Simple, partial type-inference for System F based on type-containment

Didier Rémy

INRIA-Rocquencourt

<http://pauillac.inria.fr/~remy>

Abstract

We explore partial type-inference for System F based on type-containment. We consider both cases of a purely functional semantics and a call-by-value stateful semantics. To enable type-inference, we require higher-rank polymorphism to be user-specified via type annotations on source terms. We allow implicit predicative type-containment and explicit impredicative type-instantiation. We obtain a core language that is both as expressive as System F and conservative over ML. Its type system has a simple logical specification and a partial type-reconstruction algorithm that are both very close to the ones for ML. We then propose a surface language where some annotations may be omitted and rebuilt by some algorithmically defined but logically incomplete elaboration mechanism.

Categories and Subject Descriptors D.3.3 [*Language Constructs and Features*]: Polymorphism

General Terms Design, Reliability, Languages, Theory, Verification

Keywords Type Inference, System F, Polymorphism, Type Reconstruction, Type Containment, Elaboration

Introduction

ML-style polymorphism has often been considered as a local optimal, offering one of the best compromises between simplicity and expressiveness. It is at the heart of two successful families of languages, Haskell and ML. In the last two decades, both Haskell and ML type systems evolved in surprisingly different directions while retaining their essence.

However, programmers are starting to feel restricted as programs become bigger and advanced programming patterns are used. The demand for first-class polymorphism has been increasing. First-class polymorphism is not only sometimes useful—it is quickly becoming unavoidable!

The reference calculus for first-class polymorphism is System F. Unfortunately, full type inference is undecidable in

System F [Wei94]. Adding first-class polymorphism to ML should not sacrifice type inference, one of the attractive aspects of ML. One approach to keeping type inference is to reduce it to second-order unification [Pfe88]. However, even so only provides with a semi-algorithm. Moreover, unintuitive explicit marks for type abstractions and applications are still required.

An alternate approach is to explicitly annotate source terms with their types, as in Church's presentation of System F. This turns type inference into a simple type checking algorithm. However, type annotations are quickly a burden to write and often become obtrusive, even for simple ML programs.

Of course, one wishes to have simultaneously the expressiveness of System F, the decidability of its explicitly typed version, and the *convenience* of ML-like type inference, if not full type inference. The idea is thus to bring System F and ML closer. This can be approached from two opposite directions. Starting with explicitly typed System F, one may perform some partial type inference so as to alleviate the need for some (but not all) type annotations. Or conversely, starting with ML, one may allow some explicit type annotations so as to reach most or all of System F programs.

The first approach is known as local type inference [PT00, OZZ01]. By definition, these solutions cover all of System F. However, they are not yet conservative over ML, that is, there remain ML programs that fail to type without annotations. In fact, local type inference allows getting rid of many silly and annoying type annotations, but fails to make *all* of them redundant, so some unintuitive annotations remain needed [HP99]. Here, we focus on the second approach. By construction, it leads to conservative extensions of ML.

The simplest method to extend ML with first-class polymorphism is to use *boxed* polymorphism. The idea is to embed first-class polymorphic types into simple types using explicit injection and projection functions. These coercions can be automatically attached to data constructors via algebraic data type definitions, which makes them simple and sometimes transparent to the user. This solution was originally proposed for existential types [LO94] and then applied to universal types [Ré94]. Boxed polymorphism is extremely simple, because ML never sees second-order types. The drawback is that boxes are rigid: they need to be declared and explicitly inserted. Furthermore, the history of the construction of a polymorphic value is recorded in its type, as the stacking of boxes.

Odersky and Laüfer [OL96] later observed that some types need not be boxed. They proposed a type system where the user can write $\lambda z : \forall \alpha. \alpha \rightarrow \alpha. t$, making the λ -bound variable z available within t at type $\forall \alpha. \alpha \rightarrow \alpha$,

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'05 September 26–28, 2005 Tallinn, Estonia.

Copyright © 2005 ACM 1-59593-064-7/05/0009...\$5.00.

as if it were let-bound. Another key feature in the work of Odersky and Laüfer is to generalize the instance relation of ML, allowing instantiation of toplevel quantifiers as in ML but also of inner quantifiers by recursively traversing arrow types, co-variantly on the right-hand side to instantiate their codomains and contra-variantly on the left-hand side to generalize their domain. This allows to keep types as polymorphic as possible and only instantiate them by need. For instance, $\lambda z : \forall \alpha. \alpha \rightarrow \alpha. \lambda y. z$ may be typed as $\forall \gamma. (\forall \alpha. \alpha \rightarrow \alpha) \rightarrow \gamma \rightarrow (\forall \alpha. \alpha \rightarrow \alpha)$ but still used with type $\forall \gamma. (\forall \alpha. \alpha \rightarrow \alpha) \rightarrow \gamma \rightarrow (\forall \beta. (\beta \rightarrow \beta) \rightarrow (\beta \rightarrow \beta))$. There is a restriction, however, that is essential to allow simple type-inference: when a type variable such as α above is instantiated, it can only be replaced with a monotype τ . This built-in restriction to predicative polymorphism is actually quite severe, as we shall see. Fortunately, it can be circumvented by keeping boxed types, which allow more explicit but impredicative type instantiation, coexisting with implicit predicative polymorphism [OL96].

Peyton-Jones *et al.* have recently extended Odersky and Laüfer’s proposal to propagate type annotations in source programs [PVWS05a], so that an external type annotation, such as $\lambda z. t : (\forall \alpha. \alpha \rightarrow \alpha) \rightarrow \tau$, behaves as if the argument z was also annotated with type $\forall \alpha. \alpha \rightarrow \alpha$. Their type system, which we hereafter refer to as PVWS, has been ingeniously tuned. It is also quite involved. The authors claim that: *adding higher-rank polymorphism to ML is so simple that every implementation of ML-style type inference should have it*. They argue that very few changes need to be made to adapt an implementation of ML type inference to predicative higher-rank polymorphism. While we accept their claim, we find their implementation-based demonstration unconvincing and their specification too algorithmic and unintuitive. Indeed, PVWS mixes several features that are often considered conflicting: (1) ML-like type inference; (2) a form of contravariant subsumption, and (3) propagation of user-provided annotations. Type systems that mix (1) and (2) usually exploit explicit constraints to keep principal types, which PVWS does not; (3) enforces a strict order in which type checking must be performed, which usually stand in the way when inferring principal types. Such an unusual combination of features does not imply misconception, of course—and indeed, PVWS has been carefully designed. However, it should then be studied with a lot of attention [PVWS05b].

Contributions

In this paper, we investigate partial type inference based on type containment and first-order typing constraints. Our primary goal is to bring further evidence to Peyton-Jones and Shields’ claim and thus to provide more insight to PVWS. Another goal is also to explore the design space and, in particular, how far one can go within an ML-like type-inference framework. We have at least four different contributions: A preliminary, independent result is the soundness of a variant of $F^?$ for a call-by-value stateful semantics. Our main contribution is the core language F_{ML} that brings down System F to the level of ML—it has a simple logical specification and a sound and complete first-order type inference algorithm. A secondary contribution is the surface language $F_{ML}^?$ and our proposal to split partial type inference in $F_{ML}^?$ into a composition of two separate orthogonal phases: an algorithmic elaboration process into F_{ML} , followed by type inference into F_{ML} . Our approach leaves room for variations in the definition of type-containment relations, which controls the expressiveness of the core language, and in the elaboration

process. A side contribution is also the exploration of the design space: our two-phase decomposition has advantages but also some limitations.

1. System $F(\leq)$ and some instances

System F is the canonical system for λ -calculus with second-order polymorphism. However, it may be presented in two ways. In Church’s view, terms are explicitly typed, including type annotations on λ -abstractions, and type applications. Therefore, type-checking only needs to verify that user-provided type information is correct with respect to the typing rules. On the opposite, in Curry’s view, terms are untyped and type-checking must infer the types that should be assigned to formal parameters in λ -abstractions as well as the places where type abstractions and type applications should be inserted. The two views can be reconciled by considering type inference for terms in Curry’s style as *elaborating* an explicitly typed term in Church’s style.

1.1 Terms and types

We assume given a finite or denumerable collection of constants (ranged over by c) and a denumerable collection of variables (ranged over by z and y). Constants and variables form identifiers, which we write x . Terms, ranged over by t , are those of the untyped λ -calculus, *i.e.*, identifiers x , abstractions $\lambda z. t$, or applications $t t'$. Application has higher priority than abstraction and is left-associative *i.e.* $\lambda z. \lambda z'. t t' t''$ stands for $\lambda z. (\lambda z'. ((t t') t''))$.

We assume a denumerable collection of type variables (ranged over by α). Types, ranged over by σ , are type variables, arrow types $\sigma_1 \rightarrow \sigma_2$, and polymorphic types $\forall \alpha. \sigma$. The $\forall$ symbol acts as a binder for α in σ . Free type variables in a type σ , which we write $\text{ftv}(\sigma)$, are defined accordingly. We always consider types equal modulo renaming of bound-type variables. The scope of $\forall$ -quantification extends to the right as much as possible and arrows are right-associative. That is, $\forall \alpha. \sigma_1 \rightarrow \sigma_2 \rightarrow \sigma_3$ stands for $\forall \alpha. (\sigma_1 \rightarrow (\sigma_2 \rightarrow \sigma_3))$.

A type variable α occurs positively in α . If α occurs positively (resp. negatively) in σ , then it occurs positively (resp. negatively) in $\sigma' \rightarrow \sigma$ and $\forall \beta. \sigma$ when β is distinct from α , and negatively (resp. positively) in $\sigma \rightarrow \sigma'$. We write $\text{ftv}^+(\sigma)$ (resp. $\text{ftv}^-(\sigma)$) the sets of type variables that occur positively (resp. negatively) in σ . We write $\text{ftv}^\dagger(\sigma)$ the set $\text{ftv}^+(\sigma) \setminus \text{ftv}^-(\sigma)$, which we call *non-negative free type variables*.

A relation $\mathcal{R}$ on types is *structural* if it is reflexive, transitive and closed under the following properties: if $\sigma_1 \mathcal{R} \sigma_2$, then $\forall \alpha. \sigma_1 \mathcal{R} \forall \alpha. \sigma_2$ (S-ALL), $\sigma_2 \rightarrow \sigma \mathcal{R} \sigma_1 \rightarrow \sigma$ (S-CONTRA) and $\sigma \rightarrow \sigma_1 \mathcal{R} \sigma \rightarrow \sigma_2$ (S-COV). A *congruence* is a structural equivalence relation.

Let $\approx$ be the smallest congruence that allows commutation of adjacent binders, *i.e.* $\forall \alpha. \forall \beta. \sigma \approx \forall \beta. \forall \alpha. \sigma$, and removal of redundant binders, *i.e.* $\forall \alpha. \sigma \approx \sigma$ whenever α is not free in σ . Most relations on types we will consider will be subrelations of $\approx$. Therefore, we could usually treat types equal modulo $\approx$. However, we prefer not to do so and treat $\approx$ explicitly when needed. A canonical form for $\approx$ -equivalence is one without redundant quantifiers and where adjacent quantifiers are listed in order of apparition. Canonical forms are unique (up to renaming of bound-type variables, but we treat such types as equal). We write $\text{canon}(\sigma)$ for the canonical form of σ . Subterms of types in canonical forms are also in canonical form.

Figure 1. Typing rules for $F(\leq)$

$\frac{\text{VAR} \quad x : \sigma \in \Gamma}{\Gamma \vdash x : \sigma}$	$\frac{\text{INST} \quad \Gamma \vdash t : \sigma' \quad \sigma' \leq \sigma}{\Gamma \vdash t : \sigma}$
$\frac{\text{GEN} \quad \Gamma \vdash t : \sigma \quad \alpha \notin \text{ftv}(\Gamma)}{\Gamma \vdash t : \forall \alpha. \sigma}$	$\frac{\text{FUN} \quad \Gamma, z : \sigma_2 \vdash t : \sigma_1}{\Gamma \vdash \lambda z. t : \sigma_2 \rightarrow \sigma_1}$
$\frac{\text{APP} \quad \Gamma \vdash t_1 : \sigma_2 \rightarrow \sigma_1 \quad \Gamma \vdash t_2 : \sigma_2}{\Gamma \vdash t_1 t_2 : \sigma_1}$	

We write $\bar{\alpha}$ for tuples of variables $\alpha_1, \dots, \alpha_n$ (and more generally $\bar{e}$ for a sequence of elements of the syntactic class e) and $\forall \bar{\alpha}. \sigma$ for $\forall \alpha_1 \dots \forall \alpha_n. \sigma$. We write τ for types that do not contain any universal quantifiers, called monotypes; we write ρ for types that do not have any outermost universal quantifiers. We write $\sigma[\bar{\sigma}/\bar{\alpha}]$ for the simultaneous capture-free substitution of types $\bar{\sigma}$ for variables $\bar{\alpha}$ in σ .

1.2 The generic system $F(\leq)$

Let $F(\leq)$ be the type system for the λ -calculus defined in Figure 1 and parameterized by a binary relation on types $\leq$ called *type containment*. A typing context Γ is a finite mapping from program identifiers to types. The empty environment is written $\emptyset$. We write $\Gamma, x : \sigma$ for the environment that binds x to σ and other program variables as Γ . We write $x : \sigma \in \Gamma$ to mean that Γ maps x to σ . We write $\text{ftv}(\Gamma)$ for $\bigcup_{x:\sigma \in \Gamma} \text{ftv}(\sigma)$. Typing judgments are triples of the form $\Gamma \vdash t : \sigma$. We write $t \in F(\leq)$ if (and only if) there exists a typing environment Γ and a type σ such that $\Gamma \vdash t : \sigma$. We may write $\Gamma \vdash_X t : \sigma$ when we need to remind that typing judgments refer to the type system X .

Rules VAR, APP, FUN are the same as in the simply typed λ -calculus, except that types may here have quantifiers. Rule GEN allows generalizing the type of an expression over a type variable that does not appear in the environment. The really interesting Rule is INST, which allows replacing the type of an expression with one of its instances, where the notion of instance is determined by the relation $\leq$.

System F is obtained by taking $\leq^F$ for $\leq$, which is defined as the smallest relation that satisfies the unique axiom SUB of Figure 2. (We use $\leq$ to range over arbitrary relations, while symbols $\leq, \leq^F, \leq^\eta$, etc. denotes specific relations.) As early as 1984, Mitchell proposed to replace the instance relation $\leq^F$ by a more general relation $\leq^\eta$, called *type containment*, which is defined as the smallest relation satisfying the rules of Figure 2 [Mit88]. He also showed that a term t is typable in $F(\leq^\eta)$ if and only if there exists a term t' typable in F that is η -equal to t . That is, $F(\leq^\eta)$ is the typed closure of F by η -conversion, hence $F(\leq^\eta)$ is usually called F^η . It follows that the relation $\sigma \leq^\eta \sigma'$ holds if and only if there exists a term t of type $\sigma \rightarrow \sigma'$ in F that η -reduces to the identity [Mit88]. Such terms are called *coercion* functions.

We write $\equiv$ for the kernel of $\leq$. Note that $\leq^\eta$ is a structural relation. It is also a subrelation of $\approx$. Conversely, $\leq^F$ is not structural and it is not a subrelation of $\approx$. We write $\leq_{\approx}^F$ for the composition $\leq^F \circ \approx$, which is also equal to $\approx \circ \leq^F \circ \approx$ (but not to $\approx \circ \leq^F$). The relation $\leq_{\approx}^F$ is still not structural, but it is a subrelation of $\approx$. The language $F(\leq_{\approx}^F)$ is not System F *per se*: its typing relation $\vdash_{\leq_{\approx}^F}$ is larger than

Figure 2. Containment rules for $\leq^\eta$

$\frac{\text{SUB} \quad \beta \notin \text{ftv}(\forall \bar{\alpha}. \sigma)}{\forall \bar{\alpha}. \sigma \leq \forall \bar{\beta}. \sigma[\bar{\sigma}/\bar{\alpha}]}$	$\frac{\text{TRANS} \quad \sigma \leq \sigma' \quad \sigma' \leq \sigma''}{\sigma \leq \sigma''}$
$\frac{\text{ARROW} \quad \sigma'_1 \leq \sigma_1 \quad \sigma_2 \leq \sigma'_2}{\sigma_1 \rightarrow \sigma_2 \leq \sigma'_1 \rightarrow \sigma'_2}$	$\frac{\text{ALL} \quad \sigma \leq \sigma'}{\forall \alpha. \sigma \leq \forall \alpha. \sigma'}$
$\frac{\text{DISTRIB} \quad \forall \alpha. \sigma \rightarrow \sigma' \leq (\forall \alpha. \sigma) \rightarrow \forall \alpha. \sigma'}{\quad}$	

$\vdash_{\leq^F}$. However, both languages have the same set of typable terms.

About Rule DISTRIB It is actually possible to limit uses of Rule DISTRIB to cases where $\alpha \notin \text{ftv}(\sigma)$ and furthermore $\alpha \in \text{ftv}^-(\sigma')$ without affecting the type-containment relation $\leq^\eta$. We refer to these two limited use of DISTRIB as DISTRIB-RIGHT and DISTRIB-RIGHT-NEG, respectively. Moreover, Rule DISTRIB-RIGHT is reversible, *i.e.* $\forall \alpha. \sigma' \leq \forall \alpha. \sigma \rightarrow \sigma'$ whenever $\alpha \notin \text{ftv}(\sigma)$. So we actually have $\sigma \rightarrow \forall \alpha. \sigma' \equiv^\eta \forall \alpha. \sigma \rightarrow \sigma'$ whenever $\alpha \notin \text{ftv}(\sigma)$. Types can be put in *prenex-form* by repeatedly applying the rewriting rule $\sigma \rightarrow \forall \alpha. \sigma' \rightsquigarrow \forall \alpha. \sigma \rightarrow \sigma'$ when $\alpha \notin \text{ftv}(\sigma)$ in any type context (the side-condition can always be satisfied modulo appropriate renaming). We write $\text{prf}(\sigma)$ for the prenex-form equivalent to σ , which is unique up to $\approx$. We also write $\text{prf}(\Gamma)$ for $\text{prf}(\cdot)$ composed with the mapping Γ . For any judgment $\Gamma \vdash_{F(\leq^\eta)} t : \sigma$, the judgment $\text{prf}(\Gamma) \vdash_{F(\leq^\eta)} t : \text{prf}(\sigma)$ also holds. Moreover, there exists a derivation of this judgment that uses only types in prenex-form in typing judgments. However, subderivations of $\leq^\eta$ -type-containment judgments may require the use of types not in prenex-form, *e.g.* in applications of Rule SUB.

1.3 Expressiveness of $F(\leq^\eta)$

Coercion functions allow the extrusion of quantifiers and deep type-specialization. This is convenient in a type inference context because instantiation may be performed a posteriori. For instance, the term t_K equal to $\lambda z. \lambda y. y$ can be assigned either type $\forall \alpha. \alpha \rightarrow \forall \beta. \beta \rightarrow \beta$, giving the sub-expression $\lambda y. y$ a polymorphic type, or type $\forall \alpha. \beta. \alpha \rightarrow \beta \rightarrow \beta$. Coercions make both types inter-convertible. Here, either one is actually a principal type of t_K in $F(\leq^\eta)$, *i.e.* all other types may be obtained by type-containment [Mit88]. By comparison, there is no type of t_K in System F of which all other types would be $\leq^F$ -instances. For this reason, it has been suggested that $F(\leq^\eta)$ would be a better candidate for type inference [Mit88]—before full type inference was shown to be undecidable.

The example above shows that *some terms have more types* in $F(\leq^\eta)$. One may also expect *more terms to have types* in $F(\leq^\eta)$, although we are not aware of any term of $F(\leq^\eta)$ that has been proved not to be in F . However, the additional expressive power, if any, does not seem to be very significant.

1.4 Soundness of $F(\leq)$

A typing relation $\Gamma \vdash_X t : \sigma$ and a reduction relation $\longrightarrow$ on expressions enjoy the subject reduction property if reduction preserves typing. That is, *if* $\Gamma \vdash_X t : \sigma$ *and* $t \longrightarrow t'$, *then* $\Gamma \vdash_X t' : \sigma$. Of course, this result depends on the

particular choice of X and of the reduction relation $\longrightarrow$. In the following, we will consider several type systems $F(\leq)$ where $\leq$ is a subrelation of $\leq^\eta$ (or $\leq^{\eta-}$). Type soundness for all of these systems will then follow from type soundness for $F(\leq^\eta)$.

For an effect-free semantics and in the absence of constants, the relation $\longrightarrow$ is the transitive closure of the one-step reduction $C[(\lambda z. t_1) t_2] \longrightarrow C[t_1[t_2/z]]$ for arbitrary contexts C . Subject reduction is known to hold in System F. It then easily follows from the definition of $F(\leq^\eta)$ in terms of F that subject reduction also holds in $F(\leq^\eta)$.

Subject reduction is only halfway key to type soundness. It remains to verify that any well-typed program that is not a value can be further evaluated, that is, the *progress* Lemma. Progress is trivial in the absence of constants. Otherwise, one can show that both subject reduction and progress lemmas hold under some hypotheses ensuring subject reduction and progress for reduction involving constants.

1.5 A stateful sound variant of $F(\leq^\eta)$

For sake of simplicity, we focus on the treatment of the store and references rather than general side-effecting operations. It is well-known that combining polymorphism and mutable store requires some care. Because System $F(\leq^\eta)$ is highly polymorphic, we must be even more careful. In this section we define a variant $F^v(\leq^{\eta-})$ of $F(\leq^\eta)$ that is safe when equipped with a call-by-value semantics with side effects.

One way to ensure type soundness in the presence of side effects is to give type abstraction a call-by-name semantics. That is, a polymorphic expression must be reevaluated every time it is specialized. In an explicitly typed calculus, polymorphism is introduced by a type abstraction $\Lambda \alpha. t$, which freezing the evaluation of t and every use of polymorphism is obtained by a type application $t[\sigma]$, which evaluates the whole expression t (after occurrences of α have been replaced by σ). This call-by-name semantics of polymorphism affects all expressions, whether or not their evaluation may produce side-effects. (Even in the absence of side effects, this modifies sharing of expressions and may change the complexity of computation significantly.) A variant of this solution is to restrict generalization to values, which is known as the value-restriction [Wri95]. More precisely, we may define a class of *non-expansive* expressions that includes at least values and variables, but may also include some other expressions—a formal definition is given below. This solution is now in use in most implementations of ML, although many of them still use a more restrictive definition of non-expansiveness—for no good reason.

We shall thus, as in ML, restrict Rule GEN so that it only applies when the expression t is non-expansive. However, this restriction alone is not sound in $F(\leq^\eta)$, because it can be bypassed by type-containment. For example, consider the expression t_m defined as $\lambda y. t_r$ where t_r is $\text{ref } (\lambda z. z)$. As in ML, t_m can be assigned type $\forall \alpha. \beta \rightarrow \text{ref } (\alpha \rightarrow \alpha)$. Intuitively, this is correct since then an application $t_m t_0$ will then have the monomorphic type $\text{ref } (\alpha \rightarrow \alpha)$, or more precisely, will have type $\text{ref } (\tau \rightarrow \tau)$ for any type τ . The key is that t_m should not have type $\beta \rightarrow \forall \alpha. \text{ref } (\alpha \rightarrow \alpha)$, since otherwise $t_m t_0$ would return a reference cell that could be treated as a polymorphic value. This type seems to be disallowed, since the value restriction prohibits generalization of the type of t_r . Unfortunately, type containment allows replacing the correct type $\forall \alpha. \beta \rightarrow \text{ref } (\alpha \rightarrow \alpha)$ of t_m with the unsound type $\beta \rightarrow \forall \alpha. \text{ref } (\alpha \rightarrow \alpha)$. The problem can be traced back to rule DISTRIB-RIGHT-NEG, which allows

bypassing the value restriction, and is thus unsound in the presence of side-effects.

The solution is thus both to restrict rule GEN to non-expansive expressions and invalidate Rule DISTRIB-RIGHT-NEG. Instances of Rule DISTRIB that are derivable from other rules are both harmless and useful. This is the case in particular when type variable α only occurs positively in $\text{ftv}(\sigma)$. We refer to this special instance of DISTRIB as DISTRIB-RIGHT-POS. It is actually remarkable, that Rule DISTRIB-RIGHT-POS, which is valid in $F^v(\leq^{\eta-})$ leads to the *enhanced* value restriction as defined by Garrigue [Gar04], rather than the standard value-restriction [Wri95].

Store semantics

We assume given a denumerable collection of memory locations $m \in \mathcal{M}$ and restrict constants to the following cases:

$c ::= c^+ \mid c^-$	Constants
$c^+ ::= m$	Constructors
$c^- ::= (\text{ref } \cdot) \mid (!\cdot) \mid (\cdot := \cdot)$	Destructors

Each constant c comes with a fixed arity $a(c)$. The arity of memory locations is 0. The arities of $!$, $\text{ref } \cdot$, and $\cdot := \cdot$ are, respectively 0, 1, and 2, as suggested by the notation.

By lack of space, we cannot describe the call-by-value semantics with store. We refer the reader to an extended version of this paper [Rém05] or to a detailed presentation of type-constraints [PR05].

Typing rules with side effects

Non-expansive expressions $u \in \mathcal{U}$ are defined as follows:

$$u ::= z \mid \lambda z. t \mid (\lambda z_1 \dots \lambda z_k. u) u_1 \dots u_k$$

$$\mid c^+ u_1 \dots u_k, \quad k \leq a(c)$$

$$\mid c^- u_1 \dots u_k, \quad k < a(c)$$

By construction, reduction of non-expansive expressions cannot extend the store. Remark that non-expansive expressions may contain free variables. Substituting non-expansive expressions for free variables in non-expansive expressions yields again non-expansive expressions.

We introduce a new *invariant* type symbol ref of arity 1 to classify expressions that evaluate to store locations. Invariance mean that an occurrence of a variable in a type $\text{ref } \sigma$ is both positive and negative and thus never in $\text{ftv}^\dagger(\sigma)$. Moreover, a structural relation $\leq$ must now validate the following rule:

$$\frac{\text{REF} \quad \sigma \leq \sigma' \quad \sigma' \leq \sigma}{\text{ref } \sigma \leq \text{ref } \sigma'}$$

We write $\leq^{\eta-}$ for the smallest relation $\leq$ that satisfies rules SUB, TRANS, ARROW, ALL of Figure 2 and rule REF.

We also restrict generalization to non-expansive expressions *or* to non-negative type variables.

$$\frac{\text{GEN}^v \quad \Gamma \vdash t : \sigma \quad \alpha \notin \text{ftv}(\Gamma) \quad t \in \mathcal{U} \vee \alpha \in \text{ftv}^\dagger(\sigma)}{\Gamma \vdash t : \forall \alpha. \sigma}$$

Let $F^v(\leq)$ be the type system $F(\leq)$ where Rule GEN has been replaced by Rule GEN^v . Types for constants are defined by an initial environment Γ_0 that contains exactly:

$$\begin{aligned} \text{ref } \cdot &: \quad \forall \alpha. \alpha \rightarrow \text{ref } \alpha \\ !\cdot &: \quad \forall \alpha. \text{ref } \alpha \rightarrow \alpha \\ \cdot := \cdot &: \quad \forall \alpha. \text{ref } \alpha \rightarrow \alpha \rightarrow \alpha \end{aligned}$$

With these restrictions, we still have $\Gamma_0 \vdash_{F^v(\leq^\eta-)} t_m : \forall \alpha. \text{unit} \rightarrow \text{ref}(\alpha \rightarrow \alpha)$, but not $\Gamma_0 \vdash_{F^v(\leq^\eta-)} t_m : \text{unit} \rightarrow \forall \alpha. \text{ref}(\alpha \rightarrow \alpha)$ any longer.

THEOREM 1. *The language $F^v(\leq^\eta-)$ with references is sound.*

By lack of space, we refer to [Rém05] for a more formal statement of this result and its proof. Type soundness is shown in a standard way by combination of subject reduction and progress lemmas.

1.6 Predicative fragments $F(\leq_p^F)$ and $F(\leq_p^\eta)$

The predicative fragment of System F is the system $F(\leq_p^F)$ obtained by restricting the instance relation $\leq^F$ so that only monotypes can be substituted for type variables¹. That is, $\leq_p^F$ is the smallest relation satisfying the following rule:

$$\frac{\text{SUB}_p \quad \beta \notin \text{ftv}(\forall \bar{\alpha}. \sigma)}{\forall \bar{\alpha}. \sigma \leq \forall \bar{\beta}. \sigma[\bar{\tau}/\bar{\alpha}]}$$

The predicative fragment of $F(\leq^\eta)$ can be defined either (1) as the typed-closure of $F(\leq_p^F)$ by η -conversion, or (2) as $F(\leq_p^\eta)$ where $\leq_p^\eta$ is the smallest relation satisfying all rules of Figure 2 except Rule SUB, which is replaced by Rule SUB_p . Fortunately, definitions (1) and (2) are equivalent.

The type system $F(\leq_p^\eta)$ is safe, as it is a subset of $F(\leq^\eta)$. This does not imply subject reduction. However, one may show independently that subject reduction holds in both $F(\leq_p^F)$ and in $F(\leq_p^\eta)$.

Replacing DISTRIB by DISTRIB-RIGHT in the definition of $\leq_p^\eta$ does not change the relation itself. This is not a consequence of the similar result for the impredicative relation $\leq^\eta$. Indeed, further replacing DISTRIB-RIGHT by DISTRIB-RIGHT-NEG would lead to a weaker relation than $\leq_p^\eta$. Hence, in the case of a call-by-value stateful semantics the language $F^v(\leq_p^\eta)$ could be safely enriched with a version of DISTRIB that requires α in $\text{ftv}^\dagger(\sigma') \setminus \text{ftv}(\sigma)$. We refer to this Rule as DISTRIB-RIGHT-POS.

1.7 Expressiveness of the predicative fragments

The restriction of Systems F or $F(\leq^\eta)$ to their predicative fragments comes with a significant loss of expressiveness. It means that polymorphism is limited to simple types. That is, programs can manipulate values of several types that differ in their *monomorphic* parts but must have the same *polymorphic* shape. For instance, $\text{int} \rightarrow \text{int}$ and $\text{bool} \rightarrow \text{bool}$ can be two specializations of a predicative type, but $(\forall \alpha. \alpha \rightarrow \alpha) \rightarrow (\forall \alpha. \alpha \rightarrow \alpha)$ and $\text{int} \rightarrow \text{int}$ cannot.

The restriction of rule SUB to rule SUB_p suggests a large family of counter-examples. For instance, a polymorphic function of type $\forall \alpha. \sigma \rightarrow \sigma'$ can only be used at types $\sigma[\tau/\alpha]$ where τ is a monotype. Thus, one may need as many copies of the function as there are different shapes σ at which the function needs to be used. As a particular case, the **apply** function $\lambda f. \lambda z. f z$ does not have a most general type, but as many types as there are polymorphic shapes for the type of the function f . This is a general situation that also applies to iterators such as **iter**, **map**, **fold**, *etc.*

Unsurprisingly, the well-known encoding of existential types with universal types in System F is also severely limited in the predicative fragments. A value v of an existential type $\exists \alpha. \sigma$ can be encoded as a function $\forall \beta. (\forall \alpha. \sigma \rightarrow \beta) \rightarrow \beta$.

¹ One may also consider a stratification of predicative fragments—see [Lei91]

Figure 3. Type containment rules for $\leq_p^{\eta-}$

REFL $\alpha \leq \alpha$	ARROW $\frac{\sigma'_1 \leq \sigma_1 \quad \sigma_2 \leq \sigma'_2}{\sigma_1 \rightarrow \sigma_2 \leq \sigma'_1 \rightarrow \sigma'_2}$	
ALL-I $\frac{\sigma \leq \sigma' \quad \alpha \notin \text{ftv}(\sigma)}{\sigma \leq \forall \alpha. \sigma'}$	ALL-E $\frac{\sigma[\tau/\alpha] \leq \rho}{\forall \alpha. \sigma \leq \rho}$	

The use of value v under the name x inside some term t is encoded as the application of v to $\lambda x. t$. This application can be typed in F, giving x the polymorphic type $\forall \alpha. \sigma \rightarrow \sigma'$ and instantiating β with σ' . In the predicative fragment, however, σ' must be a monotype. Hence, the restriction of existential types to the predicative fragment only allows expressions that manipulate (encoded) existential types to return monotypes. Moreover, the type hidden by the abstraction may only be a monotype. In explicit System F, the expression **pack** t as $z : \exists \beta. \sigma$, where $\sigma[\sigma'/\beta]$ is the type of t , is encoded as $\Lambda \alpha. \lambda f : \forall \beta. \sigma. f \sigma' t$ where σ' is the abstract part of the type. Hence, σ' must be a monotype τ in $F(\leq_p^F)$. As a corollary, encodings of objects [BCP99] do not work well in the predicative fragments: the object state cannot be hidden any longer as soon as it references an object.

Indeed, we may exhibit terms that are typable in F but not in the predicative fragment $F(\leq_p^F)$. One such term $(\lambda z. z y z) (\lambda z. z y z)$ is due to Pawel Urzyczyn, according to Leivant [Lei91]. Stratified versions of System F are more expressive than System $F(\leq_p^F)$ [Lei91]. However, they would not ease type inference.

2. Type inference in the language F_{ML}

Full type inference is undecidable in both System F [Wel94] and $F(\leq^\eta)$ [Wel96]—the type-containment relation $\leq^\eta$ is itself undecidable [TU02, Wel95]. To the best of our knowledge it is not known whether type inference is decidable for the predicative fragments. However, Harper and Pfenning remarked that the reduction of second-order unification to type inference for System F with explicit placeholders for type abstractions and applications also applies to the predicative fragment [Pfe88]. Still, we do not know whether full type inference for $F(\leq_p^\eta)$ is decidable, even though predicative type containment $\leq_p^\eta$ is itself decidable [PVWS05b].

We also consider the restriction $\leq_p^{\eta-}$ of $\leq_p^\eta$ obtained by removing Rule DISTRIB from the definition of $\leq_p^\eta$. That is, we define $\leq_p^{\eta-}$ as the smallest relation that satisfies rules SUB_p , TRANS, ARROW and ALL. A consequence of the removal of rule DISTRIB is that $\leq_p^{\eta-}$ has a very simple syntax-directed presentation given by rules of Figure 3.

LEMMA 1. $\leq_p^{\eta-}$ is also the smallest relation $\leq$ that satisfies the rules of Figure 3.

LEMMA 2 (Stability of $\leq_p^{\eta-}$ by substitution). *If $\sigma \leq_p^{\eta-} \sigma'$ then $\sigma[\bar{\tau}/\bar{\alpha}] \leq_p^{\eta-} \sigma'[\bar{\tau}/\bar{\alpha}]$.*

Actually, the two relations $\leq_p^\eta$ and $\leq_p^{\eta-}$ coincide on types in prenex form, as shown by the following lemma due to Peyton-Jones, Vytiniotis, Weirich and Shields [PVWS05b].

LEMMA 3 (Peyton-Jones *et al.*). $\sigma \leq_p^\eta \sigma'$ if and only if $\text{prf}(\sigma) \leq_p^{\eta-} \text{prf}(\sigma')$.

This provides us with an algorithm for testing $\leq_p^\eta$ by first projecting both types to their prenex forms and then testing for $\leq_p^{\eta-}$ (or by deriving a direct algorithm from this composition [PVWS05b]). Although the restriction $\leq_p^{\eta-}$ of $\leq_p^\eta$ is primarily introduced to ensure soundness in $F^v(\leq_p^{\eta-})$, we may also consider the language $F(\leq_p^{\eta-})$ in the absence of side-effects. One advantage is that type inference for $F(\leq_p^{\eta-})$ is closer to type inference for ML, especially in its treatment of generalization. Moreover, working with $\leq_p^{\eta-}$ is slightly more convenient than $\leq_p^\eta$ for type inference. We start with this simpler case and will consider $F(\leq_p^\eta)$ and $F^v(\leq_p^\eta)$ in sections 2.5 and 2.6.

2.1 Terms with type annotations: $F_{ML}(\leq)$

We leave the relation $\leq$ a parameter of $F_{ML}(\leq)$ so as to share some of the developments between $\leq_p^{\eta-}$ and $F_{ML}(\leq_p^\eta)$. However, the intention is that the type-containment relation $\leq$ be predicative, *i.e.* a subrelation of $F_{ML}(\leq_p^\eta)$.

Thus, the type-containment relations $\leq$ that we shall consider are decidable, and so is $\leq$ -constraint resolution, as we shall see below. However, this is not sufficient to perform type inference. In order to avoid *guessing* polymorphic types, we now extend terms with type annotations. An annotation is a type scheme σ whose free variables $\bar{\alpha}$ are locally bound; we write $\exists \bar{\alpha}. \sigma$. While σ represents the explicit (usually second-order) type information, the existentially bound variables $\bar{\alpha}$ represent the implicit type information, which will be inferred. That is, type annotations are partial, which is important to leave room for inference. As a particular case, the trivial annotations $\exists \bar{\alpha}. \alpha$ let us infer monomorphic types, as in ML. We require that type annotations be closed. This is only for the sake of simplicity. One could easily allow type annotations to have free type variables, provided these are explicitly bound somewhere in the program. This mechanism, sometimes known as *scoped type variables*, is discussed by Peyton Jones and Shields [PS03] and by Pottier and Rémy [PR]. We let θ range over type annotations.

The syntax of $F_{ML}(\leq)$ is as follows:

$$t ::= x \mid \lambda z. t \mid t_1 (t_2 : \theta) \mid \text{let } z = (t_1 : \theta) \text{ in } t_2$$

As usual, expressions may be identifiers x , abstractions $\lambda z. t$, or applications $t_1 (t_2 : \theta)$. However, arguments of applications must now always be explicitly annotated (although annotations may be trivial). The intuition for annotating the argument of applications is to avoid guessing the type of the argument. Indeed, knowing the type of the result of an application does not tell anything about the type of its argument. We also allow let-bindings $\text{let } z = (t_1 : \theta) \text{ in } t_2$, which mean $(\lambda z. t_2) (t_1 : \theta)$, but which will be typed in a special way, as in ML. As for applications, an annotation is required for the expected type of t_1 , since it cannot be deduced from the expected type of the whole expression.

2.2 Typing rules

Typing rules for $F_{ML}(\leq)$ are described in Figure 4. They use judgments of the form $\Gamma \vdash t : \sigma$ as for $F(\leq)$. These judgments should be read as checking judgments, where Γ , t and σ are given. Rules VAR, INST, GEN, and FUN are all taken from $F(\leq)$. However, as mentioned above, the intention is that $\leq$ in the right premise of Rule INST be chosen as an instance of $\leq_p^\eta$, so as to ensure predicativity. That is, the polymorphic structure of polymorphic types can only be instantiated by monotypes. Note that, as opposed to ML, rule FUN allows its argument to be polymorphic. Despite the appearance,

Figure 4. Typing rules for $F_{ML}(\leq)$

VAR	INST	GEN	FUN
APP _A			
$\Gamma \vdash t_1 : \sigma_2[\bar{\tau}/\bar{\beta}] \rightarrow \sigma_1$		$\Gamma \vdash t_2 : \sigma_2[\bar{\tau}/\bar{\beta}]$	
		$\Gamma \vdash t_1 (t_2 : \exists \bar{\beta}. \sigma_2) : \sigma_1$	
LET _A			
$\Gamma \vdash t_1 : \forall \bar{\alpha}. \sigma_1[\bar{\tau}/\bar{\beta}]$		$\Gamma, z : \forall \bar{\alpha}. \sigma_1[\bar{\tau}/\bar{\beta}] \vdash t_2 : \sigma_2$	
		$\Gamma \vdash \text{let } z = (t_1 : \exists \bar{\beta}. \sigma_1) \text{ in } t_2 : \sigma_2$	

this does not imply guessing polymorphism, as long as the expected type $\sigma_2 \rightarrow \sigma_1$ is given.

Rule APP_A (the subscript is to remind the *annotation*) is not quite the one of $F(\leq)$. We must of course take the annotation of the argument into account: the type of t_2 , which is also the domain of the type of t_1 , must be an instance of the annotation. The annotation avoids guessing the polymorphic parts of the expected type of t_1 and t_2 that cannot be deduced from the expected type of the application. When the annotation is the trivial one, $\sigma_2[\bar{\tau}/\bar{\alpha}]$ must be a monotype τ_2 , to be guessed—not a problem, since it is a monotype. More generally, only monotypes $\bar{\tau}$ need to be guessed in applications. Rule LET_A is taken from ML, except that the annotation on t_1 is used to build the expected type of t_1 , up to first-order instantiation, much as in Rule APP_A.

Expressiveness. Apart from annotations, the language $F_{ML}(\leq)$ is as expressive as $F(\leq)$ for any relation $\leq$. Formally, let the type erasure of a term t be the λ -term obtained by dropping all type annotations and replacing all let-bindings $\text{let } z = t_1 \text{ in } t_2$ by $(\lambda z. t_2) t_1$. Then, we have $t \in F(\leq)$ if and only if there exists an expression $t' \in F_{ML}(\leq)$ whose erasure is t .

Syntactic sugar. While annotations are always required in applications and let bindings, the *trivial annotations* may be used. This is not exactly an absence of annotation, since the annotation is mandatory and a trivial annotation always stands for a monomorphic type. In particular, an expression where all annotations are trivial will be typed as in ML. For convenience, we may allow $\text{let } z = t_1 \text{ in } t_2$ and $t_1 t_2$ as syntactic sugar for $\text{let } z = (t_1 : \exists \bar{\alpha}. \alpha) \text{ in } t_2$ and $t_1 (t_2 : \exists \bar{\alpha}. \alpha)$, respectively.

The language $F_{ML}(\leq)$ is a conservative extension to ML. That is, if Γ is an ML environment and t is an ML expression, *i.e.* both without any non trivial annotation, then $\Gamma \vdash_{ML} t : \tau$ if and only if $\Gamma \vdash_{F_{ML}(\leq)} t : \tau$.

The language does not allow arbitrary expressions to carry type annotations. However, we can define $(t : \theta)$ as syntactic sugar for $\text{let } z = (t : \theta) \text{ in } z$.

Example. The expression $\lambda z. z$ is well-typed, and can be given type $\tau \rightarrow \tau$, for any type τ , as in ML. It may also be given types $\forall \alpha. \alpha \rightarrow \alpha$, $(\forall \alpha. \alpha) \rightarrow (\forall \alpha. \alpha)$, $\forall \beta. ((\forall \alpha. \alpha) \rightarrow \beta) \rightarrow ((\forall \alpha. \alpha) \rightarrow \beta)$, *etc.* none of which is more general than the other in $F_{ML}(\leq_p^{\eta-})$. In $F(\leq_p^\eta)$, the first one is better than the second-one; (In $F(\leq^\eta)$, the first-one is better than all other ones on this particular example, but this is accidental.) Hence $F_{ML}(\leq_p^{\eta-})$ does not have principal types in the usual sense. However, as we shall see below, it has a principal type for any given polymorphic *shape*.

Figure 5. Syntax-directed typing rules F_A for $F_{ML}(\leq_p^{\eta-})$

$\frac{\text{VAR-INST-RHO} \quad x : \sigma \in \Gamma \quad \sigma \leq_p^{\eta-} \rho}{\Gamma \vdash x : \rho}$	GEN	FUN
$\frac{\text{APPA-RHO} \quad \Gamma \vdash t_1 : \sigma_2[\bar{\tau}/\bar{\beta}] \rightarrow \rho_1 \quad \Gamma \vdash t_2 : \sigma_2[\bar{\tau}/\bar{\beta}]}{\Gamma \vdash t_1 (t_2 : \exists \beta. \sigma_2) : \rho_1}$		
$\frac{\text{LETA-GEN-RHO} \quad \Gamma \vdash t_1 : \sigma_1[\bar{\tau}/\bar{\beta}] \quad \Gamma, z : \forall \Gamma. \sigma_1[\bar{\tau}/\bar{\beta}] \vdash t_2 : \rho_2}{\Gamma \vdash \text{let } z = (t_1 : \exists \beta. \sigma_1) \text{ in } t_2 : \rho_2}$		

As in ML, the expression $\lambda z. z z$ does not have any mono-type τ . However, we may explicitly provide the information that we expect a type of the form $(\forall \alpha. \alpha \rightarrow \alpha) \rightarrow \tau \rightarrow \tau$ for some τ . Then, the expression is well-typed (and a best type of *that shape* may be inferred). We have $\Gamma \vdash_{F_{ML}(\leq_p^{\eta-})} \lambda z. z z : \forall \beta. (\forall \alpha. \alpha \rightarrow \alpha) \rightarrow \beta \rightarrow \beta$.

Typing problems. Let φ range over first-order substitutions, simply called substitutions for short, which map type variables to monotypes. We may distinguish four different problems:

1. *Typability:* Given a term t , do there exist a context Γ and a type σ such that $\Gamma \vdash t : \sigma$?
2. *Typing inference:* Given a term t , what are the pairs of a context Γ and a type σ such that $\Gamma \vdash t : \sigma$?
3. *Type inference:* Given a term t and a context Γ , what are the types σ and substitutions φ such that $\Gamma \varphi \vdash t : \sigma$?
4. *Type checking:* Given a term t , a context Γ , and a type σ , what are the substitutions φ such that $\Gamma \varphi \vdash t : \sigma$?

Note that our definition of typechecking still allows first-order inference, which is somehow non standard. Neither typing inference, nor type inference problems have principal solutions in general in $F_{ML}(\leq_p^{\eta-})$, nor in $F_{ML}(\leq_p^{\eta-})$. However, type checking problems do.

2.3 Syntax-directed typing rules F_A

As usual, typing judgments need not use all the flexibility allowed by the typing rules. That is, typing derivations can be rearranged to form derivations that follow certain patterns. In particular, derivations that are driven by the syntax of the conclusion, called syntax-directed, are the key to type checking and type inference. Figure 5 presents an equivalent set of rules for deriving typing judgments in $F_{ML}(\leq_p^{\eta-})$. We write $\forall \Gamma. \sigma$ for the type $\forall (\text{ftv}(\sigma) \setminus \text{ftv}(\Gamma)). \sigma$.

Rules GEN and FUN are unchanged. All syntax-directed rules but GEN assign ρ -types rather than types to terms. As opposed to ML, the GEN rule is not removed in the syntax-directed system. It is the only rule that may (and thus must) be used first to derive judgments whose conclusion mentions a type scheme that is not a ρ -type. Hence, it can be used either as the last rule in a derivation, or in three other places (and only there): immediately above the left-premise of a LET rule, the premise of a FUN rule or the premise of another GEN. Strictly speaking, Rule GEN is not syntax-directed, unless we take the expected type-scheme as part of the syntax. It would be possible and obvious to inline it in the premises of both rules FUN and LET, but this would be more verbose and less readable. Rule VAR-INST-RHO behaves as

VAR followed by INST; Rule APPA-RHO is a restriction of Rule APPA so as to conclude ρ -types only (hence the -RHO suffix). Rule LETA-GEN-RHO behaves as LET (restricted to ρ -types) with a sequence of GEN above its left premise. As in ML, this is the most interesting rule. The type scheme required for t_1 is the one requested by the annotation, but where monomorphic parts may be instantiated. Hence, its free types variables that do not appear in Γ can be generalized in the type assigned to z while typechecking t_2 . This is the only place where we introduce polymorphism that is not explicitly given (in F_A , Rule GEN may only be used for polymorphism that is explicitly given). We do so, much as in ML, and without really guessing polymorphism—we just have it for free! Shapes of all other polymorphic types in the premises are never inferred and just taken from some corresponding type in the conclusion.

If we change all type annotations into $\exists \beta. \beta$ in syntax-directed rules, we obtain exactly the syntax-directed rules of ML (GEN becomes useless if the expected type is a simple type). Thus, we depart from ML only by allowing second-order polymorphic types in annotations and typing judgments. The instantiation, let-polymorphism, and inference mechanism remain the same. In particular, type inference (as defined above) relies on first order-unification, which we shall see now.

First, let us state a series of useful standard Lemmas.

LEMMA 4 (Substitution). *If $\Gamma \vdash_{F_A} t : \sigma$, then there exists a derivation of $\Gamma \varphi \vdash_{F_A} t : \sigma \varphi$ of the same length for any substitution φ .*

We write $\Gamma' \leq_p^{\eta-} \Gamma$ if $\text{dom}(\Gamma) = \text{dom}(\Gamma')$ and $\Gamma'(z) \leq_p^{\eta-} \Gamma(z)$ for each z in $\text{dom}(\Gamma)$,

LEMMA 5 (Strengthening). *If $\Gamma \vdash_{F_A} t : \sigma$ and $\Gamma' \leq_p^{\eta-} \Gamma$, then $\Gamma' \vdash_{F_A} t : \sigma$.*

The next two lemmas establish the correspondence between the two presentations of the typing rules.

LEMMA 6 (Soundness of syntax-directed rules). *Rules VAR-INST-RHO, APPA-RHO, and LETA-GEN-RHO are derivable in $F_{ML}(\leq_p^{\eta-})$, (i.e. they may be replaced by chunk of derivations using rules of $F_{ML}(\leq_p^{\eta-})$).*

LEMMA 7 (Completeness of syntax-directed rules) *Rules INST, APPA, and LETA are admissible in F_A , (i.e. adding them to the rules to the definition of F_A does not allow to derive more judgments).*

This direction is more interesting and a bit trickier than the previous one. There are a few subtleties. Admissibility of Rule INST would not be true if $\leq_p^{\eta-}$ were replaced by $\leq_p^{\eta}$ in both systems. Admissibility of Rule APPA requires to follow an application of Rule APPA-RHO with a sequence of instances of Rule GEN. This would not be true if GEN were replaced by GEN^v in both systems—see Rule APPA^v in Figure 9.

COROLLARY 8. *The syntax-directed rules and the original rules derive the same judgments.*

2.4 Type inference via type constraints

Syntax-directed typing rules open the path to a type inference algorithm. They show that a typing judgment for an expression can only hold if some judgment for its immediate subexpressions also hold. Moreover, the type judgments for sub-expressions are entirely determined from the original

Figure 6. Syntax of constraints

$$\begin{aligned}
C ::= & \text{true} \mid \text{false} \mid \tau = \tau \mid \sigma \leq \sigma \\
& \mid C \wedge C \mid \exists \alpha. C \mid \forall \alpha. C \\
& \mid x \preceq \sigma \mid \text{def } x : \forall \bar{\alpha}[C]. \sigma \text{ in } C
\end{aligned}$$

Abbreviations:

$$\begin{aligned}
(\forall \bar{\alpha}[C]. \sigma) \preceq \rho &\triangleq \exists \bar{\alpha}. (C \wedge \sigma \leq \rho) & \bar{\alpha} \notin \text{ftv}(\rho) \\
\text{let } x : \forall \bar{\alpha}[C]. \sigma \text{ in } C' &\triangleq (\exists \bar{\alpha}. C) \wedge (\text{def } x : \forall \bar{\alpha}[C]. \sigma \text{ in } C') \\
\text{let } \Gamma, x : \forall \bar{\alpha}[C]. \sigma \text{ in } C' &\triangleq \text{let } \Gamma \text{ in let } x : \forall \bar{\alpha}[C]. \sigma \text{ in } C' \\
\text{let } \emptyset \text{ in } C &\triangleq C
\end{aligned}$$

judgment except for some instantiations of type variables by monotypes. This is just as in ML and suggests an underlying first-order type-inference mechanism.

Typing constraints are a general yet simple and intuitive framework to define type inference algorithms for ML-like languages [PR05]. Below, we present a brief summary of this approach and extend it in a straightforward manner to cover type inference for $F_{ML}(\leq)$. We refer the reader to [PR05] for a more thorough presentation.

The syntax of type constraints is given in Figure 6. We use letter C to range over type constraints. Default atomic constraints are **true**, **false**, equality constraints $\tau = \tau$, and subtyping constraints $\sigma \leq \sigma'$. We build constraints by conjunction $C \wedge C'$, existential quantification $\exists \alpha. C$, or universal quantification $\forall \alpha. C$. Finally, we use two special forms of constraints for polymorphism elimination $x \preceq \rho$ and polymorphism introduction $\text{def } x : \forall \bar{\alpha}[C]. \rho \text{ in } C'$.

In order to interpret constraints, we need a model in which we may interpret free type variables. As for ML, we take the set of closed ground types (monotypes without type variables) for our model. (Here, we must assume at least one type constructor of arity 0, *e.g.* `int`.) We use letter t to range over ground types. A ground assignment ϕ is a map from all type variables to elements of the model. We also see ϕ as a mapping from types to types, by letting $(\sigma_1 \rightarrow \sigma_2)\phi = \sigma_1\phi \rightarrow \sigma_2\phi$ and $(\forall \alpha. \sigma)\phi = \forall \alpha. (\sigma\phi_\alpha)$ where ϕ_α is the restriction of ϕ to $\text{dom}(\phi) \setminus \{\alpha\}$. (Consistently with the notation for substitutions, we write $\sigma\phi$ for the application of ϕ to σ .)

We write $\phi \Vdash C$ to mean that C is valid in environment ϕ . We take $\phi \Vdash \text{true}$ and $\phi \not\Vdash \text{false}$ for any ϕ . We take $\phi \Vdash \tau = \tau'$ if and only if $\tau\phi = \tau'\phi$ and we take $\phi \Vdash \sigma \leq \sigma'$ if and only if $\sigma\phi \leq \sigma'\phi$. Technically, we could identify the equation $\tau = \tau'$ with the inequation $\tau \leq \tau'$, since they coincide. However, for sake of exposition, we prefer to explicitly transform subtyping constraints into equations when both sides are monotypes. We take $\phi \Vdash C \wedge C'$ if and only if $\phi \Vdash C$ and $\phi \Vdash C'$. We take $\phi \Vdash \exists \alpha. C$ if and only if there exists a ground type t such that $\phi[\alpha \mapsto t] \Vdash C$ (we write $\phi[\alpha \mapsto t]$ for the environment that maps α to t and otherwise coincides with ϕ). We take $\phi \Vdash \forall \alpha. C$ if and only if for every ground type t , we have $\phi[\alpha \mapsto t] \Vdash C$. A constraint C entails a constraint C' if for every ground assignment ϕ , such that $\phi \Vdash C$ we also have $\phi \Vdash C'$. We then write $C \Vdash C'$. Constraints C and C' are equivalent if both $C \Vdash C'$ and $C' \Vdash C$ holds.

Def-constraints $\text{def } x : \forall \bar{\alpha}[C]. \sigma \text{ in } C'$ can be interpreted as $C'[\forall \bar{\alpha}[C]. \sigma/x]$ (of course, the resolution of constraints will avoid this substitution). The effect of this substitution is to expose constraints of the form $\forall \bar{\alpha}[C]. \sigma \preceq \rho$. These are

Figure 7. Solving $\leq$ -constraints

$$\begin{aligned}
& \tau \leq \tau' \longrightarrow \tau = \tau' \\
& \sigma_1 \rightarrow \sigma_2 \leq \alpha \longrightarrow \exists \alpha_1 \alpha_2. (\sigma_1 \rightarrow \sigma_2 \leq \alpha_1 \rightarrow \alpha_2 \\
& \quad \wedge \alpha = \alpha_1 \rightarrow \alpha_2) \\
& \alpha \leq \sigma_1 \rightarrow \sigma_2 \longrightarrow \exists \alpha_1 \alpha_2. (\alpha_1 \rightarrow \alpha_2 \leq \sigma_1 \rightarrow \sigma_2 \\
& \quad \wedge \alpha = \alpha_1 \rightarrow \alpha_2) \\
& \sigma_1 \rightarrow \sigma_2 \leq \sigma'_1 \rightarrow \sigma'_2 \longrightarrow \sigma'_1 \leq \sigma_1 \wedge \sigma_2 \leq \sigma'_2 \\
& \quad \forall \alpha. \sigma \leq \rho \longrightarrow \exists \alpha. (\sigma \leq \rho) & \alpha \notin \text{ftv}(\rho) \\
& \sigma \leq \forall \alpha. \sigma' \longrightarrow \forall \alpha. (\sigma \leq \sigma') & \alpha \notin \text{ftv}(\sigma)
\end{aligned}$$

Figure 8. Constraint generation rules

$$\begin{aligned}
\llbracket x : \rho \rrbracket &= x \preceq \rho \\
\llbracket \lambda z. t : \alpha \rrbracket &= \exists \beta_2 \beta_1. (\llbracket \lambda z. t : \beta_2 \rightarrow \beta_1 \rrbracket \\
& \quad \wedge \alpha = \beta_2 \rightarrow \beta_1) & \beta_1, \beta_2 \neq \alpha \\
\llbracket \lambda z. t : \sigma_2 \rightarrow \sigma_1 \rrbracket &= \text{let } z : \sigma_2 \text{ in } \llbracket t : \sigma_1 \rrbracket \\
\llbracket t_1 (t_2 : \exists \beta. \sigma_2) : \rho_1 \rrbracket &= \exists \beta. (\llbracket t_1 : \sigma_2 \rightarrow \rho_1 \rrbracket \wedge \llbracket t_2 : \sigma_2 \rrbracket) \\
& & \beta \notin \text{ftv}(\rho_1) \\
\llbracket \text{let } z = (t_1 : \exists \beta. \sigma_1) \\
& \quad \text{in } t_2 : \rho_2 \rrbracket &= \text{let } z : \forall \beta. (\llbracket t_1 : \sigma_1 \rrbracket). \sigma_1 \text{ in } \llbracket t_2 : \rho_2 \rrbracket \\
\llbracket t : \forall \alpha. \sigma \rrbracket &= \forall \alpha. \llbracket t : \sigma \rrbracket
\end{aligned}$$

syntactic sugar for $\exists \bar{\alpha}. (C \wedge \sigma \leq \rho)$, provided $\bar{\alpha} \notin \text{ftv}(\rho)$. Hence, the scope in a constrained type scheme $\forall \bar{\alpha}[C]. \sigma$ is both C and σ . In fact, rather than using $\text{def } x : \forall \bar{\alpha}[C]. \rho \text{ in } C'$ directly, we often use $\text{let } x : \forall \bar{\alpha}[C]. \rho \text{ in } C'$ as an abbreviation for $\exists \bar{\alpha}. C \wedge \text{def } x : \forall \bar{\alpha}[C]. \rho \text{ in } C'$.

The solver of [PR05] need only be extended to solve subtyping constraints. These can easily be reduced to equality constraints by repeatedly applying the rules of Figure 7. (We have left it implicit in the first and second rules that variables α_1 and α_2 should not appear free on the left-hand side). Each rule preserves the meaning of constraints, indeed. When no rule applies, we are left with equality constraints of the form $\tau = \tau'$, which are treated as first-order unification constraints (the predicate $\leq$ may be interpreted as equality on monotypes). Hence, rules of Figure 7 are meant to be added to rules for solving first-order constraints as well as structural rules for rearranging constraints (see [PR05]). Indeed, the rules of Figure 7 can already be found in [OL96].

LEMMA 9. *The rewriting rules of Figure 7 preserve constraint equivalence.*

Once the (exposed) subtyping constraints have been resolved, the remaining constraints are equality constraints between monotypes, which can then be resolved by a unification algorithm, as in [PR05]. The whole simplification process ends with a constraint in solved form, which determines a most general solution to the initial problem. Remark that a type substitution φ may always be represented as the type constraint $\bigwedge_{\alpha \in \text{dom}(\varphi)} (\alpha = \alpha\varphi)$.

Constraint generation rules

We now describe a type inference algorithm by turning type checking problems into type constraint problems, which can then be simplified as described above.

Constraint generation rules are given in Figure 8. They take as input a type expression t and an expected type σ and return a constraint that describes the conditions under which t may be assigned type σ . Maybe surprisingly,

Figure 9. Syntax-directed typing rules for $F_{ML}^v(\leq_p^{\eta-})$

VAR-INST-RHO	GEN ^v	FUN
$\frac{\text{APP}_A^v \quad \Gamma \vdash t_1 : \sigma_2[\bar{\tau}/\bar{\beta}] \rightarrow \forall \bar{\alpha}. \rho_1 \quad \Gamma \vdash t_2 : \sigma_2[\bar{\tau}/\bar{\beta}] \quad \bar{\alpha} \notin \text{ftv}^{t_1, t_2}(\rho_1)}{\Gamma \vdash t_1 (t_2 : \exists \bar{\beta}. \sigma_2) : \forall \bar{\alpha}. \rho_1}$		
$\frac{\text{LET}_A\text{-GEN}^v \quad \Gamma \vdash t_1 : \sigma_1[\bar{\tau}/\bar{\beta}] \quad \Gamma, z : \forall^{t_1} \Gamma. \sigma_1[\bar{\tau}/\bar{\beta}] \vdash t_2 : \forall \bar{\alpha}. \rho_2 \quad \bar{\alpha} \notin \text{ftv}^{t_1, t_2}(\rho_2)}{\Gamma \vdash \text{let } z = (t_1 : \exists \bar{\beta}. \sigma_1) \text{ in } t_2 : \forall \bar{\alpha}. \rho_2}$		

constraint generation does not take a typing context Γ : if the program t has free program identifiers x , so will the constraint $\llbracket t : \sigma \rrbracket$. Then, the constraint may only be valid when placed in a context of the form $\text{let } \Gamma \text{ in } \cdot$ that will assign types to free type variables of t .

Most constraint generation rules can be read straightforwardly and follow syntax-directed typing rules. An identifier x has type ρ if and only if ρ is an instance of the type of x . A function $\lambda z. t$ has type α if and only if it has type $\beta_2 \rightarrow \beta_1$ where α is of the form $\beta_2 \rightarrow \beta_1$ for some types β_1, β_2 ; a function $\lambda z. t$ has a type scheme $\sigma_2 \rightarrow \sigma_1$ if and only if t has type σ_1 assuming z has type scheme σ_2 . The decomposition of applications is also straightforward. To decompose a let-binding, we first generate a constraint for typechecking its argument; this constraint is used to build a constrained type-scheme by generalizing all type variables $\bar{\beta}$ that are free in σ_1 , but free in Γ . After simplification of $\llbracket t_1 : \sigma_1 \rrbracket$ some variables of $\bar{\beta}$ will in general be further constrained, maybe so that they can only be monotypes. Generalizing them is not a problem, since the resolution algorithm will take care of such dependences, which is one of the elegance of using type constraints. Finally, an expression t has a type $\forall \bar{\alpha}. \sigma$ if and only if t has type σ for any instance of $\bar{\alpha}$.

A type checking problem $\Gamma \vdash t : \sigma$ is equivalent to the constraint $\text{let } \Gamma \text{ in } \llbracket t : \sigma \rrbracket$ (which is syntactic sugar for a sequence of let constraints, as described in Figure 6). This is stated very precisely and concisely by the following two lemmas. This uses a correspondence mapping an idempotent substitution φ to a solved type constraint $\bigwedge_{\alpha \in \text{dom}(\varphi)} \alpha = \alpha\varphi$.

LEMMA 10 (soundness). *If $\varphi \Vdash \text{let } \Gamma \text{ in } \llbracket t : \sigma \rrbracket$, then $\Gamma\varphi \vdash t : \sigma\varphi$.*

LEMMA 11 (completeness). *If $\Gamma\varphi \vdash t : \sigma\varphi$, then $\varphi \Vdash \text{let } \Gamma \text{ in } \llbracket t : \sigma \rrbracket$.*

We can read back these two theorems into a more traditional (but not really simpler) formulation.

COROLLARY 12. *Given a typing environment Γ , a program t , and a type scheme σ , a substitution φ is a solution to the type checking problem $\Gamma \vdash t : \sigma$ if and only if it is a solution to the constraint $\text{let } \Gamma \text{ in } \llbracket t : \sigma \rrbracket$.*

2.5 Type inference with side-effects

As we have seen in Section 1.5, choosing a call-by-value semantics and extending the language with side-effects also implies some changes in the typing rules. Let $F_{ML}^v(\leq_p^{\eta-})$ be the language whose expressions and typing rules are the same as those of $F_{ML}(\leq)$, except for Rule GEN, which is

Figure 10. New constraint generation rules for $F_{ML}^v(\leq_p^{\eta-})$

$\begin{aligned} \llbracket t_1 (t_2 : \exists \bar{\beta}. \sigma_2) : \forall \bar{\alpha}. \rho_1 \rrbracket = \\ \exists \bar{\beta}. (\llbracket t_1 : \sigma_2 \rightarrow \forall \bar{\alpha}. \rho_1 \rrbracket \wedge \llbracket t_2 : \sigma_2 \rrbracket) \quad \bar{\alpha} = \emptyset \text{ if } t_1, t_2 \in \mathcal{U} \\ \llbracket \text{let } z = (u_1 : \exists \bar{\beta}. \sigma_1) \text{ in } t_2 : \forall \bar{\alpha}. \rho_2 \rrbracket = \\ \text{let } z : \forall \bar{\beta}. (\llbracket u_1 : \sigma_1 \rrbracket). \sigma_1 \text{ in } \llbracket t_2 : \forall \bar{\alpha}. \rho_2 \rrbracket \quad \bar{\alpha} = \emptyset \text{ if } t_2 \in \mathcal{U} \\ \llbracket \text{let } z = (t_1 : \exists \bar{\beta}. \sigma_1) \text{ in } t_2 : \forall \bar{\alpha}. \rho_2 \rrbracket = \\ \exists \bar{\beta}. (\llbracket t_1 : \sigma_1 \rrbracket \wedge \text{let } z : \sigma_1 \text{ in } \llbracket t_2 : \forall \bar{\alpha}. \rho_2 \rrbracket) \quad t_1 \notin \mathcal{U} \\ \llbracket u : \forall \bar{\alpha}. \sigma \rrbracket = \forall \bar{\alpha}. \llbracket u : \sigma \rrbracket \end{aligned}$

replaced by Rule GEN^v. We may extend the definition of non-expansive expressions given in Section 1.5 with $\text{let } z = u_1 \text{ in } u_2$. The relation $\leq_p^{\eta-}$ is a subrelation of $\leq^{\eta-}$. Hence, $F_{ML}^v(\leq_p^{\eta-})$ is sound.

Preparing for type inference, we must also review the syntax-directed typing rules. Of course, we should change GEN to GEN^v. Since rule GEN^v is more restrictive, only non-negative type variables may be generalized a posteriori in the type of an expansive expression. As a consequence, the result of the application (see Rule APP_A) may have a polymorphic type but only if this polymorphism comes from the codomain σ_1 of the type of t_1 or from non-negative free type variables in σ_1 . Hence, we change Rule APP_A-RHO for Rule APP_A^v to allow exactly those variables $\bar{\alpha}$ that do not belong to $\text{ftv}^{t_1, t_2}(\rho_1)$ in the type of the application—others may still be generalized afterward. The notation $\text{ftv}^{\bar{t}}(\sigma)$ where $\bar{t}$ is a sequence of expressions stands for $\text{ftv}(\sigma)$ if $\bar{t} \in \mathcal{U}$ (i.e. all expressions of $\bar{t}$ are non-expansive) and $\text{ftv}^{\dagger}(\sigma)$ otherwise. Similarly, we change Rule LET_A-GEN-RHO to LET_A-GEN^v which allows $\bar{\alpha}$ to be polymorphic in the type of t_2 . The sequence of expressions t_1, t_2 used in the superscript of ftv controls the expansiveness of $\text{let } z = (t_1 : \exists \bar{\beta}. \sigma_1) \text{ in } t_2$. Consistently with the restriction of generalization, we have also replaced $\forall \Gamma. \sigma_1[\bar{\tau}/\bar{\beta}]$ by $\forall^{t_1} \Gamma. \sigma_1[\bar{\tau}/\bar{\beta}]$. The notation $\forall^{t_1} \Gamma. \sigma$ stands for $\forall \text{ftv}^{\bar{t}}(\sigma) \setminus \text{ftv}(\Gamma). \sigma$. That is, generalizable variables in the type of t_1 are $\text{ftv}(\sigma) \setminus \text{ftv}(\Gamma)$ as before if t_1 is non-expansive, and $\text{ftv}^{\dagger}(\sigma) \setminus \text{ftv}(\Gamma)$ otherwise.

Constraint generation

Let us first consider the simpler case of the standard value restriction. That is, we ignore non-negative type variables which amounts to taking the empty set for $\text{ftv}^{\bar{t}}(\sigma)$. In this case, $\text{ftv}^{\bar{t}}(\sigma)$ is $\text{ftv}(\sigma)$ if all expressions of $\bar{t}$ are non-expansive and is empty otherwise.

It is then straightforward to extend constraint generation: we need not add new forms of constraints but only test for expansiveness during constraint generation and generate different constraints for non-expansive applications and let-bindings. The modified rules are summarized in Figure 10 (capture-avoiding side-conditions have been omitted); these should replace the corresponding rules in Figure 8. Constraint resolution then proceeds as before. Lemmas 6, 7, 10, and 11 extend to $F_{ML}^v(\leq_p^{\eta-})$ with *standard* value restriction.

By lack of space, we refer to the full version [Rém05] for a discussion of the *enhanced* value restriction, which requires enriching the language of type-constraints.

2.6 Type inference for $F_{ML}(\leq_p^{\eta})$

Unsurprisingly, rules of Figure 5 do not form a syntax-directed presentation of $F_{ML}(\leq_p^{\eta})$: after replacing $\leq_p^{\eta-}$ by $\leq_p^{\eta}$ in Rule VAR-INST-RHO they would defined a typing relation that does not satisfy Rule INST. Fortunately, we may

easily reduce type inference for $F_{ML}(\leq_p^\eta)$ to type inference for $F_{ML}(\leq_p^{\eta-})$ by first putting typing problems into prenex-forms. More precisely, let $\text{prf}(\exists\beta.\sigma)$ be $\exists\beta.\text{prf}(\sigma)$ and $\text{prf}(t)$ be a copy of t where all annotations θ have been replaced by their prenex-forms $\text{prf}(\theta)$.

LEMMA 13. *Judgments $\Gamma \vdash_{F_{ML}(\leq_p^\eta)} t : \sigma$ and $\text{prf}(\Gamma) \vdash_{F_{ML}(\leq_p^{\eta-})} \text{prf}(t) : \text{prf}(\sigma)$ are equivalent.*

This shows that the choice between $\leq_p^{\eta-}$ and $\leq_p^\eta$ is unimportant in predicative systems.

2.7 Recovering impredicativity

An important limitation of $F_{ML}(\leq_p^{\eta-})$ is its predicativity. Fortunately, there is an easy way to recover impredicativity by adding an explicit impredicative decidable fragment of type-containment, say $\leq_I$ to the implicit predicative type-containment relation $\leq_p^{\eta-}$ or $\leq_p^\eta$, say $\leq_P$. Indeed, to obtain full impredicativity, *i.e.*, in the end, the expressiveness of F , it suffices that $\leq_I$ be a larger relation than $\leq^F$. We choose the relation $\leq_\approx^F$, which is better suited for type inference, as it disregards the order of quantifiers.

Formally, we may simply introduce a collection of coercion functions (*i.e.* functions that evaluate to the identity) $(_ : \exists\beta.\sigma_1 \triangleright \sigma_2)$ with types $\forall\beta.\sigma_1 \rightarrow \sigma_2$ for all pairs (σ_1, σ_2) in $\leq_I$ and with β equal to $\text{ftv}(\sigma_1) \cup \text{ftv}(\sigma_2)$. As for annotations, we require coercions to be closed, although this is only a matter of simplification. This is the purpose of the existentially bound variables whose scope extends to both sides of the coercion. We refer to this language as $F_{ML}(\leq_I, \leq_P)$.

We may here take advantage of constrained type schemes and be slightly more flexible. Let us introduce new forms of constraints $\sigma_1 \leq \sigma_2$ whose meaning is given by $\phi \Vdash \sigma_1 \leq \sigma_2$ if and only if $\sigma_1 \phi \leq \sigma_2 \phi$ (for some given relations $\leq$). Then, we may assign $(_ : \exists\beta.\sigma_1 \triangleright \sigma_2)$ the constrained type scheme $\forall\beta[\sigma_1 \leq_I \sigma_2].\sigma_1 \rightarrow \sigma_2$. This has two advantages: first, it internalizes the verification of the $\leq_I$ type-containment; second, it allows to provide types σ_1 and σ_2 up to some monomorphic instantiation. For example, taking $\leq^F$ for $\leq_I$, the coercion $(_ : \exists\beta.\forall\alpha.\alpha \rightarrow \beta \rightarrow \alpha \triangleright \sigma \rightarrow \tau \rightarrow \sigma)$ is now valid for any (closed) type τ and type scheme σ ; finding out that β must actually be τ can be left to type inference.

This extension is safe by construction, as long as $\leq_I$ is a subrelation of $\leq^\eta$, since it amounts to giving the identity a valid type in $F(\leq^\eta)$. Regarding type inference, it only remains to solve $\leq_I$ type-containment constraints so that the coercion is well-defined and Rule VAR-INST-RHO applies.

Let us now focus on the particular case of $F_{ML}(\leq_\approx^F, \leq_p^{\eta-})$, which we shall abbreviate as F_{ML} . Type constraints $\sigma_1 \leq_\approx^F \sigma_2$ can be reduced to $\text{canon}(\sigma_1) \leq^F \text{canon}(\sigma_2)$. So we are left to solving constraints of the form $\sigma_1 \leq^F \sigma_2$. Recall that all instances of the relation $\leq^F$ are of the form $\forall\bar{\alpha}.\rho \leq^F \forall\bar{\gamma}.\rho[\bar{\sigma}/\bar{\alpha}]$ with $\bar{\gamma} \notin \text{ftv}(\forall\bar{\alpha}.\rho)$. Thus, the constraint $\forall\bar{\alpha}.\rho_1 \leq^F \forall\bar{\gamma}.\rho_2$ with free type variables $\bar{\beta}$ is equivalent to the existence of type schemes $\bar{\sigma}$ such that ρ_2 and $\rho_1[\bar{\sigma}/\bar{\alpha}]$ are equal modulo α -conversion with variables $\bar{\gamma}$ not in $\text{ftv}(\forall\bar{\alpha}.\rho_1)$. That is, it is equivalent to the equality $\forall\bar{\gamma}.\exists\bar{\alpha}.\rho_1 = \rho_2$ where $\bar{\alpha}$ range over type schemes and $\bar{\alpha}$ and $\bar{\beta}$ range over monotypes. This is a (restricted) form of unification modulo α -conversion and under a mixed prefix (but no β -reduction is ever involved). Solving such constraints is folklore knowledge. We refer the reader to the full version for details.

Coercion functions are meant to be applied to terms. Formally, an application requires a type annotation, which in

this case may be exactly the domain of the coercion. We may avoid repeating the annotation by letting the syntactic sugar $(t : \exists\beta.\sigma_1 \triangleright \sigma_2)$ stand for $(_ : \exists\beta.\sigma_1 \triangleright \sigma_2) (t : \exists\beta.\sigma_1)$. When a coercion is in application position, it still needs an annotation, even though the coercion may already carry all the necessary type information. Hence, we may see $t_1 (t_2 : \exists\beta.\sigma_1 \triangleright \sigma_2)$ as syntactic sugar for $t_1 ((t_2 : \exists\beta.\sigma \triangleright \sigma') : \exists\beta.\sigma')$.

There still seem to be some redundancies in coercions since the polymorphic shape of the expected type σ'_2 must be provided when typechecking a coercion $(t : \exists\beta.\sigma_1 \triangleright \sigma_2)$. However, σ'_2 is only known up to $\leq_p^{\eta-}$. That is, in general, we only have $\sigma_2 \leq \sigma'_2$. If we were to take σ'_2 for σ_2 , we would then be left with $\leq_\approx^F \circ \leq_p^{\eta-}$ constraints. However, we have not explored yet the resolution of such constraints.

Expressiveness The set of raw terms typable in $F_{ML}(\leq_I, \leq_P)$ is exactly $F(\leq_I \circ \leq_P)$. Hence, F_{ML} contains terms of both F and $F(\leq_p^{\eta-})$ —but not all terms of $F(\leq^\eta)$. One may in fact generalize the idea of explicit coercions by providing typing evidence for any coercion, which can be done by writing coercions in F_{ML} . We could thus also allow coercions of the form $(_ : t)$ where t is a term of F_{ML} that η -reduces to the identity. This way, we would finally reach all of $F(\leq^\eta)$. However, writing coercion functions explicitly is rather heavy.

An alternative way to recover impredicativity is to use semi-explicit polymorphism [GR97], which simplifies significantly here, as types of the host language are already of arbitrary rank.

3. Propagation of annotations in $F_{ML}^?$

We consider the language F_{ML} . However, we first treat coercions as constants, as if we were in $F_{ML}(\leq_p^{\eta-})$. We shall discuss special elaboration of coercions in Section 3.5 and elaboration in $F_{ML}^v(\leq_p^{\eta-})$ in Section 3.4.

Many type annotations may seem redundant in F_{ML} . For instance, consider the expression $\text{let } f = (\lambda z.z : \sigma_{id} \rightarrow \sigma_{id}) \text{ in } f(\lambda y.y : \sigma_{id})$, which is well typed, but becomes ill-typed if we remove the annotation on the application. However, one may argue that given the type of f , it is *obvious* what the annotation on the application should be.

We first show that only the polymorphic skeletons of annotations, where the monomorphic leaves are stripped off, actually matter. We refer to them as *shapes*. Computation on shapes is simpler than computation on types, because shapes are closed. We exploit this to propose a preprocessing step that propagates shapes before typechecking in F_{ML} .

3.1 Shapes

Let *shapes* be closed polymorphic types extended with a unary type constructor $\sharp$. Shapes are considered modulo the *absorbing equation* $\sharp \rightarrow \sharp = \sharp$. A shape is in canonical form when it contains the minimum number of occurrences of $\sharp$. The shape of a polymorphic type σ , written $[\sigma]$ is obtained from σ by replacing all free type variables of σ by $\sharp$. We use $\mathcal{S}$ to range over shapes. Shapes capture the polymorphic structure of types. For example, consider the type σ_0 equal to $\forall\alpha_1.(\forall\alpha_2.(\alpha_1 \rightarrow \alpha_2) \rightarrow (\beta_0 \rightarrow \beta_0)) \rightarrow (\beta_1 \rightarrow \beta_2)$. Its shape $[\sigma]$ is $\forall\alpha_1.(\forall\alpha_2.(\alpha_1 \rightarrow \alpha_2) \rightarrow (\sharp \rightarrow \sharp)) \rightarrow (\sharp \rightarrow \sharp)$, which can be put in its canonical form $\forall\alpha_1.(\forall\alpha_2.(\alpha_1 \rightarrow \alpha_2) \rightarrow \sharp) \rightarrow \sharp$. By extension, the shape of an annotation $[\exists\beta.\sigma]$ is simply $[\sigma]$. A shape $\mathcal{S}$ may be read back as a type annotation, written $[\mathcal{S}]$. By definition $[\mathcal{S}]$ is $\exists\bar{\alpha}.\sigma$ if $\mathcal{S}$ is in canonical form and syntactically equal to $\sigma[\sharp/\bar{\alpha}]$, and each variable of $\bar{\alpha}$ occurs

exactly once in σ . For example, the annotation $\llbracket \sigma \rrbracket$ is $\exists \beta_1, \beta_2 \forall \alpha_1. (\forall \alpha_2. (\alpha_1 \rightarrow \alpha_2) \rightarrow \beta_1) \rightarrow \beta_2$. We sometimes need to strip off the front quantifiers of a shape $\mathcal{S}$. However, the resulting type may contain free type variables and must be reshaped. We write $\mathcal{S}^\flat$ for $\llbracket \rho \rrbracket$ where $\llbracket \mathcal{S} \rrbracket$ is of the form $\exists \beta. \forall \alpha. \rho$.

Our interest for shapes is that only shapes of annotations matter for type inference. If we write $\llbracket t \rrbracket$ for the term t where each annotation θ has been replaced by $\llbracket \theta \rrbracket$, this property is precisely captured by the following lemma.

LEMMA 14. *If $\Gamma \vdash t : \sigma$ then $\Gamma \vdash \llbracket t \rrbracket : \sigma$.*

Since shapes are a projection of types, one could project typing rules as well, and obtain a “shaping” relation that would assign shapes to programs. By construction, one would then expect a property such as if $\Gamma \vdash t : \sigma$ then $\llbracket \Gamma \rrbracket \vdash t : \llbracket \sigma \rrbracket$. However, such a result would not be very useful, since we already have an algorithm for typechecking.

Conversely, one could hope for some kind of shape-inference algorithm and by combination with typechecking at a given shape to solve type inference problems at unknown shapes. However, shape-inference is of course not quite realistic if we define it as finding *every* shape for which a type could be inferred. However, we may give up completeness, or even soundness, and look for some *obvious* shape for which a type might be inferred. Incompleteness may not be a problem in the context of type reconstruction where we attempt to rebuild missing type annotations in an *obvious* manner and accept to fail if there is no way to do so. *Since type reconstruction is not going to be complete with respect to some simple logical specification, let it be at least simple and intuitive!*

Thus, we shall now allow unannotated application nodes $t_1 t_2$ and let-binding nodes $\text{let } z = t_1 \text{ in } t_2$ —here, we mean no annotation at all and not the trivial annotation $\exists \beta. \beta$, which we may still write, but *explicitly*. Since the expected shape may now be missing, it becomes convenient to simultaneously allow explicit annotations on formal parameters of functions, which we write $\lambda z : \theta. t$. Precisely, let expressions of the language $\mathcal{F}_{\text{ML}}'$ be those of $\mathcal{F}_{\text{ML}}$ extended with the three constructions above. Finally, we may now define typechecking in $\mathcal{F}_{\text{ML}}'(\leq)$ by elaboration into terms of $\mathcal{F}_{\text{ML}}$ computing on shapes to fill in the missing annotations. This is a form of (incomplete) shape inference that returns both a shape and an elaborated term of $\mathcal{F}_{\text{ML}}$. We write shape inference using judgments of the form $\Gamma \vdash_\uparrow t : \mathcal{S} \Rightarrow t'$ where the $\uparrow$ is there to remember that $\mathcal{S}$ is inferred. Because $\mathcal{F}_{\text{ML}}'$ is a superset of $\mathcal{F}_{\text{ML}}$, there are still cases where the expected type, hence the expected shape, are known. In such cases, elaboration may be turned into a checking mode: it then uses the given shape and elaborates subterms from which it can return an elaborated term. Hence, we recursively define a judgment $\Gamma \vdash_\downarrow t : \mathcal{S} \Rightarrow t'$. Here, Γ , t and $\mathcal{S}$ are all given and t' is returned. The “ $\downarrow$ ” sign indicates that $\mathcal{S}$ is only checked. The idea of mixing checking and inference modes is taken from Peyton Jones and Shields [PVWS05a] but goes back to older works such as local type inference [PT00]. The direction of arrows is taken from [PVWS05a].

3.2 Elaboration

The two elaboration judgments are defined by the set of rules of Figure 11. We used variable ϵ to range over $\uparrow$ and $\downarrow$. This allows factoring some rules that could otherwise be written as pairs of rules.

Figure 11. Elaboration rules for $\mathcal{F}_{\text{ML}}'$

$\frac{\text{VAR-C} \quad x : \mathcal{S}' \in \Gamma}{\Gamma \vdash_{\downarrow} x : \mathcal{S} \Rightarrow x}$	$\frac{\text{VAR-I} \quad x : \mathcal{S} \in \Gamma}{\Gamma \vdash_{\uparrow} x : \mathcal{S} \Rightarrow x}$
$\frac{\text{FUN-C} \quad \Gamma, z : \mathcal{S}_2 \vdash_{\downarrow} t : \mathcal{S}_1 \Rightarrow t'}{\Gamma \vdash_{\downarrow} \lambda z. t : \mathcal{S}_2 \rightarrow \mathcal{S}_1 \Rightarrow \lambda z. t'}$	$\frac{\text{FUN-I} \quad \Gamma, z : \# \vdash_{\uparrow} t : \mathcal{S} \Rightarrow t'}{\Gamma \vdash_{\uparrow} \lambda z. t : \# \rightarrow \mathcal{S} \Rightarrow \lambda z. t'}$
$\frac{\text{FUN}_A\text{-C} \quad \Gamma, z : [\sigma] \vdash_{\downarrow} t : \mathcal{S}_1 \Rightarrow t'}{\Gamma \vdash_{\downarrow} \lambda z : \exists \bar{\beta}. \sigma. t : \mathcal{S}_2 \rightarrow \mathcal{S}_1 \Rightarrow \lambda z. \text{let } z = (z : \exists \bar{\beta}. \sigma) \text{ in } t'}$	$\frac{\text{FUN}_A\text{-I} \quad \Gamma, z : [\sigma] \vdash_{\uparrow} t : \mathcal{S} \Rightarrow t'}{\Gamma \vdash_{\uparrow} \lambda z : \exists \bar{\beta}. \sigma. t : [\sigma] \rightarrow \mathcal{S} \Rightarrow \lambda z. \text{let } z = (z : \exists \bar{\beta}. \sigma) \text{ in } t'}$
$\frac{\text{LET-}\epsilon \quad \Gamma \vdash_{\uparrow} t_1 : \mathcal{S}_1 \Rightarrow t'_1 \quad \Gamma, z : \mathcal{S}_1 \vdash_{\epsilon} t_2 : \mathcal{S}_2 \Rightarrow t'_2}{\Gamma \vdash_{\epsilon} \text{let } z = t_1 \text{ in } t_2 : \mathcal{S}_2 \Rightarrow \text{let } z = (t'_1 : \lfloor \mathcal{S}_1 \rfloor) \text{ in } t'_2}$	
$\frac{\text{LET}_A\text{-}\epsilon \quad \Gamma \vdash_{\downarrow} t_1 : [\sigma] \Rightarrow t'_1 \quad \Gamma, z : [\sigma] \vdash_{\epsilon} t_2 : \mathcal{S}_2 \Rightarrow t'_2}{\Gamma \vdash_{\epsilon} \text{let } z = (t_1 : \exists \bar{\beta}. \sigma) \text{ in } t_2 : \mathcal{S}_2 \Rightarrow \text{let } z = (t'_1 : \exists \bar{\beta}. \sigma) \text{ in } t'_2}$	
$\frac{\text{APPA-C} \quad \Gamma \vdash_{\downarrow} t_1 : [\sigma] \rightarrow \mathcal{S} \Rightarrow t'_1 \quad \Gamma \vdash_{\downarrow} t_2 : [\sigma] \Rightarrow t'_2}{\Gamma \vdash_{\downarrow} t_1 (t_2 : \exists \bar{\beta}. \sigma) : \mathcal{S} \Rightarrow t'_1 (t'_2 : \exists \bar{\beta}. \sigma)}$	
$\frac{\text{APPA-I} \quad \Gamma \vdash_{\uparrow} t_1 : \mathcal{S} \Rightarrow t'_1 \quad \mathcal{S}^b = \mathcal{S}_2 \rightarrow \mathcal{S}_1 \quad \Gamma \vdash_{\downarrow} t_2 : [\sigma] \Rightarrow t'_2}{\Gamma \vdash_{\uparrow} t_1 (t_2 : \exists \bar{\beta}. \sigma) : \mathcal{S}_1 \Rightarrow t'_1 (t'_2 : \exists \bar{\beta}. \sigma)}$	
$\frac{\text{APP-C} \quad \Gamma \vdash_{\uparrow} t_1 : \mathcal{S} \Rightarrow t'_1 \quad \mathcal{S}^b = \mathcal{S}_2 \rightarrow \mathcal{S}'_1 \quad \Gamma \vdash_{\uparrow} t_2 : \mathcal{S}'_2 \Rightarrow t'_2}{\Gamma \vdash_{\downarrow} t_1 t_2 : \mathcal{S}_1 \Rightarrow t'_1 (t'_2 : \lfloor \mathcal{S}_2 \rfloor)}$	
$\frac{\text{APP-I} \quad \Gamma \vdash_{\uparrow} t_1 : \mathcal{S} \Rightarrow t'_1 \quad \mathcal{S}^b = \mathcal{S}_2 \rightarrow \mathcal{S}_1 \quad \Gamma \vdash_{\uparrow} t_2 : \mathcal{S}'_2 \Rightarrow t'_2}{\Gamma \vdash_{\uparrow} t_1 t_2 : \mathcal{S}_1 \Rightarrow t'_1 (t'_2 : \lfloor \mathcal{S}_2 \rfloor)}$	

Rule VAR-C checks that there is a binding for x and ignores the shape of x . Of course, a certain relation should hold between $\mathcal{S}'$ and $\mathcal{S}$. However, this will be verified by later typechecking in $\mathcal{F}_{\text{ML}}$, so we may just ignore this condition. Rule VAR-I reads the shape from the environment Γ . Note that the best known shape of x is $\mathcal{S}$ and not an instance of $\mathcal{S}$, which would be weaker.

In Rule FUN-C the given shape of the conclusion provides us with both the shape of the parameter z and the expected shape of t . We can thus elaborate the premise in checking mode. Rule FUN-I is just the opposite, since nothing is known from the conclusion. We thus use shape $\#$ for the parameter and elaborate the premise in inference mode. In Rule FUN_A-C we actually have extra information since the shape of the parameter is known from both the annotation in the conclusion and the shape of the conclusion. Again, some relation between $[\sigma]$ and $\mathcal{S}_2$ should hold. We use $[\sigma]$, which is explicitly given and simply ignore $\mathcal{S}_2$ during elaboration. Still, the let-binding in the elaborated term, which stands for an additional pure type-annotation, ensures that the correct relation between σ and $\mathcal{S}_2$ will be verified. Rule FUN_A-I, is similar except that the premise is called in inference instead of checking mode.

Cases for let-bindings (LET- ϵ and LET_A- ϵ) are all very similar. The left-premise is called in inference mode when there is no annotation on t_1 ; the right-premise is called in inference mode when the conclusion is itself in inference

mode. The important detail is that in all cases, variable z is bound to the best-known-shape $\lceil \sigma \rceil$ or S_1 , whether it is taken from the annotation or inferred. In particular, one cannot “generalize the shape” in any meaningful way. Indeed, the type inferred for t_1 during typechecking may have free variables that could later be generalized during typechecking. However, we cannot tell during elaboration, so we may only assume for the shape of z the best known shape of t_1 . Of course, typechecking may later do better!

An annotated application is elaborated in checking mode (Rule APPA-C) by calling both premises in checking mode, since all necessary shapes are known. In inference mode (APPA-I), the left-premise must be called in inference mode because the expected shape of t_1 is not entirely known. However, only the range of the inferred shape S_1 matters and is used in the elaboration (remember that S^b is S stripped off its toplevel quantifiers and reshaped). Again the domain S_2 should be related to σ in some way, but we may ignore it here. Rules APP-I and APP-C deal with the inference mode. Both premises must be called in inference mode, since none of the expected shapes is ever entirely known. The annotation S_2 is however taken from the left-hand side. This choice is somehow arbitrary, but it seems to usually work better in practice. The difference between inference and checking modes is that the return shape S_1 is given in checking mode, while it is the range of S^b in inference mode. (In checking mode the range S'_1 of S^b , which is ignored, should be related to S_1 in some way.)

By construction, elaboration always keep existing annotations, so it is the identity on terms of F_{ML} and idempotent for terms of $F_{ML}^?$.

Variations In typing rules for unannotated applications, we have somehow made arbitrary choices, taking the annotation from the domain of the inferred shape of the function and discarding the inferred shape of the argument. In [Rém05], we propose other elaboration rules for applications that may also sometimes take the shape of the argument for building the annotation.

3.3 Typechecking in $F_{ML}^?$

Of course, elaboration may return ill-typed programs, since shapes are only an approximation of types. Hence the programs resulting from elaboration must be submitted to type inference in F_{ML} . Actually, well-typedness in $F_{ML}^?$ is simply defined by means of elaboration in F_{ML} .

DEFINITION 1. Let $\Gamma \vdash_e t : \sigma$ if and only if there exists an expression t' such that $\lceil \Gamma \rceil \vdash_e t : \lceil \sigma \rceil \Rightarrow t'$ and $\Gamma \vdash t' : \sigma$. $\square$

By construction, $F_{ML}^?$ is sound. Elaboration preserves the type erasure, hence the semantics of terms.

Elaboration rules are given in syntax-directed form and are deterministic. They straightforwardly define an algorithm that, given Γ and t as input, returns the shape of t and an elaborated term t' . Thus, type inference problems in $F_{ML}^?$ can be solved as follows: given Γ and t , let elaboration compute both a term t' and a shape S such that $\Gamma \vdash_t t : S \Rightarrow t'$; let $\exists \beta. \sigma$ be $\lceil S \rceil$; infer a principal substitution φ for the type checking problem $\Gamma \vdash t' : \sigma$ in F_{ML} and returns the substitution φ and the type $\forall \lambda \Gamma \varphi. \sigma \varphi$. By construction this type inference algorithm is sound and complete with respect to Definition 1, because propagation is deterministic and type-checking in F_{ML} is sound and complete for the logical specification of F_{ML} .

One may argue that Definition 1 does not have a logical flavor. Indeed, we agree! However, we sustain the claim that it is simple and easy to understand. Ill-typed programs may be explained by providing both the elaborated program *and* an explanation of why it does not typecheck in F_{ML} . Either the elaborated program is not as expected, and the user should easily see why since elaboration is simple or he should understand the type error with respect to the elaborated program.

3.4 Elaboration with coercions

So far, we have elaborated coercions as constants. However, we may take advantage of elaboration to allow partial type information on explicit coercions as well. Let us allow to omit any side of an explicit coercion, *i.e.* to write $(t : \exists \beta. \triangleright \sigma_2)$, $(t : \exists \beta. \sigma_1 \triangleright)$ and $(t : \triangleright)$ respectively. Elaboration must return a fully specified coercion by filling in the missing part. We do this in the obvious ways, turning a coercion into a simple type annotation or simply discarding it when insufficient information is available. By lack of space, we refer to the full version for the corresponding elaboration rules.

Conversely, we may also allow elaboration to insert coercions that are not user-specified. However, in general, this would amount to guessing polymorphism. So we should only insert obvious coercions. In particular, when otherwise elaboration would lead to a typechecking error. There are such opportunities in rules VAR-C, FUNA-C, and several application rules. We say that the annotation S_1 is compatible with the annotation S_2 , which we write $S_1 \leq_p^{\eta-} S_2$, if there exists type schemes σ_1 and σ_2 of shape S_1 and S_2 such that $\sigma_1 \leq_p^{\eta-} \sigma_2$. For instance, $(\forall \alpha. \alpha) \rightarrow \#$ is not compatible with $(\forall \alpha. \alpha \rightarrow \alpha) \rightarrow \#$.

We may restrict rule VAR-C to cases where the shape S' of x is $\#$ or is compatible with the expected shape S . Otherwise, we may elaborate x as if it were $(x : \triangleright)$. Similarly, we may restrict FUNA-C to cases where S'_2 is $\#$ or compatible with $\lceil \sigma \rceil$. Otherwise, we may elaborate $\lambda z : \exists \beta. \sigma. t$ as if it were $\lambda z. \text{let } z = (z : \exists \beta. \triangleright \sigma) \text{ in } t$. Rules for applications have several places where elaboration may obviously lead to a typechecking failure. We may introduce a coercion on the function-side or on the argument-side. Or on both sides, but this would require guessing an intermediate type scheme at which both coercions would meet. Further investigation is necessary to find a reasonable and useful strategy for inserting coercions around applications.

3.5 Elaboration with an imperative semantics

Actually, elaboration does not depend on the differences between $F_{ML}(\leq_p^{\eta-})$ and $F_{ML}^v(\leq_p^{\eta-})$. To safely elaborate programs for $F_{ML}^v(\leq_p^{\eta-})$, we only need to replace typechecking in $F_{ML}(\leq_p^{\eta-})$ by typechecking in $F_{ML}^v(\leq_p^{\eta-})$. We write $F_{ML}^?$ for the corresponding language. This language is safe by construction, since elaboration preserves the semantics and the final typechecking ensure type-safety. Of course, type-safety is not sufficient to make $F_{ML}^v(\leq_p^{\eta-})$ an interesting language. One potential problem is that elaborated programs would then too often be rejected because of the value-restriction.

For instance, consider an expression t_1 of the form $(\lambda z. \lambda y. y) t_0$ where t_0 is a well-typed expansive expression. Let σ_{id} be $\forall \alpha. \alpha \rightarrow \alpha$. In $F_{ML}^?$ we have both $\vdash_t t_1 : \sigma_{id}$ and $\vdash_\downarrow t_1 : \sigma_{id}$. In F_{ML}^v , we do not have $\vdash_t t_1 : \sigma_{id}$ any more, because this will first infer a monomorphic type and fail to generalize at the end. Fortunately, we still have $\vdash_\downarrow t_1 : \sigma_{id}$. Note that if t_0 were non-expansive, then t_1 would also be

non-expansive and no annotation would be needed—just to emphasize the benefits of using a larger class of non-expansive expressions than just values and variables.

As another example consider the expression t_2 equal to $\text{let } f = (\lambda y. y : \sigma_{id}) \text{ in } (\lambda z. f) t_0$. Here, we even have $\vdash_{\uparrow} t_2 : \sigma_{id}$ because the polymorphic shape of f can be inferred. One may wonder why an explicit annotation on f is needed, since $F_{ML}^?$ can infer that $\lambda y. y$ has polymorphic type σ_{id} . Indeed, the problem is that elaboration is performed before type inference. A solution is to use incremental elaboration as described in Section 3.6

3.6 Incremental elaboration

Annotating let-bound expressions with ML types has shown to be sometimes useful. This may seem surprising, since these types can actually be inferred. The reason is that elaboration, which uses annotations, is performed before typechecking and does not see inferred let-bound polymorphism. Incremental elaboration is a solution to this problem. Instead of performing elaboration of the whole program followed by typechecking we may elaborate and typecheck programs by smaller parts, such as toplevel phrases. Consider, for instance, the expression $\text{let } z_1 = t_1 \text{ in } \dots \text{let } z_n = t_n \text{ in } t$. It can be seen as a succession of $n+1$ phrases, each of which but the last one augmenting the initial environment with a binding $z_k : \sigma_k$ where σ_k is the type inferred for t_k .

Assuming such a mechanism, no ML-like annotation would ever be useful on the toplevel bindings of $\bar{z}$. We may push this idea further and apply it to local bindings as well using the tree-structure of let-bindings to order the sequence of small elaboration followed by typechecking steps. However, this is getting (slightly) more complicated and loosing the simplicity of the two-step mechanism so that at the end the user may not so easily understand the elaboration process. Incremental elaboration at the level of toplevel phrases seems to be a good compromise.

4. Related works and conclusions

We have already widely discussed related works in the introduction and in particular the most closely related one, PVWS, that inspired this work. Another close work, developed in parallel with ours, is a recent proposal by Peyton-Jones *et al.* [VWP05], which we refer to below as VWP. It contributes a significant improvement over PVWS.

Leaving impredicativity aside, $F_{ML}(\leq_p^-)$ and PVWS have similar, but incomparable, expressiveness. That is, both have the capability to propagate annotations in both directions. However, they also differ in small details and there are examples that one can type and the other cannot. Since our elaboration is defined as a simple preprocessing step, it can easily be modified, *e.g.* to match PVWS more closely, but not entirely.

We have chosen to elaborate applications without propagation of information from the function to the argument (*e.g.* APP-I), as opposed to PVWS. Our choice is more symmetric and keeps the flow of elaboration bottom-up, which is more intuitive for the user to guess the elaborated program. Of course, we lose some information this way, at the benefit of more predictable elaboration and typechecking. Actually, the result of elaboration is then passed to typechecking, so some information may still flow sideways, *e.g.* from the function side to the argument side. However, such information does not depend on underneath first-order unification nor on the order in which typechecking is performed. This is actually another limit of our separation of elaboration and type-

checking. Incremental elaboration re-introduces some ordering in which typechecking must be performed, but in a controlled and intuitive manner.

The language VWP shares a lot with PVWS including the monolithic algorithmic specification of typechecking. However, VWP also improves over PVWS in at least two significant ways. First, checking and inference modes are carried on types rather than on typing judgments, much as for colored local type inference [OZZ01]. This provides with a more precise control of these modes. For instance, it allows to specify that some part of a type is known and to be checked while some other part is still unknown and to be inferred—a capability that we missed in the elaboration of applications. However, as a result, VWP is also quite involved and it seems quite difficult to guess in advance whether a program is typable without running the algorithm, which can hardly be done mentally or even on paper.

Second, VWP also allows for impredicative polymorphism. The treatment of impredicative polymorphism seems better integrated in VWP. However, it is also deeply hidden into algorithmic inference rules and an ad-hoc multi-argument typechecking rule for applications. This makes it difficult to understand the interaction between impredicative instantiation and predicative type-containment, since both seem to be left implicit, which clearly cannot be the case.

The monolithic type inference process of VWP seems to be more powerful in propagating known information than $F_{ML}^?$. However, it simultaneously mixes complicated orthogonal concepts in hard-wired algorithmic rules, and as PVWS, it lacks a simple specification.

We see at least three directions for future improvements. Improving the expressiveness of the core language is to be favored. Finding a stronger relation than $\leq_p^\eta$ for which first-order constraints would be easily solvable. In particular solving constraints of the form $\leq_p^{\eta-} \circ \leq^F \circ \leq_p^{\eta-}$ would enable explicit impredicative instantiation up to predicative containment by taking this relation for $\leq_I$. We should also explore properties of elaborations in more details, which we may then advantageously exploit. Finally, elaboration could probably be made more accurate by introducing variables ranging over type schemes to represent unknown information and avoid assigning them $\sharp$ a priori. This could also help with the elaboration of impredicative coercions. Hopefully, such solutions may avoid the need for more incremental elaboration, which could be confusing and less intuitive for the user. However, we might then lose our original goal to keep type inference within the resolution of first-order constraints.

If we are to lose this simplicity, we should also compare with ML^F [LBR03, LB04], which was designed so as to allow impredicative instantiation from the beginning, and does it rather well, by comparison with impredicative instantiation in $F_{ML}^?$. Conversely, ML^F does not have type-containment—and does not allow $\leq_p^\eta$ instantiations—but this does not seem to be a problem at all in practice. Already present in ML^F is the idea of elaboration to propagate annotations from interfaces to abstractions inside expressions, although elaboration used for ML^F remains a rather trivial process. ML^F has obviously a more powerful type inference engine and seems to perform better than $F_{ML}^?$. However its meta-theory is also more involved.

Conclusions

In summary, we have proposed a core language F_{ML} with first-order type inference with predicative type containment and explicit impredicative type-instantiation. A more convenient surface language $F_{ML}^?$ may be used to alleviate some repetitions of type annotations in source programs. It is defined by a simple elaboration procedure into the core language.

We believe that the decomposition of type inference into a simple elaboration procedure that propagates second-order type annotations followed by a first-order ML-like type inference mechanism is a good compromise between a logical and an algorithmic presentation. At least, it clearly separates the algorithmic elaboration process, which is complete by definition, but incomplete by nature, from the logical order-independent specification of type inference.

As observed, small variations in the logical specification, e.g. in the type-containment relation or the application of generalization may result in rather different type inference algorithms. This confirms, if it were necessary, that algorithmic specifications of type inference are fragile. Of course, algorithmic elaboration is a remedy. The goal remains to find expressive type systems with simple logical specifications for which we have complete inference algorithms.

Acknowledgments

I would like to particularly thank François Pottier for many fruitful discussions regarding this work. I am also grateful to Dimitrios Vytiniotis, Stephanie Weirich, and Simon Peyton Jones for their useful feedback concerning this paper and, more generally, for discussions on type inference with type-containment.

References

- [BCP99] Kim B. Bruce, Luca Cardelli, and Benjamin C. Pierce. Comparing object encodings. *Information and Computation*, 155(1/2):108–133, November 1999.
- [Gar04] Jacques Garrigue. Relaxing the value restriction. In *International Symposium on Functional and Logic Programming*, volume 2998 of *Lecture Notes in Computer Science*, Nara, April 2004. Springer-Verlag.
- [GR97] Jacques Garrigue and Didier Rémy. Extending ML with semi-explicit higher-order polymorphism. In Takayasu Ito and Martín Abadi, editors, *Theoretical Aspects of Computer Software*, volume 1281 of *Lecture Notes in Computer Science*, pages 20–46. Springer-Verlag, September 1997.
- [HP99] Haruo Hosoya and Benjamin C. Pierce. How good is local type inference? Technical Report MS-CIS-99-17, University of Pennsylvania, June 1999.
- [LB04] Didier Le Botlan. *MLF: Une extension de ML avec polymorphisme de second ordre et instantiation implicite*. PhD thesis, University of Paris 7, June 2004. (english version).
- [LBR03] 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.
- [Lei91] Daniel Leivant. Finitely stratified polymorphism. *Information and Computation*, 93(1):93–113, July 1991.
- [LO94] Konstantin Läuffer and Martin Odersky. Polymorphic type inference and abstract data types. *ACM Transactions on Programming Languages and Systems*, 16(5):1411–1430, September 1994.
- [Mit88] John C. Mitchell. Polymorphic type inference and containment. *Information and Computation*, 76:211–249, 1988.
- [OL96] Martin Odersky and Konstantin Läuffer. Putting type annotations to work. In *ACM Symposium on Principles of Programming*, pages 54–67, St. Petersburg, Florida, January 21–24, 1996. ACM Press.
- [OZZ01] Martin Odersky, Christoph Zenger, and Matthias Zenger. Colored local type inference. *ACM SIGPLAN Notices*, 36(3):41–53, March 2001.
- [Pfe88] Frank Pfenning. Partial polymorphic type inference and higher-order unification. In *Proceedings of the ACM Conference on Lisp and Functional Programming*, pages 153–163. ACM Press, July 1988.
- [PR] François Pottier and Didier Rémy. The essence of ML type inference. Extended version of [PR05] (in preparation).
- [PR05] François Pottier and Didier Rémy. The essence of ML type inference. In Benjamin C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 10, pages 389–489. MIT Press, 2005.
- [PS03] Simon Peyton Jones and Mark Shields. Lexically-scoped type variables. Submitted to ICFP 2004, March 2003.
- [PT00] Benjamin C. Pierce and David N. Turner. Local type inference. *ACM Trans. Program. Lang. Syst.*, 22(1):1–44, 2000.
- [PVWS05a] Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. Practical type inference for arbitrary-rank types. Submitted to the Journal of Functional Programming, June 2005.
- [PVWS05b] Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. Practical type inference for arbitrary-rank types. technical appendix. Private communication, June 2005.
- [Rém94] Didier Rémy. Programming objects with ML-ART: An extension to ML with abstract and record types. In Masami Hagiya and John C. Mitchell, editors, *Theoretical Aspects of Computer Software*, volume 789 of *Lecture Notes in Computer Science*, pages 321–346. Springer-Verlag, April 1994.
- [Rém05] Didier Rémy. Simple, partial type-inference for System F based on type-containment. Full version, September 2005.
- [TU02] Jerzy Tiuryn and Pawel Urzyczyn. The subtyping problem for second-order types is undecidable. *Information and Computation*, 179(1):1–18, 2002.
- [VWP05] Dimitrios Vytiniotis, Stephanie Weirich, and Simon Peyton Jones. Boxy type inference for higher-rank types and impredicativity. Available electronically at , April 2005.
- [Wel94] Joe B. Wells. Typability and type checking in the second-order λ -calculus are equivalent and undecidable. In *Proceedings of the Ninth Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 176–185, 1994.
- [Wel95] Joe B. Wells. The undecidability of Mitchell’s subtyping relation. Technical Report 95-019, Computer Science Department, Boston University, December 1995.
- [Wel96] Joe B. Wells. Typability is undecidable for $F + \eta$. Technical Report 96-022, Computer Science Department, Boston University, March 1996.
- [Wri95] Andrew K. Wright. Simple imperative polymorphism. *Lisp and Symbolic Computation*, 8(4):343–355, 1995.