

Lecture Notes on Data Abstraction

15-814: Types and Programming Languages
Frank Pfenning

Lecture 17
Thursday, October 29, 2020

1 Introduction

Since we have moved from the pure λ -calculus to functional programming languages we have added rich type constructs starting from functions, disjoint sums, eager and lazy pairs, recursive types, and parametric polymorphism. The primary reasons often quoted for such a rich static type system are discovery of errors before the program is ever executed and the efficiency of avoiding tagging of runtime values. There is also the value of the types as documentation and the programming discipline that follows the prescription of types. Perhaps more important than all of these is the strong guarantees of data abstraction that the type system affords that are sadly missing from many other languages. Indeed, this was one of the original motivation in the development of ML (which stands for MetaLanguage) by Milner and his collaborators [GMM⁺78]. They were interested in developing a theorem prover and wanted to reduce its overall correctness to the correctness of a trusted core. To this end they specified an *abstract type of theorem* on which the only allowed operations are inference rules of the underlying logic. The connection between abstract types and existential types was made made Mitchell and Plotkin [MP88].

In this lecture we will first explore some more consequences of Reynolds's parametricity theorem that are used in modern compilers and then move towards questions of data abstraction and modularity.

2 Theorems for Free!

A slightly different style of application of parametricity is laid out in Philip Wadler's *Theorems for Free!* [Wad89]. Let's see what we can derive from

$$f : \forall \alpha. \alpha \rightarrow \alpha$$

for a value f . First, parametricity tells us

$$f \sim f \in [\forall \alpha. \alpha \rightarrow \alpha]$$

This time, we pick types τ and τ' and a *relation* R which is in fact a *function* $R : \tau \rightarrow \tau'$. Then

$$f[\tau] \approx f[\tau'] \in \llbracket R \rightarrow R \rrbracket$$

which means that $f[\tau] \mapsto^* f_\tau$ and $f[\tau'] \mapsto^* f_{\tau'}$ with

$$f_\tau \sim f_{\tau'} \in [R \rightarrow R]$$

Now, for arbitrary values $v : \tau$ and $v' : \tau'$, $v R v'$ actually means $R v \mapsto^* v'$. Using the definition of \sim at function type we get

$$f_\tau v \approx f_{\tau'} (R v) \in \llbracket R \rrbracket$$

but this in turn means

$$R (f_\tau v) \mapsto^* w \quad \text{and} \quad f_{\tau'} (R v) \mapsto^* w \quad \text{for some value } w$$

Wadler summarizes this by stating that for any function $R : \tau \rightarrow \tau'$,

$$R \circ f[\tau] = f[\tau'] \circ R$$

that is, f commutes with any function R . If τ is non-empty and we have $v_0 : \tau$ and choose $\tau' = \tau$ and $R = \lambda x. v_0$ we obtain

$$\begin{aligned} R (f[\tau] v_0) &\mapsto^* v_0 \\ f[\tau] (R v_0) &\mapsto^* f[\tau] v_0 \end{aligned}$$

so we find $f[\tau] v_0 \mapsto^* v_0$ which, since v_0 was arbitrary, is another way of saying that f behaves like the identity function.

3 Parametricity on Lists

For more interesting examples, we extend the notion of logical equivalence to lists. Since lists are inductively defined, we can call upon a general theory to handle them, but since we haven't discussed this theory we give the specific definition. Here, we think of lists defined with

$$\text{list } \tau = \rho\alpha. (\mathbf{nil} : 1) + (\mathbf{cons} : \tau \times \alpha)$$

even though type constructors like list haven't been formally introduced into our language. Then we use a shorthand notation for lists, that is, elaborate the left-hand side into the right-hand side:

$$[e_1, \dots, e_n] \triangleq \text{fold } \mathbf{cons} \cdot \langle e_1, \dots \text{fold } \mathbf{cons} \cdot \langle e_n, \text{fold } \mathbf{nil} \cdot \langle \rangle \rangle \rangle$$

We then extend the notion of logical equalities to values of list type *inductively* over the structure of the list, which reduces the type of the relation because each element has a smaller type.

$$e \sim e' : \text{list } \tau \text{ iff } e \mapsto^* [v_1, \dots, v_n], e' \mapsto^* [v'_1, \dots, v'_n] \text{ and } v_i \sim v'_i : \tau \text{ for all } 1 \leq i \leq n.$$

Then we have, for example, a polymorphic *map* function:

$$\begin{aligned} \text{map} &: \forall\alpha. \forall\beta. (\alpha \rightarrow \beta) \rightarrow \text{list } \alpha \rightarrow \text{list } \beta \\ \text{map} &= \Lambda\alpha. \Lambda\beta. \text{fix } m. \lambda f. \lambda l. \\ &\quad \text{case } l \text{ (fold } \mathbf{nil} \cdot \langle \rangle \Rightarrow \text{fold } \mathbf{nil} \cdot \langle \rangle \\ &\quad \quad | \text{fold } \mathbf{cons} \cdot \langle x, l' \rangle \Rightarrow \text{fold } \mathbf{cons} \cdot \langle f x, m f l' \rangle) \end{aligned}$$

The map function then satisfies (for $f : \tau \rightarrow \tau'$):

$$\text{map } [\tau] [\tau'] f [v_1, \dots, v_n] = [f v_1, \dots, f v_n]$$

where equality here is Kleene equality (both sides reduce to the same value). The example(s) are easier to understand if we isolate the special case list R for a relation $R : \tau \rightarrow \tau'$ which is actually a function. In this case we obtain

$$v \approx v' \in [\text{list } R] \text{ for an } R : \tau \rightarrow \tau' \text{ iff } (\text{map } [\tau] [\tau'] R) v = v'.$$

Returning to examples, what can the type tell us about a function

$$f : \forall\alpha. \text{list } \alpha \rightarrow \alpha \text{ list } \alpha ?$$

If the function is parametric, it should not be able to examine the list elements, or create new ones. However, it should be able to drop elements,

duplicate elements, or rearrange them. We will try to capture this equationally, just following our nose in using parametricity to see what we end up at.

We start with

$$f \sim f \in [\forall \alpha. \text{list } \alpha \rightarrow \text{list } \alpha] \text{ by parametricity.}$$

Now let $R : \tau \rightarrow \tau'$ be a function. Then $f[\tau] \mapsto^* f_\tau$, $f[\tau'] \mapsto^* f_{\tau'}$, and

$$f_\tau \sim f_{\tau'} \in [\text{list } R \rightarrow \text{list } R] \text{ by definition of } \approx.$$

Using the definition of \sim on function types, we obtain

$$\text{For any values } l : \text{list } \tau \text{ and } l' : \text{list } \tau' \text{ with } l (R \text{ list}) l' \text{ we have} \\ f_\tau l (R \text{ list}) f_{\tau'} l'$$

By the remark on the interpretation of $R \text{ list}$ when R is a function, this becomes

$$\text{If } (\text{map } [\tau] [\tau'] R) l = l' \text{ then } (\text{map } [\tau] [\tau'] R) (f l) = f l'$$

or, equivalently,

$$(\text{map } [\tau] [\tau'] R) (f [\tau] l) = f [\tau'] ((\text{map } [\tau] [\tau'] R) l).$$

In short, f commutes with $\text{map } R$. This means we can either map R over the list and then apply f to the result, or we can apply f first and then map R over the result. This implies that f could not, say, make up a new element v_0 not in l . Such an element would occur in the list returned by the right-hand side, but would occur as $R v_0$ on the left-hand side. So if we have a type with more than one element we can choose R so that $R v_0 \neq v_0$ (like a constant function) and the two sides would be different, contradicting the equality we derived.

We can use this equation to improve efficiency of code. For example, if we know that f might reduce the number of elements in the list (for example, skipping every other element), then mapping R over the list after the elements have been eliminated is more efficient than the other way around. Conversely, if f may duplicate some elements then it would be more efficient to map R over the list first and then apply f . The equality we derived from parametricity allows this kind of optimization.

We have, however, to be careful when nonterminating functions may be involved. For example, if R diverges on an element v_0 then the two sides may not be equal. For example, f might drop v_0 from the list l so the right-hand side would diverge while the left-hand side would have a value.

Here are two other similar results provided by Wadler [Wad89].

$$f : \forall \alpha. (\alpha \text{ list}) \text{ list} \rightarrow \alpha \text{ list}$$

$$(\text{map } [\tau] [\tau'] R) (f [\tau] l) = f [\tau'] ((\text{map } [\text{list } \tau] [\text{list } \tau'] (\text{map } [\tau] [\tau'] R)) l)$$

$$f : \forall \alpha. (\alpha \rightarrow \text{bool}) \rightarrow \alpha \text{ list} \rightarrow \alpha \text{ list}$$

$$(\text{map } [\tau] [\tau'] R) (f [\tau] (\lambda x. p (R x)) l) = f [\tau'] p ((\text{map } [\tau] [\tau'] R) l)$$

These theorems do not quite come “for free”, but they are fairly straightforward consequences of parametricity, keeping in mind the requirement of termination.

4 Signatures and Structures

Data abstraction in today’s programming languages is usually enforced at the level of modules (if it is enforced at all). As a running example we consider a simple module providing an implementation of a counter with constant *new* and functions *inc* and *dec* to increment and decrement the counter. We will consider two implementations and their relationship. One is using numbers in unary form (type *nat*) and numbers in binary form (type *bin*), and we will eventually prove that they are logically equivalent. We are making up some syntax (loosely based on ML) to specify interfaces between a library and its client.

Below we name *CTR* as the *signature* that describes the interface of a module.

```
CTR = {
  type ctr
  new : ctr
  inc : ctr → ctr
  dec : ctr → 1 + ctr
}
```

The value *new* will be a counter with initial value 0. The decrement function *dec* returns an optional counter with the new value, where we consider the predecessor of 0 to be undefined (returning $1 \cdot \langle \rangle$). This provides the only means for the client to observe the value of a counter. The implementations are straightforward so we elide them for now, and just assume we have type *nat* and *bin* and suitable functions on them.

```

NatCtr : CTR = {
  type ctr = nat
  new = zero
  inc = succ
  dec = pred
}

```

An interesting aspect of this definition is that, for example, $zero : nat$ while the interface specifies $new : ctr$. But this is okay because the type ctr is in fact implemented by nat in this version. Next, we show the implementation using numbers in binary representation. This time, we define some of the functions directly in the module, assuming $e : bin$ represents 0, and we have suitable functions $plus1$ and $minus1$ on binary numbers already defined.

```

BinCtr : CTR = {
  type ctr = nat
  new = e
  inc = plus1
  dec = minus1
}

```

Now what does a client look like? Assume it has an implementation $C : CTR$. It can then “open” or “import” this implementation to use its components, but it will not have any knowledge about the type of the implementation. For example, we would write

```

open C : CTR
isZero : ctr → bool
isZero = λx. case dec x (1 · ⟨ ⟩ ⇒ true
                       | r · _ ⇒ false)

```

but **not**

```

open C : CTR
isZero : ctr → bool
isZero = λn. case (unfold n) (zero · ⟨ ⟩ ⇒ true
                             | succ · _ ⇒ false)

```

because the latter supposes that the library $C : CTR$ implements the type ctr by nat , which it may not.

5 Formalizing Abstract Types

We will write a signature such as

$$CTR = \{$$

type ctr

$new : ctr$

$inc : ctr \rightarrow ctr$

$dec : ctr \rightarrow 1 + ctr$

$$\}$$

in abstract form as

$$\exists\alpha. \underbrace{\alpha}_{new} \times \underbrace{(\alpha \rightarrow \alpha)}_{inc} \times \underbrace{(\alpha \rightarrow 1 + \alpha)}_{dec}$$

where the name annotations are just explanatory and not part of the syntax. Note that α stands for ctr which is bound here by the existential quantifier.

Now what should an expression

$$e : \exists\alpha. \alpha \times (\alpha \rightarrow \alpha) \times (\alpha \rightarrow 1 + \alpha)$$

look like? It should provide a concrete implementation type (such as nat or bin) for α , as well as an implementation of the three functions. We obtain this with the following rule

$$\frac{\Gamma \vdash \sigma \text{ type} \quad \Gamma \vdash e : [\sigma/\alpha]\tau}{\Gamma \vdash \langle \sigma, e \rangle : \exists\alpha. \tau} \text{tp/exists}$$

Besides checking that σ is indeed a type with respect to all the type variables declared in Γ , the crucial aspect of this rule is that the implementation e is at type $[\sigma/\alpha]\tau$.

For example, to check that new , inc , and dec are well-typed we substitute the implementation type for α (namely nat in one case and bin in the other case) before we proceed to check the definitions.

The pair $\langle \sigma, e \rangle$ is sometimes referred to as a *package*, which is opened up by the destructor. This destructor is often called *open*, but for uniformity with all analogous cases, and to support general pattern matching, we'll write it as a case.

$$\begin{array}{ll} \text{Types} & \tau ::= \dots \mid \exists\alpha. \tau \\ \text{Expressions} & e ::= \dots \mid \langle \sigma, e \rangle \mid \text{case } e \langle \langle \alpha, x \rangle \Rightarrow e' \end{array}$$

The elimination form provides a new name α for the implementation types and a new variable x for the (eager) pair making up the implementations.

$$\frac{\begin{array}{c} (\alpha \notin \text{dom}(\Gamma) \cup \text{FTV}(\Gamma) \cup \text{FTV}(\tau')) \\ \Gamma \vdash e : \exists \alpha. \tau \quad \Gamma, \alpha \text{ type}, x : \tau \vdash e' : \tau' \end{array}}{\Gamma \vdash \text{case } e (\langle \alpha, x \rangle \Rightarrow e') : \tau'} \text{tp/casee}$$

The fact that the type α must be *new* is explicit here in the conditions that it does not already appear in Γ or τ' . Such a condition is often left implicit, relying on the well-formedness presuppositions of the judgments. For example, the presupposition that Γ may not contain any repeated variables means that if we happened to have used the name α before then we can just rename it and then apply the rule. It is crucial for data abstraction that this variable α is new because we cannot and should not be able to assume anything about what α might stand for, except the operations that are exposed in τ and are accessible via the name x . Among other things, α may not appear in τ' .

To be a little more explicit about this (because it is critical here), whenever we write $\Gamma \vdash e : \tau$ we make the following *presuppositions*:

1. All the variables and type variables in Γ are distinct.
2. $\Gamma \text{ ctx}$
3. $\Gamma \vdash \tau \text{ type}$

where the validity of context is defined by the following rules:

$\frac{}{(\cdot) \text{ ctx}} \text{ ctx/emp}$	$\frac{\Gamma \text{ ctx}}{(\Gamma, \alpha \text{ type}) \text{ ctx}} \text{ ctx/tpvar}$	$\frac{\Gamma \text{ ctx} \quad \Gamma \vdash \tau \text{ type}}{(\Gamma, x : \tau) \text{ ctx}} \text{ ctx/var}$
--	--	---

With these presuppositions the condition on α in the tp/casee rule is automatically satisfied. Whenever we write a rule we assume this presuppositions holds for the conclusion and we have to make sure they hold for all the premises. Let's look at case/exists again in this light.

1. We assume all variables in Γ are distinct, which also means they are distinct in the first premise. In the second premise they are distinct because that's how we interpret $\Gamma, \alpha \text{ type}, x : \tau$, which may include an implicit renaming of the type variable α or the variable x bound in the expression $\langle \alpha, x \rangle \Rightarrow e'$.

2. By presupposition (from the conclusion), $\Gamma \text{ ctx}$, which means that there are no *free* type variables in it, but variables declared in it can occur to their right. But what about τ ? Actually, it is okay (and in fact mostly needed) for α to appear in τ .
3. By presupposition (from the conclusion), $\Gamma \vdash \tau'$ *type*. This covers the second premise. Often, this rule is given with an explicit premise $\Gamma \vdash \tau'$ *type* to emphasize τ' must be independent of α . Indeed, the scope of α is the type of x and the expression e' .

We also see that the client e' is *parametric* in α , which means that it cannot depend on what α might actually be at runtime. It is this parametricity that will allow us to swap one implementation out for another without affecting the client as long as the two implementations are equivalent in an appropriate sense.

The dynamics is straightforward and not very interesting.

$$\frac{\frac{v \text{ value}}{\langle \sigma, v \rangle \text{ value}} \text{ val/exists} \quad \frac{e \mapsto e'}{\langle \sigma, e \rangle \mapsto \langle \sigma, e' \rangle} \text{ step/pack}_1}{\frac{e_0 \mapsto e'_0}{\text{case } e_0 (\langle \alpha, x \rangle \Rightarrow e_1) \mapsto \text{case } e'_0 (\langle \alpha, x \rangle \Rightarrow e_1)} \text{ step/casee}_0} \text{ step/casee/pack}$$

In a language with variadic sums and pattern matching, we would extend the language of patterns.

$$\text{Patterns } p ::= x \mid \langle p_1, p_2 \rangle \mid \langle \rangle \mid i \cdot p \mid \text{fold } p \mid \langle \alpha, p \rangle$$

The hypothetical open construct now corresponds a pattern match, with the scope of the openend module extending to the end of the case expression. For example, we can test an implementation of *CTR* by creating a fresh counter and then verifying that incrementing it followed by a decrement has a well-defined answer. In the definition of *test* we exploit general pattern matching so an exception is raised if a decrement of zero was attempted.

$$\begin{aligned} \text{test} &: \text{CTR} \rightarrow 1 \\ \text{test} &= \lambda c. \text{case } c (\langle \alpha, \langle \text{new}, \langle \text{inc}, \text{dec} \rangle \rangle \rangle \Rightarrow \\ &\quad \text{case } \text{dec } (\text{inc } \text{new}) (\mathbf{r} \cdot _ \Rightarrow \langle \rangle)) \end{aligned}$$

Note that the case opens the package (representing the module) and matches against its components so it can refer to them in the body of the function. The following two expressions will evaluate to unit (instead of raising an exception) if the implementation is correct (to the very limited extent that is tested here).

```
test NatCtr
test BinCtr
```

References

- [GMM⁺78] Michael J.C. Gordon, Robin Milner, L. Morris, Malcolm C. Newey, and Christopher P. Wadsworth. A metalanguage for interactive proof in LCF. In A. Aho, S. Zilles, and T. Szymanski, editors, *Conference Record of the 5th Annual Symposium on Principles of Programming Languages (POPL'78)*, pages 119–130, Tucson, Arizona, January 1978. ACM Press.
- [MP88] John C. Mitchell and Gordon D. Plotkin. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems*, 10(3):470–502, 1988.
- [Wad89] Philip Wadler. Theorems for free! In J. Stoy, editor, *Proceedings of the 4th International Conference on Functional Programming Languages and Computer Architecture (FPCA'89)*, pages 347–359, London, UK, September 1989. ACM.