

Testing properties of generic functions

Patrik Jansson

Johan Jeuring

*students of the Utrecht University
Generic Programming class*

Department of Information and Computing Sciences,
Utrecht University

Technical Report UU-CS-2006-043

www.cs.uu.nl

ISSN: 0924-3275

Testing properties of generic functions

Patrik Jansson¹, Johan Jeuring², and students of the Utrecht University
Generic Programming class³

¹ CSE, Chalmers University of Technology, Sweden, patrikj@chalmers.se

² ICS, Utrecht University, the Netherlands, johanj@cs.uu.nl

³ Contributions from Laurence Cabenda, Gerbo Engels, Jacob Kleerekoper, Sander Mak, Michiel Overeem, and Kees Visser.

Abstract Software testing is an important part of the software development process. Testing comes in many flavours: unit testing, property testing, regression testing, contract checking, etc. QuickCheck is probably one of the most advanced tools for testing properties of functional programs. It supports the definition of properties and random test-data generators in Haskell, and checks that a property passes the test cases. A datatype-generic function is parametrised by a type. Examples of generic functions are equality tests, maps and pretty printers. A generic function can be seen as a template algorithm that can be instantiated with (the structure of) a data type. Generic functions satisfy generic properties. This paper discusses testing properties of generic functions. It shows how generic properties can be formulated, and how QuickCheck can be used to test generic properties. Furthermore, it shows how to automatically generate QuickCheck generators using Generic Haskell.

1 Introduction

Software testing aims to find errors in software by means of running the software. It is an important part of the software development process. Testing comes in many flavours: unit testing, property testing, regression testing, contract checking, black-box/white-box testing, etc.

Property testing is based on the fact that functions satisfy properties. Together with a function a programmer writes one or more properties that should be satisfied by the function. For example, consider the following excerpt from a library for manipulating bits.

```
data Bit = O | I deriving (Show, Eq)
bits2int    :: [Bit] → Int
bits2int bs = bits2int' bs (length bs - 1)
  where bits2int' []      n = 0
        bits2int' (x : xs) n = bits2int' xs (n - 1) + bit2int x * 2n
bit2int     :: Bit → Int
bit2int b   = if b == O then 0 else 1
int2bits    :: Int → [Bit]
```

```

int2bits n = if n ≥ 0 then int2bits' n [] else []
  where int2bits' 0 bs = bs
        int2bits' n bs = int2bits' (n `div` 2) (int2bit (n `mod` 2) : bs)
int2bit   :: Int → Bit
int2bit n = if n == 0 then O else I

```

Functions *bits2int* and *int2bits* convert a list of bits to an integer and *vice versa*. To see if these functions are inverses we could check the following properties:

```

prop_int2bits_bits2int   :: [Bit] → Bool
prop_int2bits_bits2int bs = (int2bits . bits2int) bs == bs
prop_bits2int_int2bits   :: Int → Bool
prop_bits2int_int2bits n = (bits2int . int2bits) n == n

```

Checking with a property checker immediately reveals that they don't hold. A counterexample to the first property is $[O, I, I]$ (leading zeroes should be removed in the first property), and to the second property is -3 (negative numbers are not properly encoded in *int2bits*). Since mistakes like these are not at all uncommon in programs, we want support to either prove or test that such properties hold.

QuickCheck [1] is probably one of the most advanced tools for testing properties of functional programs. It supports the definition of properties and random test-data generators in Haskell, and checks that a property passes the test cases. Gast [7] is a tool similar to QuickCheck, but for property testing in Clean [12]. With Gast a user does not have to supply test-data generators: test data is automatically generated for arbitrary data types. For recursive data types, the test data generated by Gast is not of random size. Gast enumerates data in a breadth-first manner, so for example for lists, the first 500 test cases do not contain lists of length longer than four. QuickCheck generates lists of length up to 200 in the same situation.

A generic function is parametrised by a type. Examples of generic functions are equality, map, and pretty printers. A generic function can be seen as a template algorithm that can be instantiated with (the structure of) a data type. Generic functions satisfy generic properties. Just as a generic function, a generic property can be seen as a template property that can be instantiated with (the structure of) a data type to obtain a property. This paper

- discusses testing properties of generic functions. It shows how properties of generic functions in Generic Haskell [3] can be formulated and tested using QuickCheck.
- defines generic QuickCheck generators using Generic Haskell. This means we get the best of both worlds — we combine the strengths of QuickCheck with generic support inspired by Gast.

This paper is organised as follows. Section 2 briefly introduces and compares QuickCheck and Gast. Section 3 introduces generic programming in Generic Haskell. Section 4 shows how QuickCheck is used to check properties of generic

functions. Section 5 discusses the properties that should hold for the functions in the library of Generic Haskell. Section 6 presents different ways of generating test data for arbitrary data types. Section 7 concludes and discusses future work.

2 QuickCheck and Gast

QuickCheck and Gast are tools for checking properties of functions. This section briefly introduces the two tools.

QuickCheck

QuickCheck is an automatic testing tool for Haskell programs. The programmer provides a specification of the program, in the form of properties that functions should satisfy, and QuickCheck then tests that the properties hold in a large number of randomly generated cases. Specifications are expressed in Haskell, using combinators defined in the QuickCheck library. The library provides combinators to define properties, observe the distribution of test data, and define test data generators.

Many properties are simple Boolean functions, implicitly universally quantified over all arguments:

```
prop_PlusAssoc      :: Float → Float → Float → Bool
prop_PlusAssoc x y z = (x + y) + z == x + (y + z)
```

To test a property it is passed to the function *test*:

```
Main> test prop_PlusAssoc
Falsifiable, after 8 tests:
-4.6
-4.0
3.6
```

Here QuickCheck finds a simple counterexample illustrating that finite precision Floats don't behave like ideal real numbers.

The QuickCheck library also provides conditional properties, where tests not satisfying the precondition are discarded:

```
prop_SmallPrime   :: Integer → Property
prop_SmallPrime x = prime x ⇒ x < 100
```

```
Main> test prop_SmallPrime
OK, passed 100 successful tests.
```

Here QuickCheck has generated a few hundred test cases (randomly chosen numbers x) out of which 100 were prime and all of those were <100 . As it turns out, the default generator for integers starts with small numbers (within the interval $[-n/2..n/2]$ after n tests) which for this case means that 100 successful test

cases was not enough to find a counterexample. Fortunately, it is also possible to define a custom generator, using the infinite list of *primes*:

```
primeNumbers :: Gen Integer
primeNumbers = do n ← arbitrary
               return (primes !! abs n)

prop_SmallPrime2 :: Property
prop_SmallPrime2 = forAll primeNumbers (\x → x < 100)
```

```
Main> test prop_SmallPrime2
Falsifiable, after 26 successful tests:
107
```

QuickCheck also supports a simple but powerful way of searching for small counter examples. When a test case fails, QuickCheck tries to shrink the test case until a “local minimum” is found. As an example, for the properties of the bits example in the introduction we get the following results

```
Main> test prop_int2bits_bits2int
Falsifiable, after 2 successful tests
(shrunk failing case 3 times):
[0]
Main> test prop_bits2int_int2bits
Falsifiable, after 1 successful tests:
-3
```

Gast

Gast (Generic Automated Software Testing) [7] is a property testing tool which can be seen as a QuickCheck for Clean. Gast is implemented in the non-strict functional language Clean [12], a close relative to Haskell. From the users perspective, Gast is very similar to QuickCheck — properties can be defined as normal Boolean functions:

```
listsAreShort :: [Int] → Bool
listsAreShort xs = length xs < 5
```

and tests can be run by calling the function *test*:

```
Start = test listsAreShort
```

which in this case results in the answer

```
Passed after 500 tests.
```

This example is chosen to show that some care needs to be taken in interpreting the results from testing: Gast enumerates data in a breadth-first manner, only randomising the order “within each level”. For recursive data types this is problematic, because of the exponential growth of the search space — as we can see,

the first 500 test cases do not contain a single list with more than four elements. QuickCheck generates lists up to length around 200 in the same situation.

The enumeration approach used by Gast does have a number of advantages: it avoids generating the same test case more than once and it makes it possible to actually prove properties over finite domains within the same framework (using exhaustive testing). The Clean implementation of Gast is fast, but for recursive data types the exponential search space means that reaching reasonably sized test cases just takes too long.

Gast does not support shrinking, but there is no need to shrink failing test cases when they are generated and tested in order of increasing size.

3 Generic programming in Generic Haskell

In this section we introduce type-indexed functions by means of an example and we explain how type-indexed functions become generic in Generic Haskell.

Type-indexed functions

A type-indexed function takes an explicit type argument, and can have behaviour that depends on this type argument. For example, suppose the unit type `Unit`, sum type `:+`, and product type `:*` are defined as follows:

```
data Unit    = U
data a :+: b = Inl a | Inr b
data a :* b = a :* b.
```

We use infix type constructors `:+` and `*` and an infix value constructor `:*` to ease the presentation. The type-indexed function `eq` checks equality of two values. We define the function `eq` on booleans, the unit type, sums, and products as follows in Generic Haskell:

$$\begin{aligned} eq\{\mathbf{Bool}\} \quad n1 \quad n2 &= eqBool \ n1 \ n2 \\ eq\{\mathbf{Unit}\} \quad U \quad U &= True \\ eq\{\alpha\ :\!+\!:\ \beta\} \ (Inl\ x) \ (Inl\ y) &= eq\{\alpha\} \ x \ y \\ eq\{\alpha\ :\!+\!:\ \beta\} \ (Inr\ x) \ (Inr\ y) &= eq\{\beta\} \ x \ y \\ eq\{\alpha\ :\!+\!:\ \beta\} \ _ \quad _ &= False \\ eq\{\alpha\ :\!*:\ \beta\} \ (x1\ :\!*:\ y1) \ (x2\ :\!*:\ y2) &= eq\{\alpha\} \ x1 \ x2 \wedge eq\{\beta\} \ y1 \ y2, \end{aligned}$$

where `eqBool` is the standard equality function on Booleans. The type signature of `eq` is as follows:

$$eq\{\mathbf{a} \ :\!*\!:\ _ \} \ :: (eq\{\mathbf{a}\}) \Rightarrow \mathbf{a} \rightarrow \mathbf{a} \rightarrow \mathbf{Bool}.$$

The context $(eq\{\mathbf{a}\}) \Rightarrow$ in this signature says that `eq` has a *dependency* [9] on `eq`. A type-indexed function `f` depends on another type-indexed function `g` if `g` is used on a type argument (a *dependency variable*) `α` in the definition of `f`. The occurrences of `α` and `β` in the definition of `eq` are dependency variables.

Generic functions

A type-indexed function such as `eq` does not only work on the types that appear as its type indices. To see why `eq` is in fact *generic* and works on arbitrary data types, we give a mapping from data types to structure types such as units, sums, and products. It suffices to define a function on structure types in order to obtain a function that can be applied to values of arbitrary data types. If there is no specific case for a type in the definition of a generic function, generic behaviour is derived automatically by the compiler by exploiting the structural representation.

For example, the definition of the function `eq` that is generically derived for lists is equivalent to the following specific definition:

```
eq {α} [] [] = True
eq {α} (x : xs) (y : ys) = eq {α} x y ∧ eq {α} xs ys
eq {α} _ _ = False
```

To obtain this instance, the compiler needs to know the structural representation of lists, and how to convert between lists and their structural representation. We will describe these components in the remainder of this section.

Structure types

The structural representation (or structure type) of types is expressed in terms of units, sums, products, and base types such as integers, characters, etc. For example, for the list and tree data types defined by

```
data [a] = [] | a : [a]
data Tree a b = Tip a | Node (Tree a b) b (Tree a b),
```

we obtain the following structural representations:

```
type [a]◦ = Unit :+: a :*: List a
type Tree◦ a b = a :+: Tree a b :*: b :*: Tree a b,
```

where we assume that `:*` binds stronger than `:+` and both type constructors associate to the right. Note that the representation of a recursive type is not recursive, and refers to the recursive type itself: the representation of a type in Generic Haskell only represents the structure of the top level of the type.

Embedding-projection pairs

If a type `a` can be embedded in, or represented by, another type `b`, a witness of this property can be stored as a pair of functions converting back and forth (an embedding-projection pair):

```
data EP a b = Ep { from :: a → b, to :: b → a }.
```

A type `T` can be embedded in its structure-representation type `T◦`, witnessed by a value `convT :: EP T T◦`. For example, for lists we have that `conv[] = Ep from[] to[]`, where `from[]` and `to[]` are defined by:

$$\begin{aligned}
\text{from}_{[]} &:: \forall a. [a] \rightarrow [a]^\circ \\
\text{from}_{[]} [] &= \text{Inl } U \\
\text{from}_{[]} (x : xs) &= \text{Inr } (x :* xs) \\
\text{to}_{[]} &:: \forall a. [a]^\circ \rightarrow [a] \\
\text{to}_{[]} (\text{Inl } U) &= [] \\
\text{to}_{[]} (\text{Inr } (x :* xs)) &= x : xs.
\end{aligned}$$

The definitions of such embedding-projection pairs are automatically generated by the Generic Haskell compiler for all data types that appear in a program.

Tying the knot

Using structure-representation types and embedding-projection pairs, a call to a generic function on a data type \mathbb{T} is reduced to a call on type \mathbb{T}° . For example, for equality we obtain a function of type $a^\circ \rightarrow a^\circ \rightarrow \text{Bool}$. To convert this function back to a function of type $a \rightarrow a \rightarrow \text{Bool}$ we use the function *bimap* [2]. Hence, if the generic function is defined for view types such as `Unit`, `:+:`, and `:*:`, we do not need cases for specific data types such as `List` or `Tree` anymore. For primitive types such as `Int`, `Float`, `IO` or `→`, no structure type is available. Therefore, for a generic function to work on these types, specific cases are necessary.

Generic abstractions, local redefinitions, and default cases

Generic Haskell supports a number of extensions that simplify defining and using generic functions. First, using a *generic abstraction*, we can define a generic function in terms of another generic function instead of by induction on the structure types. For example, we can test equality of functions by means of the following generic function:

$$\begin{aligned}
\text{feq}\{b :: *\} &:: (\text{eq}\{b\}) \Rightarrow (a \rightarrow b) \rightarrow (a \rightarrow b) \rightarrow a \rightarrow \text{Bool} \\
\text{feq}\{b\} f g &= \lambda x \rightarrow \text{eq}\{b\} (f x) (g x)
\end{aligned}$$

which is a generic abstraction that is defined in terms of, and depends on, the standard generic equality function.

Generic functions may have dependencies. We can use *local redefinition* to redefine the dependencies of generic functions. For example, if we want equality on lists of characters to be case insensitive, we can write

```

equalCaseInsensitive    :: Char → Char → Bool
equalCaseInsensitive x y = toUpper x == toUpper y

let eq{α} = equalCaseInsensitive
in eq{[α]} "Generic Programming" "GENERIC programming"

```

Another way in which we may obtain this behaviour is via a so-called *default case*, which allows us to extend an existing generic function by adding new cases or overriding existing ones.

$$\begin{aligned}
\text{cieq}\{a :: *\} &:: (\text{cieq}\{a\}) \Rightarrow a \rightarrow a \rightarrow \text{Bool} \\
\text{cieq} \text{ extends } eq & \\
\text{cieq}\{\text{Char}\} x y &= \text{toUpper } x == \text{toUpper } y
\end{aligned}$$

Many more examples of these extensions, and a discussion about the merits and disadvantages of these constructs can be found in Löh’s thesis [8].

4 QuickCheck for generic functions

This section explains how we use QuickCheck for testing properties of generic functions. The biggest challenge here is to *formulate* generic properties. We start this section with a number of generic properties, and then discuss how we can use QuickCheck to test them.

Minimal and maximal values

Haskell’s prelude contains a class *Bounded* defined by

```
class Bounded a where
  minBound :: a
  maxBound :: a
```

minBound and *maxBound* are the minimal and maximal value of a type, respectively. They should satisfy

```
prop_minBound x = compare minBound x ≠ GT
prop_maxBound x = compare maxBound x ≠ LT
```

that is, *minBound* is smaller than or equal to any other value, and *maxBound* is larger than or equal to any other value. The method *compare*::*a* → *a* → Ordering, in the class *Ord* (used for totally ordered data types) allows a single comparison to determine the precise ordering of two elements:

```
data Ordering = LT | EQ | GT
```

Haskell allows to derive the bounds automatically for some user-defined data types (enumeration types and single-constructor data types whose constituent types are in *Bounded*). Generic Haskell’s library contains definitions of the generic values *gminBound*{*t*} and *gmaxBound*{*t*} for all algebraic types (not only for those types for which Haskell supports deriving). To formulate a generalisation of the property above, we also need the generic compare function *gcompare* from Generic Haskell’s library. The desired properties now read as follows:

```
prop_gminBound {t :: *} :: (gcompare {t}, gminBound {t}) ⇒ t → Bool
prop_gminBound {t} x = gcompare {t} (gminBound {t}) x ≠ GT
prop_gmaxBound {t :: *} :: (gcompare {t}, gmaxBound {t}) ⇒ t → Bool
prop_gmaxBound {t} x = gcompare {t} (gmaxBound {t}) x ≠ LT
```

Note that the properties are formulated as generic abstractions in Generic Haskell. Since generic properties describe properties of generic functions, all generic properties are defined as generic abstractions.

Properties of map

The generic equivalent *gmap* of the well known *map* function applies zero or more functions (depending on the kind of its data type argument) to the appropriate elements in a value of the data type.

$$gmap\{a :: *, b :: *\} :: (gmap\{a, b\}) \Rightarrow a \rightarrow b$$

For example, the instance of *gmap* on the data type `Tree a b` defined in Section 3 has the following type:

$$gmap\{Tree\} a\ b :: (a \rightarrow c) \rightarrow (b \rightarrow d) \rightarrow Tree\ a\ b \rightarrow Tree\ c\ d$$

Function *gmap* is defined as the deep identity function, and local redefinition is used to obtain *map*-like behaviour. For `tree :: Tree Int Char` we would write

```
let gmap\{\alpha\} = chr
    gmap\{\beta\} = ord
in  gmap\{Tree\} \alpha\ \beta\ tree
```

to convert the integers to characters, and the characters to integers.

Properties of *gmap* can be derived from properties of *map*. Function *map* on lists is a part of a functor, and satisfies the functor laws: it preserves the identity, and distributes over composition.

$$\begin{aligned} map\ id & \quad ===\ id \\ map\ (f . g) & \quad ===\ map\ f . map\ g \end{aligned}$$

where (`===`) is pointwise equality of functions on lists, implemented by `feq\{\[\]\}`, see Section 3. Generalised versions of these properties should hold for the generic *map* function *gmap*.

Function *gmap* preserves the identity, and it distributes over composition. For a type constructor `c :: * → *`, we want to obtain something like

$$prop_gmap_comp\{c\} f\ g = feq\{c\} (map\{c\} f . map\{c\} g) (map\{c\} (f . g))$$

However, we want this property to allow instantiations for types like `· :+: ·` of kind `* → * → *`, where instead of two, four functions are passed as argument. Hinze [2] shows how to generalise these properties for types of arbitrary kinds. For example, for a type constructor `d` of kind `* → * → *` we have

$$\begin{aligned} prop_gmap_id\{d\} & \quad = feq\{d\} (gmap\{d\} id\ id) id \\ prop_gmap_comp\{d\} & \quad = feq\{d\} (gmap\{d\} f\ h . gmap\{d\} g\ j) \\ & \quad \quad (gmap\{d\} (f . g) (h . j)) \end{aligned}$$

The resulting property is *kind-indexed*.

Testing generic properties

As the examples of generic properties for *gmap* show, a generic property may involve kinds, type constructors, polymorphic types, higher-order functions, and plain values. To test a property, we have to supply values for each of the above components. QuickCheck can generate values of monomorphic types and functions, but generating type constructors, let alone kinds, is out of reach. This implies, amongst others, that we have to instantiate the properties on monomorphic types

Happily, generating type constructors and kinds is not necessary. To *prove* a generic property, it suffices to prove instances of the property on the structure types [2]. Similarly, to *test* the validity of a generic property, it suffices to test the validity of a property on the structure types. To test the validity of a property on all structure types, we would have to write a separate instance of the property for each structure type. Take the property *prop_gminBound* as an example. The simplest structure type is `Unit`. For this case, the following expression would be tested:

$$gcompare\{\text{Unit}\} (gminBound\{\text{Unit}\}) U \neq GT$$

By definition of *gminBound* and *gcompare*, this test, and the equivalent test for `Int` and `Char` trivially pass. For the sum type case QuickCheck would need to test something like

$$\begin{aligned} \text{prop_gminBound_Sum } \text{cmpa } \text{cmpb } \text{mba } \text{mbb } x = \\ (\forall a . \text{cmpa } \text{mba } a \neq GT) \implies \\ (\forall b . \text{cmpb } \text{mbb } b \neq GT) \implies \\ (\text{gminBound_Sum } \text{cmpa } \text{cmpb } \text{mba } \text{mbb } x \neq GT) \end{aligned}$$

Since *gminBound* depends on *gcompare* and on itself, *prop_gminBound_Sum* takes five arguments. The last argument is a value of type `a :+: b`, and the other arguments are instances of *gcompare* and *gminBound* on the types `a` and `b`, respectively.

In general, implications $P \implies Q$ may be hard to test in QuickCheck. In particular when the condition P is often *False*, Q is only tested for a few of the generated test cases. For many of the properties this turns out to be a problem — for example, for most properties of equality the condition requires generated values to be equal. For the above property for *gminBound* the problem is even worse, because the left-hand side of the implication includes a local universal quantification which is not implementable with QuickCheck properties. We solve this problem by supplying generators: instead of testing $\lambda x \rightarrow P x \implies Q x$ we test *forAll genP* ($\lambda x \rightarrow Q x$). In general it is hard or impossible to convert a property to a generator, but to obtain testable properties this is a necessary step.

To avoid some of the problems with implications and local quantification, we define a data type which combines the structure types in a single data type, and use that data type for testing the library. The following data type combines the

most important structure types, and is easily extended with more cases for basic structure types.

```

data StructureTypes a = STUnit
                        | STInt Int
                        | STChar Char
                        | STProd (StructureTypes a) (StructureTypes a)
                        | STLabel{ anA :: a}

```

The data type `StructureTypes` contains cases for units, integers, characters, products, and labels. The cases for sums and constructors are implicit, but appear since there is a choice between constructors in the data type, and there are constructor names in the data type. To test the validity of the property `prop_gminBound` with this approach we use the QuickCheck function `test` on the data type `StructureTypes Int`:

```

test (prop_gminBound{StructureTypes Int})

```

QuickCheck generates test cases from the data type `StructureTypes Int` if we provide a generator (an element of `Gen (StructureTypes Int)`). We have used an instance of the generic generator `arb3` (defined later in Section 6).

5 Properties of the Generic Haskell Library

The Generic Haskell library consists of a number of basic generic functions that are used often in generic programs. It is a prelude for generic programs. Many functions of the Generic Haskell library are taken from Haskell’s prelude [11]. This includes functions that implement the methods that are derivable in Haskell, and generalisations of list functions such as `map`, `sum`, `prod`, `and`, etc. Another source of inspiration for the Generic Haskell library is PolyLib [5], the library of PolyP, which contains many basic generic functions.

Since generic functions from the library will often be used as basic building blocks in generic-programming applications, it is important that they are correct. Therefore, the generic functions in the Generic Haskell library are natural candidates for applying our approach to testing generic functions.

The Generic Haskell library consists of twelve modules, of which we will consider the following six: `Eq.ghs`, `Compare.ghs`, `Enum.ghs`, `Bounds.ghs`, and `ReadShow.ghs`, corresponding to the derivable Haskell classes `Eq`, `Ord`, `Enum`, `Bounded`, `Read`, and `Show`, and the module `Map.ghs`, which implements the generic map function `gmap`. We will introduce the generic functions used in this section briefly, often referring to their non-generic Haskell equivalents. More information about the functions in the Generic Haskell library can be found in the user’s guide [10].

It is impossible to give a complete set of properties for the functions defined in the six modules. Furthermore, this paper would become far too long if we would give properties for all of the functions defined in these modules. We intend to add the properties to the library of Generic Haskell though.

Properties of *gread* and *gshow*

Functions *gshow* and *gread* implement the derivable *read* and *show* functions from Haskell. Just as in Haskell, they are defined in terms of helper functions *gshowsPrec* and *greadsPrec*. Reading a value after showing it should be the identity. Showing after reading need not be the identity: the original value might contain concrete syntax (spaces, newlines) that are not generated by the *show* function.

$$\begin{aligned} \text{prop_gread_gshow}\{t :: *\} &:: (eq\{t\}, greadsPrec\{t\}, gshowsPrec\{t\}) \Rightarrow \\ & \quad t \rightarrow \text{Bool} \\ \text{prop_gread_gshow}\{t\} &= feq\{t\} (gread\{t\} . gshow\{t\}) id \end{aligned}$$

where *feq* is pointwise equality of functions, see Section 3.

It turned out that *gread* could not cope with named fields in data types. The `StructureTypes` a data type contains this constructor:

```
data StructureTypes a = ...
  | STLabel{ anA :: a }
```

The *anA* field triggered a runtime error (pattern match failure) in *gread*. QuickCheck does not trap exceptions, so when a property fails, QuickCheck fails instead of just counting this as a failed test case. Fortunately, the Haskell compiler `ghc` includes (unsafe) functions to catch exceptions in pure code, so by wrapping the property in an exception handler returning *False* for all exceptions, we have used QuickCheck to find the bug.

```
Main> test (protect prop_gread_gshow_STInt)
Falsifiable, after 3 successful tests
(shrunk failing case 3 times):
STLabel {anA = -2}
```

The problem was actually not in *gread*, but in *gshow*. There was no space character after the equality sign, so when a negative integer was shown, the two characters "`=-`" were later "parsed" by *gread* as one token. A one-character change to the source code fixed this problem, but revealed another bug, this time in *gread*. Function *gread* did not allow parentheses around `STLabel{ anA = 2 }`, while *gshow* (and the derived *show* in Haskell) printed parentheses. After this second fix, all tests passed.

Properties of *gmap*

Function *gmap* preserves the identity:

$$\begin{aligned} \text{prop_gmap_id}\{t\} &:: (eq\{t\}, gmap\{t, t\}) \Rightarrow t \rightarrow \text{Bool} \\ \text{prop_gmap_id}\{t\} &= feq\{t\} (gmap\{t\}) id \end{aligned}$$

To test this function, we instantiate it on the type `StructureTypes a`.

$$\begin{aligned} \text{prop_gmap_id_ST} &:: (Eq a) \Rightarrow \text{StructureTypes a} \rightarrow \text{Bool} \\ \text{prop_gmap_id_ST} &= \text{let } eq\{a\} = (==) \end{aligned}$$

```

    gmap{a} = id
  in prop_gmap_id{StructureTypes a}

```

Function *gmap* distributes over composition. For a type constructor $c :: * \rightarrow *$, we want to obtain something like

```

prop_gmap_comp{c} f g = feq{c} (map{c} f . map{c} g) (map{c} (f . g))

```

However, we want this property to allow instantiations for types like $\cdot :+ \cdot$ of kind $* \rightarrow * \rightarrow *$, where instead of two, four functions are passed as argument. To use a single function for expressing this property on types of different kinds, we formulate the distributivity property by means of three copies of *gmap*, of which we only define *gmap1* here. Using local redefinition we can instantiate these copies differently for different cases.

```

gmap1 {a :: *, b :: *} :: (gmap1 {a, b}) => a -> b
gmap1 extends gmap
prop_gmap_comp {a :: *, b :: *, c :: *} ::
  (eq{c}, gmap1 {b, c}, gmap2 {a, b}, gmap3 {a, c}) => a -> Bool
prop_gmap_comp {t} = feq{t} (gmap1 {t} . gmap2 {t}) (gmap3 {t})

```

To instantiate this property on the data type `StructureTypes a`, we locally redefine the *gmap* copies.

```

prop_gmap_comp_ST op f g =
  let eq{a} = op
      gmap1 {a} = f
      gmap2 {a} = g
      gmap3 {a} = f . g
  in prop_gmap_comp{StructureTypes a}

```

Properties of *enum*

Function *enum* exhaustively enumerates all possible instances of a particular data type.

```

enum{t :: *} :: (enum{t}) => [t]

```

For example, *enum{Int}* yields the list of all possible (machine-) integers. A property that should hold for this function is the following:

```

prop_enum{t} :: t -> Bool
prop_enum{t} value = value ∈ enum{t}

```

This property says that any value of type *t* should be in the enumeration of that type. Interestingly, checking this property is not really an option — at least for most real-life data types. Recursive data types often have infinitely many values, so using QuickCheck to test whether or not a value appears in the

enumeration may take infinitely long. When testing the property instantiated with the `StructureTypes Int` data type QuickCheck just looped, and at first we thought this was just to be expected. But a more careful examination revealed that the property looped already for the first test case, which should have been small enough to be found early in the enumeration list. It turned out to be a subtle bug in the definition of the generic `enum` function. The enumeration used a version of Cantor diagonalisation which was “non-productive” in the case of infinite lists. By replacing just the diagonalisation function, the generic `enum` implementation worked as expected. Still, the property remains effectively untestable — already some trees built from just seven constructors are more than 10000 elements down the list.

The problem is just another instance of the problem Gast has with coverage for recursive data types (remember that Gast also uses (randomised) enumeration): While every element is *somewhere* in the enumeration list, and will *eventually* be generated by Gast, only small elements are reachable (will be tested by Gast) within reasonable time. Testing the enumeration property with Gast (instead of QuickCheck) is possible but not very useful — it is not very surprising that values (test cases) generated from an enumeration list actually are elements of a very similar enumeration list.

Another property of `enum` relates `enum` to the generic function `empty` that returns the ‘least’ value of a type. For example, for the `List` type `empty` would return the empty list.

```
prop_enum_empty{t} :: Bool
prop_enum_empty{t} = empty{t} ∈ enum{t}
```

As the type signature reveals, this is more a unit test than a QuickCheck property. No random value is generated, so QuickCheck tests the same thing in each test. It would be more interesting to range over different types for `t`, but this does not fit the (current, non-generic) QuickCheck framework.

Properties of gcompare

Function `gcompare` generalises the derivable `compare` function from Haskell. It is often desirable that the `gcompare` function imposes a partial order on its domain. This implies that we want to test whether or not `gcompare` is anti-symmetric ($\forall a, b \in X : aRb \wedge bRa \Rightarrow a = b$), reflexive ($\forall a \in X : aRa$), and transitive ($\forall a, b, c \in X : aRb \wedge bRc \Rightarrow aRc$). Transitivity can be expressed as a QuickCheck property by:

```
prop_gcompare_trans{t :: *} :: (gcompare{t}) => t -> t -> t -> Property
prop_gcompare_trans{t} x y z = gcompare{t} x y == gcompare{t} y z ==>
                                gcompare{t} x y == gcompare{t} x z
```

We use the fact that if we know that `gcompare{t} x y == gcompare{t} y z`, then it suffices to check that `gcompare{t} x y == gcompare{t} x z`, because this implies `gcompare{t} y z == gcompare{t} x z`. We use the QuickCheck conditional operator \implies to rule out non-interesting test cases.

Reflexivity and anti-symmetry can be implemented in a similar fashion.

Another property relates function *gcompare* with the generic equality function *eq*. Function *gcompare* returns *True* for two arguments if and only function *eq* returns *True*.

$$\begin{aligned} \text{prop_gcompare_eq}\{t :: *\} &:: (gcompare\{t\}, eq\{t\}) \Rightarrow t \rightarrow t \rightarrow \text{Bool} \\ \text{prop_gcompare_eq}\{t\} \ x \ y &= (gcompare\{t\} \ x \ y == EQ) == eq\{t\} \ x \ y \end{aligned}$$

This concludes the section on properties for generic functions in the Generic Haskell library. Formulating and testing these properties has been useful: we have discovered three bugs in the library.

6 Generic generators

Normally, QuickCheck requires a user to write a test case generator for a user-defined data type on which QuickCheck is used. Generic programming allows us to automatically generate test cases for any given data type. This makes testing properties of (generic) functions easier.

Porting the Gast generator to Haskell

A generic approach to generating test cases is already available: Gast (Generic Automated Software Testing) [7], written in Clean. We have translated their implementation of pseudo random data generation [6] into Generic Haskell.

$$\text{generate}\{g :: *\} :: \text{Int} \rightarrow \text{StdGen} \rightarrow [g]$$

But we cannot immediately use this result for testing with QuickCheck. The Gast implementation provides an infinite list of elements while QuickCheck expects a generator to have type *Gen g*. To convert the provided list to a *Gen*,

$$\text{QuickCheck.elements} :: [g] \rightarrow \text{Gen } g$$

looks like a good candidate. The function *elements* requires that the input list is finite and it picks random values from the list. So, it does not preserve the order from the generated list.

The required finite list can be obtained by applying *take n* to the input list, for some suitable *n*. QuickCheck uses 100 test cases by default, so one solution could be to let *n* be 100. But this is not good enough: when using implication (\implies) in a QuickCheck property, the condition may rule out some test cases, which means that 100 test cases would not be enough. To be on the safe side we could set *n* to be, say, 100000, but that gives an other problem: picking a random element from the list gives very low probability for ‘edge cases’, like $[-1, 0, 1]$ for *Ints* which are at the beginning of the list. The best solution would be to let *n* be a parameter, start with a small *n* and let it grow while testing. Fortunately, there is a convenient combinator in the QuickCheck library which does just that — it turns a “size-parametrised generator” into a generator:

$QuickCheck.sized :: (Int \rightarrow Gen\ a) \rightarrow Gen\ a.$

Thus if $xs :: [a]$ is the list generated by the (translated) `Gast` generator, we can use $sized (\lambda n \rightarrow elements (take\ n\ xs)) :: Gen\ a$ as a generator for `QuickCheck`.

Thus we obtain a `QuickCheck` generator, written in Generic Haskell, which works for all Haskell data types. But, unfortunately, it has the same weakness for recursive types as the `Gast` generator in that it takes very long before any reasonably sized elements are generated. Worse, where `Gast` can use the systematic generation of test data for exhaustive checking for finite types, `QuickCheck` may generate duplicates (making testing less efficient) and cannot guarantee to generate all elements (incompleteness). Still, it is convenient to have a fully generic generator around, and it can be modified with default cases and local redefinitions to customise its behaviour for selected constructors or types.

Non-terminating generators

Instead of first enumerating and then selecting it should be possible to define a generic generator directly. As a first try we can define the following generic generator:

```

arb1 {a :: *}    :: (arb1 {a}) => Gen a
arb1 {Unit}     = return U
arb1 {Int}      = arbitrary
arb1 {Char}     = arbitrary
arb1 {α :+: β}  = arb_Sum (arb1 {α}) (arb1 {β})
arb1 {α :* β}  = liftM2 (:*) (arb1 {α}) (arb1 {β})

```

```

arb_Sum         :: Gen a → Gen b → Gen (a :+: b)
arb_Sum ga gb = oneof [liftM Inl ga, liftM Inr gb]

```

This generator is very simple, works for all data types and it does generate reasonably sized values, but it has at least two drawbacks: a skewed distribution and possible non-termination.

The first problem is because Generic Haskell encodes multiple-constructor data types with nested binary sums, which means that `arb1` will give a very skewed distribution of the constructors. If p_i denotes the probability of constructor C_i we get $p_i = 1/2^i$ for $i \in \{1..n - 1\}$. Here a symmetric encoding would help and the good news is that the upcoming Generic Haskell release will support this as a `Generic View`. It is possible to work around this problem already in the current version of Generic Haskell by first analysing the data type, but we have not done so.

The second problem is more subtle, but it was noted already in the first `QuickCheck` paper (for a specific `Tree` data type). For recursive data types that branch into more than one subtree, it is fairly easy to accidentally define a generator that often fails to terminate (or, actually, terminates but with an infinite tree as the result). The problem is that if a branching constructor is often generated, the final tree is only finite if all the subtrees are finite and after

a few branches the number of subtrees is high. The skewed distribution offers some degree of protection against these infinite trees, but this Bin data type is an example of the problem:

```
data Bin = B1 Bin Bin | B2 Bin Bin | L.
```

Here the probability to generate L is $1/4$ and the probability for a finite tree is only $1/3$.

A terminating generic generator

The solution to the termination problem is to use *sized* generators — we use a parameter n to limit the size of the generated trees. For a generic function it is not obvious to define what “size” should measure, but one simple choice is the number of constructors in the tree. Using a sized generator, we generate trees of size at most n . The first few cases in the definition are simple generalisations of *arb1*:

```
arb2 {a :: *}      :: (arb2 {a}, empty {a}) => Int -> Gen a
arb2 {Unit}       n = return U
arb2 {Int}        n = resize n arbitrary
arb2 {Char}       n = resize n arbitrary
arb2 {α :+: β}    n = arb_Sum (arb2 {α} n) (arb2 {β} n)
```

Our size measure tells us that we should reduce the size when passing through a constructor and distribute the size over the two subtrees in the product. In the product case it is tempting to just use

```
arb2 {α :+: β} n = liftM2 (:*) (arb2 {α} (n / 2)) (arb2 {β} (n / 2))
```

but that would generate only balanced or even total trees. Instead we divide the size randomly over the two subtree:

```
arb2 {Con c α}    n = liftM Con (arb2 {α} (n - 1))
arb2 {α :+: β}   n
  | n > 1 = do m ← choose (1, n - 1)
             x ← arb2 {α} m
             y ← arb2 {β} (n - m)
             return (x :+: y)
  | n ≤ 1 = return (empty {α} :+: empty {β})
```

This generator works for all data types, it always terminates and generates finite trees (if there are any). It still has the skewed constructor distribution and it has a similar problem with a skewed size distribution for nested products. Both these problems can be avoided with a symmetric view or with an analysis of the data type. Initial experiments are promising, but messy so we leave that for future work.

Better distribution for regular data types

A problem with all the “fully generic” generators is that they cannot treat the recursive case differently from other cases. As an example, the *arb2* generator for a normal list will distribute the size parameter evenly between the element and the tail. For lists we can include a special case in the definition, but similar problems occur also for other data types. Generic Haskell is being extended with Generic Views [4], and using the Fix view it is possible to detect the recursive case, at least for regular data types.

Using the latest (yet to be released) version of Generic Haskell we have implemented yet another (sized) generic generator:

```
arb3 {a :: *} :: Int → Gen a
```

This generator produces finite elements and has an even distribution of constructor probabilities and subtree sizes. The limitation is that it only works for regular data types (no mutual recursion and recursive occurrences must have the same parameters). The code depends on the generic function

```
children {a :: * viewed Fix} :: a → [a]
```

which is the classical example of what could be done in PolyP but cannot be done in the “old” Generic Haskell implementation.

7 Conclusions and future work

We have shown how we can formulate and test properties of generic functions using Generic Haskell and QuickCheck. Furthermore, we have defined a few generic QuickCheck generators.

Since an inductive proof of a property of a generic function only requires cases for the structure types used to represent data types, it suffices to test properties of generic functions on these structure types.

We have implemented a number of properties for generic functions in the Generic Haskell library. Formulating and testing these properties has revealed three bugs in the library. We have not yet completed the description of the properties of the functions in the library, so we expect (but do not hope) to find more bugs.

The generic QuickCheck test data generators produces test data that varies more in size than the test data generated by Gast. But they do generate duplicates.

While implementing the different tests using QuickCheck we encountered a few problems, in particular with exception handling and a better control of the size of generated test cases. It turned out that the latest version from QuickCheck (obtained from CVS) solves some of these problems.

Future work consists of finishing formulating properties for the functions in the Generic Haskell library, and further fine-tuning the generic QuickCheck test data generators. Another idea we would like to investigate is to generate random *types* instead of values, and use these randomly generated types for testing, instead of the `StructureTypes` a type.

References

1. Koen Claessen and John Hughes. QuickCheck: A lightweight tool for random testing of Haskell programs. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming*, pages 268–279. ACM Press, 2000.
2. Ralf Hinze. *Generic Programs and Proofs*. Bonn University, 2000. Habilitationsschrift.
3. Ralf Hinze and Johan Jeuring. Generic Haskell: practice and theory. In Roland Backhouse and Jeremy Gibbons, editors, *Generic Programming, Advanced Lectures*, volume 2793 of *LNCS*, pages 1–56. Springer-Verlag, 2003.
4. Stefan Holdermans, Johan Jeuring, Andres Löh, and Alexey Rodriguez. Generic views on data types. In Tarmo Uustalu, editor, *Proceedings 8th International Conference on Mathematics of Program Construction, MPC'06*, volume 4014 of *LNCS*. Springer-Verlag, 2006.
5. Patrik Jansson and Johan Jeuring. PolyLib – a polytypic function library. In *Workshop on Generic Programming, Marstrand*, June 1998.
6. Pieter Koopman and Rinus Plasmeijer. Generic generation of elements of types. In *6th Symposium on Trends in Functional Programming, TFP 2005: Proceedings*, pages 167–179. Tallinn, 2005.
7. Pieter W. M. Koopman, Artem Alimarine, Jan Tretmans, and Marinus J. Plasmeijer. Gast: Generic automated software testing. In *IFL*, pages 84–100, 2002.
8. Andres Löh. *Exploring Generic Haskell*. PhD thesis, Utrecht University, September 2004.
9. Andres Löh, Dave Clarke, and Johan Jeuring. Dependency-style Generic Haskell. In Olin Shivers, editor, *Proceedings of the International Conference on Functional Programming, ICFP'03*, pages 141–152. ACM Press, August 2003.
10. Andres Löh and Johan Jeuring (editors). The Generic Haskell user’s guide, Version 1.42 - Coral release. Technical Report UU-CS-2005-004, Utrecht University, 2005.
11. Simon Peyton Jones et al. *Haskell 98, Language and Libraries. The Revised Report*. Cambridge University Press, 2003. A special issue of the Journal of Functional Programming.
12. Rinus Plasmeijer and Marko van Eekelen. *Clean Language Report version 2.1*, 2005.