

Chapter 7

Parametricity

Write down the definition of a polymorphic function on a piece of paper. Tell me its type, but be careful not to let me see the function's definition. I will tell you a theorem that the function satisfies. – Philip Wadler, “Theorems for free!”

Polymorphism allows a single piece of code to be instantiated with multiple types. Polymorphism is *parametric* when all of the instances behave *uniformly*. This is in contrast to ad-hoc polymorphism, where values can behave differently depending on which type they are being instantiated with.

Parametricity can be thought of as the dual to abstraction. Where abstraction hides details about an implementation from the outside world, parametricity hides details about the outside world from an implementation.

7.1 Parametricity in OCaml

7.1.1 Universal types in OCaml

Universal types are the most common form of parametricity in OCaml. As discussed in Chapter 3, ML polymorphism provides simple universal types without any type annotations. For example, the polymorphic identity function, which has type $\forall\alpha.\alpha \rightarrow \alpha$:

```
let f x = x
```

However, this only provides *rank-1* or *prenex polymorphism*, which means that all universal quantifiers (\forall) must appear at the out-most position (i.e. at the very beginning of the type expression). It does not support higher-rank universal types such as $(\forall\alpha.\text{List}\alpha \rightarrow \text{Int}) \rightarrow \text{Int}$.

Type inference for higher-rank universal types is undecidable in general. As an illustration, consider the following OCaml function:

```
fun f x y -> f x + f y
```

This expression could have a number of System F ω types, including:

```
 $\forall\alpha::*. (\alpha \rightarrow \text{Int}) \rightarrow \alpha \rightarrow \alpha \rightarrow \text{Int}$ 
```

```
 $\forall\alpha::*.\forall\beta::*.\ (\forall\gamma::*.\ \gamma \rightarrow \text{Int}) \rightarrow \alpha \rightarrow \beta \rightarrow \text{Int}$ 
```

and none of these types is more general than the others, so we require some annotations in order to type-check programs involving higher-rank universals. The required annotations include explicit type abstraction and type application statements, as well as explicitly specifying the type of the universal created by a type abstraction.

Rather than directly using type application and type abstraction with type annotations, universal types in OCaml are provided through record types (similar support is also available through object types). Constructing the record type acts as a type abstraction statement, and destructing the record using a pattern or projection acts as a type application statement. The declaration of a record type includes specifying the types of its fields, which provide us with the required type annotations for type abstraction statements.

The following definition defines a type corresponding to $\forall\alpha.\text{List } \alpha \rightarrow \text{Int}$:

```
type t = { f : 'a. 'a list -> int }
```

Building a value of type `t` corresponds to a type abstraction operation:

```
let len = {f = List.length}
```

Destructing a value of type `t` using a pattern or projection corresponds to type application:

```
let g r = r.f [1; 2; 3] + r.f [1.0; 2.0; 3.0]
```

7.1.2 Functors

The most powerful form of parametricity in OCaml is provided by *functors*. Functors are functions that operate on modules. Since modules can contain any OCaml definition, functors can be parametric on any OCaml definition.

As an example, we will extend our `IntSet` module to a `Set` module that works uniformly on any module that matches the appropriate signature.

For convenience, we will define a module type for the signature that we expect of arguments to the functor:

```
module type Eq = sig
  type t
  val equal : t -> t -> bool
end
```

This signature says that we expect the module to contain a type `t` and a function `equal` for comparing two values of type `t` for equality.

We also create a signature of the functor's result:

```
module type SetS = sig
  type t
  type elt
  val empty : t
  val is_empty : t -> bool
```

```

val mem : elt -> t -> bool
val add : elt -> t -> t
val remove : elt -> t -> t
val to_list : t -> elt list
end

```

Now we define our `Set` functor as follows:

```

module Set (E : Eq) : SetS with type elt := E.t = struct

  type t = E.t list

  let empty = []

  let is_empty = function
    | [] -> true
    | _ -> false

  let rec mem x = function
    | [] -> false
    | y :: rest ->
      if (E.equal x y) then true
      else mem x rest

  let add x t =
    if (mem x t) then t
    else x :: t

  let rec remove x = function
    | [] -> []
    | y :: rest ->
      if (E.equal x y) then rest
      else y :: (remove x rest)

  let to_list t = t

end

```

Note that we have specified the result of the functor to have the signature `SetS with type elt := E.t`. This signature is the same as `SetS` but with the `elt` type component removed and all occurrences of it replaced by `E.t`. If we had wanted to leave the `elt` type in the signature, but changed from an abstract type to an alias for `E.t`, we could instead have used `SetS with type elt = E.t`.

We can apply the functor to recreate `IntSet`:

```

module IntEq = struct
  type t = int
  let equal (x : int) (y : int) =
    x = y
end

module IntSet = Set(IntEq)

```

Since the type `E.t` is abstract within the body of `Set`, the implementation of `Set` can only pass these values around and compare them using `E.equals`, it cannot use values of type `E.t` in any other way because the it cannot see the

type's definition. This means that the behaviour of `Set` cannot depend on the particular type used for `E.t`: it must behave uniformly on any `Eq` module that it is applied to.

7.2 Higher-kinded polymorphism

Modules and functors allow abstraction and parametricity for both types (with kind `*`) and type constructors (with kind `* → *`). However, the existential and universal polymorphism in the core language only supports types with kind `*`. The reason for this is that type-checking higher-kinded polymorphism, requires type annotations on type application as well as type abstraction. The module system includes such annotations, but the core language avoids all such annotations as too verbose.

As an example, consider trying to type-check applications of a function with type $\forall F :: * \rightarrow *. \forall \alpha :: *. F \alpha \rightarrow (F \alpha \rightarrow \alpha) \rightarrow \alpha$ to a value with type `List(Int × Int)`. This involves unifying the following two types:

$$F \alpha \sim \text{List}(\text{Int} \times \text{Int})$$

There are many possible solutions to this unification, including:

$$\begin{array}{ll} F = \text{List} & \alpha = \text{Int} \times \text{Int} \\ F = \Lambda \beta. \text{List}(\beta \times \beta) & \alpha = \text{Int} \\ F = \Lambda \beta. \text{List}(\text{Int} \times \text{Int}) & \end{array}$$

None of which is more general than the others.

7.2.1 Lightweight higher-kinded polymorphism

It is possible to restrict higher-kinded polymorphism so that it can be type-checked without type annotations on type applications. By restricting the type functions which can be used to a set of functions for which each function maps to a different set of types. In other words, a set \mathbf{F} of functions such that:

$$\forall F, G \in \mathbf{F}. F \neq G \Rightarrow \forall t. F(t) \neq G(t)$$

This is how Haskell provides higher-kinded polymorphism: it only supports type constructors which create fresh types. For example, Haskell has no equivalent of the OCaml type¹:

```
type 'a t = ('a * 'a) list
```

This restricted form of higher-kinded polymorphism can also be used in OCaml with an encoding based on defunctionalisation (Reynolds [1972]). For example, the following type definition:

¹A Haskell type synonym can be created which looks like this OCaml definition, however it is very different because it cannot be abstracted

```

type lst = List
type opt = Option

type ('a, 'f) app =
  | Lst : 'a list -> ('a, lst) app
  | Opt : 'a option -> ('a, opt) app

```

Allows us to represent 'a list as ('a, lst) app and 'a option as ('a, opt) app. Then we can use polymorphism in the second parameter of app to encode higher-kinded polymorphism over list and option. For example:

```

type 'f map = {
  map: 'a 'b. ('a -> 'b) -> ('a, 'f) app -> ('b, 'f) app;
}

let lmap : lst map = {map = fun f (Lst l) -> Lst (List.map f l)}
let omap : opt map = {map = fun f (Opt o) -> Opt (Option.map f o)}

let f : 'b map -> (int, 'b) app -> (string, 'b) app =
  fun m c -> m.map (fun x -> "Int: " ^ (string_of_int x)) c

let l = f lmap (Lst [1; 2; 3])
let o = f omap (Opt (Some 6))

```

The higher library ([Yallop and White \[2014\]](#)), extends this idea by providing a general app type that any type constructor can be injected into.

7.3 Parametricity in System F ω

7.3.1 Universal types

OCaml functors with abstract types in their arguments correspond to universal types in System F ω . This is not surprising: universal types are the fundamental concept of the *polymorphic* lambda calculus (System F), which is intended to capture the essence of parametric polymorphism.

As an example we will implement our Set functor using universals in System F ω .

First, for convenience, we create a type constructor and some functions for dealing with the products that represent the structures of both the argument and result structures of Set:

```

EqImpl =
  λγ::*. γ → γ → Bool;

equal = Λγ::*. λs:EqImpl γ.s;

SetImpl =
  λγ::*. λα::*.
    α
    × (α → Bool)
    × (γ → α → Bool)
    × (γ → α → α)
    × (γ → α → α)
    × (α → List γ);

```

```

empty = Λγ::*.Λα::*.λs:SetImpl γ α.π1 s;
is_empty = Λγ::*.Λα::*.λs:SetImpl γ α.π2 s;
mem = Λγ::*.Λα::*.λs:SetImpl γ α.π3 s;
add = Λγ::*.Λα::*.λs:SetImpl γ α.π4 s;
remove = Λγ::*.Λα::*.λs:SetImpl γ α.π5 s;
to_list = Λγ::*.Λα::*.λs:SetImpl γ α.π6 s;

```

Now we can create our implementation of sets, using Λ to give it a universal type corresponding to the type of the `Set` functor.

```

set_package =
  Λγ::*. λeq:EqImpl γ.
  pack List γ, <
    nil [γ],
    isempty [γ],
    λn:γ.fold [γ] [Bool]
      (λx:γ.λy:Bool.or y (equal [γ] eq n x))
      false,
    cons [γ],
    λn:γ.fold [γ] [List γ]
      (λx:γ.λl>List γ.
        if (equal [γ] eq n x) [List γ] l (cons [γ] x l))
      (nil [γ]),
    λl>List γ.l >
  as ∃α::*.SetImpl γ α;

```

If we look at the typing rules for universals (Section 2.3), we can see that the type parameter of Λ is a fresh type variable. As with abstract types in the parameters of OCaml functors, this means the body of `set_package` can only pass these values around and compare them using `eq`, it cannot use values of type γ in any other way because the it cannot see the type's definition.

This means that the behaviour of `set_package` cannot depend on the particular type used for γ : it must behave uniformly on any type that it is applied to.

7.3.2 Relational parametricity

We can also use the relational substitution to give a precise meaning to the idea that universal types provide parametricity.

Given a type T with free variables $\alpha, \beta_1, \dots, \beta_n$:

$$\begin{aligned}
& \forall \beta_1. \dots \forall \beta_n. \forall x : (\forall \alpha. T). \\
& \forall \gamma. \forall \delta. \forall \rho \subset \gamma \times \delta. \\
& T[\rho, =_{\beta_1}, \dots, =_{\beta_n}](x[\gamma], x[\delta])
\end{aligned}$$

This formula can be read as: For a value x with universal type, two types γ and δ , and any way of viewing γ and δ as representing the same thing – captured by a relation ρ – $x[\gamma]$ will behave the same as $x[\delta]$ with respect to ρ .

This is the essence of parametricity: a value with universal type will behave uniformly for any types it is applied to.

The above parametricity property for universal types can be derived from identity extension applied to universals by expanding out the relational substitution.

7.3.3 Theorems for free

In the previous chapter we applied the abstraction property to a particular kind of relation to show that abstraction guarantees code outside of a component preserves invariants. Similarly, we can apply the parametricity property to particular relations to show that parametricity guarantees code inside of a component behaves in certain ways.

For example, applying parametricity to the type $\forall\alpha.\alpha \rightarrow \alpha$, gives us the following formula:

$$\begin{aligned} \forall f : (\forall\alpha.\alpha \rightarrow \alpha). \\ \forall\gamma. \forall\delta. \forall\rho \subset \gamma \times \delta. \\ \forall u : \gamma. \forall v : \delta. \\ \rho(u, v) \Rightarrow \rho(f[\gamma]u, f[\delta]v) \end{aligned}$$

By defining a relation is_u to represent being equal to a value $u : T$:

$$\text{is}_u(x : T, y : T) = (x =_T u) \wedge (y =_T u)$$

and using it to instantiate ρ , we obtain the following formula:

$$\begin{aligned} \forall f : (\forall\alpha.\alpha \rightarrow \alpha). \\ \forall\gamma. \forall u : \gamma. \\ \text{is}_u(u, u) \Rightarrow \text{is}_u(f[\gamma]u, f[\gamma]u) \end{aligned}$$

which can be reduced to:

$$\begin{aligned} \forall f : (\forall\alpha.\alpha \rightarrow \alpha). \\ \forall\gamma. \forall u : \gamma. \\ f[\gamma]u =_\gamma u \end{aligned}$$

This shows that any value f of type $\forall\alpha.\alpha \rightarrow \alpha$ must be the identity function. Properties like this, which use parametricity to give guarantees about the behaviour of all values of a given type, are often called *free theorems* after an influential paper on the subject (Wadler [1989]).

As a second example, we can apply parametricity to the type $\forall\alpha.\text{List } \alpha \rightarrow \text{List } \alpha$:

$$\begin{aligned} \forall f : (\forall\alpha.\text{List } \alpha \rightarrow \text{List } \alpha). \\ \forall\gamma. \forall\delta. \forall\rho \subset \gamma \times \delta. \\ \forall u : \text{List } \gamma. \forall v : \text{List } \delta. \\ (\text{List } \alpha)[\rho](u, v) \Rightarrow (\text{List } \alpha)[\rho](f[\gamma]u, f[\delta]v) \end{aligned} \tag{7.1}$$

Applying the relational substitution to the System F encoding of `List` gives the following equation:

$$\begin{aligned}
(\text{List } \alpha)[\rho] = & \\
& (x : \text{List } A, y : \text{List } B). \\
& \forall \gamma. \forall \delta. \forall \rho' \subset \gamma \times \delta. \\
& \forall n : \gamma. \forall m : \delta. \\
& \forall c : A \rightarrow \gamma \rightarrow \gamma. \forall d : B \rightarrow \delta \rightarrow \delta. \\
& \rho'(n, m) \Rightarrow \\
& (\forall a : A. \forall b : B. \forall u : \gamma. \forall v : \delta. \\
& \rho(a, b) \Rightarrow \rho'(u, v) \Rightarrow \rho'(c a u, d b v)) \Rightarrow \\
& \rho'(x[\gamma] n c, y[\delta] m d)
\end{aligned}$$

which can be reduced to a more manageable form by considering the cases of `nil` and `cons` separately:

$$\begin{aligned}
(\text{List } \alpha)[\rho](x : \text{List } A, y : \text{List } B) = & \\
\left\{ \begin{array}{ll} \rho(i, j) \wedge (\text{List } \alpha)[\rho](l, k), & x = \text{cons}_A i l \wedge y = \text{cons}_B j k \\ \text{true}, & x = \text{nil}_A \wedge y = \text{nil}_B \\ \text{false}, & \text{otherwise} \end{array} \right. & (7.2)
\end{aligned}$$

By defining a family of relations $\langle g \rangle$ to represent functions $g : A \rightarrow B$:

$$\langle g \rangle(x : A, y : B) = (g x =_B y)$$

and substituting it for ρ in 7.2, we obtain the following equation:

$$\begin{aligned}
(\text{List } \alpha)[\langle g \rangle](x : \text{List } A, y : \text{List } B) = & \\
\left\{ \begin{array}{ll} g i =_B j \wedge (\text{List } \alpha)[\langle g \rangle](l, k), & x = \text{cons}_A i l \wedge y = \text{cons}_B j k \\ \text{true}, & x = \text{nil}_A \wedge y = \text{nil}_B \\ \text{false}, & \text{otherwise} \end{array} \right.
\end{aligned}$$

This recursive equation may look quite familiar. It bears a striking resemblance to the definition of the standard `map` function for lists. In fact, we can use `map` to rewrite it as:

$$\begin{aligned}
(\text{List } \alpha)[\langle g \rangle](x : \text{List } A, y : \text{List } B) = & \\
\text{map}[A][B] g x =_B y & (7.3)
\end{aligned}$$

Substituting 7.3 into 7.1 we get a free theorem for $\forall \alpha. \text{List } \alpha \rightarrow \text{List } \alpha$:

$$\begin{aligned}
\forall f : (\forall \alpha. \text{List } \alpha \rightarrow \text{List } \alpha). & \\
\forall \gamma. \forall \delta. \forall g : \gamma \rightarrow \delta & \\
\forall u : \text{List } \gamma. \forall v : \text{List } \delta. & \\
\text{map}[\gamma][\delta] g (f[\gamma] u) = f[\delta] (\text{map}[\gamma][\delta] g u) &
\end{aligned}$$

An interesting lemma of this theorem is that any function f of type $\forall \alpha. \text{List } \alpha \rightarrow \text{List } \alpha$ is a “rearrangement” function: the output of f is a list whose elements all come from the input to f .

7.4 Practical limitations

The previous chapter and this one have described how abstraction and parametricity can give guarantees about program behaviour. However, these guarantees rely on some assumptions that do not necessarily hold in real programming languages.

7.4.1 Side-effects

In a *pure* language such as System $F\omega$, a function accepts an input value and produces an output value. The relation between inputs and outputs completely defines the behaviour of a function. However, programming languages are not generally pure, they allow functions to perform *side-effects*. Such side-effects include:

- Printing to the console
- Raising exceptions
- Failing to terminate

Side-effects affect the guarantees that are provided by abstraction and parametricity. For example, we showed that parametricity ensures that all functions of type $\forall\alpha.\alpha \rightarrow \alpha$ are the identity function. However, the following three OCaml functions have that type and are not the identity function:

```
let f (x : 'a) : 'a =
  Printf.printf "Launch missiles\n";
  x

let f (x : 'a) : 'a = raise Exit

let rec f (x : 'a) : 'a = f x
```

The issue here is that the function type in OCaml is a different type from the function type in System F – since it supports side-effects. This means that the relational substitution for OCaml’s function type is different from the relational substitution that we gave for System F . The relational substitution for OCaml’s function type should ensure not only that related inputs produce related outputs, but that the side-effects of the function preserve relations.

This has a number of consequences in practice:

- When replacing one implementation with another, we must ensure that the side-effects of the two implementations are also equivalent.
- When using abstraction to preserve an invariant, we must also ensure that the side effects preserve that invariant.
- When interpreting a “free theorem” we must consider the possible affect of side-effects on a function’s behaviour.

For example, whilst we do not know that all functions of type $\forall\alpha.\alpha \rightarrow \alpha$ are the identity function, we do know that if such a function returns a value then it will be equal to the function's input.

7.4.2 Non-parametric values

Abstraction and parametricity only hold if there are no non-parametric polymorphic functions available in the environment.

For example, applying parametricity to the type $\forall\alpha.\alpha \rightarrow \alpha \rightarrow \text{Bool}$ gives us the following formula:

$$\begin{aligned} \forall f : (\forall\alpha.\alpha \rightarrow \alpha \rightarrow \text{Bool}). \\ \forall\gamma. \forall\delta. \forall\rho \subset \gamma \times \delta. \\ \forall u : \gamma. \forall v : \delta. \forall u' : \gamma. \forall v' : \delta. \\ \rho(u, v) \Rightarrow \rho(u', v') \Rightarrow \\ (f[\gamma] u u' =_{\text{Bool}} f[\delta] v v') \end{aligned}$$

If we instantiate ρ with the trivial relation that is always true, then we get the following free theorem:

$$\begin{aligned} \forall f : (\forall\alpha.\alpha \rightarrow \alpha \rightarrow \text{Bool}). \\ \forall\gamma. \forall\delta. \\ \forall u : \gamma. \forall v : \delta. \forall u' : \gamma. \forall v' : \delta. \\ (f[\gamma] u u' =_{\text{Bool}} f[\delta] v v') \end{aligned}$$

This means that any parametric function of type $\forall\alpha.\alpha \rightarrow \alpha \rightarrow \text{Bool}$ must be a *constant function*: a function that returns the same value for all inputs.

However, OCaml provides the following built-in structural equality function:

```
val (=) : 'a -> 'a -> bool
```

Even though it has this type, it is not a constant function. This means that it is not parametric.

The existence of this function and several similar ones in OCaml, can break the guarantees provided by abstraction and parametricity. However, not all guarantees are broken, for instance the preservation of invariants is not affected by any of the built-in functions available in OCaml.

In practice, this means that such functions should only be used at known types. Using them on abstract types produces non-parametric code whose correctness may rely on details of an implementation that are not exposed in its interface.