

A Typechecker Plugin for Units of Measure

Domain-specific constraint solving in GHC Haskell

Adam Gundry

Well-Typed LLP

adam@well-typed.com

Abstract

Typed functional programming and units of measure are a natural combination, as F# ably demonstrates. However, encoding statically-checked units in Haskell’s type system leads to inevitable disappointment with the usability of the resulting system. Extending the language itself would produce a much better result, but it would be a lot of work! In this paper, I demonstrate how typechecker plugins in the Glasgow Haskell Compiler allow users to define domain-specific constraint solving behaviour, making it possible to implement units of measure as a type system extension without rebuilding the compiler. This paves the way for a more modular treatment of constraint solving in GHC.

1. Introduction

Dimensions (such as length and time) and units of measure (such as metres, feet and seconds) are a highly useful mechanism to

- reduce the chances of making a costly error¹, and
- make it easier to perform calculations.

As Kennedy (2010) put it, “Units-of-measure are to science what types are to programming.” It is natural, therefore, to consider the extension of typed programming languages with support for units of measure. At a minimum, such support should allow the programmer to declare the units of quantities, and prevent them making errors such as adding incompatible quantities. There has been much work in this direction, notably by Kennedy (2010) in the context of the F# functional programming language. He has shown that units of measure fit particularly well with Hindley-Milner type inference, leading to a simple but powerful system.

For example, in F# one can write

¹It is traditional here to cite the Mars Climate Orbiter, or the Gimli Glider (<http://lamar.colostate.edu/~hillger/unit-mixups.html>).

```
> [<Measure>] type m;;
> [<Measure>] type s;;
> let time    = 3.0⟨s⟩;;
> let speed   = 5.0⟨m/s⟩;;
> let distance = time * speed;;
```

and the system will correctly infer the units of *distance*:

```
val distance : float⟨m⟩ = 15.0
```

In addition, Kennedy’s system supports *unit polymorphism*: definitions can be checked abstractly, with the concrete units being determined at the use sites. For example, one can define the function

```
> let sqr (x : float⟨_⟩) = x * x;;
val sqr : x : float⟨u⟩ → float⟨u^2⟩
```

which is polymorphic in a unit variable *u*. The type annotation on the definition is necessary because overloaded arithmetic operators in F# do not have units by default.

Modern GHC Haskell supports a range of language features (in particular, type families) that make it possible to encode quite complex properties at the type level. Correspondingly, in the Haskell world there have been various attempts to encode units of measure, in particular the robust and expressive `units` library by Muranushi and Eisenberg (2014). This allows one to write

```
time    = 3.0 % [si] s []
speed   = 5.0 % [si] m/s []
distance = time * speed
```

although the inferred type of *distance* is not F#’s `float⟨m⟩` but

```
distance :: Qu [F Length One] DefaultLCSU Double .
```

While this work is very impressive, it is inevitably limited by the features that GHC exposes to programmers, and this shows up in the inferior type inference behaviour and error messages produced by units of measure libraries as compared to genuine language extensions (as in F#).

How might we go about extending the Haskell language itself, as implemented in GHC, with units of measure? While GHC’s extensibility is impressive, writing a language extension is inevitably a drawn-out and complex process, requiring a great deal of work to specify and implement the new feature. Moreover, GHC is a moving target: its rapid pace of development makes it difficult to work on large new features without spending much time resolving merge conflicts. It would be very desirable if we could plug in support for units of measure to GHC, *without* changing GHC itself.

1.1 Type-level arithmetic

Additional motivation for such functionality comes from another desirable GHC extension: increasing the automated reasoning available to users of type-level arithmetic. For some time, thanks to the work of Iavor Diatchki², it has been possible to use natural number literals and arithmetic operators in types. For example, one can define vectors (lists indexed by their length):

```
data Vec a (n :: Nat) where
  Nil  :: Vec a 0
  Cons :: a → Vec a n → Vec a (1 + n)
myVec :: Vec Char 3
myVec = Cons 'a' (Cons 'b' (Cons 'c' Nil))
vhead :: Vec a (1 + n) → a
vhead (Cons x _) = x
```

However, further progress is stymied by the lack of support for working with numeric *variables*. While *vhead* works, this simple definition of vector tail

```
vtail :: Vec a (1 + n) → Vec a n
vtail (Cons _ xs) = xs
```

is not accepted by GHC 7.8.3, because it does not know that $(1+)$ is an injective function:

```
Could not deduce (n1 ~ n)
from the context ((1 + n) ~ (1 + n1))
  bound by a pattern with constructor Cons
  ... in an equation for vtail
```

It would be nice if the typechecker was able to prove more equations, using domain-specific knowledge about arithmetic. However, it is not entirely clear how best to implement this. One possibility is to interface GHC to an SMT solver, so that the SMT solver can solve arithmetic equations left unsolved by GHC. The *inch* typechecker (Gundry 2013) demonstrates an alternative approach based on ring normalization. In either case, the ability to provide this functionality via an external plugin, rather than directly in GHC, permits experimentation, makes development easier and enables deployment independently of GHC's release cycle.

1.2 Compiler plugins

Max Bolingbroke and Austin Seipp implemented support for compiler plugins in GHC version 7.2.1³. Inspired by a similar concept in the GNU Compiler Collection (GCC), they were originally intended for adding custom optimizations and analyses of GHC's internal Core language (System Fc). The basic idea is that a user package (distributed separately from GHC itself) contains a module *M* that exports a symbol *plugin* belonging to the type *Plugin* defined in the GHC API. Users can invoke GHC with an additional argument `-fplugin=M`, whereon the module will be dynamically linked into the running compiler, and invoked during compilation.

Crucially, plugins allow new compiler functionality to be added separately from the main development effort. This makes feature development quicker, as the entire system need not be recompiled when a plugin is changed, and makes it easier for programmers who are not compiler developers to contribute to and use plugins.

²<https://ghc.haskell.org/trac/ghc/wiki/TypeNats>

³https://downloads.haskell.org/~ghc/7.2.1/docs/html/users_guide/compiler-plugins.html

1.3 Summary

In the sequel, I will first describe `uom-plugin`, a Haskell library for units of measure, then explain the typechecker plugins feature that makes it possible. I will specify the typechecker plugins mechanism in general, and the constraint-solving algorithm used by the plugin, in terms of GHC's `OutsideIn(X)` type inference framework.

Concretely, the contributions of this paper are:

- a design for a units of measure library with good type inference properties, showing the need for domain-specific constraint solving behaviour in the typechecker (section 2);
- an explanation of the typechecker plugins interface that enables constraint solver extension, both informally in Haskell and formally by relating it to `OutsideIn(X)` (section 3);
- an algorithm for solving constraints in the equational theory of free abelian groups, which satisfies the properties required for sound and most general type inference (section 4); and
- a comparison of the resulting units of measure system to other approaches in F# and Haskell (section 5).

The `uom-plugin` library is available online.⁴ The typechecker plugin functionality on which it relies will be available in GHC 7.10.

2. Units of measure

First, let us consider how to extend our language with the syntax of units of measure, then go on to discuss its semantics.

2.1 The syntax of units

A typical approach to units of measure in programming languages is to annotate numeric types with their units, such as the `int⟨·⟩` and `float⟨·⟩` type constructors in F#. In Haskell, the natural way to do this is through the definition

```
newtype Quantity a (u :: Unit) = MkQuantity a
```

which makes `Quantity a u` use the same runtime representation as the underlying (typically numeric) type *a*, but tagged with a *phantom type parameter* (Leijen and Meijer 1999) *u* of kind `Unit`. This means that using `Quantity a u` has no runtime overhead compared to using plain *a*, but it can have additional safety guarantees.

The `Unit` datatype is lifted to the kind level via *datatype promotion* (Yorgey et al. 2012). It has no constructors, but instead is accompanied by the following type-level definitions, implemented as type families without any equations:

```
1  :: Unit
Base :: Symbol → Unit
(⊗) :: Unit → Unit → Unit -- or * : in ASCII
(⊘) :: Unit → Unit → Unit -- or / : in ASCII
```

`Base` creates base units, which are represented as type-level strings (of kind `Symbol`) for simplicity. Dimensionless quantities are represented with `1`, and the operators allow more complex units to be formed. Representing them as type families with no equations means they are essentially opaque symbols that may not be partially applied and are not injective; this avoids the equational theory of units conflicting with GHC's built-in equality rules for types.

⁴<https://github.com/adamgundry/uom-plugin>

Crucially, the `MkQuantity` constructor should be hidden from client code, so that users of library are forced to work with a unit-safe interface.⁵ If the constructor were available, users could write code like this, which would destroy all unit safety guarantees provided by the library:

```
unsafeConvertQuantity :: Quantity a u → Quantity a v
unsafeConvertQuantity (MkQuantity x) = MkQuantity x
```

Of course, users need some way to produce and consume quantities, i.e. convert between `a` and `Quantity a u`. It is fine for the library to expose

```
unQuantity :: Quantity a u → a
unQuantity (MkQuantity x) = x
```

but not

```
MkQuantity :: a → Quantity a u
```

as their composition yields `unsafeConvertQuantity`. The real problem here is that `MkQuantity` should only be *monomorphic* (sometimes known as ‘weakly polymorphic’) in its unit. It is fine for `u` to be any concrete unit, but it must not be generalized over (to become a universally quantified type variable) outside the module in which it is defined. Such variables are permitted in types in Caml (Garrigue 2004), but not in Haskell. As a workaround, the library offers a Template Haskell *quasiquote* that enables the user to write concrete quantities in a convenient syntax:

```
mass = [u| 65 kg |]
g     = [u| 9.808 m/s^2 |]
```

Moreover, omitting the numeric value yields a specialization of `MkQuantity` to the appropriate type, which is useful when units need to be attached to numeric values that are not literal constants, for example:

```
readMass :: IO (Quantity Double (Base "kg"))
readMass = fmap [u | kg |] readLn
```

The library includes the following (written `+` and `*` in ASCII):

```
(⊕) :: Num a ⇒
  Quantity a u → Quantity a u → Quantity a u
MkQuantity x ⊕ MkQuantity y = MkQuantity (x + y)

(⊗) :: Num a ⇒
  Quantity a u → Quantity a v → Quantity a (u ⊗ v)
MkQuantity x ⊗ MkQuantity y = MkQuantity (x * y)
```

The (\oplus) and (\otimes) operators on quantities are analogous to the $(+)$ and $(*)$ operators on numbers, except that the phantom parameter makes sure the units are kept in order. Quantities may be multiplied regardless of their units, but may be added only if the units match.⁶

For example, if we have some values

```
mass    :: Quantity Double (Base "kg")
distance :: Quantity Double (Base "m")
```

then we can we define

```
x :: Quantity Double (Base "kg" ⊗ Base "m")
x = mass ⊗ distance
```

⁵ In fact, the constructor is exported by a separate internal module, as it is sometimes useful in practice, e.g. to allow `Coercible` coercions of data structures containing quantities.

⁶ Unfortunately this means that `Quantity a u` cannot be an instance of the standard Haskell `Num` typeclass, which bundles addition, subtraction and multiplication together. An instance may be given only for `Quantity a $\mathbb{1}$` .

but attempting to add `mass` to `distance` gives a type error:

```
Couldn't match type 'Base "m"' with 'Base "kg"'
Expected type: Quantity Double (Base "m")
Actual type: Quantity Double (Base "kg")
In the first argument of '(+)', namely 'mass'
In the expression: mass +: distance
```

2.2 The equational theory of units

Are we done? Not quite. Our definitions so far allow us to write the *syntax* of units of measure, but we have not accounted for their *equational theory*. We would expect quantities with the units `Base "kg" ⊗ Base "m"` and `Base "m" ⊗ Base "kg"` to be interchangeable; unit multiplication should be commutative! But adding `mass ⊗ distance` to `distance ⊗ mass` gives:

```
Couldn't match type 'Base "m" *:' with 'Base "kg"
with 'Base "kg" *:' with 'Base "m"
NB: *: is a type function, and may not be injective
Expected type:
  Quantity Double ('Base "kg" *:' with 'Base "m")
Actual type:
  Quantity Double ('Base "m" *:' with 'Base "kg")
In the second argument of '(+)', namely
  (distance *: mass)
In the expression:
  (mass *: distance) +: (distance *: mass)
```

In addition to the usual GHC Haskell rules for type equality (Sulzmann et al. 2007), we would like additional equations to hold to characterise the operations. As in Kennedy’s system in F#, these equations are the standard laws of an abelian group:

$$\begin{aligned} \forall u v. ((u \otimes v) \oplus w) &\sim (u \otimes (v \oplus w)) \\ \forall u v. (u \otimes v) &\sim (v \otimes u) \\ \forall u v. (u \otimes \mathbb{1}) &\sim u \\ \forall u v. (u \otimes (\mathbb{1} \circlearrowleft u)) &\sim \mathbb{1} \end{aligned}$$

But how can we make them hold? GHC allows new axioms to be introduced using a type family, but type families (like functions) may pattern match only on *constructors*, not other type families. In any case, type families are typically useful only if they define a terminating rewrite system, and associativity and commutativity are hardly going to do so!

2.3 An attempted solution

This is the point where most Haskell units of measure libraries give up on providing exactly the desired equational theory given above. Instead, a common approach is to write a normalisation function for concrete unit expressions (implemented as a type family, of course). Provided one is very careful to talk only about equality of normal forms, not the original syntax of unit expressions, this allows some of the desired behaviour. For example, the `units` package defines addition of quantities with this type:⁷

$$\text{(⊕)} :: (d_1 @\sim d_2, \text{Num } n) \Rightarrow \text{Qu } d_1 \text{ } l \text{ } n \rightarrow \text{Qu } d_2 \text{ } l \text{ } n \rightarrow \text{Qu } d_1 \text{ } l \text{ } n$$

⁷ Instead of annotating quantities with units directly, the `units` package uses a combination of dimension (`d`) and local coherent system of units (`l`); the difference is discussed in subsection 5.2.

Boxed operators are definitions from `units`, to distinguish them from the circled operators of `uom-plugin`. \boxplus is written `|+|` in ASCII.

<code>tcPluginIO</code>	<code>:: IO a → TcPluginM a</code>	-- Perform arbitrary IO
<code>tcPluginTrace</code>	<code>:: String → SDoc → TcPluginM ()</code>	-- Print debug message
<code>findImportedModule</code>	<code>:: ModuleName → Maybe FastString → TcPluginM FindResult</code>	-- Look up details of a module
<code>tcLookupGlobal</code>	<code>:: Name → TcPluginM TyThing</code>	-- Look up a type or definition in the context
<code>newFlexiTyVar</code>	<code>:: Kind → TcPluginM TcTyVar</code>	-- Create a fresh unification variable

Figure 1. A sample of the TcPluginM interface

The constraint $d_1 @\sim d_2$ means that d_1 and d_2 should be compared up to equality of normal forms, a weaker condition than $d_1 \sim d_2$, which would require d_1 and d_2 to be equal in GHC’s equational theory. This weaker constraint means that GHC can determine that $\text{kg} * \text{m}$ and $\text{m} * \text{kg}$ have the same normal form, and hence quantities with those units may be added.

Unfortunately, this has rather drastic consequences for error messages, because they are expressed in terms of normal forms. In `units`, the erroneous addition gives this:

```
Couldn't match type [F Mass One, F Length (P Zero)]
  with []
In the expression: mass |+| distance
```

Apart from introducing yet another equivalence relation for the user to understand, and mystifying them with error messages, we lose something important in the shift from (\oplus) to (\boxplus) : well-behaved unit polymorphism. The normalisation approach breaks down when there are variables or other non-canonical unit expressions present. It cannot conclude that $u \boxtimes v$ is interchangeable with $v \boxtimes u$, because it cannot compute the normal form of variables such as u and v . If we are lucky, we may be able to postpone the constraint until we have concrete values for the variables, and hence get away with only some messy types. If not, we may not be able to write the program we want.

For example, one would like the following function to be accepted (with or without the type signature)

```
f :: Num a =>
  Quantity a u → Quantity a v → Quantity a (u ⊗ v)
f x y = (x ⊗ y) ⊕ (y ⊗ x)
```

but the `units` approach leads to an inferred type like this, involving several internal type families used to implement normalisation:

```
(Num a
, [] ~ Normalize (Normalize (d1 @@@+ Reorder d2 d1)
                        @- Normalize (d2 @@@+ Reorder d1 d2))
) => Qu d1 l a → Qu d2 l a
  → Qu (Normalize (d1 @@@+ Reorder d2 d1)) l a
```

The crucial observation of this paper is that we need to introduce support for a *domain-specific equational theory*. In the following section, we will see how this is possible.

3. Domain-specific constraint solving

Haskell type inference is essentially a problem of generating and solving constraints. These may be equalities, which arise from the typing rules (e.g. in the application $f x$, the compiler must check that f has a function type with domain equal to the type of x), or typeclass constraints, which arise from uses of overloaded functions. Similarly, standard Hindley-Milner type inference amounts to a constraint generation and solving process in which the solver performs first-order unification (Sulzmann et al. 1999).

GHC uses the OutsideIn(X) algorithm (Vytiniotis et al. 2011) to handle the constraints it generates. This is notionally parametric in the choices of constraint domain X and solving algorithm, and provides domain-independent conditions that the constraint solver must satisfy. However, in practice there is only one choice for the solver: GHC implements the solver for type equality constraints (including type families) and typeclasses also described by Vytiniotis et al. (2011). To permit domain-specific equational theories, this solver must be made user-extensible. The user is not expected to replace the solver entirely, although the capability might be interesting (e.g. to experiment with other algorithms).

In this section, I will describe how such a plugin constraint solver works, first as a practical Haskell program interfacing with GHC, then in the more formal theoretical setting of OutsideIn(X).

3.1 Plugging into GHC

Once GHC’s built-in constraint solver has finished its work, it is left with a set of constraints that it could not solve. The job of a plugin solver is to take this set of wanted constraints and either

- identify impossible constraints that GHC has failed to reject outright, for example $\text{kg} \otimes \text{kg} \sim \text{m}$; or
- solve or further simplify the constraints, perhaps generating others in the process.

When a plugin yields new constraints, the main GHC constraint solver will be re-invoked in case it can make further progress, the plugin will be called again, and so on.

To be more precise, a plugin solver is a Haskell function supplied separately with ‘given’, ‘derived’⁸ and ‘wanted’ constraints:

```
solve :: [Ct] → [Ct] → [Ct] → TcPluginM TcPluginResult
solve givens deriveds wanteds = ...
```

Here `Ct` is GHC’s internal type of constraints, `TcPluginM` is a monad providing effects suitable for plugins, and `TcPluginResult` captures possible outcomes of constraint solving:

```
data TcPluginResult
  = TcPluginOk { solved :: [(EvTerm, Ct)], new :: [Ct] }
  | TcPluginContradiction { impossible :: [Ct] }
```

The `TcPluginOk` case includes a list of *solved* constraints along with associated evidence (to be discussed in subsection 3.1.2), and a list of *new* constraints to be processed by the main solver. Note that it is possible for ‘given’ or ‘derived’ constraints to be solved, which simply means to drop them from consideration since they provide no useful information (e.g. consider $a \otimes \mathbb{1} \sim a$). The result `TcPluginOk [] []` indicates that no progress was made: no constraints could be solved and no new constraints were generated.

The details of the `TcPluginM` monad interface is not important; a few example type signatures are shown in Figure 1. These include

⁸ Derived constraints arise during the constraint solving process, e.g. from functional dependencies; they will not be considered in any detail here.

the ability to query the context (e.g. look up the definitions of types) and perform IO operations (e.g. to communicate with an external process). The ability to do IO is not used in `uom-plugin` (apart from printing debug messages), but it is useful in other plugins.

Of course, plugins should be essentially pure, but this is a matter for the plugin implementor. More generally, what does it mean for a plugin to be well-behaved? One would expect it to be:

- *pure*, i.e. producing the same result for the same inputs;
- *order-insensitive*, i.e. regarding the constraint lists passed to the `solve` function as sets (arguably the types should enforce this!);
- *sound*, i.e. claiming to solve constraints only if they can actually be solved, to be elaborated on in subsection 3.1.2;
- *most general*, i.e. solving constraints without ‘guessing’, which I will return to in section 4.

3.1.1 Plugin-aware constraint solving

The algorithm GHC uses when solving constraints in the presence of a typechecker plugin is as follows:

1. Run the built-in constraint solver, producing a set of constraints that it could neither solve nor show inconsistent.
2. Call the plugin with the remaining constraints:
 - if it returns `TcPluginContradiction`, report the impossible constraints and stop;
 - if it returns `TcPluginOk` with some new constraints, remove the solved constraints from the constraint set, add the new ones, then start again from the beginning;
 - otherwise, remove the solved constraints from the constraint set and stop.

For example, suppose GHC has arrived at a point in the typechecking process where it has some type family $F :: \text{Unit} \rightarrow *$, a given constraint $F (m \otimes s) \sim ()$, an as-yet unsolved unification variable α , and wanted constraints

$$\begin{aligned} F \alpha \sim () \\ (\alpha \otimes s) \sim m \end{aligned}$$

that have already been simplified as far as possible by the built-in constraint solver. The plugin solver can now run and output a new wanted constraint $\alpha \sim m \otimes s$, leading to the wanted constraints

$$\begin{aligned} F \alpha \sim () \\ (\alpha \otimes s) \sim m \\ \alpha \sim (m \otimes s). \end{aligned}$$

Now the built-in solver can make further progress, substituting for α and using the given constraint to discharge the first goal, leaving

$$(m \otimes s) \otimes s \sim m$$

which can be solved directly by another run of the plugin solver. Note that even this simple example involved two runs of the built-in solver and two runs of the plugin; while that could be avoided in this case if the plugin performed substitution and type family reduction itself, in general we would not want plugins to have to reimplement GHC’s entire solver!

Typeclasses	D
Data types	T
Type families	F, G
Rigid variables	a, b, c
Unification variables	α, β, γ
Type variables	$x, y, z ::= a \mid \alpha$
Types	$\tau ::= x \mid T \mid F \bar{\tau} \mid \tau_1 \tau_2 \mid \dots$
Constraints	$Q ::= \epsilon \mid Q_1 \wedge Q_2 \mid \tau_1 \sim \tau_2 \mid D \bar{\tau}$
Substitutions	$\theta, \phi ::= [\alpha \mapsto \bar{\tau}]$
Top-level axiom schemes	$\mathcal{Q} ::= Q \mid \mathcal{Q}_1 \wedge \mathcal{Q}_2 \mid \forall \bar{a}. Q \Rightarrow D \bar{\tau} \mid \forall \bar{a}. F \bar{\tau} \sim \tau$

Figure 2. Syntax of `OutsideIn(X)` types and constraints

3.1.2 Evidence of soundness

If a plugin claims to have solved a constraint, why should we believe it? It would be very easy to produce a plugin that erroneously⁹ reported constraints as solved when in fact they were not, potentially introducing type unsoundness and causing runtime crashes. Fortunately GHC already has a mechanism for detecting such errors: it does not merely typecheck code, but *elaborates* it into System F_C (Sulzmann et al. 2007), a very explicit core calculus that includes easily-checked evidence for type equality. While this does not prevent all compiler bugs, it makes constraint solver misbehaviour easier to detect.

Thus the actual implementation of plugins demands evidence for each constraint that the plugin claims to have solved. Some plugins may not be able to generate bona fide evidence, in which case they may use the equivalent of *unsafeCoerce* and assert a constraint without proof. On the other hand, the author of a plugin is free to create their own axiom schemes and build genuine evidence from them¹⁰, in which case they can be sure of the type soundness of the resulting system (provided the axioms they introduce are consistent, of course!).

In the implementation, the type `EvTerm` returned with a constraint in a `TcPluginOk` result represents terms in the evidence language. Forms of evidence include variables, axioms, typeclass dictionaries and a variety of deduction rules for equality proofs. I will not consider evidence further here.

3.2 Plugging in to `OutsideIn(X)`

Having seen how the plugin mechanism works in practice, let us step back and consider the theory justifying it. The `OutsideIn(X)` framework expects a constraint solver which takes four inputs (with the syntax given in Figure 2):

- user-defined top-level axiom schemes \mathcal{Q} (e.g. from typeclass and type family instances);
- ‘given’ constraints Q_{given} known to be true locally (e.g. from type signatures or GADT pattern matches);
- ‘touchable’ unification variables $\bar{\alpha}_{\text{tch}}$ (those for which the algorithm is allowed to solve); and
- ‘wanted’ constraints Q_{wanted} for which solutions are to be found.

⁹Or maliciously, though plugins are assumed to be trusted: they can run arbitrary IO actions from within the typechecker, which is dangerous!

¹⁰Modulo a limitation of the current implementation, which prevents custom axiom schemes being used across multiple modules.

$Q \in \mathcal{Q}$ implies $\mathcal{Q} \Vdash Q$	(R1)
$\mathcal{Q} \Vdash Q_1$ and $\mathcal{Q} \wedge Q_1 \Vdash Q_2$ implies $\mathcal{Q} \Vdash Q_2$	(R2)
$\mathcal{Q} \Vdash Q$ implies $\theta \mathcal{Q} \Vdash \theta Q$	(R3)
$\mathcal{Q} \Vdash \tau \sim \tau$	(R4)
$\mathcal{Q} \Vdash \tau_1 \sim \tau_2$ implies $\mathcal{Q} \Vdash \tau_2 \sim \tau_1$	(R5)
$\mathcal{Q} \Vdash \tau_1 \sim \tau_2$ and $\mathcal{Q} \Vdash \tau_2 \sim \tau_3$ implies $\mathcal{Q} \Vdash \tau_1 \sim \tau_3$	(R6)
$\mathcal{Q} \Vdash Q_1$ and $\mathcal{Q} \Vdash Q_2$ implies $\mathcal{Q} \Vdash Q_1 \wedge Q_2$	(R7)
$\mathcal{Q} \Vdash \tau_1 \sim \tau_2$ implies $\mathcal{Q} \Vdash [a \mapsto \tau_1] \tau \sim [a \mapsto \tau_2] \tau$	(R8)

Figure 3. Properties of entailment¹²

In response, the constraint solver must produce two outputs:

- a substitution θ for the touchable variables $\bar{\alpha}_{\text{tch}}$; and
- residual constraints Q_{residual} that could not be solved (but may have been simplified).

The behaviour of the constraint solver is described by the judgment

$$\mathcal{Q}; Q_{\text{given}}; \bar{\alpha}_{\text{tch}} \vdash^{\text{simp}} Q_{\text{wanted}} \rightsquigarrow Q_{\text{residual}}; \theta$$

which must satisfy certain conditions in order for OutsideIn(X) type inference to be sound and most general:¹¹

- (Soundness) $\mathcal{Q} \wedge Q_{\text{given}} \wedge Q_{\text{residual}} \Vdash \theta Q_{\text{wanted}}$
(Principality) $\mathcal{Q} \wedge Q_{\text{given}} \wedge Q_{\text{wanted}} \Vdash Q_{\text{residual}} \wedge \mathcal{E}_\theta$
where $\mathcal{E}_\theta = \{\alpha \sim u \mid [\alpha \mapsto u] \in \theta\}$

That is, the constraint solver must deliver a solution that is sound, i.e. the residual constraints solve the original problem, and most general, i.e. the simplifier has not ‘guessed’ any values for variables or invented constraints not entailed by the original problem. Here $\mathcal{Q} \Vdash Q$ is the constraint entailment relation, part of the X parameter of OutsideIn(X), which satisfies the properties given in Figure 3.

3.2.1 Defining a plugin constraint solver

A plugin constraint solver can be simpler than the description in the OutsideIn(X) framework, since it need not stand alone but will be combined with the built-in solver. In particular, the plugin need not deal with producing a substitution for unification variables directly. Instead, it may simply add constraints that define variables.

Suppose we have a judgment form

$$\mathcal{Q}; Q_{\text{given}}; \bar{\alpha}_{\text{tch}} \vdash^{\text{P}} Q_r \rightsquigarrow Q_s$$

meaning that the constraints Q_r can be simplified to Q_s under the given assumptions. Figure 4 shows how such a judgment can be combined with the built-in solver judgment \vdash^{simp} to produce \vdash^{psimp} , which conforms to the OutsideIn(X) interface.¹³

The basic idea is that of the implementation, discussed in subsection 3.1.1: run the main constraint solver once, then pass the residual constraints to the plugin. If the plugin generates new constraints (i.e. $Q_s \not\subseteq Q_r$), the GO rule applies and invokes the combined solver judgment again. If not (i.e. $Q_s \subseteq Q_r$), the STOP rule will simply return the remaining constraints.

¹¹ There are also some technical conditions on the domain of the substitution, which require that it substitutes only for touchable variables not occurring in the given or residual constraints.

¹² Slightly reformulated from Vytiniotis et al. (2011)

¹³ The details of how to calculate the sets of touchable variables α_1 and α_2 are omitted; it is straightforward but messy to add newly generated unification variables and remove those that have been substituted away.

Go	$\mathcal{Q}; Q_g; \bar{\alpha}_0 \vdash^{\text{simp}} Q_w \rightsquigarrow Q_r; \theta_0$
	$\mathcal{Q}; Q_g; \bar{\alpha}_1 \vdash^{\text{P}} Q_r \rightsquigarrow Q_s \quad Q_s \not\subseteq Q_r$
	$\mathcal{Q}; Q_g; \bar{\alpha}_2 \vdash^{\text{psimp}} Q_s \rightsquigarrow Q_t; \theta_1$
	$\mathcal{Q}; Q_g; \bar{\alpha}_0 \vdash^{\text{psimp}} Q_w \rightsquigarrow Q_t; \theta_0 \circ \theta_1 \mid \bar{\alpha}_0$
STOP	$\mathcal{Q}; Q_g; \bar{\alpha}_0 \vdash^{\text{simp}} Q_w \rightsquigarrow Q_r; \theta$
	$\mathcal{Q}; Q_g; \bar{\alpha}_1 \vdash^{\text{P}} Q_r \rightsquigarrow Q_s \quad Q_s \subseteq Q_r$
	$\mathcal{Q}; Q_g; \bar{\alpha}_0 \vdash^{\text{psimp}} Q_w \rightsquigarrow Q_s; \theta$

Figure 4. Plugin-extended OutsideIn(X) solver

Note that this process can be iterated, starting with the basic solver and extend it with multiple plugins.

The combined judgment \vdash^{psimp} will satisfy the OutsideIn(X) conditions on the assumption that \vdash^{simp} satisfies them, and provided that \vdash^{P} satisfies the conditions

- (Plugin soundness) $\mathcal{Q} \wedge Q_{\text{given}} \wedge Q_s \Vdash Q_r$
(Plugin principality) $\mathcal{Q} \wedge Q_{\text{given}} \wedge Q_r \Vdash Q_s$

i.e. Q_r and Q_s should be equivalent under the given constraints. In section 4.3 I will show that the units of measure plugin I am about to describe satisfies the soundness condition as-is, but satisfies only a weakened form of the principality condition.

4. Units of measure as a typechecker plugin

Having seen the general structure of typechecker plugins, let us consider a specific example. The uom-plugin constraint solver is designed to deal with equality constraints between types of kind Unit. Essentially it performs equational unification for the theory of free abelian groups. Recalling the earlier example, GHC’s built-in constraint solver might have been left with the unsolved constraint

$$\text{Base "m"} \otimes \text{Base "kg"} \sim \text{Base "kg"} \otimes \text{Base "m"}$$

but it is easy to see that this constraint is trivial simply by normalization up to the group axioms.

For constraints involving unification variables, Kennedy (1996, 2010) describes an algorithm for AG-unification that proceeds by a variant of Gaussian elimination, and shows how to extend this to types containing units of measure. For example, given the constraint

$$\alpha \otimes \alpha \sim \beta \otimes \beta \otimes \beta$$

the most general solution is

$$\alpha \sim \gamma \otimes \gamma \otimes \gamma, \beta \sim \gamma \otimes \gamma$$

for some fresh unification variable γ . Since AG-unification is decidable and possesses most general unifiers, type inference in an ML-like setting is well-behaved, though the let-generalisation step is slightly subtle (Gundry 2013).¹⁴

The situation is slightly more complex in the case of the full GHC Haskell type system, in particular because of the possible presence of universally quantified variables, type families and local constraints. Thus the plugin constraint solver may encounter constraints like

$$a \otimes a \sim b \otimes b \otimes b$$

¹⁴ GHC no longer generalises let-bindings by default in the presence of type families or GADTs, for essentially the same reason.

Unit constraints	$U ::= \epsilon \mid U_1 \wedge U_2 \mid u_1 \sim u_2$
Unit normal forms	$u ::= r \mid \mathbb{1} \mid u_1 \cdot u_2 \mid u^{-1}$
Atoms	$r ::= x \mid \mathbf{b} \mid \mathbf{F}(\bar{\tau})$
Base units	$\mathbf{b} ::= \mathbf{kg} \mid \mathbf{m} \mid \dots$

Figure 5. Syntax of unit constraints

where a and b are (universally quantified) rigid variables, or

$$\mathbf{F} a \otimes \mathbf{F} b \sim \mathbf{F} b \otimes \mathbf{F} a$$

where \mathbf{F} is a user-defined type family. Moreover, it has to deal with constraints that are ‘given’ as well as ‘wanted’, so it must simplify hypotheses as well as solving goals.

The essence of the plugin’s constraint solving algorithm is to

1. identify unsolved equality constraints between units;
2. normalize both sides of each constraint up to the group axioms;
3. incrementally simplify given constraints by rewriting them to simpler, equivalent constraints;
4. incrementally simplify wanted constraints, making use of the information from simplifying the givens.

For example, the wanted constraint $\alpha \otimes \alpha \sim (\beta \otimes \beta) \otimes \beta$ equates two types of kind `Unit`, which normalize to give $\alpha^2 \sim \beta^3$. Normal forms will be written in mathematical notation, to contrast them with Haskell type expressions. The syntax of normal forms is given in Figure 5, but they will be treated as equivalent up to the group axioms in Figure 6 (e.g. $\mathbb{1}$ and $x \cdot x^{-1}$ denote the same unit). I use U for constraints Q that include only equations between units.

Once normalized, the constraint $\alpha^2 \sim \beta^3$ can first be rewritten to $\alpha^2 \cdot \beta^{-3} \sim \mathbb{1}$. This can be simplified by substituting by $\alpha \sim \gamma \cdot \beta$ where γ is fresh, leading to $\gamma^2 \cdot \beta^{-1} \sim \mathbb{1}$. Rearranging this gives $\beta \sim \gamma^2$. Hence the solution is $\alpha \sim \gamma^3 \wedge \beta \sim \gamma^2$.

The presence of universally quantified variables or type families means that some constraints may not be solved immediately, but they may become feasible once other information has become available. This motivates a *dynamic* unification algorithm: one that makes progress on some constraints in the hope that others may become easier to solve. Since each step replaces a constraint with an equivalent constraint (up to the equational theory), it is most general, and so we can apply simplification steps in any order.

4.1 The constraint solving algorithm

For units of measure unification, the algorithm is given by the rules in Figure 7, which define the relation $U_0 \mapsto_{\bar{z}} U_1 (\theta, \phi)$. This explains how to rewrite a unit constraint U_0 into a new constraint U_1 and a pair of substitutions (θ, ϕ) that express the equivalence of the original and simplified constraints, modulo the group laws. The list of type variables \bar{z} records those that may not be modified during unification, including both rigid and unification variables.¹⁵

In the interests of simplicity, failure is not represented explicitly here, although in practice it is useful to identify obviously impossible constraints (such as $\mathbf{kg} \sim \mathbf{m}$), and the implementation does this using the `TcPluginContradiction` result (see subsection 3.1).

Rule (1) simply ensures that all unit equations are in the form $u \sim \mathbb{1}$. Rule (2) solves trivial equations; since unit normal forms

¹⁵ Parameterising by forbidden rather than touchable variables is merely a notational shortcut, to save changing the set when adding a fresh variable.

ASSOCIATIVE	IDENTITY
$\mathcal{Q} \Vdash u_1 \otimes (u_2 \otimes u_3) \sim (u_1 \otimes u_2) \otimes u_3$	$\mathcal{Q} \Vdash u \otimes \mathbb{1} \sim u$
COMMUTATIVE	INVERSE
$\mathcal{Q} \Vdash u_1 \otimes u_2 \sim u_2 \otimes u_1$	$\mathcal{Q} \Vdash u \otimes (\mathbb{1} \oslash u) \sim \mathbb{1}$
TORSION-FREE	
$\mathcal{Q} \Vdash u \otimes u \otimes \dots \otimes u \sim \mathbb{1}$	
$\mathcal{Q} \Vdash u \sim \mathbb{1}$	

Figure 6. Constraint entailment rules for units

are being considered up to the abelian group laws, this includes cases such as $\alpha \cdot \alpha^{-1} \sim \mathbb{1}$.

Rule (3) is the first to produce an output substitution, in the case where some variable can be instantiated to solve the equation. For example, $\mathbf{m}^4 \cdot \alpha^2 \sim \mathbb{1}$ is solved by substituting $[\alpha \mapsto \mathbf{m}^2]$. Of course, the variable must not belong to the list of fixed variables \bar{z} . Again this rule is interpreted up to the group laws.

The most complex rule is (4), which shows how progress can be made in cases where rule (3) does not apply and so the equation cannot immediately be solved. It relies on the fact that any unit can be expressed as a product of distinct atoms $r_1^{i_1} \dots r_n^{i_n}$. By replacing x with a fresh variable y multiplied by a suitably-chosen unit v , the exponents of the atoms can be reduced. Note that y should be a rigid variable iff x is rigid. For example, this rule introduces a fresh variable c to simplify $a^2 \cdot b^{-3} \sim \mathbb{1}$ to $c^2 \cdot b \sim \mathbb{1}$ with $\theta = [a \mapsto c \cdot b^2]$, $\phi = [c \mapsto a \cdot b^{-2}]$.

Rule (5) says that a conjunction of constraints can be simplified by simplifying one and applying the resulting substitution to the other. Just as units are considered up to the abelian group laws, conjunctions should be treated as sets, so this rule allows any constraint to be simplified.

The rules can be iterated in the obvious way to define a relation $U \mapsto_{\bar{z}}^* U' (\theta, \phi)$ that makes multiple simplification steps, composing the resulting substitutions.

4.2 Instantiating the OutsideIn(X) plugin framework

Recall that a plugin must supply a judgement

$$\mathcal{Q}; Q_g; \bar{\alpha}_{\text{tch}} \mapsto^{\text{P}} Q_r \rightsquigarrow Q_s$$

that explains how the given constraints Q_g and wanted constraints Q_r are simplified to produce the residual constraints Q_s . This judgement is defined by

$$\frac{U_g \mapsto^* U'_g (\theta_g, \phi_g) \quad \theta_g U_w \mapsto_{\bar{z}}^* U'_w (\theta_w, \phi_w)}{\bar{z} = \text{fv}(\theta_g U_w) \cup \text{fv}(\theta_w U_w) \setminus \bar{\alpha}_{\text{tch}}}$$

$$\mathcal{Q}; Q_g \wedge U_g; \bar{\alpha}_{\text{tch}} \mapsto^{\text{P}} Q_w \wedge U_w \rightsquigarrow Q_w \wedge \phi_g (U'_w \wedge \mathcal{E}_{\theta_w})$$

where Q_g and Q_w are the non-unit given and wanted constraints, respectively, and $\mathcal{E}_{\theta} = \{\alpha \sim u \mid [\alpha \mapsto u] \in \theta\}$ is the constraint form of a substitution θ .

This rule assumes without loss of generality that the unit constraints U_g and U_w are already in normal form; this is justified since every type of kind `Unit` is provably equal to its normal form according to the entailment relation.

First, the given unit constraints U_g are rewritten according to the simplification rules in Figure 7 until no more rules apply. This produces a substitution θ_g that may eliminate some rigid variables,

$u \sim v$	$\mapsto_{\bar{z}} u \cdot v^{-1} \sim \mathbb{1}$	(\cdot, \cdot)	if $v \neq \mathbb{1}$	(1)
$\mathbb{1} \sim \mathbb{1}$	$\mapsto_{\bar{z}} \epsilon$	(\cdot, \cdot)		(2)
$x^k \cdot u^k \sim \mathbb{1}$	$\mapsto_{\bar{z}} \epsilon$	$([x \mapsto u^{-1}], \cdot)$	if $x \notin \bar{z}$	(3)
$x^k \cdot r_1^{i_1} \dots r_n^{i_n} \sim \mathbb{1}$	$\mapsto_{\bar{z}} y^k \cdot r_1^{i_1 \bmod k} \dots r_n^{i_n \bmod k} \sim \mathbb{1}$	$([x \mapsto y \cdot v], [y \mapsto x \cdot v^{-1}])$	if $x \notin \bar{z}, \exists j. k \leq i_j , y$ fresh, $v = r_1^{-\lfloor i_1/k \rfloor} \dots r_n^{-\lfloor i_n/k \rfloor}$	(4)
$U_0 \wedge U_1$	$\mapsto_{\bar{z}} U'_0 \wedge \theta U_1$	(θ, ϕ)	if $U_0 \mapsto_{\bar{z}} U'_0$	(5)

Figure 7. Plugin constraint-solving algorithm

possibly generating some fresh rigid variables in the process, but with a substitution ϕ_g that relates them back to the original variables. Here \bar{z} is empty because rigid variables may be simplified using the given constraints; they will contain no unification variables. The simplified givens U'_g are discarded.

Next, the substitution θ_g is applied to the wanted unit constraints U_w (in order to eliminate rigid variables if possible), then they are rewritten according to the algorithm, producing a simplified set of constraints U'_w and a substitution θ_w . At this point, only the ‘touchable’ unification variables may be instantiated, so \bar{z} contains all the free variables that are not listed in $\bar{\alpha}_{\text{tch}}$.

Finally, the residual constraints returned by the rule consist of the unchanged non-unit wanteds Q_w , the simplified unit wanteds U'_w and the constraint form of the substitution θ_w . The substitution ϕ_g , which eliminates any fresh rigid variables introduced when simplifying the unit givens, is applied where necessary.

For example, suppose we have

$$U_g = \{a^2 \sim b^3\}, \quad U_w = \{\gamma^3 \sim a\}, \quad \bar{\alpha}_{\text{tch}} = \{\gamma\}$$

where a and b are rigid variables and γ is a unification variable. Rewriting the given constraint generates a fresh rigid variable c and produces $\theta_g = [a \mapsto c^{-3}, b \mapsto c^{-2}]$, $\phi_g = [c \mapsto a \cdot b^{-2}]$. Applying θ_g leaves us with the wanted constraint $\gamma^3 \sim c^{-3}$, which is easily solved by $\theta_w = [\gamma \mapsto c^{-1}]$. In order to eliminate the variable c introduced by simplifying the given constraint, we apply ϕ_g , so we end up with the solution $\phi_g \theta_w = [\gamma \mapsto a^{-1} \cdot b^2]$.

4.3 Soundness and generality

As discussed in subsection 3.2, `OutsideIn(X)` type inference is sound (i.e. it infers correct types for terms) and delivers principal types (i.e. any type that can be given to the term is an instance of the inferred type), under certain assumptions on the behaviour of the simplifier. These assumptions lead to conditions that the algorithm described above must satisfy.

The conditions are formulated in terms of the \Vdash relation, which satisfies the properties in Figure 3. To justify the soundness and generality of the plugin, additional inference rules are required stating that `Unit` is an abelian group, as shown in Figure 6. The `ASSOCIATIVE`, `IDENTITY`, `COMMUTATIVE` and `INVERSE` rules are the usual abelian group laws; the role of `TORSION-FREE` (characterising *free* abelian groups) will be discussed later.

Type safety depends on the fact that the \Vdash relation is consistent (i.e. it cannot prove that two observably distinct types are equal). Consistency is not threatened by the group laws, because they refer only to type families without equations (as defined in subsection 2.1). If \otimes was a constructor rather than a type family, however, it would be possible to derive a contradiction.

To simplify the technical development, in the the following I assume that the constraint entailment relation \Vdash satisfies the follow-

ing condition. This holds for the concrete entailment relation used by Vytiniotis et al. (2011).

Suppose $\mathcal{Q} \Vdash \mathcal{E}_\theta$. Then $\mathcal{Q} \Vdash \theta Q$ if and only if $\mathcal{Q} \Vdash Q$, (P) and $\mathcal{Q} \wedge \theta Q \Vdash Q'$ iff $\mathcal{Q} \wedge Q \Vdash Q'$.

The basic result about the rewrite system needed to show that solutions are both sound and most general is the following, which amounts to showing that rewriting produces equivalent constraints, assuming the appropriate substitution holds as a collection of equations in each direction. I will write $Q_0 \leftrightarrow Q_1$ to mean that the constraints Q_0 and Q_1 are equivalent in the sense that $Q_0 \Vdash Q_1$ and $Q_1 \Vdash Q_0$.

Lemma 1 (Soundness and generality of unification steps). *If $U_0 \mapsto_{\bar{z}} U_1$ (θ, ϕ) then $U_1 \wedge \mathcal{E}_\theta \leftrightarrow U_0 \wedge \mathcal{E}_\phi$.*

Proof. By induction on the definition of the \mapsto relation.

For rule (1), we need to show $u \cdot v^{-1} \sim \mathbb{1} \leftrightarrow u \sim v$, which follows straightforwardly from the group axioms. Similarly, rule (2) is trivial.

For rule (3), the interesting part is showing $x^k \cdot u^k \sim \mathbb{1} \leftrightarrow x \sim u^{-1}$. The `TORSION-FREE` rule means that $(x \cdot u)^k \sim \mathbb{1}$ implies $x \cdot u \sim \mathbb{1}$.

For rule (4), we must show that

$$y^k \cdot w \sim \mathbb{1} \wedge x \sim y \cdot v \leftrightarrow x^k \cdot r_1^{i_1} \dots r_n^{i_n} \sim \mathbb{1} \wedge y \sim x \cdot v^{-1}$$

where $v = r_1^{-\lfloor i_1/k \rfloor} \dots r_n^{-\lfloor i_n/k \rfloor}$ and $w = r_1^{i_1 \bmod k} \dots r_n^{i_n \bmod k}$, which follows from the fact that $r_1^{i_1} \dots r_n^{i_n} \sim v^{-k} \cdot w$.

For rule (5), we must show $U'_0 \wedge \theta U_1 \wedge \mathcal{E}_\theta \leftrightarrow U_0 \wedge U_1 \wedge \mathcal{E}_\phi$. By induction we have $U'_0 \wedge \mathcal{E}_\theta \leftrightarrow U_0 \wedge \mathcal{E}_\phi$, and property (P) gives $\theta U_1 \wedge \mathcal{E}_\theta \leftrightarrow U_1$. \square

In addition, the following lemma shows the relationship between the two substitutions θ and ϕ produced by the algorithm: applying ϕ to θ yields equations that follow from the input constraints U_0 .

Lemma 2. *If $U_0 \mapsto_{\bar{z}} U_1$ (θ, ϕ) then $U_0 \Vdash \mathcal{E}_{\phi \circ \theta}$.*

Proof. By induction on the definition of the \mapsto relation. The only rules that extend the substitution θ are (4), for which $x \sim u^{-1}$ follows by `TORSION-FREE`, and (5), for which the composition is $[y \mapsto x \cdot v^{-1}] \circ [x \mapsto y \cdot v] = [x \mapsto (x \cdot v^{-1}) \cdot v]$, the identity up to the group axioms. \square

From these results, which extend inductively in the obvious way to multiple reduction steps, it follows that the constraint solver is sound in the sense required by `OutsideIn(X)`.

Theorem 1 (Soundness). *If $\mathcal{Q}; Q_g; \bar{\alpha}_{\text{tch}} \vdash^P Q_w \rightsquigarrow Q_r$ then $Q_g \wedge Q_r \Vdash Q_w$.*

Proof. We need to show that

$$\mathcal{Q} \wedge (Q_g \wedge U_g) \wedge Q_w \wedge \phi_g (U'_w \wedge \mathcal{E}_{\theta_w}) \Vdash Q_w \wedge U_w$$

and since $Q_w \Vdash Q_w$ it is sufficient to show

$$U_g \wedge \phi_g (U'_w \wedge \mathcal{E}_{\theta_w}) \Vdash U_w.$$

We are justified in reasoning up to unit normal forms since if u and v are equivalent normal forms then $\epsilon \Vdash u \sim v$ so (R8) gives $\epsilon \Vdash [a \mapsto u] \tau \sim [a \mapsto v] \tau$.

Lemma 1 gives $\theta_g U_w \wedge \mathcal{E}_{\phi_w} \leftrightarrow U'_w \wedge \mathcal{E}_{\theta_w}$, so from (R3) we have $\phi_g (U'_w \wedge \mathcal{E}_{\theta_w}) \Vdash \phi_g (\theta_g U_w)$. Moreover Lemma 2 gives $U_g \Vdash \mathcal{E}_{\phi_g \circ \theta_g}$, so property (P) gives the required entailment. \square

Principality is more interesting, however. This requires that the constraint solver delivers most general solutions, which intuitively means that it makes no ‘guesses’ that are not implied by the original wanted constraints. Vytiniotis et al. (2011) define ‘guess-free solutions’ as those where the wanted constraints entail the residual constraints; for the plugin this amounts to requiring

$$\mathcal{Q} \wedge (Q_g \wedge U_g) \wedge (Q_w \wedge U_w) \Vdash Q_w \wedge \phi_g (U'_w \wedge \mathcal{E}_{\theta_w}).$$

Unfortunately, this is not true! Consider the sole wanted constraint $\alpha^2 \sim \beta^3$, which according to the algorithm in Figure 7 can be solved by $[\alpha \mapsto \gamma^3, \beta \mapsto \gamma^2]$ where γ is a fresh unification variable. The guess-free solution condition would require us to show

$$\alpha^2 \sim \beta^3 \Vdash \alpha \sim \gamma^3 \wedge \beta \sim \gamma^2,$$

and this is simply not derivable, because the fresh variable γ has been conjured out of thin air.

So what has become of the claim that the steps of the algorithm described above are all most general? The unifier delivered for $\alpha^2 \sim \beta^3$ is indeed the most general unifier, in the sense that any unifying substitution for this equation must agree with it up to the laws of an abelian group. The fresh variable γ is not really free: it is constrained by the group laws to be $\alpha \cdot \beta^{-1}$, but the definition of guess-free solutions does not allow us to make use of that knowledge.

Instead, we can prove a weaker result:

Theorem 2 (Generality). *If $\mathcal{Q}; Q_g; \bar{\alpha}_{\text{tch}} \vdash^P Q_r \rightsquigarrow Q_s$ then $Q_g \wedge Q_r \Vdash \psi Q_s$ for some substitution ψ for the freshly introduced variables, i.e. with domain $\text{fuv}(Q_s) \setminus \text{fuv}(Q_r)$.*

Proof. Taking $\psi = \phi_w$ we must show that

$$\mathcal{Q} \wedge (Q_g \wedge U_g) \wedge (Q_w \wedge U_w) \Vdash \phi_w (Q_w \wedge \phi_g (U'_w \wedge \mathcal{E}_{\theta_w}))$$

is derivable. Now $\text{dom}(\phi_w) \# \text{fuv}(Q_w)$ and $Q_w \Vdash Q_w$, so it is enough to show $U_g \wedge U_w \Vdash \phi_w \phi_g (U'_w \wedge \mathcal{E}_{\theta_w})$.

Lemma 1 gives $\theta_g U_w \wedge \mathcal{E}_{\phi_w} \leftrightarrow U'_w \wedge \mathcal{E}_{\theta_w}$, so by (R3) we have $\phi_w \phi_g (\theta_g U_w \wedge \mathcal{E}_{\phi_w}) \leftrightarrow \phi_w \phi_g (U'_w \wedge \mathcal{E}_{\theta_w})$, which implies that $\phi_w \phi_g \theta_g U_w \Vdash \phi_w \phi_g (U'_w \wedge \mathcal{E}_{\theta_w})$. Now we must have $\text{dom}(\phi_w) \# \text{fuv}(\phi_g \theta_g U_w)$ so $\phi_g \theta_g U_w \Vdash \phi_w \phi_g (U'_w \wedge \mathcal{E}_{\theta_w})$. Moreover Lemma 2 gives $U_g \Vdash \mathcal{E}_{\phi_g \circ \theta_g}$ so we can deduce the required entailment using property (P). \square

That is, the solution found by the algorithm may not be guess-free in the original sense, but there is some substitution for the fresh variables it introduces by which it can be transformed into a guess-free solution. I conjecture that this weaker property is in fact sufficient for the proof that `OutsideIn(X)` type inference (if it succeeds) delivers principal types.¹⁶

The underlying problem here is that `OutsideIn(X)` does not have a clear notion of scope for type variables: it is not the case that

$$\alpha^2 \sim \beta^3 \leftrightarrow \alpha \sim \gamma^3 \wedge \beta \sim \gamma^2,$$

¹⁶Theorems 3.2 and 5.2 of Vytiniotis et al. (2011)

but rather we must *contextualise* the variables, as in

$$\exists \alpha. \exists \beta. \alpha^2 \sim \beta^3 \leftrightarrow \exists \alpha. \exists \beta. \exists \gamma. \alpha \sim \gamma^3 \wedge \beta \sim \gamma^2.$$

In fact the same problem shows up in the algorithm described by Vytiniotis et al. (2011), which reduces the wanted constraint $F(G(x)) \sim y$ to $F(\beta) \sim y \wedge G(x) \sim \beta$ where F and G are type families and β is a fresh unification variable; it would appear that

$$F(G(x)) \sim y \not\vdash F(\beta) \sim y \wedge G(x) \sim \beta$$

contrary to their Lemma 7.2.

On another note, observe that the proofs relied on an additional rule, `TORSION-FREE`, beyond the usual laws of an abelian group. This is crucial for proving both that solutions to wanted constraints are most general, and that simplifications of given constraints are sound. It amounts to restricting models of `Unit` to being *free* abelian groups, i.e. those generated by the base units and abelian group laws but with no other equations.¹⁷

Without `TORSION-FREE`, the addition of an axiom $\text{kg} \otimes \text{kg} \sim \mathbb{1}$ would be consistent, but then it would no longer be most general to solve the wanted $\alpha \otimes \alpha \sim \mathbb{1}$ with $\alpha \sim \mathbb{1}$, nor would it be sound to simplify the given $a \otimes a \sim \mathbb{1}$ to $a \sim \mathbb{1}$, as in either case `kg` is an alternative solution.

5. Related work

The design of `uom-plugin` owes a lot to Andrew Kennedy’s implementation of units of measure in `F#`, and Richard Eisenberg’s `units Haskell` library. I compare it with each of them in turn. While there are several other Haskell libraries for units of measure¹⁸, making slightly different design choices, `units` represents the state of the art and the comparison is broadly representative.

5.1 Units of measure in F#

The plugin described in this paper provides support for units of measure that is inspired by, and broadly comparable with, Kennedy’s implementation in `F#`: constants and numeric types can be annotated with their units, units may be polymorphic, and unit equations that arise during typechecking are solved by abelian group unification. However, working in a Haskell setting introduces many new feature interactions to explore, notably with typeclasses, GADTs, type families, higher-kinded and higher-rank types. For example, Haskell allows definitions that are polymorphic in type *constructors* of kind `Unit → *`.

On the other hand, while the GHC typechecker plugins support makes some exciting new things possible, there is still more work to do before a plugin can extend GHC with a completely new language feature. In particular, Template Haskell quasiquotation allows the introduction of new syntax (e.g. for expressions containing quantities with units, or types mentioning units), but this syntax will not be used in output (such as error messages or inferred types). Thus the user can write

```
[u | 5 m/s |] :: Quantity Int [u | m/s |]
```

but the inferred type of this expression is the less easy to read

```
Quantity Int (Base "m" ⊗ Base "s")
```

¹⁷Free abelian groups are always torsion-free, and torsion-free finitely generated abelian groups are free.

¹⁸Notably `dimensional` by Björn Buckwalter (<https://dimensional.googlecode.com/>) and `unittyped` by Thijs Alkemade (<https://bitbucket.org/xnyhps/haskell-unittyped/>).

Moreover, there is no way to simplify an inferred type in a domain-specific manner. Thus a type may be presented as

$$\text{Num } a \Rightarrow \text{Quantity } a \text{ (Base "s" } \otimes \text{ (Base "m" } \otimes \text{ Base "s"))}$$

rather than the (equivalent)

$$\text{Num } a \Rightarrow \text{Quantity } a \text{ (Base "m")}$$

It should be relatively straightforward to extend GHC’s plugin support to allow extensions to pretty-printing and presentation of inferred types, but there will always be limitations of the plugin technique compared to building support into the language, as in F#.

5.2 The `units` package

Another key inspiration for this work is Richard Eisenberg’s `units` library (Muranushi and Eisenberg 2014), which is the state of the art as far as units of measure in Haskell are concerned. As discussed in section 2.3, `uom-plugin` is able to achieve better type inference behaviour and more comprehensible error messages than `units` thanks to the use of a typechecker plugin, rather than encoding everything using type families and other existing GHC features. On the other hand, since `units` does not require a plugin it is more broadly compatible and avoids the potential for plugin-introduced bugs. Moreover, it makes use of Template Haskell to permit a relatively nice input syntax.

Another crucial difference in library design is that `units` is based around working with dimensions (such as length and mass), rather than units directly. A dimension has a ‘canonical’ unit that determines how quantities are represented, but they may be introduced or eliminated using other units, with appropriate conversions performed automatically. There is even support for working with multiple local coherent systems of units (choices of canonical units for dimensions) in different parts of a single program. This allows code to be typechecked for dimension safety, but remain polymorphic in the particular units, and makes it easier to avoid numeric overflow errors when working with quantities at vastly different scales.

In the interests of simplicity, the `uom-plugin` library follows F#’s approach of indexing types by units of measure alone, not including dimensions, but the approach described in this paper should be able to scale to handle dimensions. The best way to represent them, and provide features such as automatic conversion between units of the same dimension, is a matter of ongoing work.

6. Conclusion

In this paper, I have introduced the notion of typechecker plugins both as an implementation technique in GHC and in terms of the `OutsideIn(X)` framework. I have made use of this to define a library for units of measure with good type inference properties, in particular the ability to find most general solutions to constraints arising from unit polymorphism.

Practical use of plugins is still at an early stage, as they are quite low level and closely tied to GHC’s constraint solver. There is much to do to build better abstractions on top of the low-level interface, and hence make it easier to write plugins without deep knowledge of GHC. Termination of constraint solving in the presence of plugins is a particularly tricky issue. It is quite easy for a poorly written plugin to create an infinite loop, for example by emitting a new but trivial constraint each time it is invoked. Moreover, while evidence generation gives some indication of soundness (albeit not consistency of the axiom system used to produce the evidence), it is hard to ensure that plugins deliver most general solutions to constraints.

Two main avenues for future work are extending the `uom-plugin` library itself, and adding features to GHC that make more powerful plugins possible. I will consider these, then suggest some possible other applications for the concept of typechecker plugins.

6.1 Further support for units of measure

Evidence generation The prototype units of measure plugin does not yet support evidence generation (see subsection 3.1.2); rather, it follows the method of proof by blatant assertion. In principle it should be possible to generate proofs based on the abelian group axioms from Figure 6. This would allow GHC’s `-dcore-lint` option to check that the plugin is generating correct output.

Automatic conversion inference As observed above, representing dimensions and inferring conversions between quantities of the same dimension is a matter of ongoing investigation. This is not essential, because the user can always write their own conversions, but it would be better if they were able to write something like

$$r = \text{convert } ([u \mid 10 \text{ ft/min}] \otimes [u \mid 5 \text{ s}]) \oplus [u \mid 42 \text{ m}]$$

and have the compiler automatically insert the conversion from ft/min to m/s.

One way to encode such automatic conversions is through the definition of a pair of additional type families: `Pack`, which converts a list of (base unit, integer exponent) pairs into the corresponding unit, and `Unpack`, which represents a fully known unit as a list of such pairs (in a canonical order). Thus we have:

```

type family Pack (xs :: [(Symbol, Integer)]) :: Unit where
  Pack [] = 1
  Pack ((b, i) : xs) = (Base b  $\otimes$  i)  $\otimes$  Pack xs
type family Unpack (u :: Unit) :: [(Symbol, Integer)]

Pack [("m", Pos 2), ("s", Neg 1)]
  = Base "m"  $\otimes$  2  $\otimes$  Base "s"
Unpack (Base "m"  $\otimes$  2  $\otimes$  Base "s")
  = [("m", Pos 2), ("s", Neg 1)]

```

(Here \otimes represents exponentiation for units.)

`Pack` can be defined via a standard closed type family, but `Unpack` must be defined specially by the plugin because it observes the structure of the unit. It respects the equational theory on units, and hence does not break type soundness. Together, these type families make it possible to encode the `convert` function and other advanced features using existing GHC Haskell type-level programming techniques. However, once more it becomes a challenge to make error messages simple and comprehensible. It is an interesting challenge to extend the units of measure library further while maintaining a suitable balance between features implemented in the plugin and those encoded using existing functionality.

Construction of quantities It is slightly unsatisfying that constructing literal quantities in a safe way fundamentally requires Template Haskell, rather than it providing mere syntactic sugar. One alternative is to expose the `MkQuantity` constructor to the user, and require them to follow a suitable syntactic discipline in its use: always instantiating its type to concrete units. A way to lift this restriction would be beneficial, but by no means essential.

Termination and completeness On a more theoretical note, it would be nice to prove that the plugin-extended constraint solver terminates, and is complete in an appropriate sense. Unfortunately, both of these are tricky issues in `OutsideIn(X)` even before plugins

are added, and modular reasoning about termination is particularly difficult.

Extending the algebraic structure Finally, while indexing quantities by a single abelian group of units is a reasonable point in the design space, there are other choices for the model of units and quantities. In particular, two oft-requested additional features are support for fractional units and alternative origins. Fractional units are sometimes useful, such as $\sqrt{\text{Hz}}$ (i.e. $\text{Hz}^{1/2}$), which arises when quantifying electronic noise levels. F# 4.0 has recently gained support for fractional units. Units of temperature are the most obvious example where multiple origins need to be considered, since $0C \approx 273K$. It may be possible to handle these by indexing quantities by an abelian group of *translations* (Atkey et al. 2013) in addition to units.

6.2 Extensions to the plugins mechanism

Apart from the constraint solver, there are many other points where it would be useful for typechecker plugins to be able to extend the compiler with domain-specific behaviour:

- control over simplification and presentation of inferred types, as discussed in section 5.1;
- easily defining special reduction behaviour for type families, such as the Unpack type family described in section 6.1;
- manipulating error messages, for example so that a DSL implementor can provide domain-specific guidance on likely reasons for a certain class of error, along the lines of error reflection in Idris (Christiansen 2014);
- generating new definitions for each module of the program.

6.3 Other plugins

This paper described a plugin to support units of measure in GHC, and Iavor Diatchki and Eric Seidel are working on a plugin that handles constraints involving type-level natural numbers by calling out to an SMT solver¹⁹. There are many other potential applications, however:

- Making type inference more powerful, e.g. by treating certain type families as injective, along the lines of Jan Stolarek’s ongoing work on `InjectiveTypeFamilies`²⁰;
- Effect tracking: indexing a monad by the available effects, using a solver for a theory of sets, maps or boolean rings;
- Integers (as opposed to natural numbers) in types, with a solver based on ring unification;
- Typeclasses with non-standard search strategies, rather than the usual instance search, such as `Coercible` (Breitner et al. 2014);
- The `Typeable` class for runtime type representation (Lämmel and Peyton Jones 2003);
- Adding η -laws for type-level tuples or record types;
- Record system extensions, such as extensible records via row polymorphism, or the proposed `OverloadedRecordFields` extension.²¹

¹⁹<https://github.com/yav/type-nat-solver>

²⁰<https://ghc.haskell.org/trac/ghc/wiki/InjectiveTypeFamilies>

²¹<https://ghc.haskell.org/trac/ghc/wiki/Records/OverloadedRecordFields>

In particular, it would be interesting to try factoring out an existing piece of GHC functionality (such as the `Coercible` or `Typeable` typeclasses) into a plugin, increasing modularity. Indeed, in principle one could imagine disabling the entire built-in constraint solver, allowing experimentation with alternative algorithms, although this is likely to be practically difficult as it would require the plugin to represent substitutions more directly.

More generally, the existing typechecker plugin interface is at a relatively low level, requiring the plugin implementer to have a fairly detailed knowledge of the way type inference is implemented in GHC (e.g. to generate evidence using its internal data types). A broader challenge for future work is to find a suitable interface that is both powerful enough to implement special-purpose constraint solver behaviour, and simple enough to make the creation of domain-specific constraint solvers accessible to more users. Hopefully it should be possible to build such a higher-level interface on top of the existing typechecker plugins support in GHC.

Acknowledgments

Richard Eisenberg kindly gave detailed feedback on an early draft of this paper, in addition to offering inspiration through his `units` library. My thanks go to Neil Ghani and Conor McBride for ongoing discussion of this work. The implementation of typechecker plugins in GHC was made possible through the work of Iavor Diatchki and Eric Seidel. I’m grateful to attendees at the Scottish Programming Languages Seminar for their feedback. This paper was typeset using `lhs2TeX`²². The work was funded by the University of Strathclyde EPSRC Impact Acceleration Account.

References

- Robert Atkey, Patricia Johann, and Andrew Kennedy. Abstraction and invariance for algebraically indexed types. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL ’13, pages 87–100. ACM, 2013.
- Joachim Breitner, Richard A. Eisenberg, Simon Peyton Jones, and Stephanie Weirich. Safe zero-cost coercions for Haskell. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming*, ICFP ’14, pages 189–202. ACM, 2014.
- David Raymond Christiansen. Reflect on your mistakes! Lightweight domain-specific error messages, 2014. URL <http://www.itu.dk/people/drc/drafts/error-reflection-submission.pdf>. Submitted to post-proceedings of the Symposium on Trends in Functional Programming 2014.
- Jacques Garrigue. Relaxing the value restriction. In Yukiyo Kameyama and Peter J. Stuckey, editors, *Functional and Logic Programming*, volume 2998 of *LNCS*, pages 196–213. Springer, 2004.
- Adam Gundry. *Type Inference, Haskell and Dependent Types*. PhD thesis, University of Strathclyde, 2013. URL <http://adam.gundry.co.uk/pub/thesis/>.
- Andrew Kennedy. *Programming Languages and Dimensions*. PhD thesis, University of Cambridge, 1996. URL <http://research.microsoft.com/en-us/um/people/akenn/units/ProgrammingLanguagesAndDimensions.pdf>.
- Andrew Kennedy. Types for units-of-measure: Theory and practice. In Zoltán Horváth, Rinus Plasmeijer, and Viktória Zsóka, editors, *Central European Functional Programming (CEFP ’09)*, LNCS 6299, pages 268–305. Springer, 2010.
- Ralf Lämmel and Simon Peyton Jones. Scrap your boilerplate: A practical design pattern for generic programming. In *Proceedings of the 2003 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation*, TLDI ’03, pages 26–37. ACM, 2003.

²²<http://www.andres-loeh.de/lhs2tex/>

- Daan Leijen and Erik Meijer. Domain specific embedded compilers. In *Proceedings of the 2Nd Conference on Domain-specific Languages, DSL '99*, pages 109–122. ACM, 1999.
- Takayuki Muranushi and Richard A. Eisenberg. Experience report: Type-checking polymorphic units for astrophysics research in Haskell. In *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell, Haskell '14*, pages 31–38. ACM, 2014.
- Martin Sulzmann, Martin Müller, and Christoph Zenger. Hindley/Milner style type systems in constraint form. Technical Report ACRC-99-009, University of South Australia, School of Computer and Information Science, July 1999.
- Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones, and Kevin Donnelly. System F with type equality coercions. In *Types in Language Design and Implementation (TLDI '07)*, pages 53–66. ACM, 2007.
- Dimitrios Vytiniotis, Simon Peyton Jones, Tom Schrijvers, and Martin Sulzmann. OutsideIn(X): Modular type inference with local assumptions. *Journal of Functional Programming*, 21(4–5):333–412, 2011.
- Brent A. Yorgey, Stephanie Weirich, Julien Cretin, Simon Peyton Jones, Dimitrios Vytiniotis, and José Pedro Magalhães. Giving Haskell a promotion. In *Proceedings of the 8th ACM SIGPLAN workshop on Types in Language Design and Implementation (TLDI '12)*, pages 53–66. ACM, 2012.