
    (trans
      (addS n Z)
      (app
        (refl S)
        (trans
          (sym (add_comm Z n)) -- induction!

```

```

      (addZ n))))
    (sym (addZ (S n)))

```

This proof contains the inductive argument `add_comm Z n`. □

Well-Founded Induction The main concern with inductive proofs, using the invariant directly or indirectly in the proof, is whether the induction is well-founded. In order for it to be well-founded, the use of the invariant in a proof case should be strictly smaller than the case itself, according to some norm.

Example 6. In the previous example, the following norm can be used to establish the well-foundedness of induction:

$$\begin{aligned} |\text{add_comm } x \ y| &= |x| + |y| \\ |Z| &= 0 \\ |S \ x| &= 1 + |x| \end{aligned}$$

So we have that:

$$\begin{aligned} |\text{add_comm } (S \ n) \ Z| &= 1 + |n| \\ > |\text{add_comm } Z \ n| &= |n| \end{aligned}$$

i.e., there is a proper decrease. □

We propose to use the same norm for checking well-foundedness of invariants as is used to enforce termination of type families. While this choice is restrictive, Haskell programmers are already familiar with it from type families, and it has not posed any difficulties for the proofs we have constructed so far.

A final issue is how to identify recursive uses of lemmas, and hence when to enforce the well-foundedness criterion. Using it in all cases has proven to be overly restrictive. Hence, we require the uses of lemma coercions to be annotated with whether they are non-recursive or (potentially) recursive. In the latter case, the well-foundedness criterion is (conservatively) enforced. In the former case, we have to check instead whether the annotation is justified. This requires a global inspection of the lemma dependency graph of a program, i.e. a non-modular check, to ensure that this use of a lemma indeed is not part of any cycle.

Example 7. In Example 5, the inductive argument is actually written as `rec add_comm Z n`, where `rec` indicates induction. □

Context Proofs A second complication of invariant coercions is posed by the invariant context. A invariant may only be used if its context is provided. Hence, the use of an invariant coercion must be justified by a proof that the context is indeed satisfied. This requires a second evidence language, not for equality constraints, but for type class constraints.

This time we do not have any pre-existing adequate representation available in System F_C : type class constraints are desugared to value-level dictionaries. The proof language is present only implicitly as value-level functions for dictionary construction and super class dictionary extraction. Hence, we propose a new language proof modeled after these value-level functions. It contains additional functionality, such as instance context selectors which are usually not available for dictionaries (although such selectors have been proposed (Schrijvers and Sulzmann 2008)) but do make sense at the type level (*).

Example 8. Further extending the above example, we must provide a proof of the context for the inductive argument `add_comm Z n`. In particular, we must show that both `Nat Z` and `Nat n` hold.

A proof of the former is `env (Nat Z)`, which says the proof can be found in the proof case's context. Indeed, `Nat Z` is a precondition of the proof case.

A proof of the latter is `isel 1 (env (Nat (S n)))`. Here, we know that `Nat (S N)` holds from the proof case's context. From this we can show that `Nat n` holds based on the type class instance `instance Nat n => Nat (S n)`. The proof constructor `isel i` says to look for a matching type class instance, and to take the i th type class constraint from this instance's context. In this case $i = 1$, and `isel 1` yields the desired proof for `Nat n`.

In summary, we write the inductive argument with the context proofs explicit as:

```
rec add_comm (env (Nat Z),
              isel 1 (env (Nat (S n))))
  Z n
```

□

3.5 Invariants with Class

It turns out that equational invariants are not enough. As (equational) invariant proofs require invariant coercions, and invariant coercions require type class proofs, some type class proofs also require (type) class invariants.

Example 9. The proof of the axiom for the type-preserving compiler involves among others the following auxiliary invariant:

$$\forall k. \forall t. (\text{Nat } k, \text{Type } t) \Rightarrow U \ Z \ k \ t \sim t \quad (8)$$

Here, U is a type family and $U \ i \ k \ t$ increments all De Bruijn indices no smaller than k by i in type t . The invariant expresses that if the increment is zero, then the update has no effect on the type.

This auxiliary invariant is instantiated in another proof to:

$$\forall t. \text{Type } t \Rightarrow U \ Z \ Z \ (\text{CPS } t) \sim \text{CPS } t \quad (9)$$

The use of this invariant is only justified if we can provide evidence for its context $(\text{Nat } Z, \text{Type } (\text{CPS } t))$. While the former component is trivial to show, we can only show the latter with the help of an auxiliary type class invariant of the form:

$$\forall t. \text{Type } t \Rightarrow \text{Type } (\text{CPS } t) \quad (10)$$

□

Type class invariants are also useful on their own: for enforcing dependencies between type class instances.

Example 10. For instance, we can retroactively specify that all `Num` instances are also instances of `Additive`, without having to change the `Num` class. This invariant is expressed as:

$$\forall t. \text{Num } t \Rightarrow \text{Additive } t \quad (11)$$

□

It turns out that type class invariants are in all aspects similar to equational invariants: they require proof cases based on the different possible instantiations of the context, in terms of the proof language introduced in the previous paragraph.

There is one notable difference between those two kinds of invariants: equality coercions only exist at the level of types and are completely erased at run-time, whereas type classes are generally reified as dictionaries at run-time. This implies that for an equational invariant the value-level coercion function can be trivially implemented as a no-op; whereas for class invariants the value-level coercion function needs to build the corresponding run-time evidence (i.e. dictionary).

3.6 Proof Assistance

Since theorem proving is a non-trivial matter to automate, we assume for now that the programmer writes the proof. However, we expect that at least part of the task can be automated with proof assistants and automated theorem provers. In fact, we provide a limited form of proof support already, as the coercion language of System F_C is not the most convenient form for humans.

We propose a more convenient notation, which resembles equational reasoning proofs, where a proof consists of a sequence of types:

$$t_1 \sim t_2 \sim \dots \sim t_n$$

and represents a coercion for the equality $t_1 \sim t_n$.

Each pair of subsequent types t_i and t_{i+1} should be chosen such that a coercion from t_i to t_{i+1} only requires the use of a single coercion constructor¹. Our proof assistant reconstructs an actual coercion from this notation, and also attempts to fully reconstruct type class proofs involving no more than three constructors. In case the proof assistant fails, due to its incompleteness or due to unsoundness, it reports at which step, and the programmer may intervene.

Example 11. In the new sequential notation, the proof case of Example 5 is much more compact and readable:

```
proofcase add_comm (S n) Z =
  Plus (S n) Z
  ~ S (Plus n Z)
  ~ S (Plus Z n)    -- induction!
  ~ S n
  ~ Plus Z (S n)
```

□

4. Formalization of Invariants

In this section, we formalize the type-level fragment of our Haskell language extension.

4.1 Program Syntax

Figure 1 summarizes our meta-variable naming conventions for the various syntactical categories. A *type-level program*, denoted by $\text{prog}(\overline{D_{TC}}, \overline{I_{TF}}, \overline{I_{TC}}, \overline{h}, \overline{p}, \overline{\pi_{TF}}, \overline{\pi_{TC}})$, consists of:

- type class declarations $\overline{D_{TC}}$ of the form

$$\text{class } (C_1 \ t_1, \dots, C_n \ t_n) \Rightarrow C \ a$$
- type class instances $\overline{I_{TC}}$ of the form

$$\text{instance } (C_1 \ t_1, \dots, C_n \ t_n) \Rightarrow C \ t$$
- type family instances $\overline{I_{TF}}$ of the form

$$g \ \overline{x} : s_1 \sim s_2$$
- equational invariants \overline{h} of the form

$$h \ \overline{x} : (C_1 \ y_1, \dots, C_m \ y_m) \Rightarrow s_1 \sim s_2$$
 where $\overline{y} \subseteq \overline{x}$, $\text{vars}(s_1) \subseteq \overline{x}$ and $\text{vars}(s_2) \subseteq \overline{x}$,
- type class invariants \overline{p} of the form

$$p \ \overline{x} : (C_1 \ y_1, \dots, C_m \ y_m) \Rightarrow C \ t$$
 where $\overline{y} \subseteq \overline{x}$ and $\text{vars}(t) \subseteq \overline{x}$,
- proof cases $\overline{\pi_{TF}}$ for equational invariants of the form

$$\text{proofcase } h \ \overline{t} = \gamma$$
- proof cases $\overline{\pi_{TC}}$ for type class invariants of the form.

$$\text{proofcase } p \ \overline{t} = \delta$$

The syntax for proofs is listed in Figure 2. Its meaning is explained shortly.

For reasons of simplicity we have restricted ourselves to single-parameter type classes. We believe that multi-parameter type classes (Duggan and Ophel 2002) are a straightforward extension that adds little to our discussion. We have also omitted type family declarations, because these do not contribute any useful information.

¹ Some coercion constructors do not allow a sequential notation and are not allowed

s, t, u, v	types
γ	coercions
g	equational axioms
h	equational invariants
p	type class invariants
δ	type class evidence
θ	type variable substitution
I_{TF}	type family instance
D_{TC}	type class declarations
I_{TC}	type class instance
π_{TF}	equational invariant proof case
π_{TC}	type class invariant proof case

Figure 1. Meta-Variables

γ	::=	$\text{refl } t \mid g \bar{t} \mid \text{sym } \gamma \mid \gamma_1 \circ \gamma_2 \mid F \bar{\gamma} \mid \gamma_1 \gamma_2 \mid \text{decomp}_{T_i} \gamma$
		$\mid \text{norec } \bar{\delta} \Rightarrow h \bar{t} \mid \text{rec } \bar{\delta} \Rightarrow h \bar{t}$
δ	::=	$\text{env } C t \mid \text{inst } \bar{\delta} (C t) \mid \text{ssel } i \delta \mid \text{isel } i \delta$
		$\mid \text{co } (C t) \gamma \mid \text{norec } \bar{\delta} \Rightarrow p \bar{t} \mid \text{rec } \bar{\delta} \Rightarrow p \bar{t}$

Figure 2. Proof Syntax

For the confluence and termination conditions on both type class and type family instances we refer to respectively (Sulzmann et al. 2007b) and (Schrijvers et al. 2008).

4.2 Well-Typing

The main typing rule for type-level programs is:

$$\frac{\begin{array}{l} \overline{D_{TC} \cup I_{TF} \cup I_{TC} \cup \bar{h} \cup \bar{p}} \vdash h_{i_1} : \diamond \\ \overline{D_{TC} \cup I_{TF} \cup I_{TC} \cup \bar{h} \cup \bar{p}} \vdash p_{i_2} : \diamond \\ \overline{D_{TC} \cup I_{TF} \cup I_{TC} \cup \bar{h} \cup \bar{p}} \vdash \pi_{TF, i_3} : \diamond \\ \overline{D_{TC} \cup I_{TF} \cup I_{TC} \cup \bar{h} \cup \bar{p}} \vdash \pi_{TC, i_4} : \diamond \end{array}}{\vdash \text{prog}(\overline{D_{TC}, I_{TF}, I_{TC}, \bar{h}, \bar{p}, \pi_{TF}, \pi_{TC}}) : \diamond}$$

It checks the well-typedness of proof cases, and their complete coverage of invariants. The former well-typing rules are covered by Figures 3 and 4 for equational and type class proof cases respectively. A proof case is well-typed iff its proof term is well-typed and is a proof for the invariant instance that the proof case claims to cover. The proof terms for equational proof cases are the coercions of System FC , extended with invariant coercions (INVCO). The proof term language for type class invariants is new. Its primitives reflect the common value-level functions for dictionary construction and super class selectors used in the dictionary-passing implementation of type classes. Non-standard are the (ISELTC) rule reflects the instance selector function recently proposed in (Schrijvers and Sulzmann 2008), and the (COTC) rule for coercing a type class constraint from applying to one type to another.

The environment Γ consists of equational axioms and invariants, type class instances and invariants, as well as type class constraints available in the current context.

4.3 Completeness Check

The coverage relation \sqsubseteq checks whether one set of cases is covered by another set. It is used in LEMCO and LEMTC to check the completeness of the invariant proofs, i.e. whether all the required proof cases are supplied.

DEFINITION 1. We define \sqsubseteq as follows:

$$S_1 \sqsubseteq S_2 \stackrel{\text{def}}{=} \forall t \in S_1. \exists s \in S_2. \exists \theta. \theta(s) \equiv t$$

$$\text{(ProofCo)} \quad \frac{h \bar{x} : (\overline{Cv}) \Rightarrow s \sim t \in \Gamma}{\Gamma \cup \{C \overline{[u/x]v}\}; h \bar{u} \vdash \gamma : \overline{[u/x]s} \sim \overline{[u/x]t}} \Gamma \vdash \text{proofcase } h \bar{u} = \gamma : \diamond$$

$$\text{(RefI Co)} \quad \Gamma \vdash \text{refl } t : t \sim t \quad \text{(SymCo)} \quad \frac{\Gamma \vdash \gamma : s \sim t}{\Gamma \vdash \text{sym } \gamma : t \sim s}$$

$$\text{(TransCo)} \quad \frac{\Gamma \vdash \gamma_1 : t_1 \sim t_2 \quad \Gamma \vdash \gamma_2 : t_2 \sim t_3}{\Gamma \vdash \gamma_1 \circ \gamma_2 : t_1 \sim t_3}$$

$$\text{(TFAppCo)} \quad \frac{\Gamma \vdash \gamma_i : s_i \sim t_i \quad i = 1, \dots, n}{\Gamma \vdash F \gamma_1 \dots \gamma_n : F s_1 \dots s_n \sim F t_1 \dots t_n}$$

$$\text{(AppCo)} \quad \frac{\Gamma \vdash \gamma_1 : f_1 \sim f_2 \quad \Gamma \vdash \gamma_2 : s_1 \sim s_2}{\Gamma \vdash \gamma_1 \gamma_2 : f_1 s_1 \sim f_2 s_2}$$

$$\text{(DecompT)} \quad \frac{\Gamma \vdash \gamma : T s_1 \dots s_n \sim T t_1 \dots t_n}{\Gamma \vdash \text{decomp}_{T_i} \gamma : s_i \sim t_i} \quad (i \in 1..n)$$

$$\text{(AxCo)} \quad \frac{g \bar{x} : s_1 \sim s_2 \in \Gamma}{\Gamma \vdash g \bar{t} : \overline{[t/x]s_1} \sim \overline{[t/x]s_2}}$$

$$\text{(InvCo)} \quad \frac{h \bar{x} : (\overline{Cu}) \Rightarrow s \sim t \in \Gamma \quad \Gamma \vdash \delta_i : C_i \overline{[t/x]u_i}}{\Gamma \vdash \bar{\delta} \Rightarrow \text{norec } h \bar{t} : \overline{[t/x]s} \sim \overline{[t/x]t}}$$

$$\text{(RecInvCo)} \quad \frac{h \bar{x} : (\overline{Cu}) \Rightarrow s \sim t \in \Gamma \quad \Gamma \vdash \delta_i : C_i \overline{[t/x]u_i} \quad h \bar{t} < h_{ths} \bar{t}_{ths}}{\Gamma; h_{ths} \bar{t}_{ths} \vdash \bar{\delta} \Rightarrow \text{rec } h \bar{t} : \overline{[t/x]s} \sim \overline{[t/x]t}}$$

Figure 3. Type Equation Proof System

Note that this formulation allows for multiple required proof cases to be covered by one and the same supplied proof case. In this way, a combinatorial explosion of proof cases can often be avoided.

Example 12. One of the invariants in the type preserving compiler of (Guillemette and Monnier 2008), expresses that CPS transformation before and after type variable substitution yield the same result:

$$\begin{array}{l} \text{invariant } k_s = \\ (\text{Tp } s, \text{Tp } t, \text{Nat } i) \Rightarrow \\ K (\text{Subst } s \ t \ i) \sim \text{Subst } (K \ s) \ (K \ t) \ i \end{array}$$

There are two instances of Nat and four of Tp, and hence $32 = 4 \times 4 \times 2$ required proof cases. Luckily, we can get away with writing only four proof cases, only instantiating s to a more specialized type. \square

4.4 Well-Foundedness

In order to ensure the well-foundedness of (directly or mutually) inductive proofs, we require that (potentially) inductive uses of

$$\text{(ProofTC)} \quad \frac{p \bar{x} : (\overline{C v}) \Rightarrow C t \in \Gamma \quad \Gamma \cup \{C \overline{[u/x]v}\} \vdash \delta : C \overline{[u/x]t}}{\Gamma \vdash \text{proofcase } p \bar{u} = \delta : \diamond}$$

$$\text{(EnvTC)} \quad \frac{C t \in \Gamma}{\Gamma \vdash \text{env } (C t) : C t}$$

$$\text{(InstTC)} \quad \frac{\Gamma \vdash \delta_i : C_i \theta(t_i) \quad \text{instance } (C_1 t_1, \dots, C_n t_n) \Rightarrow C t \in \Gamma}{\Gamma \vdash \text{inst } \bar{\delta} (C \theta(t)) : C \theta(t)}$$

$$\text{(SSelTC)} \quad \frac{\Gamma \vdash \delta : C \theta(a) \quad \text{class } (C_1 t_1, \dots, C_n t_n) \Rightarrow C a \in \Gamma}{\Gamma \vdash \text{sselect } i \delta : C_i \theta(t_i)}$$

$$\text{(ISelTC)} \quad \frac{\Gamma \vdash \delta : C \theta(t) \quad \text{instance } (C_1 t_1, \dots, C_n t_n) \Rightarrow C t \in \Gamma}{\Gamma \vdash \text{iselect } i \delta : C_i \theta(t_i)}$$

$$\text{(CoTC)} \quad \frac{\Gamma \vdash \gamma : t_1 \sim t_2}{\Gamma \vdash \text{co } (C t_1) \gamma : C t_2}$$

Figure 4. Type Class Proof System

$$\text{(LemCo)} \quad \frac{\{(C_1 u_1, \dots, C_n u_n) \mid (\text{instance } \dots \Rightarrow C_i u_i) \in \Gamma (1 \leq i \leq n)\} \sqsubseteq \{[\overline{v/x}](C_1 y_1, \dots, C_n y_n) \mid (\text{proofcase } h \bar{v} = \gamma) \in \Gamma\}}{\Gamma \vdash h \bar{x} : (\overline{C y}) \Rightarrow s \sim t : \diamond}$$

$$\text{(LemTC)} \quad \frac{\{(C_1 u_1, \dots, C_n u_n) \mid (\text{instance } \dots \Rightarrow C_i u_i) \in \Gamma (1 \leq i \leq n)\} \sqsubseteq \{[\overline{v/x}](C_1 y_1, \dots, C_n y_n) \mid (\text{proofcase } p \bar{v} = \gamma) \in \Gamma\}}{\Gamma \vdash p \bar{x} : (C_1 y_1, \dots, C_n y_n) \Rightarrow C t : \diamond}$$

Figure 5. Complete Proofs

invariants in proofs are annotated as such, with the keyword `rec`. In order to keep the well-foundedness check simple and modular we require that each such invariant use is strictly smaller than the covered proof case it appears in according to a well-founded partial order denoted by \prec in rule (RECINVC0).

DEFINITION 2. *The well-founded partial order \prec is defined as:*

$$h_1 \bar{t} \prec h_2 \bar{s} \stackrel{\text{def}}{\equiv} |\bar{t}| < |\bar{s}| \wedge \text{vars}(\bar{t}) \sqsubseteq \text{vars}(\bar{s})$$

where the term norm $|\cdot|$ is defined as:

$$\begin{aligned} |\bar{t}| &= \sum_i |t_i| \\ |t_1 t_2| &= |t_1| + |t_2| \\ |T| &= 1 \\ |F \bar{t}| &= 1 + |\bar{t}| \\ |a| &= 1 \end{aligned}$$

the function `vars` returns the multi-set (bag) of type variables:

$$\begin{aligned} \text{vars}(\bar{t}) &= \uplus_i \text{vars}(t_i) \\ \text{vars}(t_1 t_2) &= \text{vars}(t_1) \uplus \text{vars}(t_2) \\ \text{vars}(T) &= \emptyset \\ \text{vars}(F \bar{t}) &= \uplus_i \text{vars}(t_i) \\ \text{vars}(a) &= \{a\} \end{aligned}$$

where \uplus \sqsubseteq are the multi-set union and subset relations, i.e. taking multiplicity into account.

Observe that this definition of \prec considers induction on the combination of all invariant arguments.

Example 13. Example 5, perhaps more clearly in the higher-level notation of Example 11, involves an inductive use `add_comm Z n` in the proof for `add_comm (S n) Z`.

We have:

$$\begin{aligned} |Z \ n| &= 2 & \text{vars}(Z \ n) &= \{n\} \\ |(S \ n) \ Z| &= 3 & \text{vars}((S \ n) \ Z) &= \{n\} \end{aligned}$$

Hence, the induction is well-founded as $2 < 3 \wedge \{n\} \sqsubseteq \{n\}$, i.e. `add_comm Z n` \prec `add_comm (S n) Z`. \square

The above condition is not imposed on invariant uses that are not annotated as potentially inductive. To do so would be overly restrictive and rule out proofs for many useful properties. This is safe only if the invariant uses is truly not inductive, and the annotation is not inadvertently or maliciously omitted by the programmer. Hence, we collect all edges in the invariant “call graph”, and globally check whether no annotations have been omitted.

4.5 Type Soundness

The constructs we added to System F_C preserve its soundness. More specifically, the progress lemma still holds thanks to the completeness check: whenever we need to evaluate a call to a coercion function, the completeness check guarantees that the corresponding lemma does provide an applicable case. The type preservation lemma also still holds trivially. And the termination check is used to show that the coercion functions can indeed be implemented as no-ops.

5. Implementation and Evaluation

We have implemented a prototype well-typing checker in GHC-Haskell. This checker implements the rules formalized in Section 4. Because the formal proof notation is quite verbose and awkward for manual use, we have also implemented support for the more compact and natural notation discussed in Section 3.6. We call the former the *internal proof language*, and the latter the *external proof language*.

5.1 The External Proof Language

A proof in the external proof language consists of a sequence of types $\tau_1 \sim \dots \sim \tau_n$, and denotes a proof for the equation $\tau_1 \sim \tau_n$. Our checker contains an inferencer that reconstructs the internal language proof from this external language proof. Essentially, it reconstructs the internal proof for two subsequent types $\tau_i \sim \tau_{i+1}$ as a base case, and composes subsequent proofs γ_j and γ_{j+1} with the transitivity constructor (TRANSCO).

For the base case, a single external proof step $\tau_i \sim \tau_{i+1}$, the reconstruction proceeds top-down and only considers a finite number of possibilities:

- If τ_i and τ_{i+1} are identical, use the (REFLCO) rule.
- Otherwise:
 - repeatedly decompose the type equality into multiple proof obligations using the (APPCO) and (TFAPPCO) rules,
 - until the remaining proof obligations can be discharged with a single application of a rule from the set $\{(\text{REFLCO}), (\text{SYMCO}), (\text{AXCO}), (\text{INVC O}), (\text{RECINVC O})\}$

whose (possible) subproofs all use the (REFLCO) rule only.

When using the (INVC O) or (RECINVC O) rule, also the context proof is reconstructed. For that purpose, all possible proofs with no more than three type class proof constructors are tried, with the exception of proofs involving the (COTC) rule.

A number of proof constructors are not covered by our proof language, because they do not fit into the sequential form of the proofs. For the same reason, the proof language does not cover type class proofs. If any of these are needed, the programmer can resort to the internal proof language. Moreover, the external proof language is embedded in the internal one, so in any case only a minimal part of the proof has to be written explicitly in the internal language. Conceptually, the embedding works both ways:

- Obviously, because of the embedding a proof γ_e in the external language can be part of a larger proof γ_i in the internal language.
- The transition $\tau_i \sim \tau_{i+1}$ of an external proof $\gamma_e = \tau_1 \sim \dots \sim \tau_i \sim \tau_{i+1} \sim \dots \sim \tau_n$ can also be written as an internal proof γ_i as follows:

$$(\tau_1 \sim \dots \sim \tau_i) \circ \gamma_i \circ (\tau_{i+1} \sim \dots \sim \tau_n)$$

The experimental evaluation indicates that our external proof language is effective in reducing proof size, and that the additional expressive power of the internal language is needed only rarely.

5.2 Evaluation

As preliminary evaluation of the expressivity of our invariant language we have encoded all the invariants of Section 2 and their proof cases. Some statistics of these invariants are recorded in the table below:

main invariant	aux. invariants	cases	proof size	steps
parity	0 + 0	2 + 0	10	4
commutativity	0 + 0	4 + 0	60	19
CPS	5 + 1	21 + 3	327	101*

For each row the table lists the number of auxiliary invariants (equational + type class) the number of proof cases (equational + type class), and total proof size. As measure for proof size we use the number of coercion and type class proof constructors used. The last column denotes the number of proof steps written in the high-level notation, as a comparison to the low-level proof constructors.

There are 18 proof constructors not covered in (*). These comprise 11 constructors for the one type class invariant, and 7 con-

structors for the use of this type class invariant in the proof of an equational invariant.

The results indicate that our high-level notation is about three times more compact than the low-level one, which is a significant reduction in programmer effort. In addition to their compactness, we believe that high-level proofs can be written more quickly because the linear proof-style is much easier to grasp for the programmer.

6. Related Work

The Chameleon system allows programmers to extend the type checker with additional Constraint Handling Rules (CHRs) (Stuckey and Sulzmann 2005). These are useful for encoding additional properties, but have a more operational flavor, being left-to-right rewrite rules. The CHRs are more expressive than our invariants in that they allow existentially quantified type variables in the right-hand side. Yet, Chameleon treats the CHRs as axioms, and leaves the responsibility for soundness, completeness and termination to the programmer.

Invariants on type-level functions is something that can be done naturally in Coq (Paulin-Mohring 1993), although its type functions are closed. To go open world, one possibility may be to use the type class library of (Sozeau and Oury 2008), although it is not at all clear how to structure proofs of non-trivial lemmas using these.

The Twelf theorem prover (Pfenning and Schürmann 1999) has type families defined under an open world assumption, but these type families define relations rather than functions. Being dependently typed, it can be used to prove arbitrary invariants involving type families. The proofs take a logic programming flavor, unlike the equational proof syntax proposed here. Twelf provides coverage and termination checking, and support for proof search.

Omega (Sheard 2004) provide type-level functions similar to type families except that they are closed. One can reason about types at the term level using GADTs, but there are has no user-defined type classes or other support for type-level programs.

7. Conclusion & Future Work

We have shown the limitations of Haskell’s current type language, comprising type classes and type families. To extend the expressivity, we have proposed *type invariants*, which respect the open nature of the aforementioned type system features. Our formalization takes care of soundness, completeness and well-foundedness for type invariants and their proofs.

There are many ways in which to improve our external proof language and the reconstruction of the internal proof, e.g. covering more substantial internal proofs, more efficient search, and interactive proving to name just a few. We plan to investigate which extensions are the most likely to alleviate the programmer’s burden.

Acknowledgments

We are grateful to Brigitte Pientka, Martin Sulzmann and the anonymous reviewers for their helpful comments. Part of this work was conducted while the first author was a visitor at Université de Montréal.

References

- Ki Yung Ahn and Tim Sheard. Shared subtypes: subtyping recursive parametrized algebraic data types. In *Haskell ’08: Proceedings of the 1st ACM SIGPLAN Haskell symposium*, pages 75–86, New York, NY, USA, 2008. ACM.

- Dominic Duggan and John Ophel. Type-checking multi-parameter type classes. *J. Funct. Program.*, 12(2):133–158, 2002.
- Louis-Julien Guillemette and Stefan Monnier. A type-preserving compiler in Haskell. In *ICFP '08: Proceeding of the 13th ACM SIGPLAN international conference on Functional programming*, pages 75–86, New York, NY, USA, 2008. ACM.
- Mark P. Jones. Type classes with functional dependencies. In *Proc. of ESOP 2000*, number 1782 in Lecture Notes in Computer Science. Springer-Verlag, 2000.
- Christine Paulin-Mohring. Inductive definitions in the system Coq—rules and properties. In M. Bezem and J. Groote, editors, *International conference on Typed Lambda Calculi and Applications*. LNCS 664, Springer-Verlag, 1993.
- Frank Pfenning and Carsten Schürmann. System description: Twelf - a meta-logical framework for deductive systems. In *International Conference on Automated Deduction*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 202–206, July 1999.
- Tom Schrijvers and Martin Sulzmann. Unified Type Checking for Type Classes and Type Functions, 2008. Poster at the International Conference on Functional Programming (ICFP'08).
- Tom Schrijvers, Simon Peyton Jones, Manuel Chakravarty, and Martin Sulzmann. Type checking with open type functions. In *ICFP '08: Proceeding of the 13th ACM SIGPLAN international conference on Functional programming*, pages 51–62, New York, NY, USA, 2008. ACM.
- Tim Sheard. Languages of the future. In *OOPSLA '04: Companion to the 19th annual ACM SIGPLAN conference on Object-oriented programming systems, languages, and applications*, pages 116–119, New York, NY, USA, 2004. ACM Press. ISBN 1-58113-833-4.
- Matthieu Sozeau and Nicolas Oury. First-class type classes. In *21th International Conference on Theorem Proving in Higher Order Logics*, pages 278–293. LNCS 5170, Springer-Verlag, 2008.
- Peter J. Stuckey and Martin Sulzmann. A theory of overloading. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 27(6):1–54, 2005.
- Martin Sulzmann, Manuel Chakravarty, Simon Peyton Jones, and Kevin Donnelly. System F with type equality coercions. In *ACM SIGPLAN International Workshop on Types in Language Design and Implementation (TLDI'07)*. ACM, 2007a.
- Martin Sulzmann, Gregory J. Duck, Simon Peyton-Jones, and Peter J. Stuckey. Understanding functional dependencies via Constraint Handling Rules. *J. Funct. Program.*, 17(1):83–129, 2007b.