

A type directed translation of MLF to System F

Daan Leijen

Microsoft Research
daan@microsoft.com

Abstract

The MLF type system by Le Botlan and Rémy (2003) is a natural extension of Hindley-Milner type inference that supports full first-class polymorphism, where types can be of higher-rank and impredicatively instantiated. Even though MLF is theoretically very attractive, it has not seen widespread adoption. We believe that this partly because it is unclear how the rich language of MLF types relate to standard System F types. In this article we give the first type directed translation of MLF terms to System F terms. Based on insight gained from this translation, we also define “Rigid MLF” ($\text{MLF}^=$), a restriction of MLF where all bound values have a System F type. The expressiveness of $\text{MLF}^=$ is the same as that of boxy types, but $\text{MLF}^=$ needs fewer annotations and we give a detailed comparison between them.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features—Polymorphism

General Terms Languages, Design, Theory

Keywords First-class polymorphism, System F, MLF

1. Introduction

The MLF type system by Le Botlan and Rémy (2003) is a natural extension of Hindley-Milner type inference that supports full first-class polymorphism, where any value can have a polymorphic type. MLF has a strong theoretical foundation and requires very few type annotations. Even though MLF is very attractive for these reasons, it has not seen widespread adoption. We believe that this is partly because the type language of MLF is richer than that of System F. Besides that this is harder on the programmer, it is also harder to write a compiler for a language based on MLF: many compilers use an intermediate language based on System F, and the relation between MLF terms and System F terms far from clear.

In an attempt to remedy this situation, Leijen and Löh (2005) described a type inference algorithm that returns System F terms for well-typed MLF terms. Unfortunately, such algorithm does not give particular *insight* into the relation between MLF type rules and the translated System F terms. Furthermore, we believe that one of the reasons that a type directed translation was not given in the previous article is that it is surprisingly subtle to do so – indeed, this is the main technical contribution of this paper. Specifically:

- We are the first to give a type directed translation from MLF to System F terms, and we prove that the translation is sound. The translation of MLF types that go ‘beyond System F’ (namely flexibly bound polymorphic types) is done using evidence translation.
- We show that a translation is possible where no evidence is passed for polymorphic types with rigid bounds, which correspond naturally to inlined System F types. This greatly simplifies the translated System F terms, but we need to define the translation on slightly modified type rules. We show that the modified type rules are sound and complete with respect to the original MLF type rules.
- It is widely believed that a System F translation of impredicative types may need to traverse structures at runtime to apply coercion terms. We show conclusively that this is not the case for MLF and that there exists a concise translation that builds coerced terms directly without runtime traversals.
- The type directed System F translation leads naturally to the definition of a restriction of MLF, called “Rigid MLF” ($\text{MLF}^=$), where all bound values have a standard System F type. We show that $\text{MLF}^=$ has the same expressiveness as boxy type inference (Vytiniotis et al. 2006), but needs fewer annotations. Since $\text{MLF}^=$ hides MLF types, it can work well as a simplified version of MLF where the programmer only works with regular System F types.

In the next section, we start with an introduction to MLF and explain the MLF type rules. Section 3 first gives an overview of the difficulties of translating MLF to System F, and introduces a modified set of type rules on which we define a type directed System F translation. Section 4 defines $\text{MLF}^=$, a restriction of MLF where all values have System F types. Finally, in the related work section (Section 5) we give a detailed comparison between Hindley-Milner, MLF, $\text{MLF}^=$, and boxy types.

2. An introduction to MLF

Functional languages with type inference are almost always based on the Hindley-Milner type system (Hindley 1969; Milner 1978) – and for good reasons. Type inference based on Hindley-Milner can automatically infer most general, or *principal*, types for expressions without further type annotations. Also, the type system is *sound*, and well-typed programs cannot “go wrong”.

To achieve automatic type inference, the Hindley-Milner type system restricts polymorphism and polymorphic values are not first-class citizens. In particular, function arguments can only be monomorphic. Formally, this means that universal quantifiers can only appear at the outermost level (i.e. higher-ranked types are disallowed), and quantified variables can only be instantiated with monomorphic types (i.e. impredicative types are disallowed).

The MLF type system by LeBotlan and Remy (Le Botlan 2004; Le Botlan and Rémy 2003) is a natural extension of Hindley-Milner

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP’07, October 1–3, 2007, Freiburg, Germany.
Copyright © 2007 ACM 978-1-59593-815-2/07/0010...\$5.00

that lifts these restrictions and supports first-class polymorphism with higher-ranked and impredicative types. In contrast to Hindley-Milner, some type annotations are required in MLF to avoid guessing polymorphic types. Take for example the following program:

$$\text{poly } f = (f \ 1, f \ \text{True})$$

This program is rejected by both MLF and Hindley-Milner. The program would be accepted in MLF if we annotate the argument f with a polymorphic type, like $\forall\alpha. \alpha \rightarrow \alpha$ or $\forall\alpha. \alpha \rightarrow \text{Int}$. The type of a polymorphic argument like f cannot be inferred automatically since there exist many types for f none of which are an instance of the other – there is no principal type.

The reader might be worried that many type annotations are needed, but such is not the case. MLF has a remarkable property that the *only* required type annotations are on *arguments that are used polymorphically*. For example,

$$\text{poly } (f :: \forall\alpha. \alpha \rightarrow \alpha) = (f \ 1, f \ \text{True})$$

is a well-typed MLF program. Note that the annotation rule implies that all programs accepted by Hindley-Milner are also accepted by MLF. Moreover, the annotation rule only requires annotations on polymorphic arguments that are also *used* polymorphically and not merely ‘passed through’. For example:

$$\begin{aligned} \text{polyL} &:: [\forall\alpha. \alpha \rightarrow \alpha] \rightarrow (\text{Int}, \text{Bool}) \quad \text{-- inferred} \\ \text{polyL } xs &= \text{poly } (\text{head } xs) \end{aligned}$$

is accepted without type annotations even though xs has an (inferred) impredicative polymorphic type. We consider this an important property for abstraction and modularity, since we can (re)use standard functions like head , with type $\forall\alpha. [\alpha] \rightarrow \alpha$, for lists that contain polymorphic values. As another example, consider the application function:

$$\begin{aligned} \text{apply} &:: \forall\alpha\beta. (\alpha \rightarrow \beta) \rightarrow \alpha \rightarrow \beta \quad \text{-- inferred} \\ \text{apply } f \ x &= f \ x \end{aligned}$$

Of course, the direct application $\text{poly } id$ is well typed in MLF, but also the abstracted application $\text{apply } \text{poly } id$. No type annotations are needed, except for the argument annotation of poly . Note that to type check $\text{apply } \text{poly } id$, the α quantifier in the type of apply is impredicatively instantiated to the polymorphic type of id , namely $\forall\alpha. \alpha \rightarrow \alpha$. More generally, if an application $e_1 \ e_2$ is well typed in MLF, then the expression $\text{apply } e_1 \ e_2$ is also well typed. In general, this applies to functors (apply) applying polymorphic functions (poly) over structures holding polymorphic values (id). For example, we can use the standard map function to apply the poly function to the elements of a list of identity functions, as in $\text{map } \text{poly } [id, id]$, without any further type annotations.

2.1 Bounded types

Essential to type inference is the ability to assign principal types to expressions. In the presence of impredicative polymorphism however, we need more than standard System F types to have principal types. Consider the following program

$$\text{let } ids = [id] \text{ in } (\text{polyL } ids, ids \ ++ [inc])$$

where inc has type $\text{Int} \rightarrow \text{Int}$ and the append function ($++$) has type $\forall\alpha. [\alpha] \rightarrow [\alpha] \rightarrow [\alpha]$. In a setting with first-class polymorphism we can assign two types to ids , namely a list of polymorphic identity functions, $[\forall\alpha. \alpha \rightarrow \alpha]$, or a polymorphic list of monomorphic identity functions, $\forall\alpha. [\alpha \rightarrow \alpha]$. Indeed, in the body of the let expression, we use ids with both types: the polyL function requires the first type, while appending to an $[\text{Int} \rightarrow \text{Int}]$ list requires the second type. Unfortunately, neither of these types is an instance of the other – so what type should we give to ids to make the program well typed?

Monomorphic types		
$\tau ::= \alpha$		type variable
$\quad \ c \ \tau_1 \ \dots \ \tau_n$		constructor application
Polymorphic types		
$\sigma ::= q.\sigma$		quantified type
$\quad \ \tau$		mono type
$\quad \ \perp$		most polymorphic type
Quantifier		
$q ::= \forall(\alpha \diamond \sigma)$		a bounded quantifier
$\diamond ::= \geq \mid =$		a bound is flexible or rigid
Prefix		
$Q ::= q_1, \dots, q_n$		a prefix is a list of quantifiers
Syntactic sugar		
$\forall\alpha = \forall(\alpha \geq \perp)$		
$Q.\tau = q_1. \dots . q_n.\tau$		quantify using a prefix

Figure 1. MLF types.

MLF solves this dilemma by going beyond standard System F types and assigns a type that can be instantiated to both of the previous types. The MLF type for ids is $\forall(\beta \geq \forall\alpha. \alpha \rightarrow \alpha). [\beta]$, which is read as “a list of β , for all types β that are an instance of (or equal to) $\forall\alpha. \alpha \rightarrow \alpha$ ”. As we will see later, both of our previous types are instances of this type:

$$\begin{aligned} \forall(\beta \geq \forall\alpha. \alpha \rightarrow \alpha). [\beta] &\sqsubseteq [\forall\alpha. \alpha \rightarrow \alpha] \\ \forall(\beta \geq \forall\alpha. \alpha \rightarrow \alpha). [\beta] &\sqsubseteq \forall\alpha. [\alpha \rightarrow \alpha] \end{aligned}$$

where $\sigma_1 \sqsubseteq \sigma_2$ states that σ_2 is an instance of σ_1 . We call \geq a *flexible* bound. For regularity, MLF also introduces *rigid* bounds, written as $=$. The type $\forall(\beta = \forall\alpha. \alpha \rightarrow \alpha). [\beta]$ is read as “a list of β , where β is as polymorphic as $\forall\alpha. \alpha \rightarrow \alpha$ ”. Indeed, the type $[\forall\alpha. \alpha \rightarrow \alpha]$ is just a syntactic shorthand for the previous type and we generally inline rigid bounds for notational convenience.

2.2 MLF types formally

Figure 1 defines the grammar of MLF types. Monomorphic types τ are either a type variable α , or a type constructor application. We assume that the function constructor \rightarrow is part of the type constructors and we do not need to treat it specially. Polymorphic types σ are either a monomorphic type, the most polymorphic type \perp , or a polymorphic type quantified with a bound $\forall(\alpha \diamond \sigma_1). \sigma_2$, where the bound \diamond is either flexible (\geq) or rigid ($=$).

A quantifier with a flexible bound can be instantiated to any instance of its bound. In particular, a quantifier $\forall(\alpha \geq \perp)$ can be instantiated to any type since \perp is the most polymorphic type. For example, the full MLF type for the identity function is:

$$id :: \forall(\alpha \geq \perp). \alpha \rightarrow \alpha$$

We call $\forall(\alpha \geq \perp)$ an *unconstrained bound* and usually shorten it to $\forall\alpha$. The \perp type is defined as equivalent to $\forall(\alpha \geq \perp). \alpha$. Note that in Hindley-Milner all quantifiers are always unconstrained.

A list of quantifiers is called a *prefix* and denoted as Q . We always assume that the quantified variables are distinct and form the domain of Q , written as $\text{dom}(Q)$. If all quantifiers are unconstrained, we call Q an *unconstrained prefix*.

2.3 Type rules of MLF

The type rules for MLF are given in Figure 2. Given the expressiveness of MLF, they are surprisingly simple and very similar to the type rules of Hindley-Milner. A derivation $(Q) \Gamma \vdash e : \sigma$ means that under a prefix Q and type environment Γ the expression e has

VAR	$\frac{x : \sigma \in \Gamma}{(Q) \Gamma \vdash x : \sigma}$
APP	$\frac{(Q) \Gamma \vdash e_1 : \tau_2 \rightarrow \tau \quad (Q) \Gamma \vdash e_2 : \tau_2}{(Q) \Gamma \vdash e_1 e_2 : \tau}$
FUN	$\frac{(Q) \Gamma, x : \tau_1 \vdash e : \tau_2}{(Q) \Gamma \vdash \lambda x. e : \tau_1 \rightarrow \tau_2}$
LET	$\frac{(Q) \Gamma \vdash e_1 : \sigma_1 \quad (Q) \Gamma, x : \sigma_1 \vdash e_2 : \sigma_2}{(Q) \Gamma \vdash \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 : \sigma_2}$
GEN	$\frac{(Q, \alpha \diamond \sigma_1) \Gamma \vdash e : \sigma_2 \quad \alpha \notin \mathbf{ftv}(\Gamma)}{(Q) \Gamma \vdash e : \forall(\alpha \diamond \sigma_1). \sigma_2}$
INST	$\frac{(Q) \Gamma \vdash e : \sigma_1 \quad (Q) \sigma_1 \sqsubseteq \sigma_2}{(Q) \Gamma \vdash e : \sigma_2}$
ANN	$\frac{(Q) \Gamma \vdash e : \sigma_1 \quad (Q) \sigma \sqsubseteq \sigma_1}{(Q) \Gamma \vdash (e :: \sigma) : \sigma}$

Figure 2. Type rules

type σ . The type environment Γ contains the types of all the free *term* variables in e . The prefix Q contains the bounds of all the free *type* variables in Γ , σ , and e ¹. In Hindley-Milner, the prefix Q is usually left implicit since all the bounds of the free type variables are always unconstrained.

The application rule APP applies a function to a monomorphic argument. Remarkably, this is still expressive enough to apply functions to polymorphic arguments since the prefix Q can contain a polymorphic bound of a (monomorphic) type variable. As an example, we derive the type for the application *poly id*, where σ_{id} stands for the type $\forall \alpha. \alpha \rightarrow \alpha$, τ for $(Int, Bool)$, and $\sigma_{id} \rightarrow \tau$ as a shorthand for $\forall(\beta = \sigma_{id}). \beta \rightarrow \tau$.

$$\begin{array}{c}
\frac{\text{poly} : \sigma_{id} \rightarrow \tau \in \Gamma \quad \text{id} : \sigma_{id} \in \Gamma}{(\forall \alpha = \sigma_{id}) \Gamma \vdash \text{poly} : \sigma_{id} \rightarrow \tau \quad (\forall \alpha = \sigma_{id}) \Gamma \vdash \text{id} : \sigma_{id}} \\
\frac{(\forall \alpha = \sigma_{id}) \sigma_{id} \rightarrow \tau \sqsubseteq \alpha \rightarrow \tau \quad (\forall \alpha = \sigma_{id}) \sigma_{id} \sqsubseteq \alpha}{(\forall \alpha = \sigma_{id}) \Gamma \vdash \text{poly} : \alpha \rightarrow \tau \quad (\forall \alpha = \sigma_{id}) \Gamma \vdash \text{id} : \alpha} \\
\frac{(\forall \alpha = \sigma_{id}) \Gamma \vdash \text{poly id} : \tau}{() \Gamma \vdash \text{poly id} : \forall(\alpha = \sigma_{id}). \tau} \\
\frac{() \Gamma \vdash \text{poly id} : \forall(\alpha = \sigma_{id}). \tau}{() \Gamma \vdash \text{poly id} : (Int, Bool)}
\end{array}$$

In the above derivation it is essential that we can instantiate σ_{id} to α under the prefix $(\forall \alpha = \sigma_{id})$. The instantiation rule INST is used twice to instantiate both the argument type of *poly* and *id* to α after which can we use the application rule APP. Finally, the generalization rule GEN is used to remove the assumption from the prefix, and instantiation is applied to remove the (now) dead quantifier, deriving $(Int, Bool)$.

The annotation rule ANN defines the type rule for type annotations, using the abstraction relation \sqsubseteq . We define this relation (together with instantiation) in the next section. Note that MLF has just a single rule for type annotations, and there is no need for special type propagation rules. In particular, a lambda bound annotation $\lambda(x :: \sigma). e$ is syntactic sugar for

$$\lambda x. \mathbf{let} \ x = (x :: \sigma) \ \mathbf{in} \ e$$

As such, the essence of MLF is contained in the instance relation (\sqsubseteq) in rule INST, and the abstraction relation (\sqsubseteq) in the annotation

¹ due to annotations.

R-TRANS	$\frac{(Q) \sigma_1 \oplus \sigma_2 \quad (Q) \sigma_2 \oplus \sigma_3}{(Q) \sigma_1 \oplus \sigma_3}$
R-PREFIX	$\frac{(Q, \alpha \diamond \sigma) \sigma_1 \oplus \sigma_2 \quad \alpha \notin \mathbf{dom}(Q)}{(Q) \forall(\alpha \diamond \sigma). \sigma_1 \oplus \forall(\alpha \diamond \sigma). \sigma_2}$

Figure 3. General rules, where \oplus stands for \equiv (equivalence), \sqsubseteq (abstraction), or \sqsubseteq (instantiation).

EQ-REFL	$(Q) \sigma \equiv \sigma$
EQ-VAR	$(Q) \forall(\alpha \diamond \sigma). \alpha \equiv \sigma$
EQ-FREE	$\frac{\alpha \notin \mathbf{ftv}(\sigma_2)}{(Q) \forall(\alpha \diamond \sigma_1). \sigma_2 \equiv \sigma_2}$
EQ-MONO	$\frac{\forall(\alpha \diamond \sigma_0) \in Q \quad (Q) \sigma_0 \equiv \tau_0}{(Q) \tau \equiv \tau[\alpha \mapsto \tau_0]}$
EQ-COMM	$\frac{\alpha_1 \neq \alpha_2 \quad \alpha_1 \notin \mathbf{ftv}(\sigma_2) \quad \alpha_2 \notin \mathbf{ftv}(\sigma_1)}{(Q) \forall(\alpha_1 \diamond \sigma_1)(\alpha_2 \diamond \sigma_2). \sigma \equiv \forall(\alpha_2 \diamond \sigma_2)(\alpha_1 \diamond \sigma_1). \sigma}$
EQ-CONTEXT	$\frac{(Q) \sigma_1 \equiv \sigma_2}{(Q) \forall(\alpha \diamond \sigma_1). \sigma \equiv \forall(\alpha \diamond \sigma_2). \sigma}$

Figure 4. Equivalence

A-EQUIV	$\frac{(Q) \sigma_1 \equiv \sigma_2}{(Q) \sigma_1 \sqsubseteq \sigma_2}$
A-HYP	$\frac{\forall(\alpha = \sigma) \in Q}{(Q) \sigma \sqsubseteq \alpha}$
A-CONTEXT	$\frac{(Q) \sigma_1 \sqsubseteq \sigma_2}{(Q) \forall(\alpha = \sigma_1). \sigma \sqsubseteq \forall(\alpha = \sigma_2). \sigma}$

Figure 5. Abstraction

I-BOTTOM	$(Q) \perp \sqsubseteq \sigma$
I-ABSTRACT	$\frac{(Q) \sigma_1 \sqsubseteq \sigma_2}{(Q) \sigma_1 \sqsubseteq \sigma_2}$
I-HYP	$\frac{\forall(\alpha \geq \sigma) \in Q}{(Q) \sigma \sqsubseteq \alpha}$
I-RIGID	$(Q) \forall(\alpha \geq \sigma_1). \sigma_2 \sqsubseteq \forall(\alpha = \sigma_1). \sigma_2$
I-CONTEXT	$\frac{(Q) \sigma_1 \sqsubseteq \sigma_2}{(Q) \forall(\alpha \geq \sigma_1). \sigma \sqsubseteq \forall(\alpha \geq \sigma_2). \sigma}$

Figure 6. Instance

rule ANN. In the next section, we formally define these relations and explain the annotation rule in detail.

2.4 Instantiation, abstraction, and equivalence

Figure 6 formally defines the *instance* relation (\sqsubseteq). Through rule I-ABSTRACT, the instance relation includes the *abstraction* relation (\sqsubseteq) defined in Figure 5. Similarly, through rule A-EQUIV, the

$$\begin{aligned}
\text{nf}(\tau) &= \tau \\
\text{nf}(\perp) &= \perp \\
\text{nf}(\forall(\alpha \diamond \sigma_1). \sigma_2) &= \text{nf}(\sigma_2) && \text{iff } \alpha \notin \text{ftv}(\sigma_2) \\
\text{nf}(\forall(\alpha \diamond \sigma_1). \sigma_2) &= \text{nf}(\sigma_1) && \text{iff } \text{nf}(\sigma_2) = \alpha \\
\text{nf}(\forall(\alpha \diamond \sigma_1). \sigma_2) &= \text{nf}(\sigma_2[\alpha \mapsto \tau]) && \text{iff } \text{nf}(\sigma_1) = \tau \\
\text{nf}(\forall(\alpha \diamond \sigma_1). \sigma_2) &= \forall(\alpha \diamond \text{nf}(\sigma_1)). \text{nf}(\sigma_2)
\end{aligned}$$

Figure 7. Normal form

abstraction relation includes the *equivalence* relation (\equiv) defined in Figure 4 and we have the following inclusion: $(\sqsubseteq) \subseteq (\Xi) \subseteq (\equiv)$.

All three relations also include the general rules given in Figure 3. That is, all relations are transitive through R-TRANS, and all rules can be applied under a common prefix through rule R-PREFIX. Note that the condition $\alpha \notin \text{dom}(Q)$ in R-PREFIX can always be satisfied through alpha renaming.

2.4.1 Equivalence

The equivalence relation in Figure 4 defines an equivalence on types that abstracts from syntactical artifacts, like the order of the quantifiers. For example, rule EQ-FREE states that unbound quantifiers are not meaningful, while rule EQ-COMM allows rearrangement of independent quantifiers. Rule EQ-MONO is the only rule that reads information from the prefix, and inlines monomorphic bounds. For example, assuming $Q' = (\forall(\alpha \diamond \text{Int}), \sigma)$, we can derive:

$$\frac{\frac{\forall(\alpha \diamond \text{Int}) \in (QQ') \quad (QQ') \text{ Int} \equiv \text{Int}}{(Q, \forall(\alpha \diamond \text{Int})) \alpha \rightarrow \alpha \equiv \text{Int} \rightarrow \text{Int}} \quad \alpha \notin \text{dom}(Q)}{(Q) \forall(\alpha \diamond \text{Int}). \alpha \rightarrow \alpha \equiv \forall(\alpha \diamond \text{Int}). \text{Int} \rightarrow \text{Int}}$$

Using transitivity R-TRANS and rule EQ-FREE, we can now conclude that $\forall(\alpha \diamond \text{Int}). \alpha \rightarrow \alpha$ is equivalent to $\text{Int} \rightarrow \text{Int}$. For monomorphic types it does not matter whether the bound is flexible or rigid.

The context rule EQ-CONTEXT states we can substitute equivalent types under any bound (rigid or flexible). For example, we can derive $\forall(\beta = \forall(\alpha \geq \sigma_{id}). \alpha). [\beta] \equiv \forall(\beta = \sigma_{id}). [\beta]$ by the context rule and EQ-VAR.

2.4.2 Normal form

To abstract away from syntactical artifacts, Figure 7 defines a *normal form* function $\text{nf}(\cdot)$ that maps equivalent types to the same normal form up to rearrangement of their quantifiers (Le Botlan 2004). The normal form of a type is always equivalent to that type, i.e. $\text{nf}(\sigma) \equiv \sigma$.

2.4.3 Abstraction

The abstraction relation in Figure 5 extends the equivalence relation. In particular, it can read information from the prefix by the abstraction hypothesis rule A-HYP. If a quantifier α is rigidly bound to a type σ in the prefix, we can *abstract* a type σ by α . Note that abstraction is not symmetric, and the opposite direction is not allowed. This is essential for type inference since we can infer that a given polymorphic type can be abstracted but we cannot infer the other direction.

The abstraction relation nicely interacts with the type rules for lambda expression (FUN) and type annotations (ANN) in Figure 2. The FUN rule requires the lambda bound argument x to have a *monomorphic* type, just like Hindley-Milner – but how are we able to type a function like *poly* that requires a polymorphic argument?

The key to typing such function is to give the argument x a monomorphic type α that is rigidly bound in the prefix to the polymorphic type $\forall(\alpha = \sigma)$. The polymorphic type of α is later

revealed by the type annotation through A-HYP. Take for example the desugared version of *poly*:

$$\text{poly} = \lambda f_0. \text{let } f = (f_0 :: \sigma_{id}) \text{ in } (f \ 1, f \ \text{True})$$

We assume that e stands for the body of the lambda expression. The typing derivation for *poly* uses generalization to introduce the polymorphic bound, and assigns a monomorphic type α to the type of the lambda bound argument f_0 :

$$\frac{\vdots \quad \frac{(Q, \forall(\alpha = \sigma_{id})) \Gamma, f_0 : \alpha \vdash e : (\text{Int}, \text{Bool})}{(Q, \forall(\alpha = \sigma_{id})) \Gamma \vdash \lambda f_0. e : \alpha \rightarrow (\text{Int}, \text{Bool})} \quad \alpha \notin \text{ftv}(\Gamma)}{(Q) \Gamma \vdash \lambda f_0. e : \forall(\alpha = \sigma_{id}) (\alpha \rightarrow (\text{Int}, \text{Bool}))}$$

Given this environment, we are now able to type the body of the lambda expression. In particular, we can use f with the polymorphic type σ_{id} in the body of the *let* since we can derive that the expression $(f_0 :: \sigma_{id})$ has type σ_{id} :

$$\frac{\frac{f_0 : \alpha \in (\Gamma, f_0 : \alpha) \quad \forall(\alpha = \sigma_{id}) \in (Q, \forall(\alpha = \sigma_{id}))}{(\dots) \Gamma, f_0 : \alpha \vdash f_0 : \alpha} \quad \frac{\forall(\alpha = \sigma_{id}) \in (Q, \forall(\alpha = \sigma_{id})) \quad \sigma_{id} \Xi \alpha}{(Q, \forall(\alpha = \sigma_{id})) \sigma_{id} \Xi \alpha}}{(Q, \forall(\alpha = \sigma_{id})) \Gamma, f_0 : \alpha \vdash (f_0 :: \sigma_{id}) : \sigma_{id}}$$

The rule A-HYP applies here directly, and the type annotation rule reveals the polymorphic type of α . Note that the annotation is required since we cannot instantiate rigidly bound type variables, i.e. we cannot derive $(\forall(\alpha = \sigma_{id})) \alpha \sqsubseteq \sigma_{id}$.

2.4.4 Instantiation

The instantiation relation in Figure 6 extends abstraction and equivalence through rule I-ABSTRACT. The rule I-BOTTOM states that the most polymorphic type \perp can be instantiated to any other type. While abstraction can use rigid bounds from the prefix, instantiation can use flexible bounds from the prefix through the hypothesis rule I-HYP. If the prefix contains a quantifier $\forall(\alpha \geq \sigma)$ we know that α is an instance of σ and therefore we can always safely instantiate a type σ to α .

The context rule I-CONTEXT works *only* under flexible bounds. The context rule works nicely with I-BOTTOM to do standard instantiation of unconstrained bounds. For example:

$$\begin{aligned}
&\forall \alpha. \alpha \rightarrow \alpha \\
&= \{ \text{sugar} \} \\
&\forall(\alpha \geq \perp). \alpha \rightarrow \alpha \\
&\quad \sqsubseteq \{ \text{I-CONTEXT, I-BOTTOM} \} \\
&\forall(\alpha \geq \text{Int}). \alpha \rightarrow \alpha \\
&\quad \sqsubseteq \{ \text{I-ABSTRACT, A-EQUIV, EQ-PREFIX, EQ-MONO, EQ-FREE} \} \\
&\text{Int} \rightarrow \text{Int}
\end{aligned}$$

3. From MLF to System F

In this section, we will develop a type directed translation from MLF terms to System F terms. Such translation is very useful in practice as compilers that use MLF type inference, can use System F terms as their intermediate language. More importantly, such translation gives us a better understanding of the expressiveness and operational behaviour of MLF programs. In particular,

- It is widely believed that a System F translation of impredicative types may need to traverse structures at runtime to apply coercion terms, where further instantiations can even lead to multiple traversals. The translation we develop in this section shows that this is not the case for MLF. Instead we show that we can always build the coerced terms directly in a single pass without any traversals.

- The System F translation leads us naturally to a restriction of MLF where all bound values have a System F type and where the System F translation is particularly efficient.
- Extending MLF with other type system extensions usually requires a good understanding of the corresponding evidence translation. For example, to use MLF as a basis for languages that support qualified types, like Haskell, it is important to have an evidence translation of predicates which fits seamlessly with the presented System F translation.

3.1 Translating flexible bounds

The main difficulty in giving a System F translation of MLF terms occurs when the MLF terms have a type that goes beyond System F, i.e. polymorphic types that are flexibly bound. Consider:

```
let ids = [id] in (polyL ids, ids ++ [inc])
```

where ids has type $\forall(\alpha \geq \sigma_{id}). [\alpha]$. Unfortunately, there is no direct System F equivalent of this type, and in a naïve System F translation, ids can have two different System F types, namely $\forall\alpha \cdot [\alpha \rightarrow \alpha]$ or $[\forall\alpha \cdot \alpha \rightarrow \alpha]$. The first type is required by the append, while the second type is required by the *polyL* application. Effectively, MLF delays the instantiation and later ids is instantiated in two fundamentally different ways. In contrast, in System F this choice has to be made up front.

A potential way of translating this program to System F is to assign the type $[\forall\alpha \cdot \alpha \rightarrow \alpha]$ to ids , and explicitly coerce the type by mapping a coercion function over all element types that instantiates them $Int \rightarrow Int$ functions. In practice though this solution is unacceptable since the type inferencer might suddenly introduce coercion terms that need to traverse arbitrary structures at runtime! This is often seen as one of the fatal properties of impredicative type systems (Peyton Jones et al. 2007) and boxy type inference (Vytiniotis et al. 2006) was specifically designed to avoid this behaviour by rejecting the above program.

Fortunately, there exists a very elegant solution for MLF, first described by Leijen and Löh (2005): for every non-trivial flexible bound, we pass in evidence on how to instantiate the term – effectively delaying the instantiation to the call site. A flexible bound $\forall(\alpha \geq \sigma)$ states that α can be any instance of σ . Each such bound is translated into an explicit witness function of type $\sigma \rightarrow \alpha$ that transforms any type σ into its instantiation α . For example, the System F translation of the $ids :: \forall(\alpha \geq \sigma_{id}). [\alpha]$ value is:

```
ids :: \alpha . (\sigma_{id} -> \alpha) -> [\alpha]
ids = \Lambda\alpha . \lambda(v :: \sigma_{id} -> \alpha) . single \alpha (v id)
```

where we assume a function $single :: \forall\alpha \cdot \alpha \rightarrow [\alpha]$ that creates a singleton list. Note in particular how the witness v takes the identity function to type α . In the body of the `let` expression, ids can now be instantiated to two different types. For *polyL*, we simply pass an identity function that leaves the elements polymorphic:

```
polyL \sigma_{id} (ids \sigma_{id} (\lambda(x :: \sigma_{id}) . x))
```

For the append though, the evidence instantiates all elements to $Int \rightarrow Int$ functions:

```
ids (Int -> Int) (\lambda(x :: \sigma_{id}) . x Int)
```

For the above example, the evidence just changes the type parameters. In the presence of qualified types evidence terms can change runtime terms too. For example, the term $[inc]$ could have the type $\forall(\alpha \geq \sigma). [\alpha]$ where σ equals $\forall\alpha. Num \alpha \Rightarrow \alpha \rightarrow \alpha$. The System F translation passes an explicit runtime dictionary as evidence for $Num \alpha$ to such function. In that case the witness for ids transforms a term of type σ_{id} to σ by adding the dictionary argument:

```
ids \sigma (\lambda(x :: \sigma_{id}) . (\Lambda\alpha . \lambda(num :: Num \alpha) . x \alpha))
```

$$\begin{aligned} \sigma^* &= \text{ft}(\text{nf}(\sigma)) \\ \text{ft}(\tau) &= \tau \\ \text{ft}(\perp) &= \forall\alpha \cdot \alpha \\ \text{ft}(\forall(\alpha = \sigma_1). \sigma_2) &= \text{ft}(\sigma_2)[\alpha \mapsto \text{ft}(\sigma_1)] \\ \text{ft}(\forall(\alpha \geq \perp). \sigma) &= \forall\alpha \cdot \text{ft}(\sigma) \\ \text{ft}(\forall(\alpha \geq \sigma_1). \sigma_2) &= \forall\alpha \cdot (\text{ft}(\sigma_1) \rightarrow \alpha) \rightarrow \text{ft}(\sigma_2) \end{aligned}$$

Figure 8. MLF types to system-F types.

3.2 Translation of rigid bounds

Figure 8 defines the translation of MLF types to System F types. The translation works on normal forms to discard trivial bounds such as monotype bounds or unbound quantifiers. The translation as given is very satisfactory as the only evidence passed is for non-trivial flexible polymorphic bounds – exactly those bounds that were needed to allow type inference with first-class polymorphism in the first place.

Rigid bounds are simply inlined to get the corresponding System F type. Unfortunately, this prevents us also from defining a direct type directed translation on the standard MLF rules. In particular, since the equivalence and abstraction relation can work under rigid bounds (through EQ-CONTEXT and A-CONTEXT), it is no longer the case that when two types are equivalent, that their System F types are equivalent, i.e. $\sigma_1 \equiv \sigma_2$ does not imply $\sigma_1^* = \sigma_2^*$! Take for example the lists $xs_1 :: \forall(\gamma = \forall\alpha\beta. \alpha \rightarrow \beta \rightarrow \alpha). [\gamma]$ and $xs_2 :: \forall(\gamma = \forall\beta\alpha. \alpha \rightarrow \beta \rightarrow \alpha). [\gamma]$. The term

```
if True then xs_1 else xs_2
```

is well-typed in MLF since the types of xs_1 and xs_2 are equivalent,

$$\begin{aligned} \forall(\gamma = \forall\alpha\beta. \alpha \rightarrow \beta \rightarrow \alpha). [\gamma] \\ \equiv \{ \text{EQ-CONTEXT}, \text{EQ-COMM} \} \\ \forall(\gamma = \forall\beta\alpha. \alpha \rightarrow \beta \rightarrow \alpha). [\gamma] \end{aligned}$$

However, according to Figure 8, the corresponding System F types are not equal, and we end up with an ill-typed System F term! Nevertheless, not using an evidence translation for rigid bounds is most important in practice to avoid many ‘trivial’ coercions between equivalent types and we consider this a serious weakness of the earlier approach by Leijen and Löh (2005). Furthermore, as we will see in Section 4, the direct translation of rigid bounds enables a variation of MLF that never passes any evidence which in turns enables us gain insight in the relation between boxy type inference and MLF.

Thus, the challenge is to set up the typed translation in such a way that the above situation is prevented and where rigid bounds never need evidence translation, and we are going to tackle it head on in the following sections.

3.3 Canonical equivalence and abstraction

Both the equivalence and abstraction relation can be applied under rigid bounds through their context rules (EQ-CONTEXT and A-CONTEXT). Therefore, to ensure that the System F type of a rigid bound never changes, we need to define both a restricted equivalence and restricted abstraction relation where equivalent types have equal System F translations.

The only rule that prevents equal System F types is EQ-COMM when quantifiers are rearranged. As we see later, we cannot just remove the rearrangement rule since abstraction must be able to rearrange rigid binders. Fortunately, since the System F translation inlines rigid bounds, rearranging rigid binders does not change the translated System F type and such rearrangement is sound.

Definition 3.3.a (Canonical equivalence): We define *canonical equivalence*, written as \equiv^c , as equal to the MLF equivalence re-

lation of Figure 4, but the EQ-COMM rule is restricted to rearrangement of rigid binders:

$$\text{CEQ-COMM} \quad \frac{\alpha_1 \neq \alpha_2 \quad \alpha_1 \notin \text{ftv}(\sigma_2) \quad \alpha_2 \notin \text{ftv}(\sigma_1)}{(Q) \forall(\alpha_1 = \sigma_1)(\alpha_2 \diamond \sigma_2). \sigma \equiv^c \forall(\alpha_2 \diamond \sigma_2)(\alpha_1 = \sigma_1). \sigma}$$

Similarly, we define canonical abstraction as an extension of canonical equivalence.

Definition 3.3.b (Canonical abstraction): Canonical abstraction, written as \equiv^c , is equal to the MLF abstraction relation in Figure 5 but the rule A-EQUIV is restricted to canonical equivalence:

$$\text{CA-EQUIV} \quad \frac{(Q) \sigma_1 \equiv^c \sigma_2}{(Q) \sigma_1 \equiv^c \sigma_2}$$

Canonical equivalence satisfies many properties of MLF equivalence, for example $\text{nf}(\sigma) \equiv^c \sigma$ holds. Moreover, canonically equivalent types have equal System F types. To make this notion precise, we first define a System F substitution extraction.

Definition 3.3.c (Substitution extraction): The extraction of a monotype substitution from a prefix Q , written as Q^θ , is defined as:

$$\begin{aligned} (\emptyset)^\theta &= \text{id} \\ (\alpha \diamond \sigma, Q)^\theta &= [\alpha \mapsto \tau] \circ Q^\theta \quad \text{iff } \text{nf}(\sigma) = \tau \\ (\alpha \diamond \sigma, Q)^\theta &= Q^\theta \quad \text{otherwise} \end{aligned}$$

Similarly, we can also define the System F substitution extracted from a prefix Q , written as Q^\ominus :

$$\begin{aligned} (\emptyset)^\ominus &= \text{id} \\ (\alpha = \sigma, Q)^\ominus &= [\alpha \mapsto \sigma^*] \circ Q^\ominus \\ (\alpha \geq \sigma, Q)^\ominus &= [\alpha \mapsto \tau^*] \circ Q^\ominus \quad \text{iff } \text{nf}(\sigma) = \tau \\ (\alpha \geq \sigma, Q)^\ominus &= Q^\ominus \quad \text{otherwise} \end{aligned}$$

Note that $(Q_1 Q_2)^\ominus = Q_1^\ominus \circ Q_2^\ominus$ for any well formed prefix $Q_1 Q_2$. Furthermore, the domain of the substitution is a subset of domain of the prefix: $\text{dom}(Q^\ominus) \subseteq \text{dom}(Q)$. Using System F substitutions, we can now state our main theorems of canonical equivalence and abstraction:

Theorem 3.3.d (Canonically equivalent types have equal System F types):

$$(Q) \sigma_1 \equiv^c \sigma_2 \Rightarrow Q^\ominus(\sigma_1^*) = Q^\ominus(\sigma_2^*)$$

Theorem 3.3.e (Canonical abstraction has equal System F types):

$$(Q) \sigma_1 \equiv^c \sigma_2 \Rightarrow Q^\ominus(\sigma_1^*) = Q^\ominus(\sigma_2^*)$$

Both theorems are proved by straightforward induction over the rules of canonical equivalence and abstraction.

3.4 Canonical instance with System F translation

The canonical instance relation, written as \sqsubseteq^c , is defined in Figure 9 and Figure 10. There are two differences from normal MLF instantiation (defined in Figure 6): rule CI-ABSTRACT uses canonical abstraction, and rule CI-COMM reintroduces the ability to rearrange quantifiers with flexible bounds. Furthermore, the canonical instantiation now derives a System F witness term for the instantiation.

The derivation $(Q) \sigma_1 \sqsubseteq^c \sigma_2 \rightsquigarrow f$ states that σ_2 is a canonical instance of σ_1 under prefix Q , where the derived System F witness f has type $\sigma_1^* \rightarrow \sigma_2^*$, i.e. it instantiates a term with System F type σ_1^* into a term with type σ_2^* .

To describe such witness functions conveniently, we use the \bullet notation. An expression e with a hole \bullet stands for a function $\lambda x \cdot e$ where \bullet is replaced by the fresh variable x . For example,

$$\begin{aligned} \text{CI-TRANS} & \quad \frac{(Q) \sigma_1 \sqsubseteq^c \sigma_2 \rightsquigarrow f_1 \quad (Q) \sigma_2 \sqsubseteq^c \sigma_3 \rightsquigarrow f_2}{(Q) \sigma_1 \sqsubseteq^c \sigma_3 \rightsquigarrow f_2 (f_1 \bullet)} \\ \text{CI-PREFIX} & \quad \frac{(Q, \alpha \diamond \sigma) \sigma_1 \sqsubseteq^c \sigma_2 \rightsquigarrow f \quad \alpha \notin \text{dom}(Q)}{(Q) \forall(\alpha \diamond \sigma). \sigma_1 \sqsubseteq^c \forall(\alpha \diamond \sigma). \sigma_2} \\ & \quad \rightsquigarrow \text{gen}[\forall(\alpha \diamond \sigma). \sigma_2] (f (\text{inst}[\forall(\alpha \diamond \sigma). \sigma_1] \bullet)) \\ \text{CI-CONTEXT} & \quad \frac{(Q) \sigma_1 \sqsubseteq^c \sigma_2 \rightsquigarrow f}{(Q) \forall(\alpha \geq \sigma_1). \sigma \sqsubseteq^c \forall(\alpha \geq \sigma_2). \sigma} \\ & \quad \rightsquigarrow \text{gen}[\forall(\alpha \geq \sigma_2). \sigma] \\ & \quad \quad (\text{let } v \text{ e } = w (f \text{ e}) \text{ in } \text{inst}[\forall(\alpha \geq \sigma_1). \sigma] \bullet) \end{aligned}$$

Figure 9. General rules for canonical instance with evidence translation

$$\begin{aligned} \text{CI-BOTTOM} & \quad (Q) \perp \sqsubseteq^c \sigma \rightsquigarrow \bullet \sigma^* \\ \text{CI-ABSTRACT} & \quad \frac{(Q) \sigma_1 \sqsubseteq^c \sigma_2}{(Q) \sigma_1 \sqsubseteq^c \sigma_2 \rightsquigarrow \bullet} \\ \text{CI-HYP} & \quad \frac{\forall(\alpha \geq \sigma) \in Q}{(Q) \sigma \sqsubseteq^c \alpha \rightsquigarrow v \bullet} \\ \text{CI-RIGID} & \quad (Q) \forall(\alpha \geq \sigma_1). \sigma_2 \sqsubseteq^c \forall(\alpha = \sigma_1). \sigma_2 \\ & \quad \rightsquigarrow \text{gen}[\forall(\alpha = \sigma_1). \sigma_2] (\text{inst}[\forall(\alpha \geq \sigma_1). \sigma_2] \bullet) \\ \text{CI-COMM} & \quad \frac{\alpha_1 \neq \alpha_2 \quad \alpha_1 \notin \text{ftv}(\sigma_2) \quad \alpha_2 \notin \text{ftv}(\sigma_1)}{(Q) \forall(\alpha_1 \geq \sigma_1)(\alpha_2 \geq \sigma_2). \sigma \sqsubseteq^c \forall(\alpha_2 \geq \sigma_2)(\alpha_1 \geq \sigma_1). \sigma} \\ & \quad \rightsquigarrow \text{gen}[\forall(\alpha_2 \geq \sigma_2). \forall(\alpha_1 \geq \sigma_1). \sigma] \\ & \quad \quad (\text{inst}[\forall(\alpha_1 \geq \sigma_1). \forall(\alpha_2 \geq \sigma_2). \sigma] \bullet) \end{aligned}$$

Figure 10. Canonical instance with evidence translation

$$\begin{aligned} \bullet &= \lambda x \cdot x \\ \bullet \sigma^* &= \lambda x \cdot x \sigma^* \\ \lambda y \cdot \bullet &= \lambda x \cdot \lambda y \cdot x \end{aligned}$$

The rule CI-BOTTOM instantiates the bottom type (\perp) to an arbitrary type σ . The evidence term thus needs to instantiate a term of type $\forall \alpha \cdot \alpha$ to σ^* , which corresponds to applying the type σ^* in System F, and the derived evidence is $\bullet \sigma^*$.

Rule CI-ABSTRACT is also straightforward. Due to Theorem 3.3.e, we know that σ_1 and σ_2 have equal System F types and that the term stays the same, i.e. the evidence is the identity: \bullet .

The evidence for rule CI-HYP needs to transform a type σ^* into α – but how should it do that if nothing is known about σ ? The solution to this problem is to annotate each flexible bound $\forall(\alpha \geq \sigma)$ in the prefix with a witness term v of type $\sigma^* \rightarrow \alpha$. One can easily check that these terms are always in scope (just like the type variables in the prefix). The rule CI-HYP simply applies the witness term to transform σ^* to α , i.e. $v \bullet$.

Most of the other rules make use of two helper functions, $\text{gen}[\cdot]$ and $\text{inst}[\cdot]$, defined in Figure 11. These functions introduce evidence terms v for flexible bounds and remove trivial bounds.

The function $\text{gen}[\forall(\alpha \diamond \sigma_1). \sigma_2]$ takes a System F term of type σ_2^* and generalizes it to a term of type $(\forall(\alpha \diamond \sigma_1). \sigma_2)^*$. The $\text{gen}[\cdot]$ function ensures that in all cases the type variable α and the evidence term v of type $\sigma_1^* \rightarrow \alpha$ are defined and in scope. For convenience we use **let** bindings to define v , even though System F does not contain them, and one can read those **let** bindings as a substitution on witness terms. For rigid and trivial bounds, $\text{gen}[\cdot]$ just substitutes σ_1 for α , and the evidence term v becomes the identity.

$$\begin{aligned}
& \text{gen}[\forall(\alpha \diamond_v \sigma). \sigma_0] :: \sigma_0^* \rightarrow (\forall(\alpha \diamond \sigma). \sigma_0)^* \\
& \text{gen}[\forall(\alpha =_v \sigma). \sigma_0] \\
& \quad = (\mathbf{let} \ v \ e = e \ \mathbf{in} \ \bullet)[\alpha \mapsto \sigma^*] \\
& \text{gen}[\forall(\alpha \geq_v \sigma). \sigma_0] \\
& \quad = (\mathbf{let} \ v \ e = e \ \mathbf{in} \ \bullet)[\alpha \mapsto \sigma^*] \quad \text{iff } \alpha \notin \text{ftv}(\sigma_0) \\
& \quad = (\mathbf{let} \ v \ e = e \ \mathbf{in} \ \bullet)[\alpha \mapsto \sigma^*] \quad \text{iff } \text{nf}(\sigma_0) = \alpha \\
& \quad = (\mathbf{let} \ v \ e = e \ \mathbf{in} \ \bullet)[\alpha \mapsto \sigma^*] \quad \text{iff } \text{nf}(\sigma) = \tau \\
& \quad = \Lambda \alpha \cdot \mathbf{let} \ v \ e = e \ \mathbf{in} \ \bullet \quad \text{iff } \text{nf}(\sigma) = \perp \\
& \quad = \Lambda \alpha \cdot \lambda v : \sigma^* \rightarrow \alpha \cdot \bullet \quad \text{otherwise} \\
& \text{inst}[\forall(\alpha = \sigma). \sigma_0] \\
& \quad = \bullet \\
& \text{inst}[\forall(\alpha \geq_v \sigma). \sigma_0] \\
& \quad = \bullet \quad \text{iff } \alpha \notin \text{ftv}(\sigma_0) \\
& \quad = \bullet \quad \text{iff } \text{nf}(\sigma_0) = \alpha \\
& \quad = \bullet \quad \text{iff } \text{nf}(\sigma) = \tau \\
& \quad = \bullet \ \alpha \quad \text{iff } \text{nf}(\sigma) = \perp \\
& \quad = \bullet \ \alpha \ v \quad \text{otherwise}
\end{aligned}$$

Figure 11. Canonical type generalization and instantiation

For an unconstrained bound $\forall(\alpha \geq \perp)$, a type lambda generalizes over α . For a non-trivial flexible bound, we not only generalize over the type α , but also bind the evidence v as an argument.

Dually, the function $\text{inst}[\forall(\alpha \diamond_v \sigma_1). \sigma_2]$ takes a System F term of type $(\forall(\alpha \diamond_v \sigma_1). \sigma_2)^*$ and instantiates it to σ_2^* (or σ_1^* if $\text{nf}(\sigma_2) = \alpha$), where it is assumed that α and v are in scope and where v has type $\sigma_1^* \rightarrow \alpha$. One can easily check that this is the case for all the uses of $\text{inst}[\cdot]$. For rigid and trivial bounds nothing has to be done. For an unconstrained bound we instantiate the type argument to α , and for a non-trivial flexible bound we instantiate the type argument and pass in the evidence function v as an argument.

With these helper functions, the evidence translation for the other rules is straightforward. Rule CI-RIGID uses $\text{inst}[\cdot]$ to instantiate the flexible quantifier and immediately applies $\text{gen}[\cdot]$ with a rigid bound, which substitutes α for σ_1^* , and binds v to the identity. Another interesting rule is CI-CONTEXT. Since the bounded types themselves are instantiated, the types of the witness functions v and w are not equal and we need to build an *witness transformer*: using the evidence $f :: \sigma_1^* \rightarrow \sigma_2^*$, we can define the old witness $v :: \sigma_1^* \rightarrow \alpha$ in terms of the new witness $w :: \sigma_2^* \rightarrow \alpha$, as $v = w \circ f$.

To define soundness of the System F translation, we need to be able to refer to the witness terms in the prefix from the System F type environment, and we define the following operation to extract such environment from the prefix.

Definition 3.4.a (*System-F environment extraction*): The environment extraction from a prefix Q is written as Q^Γ and defined as:

$$\begin{aligned}
(\emptyset)^\Gamma &= \emptyset \\
(\alpha = \sigma, Q)^\Gamma &= Q^\Gamma \\
(\alpha \geq_v \sigma, Q)^\Gamma &= v : \sigma^* \rightarrow \alpha, Q^\Gamma
\end{aligned}$$

Using environment extraction, we can state that the derived System F term is well-typed in System F. The typing relation of System F ($\Gamma \vdash^F e : \sigma$) is standard and we omit it here.

Theorem 3.4.b (*Soundness of System F translation*): If two types σ_1 and σ_2 are in a canonical instance relation under a prefix Q with a witness f , then the witness f is well typed in System F under the type environment extracted from Q , with type $\sigma_1^* \rightarrow \sigma_2^*$:

$$(Q) \ \sigma_1 \sqsubseteq^c \sigma_2 \rightsquigarrow f \Rightarrow Q^\Theta(Q^\Gamma) \vdash^F f : Q^\Theta(\sigma_1^* \rightarrow \sigma_2^*)$$

$$\begin{aligned}
\text{CT-VAR} & \frac{x : \sigma \in \Gamma}{(Q) \ \Gamma \vdash^c x : \sigma \rightsquigarrow x} \\
\text{CT-APP} & \frac{(Q) \ \Gamma \vdash^c e_1 : \tau_2 \rightarrow \tau_1 \rightsquigarrow e_1 \quad (Q) \ \Gamma \vdash^c e_2 : \tau_2 \rightsquigarrow e_2}{(Q) \ \Gamma \vdash^c e_1 \ e_2 : \tau_1 \rightsquigarrow e_1 \ e_2} \\
\text{CT-FUN} & \frac{(Q) \ \Gamma, x : \tau_1 \vdash^c e : \tau_2 \rightsquigarrow e}{(Q) \ \Gamma \vdash^c \lambda x. e : \tau_1 \rightarrow \tau_2 \rightsquigarrow \lambda(x : \tau_1^*) \cdot e} \\
\text{CT-LET} & \frac{(Q) \ \Gamma \vdash^c e_1 : \sigma_1 \rightsquigarrow e_1 \quad (Q) \ \Gamma, x : \sigma_1 \vdash^c e_2 : \sigma_2 \rightsquigarrow e_2}{(Q) \ \Gamma \vdash^c \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 : \sigma_2 \rightsquigarrow (\lambda(x : \sigma_1^*) \cdot e_2) \ e_1} \\
\text{CT-GEN} & \frac{(Q, \alpha \diamond_v \sigma_1) \ \Gamma \vdash^c e : \sigma_2 \rightsquigarrow e \quad \alpha \notin \text{ftv}(\Gamma)}{(Q) \ \Gamma \vdash^c e : \forall(\alpha \diamond_v \sigma_1). \sigma_2 \rightsquigarrow \text{gen}[\forall(\alpha \diamond_v \sigma_1). \sigma_2] \ e} \\
\text{CT-INST} & \frac{(Q) \ \Gamma \vdash^c e : \sigma_1 \rightsquigarrow e \quad (Q) \ \sigma_1 \sqsubseteq^c \sigma_2 \rightsquigarrow f}{(Q) \ \Gamma \vdash^c e : \sigma_2 \rightsquigarrow f \ e} \\
\text{CT-ANN} & \frac{(Q) \ \Gamma \vdash^c e : \sigma_1 \rightsquigarrow e \quad (Q) \ \sigma \sqsubseteq^c \sigma_1}{(Q) \ \Gamma \vdash^c (e :: \sigma) : \sigma \rightsquigarrow e}
\end{aligned}$$

Figure 12. Type rules with evidence translation

3.5 Canonical type rules with System F translation

Figure 12 defines canonical type rules for MLF that derive a corresponding System F term. The type rules are equivalent to the corresponding MLF type rules (see Figure 2) except for the use of canonical instance in CT-INST and canonical abstraction in CT-ANN. The expression $(Q) \ \Gamma \vdash^c e : \sigma \rightsquigarrow e$ states that under a prefix Q and type environment Γ , the expression e is well typed with type σ , and e is a corresponding System F term with type σ^* .

The type directed translation is straightforward. Rule CT-VAR simply returns x , since the condition $x : \sigma \in \Gamma$ implies that x has type σ^* under the System F environment Γ^* , (i.e. $\Gamma^* \vdash^F x : \sigma^*$). In CT-APP, e_1 and e_2 both have monotypes and their System F types are equivalent, and we can directly apply the corresponding System F terms too. The same holds for CT-FUN that directly translates into a corresponding lambda expression. Rule CT-LET is interesting as the System F translation uses a polymorphic lambda expression since System F does not contain **let** bindings.

Generalization in rule CT-GEN uses the $\text{gen}[\cdot]$ function defined in Figure 11 to generalize the term e with type σ_2^* into $(\forall(\alpha \diamond_v \sigma_1). \sigma_2)^*$. The instantiation rule (CT-INST) uses the witness function f of type $\sigma_1^* \rightarrow \sigma_2^*$, derived with canonical instance to transform the term e to a term of type σ_2^* . Finally, type annotations in CT-ANN use canonical abstraction. By Theorem 3.3.e, their System F types are equal and no translation is necessary.

Theorem 3.5.a (*Canonical type inference derives well typed System F terms*): If under a prefix Q and type environment Γ an expression e is well typed with type σ and a translated System F term e , then there exists a System F derivation such that under a System F type environment consisting of the environment extraction of Q (namely Q^Γ) and the translated environment Γ^* , the term e is well typed with type σ^* .

$$(Q) \ \Gamma \vdash^c e : \sigma \rightsquigarrow e \Rightarrow Q^\Theta(Q^\Gamma \Gamma^*) \vdash^F Q^\Theta(e) : Q^\Theta(\sigma^*)$$

Note that under an empty prefix this simplifies to:

$$() \ \Gamma \vdash^c e : \sigma \rightsquigarrow e \Rightarrow \Gamma^* \vdash^F e : \sigma^*$$

3.6 Soundness

By carefully changing the original equivalence, abstraction, and instance relation, we were able to derive a type directed and sound System F translation for instantiation terms. But at the same time we are no longer doing MLF instantiation but canonical instantiation, and we would like to establish that the new relations are sound and complete with respect to the MLF relations. Clearly, the canonical relations are sound with respect to MLF:

Theorem 3.6.a (*Canonical equivalence, abstraction, and instance are sound*):

$$\begin{aligned} (Q) \quad \sigma_1 &\equiv^c \sigma_2 \Rightarrow (Q) \quad \sigma_1 \equiv \sigma_2 \\ (Q) \quad \sigma_1 &\in^c \sigma_2 \Rightarrow (Q) \quad \sigma_1 \in \sigma_2 \\ (Q) \quad \sigma_1 &\sqsubseteq^c \sigma_2 \Rightarrow (Q) \quad \sigma_1 \sqsubseteq \sigma_2 \end{aligned}$$

Proof of Theorem 3.6.a: This is immediate since all the new relations are sub-relations of the MLF relations. Equivalence includes all the rules of canonical equivalence, and allows rearrangement of flexible binders. Canonical abstraction equals MLF abstraction but restricts equivalence to canonical equivalence. Finally, canonical instance includes an extra rule (CI-COMM) to rearrange flexible binders, but this is also included in MLF instantiation by EQ-COMM (via A-EQUIV and I-ABSTRACT). \square

As a corollary, it follows that the canonical type rules are also sound since the only difference with the MLF type rules is the use of canonical instance and abstraction.

Theorem 3.6.b (*Canonical type inference is sound*):

$$(Q) \quad \Gamma \vdash^c e : \sigma \rightsquigarrow e \Rightarrow (Q) \quad \Gamma \vdash e : \sigma$$

3.7 Completeness

Completeness of the type rules cannot be established directly since the canonical relations are not complete with respect to the MLF relations. Take for example the following equivalent types:

$$(Q) \quad \forall \alpha \beta. \alpha \rightarrow \beta \rightarrow \alpha \equiv \forall \beta \alpha. \alpha \rightarrow \beta \rightarrow \alpha$$

These types are equivalent by using EQ-COMM to rearrange the flexible binders. Since canonical equivalence restricts rearrangement to rigid binders, the above types are *not* canonically equivalent (as it should, since the corresponding System F types are different). This seems a fatal flaw of our new rules: incompleteness would imply that our new rules can only be used to type a subset of the programs accepted by MLF.

Fortunately, the rules are incomplete only with respect to the order of the quantifiers. We can show that the canonical relations are complete with respect to the MLF relations under specific invariants where rigid types are in *canonical form*. In the next section we formalize this notion and establish a completeness theorem.

3.7.1 Canonical form

Figure 13 defines the *canonical form* of a type σ , written as $\text{can}(\sigma)$. The canonical form is a rearrangement of the normal form, where the position of a quantifier is uniquely determined by the occurrences of the bound type variable. Canonical forms satisfy the following properties:

Properties 3.7.a

- i. $(Q) \quad \sigma_1 \equiv \sigma_2 \Leftrightarrow Q^\theta(\text{can}(\sigma_1)) = Q^\theta(\text{can}(\sigma_2))$
- ii. $\text{ftv}(\text{can}(\sigma)) = \text{ftv}(\sigma)$

As a corollary of Property 3.7.a.i, we have under an unconstrained prefix that equivalent types have equal canonical forms:

$$\sigma_1 \equiv \sigma_2 \Leftrightarrow \text{can}(\sigma_1) = \text{can}(\sigma_2)$$

The rearranged form makes use of the insertion function $\text{ins}(\cdot, \cdot)$ to insert quantifiers in their canonical position. Insertion of into a

$$\begin{aligned} \text{can}(\sigma) &= \text{rf}(\text{nf}(\sigma)) \\ \text{rf}(\tau) &= \tau \\ \text{rf}(\perp) &= \perp \\ \text{rf}(\forall(\alpha \diamond \sigma_1). \sigma_2) &= \text{ins}(\alpha \diamond \text{rf}(\sigma_1), \text{rf}(\sigma_2)) \\ \text{ins}(\alpha \diamond \sigma, \tau) &= \forall(\alpha \diamond \sigma). \tau \\ \text{ins}(\alpha \diamond \sigma, \perp) &= \perp \\ \text{ins}(\alpha_1 \diamond \sigma_1, \forall(\alpha_2 \diamond \sigma_2). \sigma) &= \text{if } \alpha_1 \neq \alpha_2 \wedge \alpha_2 \notin \text{ftv}(\sigma_1) \\ &= \text{if } \alpha_1 \# (\forall(\alpha_2 \diamond \sigma_2). \sigma) < \alpha_2 \# \sigma \\ &\quad \text{then } \forall(\alpha_2 \diamond \sigma_2). \text{ins}(\alpha_1 \diamond \sigma_1, \sigma) \\ &\quad \text{else } \forall(\alpha_1 \diamond \sigma_1)(\alpha_2 \diamond \sigma_2). \sigma \end{aligned}$$

Figure 13. Canonical and rearranged form

$$\begin{aligned} \alpha \# \sigma &= \max(\text{occurrences}(\alpha, \sigma) \cup \{0\}) \\ \text{occurrences}(\alpha, \sigma) &= \{o \mid \text{select}(\text{skeleton}(\sigma), o) = \alpha\} \\ \text{select}(\perp, 1) &= \perp \\ \text{select}(\alpha, 1) &= \alpha \\ \text{select}(c \tau_1 \dots \tau_n, 1) &= c \\ \text{select}(c \tau_1 \dots \tau_n, i.o) &= \text{select}(\tau_i, o) \\ \text{skeleton}(\tau) &= \tau \\ \text{skeleton}(\perp) &= \perp \\ \text{skeleton}(\forall(\alpha \diamond \sigma_1). \sigma_2) &= \text{skeleton}(\sigma_2)[\alpha \mapsto \text{skeleton}(\sigma_1)] \end{aligned}$$

Figure 14. The position of type variable

monomorphic type or \perp has no effect, but otherwise quantifiers may be rearranged. The side condition of the last insertion case can always be satisfied by simple alpha renaming. The relative order between quantifiers is now uniquely determined by the *position* of their bound variable in the type, written as $\alpha \# \sigma$, and is defined in Figure 14.

The position of a type variable in a type is defined as the maximum of all its occurrences of in the type, or 0 if the type variable is unbound. This case never occurs for a canonical form since the rearranged form is taken over the normal form which discards unbound quantifiers.

To define the occurrences of a type, we first define the skeleton of a type ($\text{skeleton}(\sigma)$) as the full inlining of all its bounds. For example

$$\text{skeleton}(\forall(\beta \geq \forall \gamma. \alpha \rightarrow \gamma). \beta \rightarrow \alpha) = (\alpha \rightarrow \perp) \rightarrow \alpha$$

The selection function $\text{select}(\cdot, \cdot)$ takes such skeleton and a *position*, and returns an element of that skeleton. For example

$$\begin{aligned} \text{select}((\alpha \rightarrow \perp) \rightarrow \beta, 1.1.1) &= \alpha \\ \text{select}((\alpha \rightarrow \perp) \rightarrow \beta, 2.1) &= \beta \\ \text{select}((\alpha \rightarrow \perp) \rightarrow \beta, 1.2.1) &= \perp \end{aligned}$$

The occurrences of a type variable is simply the set of all positions of the type variable in the skeleton, for example

$$\begin{aligned} \text{occurrences}(\alpha, \forall(\beta \geq \forall \gamma. \alpha \rightarrow \gamma). \beta \rightarrow \alpha) &= \{1.1.1, 2.1\} \\ \text{occurrences}(\beta, \forall(\beta \geq \forall \gamma. \alpha \rightarrow \gamma). \beta \rightarrow \alpha) &= \emptyset \\ \text{occurrences}(\beta, \beta \rightarrow \alpha) &= \{1.1\} \end{aligned}$$

During insertion, we have to be careful never to insert beyond dependent binders. In particular, we need to ensure that when quantifiers are rearranged (in the **then** branch of the last case of insertion) that $\alpha_1 \notin \text{ftv}(\sigma_2)$. Indeed, together with the side conditions on insertion these are exactly the conditions under which

we can apply EQ-COMM. Fortunately, the position of type variables respects dependence:

Lemma 3.7.b (*Insertion respects dependencies*): The definition of the position of a type variable respects dependence between binders.

$$\alpha_1 \# (\forall (\alpha_2 \diamond \sigma_2). \sigma) < \alpha_2 \# \sigma \Rightarrow \alpha_1 \notin \text{ftv}(\sigma_2)$$

Proof of Lemma 3.7.b: First note that if $\alpha_2 \notin \text{ftv}(\sigma)$, then the position $\alpha_2 \# \sigma = 0$, and the position of α_1 cannot be smaller. Otherwise, if $\alpha_2 \in \text{ftv}(\sigma)$, we can prove Lemma 3.7.b by contradiction. Assume that $\alpha_1 \in \text{ftv}(\sigma_2)$ (1), where $\alpha_2 \# \sigma$ equals $\sigma_1.1$. In that case, the set of occurrences of α in $\forall (\alpha_2 \diamond \sigma_2). \sigma$, namely occurrences $(\alpha, \forall (\alpha_2 \diamond \sigma_2). \sigma)$, contains an occurrence for α with the form $\sigma_1.\sigma_2.1$, which is larger or equal to $\sigma_1.1$ (since 0 cannot be part of an occurrence). Since the position is the maximum of the occurrences, we have $\alpha \# (\forall (\alpha_2 \diamond \sigma_2). \sigma) \geq \alpha_2 \# \sigma$. By contradiction, this invalidates the assumption (1), and implies $\alpha_1 \notin \text{ftv}(\sigma_2)$. \square

3.7.2 Completeness under canonical form

Using canonical forms, we can now state a completeness theorem for canonical equivalence.

Theorem 3.7.c (*Completeness of canonical equivalence*): Whenever two types are equivalent, then their canonical forms are also canonically equivalent:

$$(Q) \sigma_1 \equiv \sigma_2 \Rightarrow (Q) \text{can}(\sigma_1) \equiv^c \text{can}(\sigma_2)$$

This can be proved by straightforward induction over the rules of equivalence.

For canonical abstraction, completeness does not hold directly since CI-HYP abstracts over rigid types in the prefix. If those types are not in canonical form the rule may not apply. To remedy this situation we are going to define the notion of *weak canonical form*.

Definition 3.7.d (*Weak canonical form*): A type σ is in weak canonical form if all its rigid bounds are in canonical form. We write σ^c for types in weak canonical form. A prefix Q is in weak canonical form if all its quantifiers are in weak canonical form, and write Q^c for such prefix.

The following properties hold for types in weak canonical form:

Properties 3.7.e

- i. $(Q^c) \sigma^c \sqsubseteq^c \text{can}(\sigma)$, a type in weak canonical form instantiates to its canonical form, and conversely;
- ii. $(Q^c) \text{can}(\sigma) \sqsubseteq^c \sigma^c$.

Using weak canonical form, we can now state a completeness theorem for canonical abstraction:

Theorem 3.7.f (*Completeness of canonical abstraction*): Whenever two types are in an abstraction relation under a prefix Q , then their canonical forms are in canonical abstraction relation under a prefix Q^c in weak canonical form:

$$(Q) \sigma_1 \sqsubseteq \sigma_2 \Rightarrow (Q^c) \text{can}(\sigma_1) \sqsubseteq^c \text{can}(\sigma_2)$$

For this theorem to hold, it is essential that canonical equivalence can rearrange rigid binders. Take for example the following abstraction:

$$\begin{aligned} & \forall (\alpha = \sigma_{id}). \forall (\gamma \geq \sigma_{id}). \forall (\beta = \sigma_{id}). [\beta] \rightarrow [\gamma] \rightarrow [\alpha] \\ & \sqsubseteq^c \{ \text{A-PREFIX, A-CONTEXT, A-HYP, and EQ-MONO} \} \\ & \forall (\alpha = \sigma_{id}). \forall (\gamma \geq \sigma_{id}). [\alpha] \rightarrow [\gamma] \rightarrow [\alpha] \end{aligned}$$

Alas, although the first type is in canonical form, the second one is not. This happens because abstraction can change the occurrences of rigidly bound types through A-HYP. Therefore, rearrangement of rigid binders is necessary to derive canonical forms.

Proof of Theorem 3.7.f: By induction over the rules of abstraction:

Case A-EQUIV: We have $(Q) \sigma_1 \equiv \sigma_2$, and therefore by Theorem 3.7.c, $(Q^c) \text{can}(\sigma_1) \equiv^c \text{can}(\sigma_2)$. By rule CA-EQUIV we can now derive $(Q^c) \text{can}(\sigma_1) \sqsubseteq^c \text{can}(\sigma_2)$.

Case A-HYP: We have $(\alpha = \sigma) \in Q$, and by therefore $(\alpha = \text{can}(\sigma)) \in Q^c$ since Q^c is in weak canonical form. By rule CA-HYP it follows that $(Q^c) \text{can}(\sigma) \sqsubseteq^c \alpha$.

Case A-CONTEXT: By induction, we have $(Q^c) \text{can}(\sigma_1) \sqsubseteq^c \text{can}(\sigma_2)$ (1). Since $\text{can}(\forall (\alpha = \sigma_1). \sigma) = \text{rf}(\text{nf}(\forall (\alpha = \sigma_1). \sigma))$ we proceed by case analysis over the normal form:

subcase $\alpha \notin \text{ftv}(\sigma)$: In this case, it follows directly that $\text{can}(\forall (\alpha = \sigma_1). \sigma) = \text{can}(\sigma) = \text{can}(\forall (\alpha = \sigma_2). \sigma)$.

subcase $\text{nf}(\sigma_1) = \tau_1$ (2): First we note that due to (2) and (1), $\text{can}(\sigma_2) = \tau_2$ (3) for some monotype τ_2 . We can now derive under a prefix Q :

$$\begin{aligned} \text{can}(\forall (\alpha = \sigma_1). \sigma) &= \{ (2) \} \\ \text{can}(\sigma)[\alpha \mapsto \tau_1] &\sqsubseteq^c \{ \text{CA-EQUIV, CEQ-MONO} \} \\ \forall (\alpha = \tau_1). \text{can}(\sigma) &\sqsubseteq^c \{ \text{CA-CONTEXT, CA-EQUIV, (2)} \} \\ \forall (\alpha = \text{can}(\sigma_1)). \text{can}(\sigma) &\sqsubseteq^c \{ \text{CA-CONTEXT, (1)} \} \\ \forall (\alpha = \text{can}(\sigma_2)). \text{can}(\sigma) &\sqsubseteq^c \{ \text{CA-CONTEXT, CA-EQUIV, (3)} \} \\ \forall (\alpha = \tau_2). \text{can}(\sigma) &\sqsubseteq^c \{ \text{CA-EQUIV, CEQ-MONO} \} \\ \text{can}(\sigma)[\alpha \mapsto \tau_2] &= \{ (3) \} \\ \text{can}(\forall (\alpha = \sigma_2). \text{can}(\sigma)) & \end{aligned}$$

subcase $\text{nf}(\sigma) = \alpha$: We have $\text{can}(\forall (\alpha = \sigma_1). \sigma) = \text{can}(\sigma_1)$ (4), and $\text{can}(\sigma_2) = \text{can}(\forall (\alpha = \sigma_2). \sigma)$. By EQ-REFL, CA-EQUIV, (1), and transitivity, we can derive $(Q^c) \text{can}(\forall (\alpha = \sigma_1). \sigma) \sqsubseteq^c \text{can}(\forall (\alpha = \sigma_2). \sigma)$.

subcase otherwise: In this case $\text{can}(\forall (\alpha = \sigma_1). \sigma)$ equals $\text{ins}(\forall (\alpha = \text{can}(\sigma_1)), \text{can}(\sigma_2))$ where the bound $\forall (\alpha = \text{can}(\sigma_1))$ is inserted into $\text{can}(\sigma_2)$. But by using CEQ-COMM to rearrange the rigid binder we can bring it back to the front and we can derive $(Q^c) \text{can}(\forall (\alpha = \sigma_1). \sigma) \sqsubseteq^c \forall (\alpha = \text{can}(\sigma_1)). \text{can}(\sigma)$. By CA-CONTEXT and (1), we know that $(Q^c) \forall (\alpha = \text{can}(\sigma_1)). \text{can}(\sigma) \sqsubseteq^c \forall (\alpha = \text{can}(\sigma_2)). \text{can}(\sigma)$. Finally, we can again apply CEQ-COMM repeatedly to derive $(Q^c) \forall (\alpha = \text{can}(\sigma_2)). \text{can}(\sigma) \sqsubseteq^c \text{ins}(\forall (\alpha = \text{can}(\sigma_2)), \text{can}(\sigma))$ where $\text{ins}(\forall (\alpha = \text{can}(\sigma_2)), \text{can}(\sigma))$ equals $\text{can}(\forall (\alpha = \sigma_2). \sigma)$.

Case A-PREFIX: Similar to A-CONTEXT.

Case A-TRANS: Follows directly by induction and CA-TRANS. \square

The above completeness theorems are rather weak as they require the types to be in canonical form. For canonical instantiation we have a stronger result where only rigid bounds need to be in canonical form:

Theorem 3.7.g (*Completeness of canonical instance*): When two types are in an instance relation under a prefix Q , then their weak canonical forms are also in a canonical instance relation under the same prefix in weak canonical form:

$$(Q) \sigma_1 \sqsubseteq \sigma_2 \Rightarrow (Q^c) \sigma_1^c \sqsubseteq^c \sigma_2^c \rightsquigarrow f$$

Proof of Theorem 3.7.g: By induction over the instance relation.

Case I-BOTTOM: We have $(Q) \perp \sqsubseteq \sigma$. Using CI-BOTTOM, we can derive $(Q^c) \perp \sqsubseteq^c \sigma^c$.

Case I-ABSTRACT: We have $(Q) \sigma_1 \sqsubseteq \sigma_2$. As a consequence of Theorem 3.7.f, $(Q^c) \text{can}(\sigma_1) \sqsubseteq^c \text{can}(\sigma_2)$ holds (1). By Property 3.7.e.i and Property 3.7.e.ii, we can also derive $(Q^c) \sigma_1^c \sqsubseteq^c \text{can}(\sigma_1)$ (2), and $(Q^c) \text{can}(\sigma_2) \sqsubseteq^c \sigma_2^c$ (3). Combining (1), (2), and (3), using CI-TRANS, we can derive $(Q^c) \sigma_1^c \sqsubseteq^c \sigma_2^c$.

Case I-HYP: We assume $(\alpha \geq \sigma) \in Q$. This implies $(\alpha \geq \sigma^c) \in Q^c$, and by CI-HYP, $(Q^c) \sigma^c \sqsubseteq^c \alpha$.

Case I-RIGID: We have $(Q) \forall (\alpha \geq \sigma_1). \sigma \sqsubseteq \forall (\alpha = \sigma_1). \sigma$. By Property 3.7.e.i, $(Q^c) \sigma_1^c \sqsubseteq^c \text{can}(\sigma_1)$ (4). The weak canonical form of $\forall (\alpha \geq \sigma_1). \sigma$ is $\forall (\alpha \geq \sigma_1^c). \sigma^c$. Using CI-CONTEXT and (4), we instantiate to $\forall (\alpha \geq \text{can}(\sigma_1)). \sigma^c$. By CI-RIGID, this instantiates to $\forall (\alpha = \text{can}(\sigma_1)). \sigma^c$ which equals $(\forall (\alpha = \sigma_1). \sigma_2)^c$.

Case I-TRANS: Immediate by induction and CI-TRANS.

Case I-PREFIX: Immediate by induction and CI-PREFIX.

Case I-CONTEXT: Immediate by induction and CI-CONTEXT. \square

Using the completeness theorems on the canonical relations, we can finally state a completeness theorem for the canonical type rules.

Theorem 3.7.h (*Canonical type inference is complete*): If we can infer a type σ for an expression in e under a certain prefix Q and environment Γ , there exists a canonical type derivation where σ , Q , and Γ are in weak canonical form:

$$(Q) \Gamma \vdash e : \sigma \Rightarrow (Q^c) \Gamma^c \vdash^c e : \sigma^c \rightsquigarrow e$$

Proof of Theorem 3.7.h: By induction of the inference rules.

Case T-VAR: We have $(x : \sigma) \in \Gamma$, and therefore, $(x : \sigma^c) \in \Gamma^c$. By rule CT-VAR, we can derive $(Q^c) \Gamma^c \vdash^c x : \sigma^c$.

Case T-APP: Immediate by induction.

Case T-FUN: Immediate by induction.

Case T-LET: Immediate by induction.

Case T-GEN: By induction and Property 3.7.a.ii, we have $\alpha \notin \text{ftv}(\Gamma^c)$, and either $(Q^c, \alpha = \text{can}(\sigma_1)) \Gamma^c \vdash^c e : \sigma_2^c$ or $(Q^c, \alpha \geq \sigma_1^c) \Gamma^c \vdash^c e : \sigma_2^c$. By rule CT-GEN, we can derive $(Q^c) \Gamma^c \vdash^c e : \forall(\alpha = \text{can}(\sigma_1)). \sigma_2^c$ or $(Q^c) \Gamma^c \vdash^c e : \forall(\alpha \geq \sigma_1^c). \sigma_2^c$ where both derived types are in weak canonical form.

Case T-INST: We have $(Q) \Gamma \vdash e : \sigma_1$ (1), and $(Q) \sigma_1 \sqsubseteq \sigma_2$ (2). By the induction hypothesis and (1), we have $(Q^c) \Gamma^c \vdash^c e : \sigma_1^c$ (3). By Theorem 3.7.g and (2), we can conclude $(Q^c) \sigma_1^c \sqsubseteq^c \sigma_2^c$, and by CT-INST and (3), we have $(Q^c) \Gamma^c \vdash^c e : \sigma_2^c$.

Case T-ANN: We have $(Q) \Gamma \vdash e : \sigma_1$ (4), and $(Q) \sigma \sqsubseteq \sigma_1$ (5). By the induction hypothesis and (4), we have $(Q^c) \Gamma^c \vdash^c e : \sigma_1^c$. Due to Property 3.7.e.i, we know that $(Q^c) \sigma_1^c \sqsubseteq^c \text{can}(\sigma_1)$, and we can use CT-INST to derive $(Q^c) \Gamma^c \vdash^c e : \text{can}(\sigma_1)$ (6). As a consequence of Theorem 3.7.f and (5), we have $(Q^c) \text{can}(\sigma_1) \sqsubseteq^c \text{can}(\sigma_2)$ (assuming the type annotation is converted to canonical form), and we can use CT-ABS with (6) to derive $(Q^c) \Gamma^c \vdash^c (e :: \sigma_1) : \text{can}(\sigma_1)$. Using CI-INST and Property 3.7.e.ii, we can instantiate to $(Q^c) \Gamma^c \vdash^c (e :: \sigma_1) : \sigma_1^c$. \square

3.8 Type inference

MLF has an effective type inference algorithm that infers principal types. It is beyond the scope of this paper to discuss type inference in detail, but we remark that the standard MLF inference algorithm can be used to infer types to canonically, as long as we maintain the invariants required for completeness (Theorem 3.7.h). In particular, type annotations must be normalized to canonical form, and the prefix, type environment, and inferred types must be in weak canonical form. Ensuring that types are in weak canonical form is easy to do by requiring in the update algorithm (Le Botlan 2004, page 123) that rigid bounds are only updated with types in canonical form, rearranging binders as necessary. No further changes are required. The translation to System F types can be done similarly to Leijen and Löh (2005) but simplified to remove translation for rigid bounds.

4. A restriction to System F types

We have defined an elegant type directed translation from MLF to System F, where only polymorphic flexible bounds require extra evidence and type annotations are needed only on lambda bound arguments that are used polymorphically. Even though this is probably the best we can hope for, it is interesting to consider a more restricted version.

In particular, the (implicit) introduction of evidence arguments for polymorphic flexible bounds may be undesirable in practice. Take for example our example from Section 2.1:

let $ids = [id]$ **in** $(polyL\ ids, ids \# [inc])$

where ids has type $\forall(\alpha \geq \sigma_{id}). [\alpha]$. In the System F translation, the ids value is transformed into a function that takes an evidence argument of type $\sigma_{id}^* \rightarrow \alpha$. This means that a programmer cannot

$$\begin{array}{l} \text{F-LET} \quad \frac{\begin{array}{l} (Q) \Gamma \vdash^c e_1 : \sigma_1 \\ \forall \sigma_0. \text{if } (Q) \Gamma \vdash e_1 : \sigma_0 \text{ then } (Q) \sigma_1 \sqsubseteq \sigma_0 \\ (Q) \sigma_1 \sqsubseteq^c \text{ftype}(\sigma_1) \\ (Q) \Gamma, x : \text{ftype}(\sigma_1) \vdash^c e_2 : \sigma_2 \end{array}}{(Q) \Gamma \vdash^c \text{let } x = e_1 \text{ in } e_2 : \sigma_2} \\ \text{F-ANN} \quad \frac{(Q) \Gamma \vdash^c e : \sigma_1 \quad (Q) \sigma \sqsubseteq^c \sigma_1 \quad \sigma \text{ is an F-type}}{(Q) \Gamma \vdash^c (e :: \sigma) : \sigma} \end{array}$$

Figure 15. New type rules for $\text{MLF}^=$

$$\text{ftype}(\sigma) = \text{ftype}_n(\text{nf}(\sigma))$$

where

$$\begin{array}{ll} \text{ftype}_n(\tau) & = \tau \\ \text{ftype}_n(\perp) & = \perp \\ \text{ftype}_n(\forall(\alpha = \sigma_1). \sigma_2) & = \forall(\alpha = \sigma_1). \text{ftype}_n(\sigma_2) \\ \text{ftype}_n(\forall(\alpha \geq \perp). \sigma) & = \forall(\alpha \geq \perp). \text{ftype}_n(\sigma) \\ \text{ftype}_n(\forall(\alpha \geq \forall Q. \tau). \sigma) & = \text{ftype}_n(\forall Q. \sigma[\alpha \mapsto \tau]) \end{array}$$

Figure 16. Force to F-Type

assume that **let** bound values are shared since they could be translated into functions! The same situation occurs with type classes too, and Haskell introduced the *monomorphism restriction* that rejects unannotated values that require evidence translation.

4.1 Rigid MLF

We consider a restriction of MLF, called ‘‘Rigid MLF’’ ($\text{MLF}^=$), that never requires evidence translation. To achieve this goal, all bound values in $\text{MLF}^=$ are restricted to standard System F types, i.e. types without polymorphic flexible bounds, and we call such types *F-types*

Definition 4.1.a (*F-types* (Le Botlan 2004)): A type in normal form is an *F-type* if and only if all its flexible bounds are of the form $\forall(\alpha \geq \perp)$. A type is an *F-type* if an only if its normal form is an F-type.

Types that are not F-types are exactly those with polymorphic flexible bounds, i.e. a type $\forall(\alpha \geq \sigma_1). \sigma_2$ where σ_1 is not equivalent to either a monotone τ or \perp .

Definition 4.1.b ($\text{MLF}^=$): We define $\text{MLF}^=$ as a restriction of MLF where both type annotations and **let** bound values are restricted to F-types.

In practice, this means that a programmer always writes standard F-type annotations and never MLF types with flexible bounds. Furthermore, the restriction of type annotations to F-types effectively restricts lambda bound arguments to F-types, and together with the restriction of **let** bound values to F-types this means that no evidence translation is ever necessary.

Figure 15 gives the new type rules for **let** expressions and type annotations where we have left out the System F translation for clarity. The type annotation rule F-ANN just requires that type annotations have an F-type. The rule for **let** expressions is more involved. We cannot just require that the inferred type for the **let** binding is an F-type since we would lose principal typings. For example, the term $[id]$ can have two F-types, namely $[\forall \alpha. \alpha \rightarrow \alpha]$, and $\forall \alpha. [\alpha \rightarrow \alpha]$ where neither is an instance of the other.

Instead, the rule F-LET forces the inferred type to an F-type using the function $\text{ftype}(\sigma)$ defined in Figure 16. This function systematically translates a shallow type to an F-type. The first four cases of $\text{ftype}_n(\cdot)$ are easy: monomorphic types, \perp , rigid bounds, and unconstrained bounds are already in the correct form.

The last case deals with non-trivial flexible polymorphic bounds. A concrete example of this case happens for the term $[id]$ with type $\forall(\beta \geq \forall\alpha. \alpha \rightarrow \alpha). [\beta]$. There are two possible strategies to convert such term to an F-type. First, we can instantiate and lift the quantifiers outside the bound, giving $\forall\alpha. [\alpha \rightarrow \alpha]$, which we call *variant HM*. This type is what the Hindley-Milner system would infer for this term, and therefore we chose this strategy for the definition of $\text{ftype}_n(\cdot)$. Another possibility is to keep the type polymorphic and instantiate to a rigid bound $\forall(\beta = \forall\alpha. \alpha \rightarrow \alpha). [\beta]$, which we call *variant F*. We did not choose this variant since we feel that compatibility with Hindley-Milner is more important.

It is not enough to just force the inferred type for the **let** bound expression to an F-type. Since both variant HM and variant F can be obtained by instantiation, we can always use the instantiation rule to derive either variant, losing principal typings again. The solution is simple though: in rule F-LET we specify that the type derived for the **let** expression must be the most general type and cannot be an instance thereof. This most general type is then forced to an F-type. Requiring **let** bindings to most general types is not new and similar solutions are explored for example by Leroy and Mauny (1993), Garrigue and Rémy (1999), and Vytiniotis et al. (2006).

Standard MLF type inference can be used to infer types for $\text{MLF}^=$ where **let** bound values are of course forced to F-types and type annotations restricted to F-types. Since lambda bound arguments are therefore F-types, this implies that rigid bounds are always F-Types too. Types where all rigid bounds are F-types are called *shallow types*. Standard type inference works for $\text{MLF}^=$ because all types in $\text{MLF}^=$ are shallow, and Le Botlan (2004, Section 10.2) showed that standard MLF type inference is sound and complete for shallow terms, deriving shallow principal types.

4.2 Expressiveness of $\text{MLF}^=$

In $\text{MLF}^=$ our original example

```
let ids = [id] in (polyL ids, ids ++ [inc]) -- rejected
```

is no longer accepted. All bound expressions now have simple F-types, and *ids* gets the standard Hindley-Milner type $\forall\alpha. [\alpha \rightarrow \alpha]$ and cannot be passed to *polyL* which expects a list of polymorphic identity functions. Of course, the expression *ids ++ [inc]* is well-typed. However, we could type the *polyL ids* application if we add a type annotation though:

```
let ids = [id] :: [\forall\alpha. \alpha \rightarrow \alpha] in polyL ids
```

This is well-typed since the annotation is a valid F-type and is not influenced by the $\text{ftype}(\cdot)$ coercion in F-LET. Of course, now the application *ids ++ [inc]* would be ill-typed. We can type check both applications if we inline the **let** binding and no longer try to share the *ids* value:

```
(polyL [id], [id] ++ [inc])
```

In order to type check this example, the type inferencer still assigns most general types with flexible polymorphic bounds to intermediate terms like $[id]$. The reader might worry that this still leads to evidence translation at runtime. However, all bound values have F-types and we can *never abstract over non F-types*. In practical terms this means that all evidence is always locally known and all remaining evidence translation can always be optimized away.

Having intermediate terms with non F-types is very important since it removes the need to annotate impredicative instantiations, and in practice that reduces the number of required type annotations significantly. For example, we did not need to annotate the application *polyL [id]* even though the list is impredicatively instantiated. The *only* type annotations necessary in $\text{MLF}^=$ are (1) on lambda bound arguments that are *used* polymorphically, and (2) on **let** bound expressions that contain flexible polymorphic types

that should stay polymorphic. Admittedly, the second condition is harder to explain. On the bright side, $\text{MLF}^=$ stays fully compatible with Hindley-Milner and in practice few of such annotations are necessary.

5. Related work

There have been many proposed extensions to standard Hindley-Milner type inference to support higher-ranked or impredicative types. Jones (1997) describes a system that retains a stratification between monomorphic types and type schemes and embed polymorphic types inside type constructors, where constructor application correspond to type abstraction, and pattern matching to type application.

Odersky and Läufer (1996) describe a type system for higher-ranked types where higher-ranked arguments need to be annotated with their type. Peyton Jones et al. (2007) extend this work where the type rules propagate known type information to reduce the annotation burden. Dijkstra (2005) describes a further extension that supports impredicative types and has bidirectional propagation of type annotations. In earlier work, Garrigue and Rémy (1999) describe a system where polymorphic types are embedded inside monotypes where polytypes are marked whether they are annotated or declared.

MLF (Le Botlan and Rémy 2003) goes beyond System F types to support type inference with first-class polymorphism. Several variants of MLF and their relation to other systems are explored in (Le Botlan and Rémy 2007). Rémy and Jakobowski (2007) present an efficient graph based inference algorithm for MLF that has the same complexity as normal Hindley-Milner type inference.

The work most closely related to this article is by Leijen and Löh (2005), who describe how MLF can be used with the theory of qualified types (Jones 1994) and give an algorithm that translates MLF terms to System F where evidence is passed for all bounded polymorphic types, including rigidly bound types.

5.1 A comparison with boxy types

Recently, Vytiniotis et al. (2006) introduce boxy type inference where inferred and annotated types are elegantly distinguished through boxes in the types. Just like MLF, boxy types support first-class polymorphism. Moreover, it has proved its value in practice since it has been implemented in version 6.6 of the Glasgow Haskell Compiler (GHC).

Directly comparing boxy types and MLF is difficult since their formulations differ substantially, and since MLF is strictly more expressive than boxy types (due to flexible polymorphic bounds). $\text{MLF}^=$ on the other hand is just *as expressive* as boxy types where all values can be given System F types. The set of programs accepted by boxy types is exactly the same as for $\text{MLF}^=$ modulo type annotations. A interesting way of comparing the systems would therefore be to study how type checking behaves under common small program transformations, for example, if $f\ x$ is well typed, is $\text{apply } f\ x$ also well typed?

Let us first look at inlining of **let** bindings: if $\text{let } x = e_1 \text{ in } e_2$ is accepted, is the inlining $e_2[x \mapsto e_1]$ also accepted?

binding type	HM	MLF	$\text{MLF}^=$	boxy
monomorphic	✓	✓	✓	✓
polymorphic	✓	✓	✓	×
higher-rank	-	✓	✓	×
impredicative	-	✓	✓	×

In the above table HM stands for Hindley-Milner type inference. We make a distinction for each kind of binding type, where “polymorphic” stands for a rank-1 polymorphic Hindley-Milner type. In Hindley-Milner, MLF, and $\text{MLF}^=$ we can always inline a **let** bind-

ing and still have a well-typed program. Surprisingly, this does not hold for boxy type inference – not even for rank-1 types. This is because the boxy type system relies crucially on type generalization at **let** bindings. Let’s assume that *ids* has type $[\forall\alpha. \alpha \rightarrow \alpha]$, and that we have a function *choose* of type $\forall\alpha. \alpha \rightarrow \alpha \rightarrow \alpha$. Now take the following program:

```
let f = choose [] in f ids
```

This program is well-typed in all systems (except HM of course) where the type of *f* is $\forall\alpha. [\alpha] \rightarrow [\alpha]$ (i.e. a standard Hindley-Milner type). However, if we inline the binding to *choose [] ids*, the boxy type system rejects this program since it fails to generalize the intermediate *choose []* term (and annotations do not help’).

We now take a look at the inverse of inlining: if $e_2[x \mapsto e_1]$ is accepted, is the abstraction **let** $x = e_1$ **in** e_2 also accepted?

binding type	HM	MLF	MLF ⁼	boxy
monomorphic	✓	✓	✓	✓
polymorphic	✓	✓	✓	✓
higher-rank	-	✓	ann	ann
impredicative	-	✓	×	×

In the above table ‘ann’ means that the abstracted **let** binding might need an annotation. Of all systems, only MLF allows this transformation. HM, MLF⁼, and boxy types allow it for standard Hindley-Milner types but higher-rank types may need an annotation. This happens specifically when a function returns a higher order function. Suppose *g* has type $(Int \rightarrow (\forall\alpha. \alpha \rightarrow \alpha)) \rightarrow Int$, then the expression $g(\lambda x. id)$ is well typed in MLF, MLF⁼, and boxy types. The abstraction however is only accepted in MLF:

```
let f =  $\lambda x. id$  in g f
```

In both boxy types and MLF⁼ we need to annotate the binding. Without annotation, both systems assign the type $\forall\alpha\beta. \alpha \rightarrow (\beta \rightarrow \beta)$ which is not polymorphic enough. The binding *f* must be annotated with the type $\forall\beta. \beta \rightarrow (\forall\alpha. \alpha \rightarrow \alpha)$,

Impredicative types cannot be abstracted in general in MLF⁼ and boxy types. Whenever such type requires a non-trivial flexible polymorphic bound in MLF, both MLF⁼ and boxy types fail to type it and with good reason: without evidence translation, there is not even a possible System F translation (See Section 3.1).

The final transformation we study is the *apply* transformation: if $e_1 e_2$ is well typed, is *apply* $e_1 e_2$ also accepted?

	HM	MLF	MLF ⁼	boxy
apply	✓	✓	✓	ann

In Section 2 we argued that this a particularly important transformation since it enables general abstraction over polymorphic values. Unfortunately, this property does not hold for boxy types as impredicative instantiation needs a type annotation. Take for example the application of *runST* (*return* 1) where *runST* has type $\forall s. ST s \alpha$. To type *apply runST (return 1)*, we need to annotate the full instantiation type of *apply*:

$$(apply :: ((\forall s. ST s Int) \rightarrow Int) \rightarrow (\forall s. ST s Int) \rightarrow Int)$$

runST (return 1)

A heavy burden indeed!

6. Conclusion and future work

We presented a type directed translation of MLF to System F terms. This is important in practice in order to compile MLF typed programs efficiently. In particular, when MLF is extended with qualified types it is essential to have a translation scheme to System F that can accomodate evidence translation for qualified types, and we hope that the presented translation increases adoption of MLF.

Even though MLF is very attractive as an inference system for first-class polymorphism, it is a drawback that users are exposed more complicated MLF types. Inspired by MLF⁼, we recently discovered a type system for first-class polymorphism, called HMF, which is much simpler than MLF. It uses just regular System F types and is still very expressive. For example, if $e_1 e_2$ is well typed, then *apply* $e_1 e_2$ is well typed too. We are currently studying the properties of this system (Leijen 2007).

References

- Atze Dijkstra. *Stepping through Haskell*. PhD thesis, Universiteit Utrecht, November 2005.
- Jaques Garrigue and Didier Rémy. Semi-explicit first-class polymorphism for ML. *Journal of Information and Computation*, 151:134–169, 1999.
- J.R. Hindley. The principal type scheme of an object in combinatory logic. *Transactions of the American Mathematical Society*, 146:29–60, December 1969.
- Mark P. Jones. First-class polymorphism with type inference. In *Proceedings of the Twenty Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press, January 1997.
- Mark P. Jones. *Qualified types in Theory and Practice*. Distinguished Dissertations in Computer Science. Cambridge University Press, 1994.
- Didier Le Botlan. *ML^F: Une extension de ML avec polymorphisme de second ordre et instanciation implicite*. PhD thesis, INRIA Rocquencourt, May 2004. Available in English.
- Didier Le Botlan and Didier Rémy. MLF: Raising ML to the power of System-F. In *Proceedings of the International Conference on Functional Programming (ICFP 2003)*, Uppsala, Sweden, pages 27–38. ACM Press, aug 2003.
- Didier Le Botlan and Didier Rémy. Recasting MLF. Technical Report 6228, INRIA, Rocquencourt, June 2007. URL <http://gallium.inria.fr/~remy/project/mlf>.
- Daan Leijen. HMF: simple type inference for first-class polymorphism. Work in progress, July 2007.
- Daan Leijen and Andres Löf. Qualified types for MLF. In *The International Conference on Functional Programming (ICFP’05)*. ACM Press, September 2005.
- Xavier Leroy and Michel Mauny. Dynamics in ML. *Journal of Functional Programming*, 3(4):431–463, 1993.
- Robert Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17:248–375, 1978.
- Martin Odersky and Konstantin Läufer. Putting type annotations to work. In *23th ACM Symp. on Principles of Programming Languages (POPL’96)*, pages 54–67, January 1996.
- Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. Practical type inference for arbitrary-rank types. *Journal of Functional Programming*, 17(1):1–82, 2007.
- Didier Rémy and Boris Yakobowski. A graphical presentation of MLF types with a linear-time unification algorithm. In *TLDI’07: Proc. of the 2007 ACM SIGPLAN int. workshop on Types in languages design and implementation*, pages 27–38, 2007.
- Dimitrios Vytiniotis, Stephanie Weirich, and Simon Peyton Jones. Boxy types: type inference for higher-rank types and impredicativity. In *The International Conference on Functional Programming (ICFP’06)*, September 2006.