

Lecture Notes 11

System F: Polymorphism and Abstraction

Carlo Angiuli

B522: PL Foundations

April 14, 2025

these notes are a bit short on explanations...

In this lecture, we will discuss System F [Gir72; Rey74], also known as the *polymorphic λ -calculus* or the *second-order λ -calculus*, a famous core calculus that goes beyond “simple” types by introducing universal and existential quantification to our syntax of types (or propositions, via Curry–Howard). These quantifiers extend the simply-typed λ -calculus with two important type-based abstraction mechanisms, namely *parametric polymorphism* and *abstract types*.

System F and abstract types are covered in Chapters 16 and 17 of Harper [Har16] respectively. Students may also wish to consult Sections 4 and 5 of Lau Skorstengaard’s notes from Amal Ahmed’s OPLSS lectures on logical relations. For those who wish to dig deeper, the seminal research papers on these topics are surprisingly accessible:

- “Types, Abstraction and Parametric Polymorphism” by Reynolds [Rey83]
- “Theorems for Free!” by Wadler [Wad89]
- “Representation Independence and Data Abstraction” by Mitchell [Mit86]
- “Abstract types have existential type” by Mitchell and Plotkin [MP88]

1 System F

In our lecture on type isomorphisms we talked about “*the* identity function I,” but the STLC actually has many different identity functions, one for every type:

$$\begin{aligned} \cdot &\vdash \lambda x : \text{unit}. x : \text{unit} \rightarrow \text{unit} \\ \cdot &\vdash \lambda x : \text{bool}. x : \text{bool} \rightarrow \text{bool} \\ &\vdots \end{aligned}$$

The STLC does not let us use a single function at all of these types; remember, it has uniqueness of types! System F also has uniqueness of types, but what it adds to the STLC is the ability to express that the identity function $\lambda x : \alpha. x$ has type $\alpha \rightarrow \alpha$ *generically for any type α* , by giving it the *polymorphic* type $\forall \alpha. \alpha \rightarrow \alpha$. This polymorphic type tells us that we can instantiate the α in the function’s type with any concrete type of our choice, such as `unit` or `bool`.

This is call-by-name System F:

Syntax:

<i>Types</i>	$\tau ::=$	<code>arr</code> (τ_1, τ_2)	$\tau_1 \rightarrow \tau_2$	function type
		<code>all</code> ($\alpha. \tau$)	$\forall \alpha. \tau$	polymorphic type
<i>Terms</i>	$e ::=$	<code>lambda</code> ($\tau, x. e$)	$\lambda x : \tau. e$	λ -abstraction
		<code>app</code> (e_1, e_2)	$e_1 e_2$	application
		<code>Lambda</code> ($\alpha. e$)	$\Lambda \alpha. e$	type abstraction
		<code>App</code> (e, τ)	$e @ \tau$	type application

The $\Delta \vdash \tau$ ty judgment:

$$\frac{}{\Delta, \alpha \text{ ty} \vdash \alpha \text{ ty}} \quad \frac{\Delta \vdash \tau_1 \text{ ty} \quad \Delta \vdash \tau_2 \text{ ty}}{\Delta \vdash \tau_1 \rightarrow \tau_2 \text{ ty}} \quad \frac{\Delta, \alpha \text{ ty} \vdash \tau \text{ ty}}{\Delta \vdash \forall \alpha. \tau \text{ ty}}$$

Since terms can now contain type variables, the typing judgment must be indexed by both a type variable context Δ and a term variable context Γ , where the types in Γ must be well-formed with respect to the context Δ .

The $\Delta; \Gamma \vdash e : \tau$ judgment:

$$\frac{}{\Delta; \Gamma, x : \tau \vdash x : \tau} \text{VAR} \quad \frac{\Delta \vdash \tau_1 \text{ ty} \quad \Delta; \Gamma, x : \tau_1 \vdash e_2 : \tau_2}{\Delta; \Gamma \vdash \lambda x : \tau_1. e_2 : \tau_1 \rightarrow \tau_2} \rightarrow\text{-INTRO}$$

$$\frac{\Delta; \Gamma \vdash f : \tau_1 \rightarrow \tau_2 \quad \Delta; \Gamma \vdash e_1 : \tau_1}{\Delta; \Gamma \vdash f e_1 : \tau_2} \rightarrow\text{-ELIM}$$

$$\frac{\Delta, \alpha \text{ ty}; \Gamma \vdash e : \tau}{\Delta; \Gamma \vdash \Lambda \alpha. e : \forall \alpha. \tau} \forall\text{-INTRO} \quad \frac{\Delta; \Gamma \vdash e : \forall \alpha. \tau \quad \Delta \vdash \tau' \text{ ty}}{\Delta; \Gamma \vdash e @ \tau' : \tau[\tau'/\alpha]} \forall\text{-ELIM}$$

Remark 11.1. Our conventions for judgments hide an implicit side condition of the \forall -INTRO rule: it can only be applied when the context Γ is well-formed in Δ , i.e., when α does not occur free in Γ .

The v val judgment:

$$\frac{}{\lambda x : \tau. e \text{ val}} \qquad \frac{}{\Lambda \alpha. e \text{ val}}$$

The $e \mapsto e'$ judgment:

$$\frac{f \mapsto f'}{f e_1 \mapsto f' e_1} \qquad \frac{}{(\lambda x : \tau_1. e_2) e_1 \mapsto e_2[e_1/x]}$$

$$\frac{e \mapsto e'}{e @ \tau \mapsto e' @ \tau} \qquad \frac{}{(\Lambda \alpha. e) @ \tau \mapsto e[\tau/\alpha]}$$

Example 11.2. Returning to the polymorphic identity function, we define:

$$\mathbf{I} : \forall \alpha. \alpha \rightarrow \alpha$$

$$\mathbf{I} := \Lambda \alpha. \lambda x : \alpha. x$$

Then $\mathbf{I}@unit : unit \rightarrow unit$ and $\mathbf{I}@bool : bool \rightarrow bool$ and even

$$\mathbf{I} @ (\forall \alpha. \alpha \rightarrow \alpha) : (\forall \alpha. \alpha \rightarrow \alpha) \rightarrow (\forall \alpha. \alpha \rightarrow \alpha)$$

The ability to use a single piece of code at multiple types is called *polymorphism*. System F supports a specific kind of polymorphism called *parametric polymorphism*, in which polymorphic terms must behave uniformly at every type. For example, a term of type $\forall \alpha. \alpha \rightarrow \alpha$ is not simply a term that happens to work at type $\tau \rightarrow \tau$ for any τ ; it has a single definition that is uniform or generic in the type α .

In contrast, *ad hoc polymorphism* refers to terms that are available at every type but may be defined differently at each one. For example, many languages have a $+$ function which can operate on any type, but given two numbers performs numerical addition, given two strings performs string concatenation, etc. Similarly, `equal?` compares numbers for numerical equality, lists for structural equality, functions for pointer equality, etc.

In object-oriented languages, mechanisms for parametric polymorphism are often called *generics* and ad hoc polymorphism is often called *operator overloading*.

Remark 11.3.

type application is weird, usually it's inferred

Remark 11.4. The notation \forall makes good sense from the type system perspective: a term with type $\forall\alpha.\alpha \rightarrow \alpha$ has type $\alpha \rightarrow \alpha$ “for all” types α . In addition, the typing rules extend the Curry–Howard correspondence to universal quantification. In logic, to prove (introduce) a \forall statement, one must construct a proof of that statement for an arbitrary (variable) term; to use (eliminate) a \forall statement, one may instantiate the quantified variable with any term whatsoever.

1.1 Church encodings

Note that System F has no base types, but the type grammar is not empty because of type variables.

We can extend System F with all the other types we have considered so far (in fact, roughly speaking, we can think of typed functional programming languages like OCaml as some combination of System F and PCF with isorecursive types.) But it turns out that we do not need them. (An aside.)

$$\begin{aligned} \text{bool} &:= \forall\alpha.\alpha \rightarrow \alpha \rightarrow \alpha \\ \text{true} &:= \Lambda\alpha.\lambda x : \alpha.\lambda y : \alpha.x \\ \text{false} &:= \Lambda\alpha.\lambda x : \alpha.\lambda y : \alpha.y \\ \text{if}(e_1, e_2, e_3) : \tau &:= e_1@_\tau e_2 e_3 \end{aligned}$$

$$\begin{aligned} \text{nat} &:= \forall\alpha.\alpha \rightarrow (\alpha \rightarrow \alpha) \rightarrow \alpha \\ \text{zero} &:= \Lambda\alpha.\lambda z : \alpha.\lambda s : \alpha \rightarrow \alpha.z \\ \text{suc}(e) &:= \Lambda\alpha.\lambda z : \alpha.\lambda s : \alpha \rightarrow \alpha.s (e@_\alpha z s) \\ \text{natrec}(e, e_z, e_s) : \tau &:= e@_\tau e_z e_s \end{aligned}$$

Exercise 11.5. Define $\text{not} : \text{bool} \rightarrow \text{bool}$.

These encodings also work in the untyped λ -calculus, but they don’t work in STLC because we would have to fix the type that we map out into.

1.2 Free theorems

See Wadler [Wad89]. Parametric polymorphism is very strong:

- There are no closed terms of type $\forall\alpha.\alpha$.
- The polymorphic identity function is the only term (up to observational equivalence) of type $\forall\alpha.\alpha \rightarrow \alpha$.
- There is no closed term of the “fixpoint” type $\forall\alpha.(\alpha \rightarrow \alpha) \rightarrow \alpha$.

- Any $\text{equal?} : \forall \alpha. \alpha \rightarrow \alpha \rightarrow \text{bool}$ must be constant.
- Any $\text{map} : \forall \alpha. \forall \beta. (\alpha \rightarrow \beta) \rightarrow \text{list}(\alpha) \rightarrow \text{list}(\beta)$ must satisfy the equations

$$\begin{aligned} & \text{map@}\tau\text{@}\tau' f \\ & \cong \lambda \ell : \text{list}(\tau). \text{map@}\tau\text{@}\tau' f (\text{map@}\tau\text{@}\tau (\mathbf{I@}\tau) \ell) \\ & \cong \lambda \ell : \text{list}(\tau). \text{map@}\tau'\text{@}\tau' (\mathbf{I@}\tau') (\text{map@}\tau\text{@}\tau' f \ell) \end{aligned}$$

2 Abstract types

An even more important application of type genericity arises when considering abstract interfaces, e.g. in data structures.

An implementation of a *queue* (of numbers) must provide:

- a representation type τ_r for queues,
- a queue $\text{empty} : \tau_r$,
- a function $\text{enqueue} : \text{nat} \rightarrow \tau_r \rightarrow \tau_r$, and
- a function $\text{dequeue} : \tau_r \rightarrow (\text{unit} + (\text{nat} \times \tau_r))$.

There are various ways to implement queues, including the simple `ListQueue` implementation in which $\tau_r = \text{list}(\text{nat})$, and the more efficient `BatchedQueues` in which $\tau_r = \text{list}(\text{nat}) \times \text{list}(\text{nat})$ [Oka99, Section 5.2].

There are many reasons why a program may wish to use a queue as part of some other computation. We would like to arrange that such programs, which we will call *clients* of the queue library, not only do not need to understand how queues are implemented but in fact are *prohibited* from knowing the implementation.

To this end we will define an abstract interface (here, a type and three terms) that all queue implementations will implement, and with respect to which all clients will be implemented. Our type discipline will then enforce that clients may not take the length of a queue even if we happen to link them against the `ListQueue` implementation. (Indeed, that would prevent us from swapping the `ListQueues` for `BatchedQueues`.)

Maintaining a strict separation of concerns between libraries and clients is crucial to programming in the large. Object-oriented languages call this separation *encapsulation*; non-OO programming language theorists typically call it *data abstraction*, and the type of queues an *abstract data type*.

“Type structure is a syntactic discipline for enforcing levels of abstraction.” —John C. Reynolds

Extend System F with existential types:

The $\Delta \vdash \tau$ ty judgment:

$$\dots \quad \frac{\Delta, \alpha \text{ ty} \vdash \tau \text{ ty}}{\Delta \vdash \exists \alpha. \tau \text{ ty}}$$

The $\Delta; \Gamma \vdash e : \tau$ judgment:

$$\dots \quad \frac{\Delta \vdash \tau_r \text{ ty} \quad \Delta, \alpha \text{ ty} \vdash \tau_i \text{ ty} \quad \Delta; \Gamma \vdash e : \tau_i[\tau_r/\alpha]}{\Delta; \Gamma \vdash \text{pack} \langle \tau_r, e \rangle \text{ as } \exists \alpha. \tau_i : \exists \alpha. \tau_i} \exists\text{-INTRO}$$

$$\frac{\Delta; \Gamma \vdash e : \exists \alpha. \tau_i \quad \Delta \vdash \tau' \text{ ty} \quad \Delta, \alpha \text{ ty}; \Gamma, x : \tau_i \vdash e' : \tau'}{\Delta; \Gamma \vdash \text{unpack} \langle \alpha, x \rangle = e \text{ in } e' : \tau'} \exists\text{-ELIM}$$

The v val judgment:

$$\dots \quad \frac{}{\text{pack} \langle \tau_r, e \rangle \text{ as } \exists \alpha. \tau_i \text{ val}}$$

The $e \mapsto e'$ judgment:

$$\frac{e \mapsto e''}{\text{unpack} \langle \alpha, x \rangle = e \text{ in } e' \mapsto \text{unpack} \langle \alpha, x \rangle = e'' \text{ in } e'}$$

$$\frac{}{\text{unpack} \langle \alpha, x \rangle = (\text{pack} \langle \tau_r, e \rangle \text{ as } \exists \alpha. \tau_i) \text{ in } e' \mapsto e'[\tau_r/\alpha][e/x]}$$

Remark 11.6. In fact it is possible to Church encode existential types, so everything we say about System F with existential types actually applies directly to plain System F. However, it will be much clearer for us to work with existential types without going through an encoding.

Remark 11.7. The notation \exists may seem somewhat odd at first, although it makes some sense: to construct a term of type $\exists \alpha. \alpha \times \dots$ there must “exist” some representation type τ_r for which we can construct a term of type $\tau_r \times \dots$. In fact this notation is chosen because the typing rules for abstract types extend the Curry–Howard correspondence to existential quantification. In logic, to prove (introduce) an \exists statement, one must exhibit a particular witness along with a proof of the statement for that witness. To use (eliminate) an \exists statement, one can assume that a *generic* witness satisfying the relevant property exists.

Returning to the queue example,

$$\begin{aligned} \text{QueueImpl} &:= \exists \alpha. \alpha \times (\text{nat} \rightarrow \alpha \rightarrow \alpha) \times (\alpha \rightarrow \text{unit} + \text{nat} \times \alpha) \\ \text{ListQueue} &:= \text{pack } \langle \text{list}(\text{nat}), (\text{nil}, [\text{enq}], [\text{deq}]) \rangle \text{ as } \exists \alpha. \alpha \times \dots \\ \text{BatchedQueue} &:= \text{pack } \langle \text{list}(\text{nat}) \times \text{list}(\text{nat}), \dots \rangle \text{ as } \exists \alpha. \alpha \times \dots \end{aligned}$$

Client:

$$\text{unpack } \langle \tau_q, \text{impl} \rangle = \text{ListQueue in } \dots \text{fst}(\text{impl}) \dots$$

The type of the unpack has to not contain τ_q , and the client code $\dots \text{fst}(\text{impl}) \dots$ has to be parametrically polymorphic in τ_q . This is how we ensure data abstraction. Note however that at runtime there is no longer any data abstraction.

2.1 Representation independence

What happens when we replace `ListQueue` with `BatchedQueue`?

$$\text{unpack } \langle \tau_q, \text{impl} \rangle = \text{BatchedQueue in } \dots \text{fst}(\text{impl}) \dots$$

This program is still well-typed, so it does not go wrong. But does it compute the same thing as the previous program?

In general, two implementations of an interface might behave completely differently. For example, the new queue’s dequeue function may always return `inl(())`. In this case, we might imagine that despite having different runtime representations of queues under the hood, `ListQueues` and `BatchedQueues` “have the same extensional behavior,” and thus expect that the two programs should compute the same result—even though they run different code!

One way to formalize the notion of “having the same extensional behavior” is by exhibiting what is known as a *bisimulation* between the two queue implementations. Given two implementations

$$\begin{aligned} \text{QueueImpl}_1 &:= \text{pack } \langle \tau_1, (\text{emp}_1, \text{enq}_1, \text{deq}_1) \rangle \text{ as } \exists \alpha. \alpha \times \dots \\ \text{QueueImpl}_2 &:= \text{pack } \langle \tau_2, (\text{emp}_2, \text{enq}_2, \text{deq}_2) \rangle \text{ as } \exists \alpha. \alpha \times \dots \end{aligned}$$

a *queue bisimulation* between `QueueImpl1` and `QueueImpl2` is a binary relation R between closed terms of type τ_1 and closed terms of type τ_2 , such that

- $R(\text{emp}_1, \text{emp}_2)$ holds,
- for all $n : \text{nat}$, $q_1 : \tau_1$, and $q_2 : \tau_2$ satisfying $R(q_1, q_2)$, $R(\text{enq}_1 \ n \ q_1, \text{enq}_2 \ n \ q_2)$ holds, and

- for all $q_1 : \tau_1$ and $q_2 : \tau_2$ satisfying $R(q_1, q_2)$, either
 - $\text{deq } q_1 \cong \text{inl}()$ and $\text{deq } q_2 \cong \text{inl}()$, or
 - $\text{deq } q_1 \cong \text{inr}(n_1, q'_1)$ and $\text{deq } q_2 \cong \text{inr}(n_2, q'_2)$ where $n_1 \cong n_2$ and $R(q'_1, q'_2)$.

The *representation independence* theorem [Mit86] states that if we have a queue bisimulation between QueueImpl_1 and QueueImpl_2 , then QueueImpl_1 and QueueImpl_2 are observationally equivalent at type QueueImpl . That is, no program’s result can be affected by swapping QueueImpl_1 for QueueImpl_2 . In particular, every client computes the same answers with respect to both implementations.

We emphasize that this theorem holds *with no conditions whatsoever* on the client code: our type system guarantees that no program can tell apart bisimilar terms of existential type.

In the next lecture we will use logical relations to prove the *parametricity theorem* for System F, a powerful result from which we can obtain all of the free theorems and representation independence theorems discussed in this lecture.

References

- [Gir72] Jean-Yves Girard. “Interprétation fonctionnelle et élimination des coupures de l’arithmétique d’ordre supérieur”. PhD thesis. Université Paris VII, June 1972. URL: <https://girard.perso.math.cnrs.fr/These.pdf>.
- [Har16] Robert Harper. *Practical Foundations for Programming Languages*. Second Edition. Cambridge University Press, 2016. ISBN: 9781107150300. DOI: [10.1017/CB09781316576892](https://doi.org/10.1017/CB09781316576892).
- [Mit86] John C. Mitchell. “Representation Independence and Data Abstraction”. In: *Proceedings of the 13th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. POPL ’86. New York, NY, USA: ACM, 1986, pp. 263–276. ISBN: 9781450373470. DOI: [10.1145/512644.512669](https://doi.org/10.1145/512644.512669).
- [MP88] John C. Mitchell and Gordon D. Plotkin. “Abstract types have existential type”. In: *ACM Transactions on Programming Languages and Systems* 10.3 (July 1988), pp. 470–502. ISSN: 0164-0925. DOI: [10.1145/44501.45065](https://doi.org/10.1145/44501.45065).
- [Oka99] Chris Okasaki. *Purely functional data structures*. Cambridge University Press, 1999. DOI: [10.1017/CB09780511530104](https://doi.org/10.1017/CB09780511530104).

- [Rey74] John C. Reynolds. “Towards a theory of type structure”. In: *Programming Symposium*. Ed. by B. Robinet. Berlin, Heidelberg: Springer Berlin Heidelberg, 1974, pp. 408–425. ISBN: 978-3-540-37819-8.
- [Rey83] John C. Reynolds. “Types, Abstraction and Parametric Polymorphism”. In: *Information Processing '83: Proceedings of the IFIP 9th World Computer Congress*. Ed. by R. E. A. Mason. North-Holland, 1983, pp. 513–523. ISBN: 0444867295.
- [Wad89] Philip Wadler. “Theorems for Free!” In: *Proceedings of the Fourth International Conference on Functional Programming Languages and Computer Architecture*. FPCA '89. New York, NY, USA: ACM, 1989, pp. 347–359. ISBN: 0897913280. DOI: [10.1145/99370.99404](https://doi.org/10.1145/99370.99404).