MPRI 2.4, Functional programming and type systems
Metatheory of System F

Didier Rémy

September 29, 2021

Plan of the course

Metatheory of System F

ADTs, Recursive types, Existential types, GATDs

Going higher order with F^ω!

Logical relations

Abstract Data types, Existential types, GADTs

Contents

- Application to typed closure conversion

  - Environment passing

  - Closure passing

Type-preserving compilation

Compilation is type-preserving when each intermediate language is *explicitly typed*, and each compilation phase transforms a typed program into a typed program in the next intermediate language.

Why *preserve types* during compilation?

- it can help debug the compiler;
- types can be used to drive optimizations;
- types can be used to produce *proof-carrying code;*
- proving that types are preserved can be the first step towards proving that the *semantics* is preserved [Chlipala, 2007].

Type-preserving compilation

Type-preserving compilation exhibits an encoding of programming constructs into programming languages with usually richer type systems.

The encoding may sometimes be used directly as a programming idiom in the source language.

For example:

- Closure conversion requires an extension of the language with existential types, which happens to be very useful on their own.
- Closures are themselves a simple form of objects, which can also be explained with existential types.
- Defunctionalization may be done manually on some particular programs, *e.g.* in web applications to monitor the computation.

Type-preserving compilation

A classic paper by Morrisett *et al.* [1999] shows how to go from System F to Typed Assembly Language, while preserving types along the way. Its main passes are:

- *CPS conversion* fixes the order of evaluation, names intermediate computations, and makes all function calls tail calls;
- *closure conversion* makes environments and closures explicit, and produces a program where all functions are closed;
- allocation and initialization of tuples is made explicit;
- the calling convention is made explicit, and variables are replaced with (an unbounded number of) machine registers.

Translating types

In general, a type-preserving compilation phase involves not only a translation of *terms,* mapping M to $[\![M]\!]$, but also a translation of *types,* mapping τ to $[\![\tau]\!]$, with the property:

$$\Gamma \vdash M : \tau \quad \text{implies} \quad [\![\Gamma]\!] \vdash [\![M]\!] : [\![\tau]\!]$$

The translation of types carries a lot of information: examining it is often enough to guess what the translation of terms will be.

See the old lecture on type closure conversion.

Closure conversion

First-class functions may appear in the body of other functions. hence, their own body may contain free variables that will be bound to values during the evaluation in the execution environment.

Because they can be returned as values, and thus used outside of their definition environment, they must store their execution environment in their value.

A *closure* is the packaging of the code of a first-class function with its runtime environment, so that it becomes closed, *i.e.* independent of the runtime environment and can be moved and applied in another runtime environment.

Closures can also be used to represent recursive functions and objects (in the object-as-record-of-methods paradigm).

Source and target

In the following,

- the *source* calculus has *unary* λ-abstractions, which can have free variables;

- the *target* calculus has *binary* λ-abstractions, which must be *closed.*

Closure conversion can be easily extended to n-ary functions, or n-ary functions may be *uncurried* in a separate, type-preserving compilation pass.

Variants of closure conversion

There are at least two variants of closure conversion:

- in the *closure-passing variant,*
  the closure and the environment are a single memory block;

- in the *environment-passing variant,*
  the environment is a separate block, to which the closure points.

The impact of this choice on the translation of terms is minor.

Its impact on the translation of types is more important:
the closure-passing variant requires more type-theoretic machinery.

Closure-passing closure conversion

Let $\{x_1, \ldots, x_n\}$ be $\mathrm{fv}(\lambda x. a)$:

$$
\begin{aligned}
[\![\lambda x. a]\!] &= \text{ let } code = \lambda(clo, x). \\
&\qquad \text{ let } (\_, x_1, \ldots, x_n) = clo \text{ in } [\![a]\!] \text{ in} \\
&\quad (code, x_1, \ldots, x_n) \\[1em]
[\![a_1\ a_2]\!] &= \text{ let } clo = [\![a_1]\!] \text{ in} \\
&\quad \text{ let } code = \mathsf{proj}_0\ clo \text{ in} \\
&\quad code\ (clo, [\![a_2]\!])
\end{aligned}
$$

(The variables $code$ and clo must be suitably fresh.)

Important! The layout of the environment must be known only at the closure allocation site, not at the call site. In particular, $\mathsf{proj}_0\ clo$ need not know the size of clo.

Environment-passing closure conversion

Let $\{x_1, \ldots, x_n\}$ be $\mathrm{fv}(\lambda x.\, a)$:

$$
\begin{aligned}
[\![\lambda x.\, a]\!] \;&=\; \textit{let code} = \lambda(\mathit{env}, x).\\
&\qquad \textit{let } (x_1, \ldots, x_n) = \mathit{env} \textit{ in } [\![a]\!] \textit{ in}\\
&\qquad (\mathit{code}, (x_1, \ldots, x_n))\\[6pt]
[\![a_1\; a_2]\!] \;&=\; \textit{let } (\mathit{code}, \mathit{env}) = [\![a_1]\!] \textit{ in}\\[6pt]
&\qquad \mathit{code}\; (\mathit{env}, [\![a_2]\!])
\end{aligned}
$$

Questions: How can closure conversion be made *type-preserving?*

The key issue is to find a sensible definition of the type translation. In particular, what is the translation of a function type, $[\![\tau_1 \to \tau_2]\!]$?

Environment-passing closure conversion

Let $\{x_1, \ldots, x_n\}$ be $\mathrm{fv}(\lambda x.\, a)$:

$$
\begin{aligned}
[\![\lambda x.\, a]\!] \;=\; &\textsf{let } code = \lambda(env, x). \\
&\quad \textsf{let } (x_1, \ldots, x_n) = env \textsf{ in } [\![a]\!] \textsf{ in} \\
&(code, (x_1, \ldots, x_n))
\end{aligned}
$$

Assume $\Gamma \vdash \lambda x.\, a : \tau_1 \to \tau_2$.

Assume, *w.l.o.g.*. $\mathrm{dom}(\Gamma) = \mathrm{fv}(\lambda x.\, a) = \{x_1, \ldots, x_n\}$.

Write $[\![\Gamma]\!]$ for the tuple type $x_1 : [\![\tau_1']\!]; \ldots; x_n : [\![\tau_n']\!]$ where Γ is $x_1 : \tau_1'; \ldots; x_n : \tau_n'$. We also use $[\![\Gamma]\!]$ as a type to mean $[\![\tau_1']\!] \times \ldots \times [\![\tau_n']\!]$.

We have $\Gamma, x : \tau_1 \vdash a : \tau_2$, so in environment $[\![\Gamma]\!], x : [\![\tau_1]\!]$, we have

- env has type $[\![\Gamma]\!]$,
- $code$ has type $([\![\Gamma]\!] \times [\![\tau_1]\!]) \to [\![\tau_2]\!]$, and
- the entire closure has type $(([\![\Gamma]\!] \times [\![\tau_1]\!]) \to [\![\tau_2]\!]) \times [\![\Gamma]\!]$.

Now, *what should be the definition of $[\![\tau_1 \to \tau_2]\!]$?*

Towards a type translation

Can we adopt this as a definition?

$$\llbracket \tau_1 \to \tau_2 \rrbracket \quad = \quad ((\llbracket \Gamma \rrbracket \times \llbracket \tau_1 \rrbracket) \to \llbracket \tau_2 \rrbracket) \times \llbracket \Gamma \rrbracket$$

Naturally not. This definition is mathematically ill-formed: we cannot use Γ out of the blue.

That is, this definition is not uniform: it depends on Γ, *i.e.* the size and layout of the environment.

Do we really need to have a uniform translation of types?

Towards a type translation

Yes, we do.

We need a uniform translation of types, not just because it is nice to have one, but because it describes a *uniform calling convention.*

If closures with distinct environment sizes or layouts receive distinct types, then we will be unable to translate this well-typed code:

$$if \ldots \text{ then } \lambda x.\, x + y \text{ else } \lambda x.\, x$$

Furthermore, we want function invocations to be translated uniformly, without knowledge of the size and layout of the closure's environment.

So, *what could be the definition of $[\![\tau_1 \rightarrow \tau_2]\!]$?*

The type translation

The only sensible solution is:

$$\llbracket \tau_1 \rightarrow \tau_2 \rrbracket \;\; = \;\; \exists \alpha.((\alpha \times \llbracket \tau_1 \rrbracket) \rightarrow \llbracket \tau_2 \rrbracket) \times \alpha$$

An *existential quantification* over the type of the environment abstracts away the differences in size and layout.

Enough information is retained to ensure that the application of the code to the environment is valid: this is expressed by letting the variable α occur twice on the right-hand side.

The type translation

The existential quantification also provides a form of *security*: the caller cannot do anything with the environment except pass it as an argument to the code; in particular, it cannot inspect or modify the environment.

For instance, in the source language, the following coding style guarantees that x remains even, no matter how f is used:

$$\text{let } f = \text{let } x = \text{ref } 0 \text{ in } \lambda().\, x := (x + 2); !\, x$$

After closure conversion, the reference x is reachable via the closure of f.
A malicious, untyped client could write an odd value to x.
However, a *well-typed* client is unable to do so.

This encoding is not just type-preserving, but also *fully abstract:* it preserves (a typed version of) observational equivalence [Ahmed and Blume, 2008].

- Application to typed closure conversion
  - Environment passing
  - Closure passing

Typed closure conversion

Everything is now set up to prove that, in System F with existential types:

$$\Gamma \vdash M : \tau \quad \text{implies} \quad [\![\Gamma]\!] \vdash [\![M]\!] : [\![\tau]\!]$$

Environment-passing closure conversion

Assume $\Gamma \vdash \lambda x.\, M : \tau_1 \to \tau_2$ and $\mathrm{dom}(\Gamma) = \{x_1, \ldots, x_n\} = \mathrm{fv}(\lambda x.\, M)$.

$$
\begin{aligned}
[\![\lambda x{:}\tau_1.\, M]\!] \quad = \quad & \textit{let code} : ([\![\Gamma]\!] \times [\![\tau_1]\!]) \to [\![\tau_2]\!] = \\
& \quad \lambda(\textit{env} : [\![\Gamma]\!], x : [\![\tau_1]\!]). \\
& \qquad \textit{let } (x_1, \ldots, x_n : [\![\Gamma]\!]) = \textit{env in} \\
& \qquad [\![M]\!] \\
& \textit{in} \\
& \textit{pack } [\![\Gamma]\!], (\textit{code}, (x_1, \ldots, x_n)) \\
& \textit{as } \exists \alpha.((\alpha \times [\![\tau_1]\!]) \to [\![\tau_2]\!]) \times \alpha
\end{aligned}
$$

We find $[\![\Gamma]\!] \vdash [\![\lambda x{:}\tau_1.\, M]\!] : [\![\tau_1 \to \tau_2]\!]$, as desired.

Environment-passing closure conversion

Assume $\Gamma \vdash M : \tau_1 \to \tau_2$ and $\Gamma \vdash M_1 : \tau_1$.

$$
\begin{aligned}
[\![M\ M_1]\!] \ =\ & \text{let } \alpha, (\,code : (\alpha \times [\![\tau_1]\!]) \to [\![\tau_2]\!], env : \alpha) = \\
& \quad \text{unpack } [\![M]\!] \text{ in} \\
& code\,(\,env, [\![M_1]\!])
\end{aligned}
$$

We find $[\![\Gamma]\!] \vdash [\![M\ M_1]\!] : [\![\tau_2]\!]$, as desired.

Environment-passing closure conversion   recursion

Recursive functions can be translated in this way, known as the "fix-code" variant [Morrisett and Harper, 1998] (leaving out type information):

$$\llbracket \mu f.\lambda x.M \rrbracket = \text{let rec } code \, (env, x) =$$
$$\text{let } f = pack \, (code, env) \text{ in}$$
$$\text{let } (x_1, \ldots, x_n) = env \text{ in}$$
$$\llbracket M \rrbracket \text{ in}$$
$$pack \, (code, (x_1, \ldots, x_n))$$

where $\{x_1, \ldots, x_n\} = \text{fv}(\mu f.\lambda x.M)$.

The translation of applications is unchanged: recursive and non-recursive functions have an identical calling convention.

What is the weak point of this variant?

A new closure is allocated at every call.

Environment-passing closure conversion   recursion

Instead, the "fix-pack" variant [Morrisett and Harper, 1998] uses an extra field in the environment to store a back pointer to the closure:

$$
\begin{aligned}
\llbracket \mu f.\lambda x.M \rrbracket \;=\; & \text{let } code\,(env, x) = \\
& \quad \text{let } (f, x_1, \ldots, x_n) = env \text{ in} \\
& \quad \llbracket M \rrbracket \\
& \text{in} \\
& \text{let rec } clo = (\,code, (\,clo, x_1, \ldots, x_n)\,) \text{ in} \\
& clo
\end{aligned}
$$

where $\{x_1, \ldots, x_n\} = \mathrm{fv}(\mu f.\lambda x.M)$.

This requires general, recursively-defined *values.* Closures are now *cyclic* data structures.

Environment-passing closure conversion   recursion

Here is how the "fix-pack" variant is type-checked. Assume
$\Gamma \vdash \mu f.\lambda x.M : \tau_1 \to \tau_2$ and $\mathrm{dom}(\Gamma) = \{x_1, \ldots, x_n\} = \mathrm{fv}(\mu f.\lambda x.M)$.

$$\llbracket \mu f : \tau_1 \to \tau_2.\lambda x.M \rrbracket =$$
$$\textit{let code} : (\llbracket f : \tau_1 \to \tau_2; \Gamma \rrbracket \times \llbracket \tau_1 \rrbracket) \to \llbracket \tau_2 \rrbracket =$$
$$\lambda(\textit{env} : \llbracket f : \tau_1 \to \tau_2, \Gamma \rrbracket, x : \llbracket \tau_1 \rrbracket).$$
$$\textit{let } (f, x_1, \ldots, x_n) : \llbracket f : \tau_1 \to \tau_2, \Gamma \rrbracket = \textit{env in}$$
$$\llbracket M \rrbracket \textit{ in}$$
$$\textit{let rec clo} : \llbracket \tau_1 \to \tau_2 \rrbracket =$$
$$\textit{pack } \llbracket f : \tau_1 \to \tau_2, \Gamma \rrbracket, (\textit{code}, (\textit{clo}, x_1, \ldots, x_n))$$
$$\textit{as } \exists \alpha((\alpha \times \llbracket \tau_1 \rrbracket) \to \llbracket \tau_2 \rrbracket) \times \alpha)$$
$$\textit{in clo}$$

Problem?

Environment-passing closure conversion   recursion

The recursive function may be polymorphic, but recursive calls are monomorphic...

We can generalize the encoding afterwards,

$$[\![\Lambda\vec{\beta}.\,\mu f : \tau_1 \to \tau_2.\lambda x.M]\!] = \Lambda\vec{\beta}.\,[\![\mu f : \tau_1 \to \tau_2.\lambda x.M]\!]$$

whenever the right-hand side is well-defined.

This allows the *indirect* compilation of polymorphic recursive functions as long as the recursion is monomorphic.

Fortunately, the encoding can be straightforwardly adapted to *directly* compile polymorphically recursive functions into polymorphic closure.

Environment-passing closure conversion   recursion

$$\llbracket \mu f : \forall \vec{\beta}. \tau_1 \to \tau_2. \lambda x.M \rrbracket =$$

> *let code* $: \forall \vec{\beta}. (\llbracket f : \forall \vec{\beta}. \tau_1 \to \tau_2; \Gamma \rrbracket \times \llbracket \tau_1 \rrbracket) \to \llbracket \tau_2 \rrbracket =$
>> $\lambda(env : \llbracket f : \forall \vec{\beta}. \tau_1 \to \tau_2, \Gamma \rrbracket, x : \llbracket \tau_1 \rrbracket).$
>>> *let* $(f, x_1, \ldots, x_n) : \llbracket f : \forall \vec{\beta}. \tau_1 \to \tau_2, \Gamma \rrbracket = env$ *in*
>>> $\llbracket M \rrbracket$ *in*
>
> *let rec clo* $: \llbracket \forall \vec{\beta}. \tau_1 \to \tau_2 \rrbracket =$
>> $\Lambda \vec{\beta}.$ *pack* $\llbracket f : \forall \vec{\beta}. \tau_1 \to \tau_2, \Gamma \rrbracket, (code \ \vec{\beta}, (clo, x_1, \ldots, x_n))$
>>> *as* $\exists \alpha ((\alpha \times \llbracket \tau_1 \rrbracket) \to \llbracket \tau_2 \rrbracket) \times \alpha)$
>
> *in clo*

The encoding is simple.

However, this requires the introduction of recursive non-functional values "let rec x = v". While this is a useful construct, it really alters the operational semantics and requires updating the type soundness proof.

- Application to typed closure conversion
  - Environment passing
  - Closure passing

Closure-passing closure conversion

$$\llbracket \lambda x.\, M \rrbracket = \begin{aligned}[t] &\textit{let } code = \lambda(clo, x).\\ &\quad \textit{let } (\_, x_1, \ldots, x_n) = clo \textit{ in}\\ &\quad \llbracket M \rrbracket\\ &\textit{in } (code, x_1, \ldots, x_n) \end{aligned}$$

$$\llbracket M_1\ M_2 \rrbracket = \begin{aligned}[t] &\textit{let } clo = \llbracket M_1 \rrbracket \textit{ in}\\ &\textit{let } code = \textsf{proj}_0\ clo \textit{ in}\\ &code\ (clo, \llbracket M_2 \rrbracket) \end{aligned}$$

There are two difficulties:

- a closure is a tuple, whose *first* field should be *exposed* (it is the code pointer), while the number and types of the remaining fields should be abstract;

- the first field of the closure contains a function that expects *the closure itself* as its first argument.

Closure-passing closure conversion

There are two difficulties:

- a closure is a tuple, whose *first* field should be *exposed* (it is the code pointer), while the number and types of the remaining fields should be abstract;

- the first field of the closure contains a function that expects *the closure itself* as its first argument.

What type-theoretic mechanisms could we use to describe this?

- existential quantification over the *tail* of a tuple (a.k.a. a *row*);

- *recursive types.*

Tuples, rows, row variables

The standard tuple types that we have used so far are:

$$\begin{array}{llll} \tau & ::= & \dots \mid \Pi\,R & \text{– types} \\ R & ::= & \epsilon \mid (\tau; R) & \text{– rows} \end{array}$$

The notation $(\tau_1 \times \dots \times \tau_n)$ was sugar for $\Pi\,(\tau_1; \dots; \tau_n; \epsilon)$.

Let us now introduce *row variables* and allow *quantification* over them:

$$\begin{array}{llll} \tau & ::= & \dots \mid \Pi\,R \mid \forall \rho.\,\tau \mid \exists \rho.\tau & \text{– types} \\ R & ::= & \rho \mid \epsilon \mid (\tau; R) & \text{– rows} \end{array}$$

This allows reasoning about the first few fields of a tuple whose length is not known.

Typing rules for tuples

The typing rules for tuple construction and deconstruction are:

$$\text{TUPLE} \qquad \frac{\forall i. \in [1,n] \quad \Gamma \vdash M_i : \tau_i}{\Gamma \vdash (M_1, \ldots, M_n) : \Pi\,(\tau_1; \ldots; \tau_n; \epsilon)}$$

$$\text{PROJ} \qquad \frac{\Gamma \vdash M : \Pi\,(\tau_1; \ldots; \tau_i; R)}{\Gamma \vdash \textit{proj}_i\, M : \tau_i}$$

These rules make sense with or without row variables

Projection does not care about the fields beyond i. Thanks to row variables, this can be expressed in terms of *parametric polymorphism:*

$$\textit{proj}_i : \forall \alpha._1 \ldots \alpha_i \rho.\ \Pi\,(\alpha_1; \ldots; \alpha_i; \rho) \to \alpha_i$$

About Rows

Rows were invented by Wand and improved by Rémy in order to ascribe precise types to operations on *records.*

The case of tuples, presented here, is simpler.

Rows are used to describe *objects* in Objective Caml [Rémy and Vouillon, 1998].

Rows are explained in depth by Pottier and Rémy [Pottier and Rémy, 2005].

Closure-passing closure conversion

Rows and recursive types allow to define the translation of types in the closure-passing variant:

$$
\begin{aligned}
&\quad [\![\tau_1 \to \tau_2]\!] \\
=\ &\exists \rho. && \rho \text{ describes the environment} \\
&\quad \mu\alpha. && \alpha \text{ is the concrete type of the closure} \\
&\qquad \Pi\ (&& \text{a tuple...} \\
&\qquad\quad (\alpha \times [\![\tau_1]\!]) \to [\![\tau_2]\!]; && \text{...that begins with a code pointer...} \\
&\qquad\quad \rho && \text{...and continues with the environment} \\
&\qquad)
\end{aligned}
$$

See Morrisett and Harper's "fix-type" encoding [1998].

Question: Why is it $\exists \rho.\ \mu\alpha.\ \tau$ and not $\mu\alpha.\ \exists \rho.\ \tau$

The type of the environment is fixed once for all and does not change at each recursive call.

Question: Notice that ρ appears only once. Any comments?

Closure-passing closure conversion

Let $Clo(R)$ abbreviate $\mu\alpha.\Pi\left((\alpha \times [\![\tau_1]\!]) \to [\![\tau_2]\!]; R\right)$.

Let $UClo(R)$ abbreviate its unfolded version,
$\Pi\left((Clo(R) \times [\![\tau_1]\!]) \to [\![\tau_2]\!]; R\right)$.

We have $[\![\tau_1 \to \tau_2]\!] = \exists\rho.Clo(\rho)$.

$$
\begin{aligned}
[\![\lambda x\!:\![\![\tau_1]\!].\, M]\!] \;=\; & \text{let } code : (Clo([\![\Gamma]\!]) \times [\![\tau_1]\!]) \to [\![\tau_2]\!] = \\
& \quad \lambda(\,clo : Clo([\![\Gamma]\!], x : [\![\tau_1]\!]). \\
& \qquad \text{let } (\_, x_1, \ldots, x_n) : UClo[\![\Gamma]\!] = \text{unfold } clo \text{ in} \\
& \qquad [\![M]\!] \text{ in} \\
& \quad \text{pack } [\![\Gamma]\!], (\text{fold } (code, x_1, \ldots, x_n)) \\
& \quad \text{as } \exists\rho.\,Clo(\rho)
\end{aligned}
$$

$$
\begin{aligned}
[\![M_1\ M_2]\!] \;=\; & \text{let } \rho, clo = \text{unpack } [\![M_1]\!] \text{ in} \\
& \text{let } code : (Clo(\rho) \times [\![\tau_1]\!]) \to [\![\tau_2]\!] = \\
& \quad \text{proj}_0\,(\text{unfold } clo) \text{ in} \\
& \text{code } (clo, [\![M_2]\!])
\end{aligned}
$$

Closure-passing closure conversion   recursive functions

In the closure-passing variant, recursive functions can be translated as:

$$\llbracket \mu f.\lambda x.M \rrbracket = \begin{array}{l} \textit{let code} = \lambda(\textit{clo}, x). \\ \qquad \textit{let } f = \textit{clo in} \\ \qquad \textit{let } (\_, x_1, \ldots, x_n) = \textit{clo in} \\ \qquad \llbracket M \rrbracket \\ \textit{in } (\textit{code}, x_1, \ldots, x_n) \end{array}$$

where $\{x_1, \ldots, x_n\} = \text{fv}(\mu f.\lambda x.M)$.

No extra field or extra work is required to store or construct a representation of the free variable f: the closure itself plays this role.

However, this untyped code can only be typechecked when recursion is monomorphic.

Exercise:

Check well-typedness with monomorphic recursion.

Closure-passing closure conversion   recursive functions

The problem to adapt this encoding to polymorphic recursion is that recursive occurrences of f are rebuilt from the current invocation of the closure, *i.e.* is monomorphic since the closure is invoked after type specialization.

By contrast, in the environment passing encoding, the environment contained a polymorphic binding for the recursive calls that was filled with the closure before its invokation, *i.e.* with a polymorphic type.

Fortunately, we may slightly change the encoding, using a recursive closure as in the type-passing version, to allow typechecking in System F.

Closure-passing closure conversion   recursive functions

Let τ be $\forall\vec{\alpha}.\,\tau_1 \to \tau_2$ and Γ_f be $f : \tau, \Gamma$ where $\vec{\beta} \,\#\, \Gamma$

$$[\![\mu f\!:\!\tau.\,\lambda x.M]\!] = \textit{let code} =$$
$$\Lambda\vec{\beta}.\,\lambda(clo : Clo[\![\Gamma_f]\!], x : [\![\tau_1]\!]).$$
$$\textit{let } (\_code, f, x_1, \ldots, x_n) : \forall\vec{\beta}.\,UClo([\![\Gamma_f]\!]) =$$
$$\textit{unfold clo in}$$
$$[\![M]\!] \textit{ in}$$
$$\textit{let rec clo} : \forall\vec{\beta}.\,\exists\rho.\,Clo(\rho) = \Lambda\vec{\beta}.$$
$$\textit{pack } [\![\Gamma]\!], (\textit{fold } (code \; \vec{\beta}, clo, x_1, \ldots, x_n)) \textit{ as } \exists\rho.\,Clo(\rho)$$
$$\textit{in clo}$$

Remind that $Clo(R)$ abbreviates $\mu\alpha.\Pi\left((\alpha \times [\![\tau_1]\!]) \to [\![\tau_2]\!]; R\right)$. Hence, $\vec{\beta}$ are free variables of $Clo(R)$.

Here, a polymorphic recursive function is *directly* compiled into a polymorphic recursive closure. Notice that the type of closures is unchanged so the encoding of applications is also unchanged.

Mutually recursive functions   Environment passing

Can we compile mutually recursive functions?

$$M \triangleq \mu(f_1, f_2).(\lambda x_1. M_1, \lambda x_2. M_2)$$

Environment passing:

$$
\begin{aligned}
[\![M]\!] \ = \ &\textit{let } code_i = \lambda(env, x). \\
&\quad \textit{let } (f_1, f_2, x_1, \ldots, x_n) = env \textit{ in} \\
&\quad [\![M_i]\!] \\
&\textit{in} \\
&\textit{let rec } clo_1 = (code_1, (clo_1, clo_2, x_1, \ldots, x_n)) \\
&\quad \textit{and } clo_2 = (code_2, (clo_1, clo_2, x_1, \ldots, x_n)) \textit{ in} \\
&clo_1, clo_2
\end{aligned}
$$

Mutually recursive functions   Environment passing

Can we compile mutually recursive functions?

$$M \triangleq \mu(f_1, f_2).(\lambda x_1. M_1, \lambda x_2. M_2)$$

Environment passing:

$$
\begin{aligned}
[\![M]\!] \;=\; & \mathsf{let}\ code_i = \lambda(env, x). \\
& \quad \mathsf{let}\ (f_1, f_2, x_1, \ldots, x_n) = env\ \mathsf{in} \\
& \quad [\![M_i]\!] \\
& \mathsf{in} \\
& \mathsf{let\ rec}\ env = (clo_1, clo_2, x_1, \ldots, x_n) \\
& \quad \mathsf{and}\ clo_1 = (code_1, env) \\
& \quad \mathsf{and}\ clo_2 = (code_2, env)\ \mathsf{in} \\
& clo_1, clo_2
\end{aligned}
$$

Mutually recursive functions   Environment passing

Can we compile mutually recursive functions?

$$M \triangleq \mu(f_1, f_2).(\lambda x_1. M_1, \lambda x_2. M_2)$$

Environment passing:

let $code_i = \lambda(clo, x).$
    let $(\_, f_1, f_2, x_1, \ldots, x_n) = clo$ *in* $[\![M_i]\!]$
 in
let rec $clo_1 = (code_1, clo_1, clo_2, x_1, \ldots, x_n)$
    and $clo_2 = (code_2, clo_1, clo_2, x_1, \ldots, x_n)$
 in clo_1, clo_2

Question: Can we share the closures c_1 and c_2 in case n is large?

Mutually recursive functions   Environment passing

Can we compile mutually recursive functions?

$$M \triangleq \mu(f_1, f_2).(\lambda x_1. M_1, \lambda x_2. M_2)$$

Environment passing:

> *let* $code_1 = \lambda(clo, x).$
>   *let* $(\_code_1, \_code_2, f_1, f_2, x_1, \ldots, x_n) = clo$ *in* $[\![M_1]\!]$ *in*
> *let* $code_2 = \lambda(clo, x).$
>   *let* $(\_code_2, f_1, f_2, x_1, \ldots, x_n) = clo$ *in* $[\![M_2]\!]$ *in*
> *let rec* $clo_1 = (code_1, code_2, clo_1, clo_2, x_1, \ldots, x_n)$ *and* $clo_2 = clo_1.tail$
> *in* clo_1, clo_2

- $clo_1.tail$ returns a pointer to the tail $(code_2, clo_1, clo_2, x_1, \ldots, x_n)$
  of clo_1 without allocating a new tuple.
- This is only possible with some support from the GC (and
  extra-complexity and runtime cost for GC)

Optimizing representations

Can closure passing and environment passing be mixed?

No because the calling-convention (*i.e.*, the encoding of application) must be uniform.

However, their is some flexibility in the representation of the closure. For instance, the following change is completely local:

$$
\begin{aligned}
[\![\lambda x.\, M]\!] \quad &= \quad let\ code = \lambda(clo, x). \\
&\qquad\quad let\ (\_,\ (\,x_1, \ldots, x_n\,)\,) = clo\ in\ [\![M]\!]\ in \\
&\qquad\quad (code,\ (\,x_1, \ldots, x_n\,)\,) \\
[\![M_1\ M_2]\!] \quad &= \quad let\ clo = [\![M_1]\!]\ in \\
&\qquad\quad let\ code = proj_0\ clo\ in \\
&\qquad\quad code\ (clo, [\![M_2]\!])
\end{aligned}
$$

Applications? When many definitions share the same closure, the closure (or part of it) may be shared.

Encoding of objects

The closure-passing representation of mutually recursive functions is similar to the representations of objects in the object-as-record-of-functions paradigm:

A class definition is an object generator:

$$
\begin{aligned}
&class\ c\ (x_1, \ldots x_q)\{\\
&\quad meth\ m_1 = M_1\\
&\quad \ldots\\
&\quad meth\ m_p = M_p\\
&\}
\end{aligned}
$$

Given arguments for parameter $x_1, \ldots x_1$, it will build recursive methods $m_1, \ldots m_n$.

Encoding of objects

A class can be compiled into an object closure:

$$\begin{aligned}
&let\ m = \\
&\quad let\ m_1 = \lambda(m, x_1, \ldots, x_q).\, M_1\ in \\
&\quad \ldots \\
&\quad let\ m_p = \lambda(m, x_1, \ldots, x_q).\, M_p\ in \\
&\quad \{m_1, \ldots, m_p\}\ in \\
&\lambda x_1 \ldots x_q.\,(m, x_1, \ldots x_q)
\end{aligned}$$

Each m_i is bound to the code for the corresponding method.
The code of all methods are combined into a record of methods,
which is shared between all objects of the same class.

Calling method m_i of an object p is

$$(\mathit{proj}_0\ p).m_i\ p$$

How can we type the encoding?

Typed encoding of objects

Let τ_i be the type of M_i, and row R describe the types of $(x_1, \ldots x_q)$.

Let $Clo(R)$ be $\mu\alpha.\Pi(\{(m_i : \alpha \to \tau_i)^{i \in 1..n}\}; R)$ and $UClo(R)$ its unfolding.

Fields R are hidden in an existential type $\exists\rho.\,\mu\alpha.\Pi(\{(m_i : \alpha \to \tau_i)^{i \in I}\}; \rho)$:

$$
\begin{aligned}
&\text{let } m = \{ \\
&\quad m_1 = \lambda(m, x_1, \ldots x_q : UClo(R)).\, [\![M_1]\!] \\
&\quad \ldots \\
&\quad m_p = \lambda(m, x_1, \ldots x_q : UClo(R)).\, [\![M_p]\!] \\
&\} \text{ in} \\
&\lambda x_1. \ldots \lambda x_q.\, \text{pack } R, \text{fold } (m, x_1, \ldots x_q) \text{ as } \exists\rho.\, (M, \rho)
\end{aligned}
$$

Calling a method of an object p of type M is

$$p\#m_i \triangleq \text{let } \rho, z = \text{unpack } p \text{ in } (\text{proj}_0 \text{ unfold } z).m_i \, z$$

An object has a recursive type but it is *not* a recursive value.

Typed encoding of objects

Typed encoding of objects were first studied in the 90's to understand what objects really are in a type setting.

These encodings are in fact type-preserving compilation of (primitive) objects.

There are several variations on these encodings. See [Bruce et al., 1999] for a comparison.

See [Rémy, 1994] for an encoding of objects in (a small extension of) ML with iso-existentials and universals.

See [Abadi and Cardelli, 1996, 1995] for more details on primitive objects.

Moral of the story

Type-preserving compilation is rather *fun.* (Yes, really!)

It forces compiler writers to make the structure of the compiled program *fully explicit,* in type-theoretic terms.

In practice, building explicit type derivations, ensuring that they remain small and can be efficiently typechecked, can be a lot of work.

Optimizations

Because we have focused on type preservation, we have studied only naïve closure conversion algorithms.

More ambitious versions of closure conversion require program analysis: see, for instance, Steckler and Wand [1997]. These versions *can* be made type-preserving.

Other challenges

Defunctionalization, an alternative to closure conversion, offers an interesting challenge, with a simple solution [Pottier and Gauthier, 2006].

Designing an efficient, type-preserving compiler for an *object-oriented language* is quite challenging. See, for instance, Chen and Tarditi [2005].

Bibliography I

(Most titles have a clickable mark "▷" that links to online versions.)

▷ Martín Abadi and Luca Cardelli. A theory of primitive objects: Untyped and first-order systems. *Information and Computation*, 125(2):78–102, March 1996.

▷ Martín Abadi and Luca Cardelli. A theory of primitive objects: Second-order systems. *Science of Computer Programming*, 25(2–3):81–116, December 1995.

▷ Amal Ahmed and Matthias Blume. Typed closure conversion preserves observational equivalence. In *ACM International Conference on Functional Programming (ICFP)*, pages 157–168, September 2008.

▷ Michael Brandt and Fritz Henglein. Coinductive axiomatization of recursive type equality and subtyping. *Fundamenta Informaticæ*, 33:309–338, 1998.

Bibliography II

▷ Kim B. Bruce, Luca Cardelli, and Benjamin C. Pierce. Comparing object encodings. *Information and Computation*, 155(1/2):108–133, November 1999.

▷ Juan Chen and David Tarditi. A simple typed intermediate language for object-oriented languages. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 38–49, January 2005.

▷ Adam Chlipala. A certified type-preserving compiler from lambda calculus to assembly language. In *ACM Conference on Programming Language Design and Implementation (PLDI)*, pages 54–65, June 2007.

Julien Cretin and Didier Rémy. System F with Coercion Constraints. In *Logics In Computer Science (LICS)*. ACM, July 2014.

Jacques Garrigue and Didier Rémy. Ambivalent Types for Principal Type Inference with GADTs. In *11th Asian Symposium on Programming Languages and Systems*, Melbourne, Australia, December 2013.

Bibliography III

▷ Nadji Gauthier and François Pottier. Numbering matters: First-order canonical forms for second-order recursive types. In *Proceedings of the 2004 ACM SIGPLAN International Conference on Functional Programming (ICFP'04)*, pages 150–161, September 2004. doi: http://doi.acm.org/10.1145/1016850.1016872.

▷ Neal Glew. A theory of second-order trees. In Daniel Le Métayer, editor, *Programming Languages and Systems, 11th European Symposium on Programming, ESOP 2002, held as Part of the Joint European Conference on Theory and Practice of Software, ETAPS 2002, Grenoble, France, April 8-12, 2002, Proceedings*, volume 2305 of *Lecture Notes in Computer Science*, pages 147–161. Springer, 2002. doi: $10.1007/3\text{-}\sigma_2 540\text{-}\sigma_2 45927\text{-}\sigma_2 8\backslash\_11$.

Robert Harper and Benjamin C. Pierce. Design considerations for ML-style module systems. In Benjamin C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 8, pages 293–345. MIT Press, 2005.

Bibliography IV

▷ Konstantin Läufer and Martin Odersky. Polymorphic type inference and abstract data types. *ACM Transactions on Programming Languages and Systems*, 16 (5):1411–1430, September 1994.

Fabrice Le Fessant and Luc Maranget. Optimizing pattern-matching. In *Proceedings of the 2001 International Conference on Functional Programming*. ACM Press, 2001.

Luc Maranget. Warnings for pattern matching. *Journal of Functional Programming*, 17, May 2007.

▷ John C. Mitchell and Gordon D. Plotkin. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems*, 10(3):470–502, 1988.

▷ Benoît Montagu and Didier Rémy. Modeling abstract types in modules with open existential types. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 63–74, January 2009.

Bibliography V

▷ Greg Morrisett and Robert Harper. Typed closure conversion for recursively-defined functions (extended abstract). In *International Workshop on Higher Order Operational Techniques in Semantics (HOOTS)*, volume 10 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science, 1998.

▷ Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From system F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):528–569, May 1999.

▷ Chris Okasaki. *Purely Functional Data Structures*. Cambridge University Press, 1999.

▷ François Pottier and Nadji Gauthier. Polymorphic typed defunctionalization and concretization. *Higher-Order and Symbolic Computation*, 19:125–162, March 2006.

▷ François Pottier and Yann Régis-Gianas. Stratified type inference for generalized algebraic data types. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 232–244, January 2006.

▷ François Pottier and Yann Régis-Gianas. Towards efficient, typed LR parsers. In *ACM Workshop on ML*, volume 148-2 of *Electronic Notes in Theoretical Computer Science*, pages 155–180, March 2006.

▷ François Pottier and Didier Rémy. The essence of ML type inference. In Benjamin C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 10, pages 389–489. MIT Press, 2005.

▷ Didier Rémy. Programming objects with ML-ART: An extension to ML with abstract and record types. In *International Symposium on Theoretical Aspects of Computer Software (TACS)*, pages 321–346. Springer, April 1994.

▷ Didier Rémy and Jérôme Vouillon. Objective ML: An effective object-oriented extension to ML. *Theory and Practice of Object Systems*, 4(1):27–50, 1998.

▷ John C. Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing 83*, pages 513–523. Elsevier Science, 1983.

▷ Andreas Rossberg, Claudio V. Russo, and Derek Dreyer. F-ing modules. *J. Funct. Program.*, 24(5):529–607, 2014. doi: 10.1017/S0956796814000264.

Bibliography VII

▷ Gabriel Scherer and Didier Rémy. Full reduction in the face of absurdity. In *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, pages 685–709, 2015. doi: $10.1007/978$-$\sigma_2 3$-$\sigma_2 662$-$\sigma_2 46669$-$\sigma_2 8\_28$.

▷ Vincent Simonet and François Pottier. A constraint-based approach to guarded algebraic data types. *ACM Trans. Program. Lang. Syst.*, 29(1), January 2007. ISSN 0164-0925. doi: $10.1145/1180475.1180476$.

▷ Paul A. Steckler and Mitchell Wand. Lightweight closure conversion. *ACM Transactions on Programming Languages and Systems*, 19(1):48–86, 1997.

▷ Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones, and Kevin Donnelly. System f with type equality coercions. In *Proceedings of the 2007 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation*, TLDI '07, pages 53–66, New York, NY, USA, 2007. ACM. ISBN 1-59593-393-X. doi: $10.1145/1190315.1190324$.

Bibliography VIII

▷ Dimitrios Vytiniotis, Simon Peyton jones, Tom Schrijvers, and Martin Sulzmann. Outsidein(x) modular type inference with local assumptions. *J. Funct. Program.*, 21(4-5):333–412, September 2011. ISSN 0956-7968. doi: 10.1017/S0956796811000098.