

# GEC: A Toolkit for Generic Rapid Prototyping of Type Safe Interactive Applications

Peter Achten, Marko van Eekelen,  
Rinus Plasmeijer, and Arjen van Weelden

Nijmeegs Instituut voor Informatica en Informatiekunde,  
Radboud Universiteit Nijmegen, Toernooiveld 1,  
6525 ED Nijmegen, The Netherlands  
{P.Achten, rinus, arjenw}@cs.ru.nl, marko@niii.ru.nl

**Abstract.** Programming GUIs with conventional GUI APIs is notoriously tedious. In these notes we present the GEC toolkit in which the programmer can create user interfaces without any knowledge of low-level I/O handling. Instead, he works with Graphical Editor Components (GEC). A GEC is an interactive component that is automatically derived from an arbitrary monomorphic data type, including higher order types. It contains a value of that data type, visualizes it, and allows the user to manipulate it in a type-safe way. The toolkit has a library of data types that represent standard GUI elements such as buttons, text fields, and so on. As a consequence, the programmer works with data types that model the interactive system that he is interested in. Programs are constructed as a collection of communicating GECs. This communication can be arranged in an ad-hoc way, or in a disciplined way, using a combinator library based on arrows. GECs are suitable for rapid prototyping of real world applications, for teaching and for debugging. These notes focus on the *use* of the GEC toolkit for functional programmers, only briefly explaining its inner workings and underlying principles.

## 1 Introduction

In the last decade, Graphical User Interfaces (GUIs) have become *the* standard for user interaction. Programming these interfaces can be done without much effort when the interface is rather static, and for many of these situations excellent tools are available. However, when there is more dynamic interaction between interface and application logic, such applications require tedious manual programming in any programming language. Programmers need to be skilled in the use of a large programming toolkit.

The goal of the *Graphical Editor* project is to obtain a concise programming toolkit that is *abstract*, *compositional*, and *type-directed*. Abstraction is required to reduce the size of the toolkit, compositionality reduces the effort of putting together (or altering) GUI code, and type-directed automatic creation of GUIs allows the programmer to focus on the data model. In contrast to visual programming environments, programming toolkits can provide ultimate flexibility,

type safety, and dynamic behavior within a single framework. We use a *pure functional* programming language (Clean [22]) because functional programming languages have proven to be very suitable for creating abstraction layers on top of each other. Additionally, they have strong support for type definitions and type safety.

Our programming toolkit utilizes the *Graphical Editor Component (GEC)* [6] as universal building block for constructing GUIs. A  $GEC_{\tau}$  is a graphical editor for values of any *monomorphic first-order* type  $\tau$ . This type-directed creation of *GECs* has been obtained by *generic programming* techniques [9,16,15]. With generic programming one defines a family of functions that depend on the structure of types. The *GEC* toolkit project is to our knowledge the first project in which generic programming techniques are used for the creation of GUI applications. It is not the purpose of these notes to explain the inner workings of the *GEC* building blocks. The reader is referred to [6] for that. Instead we focus on the *use* of these building blocks and on *how* the toolkit is built using the basic blocks.

The basic first order *GEC* building blocks from [6] have been extended in two ways, such that we *can* construct higher-order value editors [8]. The first extension uses run-time *dynamic typing* [1,21], which allows us to include them in the *GEC* toolkit, but this does not allow type-directed GUI creation. It does, however, enable the toolkit to use polymorphic higher-order functions and data types. The second extension uses compile-time static typing, in order to gain monomorphic higher-order type-directed GUI creation of *abstract* types. It uses the *abstraction mechanism* of the *GEC* toolkit [7].

Apart from putting all the earlier published work together in a single context, focusing on the use of the toolkit and explaining the extensions using the basic building blocks, these notes also introduce a library for composing *GECs* which is based on the arrows [17] concept. Furthermore, these notes contain exercises at the end of each section to encourage the reader to get familiar with the treated *GEC* concepts.

These notes are structured as follows. Section 2 contains an overview of the basic first-order *GEC* toolkit. In Sect. 3 it is explained how *GECs* can be composed to form larger applications both using *GECs* directly as well as using a new arrows library. The *GEC*-abstraction for model-view programming is treated in Sect. 4. Extensions for working with higher order types, dynamically and statically are covered in Sect. 5. Related work is discussed in Sect. 6 and conclusions are summarized in Sect. 7.

A note on the implementation and the examples in this paper. The project has been realized in Clean. Familiarity with Haskell is assumed, relevant differences between Haskell and Clean are explained in footnotes. The GUI code is mapped to Object I/O [4], which is Clean's library for GUIs. Given sufficient support for dynamic types, the results of this project can be transferred to *Generic Haskell* [19], using the Haskell [20] port of Object I/O [3]. The complete code

of all examples (including the complete *GEC* implementation in Clean) can be downloaded from <http://www.cs.ru.nl/~clean/gec>.

Finally, we need to point out that Clean uses an explicit multiple environment passing style [2] for I/O programming. As *GECs* are integrated with Clean Object I/O, the I/O functions that are presented in these notes are state transition functions on the program state (**PSt ps**). The program state represents the external world of an interactive program, tailored for GUI operations. In these notes the identifier **env** is a value of this type. The uniqueness type system [10] of Clean ensures single threaded use of the environment. To improve the readability, uniqueness type attributes that actually appear in the type signatures are not shown. Furthermore, the code has been slightly simplified, leaving out a few details that are irrelevant for these notes.

## 2 The Basic *GEC* Programming Toolkit

With the *GEC* programming toolkit [6], one constructs GUI applications in a *compositional* way using a high level of *abstraction*. The basic building block is the Graphical Editor Component (*GEC*).

**Graphical Editor Components.** A  $GEC_t$  is an editor for values of type **t**. It is generated with a generic function. The power of a generic scheme is that we obtain an editor for free for any data type. This makes the approach particularly suitable for *rapid prototyping*.

The standard appearance of a *GEC* is illustrated by the following example that uses many functions and types that will be explained below:

```

module Editor
import StdEnv, StdIO, StdGEC

Start :: *World → *World1
Start world = startIO MDI2 Void3 myEditor world

myEditor = generateEditor ("List", [1])

generateEditor :: (String, t) (PSt ps) → PSt ps4 |5 gGEC{*}6 t
generateEditor (windowName, initialValue) env = newenv
where
    (gecInterface, newenv)
      = gGEC{*} (windowName, initialValue, const id) env

```

<sup>1</sup> This function is equivalent with Haskell `main::IO ()`.

<sup>2</sup> MDI selects Object I/O's Multiple Document Interface.

<sup>3</sup> Void is equivalent with Haskell `()`.

<sup>4</sup> Clean separates the types of function arguments by whitespace, instead of `→`.

<sup>5</sup> In a function type, `|` introduces all overloading class restrictions.

<sup>6</sup> Use the generic instance of kind `*` of `gGEC`.



**Fig. 1.** Generated editor for the standard list type, initially with value [1]

The complete GUI application is shown in Fig. 1.

The generic function `gGEC` creates *GECs*. The way it is defined is explained in [6]. Here, we will focus on its use. It takes a *definition* (`GECDef t env`) of a  $GEC_t$  and *creates* the  $GEC_t$  object in the environment. It returns an *interface* (`GECInterface t env`) to that  $GEC_t$  object. The environment `env` is in this case (`PSt ps`), since `gGEC` uses Object I/O.

**generic**<sup>7</sup> `gGEC t :: GECFunction t (PSt ps)`

```
:: GECFunction t env
  ::= 8 (GECDef t env) env → (GECInterface t env, env)
```

The (`GECDef t env`) consists of three elements. The first is a string that identifies the top-level Object I/O element (window or dialog) in which the editor must be created. The second is the initial value of type `t` of the editor. The third is a callback function of type `t → env → env`. This callback function tells the editor which parts of the program need to be informed of user actions. The editor uses this function to respond to changes to the value of the editor by the application user.

```
:: GECDef t env ::= (String, t, CallbackFunction t env)
:: CallbackFunction t env ::= t env → env
```

The (`GECInterface t env`) is a record that contains all *methods* of the newly created  $GEC_t$ .

```
:: GECInterface t env = { gecGetValue :: GecGet t env
                        , gecSetValue :: GecSet t env
                        }9
:: GecGet      t env ::= env → (t, env)
:: GecSet      t env ::= Update t env → env
:: Update      = YesUpdate | NoUpdate
```


The `gecGetValue` method returns the current value, and `gecSetValue` sets the current value of the associated  $GEC_t$  object. The `gecSetValue` method has an argument of type `Update` indicating whether or not the call-back function has to be called propagating the change of the value through the system.

In Fig. 2 the basic use of the function `gGEC` is illustrated by showing the corresponding GUI for several alternative definitions of `myEditor` (as in the example

<sup>7</sup> **generic**  $f t :: T(t)$  introduces a generic function  $f$  with type scheme  $T(t)$ .

<sup>8</sup> `::=` introduces a type synonym.

<sup>9</sup> A record type with fields  $f_i$  of types  $t_i$  is denoted as  $\{f_i :: t_i\}$ .



**Fig. 2.** Automatically derived editors for standard types

above). This generates an editor for the argument data type. All you have to specify is the name of the window and an initial value. On the right the editor is shown.

For standard types a version of `gGEC` is derived automatically. For user-defined types it is required that a version of `gGEC` is explicitly derived for the given type. For the type `Tree` this is explicitly done in the example. In the rest of these notes these **derives** are not shown.

Programs can consist of several editors. Editors can communicate with each other by tying together the various `gECSetValue` and `gECGetValue` methods. In Sect. 3.2 it is shown how an arrow combinator library [5] can be used for the necessary plumbing. In this section we use the function `selfGEC` (explained in Sect. 3.1) to create ‘self-correcting’ editors:

```
selfGEC ::
  String (t -> t) t (PSt ps) -> (PSt ps) | gGEC{[*]} t & bimap{[*]}10 ps
```

Given function `f` of type `t -> t` on the data model of type `t` and an initial value `v` of type `t`, `selfGEC gui f v` creates the associated `GECt` using `gGEC` (hence the context restriction). `selfGEC` creates a feedback loop that sends every edited output value back as input to the same editor, after applying `f`.

An example of the use of `selfGEC` is given by the following program that creates an editor for a *self-balancing* binary tree:

<sup>10</sup> The generic `gGEC` function requires an instantiation of the predefined generic function `bimap`.

```

myEditor    = selfGEC "Tree" balanceTree (Node Leaf 1 Leaf)

balanceTree = fromListToBalTree o11 fromTreeToList

fromTreeToList Leaf = []
fromTreeToList (Node l x r)
    = fromTreeToList l ++12 [x:fromTreeToList r]13

fromListToBalTree = balance o sort14
where balance []       = Leaf
      balance [x]      = Node Leaf x Leaf
      balance xs       = Node (balance bs) b (balance as)
      where (as, [b:bs]) = splitAt (length xs / 2) xs

```

In this example, we create a  $GEC_{(\text{Tree Int})}$  which displays the indicated initial value `Node Leaf 1 Leaf` (left screen shot in Fig. 3). The user can manipulate this value in any desired order, producing new values of type `Tree Int` (e.g., turning the upper `Leaf` into a `Node` with the pull-down menu, the result of which is shown in the right screen shot in Fig. 3). Each time a new value is created or edited, the feedback function `balanceTree` is applied. `balanceTree` takes an argument of type `Tree a` and returns the tree after balancing it. The shape and lay-out of the tree being displayed adjusts itself automatically. Default values are generated by the editor when needed.




Fig. 3. A self-correcting editor for balanced trees

Note that the only things that need to be specified by the programmer are the initial value of the desired type, and the feedback function.

**Customizing Types.** Clean allows generic functions to be overruled by custom definitions for arbitrary types. `gGEC` is no exception to this rule. The left screenshot in Fig. 4 shows the default interface of the definition below for the ubiquitous *counter* example, and the code that creates it:

Although the definition of the counter is a sensible one, its visual interface clearly is not. In [6] we show how to change the representation of all values

<sup>11</sup> `o` is the function composition operator.


<sup>12</sup> `++` is the list concatenation operator.

<sup>13</sup> In Clean, list denotations are always delimited by `[` and `]`.

<sup>14</sup> `sort :: [a] → [a] | Ord a`.

```
myEditor = selfGEC "Counter" updCtr (0,Neutral)
```

```
updCtr :: Counter → Counter
updCtr (n,Up)   = (n+1,Neutral)
updCtr (n,Down) = (n-1,Neutral)
updCtr any     = any
```



```
:: Counter ::= (Int,UpDown)
:: UpDown  = UpPressed | DownPressed | Neutral
```

**Fig. 4.** The default (left) and customized (right) editor of the counter example

of type `Counter` to the screenshot shown at the right in Fig. 4. Because it has been explained in detail in [6], we will not repeat the code, but point out the important points:

- In this particular example, only the definitions of `(,)` (hide the constructor and place its arguments next to each other) and `UpDown` (display  instead of `(Neutral)`) need to be changed.
- Normally `gGEC` creates the required logical (value passing) and visual infrastructure (GUI components). The programmer, when customizing `gGEC`, only needs to define the visual infrastructure. The programmer must be knowledgeable about Object I/O programming.
- The overruled instance works not only at the top-level. Every nested occurrence of the `Counter` type is now represented as shown right in Fig. 4.

For the creation of GUI applications, we need to model both specific GUI elements (such as buttons) and layout control (such as horizontal, vertical layout). In a way similar to the one shown above for the spin button, this has also been done by specializing `gGEC` for a number of other types that either represent GUI elements or layout. Below the predefined specialized editors are shown for a number of types. The specialized editor for `Display` creates a non-editable GUI; for `Button` a button is created; for `<|>` and `<->` two editors are created below each other, respectively next to each other; and finally `Hide` creates no GUI at all which is useful for remembering state.

For large data structures it may be infeasible to display the complete data structure. Customization can be used to define a  $GEC_t$  that creates a view on a finite subset of such a large data structure with buttons to browse through the rest of the data structure. This same technique can also be used to create  $GECs$  for lazy infinite data structures. For these infinite data structures customization is a must since clearly they can never be fully displayed.

*Exercise 1. A single address GEC.* Write a  $GEC$  for editing a single record containing standard data base data such as name, address and location.

*Exercise 2. A List GEC (advanced).* Write a specialized  $GEC$  that edits a lazy list with buttons to go to editing the next element.

Type of the value given to myEditor:	Corresponding GUI:
<code>:: Display a = Display a</code>	
<code>:: Button = Button String             Pressed</code>	
<code>:: &lt; &gt; a b = a &lt; &gt; b</code>	
<code>:: &lt;-&gt; a b = a &lt;-&gt; b</code>	
<code>:: Hide a = Hide a</code>	

**Fig. 5.** Effect of some predefined customized editors on “Hello World!”

*Exercise 3. An address data base GEC.* Combine the applications above to write a GEC that edits a list of addresses.

### 3 Composition of GECs

In this section we present a number of examples to show how *GECs* can be combined using the callback mechanism and method invocation (Sect. 3.1). In Sect. 3.2 we show how these examples can be expressed using arrow combinators.

#### 3.1 Manual Composition of GECs

**Functionally Dependent GECs.** The first composition example establishes a functional dependency of type  $a \rightarrow b$  between a source editor  $GEC_a$  and destination editor  $GEC_{Display\ b}$ :

```
applyGECs :: (String,String) (a → b) a (PSt ps) → PSt ps
           | gGEC{★} a & gGEC{★} b & bimap{★} ps
applyGECs (sa,sb) f va env
  #15 (gec_b, env) = gGEC{★} (sb, Display (f va), const id) env
  # (gec_a, env) = gGEC{★} (sa, va, set gec_b f) env
  = env
```

<sup>15</sup> The #-notation of Clean has a special scope rule such that the same variable name can be used for subsequent non-recursive #-definitions. For mutually recursive definitions (as in `apply2GECs`) a standard **where**-definition has to be used with a different name for each variable.

```
set :: (GECInterface b (PSt ps)) (a → b) a (PSt ps) → (PSt ps)
set gec f va env = gec.16gecSetValue NoUpdate (Display (f va)) env
```

The callback function of  $GEC_a$  uses the `gecSetValue` interface method of  $GEC_b$  to update the current  $b$  value whenever the user modifies the  $a$  value. As a simple example, one can construct an interactive editor for lists (see Fig. 6) that are mapped to balanced trees by:

```
myEditor
  = applyGECs ("List", "Balanced Tree") fromListToBalTree [1,5,2]
```




Fig. 6. Turning lists into balanced binary trees

Of course, the same can be done for binary functions with slightly more effort:

```
apply2GECs :: (String,String,String) (a → b → c) a b (PSt ps)
  → (PSt ps)
  | gGEC{*} a & gGEC{*} b & gGEC{*} c & bimap{*} ps
apply2GECs (sa,sb,sc) f va vb env = env3
```

where

```
(gec_c,env1) = gGEC{*} (sc,Display (f va vb),const id) env
(gec_b,env2) = gGEC{*} (sb,vb,combine gec_a gec_c (flip f)) env1
(gec_a,env3) = gGEC{*} (sa,va,combine gec_b gec_c f) env2
```

```
combine :: (GECInterface y (PSt ps)) (GECInterface z (PSt ps))
  (x → y → z) x (PSt ps) → PSt ps
```

```
combine gy gc f x env
  # (y,env) = gy.gecGetValue env
  # env     = gc.gecSetValue NoUpdate (Display (f x y)) env
  = env
```

Notice that, due to the explicit environment passing style, it is trivial in Clean to connect  $GEC_b$  with  $GEC_a$  and vice versa. In Haskell's monadic I/O one needs to tie the knot with `fixIO`.

As an example, one can construct two interactive list editors, that are merged and put into a balanced tree (Fig. 7 shows the result):

<sup>16</sup>  $r.f$  denotes the selection of field  $f$  of record  $r$ .




Fig. 7. Merging two lists into a balanced binary tree

```
myEditor
  = apply2GECs ("List1", "List2", "Balanced Tree") makeBalTree [1] [1]
where
  makeBalTree l1 l2 = fromListToBalTree (l1 ++ l2)
```

**Self-correcting GECs.** In this example we give the implementation of the *self-correcting* editor function `gGEC` that was already used in Sect. 2. Self-correcting editors update *themselves* in response to user edit operations. The function definition is concise:

```
selfGEC :: String (a → a) a (PSt ps) → (PSt ps)
         | gGEC{*} a & bimap{*} ps
selfGEC s f v env = env1
where
  (gEC, env1) = gGEC{*} (s, f v, λx → gEC.gECSetValue NoUpdate (f x)) env
```

As an example, one can now construct a *self-sorting* list as follows:

```
myEditor = selfGEC "Self Sorting List" sort [5, 1, 2]
```

It is impossible for a user of this editor to create a stable non-sorted list value.

**Mutually Dependent GECs.** In a similar way one can define mutually dependent *GECs*. Take the following definition of `mutualGEC`.

```
mutualGEC :: (String, String) a (a → b) (b → a) (PSt ps) → (PSt ps)
          | gGEC{*} a & gGEC{*} b & bimap{*} ps
mutualGEC (gui1, gui2) va a2b b2a env = env2
where (gEC_b, env1) = gGEC{*} (gui1, a2b va, set gEC_a b2a) env
      (gEC_a, env2) = gGEC{*} (gui2, va, set gEC_b a2b) env1
```

This function displays two *GECs*. It is given an initial value `va` of type `a`, a function `a2b :: a → b`, and a function `b2a :: b → a`. The `gEC_a` initially displays `va`, while `gEC_b` initially displays `a2b va`. Each time one of the *GECs* is changed, the other is updated automatically. The order in which changes are made is irrelevant. For example, the application `mutualGEC ("Euros", "Pounds")`

```
exchangerate = 1.4
```

```
:: Pounds = {pounds :: Real}
:: Euros = {euros :: Real}
```



```
toPounds :: Euros → Pounds
toPounds {euros} = {pounds = euros / exchangerate}

toEuros :: Pounds → Euros
toEuros {pounds} = {euros = pounds * exchangerate}
```

Fig. 8. Mutually dependent  $GEC_{\text{Pounds}}$  and  $GEC_{\text{Euros}}$

$\{\text{euros} = 3.5\}$  `toPounds toEuros` results in an editor that calculates the exchange between pounds and euros (see Fig. 8) and vice versa.

The example of Fig. 8 may look a bit like a tiny spreadsheet, but it is essentially different since standard spreadsheets do not allow mutual dependencies between cells. Notice also the separation of concerns: the way  $GEC$ s are coupled is defined completely separate from the actual functionality.

### 3.2 Combinators for GEC Composition

The examples in Sect. 3.1 show that  $GEC$ s can be composed by writing appropriate callback functions that use the `GECInterface` methods `gecGetValue` (get the value of a  $GEC$ ) and `gecSetValue` (set its value). This explicit plumbing can become cumbersome when larger and more complex situations must be specified. What is needed, is a disciplined, and more abstract way of combining components. *Monads* [26] and *arrows* [17] are the main candidates for such a discipline. Monads abstract from computations that produce a value, whereas arrows abstract from computations that, *given certain input*, produce values. Because  $GEC$ s also have input and produce values, arrows are the best match. In this section we show how arrows can be used successfully for the composition of  $GEC$ s, resulting in structures that resemble *circuits of GECs*.

It is the task of our arrow model to introduce a standardized way of combining  $GEC$ s. As explained in Sect. 2, one uses a  $GEC_t$  through its interface of type `GECInterface t env`. Method `gecSetValue :: GecSet t env` sets a new value of type `t` in the associated  $GEC_t$ , and `gecGetValue :: GecGet t env` reads its current value of type `t`.

If we generalize these types, then we can regard a *GEC-to-be-combined* as a component that has input `a` and output `b` (where `a = b = t` in case of a ‘pure’  $GEC_t$ ). This generalization of a *GEC-to-be-combined* has type `GecCircuit a b` because of its resemblance with electronic circuits. Consequently, this `GecCircuit a b` has a slightly more general interface, namely a method to *set* values of type `GecSet a env`, and a method to *get* values of type `GecGet b env`. This generalized flow of control of a circuit is visualized in Fig. 9.




Fig. 9. A GEC Circuit (external view)




Fig. 10. A GEC Circuit (internal view)

When circuits are combined this will yield a double connection (one forward *set* and one backward *get* for each circuit). It is essential to realize that usage of the *set* method is restricted to the circuit that produces that input, and, likewise, usage of the *get* method is restricted to the circuit that needs that output.

Moreover, a *GEC-to-be-combined* of type `GecCircuit a b` needs to know where to send its output to, and where to obtain its input from. More precisely, it is only completely defined if it is provided with a corresponding *set* method (of type `GecSet b env`) and a *get* method (of type `GecGet a env`). These methods correspond exactly with the ‘missing’ methods in Fig. 9. Put in other words, a `GecCircuit a b` behaves as a *function*. Indeed, the way we obtain the restricted communication is by passing *continuation functions*. Through these continuations values are passed and set throughout the circuit. Each `GecCircuit a b` is a function that takes two continuations as arguments (one for the input and one for the output) and produces two continuations. The way a circuit takes its continuation arguments, creates a circuit and produces new continuations, can be visualized with the internal view of a circuit (see Fig. 10).

A `GecCircuit` is not only a continuation pair transformation function but it also transforms an Object I/O environment since it has to be able to incorporate the environment functions for the creation of graphical editor components. These environment functions are of type  $(PSt\ ps) \rightarrow (PSt\ ps)$ .

The global idea sketched above motivates the following full definition of the `GecCircuit a b` type:

```

:: GecCircuit a b
  = GecCircuit (∀ ps:
    (GecSet b (PSt ps),GecGet a (PSt ps),PSt ps)
    → (GecSet a (PSt ps),GecGet b (PSt ps),PSt ps))

```

The circuits do not depend on the program state `ps`. This is expressed elegantly using a rank-2 polymorphic function type.

A `GecCircuit a b` generalizes *GECs* by accepting input values of type `a` and produces output values of type `b`. Clearly, for every *GEC<sub>a</sub>* there exists a `GecCircuit a a`. This relation is expressed concisely with the function `edit`:

```

edit :: String → GecCircuit a a | gGEC{[*]} a

```

We will provide an instantiation of the standard arrow class for our *GEC* arrows of type `GecCircuit`. This standard arrow class definition is given below. It describes the basic combinators `>>>` (serial composition), `arr` (function lifting), and `first` (saving values across computations). The other definitions below can all be derived in the standard way from these basic arrow combinators. They are repeated here because we use them in our examples.

```
class Arrow arr where
  arr  :: (a → b) → arr a b
  (>>>) :: (arr a b) → (arr b c) → arr a c
  first :: (arr a b) → arr (a,c) (b,c)

// Combinators for free:
second :: (arr a b) → arr (c, a) (c, b)
second gec = arr swap >>> first gec >>> arr swap
where swap t = (snd t, fst t)

returnA :: arr a a
returnA = arr id

(<<<) infixr 1 :: (arr b c) (arr a b) → arr a c
(<<<) l r = r >>> l

(***) infixr 3 :: (arr a b) (arr c d) → arr (a,c) (b,d)
(***) l r = first l >>> second r

(&&&) infixr 3 :: (arr a b) (arr a c) → arr a (b,c)
(&&&) l r = arr (λx → (x,x)) >>> (l *** r)
```

We use the arrow combinator definitions in the examples below. For each example of Sect. 3.1, we give the definition using arrow combinators, and some of the circuit structures as figures.

However, we first need to show how such a circuit comes to life in Object I/O. This is done with the function `startCircuit` which basically turns a circuit into an Object I/O state transition function. As such it can be used in the `myEditor` function of Sect. 2.

```
startCircuit :: (GecCircuit a b) a (PSt ps) → PSt ps
startCircuit (GecCircuit k) a env
  = let (_,_,env1) = k (setb,geta,env) in env1
where geta    env = (a,env)
      setb _ _ env = env
```

Upon creation, the circuit function is applied to a `geta` function producing the initial argument and a dummy `set` function that just passes the environment.

**Functionally Dependent GECs.** The first arrow example (of which the external view is given in Fig. 11) implements `applyGECs` of Sect. 3.1.




Fig. 11. applyGECs using arrows balancing a tree, external view




Fig. 12. apply2GECs using arrows creating a balanced tree from two lists, external view

```
myEditor
= startCircuit (applyGECs ("List", "Balanced Tree")
                (Display o fromListToBalTree)) [1,5,2]
```

```
applyGECs :: (String,String) (a -> b) -> GecCircuit a b
           | gGEC{*} a & gGEC{*} b
```

```
applyGECs (sa, sb) f = edit sa >>> arr f >>> edit sb
```

Two visual editors are shown. The first allows the user to edit the (initial) list, and the second shows (and allows the user to edit) the resulting balanced tree. In the hand coded examples, the initial value of a *GEC* was specified at the time of its creation. Using the arrow combinators to construct a *GecCircuit*, we specify the initial values for all *GECs* when we start the circuit.

The next example shows the arrow combinator version of `apply2GECs` (see Fig. 12 for its external view):

```
myEditor = startCircuit (apply2GECs ("List1", "List2", "Balanced Tree")
                              makeBalTree) ([1],[2])
```

where

```
makeBalTree (l1,l2) = Display (fromListToBalTree (l1 ++ l2))
```

```
apply2GECs :: (String,String,String) ((a,b) -> c) -> GecCircuit (a,b) c
           | gGEC{*} a & gGEC{*} b & gGEC{*} c
```

```
apply2GECs (sa, sb, sc) f = edit sa *** edit sb >>> arr f >>> edit sc
```

The initial values for the input lists are paired, to allow the delayed initialization using `startCircuit`. The example clearly shows that combining *GECs* using arrow combinators is much more readable than the (often) recursive hand-written functions. The linear flow of information between *GECs*, using the `>>>` combinator, corresponds directly with the code. Although splitting points in flow of information, using the `***` combinator, is less clear, it is still easier on the eyes than the examples of Sect. 3.1.




Fig. 13. selfGEC using arrows, self balancing a tree, external view

**Self-correcting GECs.** The example below shows the arrow combinator version of the selfGEC example (see its external view in Fig. 13).

```
myEditor = startCircuit selfGEC Leaf
```

```
selfGEC :: String (a -> a) -> GecCircuit a a | gGEC{*} a
selfGEC s f = feedback (arr f >>> edit s)
```

The way the feedback combinator constructs a feedback circuit is by taking the value of the circuit and feeding it back again into the circuit. This is done in such a way that it will not be propagated further when it arrives at a GEC editor.

When a feedback circuit contains no editor at all, the meaning of the circuit is undefined since in that case the calculation of the result would depend on itself in a circular way. A feedback circuit in which each path of the circuit contains an editor, is called *well-formed*. It is easy to check syntactically whether feedback circuits are well-formed. Consider the following examples of non well-formed and well-formed feedback circuits.

```
nonWellFormed1 = feedback (arr id >>> arr ((+) 1))
nonWellFormed2 = feedback (arr id &&& edit "Int" >>>
                           arr (\(x, y) -> x + y) )
wellFormed = feedback (edit "Int" >>> arr ((+) 1))
```

It should be clear that the selfGEC function is well-formed. This completes the arrow combinator versions of the examples of Sect. 3.1. The counter example (Sect. 2) is also conveniently, and concisely, expressed below using arr and >>>.

```
myEditor = startCircuit (selfGEC "Counter" updCntr) (0,Neutral)
```

*Exercise 4. An intelligent form.* Write using the combinators a GEC for an intelligent form which calculates some of its values from others (VAT e.g.).

*Exercise 5. An editor for editors (advanced exercise).* Write an editor with which it is possible to create editors. An optional basic design scheme of such a GEC for GECs is shown below.

```
myEditor = startCircuit (designEditor >>>
                        arr convert >>>
                        applicationEditor) initvalue
```

```
designEditor :: GecCircuit DesignEditor DesignEditor
```

```

designEditor = feedback (
    toDesignEditor          >>>
    edit "design"           >>>
    arr (updateDesign o fromDesignEditor))

applicationEditor :: GecCircuit ApplicationEditor ApplicationEditor
applicationEditor = feedback (
    arr (toApplicEditor o updateApplication) >>>
    edit "application"                      >>>
    arr fromApplicEditor                    )

```

It uses two quite complex *GECs* that allow the user to edit the type and visual appearance of another *GEC*: an editor for designing a *GEC*, as well as an editor that displays, and allows the designer to interact with the designed *GEC*. Note that the information flow between these editors is nicely expressed using the arrow combinator `>>>` and that both feedback loops are well-formed.

## 4 Compositional Model-View Programming

When constructing a GUI application, the need arises to incrementally build and modify the GUI under construction. From what is explained in the previous sections, this means that one needs to modify the data structures, and hence also the dependent code, when changing the application. In this section we explain how to obtain a good separation between the *logic* of a GUI application (the *model*) and the way it is presented to the user (the *view*). This way of separating concerns is an instance of the *model-view* paradigm [18]. We show that it can be incorporated smoothly within our toolkit by inventing a new way to realize abstraction and composition based on the specialization mechanism that is used by the generic framework of the GEC toolkit.

The technique is illustrated by means of the following running example of a record with three fields, one of which contains the sum of the other two fields.

The code fragment below shows the *data model*. The data model is a record of type `MyDataModel`. The intention is that whenever the user edits one of the fields `d_value1` or `d_value2`, then these new values are summed and displayed in field `d_sum`. This behavior is defined by `updDataModel`. We want to emphasize that the types and code shown in Fig. 14 are ‘carved in stone’: they do not change in the rest of this section.

As a trivial starting point we take the view model equal to the data model resulting in the following code and corresponding window (see Fig. 15).

The aim of this section is to show how the *view* on this data model can be varied without any modifications to this data model. Some of the kind of variations we would like to make easily are shown in Fig. 16 from left to right: a non-editable field for the result, additionally one of the fields implemented as a

<sup>17</sup> The record update  $\{r \& f_1 = v_1, \dots, f_n = v_n\}$  denotes a new record equal to  $r$ , but with fields  $f_i$  having values  $v_i$ .

```

:: MyDataModel
  = { d_value1 :: Int, d_value2 :: Int, d_sum :: Int }

initDataModel :: (Int,Int) → MyDataModel
initDataModel (v1,v2)
  = { d_value1 = v1, d_value2 = v2, d_sum = v1 + v2 }

updDataModel :: MyDataModel → MyDataModel
updDataModel d = { d &17 d_sum = d.d_value1 + d.d_value2 }

myEditor = selfGEC "View on Data"
           (toMyViewModel o updDataModel o fromMyViewModel)
           (toMyViewModel (initDataModel (0,0)))

```

**Fig. 14.** The code that is carved in stone for the running sum-example of this section


```

:: MyViewModel
  = { v_value1 :: Int, v_value2 :: Int, v_sum :: Int }


toMyViewModel :: MyDataModel → MyViewModel
toMyViewModel d = { v_value1 = d.d_value1
                  , v_value2 = d.d_value2
                  , v_sum     = d.d_sum }

fromMyViewModel :: MyViewModel → MyDataModel
fromMyViewModel v = { d_value1 = v.v_value1
                    , d_value2 = v.v_value2
                    , d_sum     = v.v_sum }

```



**Fig. 15.** The running sum-example, with trivial view



**Fig. 16.** 3 Model-View Variations of the sum-example of Fig. 15

counter to see variations and all editable fields as counter fields with one of the fields itself calculated as a sum of other fields.

In the following sections we show step by step how this can be accomplished. First, in Sect. 4.1 we show how to construct self-contained editors that take care of their own update and conversion behavior. In Sect. 4.2 we turn these self-contained editors into *abstract reusable* editors, thus encapsulating all information about their implementation and behaviour. We show how this is possible

although *abstract* types seem to be at odds with generic programming. Finally, we show in Sect. 4.3 that these self-contained editors are truly reusable elements themselves. In that section we will also show (in Fig. 17) how to produce the variations (given in Fig. 16) of the running example (shown in Fig. 15).

#### 4.1 Defining Self-contained Editors

If we want to reuse an existing editor, it is not enough to reuse its type. We also want to reuse its functionality: each editor should take care of its own update. For this purpose we need a type in which we can store the functionality of an editor. If want to create a view  $v$  on a domain model  $d$ , we need to be able to replace a standard editor for type  $d$  by a self-contained editor for some isomorphic view type  $v$ . Furthermore, since we generally also have to perform conversions between these types, we like to store them as well, such that each editor can take care of its own conversions. Finally, it is generally useful to take into account the old value of  $v$  when converting from  $d$  since editors may have an internal state. Therefore we define a new type, `ViewGEC d v` (and a corresponding creation function `mkViewGEC`), in which we can store the update and conversion functions:

```
:: ViewGEC d v = { d_val      :: d
                  , d_to_v    :: d → (Maybe v) → v
                  , update_v  :: v → v
                  , v_to_d    :: v → d }
```

```
mkViewGEC :: d (d → (Maybe b) → v) (v → v) (v → d) → ViewGEC d v
mkViewGEC d fdvv fvv fvd
    = { d_val      = d
      , d_to_v    = fdvv
      , update_v  = fvv
      , v_to_d    = fvd }
```

For convenience, we define the creation function `mkSimpleViewGEC` that has a slightly simpler interface. In Exercise 6 you can practice with the more general interface for managing state.

```
mkSimpleViewGEC :: d (d → v) (v → v) (v → d) → ViewGEC d v
mkSimpleViewGEC d fdv fvv fvd = mkViewGEC d fdvv fvv fvd
where fdvv d Nothing = fdv d
      fdvv _ (Just v) = v
```

Next, we define a specialized version of our generic editor `gGEC` for this type. The top-level definition is given below. Notice that in `gGEC{ViewGEC}` two additional parameters appear: `gGECd` and `gGECv`. This is caused by the fact that generic functions in `Clean` are kind-indexed functions. As `ViewGEC d v` is of kind  $\star \rightarrow \star \rightarrow \star$ , the generic function has two additional parameters, one for type  $d$  and one for type  $v$ .

```

gGEC{ViewGEC} gGECd gGECv (vName, viewGEC, viewGECallback) env
  = ({ gecSetValue = viewSetValue vInterface
      , gecGetValue = viewGetValue vInterface }, new_env)
where (vInterface,new_env)
      = gGECv ( vName
                , viewGEC.d_to_v viewGEC.d_val Nothing
                , viewCallback vInterface
                ) env

```

The `ViewGEC` editor does the following. The value of type `d` is stored in the `ViewGEC` record, but a `d`-editor (`gGECd`) for it is not created. Taking the old value of `v` into account, the `d`-value is converted to a `v`-value using the conversion function `d_to_v :: d → (Maybe v) → v`. For this `v`-value we do generate a generic `v`-editor (`gGECv`) to store and edit the `v`-value.

What remains to be defined are the callback function for the view editor (`viewCallback`) and the `GECInterface` (`ViewGEC d v`) methods (`viewSetValue` and `viewGetValue`). We discuss the most complicated one, the callback function, first. Whenever the application user creates a new `v`-value with this editor, the call-back function of the `v`-editor is called (`viewCallback`) and the `update_v :: v → v` function is applied. This is similar to applying (`selfGEC update_v`) to the corresponding new value of type `v`. The resulting new `v`-value is shown in the `v`-editor again, and it is converted back to a `d`-value as well, using the function `v_to_d :: v → d`. This new `d`-value is then stored in the `ViewGEC` record in the `d_val` field, and the call-back function for the `ViewGEC` editor is called (`viewGECCallback`).

```

viewCallback vInterface new_v env
  = viewGECCallback {viewGEC & d_val = new_d} new_env
where new_upd_v = viewGEC.update_v new_v
      new_env   = vInterface.gecSetValue new_upd_v env
      new_d     = viewGEC.v_to_d new_upd_v

```

The two interface methods to write (`viewSetValue`) and read (`viewGetValue`) are fairly straightforward. Writing a value of type `ViewGEC d v` amounts to writing a value of type `v` using the current old value of type `v` and the new value of type `d` that is stored in the `ViewGEC` record. Reading a value of type `ViewGEC d v` amounts to reading the current `v` value and wrap it up in the record after converting it to a `d` value.

```

viewSetValue vInterface new_viewGEC env
  = vInterface.gecSetValue new_v new_env
where new_v = new_viewGEC.d_to_v new_viewGEC.d_val (Just old_v)
      (old_v,new_env) = vInterface.gecGetValue env

```

```

viewGetValue vInterface env
  = ({viewGEC & d_val = viewGEC.v_to_d current_v},new_env)
where (current_v,new_env) = vInterface.gecGetValue env

```

The concrete behavior of the generated `ViewGEC` editor now not only depends on the type, but also on the concrete information stored in a value of type `ViewGEC`. A self-contained reusable editor, such as `counterGEC` below, is now quickly constructed. The corresponding editor takes care of the conversions and the update. The `displayGEC` does a trivial update (`identity`) and also takes care of the required conversions.

```
counterGEC :: Int → ViewGEC Int Counter
counterGEC i = mkViewGEC i toCounter updCntr fromCounter
```

```
displayGEC :: a → ViewGEC a (Display a)
displayGEC x = mkViewGEC x toDisplay id fromDisplay
```

Making use of these new self-contained editors we can attach a *view model* to the *data model* that was presented in the start of this section by giving appropriate definitions of the conversion functions `toMyViewModel` and `fromMyViewModel`. All other definitions remain the same.

```
:: MyViewModel = { v_value1 :: ViewGEC Int Counter
                  , v_value2 :: ViewGEC Int Counter
                  , v_sum     :: ViewGEC Int (Display Int) }
```

```
toMyViewModel :: MyDataModel → MyViewModel
toMyViewModel d = { v_value1 = counterGEC d.d_value1
                  , v_value2 = counterGEC d.d_value2
                  , v_sum     = displayGEC d.d_sum }
```

```
fromMyViewModel :: MyViewModel → MyDataModel
fromMyViewModel v = { d_value1 = v.v_value1.d_val
                    , d_value2 = v.v_value2.d_val
                    , d_sum     = v.v_sum.d_val }
```

In the definition of `toMyViewModel` we can now choose any suitable self-contained editor. Each editor handles the needed conversions and updates itself automatically. To obtain the value we are interested in, we just have to address the `d_val` field.

The example shows that we have indeed obtained proper compositional behavior but the programming method is not truly compositional. If we would replace a self-contained editor by another in `toMyViewModel`, all other code remains the same. However, we *do* have to change the type of `MyViewModel`. In this type it is completely visible what kind of editor has been used. In the following section it is shown how to create a complete compositional abstraction using a special abstract data type for our self-contained editors.

## 4.2 Abstract Self-contained Editors

The concrete value of type `ViewGEC d v` is used by the generic mechanism to generate the desired self-contained editors. The `ViewGEC d v` type depends on

the type of the editor  $v$  that is being used. Put in other words, the type still reveals information about the implementation of editor  $v$ . This is undesirable for two reasons: one can not exchange views without changing types, and the type of composite views reflects their composite structure. For these reasons, a type is required that *abstracts* from the concrete editor type  $v$ .

However, if we manage to hide these types, how can the generic mechanism generate the editor for it? The generic mechanism can only generate an editor for a given concrete type, not for an abstract type of which the content is unknown. The solution is as follows. When the abstraction is being made, we *do* know the contents and its type. Hence, we can store the *generic editor function* (of type `GECFunction`, see Sect. 2) in the abstract data structure itself where the abstraction is being made. The stored editor function can be applied later when we really need to construct the editor. Therefore, it is possible to define an abstract data structure (`AGEC d`) in which the `ViewGEC d v` is stored with its corresponding generic `gGEC` function for  $v$ . Technically this requires a type system that supports existentially quantified types as well as rank-2 polymorphism.

```
:: AGECEC d
  =  $\exists v$ : AGECEC (ViewGEC d v) ( $\forall ps$ : GECFunction (ViewGEC d v) (PSt ps))
```

```
mkAGECEC :: (ViewGEC d v)  $\rightarrow$  AGECEC d | gGEC{ $\star$ } v
mkAGECEC viewGEC = AGECEC viewGEC (gGEC{ $\star \rightarrow \star \rightarrow \star$ } undef gGEC{ $\star$ })
```

`gGEC{AGECEC}` = ... // similar to `gGEC{ViewGEC}`, but apply function in `AGECEC`

The function `mkAGECEC` creates the desired `AGECEC` given a `viewGEC`. Looking at the type of `AGECEC`, the generic system can deduce that the editor to store has to be a generic editor for type `ViewGEC d v`. To generate this editor, the generic system by default requires an editor for type `d` and type `v` as well. We know that in this particular case we do not use the `d`-editor at all. We can tell this to the generic system by making use of the fact that generic functions in `Clean` are kind indexed. The system allows us, if we wish, to explicitly specify the editors for type `d` (`undef`) and type `v` (`gGEC{ $\star$ }`) to be used by the editor for `ViewGEC (gGEC{ $\star \rightarrow \star \rightarrow \star$ })`. In this case we know that we do not need an editor for type `d` (hence the `undef`), and use the standard generic editor for type `v`. The overloading context restriction in the type of `mkAGECEC ( | gGEC{ $\star$ } v)` states that for making an `AGECEC d` out of a `ViewGEC d v` only an editor for type `v` is required.

We also have to define a specialized version of `gGEC` for the `AGECEC` type. The corresponding generated editor applies the stored editor to the stored `ViewGEC`.

The types and kind indexed generic programming features we have used here may look complicated, but for the programmer an abstract editor is easy to make. To use a self-contained editor of type `v` as editor for type `d`, a `ViewGEC d v` has to be defined. Note that the editor for type `v` is automatically derived for the programmer by the generic system! The function `mkAGECEC` stores them both into an `AGECEC`. The functions `counterAGECEC` and `displayAGECEC` show how easy `AGECEC`'s can be made. One might be surprised that the overloading context for `displayAGECEC`

still requires a *d*-editor ( `| gGEC{*} d`). This is caused by the fact that in this particular case type *d* is used in the definition of type `Display`.

```
counterAGEC :: Int → AGECE Int
counterAGEC i = mkAGECE (counterGEC i)
```

```
displayAGECE :: d → AGECE d | gGEC{*} d
displayAGECE x = mkAGECE (displayGEC x)
```

We have chosen to export `AGECE d` as a Clean abstract data type. This implies that code that uses such an abstract value can not apply record selection to access the *d* value. For this purpose we provide the following obvious projection functions to retrieve the *d*-value from an `AGECE d` (`^^`) and to store a new *d*-value in an existing `AGECE d` (the infix operator `^=`).

```
(^^) :: (AGECE d) → d                // Read current value
(^^) (AGECE viewGEC gGEC) = viewGEC.d_val
```

```
(^=) infix1 :: (AGECE d) d → (AGECE d) // Set new value
(^=) (AGECE viewGEC gGEC) nval = AGECE {viewGEC & d_val=nval} gGEC
```

Using abstract editors we can refine the *view model* data type and conversion functions:

```
:: MyViewModel = { v_value1 :: AGECE Int
                  , v_value2 :: AGECE Int
                  , v_sum    :: AGECE Int }
```

```
toMyViewModel :: MyDataModel → MyViewModel
toMyViewModel d = { v_value1 = counterAGECE d.d_value1
                  , v_value2 = counterAGECE d.d_value2
                  , v_sum    = displayAGECE d.d_sum }
```

```
fromMyViewModel :: MyViewModel → MyDataModel
fromMyViewModel v = { d_value1 = ^^ v.v_value1
                   , d_value2 = ^^ v.v_value2
                   , d_sum    = ^^ v.v_sum }
```

The advantage of the use of the `AGECE`-type is that, if we want to pick another editor, we only have to tell which one to pick in the definition of `toMyViewModel`. The types used in `MyViewModel` all remain the same (`AGECE Int`), no matter which editor is chosen. Also the definition of `fromMyViewModel` remains unaffected.

### 4.3 Abstract Editors Are Compositional

In order to show the compositional nature of abstract editors, we first turn the running example into an abstract editor: `sumAGECE :: AGECE Int`. It can be used *itself* as an `Int`-editor. We follow the scheme introduced above:

```

sumAGEC :: Int → AGECE Int                                     // see counterAGEC (4.2)
sumAGEC i = mkAGECE (sumGEC i)
where sumGEC :: Int → ViewGEC Int MyViewModel // see counterGEC (4.1)
      sumGEC i = mkSimpleViewGEC i toV updV fromV
      where toV  = toMyViewModel o toMyData
            fromV = fromMyData    o fromMyViewModel
            updV  = toMyViewModel o updDataModel o fromMyViewData

            toMyData i = {d_value1 = 0, d_value2 = i, d_sum = i}
            fromMyData d = d.sum

```

Now `sumAGEC`, `counterAGEC`, and `displayAGEC` are interchangeable components. If we want to experiment with variants of the running example, we pick the instance of our choice in the `toMyViewModel` function (in this way achieving the wanted examples of the beginning of this section: see Fig. 17).

**Alternative definition of `toMyViewModel`:**

```


toMyViewModel1 d
  = { v_value1 = idAGEC      d.d_value1
      , v_value2 = idAGEC      d.d_value2
      , v_sum    = displayAGEC d.d_sum }

toMyViewModel2 d
  = { v_value1 = idAGEC      d.d_value1
      , v_value2 = counterAGEC d.d_value2
      , v_sum    = displayAGEC d.d_sum }

toMyViewModel3 d
  = { v_value1 = counterAGEC d.d_value1
      , v_value2 = sumAGECE  d.d_value2
      , v_sum    = displayAGEC d.d_sum }

```

**Corresponding GUI:**



**Fig. 17.** *Plug-and-play* your favorite abstract editors. The only code that changes is the function `toMyViewModel`. The values have been edited by the user.

We have set up a library of abstract components. One of these library functions `idAGEC` (which takes a value of any type and promotes it to an abstract editor component for that type) is used in the example above. With this library it is possible to rapidly create GUIs in a declarative style. This is useful e.g. for prototyping, education, tracing and debugging purposes.

Below, we summarize only those functions of the collection that are used in the examples in these notes:

```

vertlistAGEC :: [a] → AGECE [a] | gGEC{[*]} a
// all elements displayed in a column
counterAGEC  :: a → AGECE a | gGEC{[*]}, IncDec a
// a special number editor
hidAGEC      :: a → AGECE a // identity, no editor
displayAGEC  :: a → AGECE a | gGEC{[*]} a // identity, non-editable editor

```

**Exercise 6. Abstract Lists (advanced).** Modify the solution of Exercise 2 in such a way that it is an abstract editor for lists. Make use of the general `mkViewGEC` function that manipulates the state of the view.

**Exercise 7. Abstract Intelligent Form.** Experiment with your solution for Exercise 4. Replace parts of it by *AGECs* and experiment with different editors at the basic level.

## 5 Higher-Order GECs

In the previous sections all functionality had to be encoded in the program. A user could only edit data values, no functions.

In this section we explain how basic *GECs* have been extended with the ability to deal with functions and expressions allowing the user to edit functional values.

Consider as a motivating example the running example of the previous section. It had the function `sum` encoded in the program. In order to change it to product e.g. the user would have to ask a programmer to change the program. Of course, a user would prefer to have the ability to change (and add) functionality.

A user would like to be able to type in expressions using e.g. `twice` and `map` and let the *GEC* react accordingly. In Fig. 18 we show how such an editor could look like. On the left there is a *GEC* in which a user has typed in a partial application of the `map` function, on the right the `twice` function has been used.

Looking carefully one can imagine that it is not an easy task to achieve such functionality. Depending on what functional value is typed in, the number of other fields will vary.

Suppose the expression field is the `twice` function  $\lambda f\ x \rightarrow f\ (f\ x)$  as in the right example of Fig. 18. If the first argument is the increment function `((+) 1)`,




Fig. 18. Two *GECs* editing functional values

there is one further argument, an `Int`, and of course a field for the result. However, if the first argument would also be the twice function then an extra argument would be required!

The number of argument fields depend on the type of the functional value that is typed in by the user (or of the actual arguments that are given to it by the user). These examples of Fig. 18 are created in the sections below.

Because functions are opaque, the solution requires a means of interpreting functional expressions as functional values. Instead of writing our own parser/interpreter/type inference system we use the *Esther* shell [24]. Esther enables the user to enter expressions (using a subset of `Clean`) that are dynamically typed, and transformed to values and functions using compiled code. It is also possible to reuse earlier created functions, which are stored on disk. Its implementation relies on the *dynamic type system* [1,21,25] of `Clean`.

The shell uses a text-based interface, and hence it makes sense to create a special *string*-editor (Sect. 5.1), which converts any string into the corresponding dynamically typed value. This special editor has the same power as the Esther command interpreter and can deliver any dynamic value, including higher-order polymorphic functions. In addition we show that the actual *content* of a dynamic value can be influenced by the very same generic mechanism, using type dependent functions (Sect. 5.2). With this mechanism, dynamics can be used in a type-directed way, but only for monomorphic types in dynamics.

## 5.1 Dynamically Typed Higher-Order GECs

We first introduce the foundations of the Esther shell, and proceed by showing how to construct an editor for functions.

**Dynamics in Clean.** A *dynamic* is a value of static type `Dynamic`, which contains an expression as well as a representation of its static type, e.g., `dynamic 42 :: Int`, `dynamic map fst :: ∀a b: [(a, b)] → [a]`. Basically, dynamic types turn every (first and higher-order) type into a first-order type, while providing run-time access to the original type and value.

Function alternatives and case patterns can match on values of type `Dynamic`. Such a pattern match consists of a value pattern and a type pattern, e.g., `[4, 2] :: [Int]`. The compiler translates a pattern match on a type into run-time type unification. If the unification is successful, type variables in a type pattern are bound to the offered type. Applying dynamics at run-time will be used to create an editor that changes according to the type of entered expressions (Sect. 5.1, Example 2).

```
dynamicApply :: Dynamic Dynamic → Dynamic
dynamicApply (f :: a → b) (x :: a) = dynamic f x      :: b
dynamicApply      df          dx    = dynamic "Error" :: String
```

`dynamicApply` tests if the argument type of the function `f`, inside its first argument, can be unified with the type of the value `x`, inside the second argument. `dynamicApply` can safely apply `f` to `x`, if the type pattern match succeeds. It yields

a value of the type that is bound to the type variable  $\mathbf{b}$  by unification, wrapped in a dynamic. If the match fails, it yields a string in a dynamic.

Type variables in type patterns can also relate to type variables in the static type of a function. A  $\hat{\phantom{x}}$  behind a variable in a pattern associates it with the same type variable in the static type of the function.

```
matchDynamic :: Dynamic →  $\mathbf{t}$  | TC  $\mathbf{t}$ 
matchDynamic (x ::  $\mathbf{t}^{\hat{\phantom{x}}}$ ) = x
```

The static type variable  $\mathbf{t}$ , in the example above, is determined by the static context in which it is used, and imposes a restriction on the actual type that is accepted at run-time by `matchDynamic`. The function becomes overloaded in the predefined TC (type code) class. This makes it a type dependent function [21].

The dynamic run-time system of Clean supports writing dynamics to disk and reading them back again, possibly in another program or during another execution of the same program. This provides a means of type safe communication, the ability to use compiled plug-ins in a type safe way, and a rudimentary basis for mobile code. The dynamic is read in lazily after a successful run-time unification. The amount of data and code that the dynamic linker links is, therefore, determined by the evaluation of the value inside the dynamic.

```
writeDynamic :: String Dynamic env → (Bool,env) | FileSystem env
readDynamic  :: String env → (Bool,Dynamic,env) | FileSystem env
```

Programs, stored as dynamics, have Clean types and can be regarded as a typed file system. We have shown that `dynamicApply` can be used to type check any function application at run-time using the static types stored in dynamics. Combining both in an interactive ‘read expression – apply dynamics – evaluate and show result’ loop, already gives a simple shell that supports the type checked run-time application of programs to documents. The `composeDynamic` function below, taken from the Esther shell, applies dynamics and infers the type of an expression.

```
composeDynamic  :: String env → (Dynamic,env) | FileSystem env
showValueDynamic :: Dynamic → String
```

`composeDynamic expr env` parses *expr*. Unbound identifiers in *expr* are resolved by reading them from the file system. In addition, overloading is resolved. Using the parse tree of *expr* and the resolved identifiers, the `dynamicApply` function is used to construct the (functional) value *v* and its type  $\tau$ . These are packed in a **dynamic**  $v :: \tau$  and returned by `composeDynamic`. In other words, if  $env \vdash expr :: \tau$  and  $\llbracket expr \rrbracket_{env} = v$  then `composeDynamic expr env = (v ::  $\tau$ , env)`. The `showValueDynamic` function yields a string representation of the value inside a dynamic.

**Creating a GEC for the Type Dynamic.** With the `composeDynamic` function, an editor for dynamics can be constructed. This function needs an appropriate environment to access the dynamic values and functions (plug-ins) that are stored on disk. The standard (PSt ps) environment used by the generic

`gGEC` function (Sect. 2) is such an environment. This means that we can use `composeDynamic` in a specialized editor to offer the same functionality as the command line interpreter. Instead of Esther’s console we use a `String` editor as interface to the application user. In addition we need to convert the provided string into the corresponding dynamic. We therefore define a composite data type `DynString` and a specialized `gGEC`-editor for this type (a  $GEC_{\text{DynString}}$ ) that performs the required conversions.

```
:: DynString = DynStr Dynamic String
```

The choice of the composite data type is motivated mainly by simplicity and convenience: the string can be used by the application user for typing in the expression. It also stores the original user input, which cannot be extracted from the dynamic when it contains a function.

Now we specialize `gGEC` for this type `DynString`. The complete definition of `gGEC{DynString}` is given below.

```
gGEC{DynString} (gui, DynStr _ expr, dynStringUpd) env
  # (stringGEC, env) = gGEC{*} (gui, expr, stringUpd dynStringUpd) env
  = ( { gecSetValue = dynSetValue stringGEC.gecSetValue
      , gecGetValue = dynGetValue stringGEC.gecGetValue } , env)
where dynSetValue stringSetValue (DynStr _ expr) env
  = stringSetValue expr env
  dynGetValue stringGetValue env
  # (nexpr, env) = stringGetValue env
  # (ndyn, env) = composeDynamic nexpr env
  = (DynStr ndyn nexpr, env)
  stringUpd dynStringUpd nexpr env
  # (ndyn, env) = composeDynamic nexpr env
  = dynStringUpd (DynStr ndyn nexpr) env
```

The created  $GEC_{\text{DynString}}$  displays a box for entering a string by calling the standard generic `gGEC{*}` function for the value `expr` of type `String`, yielding a `stringGEC`. The `DynString`-editor is completely defined in terms of this `String`-editor. It only has to take care of the conversions between a `String` and a `DynString`. This means that its `gecSetValue` method `dynSetValue` sets the string component of a new `DynString` in the underlying `String`-editor. Its `gecGetValue` method `dynGetValue` retrieves the string from the `String`-editor, converts it to the corresponding `Dynamic` by applying `composeDynamic`, and combines these two values in a `DynString`-value. When a new string is created by the application user, this will call the callback function `stringUpd`. It invokes the callback function `dynStringUpd` (provided as an argument upon creation of the `DynString`-editor), after converting the `String` to a `DynString`.

It is convenient to define a constructor function `mkDynStr` that converts any input `expr`, that has value `v` of type  `$\tau$` , into a value of type `DynString` guaranteeing that if  $v :: \tau$  and  $\llbracket expr \rrbracket = v$ , then  $(\text{DynStr } (v::\tau) \text{ expr}) :: \text{DynString}$ .

```
mkDynStr :: a → DynString | TC a
mkDynStr x = let dx = dynamic x in DynStr dx (showValueDynamic dx)
```



**Function Test Example.** We construct an interactive editor that can be used to test functions. It can be a newly defined function, say  $\lambda x \rightarrow x^2$ , or any existing function stored on disk as a `Dynamic`. Hence the tested function can vary from a small function, say `factorial`, to a large complete application.

```

:: MyRecord = { function :: DynString
               , argument :: DynString
               , result   :: DynString }

myEditor
= selfGEC "test" guiApply (initval id 0)
where
  initval f v
    = { function = mkDynStr f
        , argument = mkDynStr v
        , result   = mkDynStr (f v) }
  guiApply
    r=:18{ function = DynStr (f::a → b) _
           , argument = DynStr (v::a)   _ }
    = {r & result = mkDynStr (f v)}
  guiApply r = r

```



`MyRecord` has three `DynString` fields: `function`, `argument`, and `result`. The user can use this editor to enter a function definition and its argument. The `selfGEC` function ensures that each time a new string is created with the editor "test", the function `guiApply` is applied that provides a new value of type `MyRecord` to the editor. The function `guiApply` tests, in a similar way as the function `dynamicApply` (see Sect. 5.1), whether the type of the supplied function and argument match. If so, a new result is calculated. If not, nothing happens.

This editor can only be used to test functions with one argument. What happens if we edit the function and the argument in such a way that the result is not a plain value but a function itself? Take, e.g., as function the twice function  $\lambda f x \rightarrow f (f x)$ , and as argument the increment function  $((+) 1)$ . Then the result is also a function  $\lambda x \rightarrow ((+) 1) ((+) 1 x)$ . The editor displays `<function>` as result. There is no way to pass an argument to the resulting function.

With an editor like the one above, the user can enter expressions that are automatically converted into the corresponding `Dynamic` value. As in the shell, unbound names are expected to be dynamics on disk. Illegal expressions result in a `Dynamic` containing an error message.

To have a properly higher-order dynamic application example one needs an editor in which the user can type in functions of arbitrary arity, and subsequently enter arguments for this function. The result is then treated such that, if it is a function, editors are added dynamically for the appropriate number of arguments. This is explained in the following example.

<sup>18</sup>  $x =: e$  binds  $x$  to  $e$ .

**Expression Test Example.** We construct a test program that accepts arbitrary expressions and adds the proper number of argument editors, which again can be arbitrary expressions. The number of arguments cannot be statically determined and has to be recalculated each time a new value is provided. Instead of an editor for a record we therefore create an editor for a list of tuples. Each tuple consists of a string used to prompt to the user, and a `DynString`-value. The tuple elements are displayed below each other using the predefined list editor `vertlistAGEC` (Sect. 4.3) and access operator `^^` (Sect. 4.2). The `selfGEC` function is used to ensure that each change made with the editor is tested with the `guiApply` function and the result is shown in the editor.

```
myEditor
= selfGEC "test" (guiApply o (^^))
  (vertlistAGEC [show "expression " 0])
```



where

```
guiApply [df:=(_,DynStr f _):args]
  = vertlistAGEC [df:check f args]
```

where

```
check (f::a → b)
  [arg:=(_,DynStr (x::a) _):args]
  = [arg : check (dynamic f x) args]
check (f::a → b) _
  = [show "argument " "??"]
check (x::a) _
  = [show "result " x]
```

```
show s v = (Display s,mkDynStr v)
```



The key part of this example is formed by the function `check` which calls itself recursively on the result of the dynamic application. As long as function and argument match, and the resulting type is still a function, it requires another argument which is checked for type consistency. If the resulting type is a plain value, it is evaluated and shown using the predefined function `display`, which creates a non-editable editor that just displays its value. As soon as a type mismatch is detected, a question mark is displayed to prompt the user to try again. With this editor, any higher-order polymorphic function can be entered and tested.

## 5.2 Statically Typed Higher-Order *GECs*

The editors presented in the previous section are flexible because they deliver a `Dynamic` (packed into the type `DynString`). They have the disadvantage that the programmer has to program a check, such as the `check` function in the previous example, on the type consistency of the resulting `Dynamics`.

In many applications it is statically known what the type of a supplied function must be. In this section we show how the run-time type check can be replaced by a compile-time check, using the abstraction mechanism for *GECs*. This gives

the programmer a second solution for higher-order types that is statically typed, which allows, therefore, type-directed generic GUI creation.

**Adding Static Type Constraints to Dynamic GECs.** The abstraction mechanism provided by *AGECs* is used to build type-directed editors for higher-order types, which check the type of the entered expressions dynamically. These statically typed higher-order editors are created using the function `dynamicAGEC`. The full definition of this function is specified and explained below.

```
dynamicAGEC :: d → AGEC d | TC d
dynamicAGEC x
  = mkAGEC (mkSimpleViewGEC x toView (updView x) (fromView x))
where
  toView :: d → (DynString, AGEC DynString)
  toView newX = let dx = mkDynStr newX in (dx, hidAGEC dx)

  fromView :: d (DynString, AGEC DynString) → d | TC d
  fromView _ (_, oldx) = case ^^oldx of DynStr (x::d^ ) _ → x

  updView :: d (DynString, AGEC DynString)
           → (DynString, AGEC DynString) | TC d
  updView _ (newx=(DynStr (x::d^ ) _), _) = (newx, hidAGEC newx)
  updView _ (_, oldx) = (^^oldx, oldx)
```

The abstract Dynamic editor, which is the result of the function `dynamicAGEC` initially takes a value of some statically determined type `d`. It converts this value into a value of type `DynString`, such that it can be edited by the application user as explained in Sect. 5.1. The application user can enter an expression of arbitrary type, but now it is ensured that only expressions of type `d` are approved.


The function `updView`, which is called in the abstract editor after any edit action, checks, using a type pattern match, whether the newly created dynamic can be unified with the type `d` of the initial value (using the `^`-notation in the pattern match as explained in Sect. 5.1). If the type of the entered expression is different, it is rejected and the previous value is restored and shown. To do this, the abstract editor has to remember in its internal state also the previously accepted correctly typed value. Clearly we do not want to show this part of the internal state to the application user. This is achieved using the abstract editor `hidAGEC` (Sect. 4.3), which creates an invisible editor, i.e., a store, for any type.

**Function Test Example, Revisited.** Consider the following variation of the function test example on page 237:

```
:: MyRecord a b = { function :: AGEC (a → b)
                  , argument :: AGEC a
                  , result   :: AGEC b }
```

```
myEditor = selfGEC "test" guiApply (initval ((+) 1.0) 0.0)
where
```

```
  initval f v
    = { function = dynamicAGEC f
        , argument = dynamicAGEC v
        , result   = displayAGEC (f v) }
```




```
  guiApply myrec={ function = af, argument
                  = {myrec & result = displayAGEC ((~^af) (~^av))}
```

The editor above can be used to test functions of a certain statically determined type. Due to the particular choice of the initial values  $((+) 1.0 :: \text{Real} \rightarrow \text{Real}$  and  $0.0 :: \text{Real}$ ), the editor can only be used to test functions of type  $\text{Real} \rightarrow \text{Real}$  applied to arguments of type  $\text{Real}$ . Notice that it is now statically guaranteed that the provided dynamics are correctly typed. At runtime the `dynamicAGEC`-editors take care of the required checks and they reject ill-typed expressions. The programmer therefore does not have to perform any checks anymore. The abstract `dynamicAGEC`-editor delivers a value of the proper type just like any other abstract editor.

The code in the above example is not only simple and elegant, but it is also very flexible. The `dynamicAGEC` abstract editor can be replaced by any other abstract editor, provided that the statically derived type constraints (concerning `f` and `v`) are met. This is illustrated by the next example.

**Function Test Example, Once Again Revisited.** If one prefers a counter as input editor for the argument value, one only has to replace `dynamicAGEC` by `counterAGEC` in the definition of `initval`:

```
initval f v
  = { function = dynamicAGEC f
      , argument = counterAGEC v
      , result   = displayAGEC (f v) }
```



The `dynamicAGEC` is typically used when *expression* editors are preferred over *value* editors of a type, and when application users need to be able to enter functions of a statically fixed monomorphic type.

One can create an editor for any higher-order type  $\mathbf{t}$ , even if it contains polymorphic functions. It is required that all higher-order parts of  $\mathbf{t}$  are abstracted, by wrapping them with an *AGEC* type. Basically, this means that each part of  $\mathbf{t}$  of the form  $\mathbf{a} \rightarrow \mathbf{b}$  must be changed into *AGEC*  $(\mathbf{a} \rightarrow \mathbf{b})$ . For the resulting type  $\mathbf{t}'$  an edit dialog *can* be automatically created, e.g., by applying `selfGEC`. However, the initial value that is passed to `selfGEC` must be monomorphic, as usual for any instantiation of a generic function. Therefore, editors for polymorphic types cannot be created automatically using this statically typed generic technique. As explained in Sect. 5.1 polymorphic types can be handled with dynamic type checking.

Summarizing, we have shown two ways to create editors that can deal with higher order types. Firstly, one can create dynamically typed higher-order editors, which have the advantages that we can deal with polymorphic higher order types and overloading. This has the disadvantage that the programmer has to check type safety in the editor. Secondly, we have treated a method in which the compiler can ensure type correctness of higher-order types in statically typed editors, but then the resulting editors can only edit monomorphic types.

*Exercise 8. Dynamically Adaptable Intelligent Form.* Change your solutions from exercises 4 and 7 such that the intelligence can be dynamically changed by typing in the function that is applied.

## 6 Related Work

We distinguish three areas of related work:

**Grammars Instead of Types:** Taking a different perspective on the type-directed nature of our approach, one can argue that it is also possible to obtain editors by starting from a grammar specification instead of a type. Such toolkits require a grammar as input and yield an editor GUI as result. Projects in this flavor are for instance the recent *Proxima* project [23], which relies on XML and its DTD (Document Type Definition language), and the *Asf+Sdf Meta-Environment* [11] which uses an *Asf* syntax specification and *Sdf* semantics specification. The major difference with such an approach is that these systems need both a grammar and some kind of interpreter. In our system higher-order elements are immediately available as a functional value that can be applied and passed to other components.

**GUI Programming Toolkits:** From the abstract nature of the *GEC* toolkit it is clear that we need to look at GUI toolkits that also offer a high level of abstraction. Most GUI toolkits are concerned with the low level management of widgets in an imperative style. One well-known example of an abstract, compositional GUI toolkit based on a combinator library is *Fudgets* [12]. These combinators are required for plumbing when building complex GUI structures from simpler ones. In our system far less plumbing is needed. Most work is done automatically by the generic function *gGEC*. The only plumbing needed in our system is for combining the *GEC*-editors themselves. Any application will only have a very limited number of *GEC*-editors. Furthermore, the *Fudget* system does not provide support for editing function values or expressions.

A  $GEC_{\tau}$  is a  $\tau$ -stateful object, hence it makes sense to look at object oriented approaches. The power of abstraction and composition in our functional framework is similar to *mixins* [13] in object oriented languages. One can imagine an OO GUI library based on compositional and abstract mixins in order to obtain a similar toolkit. Still, such a system lacks higher-order data structures.

**Visual Programming Languages:** Due to the extension of the *GEC* programming toolkit with higher-order types, *visual programming languages* have come within reach as *application domain*. One interesting example is the *Vital* system [14] in which *Haskell*-like scripts can be edited. Both systems allow direct manipulation of expressions and custom types, allow customization of views, and have guarded data types (the `selfGEC` function). In contrast with the *Vital* system, which is a dedicated system and has been implemented in *Java*, our system is a general purpose toolkit. We could use our toolkit to construct a visual environment in the spirit of *Vital*.

## 7 Conclusions

We have presented the *GEC* toolkit for rapid prototyping of type safe interactive applications. The toolkit

1. *produces type-safe interactive applications* composed from *Graphical Editor Components*;
2. is highly *automatic* due to generic generative programming techniques;
3. can be used for first order and *higher order* types;
4. can be *customized* to create any kind of user interface;
5. allows *abstraction* using model-view programming to hide details and allow type-safe view changes;
6. is *compositional* on various levels:
  - Types** standard composition of types lead to composition of corresponding graphical editor components;
  - Expressions** the user can enter expressions in which values and functions can be defined/used compositionally; these functions can even be compiled functions (possibly taken from complete applications) that are read from disk, linked in dynamically and applied in a compositional way;
  - GECs** *GECs* can be composed in an ad-hoc way by standard functional programming or in a structured way using arrow combinators;
  - AGECs** *AGECs* can be composed in a statically type safe way.
7. enables the programmer to *focus on a data type* representing the interaction with the user instead of on the many nasty details of a graphical toolkit;
8. can be *downloaded* from <http://www.cs.ru.nl/~clean/gec>.

## Acknowledgements

The authors would like to thank the referees for their detailed comments.

## References

1. M. Abadi, L. Cardelli, B. Pierce, G. Plotkin, and D. Remy. Dynamic typing in polymorphic languages. In *Proc. of the ACM SIGPLAN Workshop on ML and its Applications*, San Francisco, June 1992.

2. P. Achten. *Interactive Functional Programs - models, methods, and implementations*. PhD thesis, University of Nijmegen, The Netherlands, 1996.
3. P. Achten and S. Peyton Jones. Porting the Clean Object I/O library to Haskell. In M. Mohnen and P. Koopman, editors, *Proc. of the 12th International Workshop on the Implementation of Functional Languages, IFL'00, Selected Papers*, volume 2011 of *LNCS*, pages 194–213. Aachen, Germany, Springer, Sept. 2001.
4. P. Achten and R. Plasmeijer. Interactive Functional Objects in Clean. In C. Clack, K. Hammond, and T. Davie, editors, *Proc. of the 9th International Workshop on the Implementation of Functional Languages, IFL 1997, Selected Papers*, volume 1467 of *LNCS*, pages 304–321. St. Andrews, UK, Springer, Sept. 1998.
5. Achten, Peter and van Eekelen, Marko and Plasmeijer, Rinus and van Weelden, Arjen. Arrows for Generic Graphical Editor Components. Technical report NIII-R0416, Nijmegen Institute for Computing and Information Sciences, University of Nijmegen, The Netherlands, 2004. available at <http://www.niii.kun.nl/research/reports/full/NIII-R0416.pdf>.
6. Achten, Peter, van Eekelen, Marko and Plasmeijer, Rinus. Generic Graphical User Interfaces. In Greg Michaelson and Phil Trinder, editors, *Selected Papers of the 15th Int. Workshop on the Implementation of Functional Languages, IFL03*, volume 3145 of *LNCS*. Edinburgh, UK, Springer, 2003.
7. Achten, Peter, van Eekelen, Marko and Plasmeijer, Rinus. Compositional Model-Views with Generic Graphical User Interfaces. In *Practical Aspects of Declarative Programming, PADL04*, volume 3057 of *LNCS*, pages 39–55. Springer, 2004.
8. Achten, Peter, van Eekelen, Marko, Plasmeijer, Rinus and van Weelden, Arjen. Automatic Generation of Editors for Higher-Order Data Structures. In Wei-Ngan Chin, editor, *Second ASIAN Symposium on Programming Languages and Systems (APLAS 2004)*, volume 3302 of *LNCS*, pages 262–279. Springer, 2004.
9. A. Alimarine and R. Plasmeijer. A Generic Programming Extension for Clean. In T. Arts and M. Mohnen, editors, *The 13th International workshop on the Implementation of Functional Languages, IFL'01, Selected Papers*, volume 2312 of *LNCS*, pages 168–186. Älvsjö, Sweden, Springer, Sept. 2002.
10. E. Barendsen and S. Smetsers. *Graph Rewriting Aspects of Functional Programming*, chapter 2, pages 63–102. World scientific, 1999.
11. M. v. d. Brand, A. van Deursen, J. Heering, H. de Jong, M. de Jonge, T. Kuipers, P. Klint, L. Moonen, P. Olivier, J. Scheerder, J. Vinju, E. Visser, and J. Visser. The Asf+Sdf Meta-Environment: a Component-Based Language Development Environment. In R. Wilhelm, editor, *Compiler Construction 2001 (CC'01)*, pages 365–370. Springer-Verlag, 2001.
12. M. Carlsson and T. Hallgren. FUDGETS - a graphical user interface in a lazy functional language. In *Proc. of the ACM Conference on Functional Programming and Computer Architecture, FPCA '93*, Copenhagen, Denmark, 1993.
13. M. Flatt, S. Krishnamurthi, and M. Felleisen. Classes and mixins. In *The 25TH ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL98)*, pages 171–183, San Diego, California, 1998. ACM, New York, NY.
14. K. Hanna. Interactive Visual Functional Programming. In S. P. Jones, editor, *Proc. Intl Conf. on Functional Programming*, pages 100–112. ACM, October 2002.
15. R. Hinze. A new approach to generic functional programming. In *The 27th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 119–132. Boston, Massachusetts, January 2000.
16. R. Hinze and S. Peyton Jones. Derivable Type Classes. In G. Hutton, editor, *2000 ACM SIGPLAN Haskell Workshop*, volume 41(1) of *ENTCS*. Montreal, Canada, Elsevier Science, 2001.

17. J. Hughes. Generalising Monads to Arrows. *Science of Computer Programming*, 37:67–111, May 2000.
18. G. Krasner and S. Pope. A cookbook for using the Model-View-Controller user interface paradigm in Smalltalk-80. *Journal of Object-Oriented Programming*, 1(3):26–49, August 1988.
19. A. Löh, D. Clarke, and J. Jeuring. Dependency-style generic haskell. In *ICFP '03: Proc. of the eighth ACM SIGPLAN international conference on Functional programming*, pages 141–152. ACM Press, 2003.
20. S. Peyton Jones and Hughes J. et al. *Report on the programming language Haskell 98*. University of Yale, 1999. <http://www.haskell.org/definition/>.
21. M. Pil. Dynamic types and type dependent functions. In K. Hammond, T. Davie, and C. Clack, editors, *Implementation of Functional Languages (IFL '98)*, LNCS, pages 169–185. Springer Verlag, 1999.
22. R. Plasmeijer and M. van Eekelen. *Concurrent CLEAN Language Report (version 2.0)*, December 2001. <http://www.cs.kun.nl/~clean/contents/contents.html>.
23. M. Schrage. *Proxima, a presentation-oriented editor for structured documents*. PhD thesis, University of Utrecht, 2004.
24. A. van Weelden and R. Plasmeijer. A functional shell that dynamically combines compiled code. In P. Trinder and G. Michaelson, editors, *Selected Papers Proc. of the 15th International Workshop on Implementation of Functional Languages, IFL'03*. Heriot Watt University, Edinburgh, Sept. 2003.
25. M. Vervoort and R. Plasmeijer. Lazy dynamic input/output in the lazy functional language Clean. In R. Peña and T. Arts, editors, *The 14th International Workshop on the Implementation of Functional Languages, IFL'02, Selected Papers*, volume 2670 of *LNCS*, pages 101–117. Springer, Sept. 2003.
26. P. Wadler. Comprehending Monads. In *Proc. of the 1990 ACM Conference on Lisp and Functional Programming*, pages 61–77, Nice, France, 1990.