

A Church-Style Intermediate Language for ML^F

Didier Rémy¹ Boris Yakobowski²

¹ INRIA Paris - Rocquencourt

² CEA, LIST, Laboratoire Sûreté des Logiciels,
Boîte 94, 91191 Gif-sur-Yvette Cedex, France.

Abstract ML^F is a type system that seamlessly merges ML-style implicit but second-class polymorphism with System-F explicit first-class polymorphism. We present $x\text{ML}^F$, a Church-style version of ML^F with full type information that can easily be maintained during reduction. All parameters of functions are explicitly typed and both type abstraction and type instantiation are explicit. However, type instantiation in $x\text{ML}^F$ is more general than type application in System F. We equip $x\text{ML}^F$ with a small-step reduction semantics that allows reduction in any context and show that this relation is confluent and type preserving. We also show that both subject reduction and progress hold for weak-reduction strategies, including call-by-value with the value-restriction.

Introduction

ML^F (Le Botlan and Rémy 2003, 2007; Rémy and Yakobowski 2008b) is a type system that seamlessly merges ML-style implicit but second-class polymorphism with System-F explicit first-class polymorphism. This is done by enriching System-F types. Indeed, System F is not well-suited for partial type inference, as illustrated by the following example. Assume that a function, say *choice*, of type $\forall(\alpha) \alpha \rightarrow \alpha \rightarrow \alpha$ and the identity function *id*, of type $\forall(\beta) \beta \rightarrow \beta$, have been defined. How can the application *choice* to *id* be typed in System F? Should *choice* be applied to the type $\forall(\beta) \beta \rightarrow \beta$ of the identity that is itself kept polymorphic? Or should it be applied to the monomorphic type $\gamma \rightarrow \gamma$, with the identity being applied to γ (where γ is bound in a type abstraction in front of the application)? Unfortunately, these alternatives have incompatible types, respectively $(\forall(\alpha) \alpha \rightarrow \alpha) \rightarrow (\forall(\alpha) \alpha \rightarrow \alpha)$ and $\forall(\gamma) (\gamma \rightarrow \gamma) \rightarrow (\gamma \rightarrow \gamma)$: none is an instance of the other. Hence, in System F, one is forced to irreversibly choose between one of the two explicitly typed terms.

However, a type inference system cannot choose between the two, as this would sacrifice completeness and be somehow arbitrary. This is why ML^F enriches types with instance-bounded polymorphism, which allows to write more expressive types that factor out in a single type all typechecking alternatives in such cases as the example of *choice*. Now, the type $\forall(\alpha \geq \tau) \alpha \rightarrow \alpha$, which should be read “ $\alpha \rightarrow \alpha$ where α is any instance of τ ”, can be assigned to *choice* *id*, and the two previous alternatives can be recovered *a posteriori* by choosing different instances for α .

Currently, the language ML^F comes with a Curry-style version $i\text{ML}^F$ where no type information is needed and a type-inference version $e\text{ML}^F$ that requires partial type information (Le Botlan and Rémy 2007). However, $e\text{ML}^F$ is not quite in Church’s style, since a large amount of type information is still implicit and partial type information cannot be easily maintained during reduction. Hence, while $e\text{ML}^F$ is a good surface language, it is not a good candidate for use as an internal language during the compilation process, where some program transformations, and perhaps some reduction steps, are being performed. This has been a problem for the adoption of ML^F in the Haskell community (Peyton Jones 2003), as the Haskell compilation chain uses an explicitly-typed internal language.

This is also an obstacle to proving subject reduction, which does not hold in $e\text{ML}^F$. In a way, this is unavoidable in a language with non-trivial partial type inference. Indeed, type annotations cannot be completely dropped, but must at least be transformed and reorganized during reduction. Still, one could expect that $e\text{ML}^F$ be equipped with reduction rules for type annotations. This has actually been considered in the original presentation of ML^F , but only with limited success. The reduction kept track of annotation sites during reduction; this showed, in particular, that no new annotation site needs to be introduced during reduction. Unfortunately, the exact form of annotations could not be maintained during reduction, by lack of an appropriate language to describe their computation. As a result, it has only been shown that some type derivation can be rebuilt after the reduction of a well-typed program, but without exhibiting an algorithm to compute them during reduction.

Independently, Rémy and Yakobowski (2008b) have introduced graphic constraints, both to simplify the presentation of ML^F and to improve its type inference algorithm. This also lead to a simpler, more expressive definition of ML^F .

In this paper, we present $x\text{ML}^F$, a Church-style version of ML^F that contains full type information. In fact, type checking becomes a simple and local verification process—by contrast with type inference in $e\text{ML}^F$, which is based on unification. In $x\text{ML}^F$, type abstraction, type instantiation, and all parameters of functions are explicit, as in System F. However, type instantiation is more general and more atomic than type application in System F: we use explicit type instantiation expressions that are proof evidences for the type instance relations.

In addition to the usual β -reduction, we give a series of reduction rules for simplifying type instantiations. These rules are confluent when allowed in any context. Moreover, reduction preserves typings, and is sufficient to reduce all typable expressions to a value when used in either a call-by-value or call-by-name setting. This establishes the soundness of ML^F for a call-by-name semantics for the first time. Notably, $x\text{ML}^F$ is a conservative extension of System F.

The paper is organized as follows. We present $x\text{ML}^F$, its syntax and its static and dynamic semantics in §1. We study its main properties, including type soundness for different evaluations strategies in §2. We discuss possible variations, as well as related and future works in §3. All proofs are omitted, but can be found in (Yakobowski 2008, Chapters 14 & 15).

$\alpha, \beta, \gamma, \delta$	Type variable	x, y, z	Term variable
$\tau ::=$	Type	$a ::=$	Term
α	Type variable	x	Variable
$\tau \rightarrow \tau$	Arrow type	$\lambda(x : \tau) a$	Function
$\forall(\alpha \geq \tau) \tau$	Quantification	$a a$	Application
\perp	Bottom type	$\Lambda(\alpha \geq \tau) a$	Type function
$\phi ::=$	Instantiation	$a \phi$	Instantiation
τ	Bottom	$\text{let } x = a \text{ in } a$	Let-binding
$!\alpha$	Abstract	$\Gamma ::=$	Environment
$\forall(\geq \phi)$	Inside	\emptyset	Empty
$\forall(\alpha \geq) \phi$	Under	$\Gamma, \alpha \geq \tau$	Type variable
$\&$	\forall -elimination	$\Gamma, x : \tau$	Term variable
\wp	\forall -introduction		
$\phi; \phi$	Composition		
$\mathbb{1}$	Identity		

Figure 1. Grammar of types, instantiations, and terms

1 The calculus

Types, instantiations, terms, and typing environments All the syntactic definitions of $x\text{ML}^F$ can be found in Figure 1. We assume given a countable collection of variables ranged over by letters α , β , γ , and δ . As usual, types include type variables and arrow types. Other type constructors will be added later—straightforwardly, as the arrow constructor receives no special treatment. Types also include a bottom type \perp that corresponds to the System-F type $\forall\alpha.\alpha$. Finally, a type may also be a form of bounded quantification $\forall(\alpha \geq \tau) \tau'$, called *flexible* quantification, that generalizes the $\forall\alpha.\tau$ form of System F and, intuitively, restricts the variable α to range only over instances of τ . The variable α is bound in τ' but not in τ . (We may write $\forall(\alpha) \tau'$ when the bound τ is \perp .)

In Church-style System F, type instantiation inside terms is simply type application, of the form $a \tau$. By contrast, type instantiation $a \phi$ in $x\text{ML}^F$ details every intermediate instantiation step, so that it can be checked locally. Intuitively, the *instantiation* ϕ transforms a type τ into another type τ' that is an instance of τ . In a way, ϕ is a witness for the instance relation that holds between τ and τ' . It is therefore easier to understand instantiations altogether with their static semantics, which will be explained in the next section.

Terms of $x\text{ML}^F$ are those of the λ -calculus enriched with **let** constructs, with two small differences. Type instantiation $a \phi$ generalizes System-F type application. Type abstractions are extended with an instance bound τ and written $\Lambda(\alpha \geq \tau) a$. The type variable α is bound in a , but not in τ . We abbreviate $\Lambda(\alpha \geq \perp) a$ as $\Lambda(\alpha) a$, which simulates the type abstraction $\Lambda\alpha.a$ of System F. We write $\text{ftv}(\tau)$ and $\text{ftv}(a)$ the set of type variables that appear free in τ and a .

We identify types, instantiations, and terms up to the renaming of bound variables. The capture-avoiding substitution of a variable v inside an expression s by an expression s' is written $s\{v \leftarrow s'\}$.

$\frac{\text{INST-BOT}}{\Gamma \vdash \tau : \perp \leq \tau}$	$\frac{\text{INST-UNDER}}{\Gamma, \alpha \geq \tau \vdash \phi : \tau_1 \leq \tau_2}$	$\frac{\text{INST-ABSTR}}{\alpha \geq \tau \in \Gamma}$
$\frac{\text{INST-INSIDE}}{\Gamma \vdash \phi : \tau_1 \leq \tau_2}$	$\frac{\text{INST-INTRO}}{\alpha \notin \text{ftv}(\tau)}$	
$\frac{\Gamma \vdash \phi : \tau_1 \leq \tau_2}{\Gamma \vdash \forall(\geq \phi) : \forall(\alpha \geq \tau_1) \tau \leq \forall(\alpha \geq \tau_2) \tau}$	$\frac{\Gamma \vdash \forall(\alpha \geq \tau) \phi : \forall(\alpha \geq \tau) \tau_1 \leq \forall(\alpha \geq \tau) \tau_2}{\Gamma \vdash \forall(\alpha \geq \tau) \tau_1 \leq \forall(\alpha \geq \tau) \tau_2}$	$\frac{\Gamma \vdash !\alpha : \tau \leq \alpha}{\Gamma \vdash !\alpha : \tau \leq \alpha}$
$\frac{\text{INST-COMP}}{\Gamma \vdash \phi_1 : \tau_1 \leq \tau_2}$	$\frac{\text{INST-ELIM}}{\Gamma \vdash \& : \forall(\alpha \geq \tau) \tau' \leq \tau' \{ \alpha \leftarrow \tau \}}$	$\frac{\text{INST-ID}}{\Gamma \vdash \mathbb{1} : \tau \leq \tau}$
$\frac{\Gamma \vdash \phi_2 : \tau_2 \leq \tau_3}{\Gamma \vdash \phi_1; \phi_2 : \tau_1 \leq \tau_3}$		

Figure 2. Type instance

As usual, type environments assign types to program variables. However, instead of just listing type variables, as is the case in System F, type variables are also assigned a bound in a binding of the form $\alpha \geq \tau$. We write $\text{dom}(\Gamma)$ for the set of all terms and type variables that are bound by Γ . We also assume that typing environments are *well-formed*, *i.e.* they do not bind twice the same variable and free type variables appearing in a type of the environment Γ must be bound earlier in Γ . Formally, the empty environment is well-formed and, given a (well-formed) environment Γ , the relations $\alpha \notin \text{dom}(\Gamma)$ and $\text{ftv}(\tau) \subseteq \text{dom}(\Gamma)$ must hold to form environments $\Gamma, \alpha \geq \tau$ and $\Gamma, x : \tau$.

Instantiations Instantiations ϕ are defined in Figure 1. Their typing, described in Figure 2, are *type instance* judgments of the form $\Gamma \vdash \phi : \tau \leq \tau'$, stating that in environment Γ , the instantiation ϕ transforms the type τ into the type τ' .

The *bottom* instantiation τ expresses that (any) type τ is an instance of the bottom type. The *abstract* instantiation $!\alpha$, which assumes that the hypothesis $\alpha \geq \tau$ is in the environment, abstracts the bound τ of α as the type variable α . The *inside* instantiation $\forall(\geq \phi)$ applies ϕ to the bound τ' of a flexible quantification $\forall(\alpha' \geq \tau') \tau$. Conversely, the *under* instantiation $\forall(\alpha \geq) \phi$ applies ϕ to the type τ under the quantification. The type variable α is bound in ϕ ; the environment in the premise of the rule INST-UNDER is increased accordingly. The *quantifier introduction* \wp^3 introduces a fresh trivial quantification $\forall(\alpha \geq \perp)$. Conversely, the *quantifier elimination* $\&$ eliminates the bound of a type of the form $\forall(\alpha \geq \tau) \tau'$ by substituting τ for α in τ' . This amounts to definitely choosing the present bound τ for α , while the bound before the application could be further instantiated by some inside instantiation. The *composition* $\phi; \phi'$ witnesses the transitivity of type instance, while the *identity* instantiation $\mathbb{1}$ witnesses reflexivity.

³ The choice of \wp is only by symmetry with the elimination form $\&$ described next, and has no connection at all with linear logic.

$$\begin{array}{lcl}
\tau (!\alpha) & = & \alpha \\
\perp \tau & = & \tau \\
\tau \mathbb{1} & = & \tau \\
\tau (\phi_1; \phi_2) & = & (\tau \phi_1) \phi_2 \\
\tau \wp & = & \forall (\alpha \geq \perp) \tau \quad \alpha \notin \text{ftv}(\tau) \\
(\forall (\alpha \geq \tau) \tau') \& & = \tau' \{ \alpha \leftarrow \tau \} \\
(\forall (\alpha \geq \tau) \tau') (\forall (\geq \phi)) & = & \forall (\alpha \geq \tau \phi) \tau' \\
(\forall (\alpha \geq \tau) \tau') (\forall (\alpha \geq) \phi) & = & \forall (\alpha \geq \tau) (\tau' \phi)
\end{array}$$

Figure 3. Type instantiation (on types)

Example Let τ_{\min} , τ_{cmp} , and τ_{and} be the types of the parametric minimum and comparison functions and of the conjunction of boolean formulas:

$$\begin{array}{lcl}
\tau_{\min} & \triangleq & \forall (\alpha \geq \perp) \alpha \rightarrow \alpha \rightarrow \alpha \\
\tau_{\text{cmp}} & \triangleq & \forall (\alpha \geq \perp) \alpha \rightarrow \alpha \rightarrow \text{bool} \\
\tau_{\text{and}} & \triangleq & \text{bool} \rightarrow \text{bool} \rightarrow \text{bool}
\end{array}$$

Let ϕ be the instantiation $\forall (\geq \text{bool}); \&$. Then, $\vdash \phi : \tau_{\min} \leq \tau_{\text{and}}$ and $\vdash \phi : \tau_{\text{cmp}} \leq \tau_{\text{and}}$ hold. Let τ_K be the type $\forall (\alpha \geq \perp) \forall (\beta \geq \perp) \alpha \rightarrow \beta \rightarrow \alpha$ (e.g. of the λ -term $\lambda(x) \lambda(y) x$) and ϕ' be the instantiation⁴ $\forall (\alpha \geq) (\forall (\geq \alpha)); \&$. Then, $\phi' : \tau_K \leq \tau_{\min}$.

Type application As above, we often instantiate a quantification over \perp and immediately substitute the result. Moreover, this pattern corresponds to the System-F unique instantiation form. Therefore, we define $\langle \tau \rangle$ as syntactic sugar for $(\forall (\geq \tau); \&)$. The instantiations ϕ and ϕ' can then be abbreviated as $\langle \text{bool} \rangle$ and $\forall (\alpha \geq) \langle \alpha \rangle$. More generally, we write $\langle \phi \rangle$ for the computation $(\forall (\geq \phi); \&)$.

Properties of instantiations Since instantiations make all steps in the instance relation explicit, their typing is deterministic.

Lemma 1. *If $\Gamma \vdash \phi : \tau \leq \tau_1$ and $\Gamma' \vdash \phi : \tau \leq \tau_2$, then $\tau_1 = \tau_2$.*

The use of Γ' instead of Γ may be surprising. However, Γ does not contribute to the instance relation, except in the side condition of rule INST-ABSTR. Hence, the type instance relation defines a partial function, called *type instantiation*⁵ that, given an instantiation ϕ and a type τ , returns (if it exists) the unique type $\tau \phi$ such that $\vdash \phi : \tau \leq \tau \phi$. An inductive definition of this function is given in Figure 3. Type instantiation is complete for type instance:

Lemma 2. *If $\Gamma \vdash \phi : \tau \leq \tau'$, then $\tau \phi = \tau'$.*

However, the fact that $\tau \phi$ may be defined and equal to τ' does not imply that $\Gamma \vdash \phi : \tau \leq \tau'$ holds for some Γ . Indeed, type instantiation does not check the premise of rule INST-ABSTR. This is intentional, as it avoids parametrizing type instantiation over the type environment. This means that type instantiation is not sound *in general*. This is never a problem, however, since we only use type instantiation originating from well-typed terms for which there always exists some context Γ such that $\Gamma \vdash \phi : \tau \leq \tau'$.

⁴ The occurrence of α in the inside instantiation is bound by the under instantiation.

⁵ There should never be any ambiguity with the operation $a \phi$ on expressions; moreover, both operations have strong similarities.

$$\begin{array}{c}
\text{VAR} \\
\frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau} \\
\\
\text{LET} \\
\frac{\Gamma \vdash a : \tau \quad \Gamma, x : \tau \vdash a' : \tau'}{\Gamma \vdash \text{let } x = a \text{ in } a' : \tau'} \\
\\
\text{APP} \\
\frac{\Gamma \vdash a_1 : \tau_2 \rightarrow \tau_1 \quad \Gamma \vdash a_2 : \tau_2}{\Gamma \vdash a_1 a_2 : \tau_1} \\
\\
\text{ABS} \\
\frac{\Gamma, x : \tau \vdash a : \tau'}{\Gamma \vdash \lambda(x : \tau) a : \tau \rightarrow \tau'} \\
\\
\text{TABS} \\
\frac{\Gamma, \alpha \geq \tau' \vdash a : \tau \quad \alpha \notin \text{ftv}(\Gamma)}{\Gamma \vdash \Lambda(\alpha \geq \tau') a : \forall(\alpha \geq \tau') \tau} \\
\\
\text{TAPP} \\
\frac{\Gamma \vdash a : \tau \quad \Gamma \vdash \phi : \tau \leq \tau'}{\Gamma \vdash a \phi : \tau'}
\end{array}$$

Figure 4. Typing rules for $x\text{ML}^F$

We say that types τ and τ' are equivalent in Γ if there exist ϕ and ϕ' such that $\Gamma \vdash \phi : \tau \leq \tau'$ and $\Gamma \vdash \phi' : \tau' \leq \tau$. Although types of $x\text{ML}^F$ are *syntactically* the same as the types of $i\text{ML}^F$ —the Curry-style version of ML^F (Le Botlan and Rémy 2007)—they are richer, because type equivalence in $x\text{ML}^F$ is finer than type equivalence in $i\text{ML}^F$, as will be explained in §3.

Typing rules for $x\text{ML}^F$ Typing rules are defined in Figure 4. Compared with System F, the novelties are type abstraction and type instantiation, unsurprisingly. The typing of a type abstraction $\Lambda(\alpha \geq \tau) a$ extends the typing environment with the type variable α bound by τ . The typing of a type instantiation $a \phi$ resembles the typing of a coercion, as it just requires the instantiation ϕ to transform the type of a into the type of the result. Of course, it has the full power of the type application rule of System F. For example, the type instantiation $a \langle \tau \rangle$ has type $\tau' \{ \alpha \leftarrow \tau \}$ provided the term a has type $\forall(\alpha) \tau'$. As in System F, a well-typed closed term has a unique type—in fact, a unique typing derivation.

A let-binding $\text{let } x = a_1 \text{ in } a_2$ cannot entirely be treated as an abstraction for an immediate application $(\lambda(x : \tau_1) a_2) a_1$ because the former does not require a type annotation on x whereas the latter does. This is nothing new, and the same as in System F extended with let-bindings. (Notice however that τ_1 , which is the type of a_1 , is fully determined by a_1 and could be synthesized by a typechecker.)

Example Let id stand for the identity $\Lambda(\alpha \geq \perp) \lambda(x : \alpha) x$ and τ_{id} for the type $\forall(\alpha \geq \perp) \alpha \rightarrow \alpha$. We have $\vdash \text{id} : \tau_{\text{id}}$. The function `choice` mentioned in the introduction, may be defined as $\Lambda(\beta \geq \perp) \lambda(x : \beta) \lambda(y : \beta) x$. It has type $\forall(\beta \geq \perp) \beta \rightarrow \beta \rightarrow \beta$. The application of `choice` to `id`, which we refer to below as `choice_id`, may be defined as $\Lambda(\beta \geq \tau_{\text{id}}) \text{choice } \langle \beta \rangle (\text{id } (!\beta))$ and has type $\forall(\beta \geq \tau_{\text{id}}) \beta \rightarrow \beta$. The term `choice_id` may also be given weaker types by type instantiation. For example, `choice_id &` has type $(\forall(\alpha \geq \perp) \alpha \rightarrow \alpha) \rightarrow (\forall(\alpha \geq \perp) \alpha \rightarrow \alpha)$ as in System F, while `choice_id (&; \forall(\gamma \geq) (\forall(\geq \langle \gamma \rangle); &))` has the ML type $\forall(\gamma \geq \perp) (\gamma \rightarrow \gamma) \rightarrow \gamma \rightarrow \gamma$.

Reduction The semantics of the calculus is given by a small-step reduction semantics. We let reduction occur in any context, including under abstractions.

$(\lambda(x : \tau) a_1) a_2$	$\longrightarrow a_1\{x \leftarrow a_2\}$	(β)
$\text{let } x = a_2 \text{ in } a_1$	$\longrightarrow a_1\{x \leftarrow a_2\}$	(β_{let})
$a \mathbb{1}$	$\longrightarrow a$	(ι -ID)
$a(\phi; \phi')$	$\longrightarrow a\phi(\phi')$	(ι -SEQ)
$a \wp$	$\longrightarrow \Lambda(\alpha \geq \perp) a$	$\alpha \notin \text{ftv}(a)$ (ι -INTRO)
$(\Lambda(\alpha \geq \tau) a) \&$	$\longrightarrow a\{!\alpha \leftarrow \mathbb{1}\}\{\alpha \leftarrow \tau\}$	(ι -ELIM)
$(\Lambda(\alpha \geq \tau) a) (\forall(\alpha \geq) \phi)$	$\longrightarrow \Lambda(\alpha \geq \tau) (a\phi)$	(ι -UNDER)
$(\Lambda(\alpha \geq \tau) a) (\forall(\geq) \phi)$	$\longrightarrow \Lambda(\alpha \geq \tau \phi) a\{!\alpha \leftarrow \phi; !\alpha\}$	(ι -INSIDE)
$E[a] \longrightarrow E[a']$	$\text{if } a \longrightarrow a'$	(CONTEXT)

Figure 5. Reduction rules

That is, the evaluation contexts are single-hole contexts, given by the grammar:

$$\begin{aligned}
 E ::= & [\cdot] \mid E\phi \mid \lambda(x : \tau) E \mid \Lambda(\alpha \geq \tau) E \\
 & \mid E a \mid a E \mid \text{let } x = E \text{ in } a \mid \text{let } x = a \text{ in } E
 \end{aligned}$$

The reduction rules are described in Figure 5. As usual, basic reduction steps contain β -reduction, with the two variants (β) and (β_{let}). Other basic reduction rules, related to the reduction of type instantiations and called ι -steps, are described below. The one-step reduction is closed under the context rule. We write \longrightarrow_{β} and \longrightarrow_{ι} for the two subrelations of \longrightarrow that contains only CONTEXT and β -steps or ι -step, respectively. Finally, the reduction is the reflexive and transitive closure $\longrightarrow^{\rightarrow}$ of the one-step reduction relation.

Reduction of type instantiation Type instantiation redexes are all of the form $a\phi$. The first three rules do not constrain the form of a . The identity type instantiation is just dropped (Rule ι -ID). A type instantiation composition is replaced by the successive corresponding type instantiations (Rule ι -SEQ). Rule ι -INTRO introduces a new type abstraction in front of a ; we assume that the bound variable α is fresh in a . The other three rules require the type instantiation to be applied to a type abstraction $\Lambda(\alpha \geq \tau) a$. Rule ι -UNDER propagates the type instantiation under the bound, inside the body a . By contrast, Rule ι -INSIDE propagates the type instantiation ϕ inside the bound, replacing τ by $\tau\phi$. However, as the bound of α has changed, the domain of the type instantiations $!\alpha$ is no more τ , but $\tau\phi$. Hence, in order to maintain well-typedness, all the occurrences of the instantiation $!\alpha$ in a must be simultaneously replaced by the instantiation $(\phi; !\alpha)$. Here, the instantiation $!\alpha$ is seen as atomic, *i.e.* all occurrences of $!\alpha$ are substituted, but other occurrences of α are left unchanged (see the appendix for the formal definition). For instance, if a is the term

$$\Lambda(\alpha \geq \tau) \lambda(x : \alpha \rightarrow \alpha) \lambda(y : \perp) y(\alpha \rightarrow \alpha) (z(!\alpha))$$

then, the type instantiation $a(\forall(\geq)\phi)$ reduces to:

$$\Lambda(\alpha \geq \tau\phi) \lambda(x : \alpha \rightarrow \alpha) \lambda(y : \perp) y(\alpha \rightarrow \alpha) (z(\phi; !\alpha))$$

Rule ι -ELIM eliminates the type abstraction, replacing all the occurrences of α inside a by the bound τ . All the occurrences of $!\alpha$ inside τ (used to instantiate τ

into α) become vacuous and must be replaced by the identity instantiation. For example, reusing the term a above, $a \&$ reduces to $\lambda(x : \tau \rightarrow \tau) \lambda(y : \perp) y (\tau \rightarrow \tau) (z \mathbb{1})$. Notice that type instantiations $a \tau$ and $a (!\alpha)$ are irreducible.

Examples of reduction Let us reuse the term `choice_id` defined in §1 as $\Lambda(\beta \geq \tau_{\text{id}}) \text{choice } \langle \beta \rangle (\text{id } (!\beta))$. Remember that $\langle \tau \rangle$ stands for the System-F type application τ and expands to $(\forall (\geq \tau); \&)$. Therefore, the type instantiation `choice` $\langle \beta \rangle$ reduces to the term $\lambda(x : \beta) \lambda(y : \beta) x$ by ι -SEQ, ι -INSIDE and ι -ELIM. Hence, the term `choice_id` reduces by these rules, CONTEXT, and $\langle \beta \rangle$ to the expression $\Lambda(\beta \geq \tau_{\text{id}}) \lambda(y : \beta) \text{id } (!\beta)$.

Below are three specialized versions of `choice_id` (remember that $\forall(\alpha) \tau$ and $\Lambda(\alpha) a$ are abbreviations for $\forall(\alpha \geq \perp) \tau$ and $\Lambda(\alpha \geq \perp) a$). Here, all type instantiations are eliminated by reduction, but this is not always possible in general.

$$\begin{aligned} \text{choice_id } \langle \langle \text{int} \rangle \rangle & : (\text{int} \rightarrow \text{int}) \rightarrow (\text{int} \rightarrow \text{int}) \\ & \longrightarrow \lambda(y : \text{int} \rightarrow \text{int}) (\lambda(x : \text{int}) x) \\ \text{choice_id } \& & : (\forall(\alpha) \alpha \rightarrow \alpha) \rightarrow (\forall(\alpha) \alpha \rightarrow \alpha) \\ & \longrightarrow \lambda(y : \forall(\alpha) \alpha \rightarrow \alpha) (\Lambda(\alpha) \lambda(x : \alpha) x) \\ \text{choice_id } (\& ; \forall(\gamma \geq) (\forall(\geq \langle \gamma \rangle); \&)) & : \forall(\gamma) (\gamma \rightarrow \gamma) \rightarrow (\gamma \rightarrow \gamma) \\ & \longrightarrow \Lambda(\gamma) \lambda(y : \gamma \rightarrow \gamma) (\lambda(x : \gamma) x) \end{aligned}$$

System F as a subsystem of $x\text{MLF}$ System F can be seen as a subset of $x\text{MLF}$, using the following syntactic restrictions: all quantifications are of the form $\forall(\alpha) \tau$ and \perp is not a valid type anymore (however, as in System F, $\forall(\alpha) \alpha$ is); all type abstractions are of the form $\Lambda(\alpha) a$; and all type instantiations are of the form $a \langle \tau \rangle$.

The derived typing rule for $\Lambda(\alpha) a$ and $a \langle \tau \rangle$ are exactly the System-F typing rules for type abstraction and type application. Hence, typechecking in this restriction of $x\text{MLF}$ corresponds to typechecking in System F.

Moreover, the reduction in this restriction also corresponds to reduction in System F. Indeed, a reducible type application is necessarily of the form $(\Lambda(\alpha) a) \langle \tau \rangle$ and can always be reduced to $a\{\alpha \leftarrow \tau\}$ as follows:

$$\begin{aligned} (\Lambda(\alpha) a) \langle \tau \rangle & = (\Lambda(\alpha \geq \perp) a) (\forall(\geq \tau); \&) & (1) \\ & \longrightarrow (\Lambda(\alpha \geq \perp) a) (\forall(\geq \tau)) (\&) & (2) \\ & \longrightarrow (\Lambda(\alpha \geq \perp) a) a\{!\alpha \leftarrow \tau; !\alpha\} (\&) = (\Lambda(\alpha \geq \tau) a) (\&) & (3) \\ & \longrightarrow a\{!\alpha \leftarrow \mathbb{1}\} \{\alpha \leftarrow \tau\} = a\{\alpha \leftarrow \tau\} & (4) \end{aligned}$$

Step (1) is by definition; step (2) is by ι -SEQ; step (3) is by ι -INSIDE, step (4) is by ι -ELIM and equality steps (3) and (4) are by type instantiation and by assumption as a is a term of System F, thus in which $!\alpha$ does not appear.

2 Properties of reduction

The reduction has been defined so that the type erasure of a reduction sequence in $x\text{MLF}$ is a reduction sequence in the untyped λ -calculus. Formally, the type

erasure of a term a of $x\text{ML}^F$ is the untyped λ -term $[a]$ defined inductively by

$$\begin{array}{ll} [x] = x & [\text{let } x = a_1 \text{ in } a_2] = \text{let } x = [a_1] \text{ in } [a_2] \\ [a \phi] = [a] & [\lambda(x : \tau) a] = \lambda(x) [a] \\ [a_1 a_2] = [a_1] [a_2] & [A(\alpha \geq \tau) a] = [a] \end{array}$$

It is immediate to verify that two terms related by ι -reduction have the same type erasure. Moreover, if a β -reduces to a' , then the type erasure of a β -reduces to the type erasure of a' in one step in the untyped λ -calculus.

2.1 Subject reduction

Reduction of $x\text{ML}^F$, which can occur in any context, preserves typings. This relies on weakening and substitution lemmas for both instance and typing judgments.

Lemma 3 (Weakening). *Let $\Gamma, \Gamma', \Gamma''$ be a well-formed environment.*

If $\Gamma, \Gamma'' \vdash \phi : \tau_1 \leq \tau_2$, then $\Gamma, \Gamma', \Gamma'' \vdash \phi : \tau_1 \leq \tau_2$.

If $\Gamma, \Gamma'' \vdash a : \tau'$, then $\Gamma, \Gamma', \Gamma'' \vdash a : \tau'$.

Lemma 4 (Term substitution). *Assume that $\Gamma \vdash a' : \tau'$ holds.*

If $\Gamma, x : \tau', \Gamma' \vdash \phi : \tau_1 \leq \tau_2$ then $\Gamma, \Gamma' \vdash \phi : \tau_1 \leq \tau_2$.

If $\Gamma, x : \tau', \Gamma' \vdash a : \tau$, then $\Gamma, \Gamma' \vdash a\{x \leftarrow a'\} : \tau$

The next lemma, which expresses that we can substitute an instance bound inside judgments, ensures the correctness of Rule ι -ELIM.

Lemma 5 (Bound substitution). *Let φ and θ be respectively the substitutions $\{\alpha \leftarrow \tau\}$ and $\{!\alpha \leftarrow \mathbb{1}\}\{\alpha \leftarrow \tau\}$.*

If $\Gamma, \alpha \geq \tau, \Gamma' \vdash \phi : \tau_1 \leq \tau_2$ then $\Gamma, \Gamma' \varphi \vdash \phi \theta : \tau_1 \varphi \leq \tau_2 \varphi$.

If $\Gamma, \alpha \geq \tau, \Gamma' \vdash a : \tau'$ then $\Gamma, \Gamma' \varphi \vdash a \theta : \tau' \varphi$.

Finally, the following lemma ensures that an instance bound can be instantiated, proving in turn the correctness of the rule ι -INSIDE.

Lemma 6 (Narrowing). *Assume $\Gamma \vdash \phi : \tau \leq \tau'$. Let θ be $\{!\alpha \leftarrow \phi; !\alpha\}$.*

If $\Gamma, \alpha \geq \tau, \Gamma' \vdash \phi' : \tau_1 \leq \tau_2$ then $\Gamma, \alpha \geq \tau', \Gamma' \vdash \phi' \theta : \tau_1 \leq \tau_2$.

If $\Gamma, \alpha \geq \tau, \Gamma' \vdash a : \tau''$ then $\Gamma, \alpha \geq \tau', \Gamma' \vdash a \theta : \tau''$

Subject reduction is an easy consequence of all these results.

Theorem 1 (Subject reduction). *If $\Gamma \vdash a : \tau$ and $a \longrightarrow a'$ then, $\Gamma \vdash a' : \tau$.*

2.2 Confluence

Theorem 2. *The relation \longrightarrow_β is confluent. The relations \longrightarrow_ι and \longrightarrow are confluent on the terms well-typed in some context.*

This result is proved using the standard technique of parallel reductions (Barendregt 1984). Thus β -reduction and ι -reduction are independent; this allows for instance to perform ι -reductions under λ -abstractions as far as possible while keeping a weak evaluation strategy for β -reduction.

The restriction to well-typed terms for the confluence of ι -reduction is due to two things. First, the rule ι -INSIDE is not applicable to ill-typed terms in which $\tau \phi$ cannot be computed (for example $(\Lambda(\alpha \geq \text{int}) a) (\forall (\geq \&))$). Second, $\tau \phi$ can sometimes be computed, even though $\Gamma \vdash \phi : \tau \leq \tau'$ never holds (for example if ϕ is $!\alpha$ and τ is not the bound of α in Γ). Hence, type errors may be either revealed or silently reduced and perhaps eliminated, depending on the reduction path. As an example, let a be the term

$$(\Lambda(\alpha \geq \forall(\gamma) \gamma) ((\Lambda(\beta \geq \text{int}) x) (\forall (\geq !\alpha)))) (\forall (\geq \&))$$

We have both $a \rightarrow \Lambda(\alpha \geq \perp) ((\Lambda(\beta \geq \text{int}) x) (\forall (\geq \& !\alpha))) \not\rightarrow$, and $a \rightarrow (\Lambda(\alpha \geq \forall(\gamma) \gamma) \Lambda(\beta \geq \alpha) x) (\forall (\geq \&)) \rightarrow \Lambda(\alpha \geq \perp) \Lambda(\beta \geq \alpha) x \not\rightarrow$.

The fact that ill-typed terms may not be confluent is not new: for instance, this is already the case with η -reduction in System F. We believe this is not a serious issue. In practice, this means that typechecking should be performed before any program simplification, which is usually the case anyway.

2.3 Strong normalization

We conjecture, but have not proved, that all reduction sequences are finite.

2.4 Accommodating weak reduction strategies and constants

In order to show that the calculus may also be used as the core of a programming language, we now introduce constants and restricts the semantics to a weak evaluation strategy.

We let the letter c range over constants. Each constant comes with its arity $|c|$. The dynamic semantics of constants must be provided by primitive reduction rules, called δ -rules. However, these are usually of a certain form. To characterize δ -rules (and values), we partition constants into *constructors* and *primitives*, ranged over by letters C and f , respectively. The difference between the two lies in their semantics: primitives (such as $+$) are reduced when fully applied, while constructors (such as `cons`) are irreducible and typically eliminated when passed as argument to primitives.

In order to classify constructed values, we assume given a collection of type constructors κ , together with their arities $|\kappa|$. We extend types with constructed types $\kappa (\tau_1, \dots, \tau_{|\kappa|})$. We write $\bar{\alpha}$ for a sequence of variables $\alpha_1, \dots, \alpha_k$ and $\forall (\bar{\alpha}) \tau$ for the type $\forall (\alpha_1) \dots \forall (\alpha_k) \tau$. The static semantics of constants is given by an initial typing environment Γ_0 that assigns to every constant c a type τ of the form $\forall (\bar{\alpha}) \tau_1 \rightarrow \dots \tau_n \rightarrow \tau_0$, where τ_0 is a constructed type whenever the constant c is a constructor.

We distinguish a subset of terms, called values and written v . Values are term abstractions, type abstractions, full or partial applications of constructors, or partial applications of primitives. We use an auxiliary letter w to characterize the arguments of functions, which differ for call-by-value and call-by-name strategies. In values, an application of a constant c can involve a series of type instantiations, but only evaluated ones and placed before all other arguments.

Moreover, the application may only be partial whenever c is a primitive. Evaluated instantiations θ may be quantifier eliminations or either inside or under (general) instantiations. In particular, $a\tau$ and $a(!\alpha)$ are *never* values. The grammar for values and evaluated instantiations is as follows:

$$\begin{aligned}
v &::= \lambda(x : \tau) a \\
&| A(\alpha : \tau) a \\
&| C \theta_1 \dots \theta_k w_1 \dots w_n & n \leq |C| \\
&| f \theta_1 \dots \theta_k w_1 \dots w_n & n < |f| \\
\theta &::= \forall(\geq \phi) \mid \forall(\alpha \geq) \phi \mid \&
\end{aligned}$$

Finally, we assume that δ -rules are of the form $f \theta_1 \dots \theta_k w_1 \dots w_{|f|} \longrightarrow_f a$ (that is, δ -rules may only reduce fully applied primitives).

In addition to this general setting, we make further assumptions to relate the static and dynamic semantics of constants.

SUBJECT REDUCTION: δ -reduction preserves typings, *i.e.*, for any typing context Γ such that $\Gamma \vdash a : \tau$ and $a \longrightarrow_f a'$, the judgment $\Gamma \vdash a' : \tau$ holds.

PROGRESS: Well-typed, full applications of primitives can be reduced, *i.e.*, for any term a of the form $f \theta_1 \dots \theta_k w_1 \dots w_n$ verifying $\Gamma_0 \vdash a : \tau$, there exists a term a' such that $a \longrightarrow_f a'$.

Call-by-value reduction We now specialize the previous framework to a call-by-value semantics. In this case, arguments of applications in values are themselves restricted to values, *i.e.* w is taken equal to v . Rules (β) and (β_{let}) are limited to the substitution of values, that is, to reductions of the form $(\lambda(x : \tau) a) v \longrightarrow a\{x \leftarrow v\}$ and $\text{let } x = v \text{ in } a \longrightarrow a\{x \leftarrow v\}$. Rules ι -ID, ι -COMP and ι -INTRO are also restricted so that they only apply to values (*e.g.* a is textually replaced by v in each of these rules). Finally, we restrict rule CONTEXT to call-by-value contexts, which are of the form

$$E_v ::= [\cdot] \mid E_v a \mid v E_v \mid E_v \phi \mid \text{let } x = E_v \text{ in } a$$

We write \longrightarrow_v the resulting reduction relation. It follows from the above restrictions that the reduction is deterministic. Moreover, since δ -reduction preserves typings, by assumption, the relation \longrightarrow_v also preserves typings by Theorem 1.

Progress holds for call-by-value. In combination with subject-reduction, this ensures that the evaluation of well-typed terms “cannot go wrong”.

Theorem 3. *If $\Gamma_0 \vdash a : \tau$, then either a is a value or $a \longrightarrow_v a'$ for some a' .*

Call-by-value reduction and the value restriction The value-restriction is the most standard way to add side effects in a call-by-value language. It is thus important to verify that it can be transposed to $x\text{MLF}$.

Typically, the *value restriction* amounts to restricting type generalization to non-expansive expressions, which contain at least value-forms, *i.e.* values and

term variables, as well as their type-instantiations. Hence, we obtain the following revised grammar for expansive expressions b and for non-expansive expressions u .

$$\begin{array}{l}
b ::= u \mid b b \mid \text{let } x = u \text{ in } b \\
u ::= x \mid \lambda(x : \tau) b \mid \Lambda(\alpha : \tau) u \mid u \phi \\
\quad \mid C \theta_1 \dots \theta_k u_1 \dots u_n \qquad n \leq |C| \\
\quad \mid f \theta_1 \dots \theta_k u_1 \dots u_n \qquad n < |f|
\end{array}$$

As usual, we restrict let-bound expressions to be non-expansive, since they implicitly contain a type generalization. Notice that, although type instantiations are restricted to non-expansive expressions, this is not a limitation: $b \phi$ can always be written as $(\lambda(x : \tau) x \phi) b$, where τ is the type of a , and similarly for applications of constants to expansive expressions.

Theorem 4. *Expansive and non-expansive expressions are closed by call-by-value reduction.*

Corollary 1. *Subject reduction holds with the value restriction.*

It is then routine work to extend the semantics with a global store to model side effects and verify type soundness for this extension.

Call-by-name reduction For call-by-name reduction semantics, we can actually increase the set of values, which may contain applications of constants to arbitrary expressions; that is, we take a for w . The ι -reduction is restricted as for call-by-value. However, evaluation contexts are now $E_n ::= [\cdot] \mid E_n a \mid E_n \phi$. We write \longrightarrow_n the resulting reduction relation. As for call-by-value, it is deterministic by definition and preserves typings. It may also always progress.

Theorem 5. *If $\Gamma_0 \vdash a : \tau$, then either a is a value or $a \longrightarrow_n a'$ for some a' .*

3 Discussion

Elaboration of graphical $e\text{ML}^F$ into $x\text{ML}^F$ To verify that, as expected, $x\text{ML}^F$ can be used as an internal language for $e\text{ML}^F$, we have exhibited a type-preserving type-erasure-preserving translation from $e\text{ML}^F$ to $x\text{ML}^F$. Technically, this translation is based on the presolutions of type inference problems in the graphic constraint framework of ML^F . An important corollary is the type soundness of $e\text{ML}^F$ —in its most expressive⁶ version (Rémy and Yakobowski 2008b). By lack of space, this translation is however not presented in this paper, but can be found in (Rémy and Yakobowski 2008a). We also expect that $x\text{ML}^F$ could be used as an internal language for **HML**, another less expressive but simpler surface language for $i\text{ML}^F$ that has been recently proposed (Leijen 2009).

⁶ So far, type-soundness has only been proved for the original, but slightly weaker variant of ML^F (Le Botlan 2004) and for the shallow, recast version of ML^F (Le Botlan and Rémy 2007).

Expressiveness of $x\text{ML}^F$ The translation of $e\text{ML}^F$ into $x\text{ML}^F$ shows that $x\text{ML}^F$ is at least as expressive as $e\text{ML}^F$. However, and perhaps surprisingly, the converse is not true. That is, there exist programs of $x\text{ML}^F$ that cannot be typed in ML^F . While, this is mostly irrelevant when using ML^F as an internal language for $e\text{ML}^F$, the question is still interesting from a theoretical point of view, as understanding $x\text{ML}^F$ on its own, *i.e.* independently of the type inference constraints of $e\text{ML}^F$, could perhaps suggest other useful extensions of $x\text{ML}^F$.

For the sake of simplicity, we explain the difference between $x\text{ML}^F$ and $i\text{ML}^F$, the Curry-style version of ML^F (which has the same expressiveness as $e\text{ML}^F$). Although syntactically identical, the types of $x\text{ML}^F$ and of syntactic $i\text{ML}^F$ differ in their interpretation of alias bounds, *i.e.* quantifications of the form $\forall (\beta \geq \alpha) \tau$. Consider, for example, the two types τ_0 and τ_{id} defined as $\forall (\alpha \geq \tau) \forall (\beta \geq \alpha) \beta \rightarrow \alpha$ and $\forall (\alpha \geq \tau) \alpha \rightarrow \alpha$. In $i\text{ML}^F$, alias bounds can be expanded and τ_0 and τ_{id} are equivalent. Roughly, the set of their instances (stripped of toplevel quantifiers) is $\{\tau' \rightarrow \tau' \mid \tau \leq \tau'\}$. In contrast, the set of instances of τ_0 is larger in $x\text{ML}^F$ and at least a superset of $\{\tau'' \rightarrow \tau' \mid \tau \leq \tau' \leq \tau''\}$. This level of generality cannot be expressed in $i\text{ML}^F$.

The current treatment of alias bounds in $x\text{ML}^F$ is quite natural in a Church-style presentation. Surprisingly, it is also simpler than treating them as in $e\text{ML}^F$. A restriction of $x\text{ML}^F$ without alias bounds that is closed under reduction and in closer correspondence with $i\text{ML}^F$ can still be defined a posteriori, by constraining the formation of terms, but the definition is contrived and unnatural. Instead of restricting $x\text{ML}^F$ to match the expressiveness of $i\text{ML}^F$, a question worth further investigation is whether the treatment of alias bounds could be enhanced in $i\text{ML}^F$ and $e\text{ML}^F$ to match the one in $x\text{ML}^F$ without compromising type inference.

Related works A strong difference between $e\text{ML}^F$ and $x\text{ML}^F$ is the use of explicit coercions to trace the derivation of type instantiation judgments. A similar approach has already been used in a language with subtyping and intersection types, proposed as a target for the compilation of bounded polymorphism (Crary 2000). In both cases, coercions are used to make typechecking a trivial process. In our case, they are also exploited to make subject reduction easy—by introducing the language to describe how type instance derivations must be transformed during reduction. (We believe that the use of explicit coercions for simplifying subject-reduction proofs has been neglected.) In both approaches, reduction is split into a standard notion of β -reduction and a new form of reduction (which we call ι -reduction) that only deals with coercions, preserves type-erasures, and is (conjectured to be) strongly normalizing. There are also important differences. While both coercion languages have common forms, our coercions intend to keep the instance-bounded polymorphism form $\forall (\alpha \geq \tau) \tau'$. On the opposite, coercions are used to eliminate the subtype-bounded polymorphism form $\forall (\alpha \leq \tau) \tau'$ in (Crary 2000), using intersection types and contravariant arrow coercions instead, which we do not need. It would be worth checking whether union types, which are proposed as an extension in (Crary 2000), could be used to encode away our instance-bounded polymorphism form.

Besides this work and the several papers that describe variants of ML^F , there are actually few other related works. Both Leijen and Löh (2005) and Leijen (2007) have studied the extension of ML^F with qualified types, and as a subcase, the translation of ML^F without qualified types into System F. However, in order to handle type instantiations, a term a of type $\forall(\alpha \geq \tau') \tau$ is elaborated as a function of type $\forall(\alpha) (\tau'_* \rightarrow \alpha) \rightarrow \tau_*$, where τ_* is a runtime representation of τ . The first argument is a *runtime coercion*, which bears strong similarities with our instantiations. However, an important difference is that their coercions are at the level of terms, while our instantiations are at the level of types. In particular, although coercion functions should not change the semantics, this critical result has not been proved so far, while in our settings the type-erasure semantics comes for free by construction. The impact of coercion functions in a call-by-value language with side effects is also unclear. Perhaps, a closer connection between their coercion functions and our instantiations could be established and used to actually prove that their coercions do not alter the semantics. However, even if such a result could be proved, coercions should preferably remain at the type level, as in our setting, than be intermixed with terms, as in their proposal.

Future works The demand for an internal language for ML^F was first made in the context of using the $e\text{ML}^F$ type system for the Haskell language. We expect $x\text{ML}^F$ to better accommodate qualified types than $e\text{ML}^F$ since no evidence function would be needed for flexible polymorphism, but it remains to be verified.

Type instantiation, which changes the type of an expression without changing its meaning, goes far beyond type application in System F and resembles retyping functions in System F^η —the closure of F by η -conversion (Mitchell 1988). Those functions can be seen either at the level of terms, as expressions of System F that $\beta\eta$ -reduces to the identity, or at the level of types as a *type conversion*. Some loose parallel can be made between the encoding of ML^F in System F by Leijen and Löh (2005) which uses term-level coercions, and $x\text{ML}^F$ which uses type-level instantiations. Additionally, perhaps F^η could be extended with a form of abstraction over retyping functions, much as type abstraction $\forall(\alpha \geq \tau)$ in $x\text{ML}^F$ amounts to abstracting over the instantiation $!\alpha$ of type $\tau \rightarrow \alpha$. (Or perhaps, as suggested by the work of Cray (2000), intersection and union types could be added to F^η to avoid the need for abstracting over coercion functions.)

Regarding type soundness, it is also worth noticing that the proof of subject reduction in $x\text{ML}^F$ does not subsume, but complements, the one in the original presentation of ML^F . The latter does not explain how to transform type annotations, but shows that annotation sites need not be introduced (only transformed) during reduction. Because $x\text{ML}^F$ has full type information, it cannot say anything about type information that could be left implicit and inferred. Given a term in $x\text{ML}^F$, can we rebuild a term in $i\text{ML}^F$ with minimal type annotations? While this should be easy if we require that corresponding subterms have identical types in $x\text{ML}^F$ and $i\text{ML}^F$, the answer is unclear if we allow subterms to have different types.

The semantics of $x\text{ML}^F$ allows reduction (and elimination) of type instantiations $a\phi$ through ι -reduction but does not operate reduction (and simplification) of instantiations ϕ alone. It would be possible to define a notion of reduction on instantiations $\phi \rightarrow \phi'$ (such that, for instance, $\forall(\geq \phi_1; \phi_2) \rightarrow \forall(\geq \phi_1); \forall(\geq \phi_2)$, or conversely?) and extend the reduction of terms with a context rule $a\phi \rightarrow a\phi'$ whenever $\phi \rightarrow \phi'$. This might be interesting for more economical representations of instantiation. However, it is unclear whether there exists an interesting form of reduction that is both Church-Rosser and large enough for optimization purposes. Perhaps, one should rather consider instantiation transformations that preserve observational equivalence, which would leave more freedom in the way one instantiation could be replaced by another.

Extending $x\text{ML}^F$ to allow higher-order polymorphism is another interesting research direction for the future. Such an extension is already under investigation for the type inference version $e\text{ML}^F$ (Herms 2009).

Conclusion We have completed the ML^F trilogy by introducing the Church-style version $x\text{ML}^F$, that was still desperately missing for type-aware compilation and from a theoretical point of view. The original type-inference version $e\text{ML}^F$, which requires partial type annotations but does not tell how to track them during reduction, now lies between the Curry-style presentation $i\text{ML}^F$ that ignores all type information and $x\text{ML}^F$ that maintains it during reduction. We have shown that $x\text{ML}^F$ is well-behaved: reduction preserves well-typedness, and the calculus is sound for both call-by-value and call-by-name semantics.

Hence, $x\text{ML}^F$ can be used as an internal language for ML^F , with either call-by-value or call-by-name semantics, and also for the many restrictions of ML^F that have been proposed, including HML. Indeed, the translation of partially typed $e\text{ML}^F$ programs into fully typed $x\text{ML}^F$ ones, presented in (Rémy and Yakobowski 2008a), preserves well-typedness and the type erasure of terms, and therefore ensures the type soundness of $e\text{ML}^F$. Hopefully, this will help the adoption of ML^F and maintain a powerful form of type inference in modern programming languages that will necessarily feature first-class polymorphism.

Independently, the idea of enriching type applications to richer forms of type transformations might also be useful in other contexts.

References

- Henk P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North-Holland, 1984. ISBN 0-444-86748-1.
- Karl Crary. Typed compilation of inclusive subtyping. In *ICFP '00: Proceedings of the fifth ACM SIGPLAN international conference on Functional programming*, pages 68–81, New York, NY, USA, 2000. ACM. ISBN 1-58113-202-6. doi: <http://doi.acm.org/10.1145/351240.351247>.
- Paolo Herms. Partial Type Inference with Higher-Order Types. Master's thesis, University of Pisa and INRIA, 2009. To appear.

- Didier Le Botlan. *MLF : An extension of ML with second-order polymorphism and implicit instantiation*. PhD thesis, Ecole Polytechnique, June 2004. english version.
- Didier Le Botlan and Didier Rémy. MLF: Raising ML to the power of System F. In *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming*, pages 27–38, August 2003.
- Didier Le Botlan and Didier Rémy. Recasting MLF. Research Report 6228, INRIA, Rocquencourt, BP 105, 78 153 Le Chesnay Cedex, France, June 2007.
- Daan Leijen. A type directed translation of MLF to System F. In *The International Conference on Functional Programming (ICFP'07)*. ACM Press, October 2007.
- Daan Leijen. Flexible types: robust type inference for first-class polymorphism. In *Proceedings of the 36th annual ACM Symposium on Principles of Programming Languages (POPL'09)*, pages 66–77, New York, NY, USA, 2009. ACM.
- Daan Leijen and Andres Löh. Qualified types for MLF. In *ICFP '05: Proceedings of the tenth ACM SIGPLAN international conference on Functional programming*, pages 144–155, New York, NY, USA, September 2005. ACM Press. ISBN 1-59593-064-7.
- John C. Mitchell. Polymorphic type inference and containment. *Information and Computation*, 2/3(76):211–249, 1988.
- Simon Peyton Jones. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, May 2003. ISBN 0521826144.
- Didier Rémy and Boris Yakobowski. A church-style intermediate language for MLF (extended version). Available at <http://gallium.inria.fr/~remy/mlf/xmlf.pdf>, September 2008a.
- Didier Rémy and Boris Yakobowski. From ML to MLF: Graphic type constraints with efficient type inference. In *The 13th ACM SIGPLAN International Conference on Functional Programming (ICFP'08)*, pages 63–74, Victoria, BC, Canada, September 2008b.
- Boris Yakobowski. *Graphical types and constraints: second-order polymorphism and inference*. PhD thesis, University of Paris 7, December 2008.

A Appendix: definition of $a\{\!\alpha_0 \leftarrow \phi_0\}$

Formally, $a\theta$ where θ is $\{\!\alpha_0 \leftarrow \phi_0\}$ is defined recursively as follows (we assume $\alpha \neq \alpha_0$). The interesting lines are the two first ones of the second column; other lines are just lifting the substitution from the leaves to types, type instantiations,

and terms in the usual way.

Types

$$\tau \theta = \tau$$

Terms

$$x \theta = x$$

$$(a_1 a_1) \theta = (a_1 \theta) (a_1 \theta)$$

$$(a \phi) \theta = a (\phi \theta)$$

$$(\lambda(x : \tau) a) \theta = \lambda(x : \tau \theta) a \theta$$

$$(\Lambda(\alpha \geq \tau) a) \theta = \Lambda(\alpha : \tau \theta) a \theta$$

Type instantiations

$$!\alpha \theta = !\alpha$$

$$!\alpha_0 \theta = \phi_0$$

$$(\tau) \theta = (\tau \theta)$$

$$(\forall (\geq \phi)) \theta = \forall (\geq \phi \theta)$$

$$(\forall (\alpha \geq) \phi) \theta = \forall (\alpha \geq) (\phi \theta)$$

$$(\phi; \phi') \theta = (\phi \theta); (\phi' \theta)$$

$$\& \theta = \&$$

$$\wp \theta = \wp$$

$$\mathbb{1} \theta = \mathbb{1}$$