

COMP3161/COMP9164 Supplementary Lecture Notes

Parametric Polymorphism

Liam O'Connor

Johannes Åman Pohjola*

November 3, 2025

Polymorphism is a prominent part of most modern programming languages. It allows some form of *generic* programming, where values of *different types* can be manipulated by the *same function*.

Parametric polymorphism, sometimes called *generics* in OO languages, is the simplest form of polymorphism, where a function can be declared to operate over *any type at all*. For example, suppose we had a swap function:

$$\text{swap } (x, y) = (y, x)$$

What would the type of this function be? In a *monomorphic* language like MinHS, we couldn't write this function generically. We would have to have a variety of functions, $\text{swapBI} : \text{Bool} \times \text{Int} \rightarrow \text{Int} \times \text{Bool}$, $\text{swapIB} : \text{Int} \times \text{Bool} \rightarrow \text{Bool} \times \text{Int}$, and so on — a total of T^2 functions where T is the number of types in the language¹. This is obviously highly impractical, seeing as all these functions have the same implementation. What we want is to express a type $\text{swap} : a \times b \rightarrow b \times a$ for *all* types a and b . That is what parametric polymorphism gives us.

1 Type Parameters

Currently, all functions in MinHS take some values of a concrete type, and return a value of a concrete type. The literature on typed λ -calculus denotes these functions from *values* to *values* with the notation $\lambda(x : \tau). y$, which is analogous to $\text{recfun } (f : \tau \rightarrow \tau') x = y$ in MinHS. A function that constructs a pair of two integers could be written with this notation as follows:

$$\text{mkIntIntPair} = (\lambda(x : \text{Int}). (\lambda(y : \text{Int}). (x, y)))$$

This function takes an argument x of type Int , and returns a *function*, which, given a y of type Int , will produce a pair of x and y . This nesting of functions is how we achieve n -ary functions in Haskell and similar languages, and is called *currying*.

In order to get parametric polymorphism, we extend functions slightly. In addition to having functions from *values* to *values*, like above, we include functions from *types* to *values*, usually written like $\Lambda t. v$. These Λ binders introduce *type variables* which can be used in type signatures for values wherever it is in scope. For example, a generic mkPair function could be written like this:

$$\text{mkPair} = \Lambda a. \Lambda b. \lambda(x : a). \lambda(y : b). (x, y)$$

Applying a type to one of these generic functions is called *type application* or *specialisation*, and is written a variety of ways in the literature, including $\text{mkPair}@Int@Int$, $\text{mkPair } [Int] [Int]$ or $\text{mkPair } \{Int\} \{Int\}$.

*Minor edits.

¹Since we have products and sums, $T = \infty$

Universal Quantification

To give a type to our new $\Lambda t. e$ form, and thus to our `mkPair` function, we need to *reflect* the type variables introduced by the Λ on to the type level, where they are introduced by the *universal quantifier*, \forall :

$$\frac{\Gamma \vdash e : \tau \quad t \notin FV(\Gamma)}{\Gamma \vdash \Lambda t. e : \forall t. \tau}$$

Note that we add a requirement that the introduced variable t is not in the set of free type variables in the environment Γ . This ensures that we don't have clashing type variable names.

This *generalisation rule* allows us to provide a type to our `mkPair` function.

$$\frac{\frac{\frac{x : a; y : b \vdash x : a \quad x : a; y : b \vdash y : b}{x : a; y : b \vdash \text{pair}(x, y) : a \times b}}{x : a \vdash \lambda(y : b). \text{pair}(x, y) : b \rightarrow a \times b}}{\vdash \lambda(x : a). \lambda(y : b). \text{pair}(x, y) : a \rightarrow b \rightarrow a \times b}}{\vdash \Lambda a. \Lambda b. \lambda(x : a). \lambda(y : b). \text{pair}(x, y) : \forall b. a \rightarrow b \rightarrow a \times b}}{\vdash \Lambda a. \Lambda b. \lambda(x : a). \lambda(y : b). \text{pair}(x, y) : \forall a. \forall b. a \rightarrow b \rightarrow a \times b}$$

To type the application of our `mkPair` function, we need a type for the *type application* form, $e@t$:

$$\frac{e : \forall a. \tau'}{e@t : \tau'[a := t]}$$

This states that we can substitute the type variable a in the type for e with the type after the $@$ and get a well-typed result. Now we can type a term like `mkPair@Int@Bool 3 True`:

$$\frac{\frac{\frac{\dots \vdash \text{mkPair} : \forall ab. a \rightarrow b \rightarrow a \times b}{\dots \vdash \text{mkPair}@Int : \forall b. Int \rightarrow b \rightarrow Int \times b}}{\dots \vdash \text{mkPair}@Int@Bool : Int \rightarrow Bool \rightarrow Int \times Bool} \quad \dots \vdash 3 : Int}{\dots \vdash \text{mkPair}@Int@Bool 3 : Bool \rightarrow Int \times Bool} \quad \dots \vdash \text{True} : Bool}{\dots \vdash \text{mkPair}@Int@Bool 3 \text{ True} : Int \times Bool}$$

This lets us define functions that are *generic* over their arguments, as required, so now let us examine what extensions we need to add to MinHS to make parametric polymorphism possible in MinHS programs.

2 Applying to MinHS

2.1 New Syntax

We introduce two new forms of expression syntax: `type a. e`, which is analogous to the $\Lambda a. e$ notation from earlier, and $e@t$ for type application.

We also extend type syntax with the universal quantifier `forall a. τ` and type variables `a`, `b`, etc. This means our static semantics must ensure that types are well formed – that all type variables have an accompanying quantifier:

$$\frac{\frac{t \text{ bound} \in \Delta}{\Delta \vdash t \text{ Ok}} \quad \frac{}{\Delta \vdash \text{Int} \text{ Ok}} \quad \frac{}{\Delta \vdash \text{Bool} \text{ Ok}} \quad \frac{\Delta \vdash \alpha \text{ Ok} \quad \Delta \vdash \beta \text{ Ok}}{\Delta \vdash \alpha \rightarrow \beta \text{ Ok}}}{\frac{\Delta, a \text{ bound} \vdash \tau \text{ Ok}}{\Delta \vdash \text{forall } a. \tau \text{ Ok}}}$$

2.2 Typing Rules

The typing rules for `type a. e` are the same as the typing rules for $\Lambda a. e$, only using our explicit wellformedness checking for types (which necessitates an additional context of type variables):

$$\frac{a \text{ bound}, \Delta; \Gamma \vdash e : \tau}{\Delta; \Gamma \vdash \text{type } a. e : \forall a. \tau}$$

Similarly for type application:

$$\frac{\Delta; \Gamma \vdash e : \text{forall } a. \tau \quad \Delta \vdash \rho \text{ ok}}{\Delta; \Gamma \vdash e@_\rho : \tau[a := \rho]}$$

2.3 Dynamic Semantics

Firstly, we evaluate the expression in the left-hand side as much as possible:

$$\frac{e \mapsto_M e'}{e@_\tau \mapsto_M e'@_\tau}$$

Once it has fully evaluated, we expect to see a **type**-abstraction remaining. Then we apply a substitution, similarly to normal function application (i.e. β -reduction):

$$\overline{(\text{type } a. e)@_\tau \mapsto_M e[a := \tau]}$$

3 Implementation Techniques

Our simple dynamic semantics belies a complex implementation headache.

While we can easily define functions that operate uniformly on multiple types, when this is compiled to machine code the results may differ depending on the *size* of the type in question on the machine stack. For example, a function that takes a `int` parameter will allocate a different amount of stack space than a function that takes a `long` parameter.

There are two main approaches to solve this problem:

Template Specialisation In this approach, we automatically generate monomorphic copies of each polymorphic function, one for each of the different types applied to it. For example, if we defined our polymorphic swap function:

$$\begin{aligned} \text{swap} &= \text{type } a. \text{ type } b. \\ &\quad \text{recfun } \text{swap} :: (a \times b) \rightarrow (b \times a) \\ &\quad \quad p = (\text{snd } p, \text{fst } p) \end{aligned}$$

Then a type application like `swap@Int@Bool` would be replaced *statically* by the compiler with the monomorphic version:

$$\begin{aligned} \text{swap}_{\text{IB}} &= \text{recfun } \text{swap} :: (\text{Int} \times \text{Bool}) \rightarrow (\text{Bool} \times \text{Int}) \\ &\quad \quad p = (\text{snd } p, \text{fst } p) \end{aligned}$$

A new copy is made for each unique type application.

This approach is not too difficult to implement, produces code with minimal run-time overhead, and is relatively easy for a programmer to understand and debug. It also allows the compiler to make optimisations based on the specific types used for a polymorphic function. For example, a list of booleans could be represented more efficiently as a bit vector.

However, as a copy is made for each type application, the binary sizes end up rather large, which can slow compilation times. More importantly, it imposes a severe restriction on the kinds of polymorphism that can be allowed — we must be able to determine all of the possible type instantiations *statically*, which rules out *polymorphic recursion* and other situations where the type instantiations is dependent on run-time information.

Boxing An alternative to the copy-paste-heavy template instantiation approach is to make *all types* represented on the stack in the same way. Thus, a polymorphic function only requires one function in the generated code.

Typically this is done by *boxing* each type. That is, all data types are represented as a *pointer* to a data structure on the heap. If everything is a pointer, then all values use exactly 32 (or 64) bits of stack space.

The extra indirection has a run-time penalty, and it can make garbage collection more necessary, but it results in smaller binaries and unrestricted polymorphism.

4 Generality and Parametricity

If we need a function of type $\text{Int} \rightarrow \text{Int}$, a polymorphic function of type $\text{forall } a. a \rightarrow a$ will do just fine, we can just instantiate the type variable to Int . But the reverse is not true. This gives rise to an ordering.

We say that a type τ is *more general* than a type ρ , often written $\rho \sqsubseteq \tau$, if type variables in τ can be instantiated to give the type ρ . For example, in the below example, the *most general* type is $\forall a. a$, as a could be instantiated to any type, therefore this type is more general than all other types.

$$\text{Int} \rightarrow \text{Int} \sqsubseteq \forall z. z \rightarrow z \sqsubseteq \forall x y. x \rightarrow y \sqsubseteq \forall a. a$$

As we get more and more general, the number of possible total, terminating implementations of the type gets smaller and smaller:

Type	Size
$\text{Int} \rightarrow \text{Int}$	$(2^{32})^{2^{32}}$
$\forall z. z \rightarrow z$	1
$\forall a b. a \rightarrow b$	0
$\forall a. a$	0

This means that the more general we make our type, the more constrained we are in implementation. This allows us to conclude some things about a function just by looking at its type signature.

For example, consider this list function *foo*:

$$\text{foo} :: \forall a. [a] \rightarrow [a]$$

This function has a type that immediately lets me conclude that the output list consists only of elements in the input list. This is because *foo* is defined regardless of what the type a is. Therefore, the only values of type a that *foo* knows about are those values that are in the input. Therefore, the output list can only be made from elements in the input.

The formal principle that captures this intuition is called *parametricity*.

Formally, we can say that if we apply any arbitrary function f to the polymorphically-typed inputs of a polymorphic function, then that is equivalent to applying that same function to the polymorphically-typed outputs of the same function. For example, for *foo* above, the parametricity theorem is, for any f and ls :

$$\text{map } f (\text{foo } ls) = \text{foo } (\text{map } f ls)$$

The identity function, $\text{id} :: \forall a. a \rightarrow a$, has an even simpler theorem:

$$\text{id } (f x) = f (\text{id } x)$$

There are many other examples, like *head*:

$$\text{head} :: \forall a. [a] \rightarrow a$$

Which has the following theorem:

$$f (\text{head } \ell) = \text{head } (\text{map } f \ell)$$

The list appending function:

$$(++) :: \forall a. [a] \rightarrow [a] \rightarrow [a]$$

Which has the following theorem:

$$\text{map } f (a ++ b) = \text{map } f a ++ \text{map } f b$$

Concatenating a list of lists:

$$\text{concat} :: \forall a. [[a]] \rightarrow [a]$$

Has the following theorem:

$$\text{map } f (\text{concat } ls) = \text{concat } (\text{map } (\text{map } f) ls)$$

It even works for higher order functions, like *filter*:

$$\text{filter} :: \forall a. (a \rightarrow \text{Bool}) \rightarrow [a] \rightarrow [a]$$

Which has the following theorem:

$$\text{filter } p (\text{map } f ls) = \text{map } f (\text{filter } (p \circ f) ls)$$