

Experience Report: Type-checking Polymorphic Units for Astrophysics Research in Haskell

Takayuki Muranushi

The Hakubi Center for Advanced Research
Kyoto University
muranushi.takayuki.3r@kyoto-u.ac.jp

Richard A. Eisenberg

University of Pennsylvania
eir@cis.upenn.edu

Abstract

Many of the bugs in scientific programs have their roots in mistreatment of physical dimensions, via erroneous expressions in the quantity calculus. Now that the type system in the Glasgow Haskell Compiler is rich enough to support type-level integers and other promoted datatypes, we can type-check the quantity calculus in Haskell. In addition to basic dimension-aware arithmetic and unit conversions, our units library features an extensible system of dimensions and units, a notion of dimensions apart from that of units, and unit polymorphism designed to describe the laws of physics. We demonstrate the utility of units by writing an astrophysics research paper. This work is free of unit concerns because every quantity expression in the paper is rigorously type-checked.

Categories and Subject Descriptors J.2 [Physical Sciences and Engineering]: Astronomy; Physics; D.3.3 [Programming Languages]: Language Constructs and Features—Data types and structures; D.3.2 [Programming Languages]: Language Classifications—Haskell

Keywords Type families; type-level computation; Haskell; quantity calculus

1. Introduction

In 1983, when Canada was moving from the Imperial to the metric system, a Boeing 767 ran out of fuel in the midst of a flight and was on the verge of crashing, all because the airplane was given 22,300 pounds of fuel where 22,300 kg was actually needed [15]. In 1999, a NASA spacecraft Mars Climate Orbiter was lost after it entered too low an orbit, because the software involved mistakenly interpreted the thruster force in pounds of force when a measurement in Newtons was assumed [17]. These are just samples of the many incidents that have been caused by a mistake in units and dimensions.* Because of the possibility of such mistakes, scientists routinely check for both dimensions and units as they perform cal-

* See, <http://lamar.colostate.edu/~hillger/unit-mixups.html>, from where the anecdotes above are adapted.

culations. Kennedy sums it up well in saying that “units-of-measure are to science what types are to programming” [9].

Type systems that enforce the correct handling of dimensions and units have been around for some time [6, 8, 16]. Haskell, a language renowned for its rich type system, naturally has multiple packages that enable type-checking units-of-measure.[†] However, as we attempt to describe real astrophysical reasoning in Haskell, we became aware of an important concept that no existing library supports — unit polymorphism.

Any meaningful law of physics holds in any system of units. For example, the *mass* of fuel consumed in a flight equals the *length* of the air route divided by the economy of the aircraft (*length* travelled per *volume* of fuel) times *density* (*mass* per *volume*) of the fuel. Slanted words all refer to *dimensions*, and the equation holds no matter whether specific *units* are kilograms or pounds. Such unit-independent laws are common in the quantity calculus, and unit polymorphism naturally arises when we type-check our work. We demonstrate three distinct merits of unit polymorphism (§ 3.4):

1. We can faithfully express unit-independence of laws of physics.
2. We can write libraries that deal with quantities without forcing users to adopt specific choice of units.
3. We can perform calculations that would have resulted in overflows/underflows, had we used the default choice of units.

This report addresses the implementation of a unit-polymorphic quantity calculus system not by extending Haskell’s type system directly, but by defining a library built with the features that Haskell already supports. We describe the package units, and our experiences in using this package to support astrophysics research.

Our contributions are as follows:

- In order to type-check the quantity calculus in Haskell’s type system, we summarize the design principles and specifics of the units package, which is available on Hackage (§ 3).
- We compare units to unittyped, a similar library that does not support unit polymorphism (§ 4).
- With units, we can now write functions that deal with abstract dimensions in Haskell. We demonstrate the use of units library by writing an astrophysics paper in Haskell (§ 5).

2. Background: terminology

We describe here some terminology taken from the quantity calculus, which we will use throughout this experience report. These definitions are somewhat informal; please see “International Vocabulary of Metrology” ([1], hereafter VIM) for the full details.

[†] dimensional, dimensional-tf and unittyped among others. For a list of these, see http://www.haskell.org/haskellwiki/Physical_units

Quantity [VIM 1.1] A quantity can be thought of something that can be measured. A quantity has an associated dimension and is expressed in terms of a magnitude and a unit. For example, *diameter* and *wavelength* are both quantities of the dimension *length*. Examples of quantity values are 3.2 kg and 9.8 m/s².

Dimension [VIM 1.7] A dimension is a name given to a group of quantities that differ only by numerical factors. Examples include *length*, *mass*, or *velocity*. Dimensions can be composed of other dimensions; a *velocity* is a *length* over a *time*.

Unit [VIM 1.9] A unit is a particular amount of a quantity, with which all other quantities of the same dimension can be measured. Examples of units include meters, feet, and kilograms. Note that a unit always measures a unique dimension, but that a dimension may have several different units available.

Quantities of the same units can be added. Any two quantities can be multiplied, with their units and dimensions changing accordingly. For example, 50 mph = 100 mile/2 hr, at dimension $velocity = length^1 \times time^{-1}$.

Coherent System of Units (CSU) [VIM 1.14] A coherent system of units consists of a chosen set of base units and base dimensions. Other units are derived by taking the product of powers of base units. *Coherent* units are direct products of powers of base units, while *off-system* units require conversion factors other than 1.

For example, the International System of Units (SI) is a CSU that consists of seven base units and dimensions. A Joule (=1 kg · m²/s²) is the coherent derived unit of *energy* in SI, while a liter (=10⁻³ m³) is an off-system unit of *volume* with respect to SI.

Numerical Value [VIM 1.20] The numerical value of a quantity is the number in the quantity represented in a certain unit. For example, the numerical value of 9.8 m/s² is 9.8.

3. The units package

3.1 Design goals

What would distinguish units from other quantity calculus packages? What design principles were we aiming for? These ideals are presented roughly in order of importance.

Strong type-checking: Under no circumstances should a quantity representing a length be allowed to be added to a quantity representing a time, for example. Multiplication of these quantities is allowed, of course, yielding a multiplicative dimension of $length \times time$.

Extensibility: A key attribute of units is that it is fully extensible. The core system defines only one dimension, *Dimensionless*, with its unit *Number*. All other units are defined purely using the interface exported by units. This makes units easily applicable to domains beyond physics, such as finance and management (Imagine *ManMonth*, a unit of dimension *Labor*.)

Simple user-visible types: After a small amount of work to set up the dimensions and the units, users should be able to write simple type signatures on their code. Additionally, declaring new dimensions and units should be formulaic and not require expert knowledge of Haskell's type system.

Unit polymorphism: Definitions should not be tied to a user's particular choice of units. For example, take the following definition of kinetic energy E_k :

```
kineticEnergy :: Mass -> Velocity -> Energy
kineticEnergy m v = 0.5 * |m| * |v| * |v|
```

The property that $E_k = \frac{1}{2}mv^2$ remains true whether we're measuring mass in kilograms, velocity in meters per second, and energy in Joules, or not. Thus, we want enough flexibility in our types to be able to give *kineticEnergy* a type more like

```
kineticEnergy :: Mass u -> Velocity u -> Energy u
```

where the type variable *u* encodes details of the particular units at hand. Here, the *dimensions*, but not the *units* are fixed – *kineticEnergy* is a dimension-monomorphic, unit-polymorphic definition.

Unit conversions: A user should be able to specify known quantities in any convenient choice of units and request a result in any convenient choice of units. How many years does the Sun take to complete its galactic orbit of 27,200 light-years radius, given its speed 220 km/s? Such a calculation should not require the user to think about unit conversions — the type system should handle it for him.

3.2 Interface

How would one keep track of dimensions and units in a type? Let's focus on dimensions, first. A dimension can be thought of as a set of base dimensions paired with integer exponents. For example, a *velocity* is a $length^1 \times time^{-1}$, and *energy* is $mass^1 \times length^2 \times time^{-2}$. To make these pairs easier to manage, we introduce a *Factor* kind*

```
data kind Factor = F * Z
```

where *Z* is our kind for type-level integers. Here, and throughout our implementation, we make liberal use of algebraic data kinds, promoted from datatypes, as introduced by Yorgey et al. [20].

Is the kind of a dimension *Dimension*? No — we choose to use $*$ because of its extensibility and convenience. Users can, in any module, freely declare new datatypes with **data**, and these datatypes serve as dimensions. For example, we might say

```
data Length = Length
```

and use the type *Length* as a dimension in a factor.

For these reasons, the first argument to our quantity type, of kind [*Factor*], represents the set of dimension factors.

But, what about units? Units control exactly what number is internally stored. If we are storing the length of a day in seconds, we would store 86,400. If we were storing in hours, we would store 24. However, we still want to keep the dimensions primary, not the units, in order to allow for unit-polymorphic programming.

The solution: store a finite map from dimensions to units in the type. We call this mapping a local CSU, or LCSU. It is local because, as we will see, it can vary in different places in our program.

Finally, we wish to allow users to choose their own underlying numerical type, so we index by this choice of representation.

We thus have the *Qu* type, defined as follows:

```
newtype Qu (dimFactors :: [Factor])
          (lcsu :: LCSU) (n :: *) = Qu n
```

The *dimFactors* parameter is our type-level list of *Factors*. The parameter *n* is the underlying numerical representation; functions creating and eliminating *Qus* require *n* to be in the *Fractional*

* Throughout this report, we imagine that Haskell has a richer syntax for dealing with kinds. Here, we declare a datakind, which is a datatype that can be used only at the kind level. Alas, this is not yet possible — one must declare a datatype to get a datakind. Because $*$ is only a kind, not a type, using $*$ as we have done here is impossible. The true definition of *Factor* is **data Factor star = F star Z**, and whenever we use *Factor*, we must write *Factor star*. These details are distracting, so we ignore them from here on out.

class. Note that *Qu* is a **newtype** — there is no runtime cost to type-checking your quantities.

An *LCSU* is stored as a type-level association list, mapping dimensions to units:

```
data kind LCSU = MkLCSU [(*, *)]
```

Here, we see that units, like dimensions, are denoted with elements of kind \star . The situation for units is analogous for that for dimensions — we want the set of units to be modular and extensible; hooking into Haskell’s ability to do this with datatype definitions is the easiest way of achieving these goals.

Dimension and Unit classes When a user wishes to declare a new dimension or unit, she must also make instances of the *Dimension* or *Unit* class, as appropriate. The details of how these work internally are beyond the scope of this report, but they are critically necessary for type-checking and automatic unit conversion. It is worth noting, however, that a *Unit* instance must declare the related dimension and define the relationship to a convertible unit, if any — the rest of the unit conversions are inferred from this one known link.

Unit and dimension combinators The units package defines a small set of combinators used to build compound units and dimensions from simpler ones. These are $*$ for multiplication, $:/$ for division, $^{\wedge}$ for exponentiation, and $:@$ for prefixing. Thus, the following definitions are all quite logical:

```
type Velocity = Length :/ Time -- divide two units
type Kiloqram = Kilo :@ Gram -- Kilo is prefix, not unit
type Area = Length :^ Two -- Two is type-level 2
```

We must distinguish $:@$ (prefixing) from $*$ (unit multiplication) because *Kilo :@ Gram* can, for example, be added to something expressed in *Grams*, but *Meter :* Gram* cannot.

Setting up the Qu type Getting all the details correct for the *Qu* type is challenging. The units package exports a number of type synonyms to make this easier. Here is a more complete example of how it all fits together:

```
data LengthDim = LengthDim
instance Dimension LengthDim
data TimeDim = TimeDim
instance Dimension TimeDim
data Meter = Meter
instance Unit Meter where
  -- there is no known convertible unit
  type BaseUnit Meter = Canonical
  type DimOfUnit Meter = LengthDim
data Second = Second
instance Unit Second where
  type BaseUnit Second = Canonical
  type DimOfUnit Second = TimeDim
type MyLCSU = MkLCSU `[(LengthDim, Meter),
                    (TimeDim, Second)]
type Length = MkQu_DLN LengthDim MyLCSU Double
type Time = MkQu_DLN TimeDim MyLCSU Double
type Velocity = MkQu_DLN (LengthDim :/ TimeDim)
                    MyLCSU Double
distance :: Velocity -> Time -> Length
distance v t = v |*| t
```

The *MkQu_DLN* type synonym creates an appropriately instantiated *Qu* type, suitable for use in the type signature for *distance*. The *DLN* suffix means that *MkQu_DLN* expects a dimension, an *LCSU*, and a numerical type. Other combinations are also exported.

Introducing and eliminating quantities Quantities are made using the (%) operator and eliminated using the (#) operator:

```
(%) :: (ValidDLU dim lcsu unit, Fractional n)
     => n -> unit -> Qu dim lcsu n
(#) :: (ValidDLU dim lcsu unit, Fractional n)
     => Qu dim lcsu n -> unit -> n
```

Both of these functions take a unit parameter. This is why the unit (and dimension) types are not empty; it is convenient to use the one data constructor to represent the type. Continuing the example above, we could write

```
speedOfLight :: Velocity
speedOfLight = 299792458 % (Meter :/ Second)
```

The operator (%) treats its numerical parameter as being measured in terms of the unit specified; the value is converted to be measured in the unit specified in the *LCSU*. The runtime information about an *LCSU* necessary for this conversion is buried inside the *ValidDLU* definition. (*ValidDLU* asserts that a dimension can be used with a certain *LCSU* and choice of unit.) The operator (#) is analogous, but operates in the other direction.

Quantity combinators All the standard mathematical operations are available on quantities: addition, subtraction, multiplication and division (both among quantities and with scalars), exponentiation, comparisons, and roots. Operator names are decorated by vertical bar(s) on the side(s) quantities are applied.

Because exponentiation and root-taking change the type-level integers, the exponent must be known at compile time. This reflects the fact that only nondimensional quantities can be raised to variable powers. We use singletons (as generated by the singletons package [4]) to represent the exponents.

The types of these combinators ensure the correctness of the operations. For example, the type of $|+|$, the addition operator, checks to make sure the dimensions of both operands are equivalent.

3.3 Implementation

The gritty details of how the implementation tracks these types and makes all the pieces fit together are beyond the scope of this report. We pause to note a few points that may be of general interest:

Type-level integers We use our own type-level integer kind *Z* extensively. This kind uses a unary (Peano-style) representation with type families implementing standard operations. There is no enhanced reasoning/solving capability with our *Z* kind, but this is not a problem in practice. We chose not to use GHC’s built-in type-level naturals for two reasons: 1) we require integers, not natural numbers; and 2) there is not yet enough support for runtime access to the build-in naturals. Specifically, if n_1 and n_2 are known at runtime (that is, we have *KnownNat* n_1 and *KnownNat* n_2), we cannot get $n_1 + n_2$ at runtime (that is, GHC cannot solve for *KnownNat* ($n_1 + n_2$)).

Closed type families Our implementation depends critically on closed type families [5] to be able to consider a type-level list of *Factors* to be a set. Closed type families give us the ability to write a Boolean type-level equality predicate on types of kind \star . This predicate, in turn, is necessary when defining the set-contains operation; we need to know when two types do *not* equal in order to recur. The whole approach of units — using types of kind \star to encode an extensible set — would be infeasible without closed type families.

3.4 Examples of unit polymorphism

Writing unit-independent quantity expressions So, how many pounds of gasoline did the pilot actually need? Given the travel

distance *dist* (in miles), the fuel economy of the aircraft *eco* (miles per gallon) and gasoline density *gasden* (pounds per gallon), it is easy to write the function using units:

```
-- Length, Density, and Mass are from from an SI library
-- D.Area refers to the dimension Area, not the type Area
-- MOne is a type-level number for (-1)
type PerArea = MkQu_DLN (D.Area :^ MOne)
gasMass :: Fractional n => Length l n -> PerArea l n
        -> Density l n -> Mass l n
gasMass dist eco gasden = dist | / | eco | * | gasden
```

Note that the type is parametric in *l*, the LCSU. This means that the function will work regardless of the units we wish to use to express the given dimensions. Also, because no conversions need to be done during this computation, we do not need any constraints on *l* at all – run-time knowledge about *l* is needed only when creating a quantity or extracting a quantity’s numerical value.

Avoiding over/underflows by domain-specific scaling One crucial application of the LCSU mechanism is to provide consistent scaling of quantities for astronomers or chemists, whose research objects are much larger or smaller compared to daily scales. The use of SI units in such fields might cause overflows / underflows and produce nonsensical results.

For example, the Lennard-Jones potential is a model of molecular interaction commonly used in chemistry. The model takes two parameters ϵ and σ and gives the attractive force *F* as a function of the distance *r* between two molecules:

$$F = \frac{24\epsilon\sigma^6}{r^7} - \frac{48\epsilon\sigma^{12}}{r^{13}}$$

units can readily model this formula:

```
ljForce :: Energy l Float -> Length l Float
        -> Length l Float -> Force l Float
ljForce eps sigma r
= (24 * | eps | * | sigma | ^ pSix) | / | (r | ^ pSeven)
  | - | (48 * | eps | * | sigma | ^ pTwelve) | / | (r | ^ pThirteen)
```

Then, how many Newtons is the attractive force between two Argon atoms at distance 4 Å? If we type the following into GHCi

```
λ> let sigmaAr = 3.4e-8 % Meter
     epsAr     = 1.68e-21 % Joule
     r         = 4.0e-8 % Meter
λ> (ljForce epsAr sigmaAr r :: Force SI Float) # Newton
```

we get a blunt *NaN* as answer. This is because the Lennard-Jones model involves the inverse 13th to 6th powers of lengths, and $(10^{-8})^6$ is already out of range of the single precision float. (Recall 1 Å = 10^{-8} m.)

Define an LCSU that is more suitable for chemistry, say

```
type CU = MkLCSU '[ '(Length, Angstrom),
                    '(Mass, ProtonMass), '(Time, Pico :@ Second)]
```

and you will get a meaningful response just by replacing *SI* with *CU*.

```
λ> (ljForce epsAr sigmaAr r :: Force CU Float) # Newton
9.3407324e-14
```

Laws with such powers are very common, and there are many applications so computationally heavy that the use of double precision is not an option. Most CPUs can perform single precision float arithmetic twice as fast as double precision, and GPUs (Graphic Processing Units) can do even more (4–20 times). As an extreme case, GRAPE series [3, 12], the gravity-force-specific computers

widely used by astrophysicists, adopt fixed-point real number representations for speed. On such cutting-edge hardware where real number range is traded off for better performance, users are responsible for choosing the right scale, and units supports doing so.

4. The untyped package

We evaluate the success of units by comparing it with untyped, a different approach to type-checking units-of-measure in Haskell. Indeed, before discovering units, one of the authors (TM) had contributed to untyped with physics applications in mind. We give a brief overview of untyped in this section. The differences from a user’s point of view are summarized in Table 1.

untyped, like units, expresses compound dimensions and units as lists of pairs of key types and type-level integers. The type for quantity in untyped is:

```
Value (dim :: [(*, Z)]) (uni :: [(*, Z)]) (n :: *)
```

which resembles units’s counterpart*:

```
Qu (dim :: [Factor]) (lcsu :: LCSU) (n :: *)
```

However, the two libraries treat dimension-unit correspondence in different ways. The *Value* type constructor of untyped takes both the dimension and unit as arguments. A type class *Convertible*’ asserts that *uni* is of the dimension *dim*, and then provides the conversion factor for the unit. Every arithmetic operation requires a *Convertible*’ constraint.

On the other hand, *Qu* of units takes only the dimensions. The corresponding unit can be computed on demand by referring to the LCSU dictionary. A *Value* can equally represent coherent and off-system units, while *Qu* is designed to express only coherent unit quantities. Off-system units in units appear only at conversion to/from numerical values. Various type constraints are still needed for unit conversions, but none is needed for arithmetic. This is because coherent units can be operated on without conversion.

Another difference is in the method of type-level computation. untyped uses functional dependencies while units uses closed type families (Table 1, (2)). Critically, the way untyped uses functional dependencies appears to take advantage of a long-standing GHC bug,[†] which is fixed in GHC 7.8. (Table 1, (3)) Essentially, the functional dependencies in untyped overlap in much the same way as the type family equations do in units, as needed to implement type-level sets. However, overlapping equations in a closed type family has a reasonable, well-defined semantics; such overlap over unordered functional dependencies does not. It is an open question whether the untyped approach can be brought up to date with GHC 7.8. The authors conjecture that this is indeed not possible.

The discontinuity of the build is one reason why TM stopped using untyped. The other is the lack of unit polymorphism, discovered while using untyped. Filling this gap motivated units’ design.

4.1 The gasMass example in untyped

Let’s see what the *gasMass* example would look like when written in untyped. We still want the definition to be unit-polymorphic. A natural starting point is to put type variables in places of unit types.

* Note that *Factor* is equivalent to $(*, Z)$

[†] Fixing the bug (<https://ghc.haskell.org/trac/ghc/ticket/2247>) is by implementing the liberal coverage condition [18].

(1) quantity calculus package	untyped	units
(2) means of type level computation	functional dependencies	closed type families
(3) current status	buildable on GHC 7.6.3 but not on 7.8.1	buildable on GHC 7.8.2
(4) type signature of a quantity	$Value (dim :: [(*, Z)]) (uni :: [(*, Z)]) (n :: *)$	$Qu (dim :: [Factor]) (lcsu :: LCSU) (n :: *)$
(5) convert to quantity	$x = mkVal 5.2 :: Gram Double$	$x = 5.2 \% Gram$
(6) convert to numerical value	$val x$	$x \# Gram$
(7) extract x (in g) in kg	$val (autoc x :: Value [(Mass, POne)] [(Kilo Gram, POne)] Double)$	$x \# kilo Gram$
(8) pretty-print x (in g) in kg	$pp (fmap (0.001*) x) ++ "kg"$	$ppIn (kilo Gram) x$
(9) define a type synonym for a compound unit (spectral radiance)	$type SR = Value [(Mass, POne), (Time, NTwo)] [(Watt, POne), (Meter, NTwo), (Hertz, NOne)]$	$type SR = Watt :/(Meter :^ Two) :/ Hertz$

Table 1: Comparison of features in untyped and units. Code fragments are abbreviated for the sake of readability.

```
gasMass :: Value '[ '(Length, POne)] len n
  → Value '[ '(Length, NTwo)] fe n
  → Value '[ '(Mass, POne), '(Length, NThree)] den n
  → Value '[ '(Mass, POne)] mass n
gasMass dist eco gasden = dist | / | eco | * | gasden
```

Alas, it doesn't compile! The correct program is as follows.

```
gasMass :: (Fractional n
, Convertible '[ '(Length, POne)] len
, Convertible '[ '(Length, NTwo)] fe
, Convertible '[ '(Length, PThree)] vol
, Convertible '[ '(Mass, POne), '(Length, NThree)] den
, Convertible '[ '(Mass, POne)] mass
, MapNeg fe nfe -- nfe = 1 / fe
, MapMerge len nfe vol -- vol = len * nfe
, MapMerge vol den mass -- mass = vol * den
) ⇒ Value '[ '(Length, POne)] len n
  → Value '[ '(Length, NTwo)] fe n
  → Value '[ '(Mass, POne), '(Length, NThree)] den n
  → Value '[ '(Mass, POne)] mass n
gasMass dist eco gasden = dist | / | eco | * | gasden
```

We need 9 lines of type constraints and 4 lines of types for just one line of value-level computation. This result does not argue in favor of strong typing.

Here, the result unit type *mass* is not just universally quantified, but actually depends on *len*, *fe*, and *den* in a complicated way — hence all the constraints above. The number of constraints required are at least twice the number of arithmetic operations in the function body. This quickly renders unit polymorphism in untyped impractical.

5. Writing an Astrophysics Research Paper in untyped and units

In this section, we report the experience of writing an astrophysics paper, titled “Observation of Lightning in Protoplanetary Disks by Ion Lines” [14]. Notably, the paper draft is written in the form of a Haskell program that performs calculations, plots figures, generates \LaTeX sources for the slides and the paper. In addition to the use of units, we aim to further assure the consistency of the paper by using the same program both to generate the \LaTeX formulae in the paper and to perform the actual computation.

The development history of the paper consists of two versions, broadly speaking. The paper draft was initially written using untyped and later rewritten using units. It took 5 working days to

rewrite the draft. The untyped and units version are 2447 and 2462 lines of Haskell code, respectively. As a result, we have two programs that achieve the same goal using different libraries, ideal material for a comparative study.

Dimension and unit as independent concepts Working with untyped, we had difficulty dealing with dimensions and units separately. In untyped, the kind of *dim* and *uni* are both $[(*, Z)]$. Writing “functions” for kinds other than $*$ requires type-level programming, which is still awkward in Haskell. In untyped, the library-supported way to wrap *dim*, *unit* $:: [(*, Z)]$ into kind $*$ is by *Value*, which takes a dimension and a unit in pair.

In units, on the other hand, there are two ways to refer to units and dimensions. One is by unit and dimension combinators (§ 3.2) that produces values of kind $*$. The other is by type-level list of kind $[Factor]$. The ability to refer to dimensions or units separately, and the existence of two kinds of representations makes units more expressive.

Defining type synonyms for compound units (c.f. Table 1, (9)) is a frequently used technique in our work (28 such synonyms are defined). In untyped, we must specify the destination unit in the type signature, along with its dimensions. Unit conversions (c.f. Table 1, (7)) are another place where we want to specify only units. In units, we specify only units; the dimension for the unit is inferred. The combinators over kind $*$ offer ease to the programmer by allowing value-level units (for example $Meter :^ pTwo$) that are quite similar to their type-level counterparts ($Meter :^ Two$). The more-internal types of kind $[Factor]$ make the implementation tractable.

Numerical value accessibility In untyped, users can convert a numerical value to a quantity, or vice versa, using two functions $mkVal :: n \rightarrow Value d u n$ and $val :: Value d u n \rightarrow n$ (c.f. Table 1, (5,6)). Moreover, the type constructor $Value d u$ is a *Functor*, so users can apply any numerical function to the numerical value of a quantity.

On the other hand, units is designed to limit user access to numerical values, in order to protect unit consistency. The recommended way for converting between quantities and numerical values is by use of the $(\%)$ and $(\#)$ operators (c.f. § 3.2), which explicitly take a unit expression as an argument. The direct use of the *Qu* constructor is discouraged.* $Qu d \ell$ is not a *Functor*. $Qu d \ell n$ is in the *Floating* class — the type class that allows use of transcendental functions such as *sin* and *exp* — only if *n* is *Floating* and *d* is dimensionless. This meets with our physical expectations of when these functions are applicable.

* units still exports the *Qu* constructor in module *Data.Metrology.Unsafe*

The expression $x \# \text{Number}$ converts x to numerical value with type-level assertion that x is a dimensionless quantity (c.f. Table 1, (6)). The untyped counterpart, $\text{val } x$, has no such assertion. Users can still add the assertion by adding a type signature to val , but since it makes the expression longer, users cease to use it. Importantly, the shortest possible expression for conversion function in units, $(\# \text{Number})$, comes *with* an assertion, while in untyped val , is *without*.

In untyped, 94 mkVals and 19 vals were used to convert quantities of various units to numbers, and vice versa. They are changed to $(\#)$ s and $(\%)$ s with corresponding units in units. Each update represents more static checking for the code powering the paper.

Pretty Printing Pretty printing is important because it is where the results are presented to outside world. We symbolically refer to a pretty-printer as pp . Several flavor of pps are used in the paper [14], 63 times in total.

Working with untyped, the pretty printer converted only the numerical value to L^AT_EX expressions, and the unit symbols were typeset “by hand” outside of the checked system. Moreover, the combined use of the pretty printer and fmap , such as in $\text{pp } (\text{fmap } (0.001\star) x) ++ \text{"kg"}$ (Table 1, (8)), occurs 3 times in the untyped-based source code. Here, the user avoids the use of in-place unit conversion (which is tedious in untyped) and manipulates the numerical value and the unit symbol directly. With the pretty printer controlling only the numerical value, we cannot assert the correctness of the quantity expression at the L^AT_EX-level.

In the units version, we redefined the pretty-printers so that each takes a unit and a quantity as arguments, and prints the quantity in the given unit, followed by the unit symbol. This new design allows mapping of quantities from Haskell to L^AT_EX in the correct sense, since a quantity consists of a magnitude and a measurement unit. In untyped, it was difficult to write pretty-printers that take the unit as an argument, since the value-level representation of units is absent.

Error Messages The readability of the error messages is important to help the users locate and fix unit mistakes.

Error messages of units are relatively instructive, compared to those of untyped. The separation of dimensions and units is one factor that contributes to the readability because it effectively halves the length of the printout of corresponding types. The error messages, such as `Couldn't match type 'Time' with 'Length'`, are quite suggestive.

Two major shortcomings in the clarity of error messages are with type-level integers (which print in unary), and the SI LCSU (which prints in seven lines). It is unclear how to improve this without more control over the error message generation system.

Unit Coherence Writing with untyped, the quantities in the paper ranged over multiple systems of units: the cgs (centimeter-gram-second) units, SI units, and domain-specific units such as astronomical units (AUs) and electron volts. Several components of the paper depend on multiple systems of units. At each such place, we have to decide on which system to use and to convert other units. The resultant paper draft was chimera of unit systems.

This undesirable situation disappeared when working with in units. We choose the SI as the LCSU in all parts of the paper. We could still define and pretty-print quantities using arbitrary units. The automatic conversion afforded by $(\%)$ and $(\#)$ allowed us to use our units of choice.

Summary of Comparison We ended our experience with a few takeaways:

- An effective way to achieve unit polymorphism is by representing a quantity as triples of its dimension, the system of units it belongs to, and its numerical value type.
- Having units and dimensions be representable as types of kind \star makes them easier to work with, as we do not have to bother with proxies.
- A library can be designed to hide the value constructor for the quantity type, to refrain users from unit-unsafe computations, and yet remain user friendly. It is important that the shortest ways to write quantity calculus are unit-safe, and are short enough.

6. Related work: Comparison to F#

The units package provides an embedded domain-specific type system in Haskell, and we have compared it against another, similar approach. Here, we compare units against F#'s approach to type-checking units.

F# has built-in support for type-checking quantities, as derived from the work by Kennedy [8]. Kennedy's work does not include a separation between dimensions and units, and thus does not support the unit polymorphism as described in this paper.

Here is some F# code that converts between temperature units:*

```
[<Measure>] type degC // temperature, Celsius
[<Measure>] type degF // temperature, Fahrenheit

let convertCtoF (temp : float<degC>) =
    9.0<degF> / 5.0<degC> * temp + 32.0<degF>
let convertFtoC (temp : float<degF>) =
    5.0<degC> / 9.0<degF> * (temp - 32.0<degF>)
```

There are several immediate advantages to F#'s approach over units's. Computation uses normal arithmetic operators, such as \star and $/$; no need for the more verbose $|*$ and $|/$. We also see that declaring new units is much easier in F# than it is with units – no need for a *Unit* and *Dimension* classes.

F# falls short in one substantial way, however: it does not support dimension-monomorphic unit polymorphism. It is possible, for example, to express gasMass in F# while remaining polymorphic over the units:

```
let gasMass (dist : float<'len>)
    (eco : float<1 / 'len^2>)
    (gasden : float<'mass / 'len^3>)
    : float<'mass> = dist / eco * gasden
```

However, this definition is dimension-*polymorphic*. Nothing constrains the units len or mass is bound to – they might not represent length and mass. The power of units's implementation of polymorphism allows the writer of gasMass to restrict calls of that function to have appropriate units.

The lack of notion of dimensions also means that F# is unable to do any implicit unit conversion. Any conversion must be specified by hand-written multiplication by conversion factors, and they are not distinguished from ad-hoc conversion factors between units of different dimensions. The following three functions are equally accepted by F#. Note the bad conversion factor in the second line:

```
let convg2kg (x : float<g>) = x / 1000.0<g / kg>
let convg2kg' (x : float<g>) = x / 10.0<g / kg>
let convmm2kg (x : float<mm>) = x / 39.17<mm / kg>
```

Error messages As expected, the built-in nature of F#'s support for type-checking quantities leads to better error messages. For

* Adapted from <http://msdn.microsoft.com/en-us/library/dd233243.aspx>

example, the following lines of F# produce the error messages below, respectively:

```
let x = 4<m> + 2<s>
let vel : float<m / s> = 5.0<m> * 3.0<s>
```

The unit of measure 's' does not match the unit of measure 'm'
The unit of measure 'm/s' does not match the unit of measure 'm s'

Contrast to equivalent code in Haskell:

```
x = (4 % Meter) |+|(2 % Second)
vel = (5 % Meter) |*|(3 % Second) :: Velocity
```

Couldn't match type 'D.Length' with 'D.Time'
Couldn't match type 'S 'Zero' with 'P 'Zero'

The first Haskell error message is quite helpful. The second is less so, for two reasons: numbers are printed in unary and the user has to deduce that the error message is discussing the exponent assigned to the time dimension.

Let's try constructions that are dimension-polymorphic. We'll write a function *prod* that multiplies a list of quantities together. The result must accordingly be the product of all of the units, but we'll say that the result has the same measure as the input list. Here is the F#:

```
let prod (qs : float<'u> list) : float<'u> =
    List.fold (*) 1.0 <-> qs
```

warning: This construct causes code to be less generic than indicated by the type annotations. The unit-of-measure variable 'u' has been constrained to be measure '1'.

And now the Haskell:

```
prod :: [Qu d l n] -> Qu d l n
prod = foldr (|*|) (1 % Number)
```

Couldn't match type 'd' with '' []'

Interestingly, this is *not* an error in F#. Like its ancestor ML, F# does not issue a hard error when a declared type is too polymorphic. Haskell's error more cryptically suggests the same problem as F#'s warning, that the function makes sense only when the quantities are in fact dimensionless. However, it would take a good measure of experience for the units programmer to figure this out from the error message.

Malformed environments One area where F# clearly is superior is in the creation of the unit environment. As we've seen above, declaring new units in F# is easy with a [**<Measure>**] type declaration. However, declaring new units in Haskell with units is not so easy: one must declare dimension and unit types, with *Dimension* and *Class* instances. This process is easy to get wrong, especially because dimensions, units, and ordinary types all share the kind ***. The error messages when working in a malformed environment are horrific.

For example, suppose a novice user just declared a *Meter* type, with no dimension and no *Unit* instance:

```
data Meter = Meter
type Length = MkQu 'U Meter
add1 :: Length -> Length
add1 x = x |+|(1 % Meter)
```

```
Could not deduce
(Subset
 (CanonicalUnitsOfFactors (UnitFactorsOf Meter))
 (CanonicalUnitsOfFactors
 (LookupList (DimFactorsOf (DimOfUnit Meter))
 'DefaultLCSU))
 ,Subset
 (CanonicalUnitsOfFactors
 (LookupList (DimFactorsOf (DimOfUnit Meter))
 'DefaultLCSU))
 (CanonicalUnitsOfFactors (UnitFactorsOf Meter)))
 arising from a use of ''
```

Even the designers of units cannot easily deduce the fix – add a proper *Unit* instance – from this message. F#'s built-in syntax for unit declarations avoids this problem entirely. Future work with the units package includes some Template Haskell functions to produce correct dimension and unit definitions, as well as a Template Haskell function to check the sanity of the environment.

7. Conclusion

The units library allows us to work with a domain-specific system of units, like the chemists' units above. We can use multiple LCSUs in the case of interdisciplinary study such as astrochemistry, units's design of CSU-local computation and manual inter-CSU conversion encourages users to create large LCSU blocks and minimize conversions between them.

The assertion that all unit conversions are eliminated under each LCSU helps us to optimize the underlying computation. Hardware accelerators such as GPUs are becoming popular today, and multiple Haskell libraries for parallel array computation on CPUs and GPUs have been proposed [2, 7, 10, 11, 19]. Expressing laws of hydrodynamics in Haskell has been demonstrated [13]. Such libraries, if used together with easy, consistent scaling and the correctness check of units provided by units, constitute a powerful development environment for computational science.

What is more, we can write a well-documented library with the basic equations of high-school physics without choosing any particular system of units — the library would take the user's choice as a parameter. Putting all of this together, we will be able to teach and study physics in Haskell.

Acknowledgments

We thank Thijs Alkemade for the development of the untyped library that greatly helped in writing the untyped version of the astrophysics paper; Iavor Diatchki for the improvements of the type-level naturals in GHC upon our request; and the anonymous reviewers for detailed and helpful feedback. This material is based upon work supported by the National Science Foundation under Grant No. 1116620.

References

- [1] Bureau International des Poids et Mesures. International vocabulary of metrology: Basic and general concepts and associated terms. *JCGM*, pages 1–91, 2012.
- [2] M. M. T. Chakravarty, G. Keller, S. Lee, T. L. McDonell, and V. Grover. Accelerating Haskell array codes with multicore GPUs. In *Declarative Aspects of Multicore Programming (DAMP '11)*, pages 3–14, 2011.
- [3] T. Ebisuzaki, J. Makino, T. Fukushige, M. Taiji, D. Sugimoto, T. Ito, and S. K. Okumura. GRAPE Project: an Overview. *Publications of the Astronomical Society of Japan*, 45:269–278, June 1993.
- [4] R. A. Eisenberg and S. Weirich. Dependently typed programming with singletons. In *Haskell Symposium '12*, pages 117–130, 2012.

- [5] R. A. Eisenberg, D. Vytiniotis, S. Peyton Jones, and S. Weirich. Closed type families with overlapping equations. In *Principles of Programming Languages (POPL '14)*, pages 671–683, 2014.
- [6] P. Guo and S. McCamant. Annotation-less unit type inference for C. *Final Project, 6.883: Program Analysis*, 2005.
- [7] G. Keller, M. M. T. Chakravarty, R. Leshchinskiy, S. Peyton Jones, and B. Lippmeier. Regular, shape-polymorphic, parallel arrays in Haskell. In *SIGPLAN Not.*, volume 45, pages 261–272. ACM, 2010.
- [8] A. Kennedy. *Programming Languages and Dimensions*. PhD thesis, University of Cambridge, 1996.
- [9] A. J. Kennedy. Types for units-of-measure: Theory and practice. In Z. Horváth, R. Plasmeijer, and V. Zsók, editors, *Central European Functional Programming School*, volume 6299 of *Lecture Notes in Computer Science*, pages 268–305. Springer, 2010. .
- [10] B. Larsen. Simple optimizations for an applicative array language for graphics processors. In *Declarative Aspects of Multicore Programming (DAMP '11)*, pages 25–34, 2011.
- [11] G. Mainland and G. Morrisett. Nikola: Embedding compiled GPU functions in Haskell. *SIGPLAN Not.*, 45(11):67–78, Sept. 2010.
- [12] J. Makino, T. Fukushige, M. Koga, and K. Namura. GRAPE-6: Massively-Parallel Special-Purpose Computer for Astrophysical Particle Simulations. *Publications of the Astronomical Society of Japan*, 55:1163–1187, Dec. 2003.
- [13] T. Muranushi. Paraiso: an automated tuning framework for explicit solvers of partial differential equations. *Computational Science & Discovery*, 5(1):015003, 2012.
- [14] T. Muranushi, E. Akiyama, S. Inutsuka, N. Hideko, and S. Okuzumi. Observation of lightning in protoplanetary disks by ion lines. *The Astrophysical Journal, to be Submitted*, 2014.
- [15] W. H. Nelson. The Gimli glider. *Soaring Magazine*, 1997.
- [16] P. Roy and N. Shankar. Simcheck: An expressive type system for simulink. In *NASA Formal Methods*, pages 149–160, 2010.
- [17] A. G. Stephenson, D. R. Mulville, F. H. Bauer, G. A. Dukeman, P. Norvig, L. S. LaPiana, P. J. Rutledge, D. Folta, and R. Sackheim. Mars climate orbiter mishap investigation board Phase I report, 44 pp. NASA, Washington, DC, 1999.
- [18] M. Sulzmann, G. J. Duck, S. Peyton-Jones, and P. J. Stuckey. Understanding functional dependencies via constraint handling rules. *Journal of Functional Programming*, 17:83–129, 2007.
- [19] J. Svensson, M. Sheeran, and K. Claessen. Obsidian: A domain specific embedded language for parallel programming of graphics processors. In *Implementation and Application of Functional Languages (IFL '08)*, pages 156–173, 2011.
- [20] B. A. Yorgey, S. Weirich, J. Cretin, S. Peyton Jones, D. Vytiniotis, and J. P. Magalhães. Giving Haskell a promotion. In *Types in Language Design and Implementation (TLDI '12)*, pages 53–66, 2012.