

A Certified Implementation of ML with Structural Polymorphism

Jacques Garrigue

Graduate School of Mathematical Sciences,
Nagoya University, Chikusa-ku, Nagoya 464-8602
garrigue@math.nagoya-u.ac.jp

Abstract. The type system of Objective Caml has many unique features, which make ensuring the correctness of its implementation difficult. One of these features is structurally polymorphic types, such as polymorphic object and variant types, which have the extra specificity of allowing recursion. We implemented in Coq a certified interpreter for Core ML extended with structural polymorphism and recursion. Along with type soundness of evaluation, soundness and principality of type inference are also proved.

1 Introduction

While many results have already been obtained in the mechanization of metatheory for ML [13,6,5,11,19] and pure type systems [3,1], Objective Caml [12] has unique features which are not covered by existing works. For instance, polymorphic object and variant types require some form of structural polymorphism [8], combined with recursive types, and both of them do not map directly to usual type systems. Among the many other features, let us just cite the relaxed valued restriction [9], which accommodates side-effects in a smoother way, first class polymorphism [10] as used in polymorphic methods, labeled arguments [7], structural and nominal subtyping (the latter obtained through private abbreviations). There is plenty to do, and we are interested not only in type safety, but also in the correctness of type inference, as it gets more and more involved with each added feature.

Since it seems difficult to ensure the correctness of the current implementation, it would be nice to have a fully certified reference implementation at least for a subset of the language, so that one could check how it is supposed to work. As a first step, we certified type inference and evaluation for Core ML extended with local constraints, a form of structural polymorphism which allows inference of recursive types, such as polymorphic variants or objects. The formal proofs cover soundness of evaluation, both through rewriting rules and using a stack-based abstract machine, and soundness and completeness of the type inference algorithm.

While we based our developments on the “Engineering metatheory” methodology [1], our interest is in working on a concrete type system, with advanced typing features, like in the mechanized metatheory of Standard ML [11]. We are not so much concerned about giving a full specification of the operational semantics, as in [15].

The contribution of this paper is two-fold. First, the proofs presented here are original, and in particular it is to our knowledge the first proof of correctness of type inference for a type system containing recursive types, and even of type soundness for a system combining recursive types and a form of structural subsumption. Second, we have used extensively the techniques proposed in [1] to handle binding, and it is interesting to see how they fare in a system containing recursion, or when working on properties other than soundness. On the one hand we have been agreeably surprised by the compatibility of these techniques with explicit renaming (as necessary for type inference), but on the other hand one can easily get entangled in the plethora of quantifiers.

The Coq proof scripts and the extracted code can be found at:

<http://www.math.nagoya-u.ac.jp/~garrigue/papers/#certint1009>

Having them at hand while reading this paper should clarify many points. In particular, due to the size of some definitions, we could only include part of them in this paper, and we refer the reader to the proof scripts for all the missing details.

2 Structural Polymorphism

Structural polymorphism, embodied by polymorphic variants and objects, enriches types with both a form of width subsumption, and mutual recursive types. A type system for structural polymorphism was introduced in [8]. To help understand what we are working with, we repeat here the basic definitions, but please refer to the above paper for details.

Terms are the usual ones: variables, constants, functions, application and let-binding. We intend to provide all other constructs through constants and δ -rules.

$$e ::= x \mid c \mid \lambda x.e \mid e e \mid \text{let } x = e \text{ in } e$$

Types are less usual.

$$\begin{array}{ll} \tau ::= \alpha \mid \tau_1 \rightarrow \tau_2 & \text{type} \\ \kappa ::= \bullet \mid (C, \{l_1 \mapsto \tau_1, \dots, l_n \mapsto \tau_n\}) & \text{kind} \\ K ::= \alpha_1 :: \kappa_1, \dots, \alpha_n :: \kappa_n & \text{kinding environment} \\ \sigma ::= \forall \tilde{\alpha}. K \triangleright \tau & \text{polytype} \end{array}$$

A *type* is either a type variable or a function type. This may seem not expressive enough, but in this system type variables need not be abstract, as a *kinding environment* associates them with their respective *kinds*. When they are associated with a concrete kind, they actually denote structural types, like records or variants. Such types are described by a pair (C, R) of a local constraint C and a mapping¹ R from labels to types. On the other hand \bullet just denotes an (abstract) type variable. As you can see, type variables may appear inside kinds, and since kinding environments are allowed to be recursive,

¹ In order to make type inference principal, this “mapping” is not always a function (*i.e.* the same l may, under some conditions, be related to several τ 's), but this should not matter at the level of detail of this paper.

<p>VARIABLE $\frac{K, K_0 \vdash \theta : K \quad \text{dom}(\theta) \subset B}{K; \Gamma, x : \forall B. K_0 \triangleright \tau \vdash x : \theta(\tau)}$</p> <p>ABSTRACTION $\frac{K; \Gamma, x : \tau \vdash e : \tau'}{K; \Gamma \vdash \lambda x. e : \tau \rightarrow \tau'}$</p> <p>LET $\frac{K; \Gamma \vdash e_1 : \sigma \quad K; \Gamma, x : \sigma \vdash e_2 : \tau}{K; \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau}$</p>	<p>CONSTANT $\frac{K_0 \vdash \theta : K \quad T\text{const}(c) = K_0 \triangleright \tau}{K; \Gamma \vdash c : \theta(\tau)}$</p> <p>APPLICATION $\frac{K; \Gamma \vdash e_1 : \tau \rightarrow \tau' \quad K; \Gamma \vdash e_2 : \tau}{K; \Gamma \vdash e_1 e_2 : \tau'}$</p> <p>GENERALIZE $\frac{K; \Gamma \vdash e : \tau \quad B = \text{FV}_K(\tau) \setminus \text{FV}_K(\Gamma)}{K _{K B}; \Gamma \vdash e : \forall B. K _B \triangleright \tau}$</p>
---	--

Fig. 1. Typing rules (original)

<p>VARIABLE $\frac{K \vdash \bar{\tau} :: \bar{\kappa}^{\bar{\tau}}}{K; \Gamma, x : \bar{\kappa} \triangleright \tau_1 \vdash x : \tau_1^{\bar{\tau}}}$</p> <p>ABSTRACTION $\frac{\forall x \notin L \quad K; \Gamma, x : \tau \vdash e^x : \tau'}{K; \Gamma \vdash \lambda e : \tau \rightarrow \tau'}$</p> <p>LET $\frac{K; \Gamma \vdash e_1 : \sigma \quad \forall x \notin L \quad K; \Gamma, x : \sigma \vdash e_2^x : \tau}{K; \Gamma \vdash \text{let } e_1 \text{ in } e_2 : \tau}$</p>	<p>CONSTANT $\frac{K \vdash \bar{\tau} :: \bar{\kappa}^{\bar{\tau}} \quad T\text{const}(c) = \bar{\kappa} \triangleright \tau_1}{K; \Gamma \vdash c : \tau_1^{\bar{\tau}}}$</p> <p>APPLICATION $\frac{K; \Gamma \vdash e_1 : \tau \rightarrow \tau' \quad K; \Gamma \vdash e_2 : \tau}{K; \Gamma \vdash e_1 e_2 : \tau'}$</p> <p>GENERALIZE $\frac{\forall \bar{\alpha} \notin L \quad K, \bar{\alpha} :: \bar{\kappa}^{\bar{\alpha}}; \Gamma \vdash e : \tau^{\bar{\alpha}}}{K; \Gamma \vdash e : \bar{\kappa} \triangleright \tau}$</p>
--	---

Fig. 2. Typing rules using cofinite quantification

we can use them to define recursive types (where the recursion must necessarily go through kinds.) Since type variables only make sense in presence of a kinding environment, *polytypes* have to include a kinding environment for the variables they quantify; *i.e.*, in $\forall \bar{\alpha}. K \triangleright \tau$, K is such that $\text{dom}(K) = \{\bar{\alpha}\}$, and the variables of $\bar{\alpha}$ may appear both inside the kinds of K and in τ . A good way to understand these definitions is to see types as directed graphs, where variables are just labels for nodes.

This type system is actually a framework, where the concrete definition of local constraints, and how they interact with types, is kept abstract. One can then apply this framework to an appropriate *constraint domain* to implement various flavours of polymorphic variants and records. A constraint domain \mathcal{C} is a set of constraints combined with an entailment relation \models on these constraints, and a predicate $\text{unique}(C, l)$ telling whether l may map to several types, satisfying some properties. By extension we also use the notation $\kappa' \models \kappa$ for kinds, *i.e.* $(C', R') \models (C, R)$ iff $C' \models C$ and $R \subset R'$.

Kinding environments are used in two places: in polytypes where they associate kinds to quantified type variables, and in typing judgments, which are of the form $K; \Gamma \vdash e : \tau$, where the variables kinded in K may appear in both Γ and τ . The typing rules are given in Fig. 1. $K \vdash \theta : K'$ means that the substitution θ (defined as usual) preserves kinds between K and K' (it is *admissible* between K and K'). Formally, if α has a concrete kind in K ($\alpha :: \kappa \in K$, $\kappa \neq \bullet$), then $\theta(\alpha) = \alpha'$ is a variable, and it has a more concrete kind in K' ($\alpha' :: \kappa' \in K'$ and $\kappa' \models \theta(\kappa)$). The main difference with Core ML is that GENERALIZE has to split the kinding environment into a generalized part, which contains the kinds associated to generalized type variables (denoted by $K|_B$), and a non-generalized part for the rest (denoted by $K|_{K|B}$). When determining which type

<p>R-ABS $(\lambda e_1) v_2 \longrightarrow e_1^{v_2}$</p> <p>R-LET $\text{let } v_1 \text{ in } e_2 \longrightarrow e_2^{v_1}$</p>	<p>R-DELTA $\frac{e = \text{Delta.reduce } c [v_1; \dots; v_n]}{c v_1 \dots v_n \longrightarrow e}$</p> <p>R-LET₁ $\frac{e_1 \longrightarrow e'_1}{\text{let } e_1 \text{ in } e_2 \longrightarrow \text{let } e'_1 \text{ in } e_2}$</p>	<p>R-APP₁ $\frac{e_1 \longrightarrow e'_1}{e_1 e_2 \longrightarrow e'_1 e_2}$</p> <p>R-APP₂ $\frac{e_2 \longrightarrow e'_2}{v_1 e_2 \longrightarrow v_1 e'_2}$</p>
---	--	---

Fig. 3. Reduction rules

variables can be generalized, we must be careful that for any type variable accessible from Γ , the type variables appearing in its kind (inside K) are also accessible. For this reason FV takes K as parameter; if $\alpha :: \kappa \in K$, then $FV_K(\alpha) = \{\alpha\} \cup FV_K(\kappa)$.

It may be difficult to understand this type system in abstract form. Concrete constraint domains and constants are given in Fig. 6 and 7, and an example appears in Section 7 of this paper.

3 Type Soundness

The first step of our mechanical proof, using Coq [17], was to prove type soundness for the system described in the previous section, starting from Aydemir and others' proof for Core ML included in [1], which uses *locally nameless cofinite quantification*. This proof uses de Bruijn indices for local quantification inside terms and polytypes, and quantifies over an abstract avoidance set for avoiding name conflicts.

Fig. 2 contains the typing rules adapted to locally nameless cofinite quantification, and the reduction rules are in Fig. 3. They both use locally nameless terms and types.

$e ::= n \mid x \mid c \mid \lambda e \mid e e \mid \text{let } e \text{ in } e$	term
$\tau ::= n \mid \alpha \mid \tau_1 \rightarrow \tau_2$	type
$\kappa ::= \bullet \mid (C, \{l_1 \mapsto \tau_1, \dots, l_n \mapsto \tau_n\})$	kind
$\sigma ::= \bar{\kappa} \triangleright \tau$	polytype

$\bar{\tau}$ and $\bar{\kappa}$ represent sequences of types and kinds. When we write $\bar{\alpha}$, we also assume that all type variables inside the sequence are distinct. Polytypes are now written $\bar{\kappa} \triangleright \tau$, where the length of $\bar{\kappa}$ is the number of generalized type variables, represented as de Bruijn indices $1 \dots n$ inside types². $\tau_1^{\bar{\tau}}$ is τ_1 where de Bruijn indices were substituted with types of $\bar{\tau}$, accessed by their position. Similarly $\bar{\kappa}^{\bar{\tau}}$ substitute all the indices inside the sequence $\bar{\kappa}$. e^x only substitutes x for the index 1. $K \vdash \tau :: \kappa$ is true when either $\kappa = \bullet$, or $\tau = \alpha$, $\alpha :: \kappa' \in K$ and $\kappa' \models \kappa$. $K \vdash \bar{\tau} :: \bar{\kappa}$ enforces this for every member of $\bar{\tau}$ and $\bar{\kappa}$ at identical positions, which is just equivalent to our condition $K \vdash \theta : K'$ for the preservation of kinds.

$\forall x \notin L$ and $\forall \alpha \notin L$ are cofinite quantifications, with scope the hypotheses on the right of the quantifier. Each L appearing in a derivation is existentially quantified (*i.e.* one chooses a concrete L when building the derivation), but has to be finite, to allow an infinite number of variables outside of L . At first, the rules may look very different from

² The implementation has indices starting from 0, but we will start from 1 in this explanation.

```

Module Type CstrIntf.
  Parameter cstr attr : Set.                (* types for abstract constraints and labels *)
  Parameter valid : cstr → Prop.           (* validity of a constraint *)
  Parameter unique : cstr → attr → bool.  (* uniqueness of a label *)
  Parameter  $\sqcup$  : cstr → cstr → cstr.    (* least upper bound *)
  Parameter  $\models$  : cstr → cstr → Prop.    (* entailment between constraints *)
  ...                                       (* some properties of these definitions *)

Module Type CstIntf.
  Parameter const : Set.                   (* constants *)
  Parameter arity : const → nat.          (* their arity *)

```

Fig. 4. Interfaces for constraints and constants

those in Fig. 1, but they coincide if we instantiate L appropriately. For instance, if we use $\text{dom}(\Gamma)$ for L in $\forall x \notin L$, this just amounts to ensuring that x is not already bound. Inside *GENERALIZE*, we could use $\text{dom}(\mathbf{K}) \cup \text{FV}_{\mathbf{K}}(\Gamma)$ for L to ensure that the newly introduced variables are locally fresh. This may not be intuitive, but this is actually a very clever way to encode naming constraints implicitly. Moreover, when we build a new typing derivation from an old one, we can avoid renaming variables by just enlarging the avoidance sets.

Starting from an existing proof was a tremendous help, but many new definitions were needed to accommodate kinds, and some existing ones had to be modified. For instance, in order to accommodate the mutually recursive nature of kinding environments, we need simultaneous type substitutions, rather than the iterated ones of the original proof. The freshness of individual variables (or sequences of variables: $\bar{\alpha} \notin L$) becomes insufficient, and we need to handle disjointness conditions on sets ($L_1 \cap L_2 = \emptyset$). As a result, the handling of freshness, which was almost fully automatized in the proof of Core ML, required an important amount of work with kinds, even after developing some tactics for disjointness.

We also added a formalism for constants and δ -rules, which are needed to give an operational semantics to structural types. Overall, the result was a doubling of the size of the proof, from 1000 lines to more than 2000, but the changes were mostly straightforward. This does not include the extra metatheory lemmas and set inclusion tactics that we use for all proofs.

The formalism of local constraints was defined as a framework, able to handle various flavours of variant and object types, just by changing the constraint part of the system. This was formalized through the use of functors. The signature for constraints and constants is in Fig. 4, and an outline of the module structure of the soundness proof (including the statements proved) is in Fig. 5. We omit here the definitions of terms, types, typing derivations, and reduction, as they just implement the locally nameless definitions we described above. A value is either a λ -abstraction, or a constant applied to a list of values of length less than its arity.

This approach worked well, but there are some drawbacks. One is that since some definitions depend on parameters of the framework, and some of the proofs required by the framework depend on those definitions, we need nested functors, and the instantiation of the framework with a *constraint domain* looks like a “dialogue”: we repeatedly

```

Module MkDefs (Cstr : CstrIntf) (Const : CstIntf).
  Inductive typ : Set := ... (* our types *)
  Inductive type : typ → Prop := ... (* well-formed types *)
  Inductive trm : Set := ... (* our terms *)
  ...

Module Type DeltaIntf.
  Parameter type : Const.const → sch. (* types of constants *)
  Parameter reduce : ∀ c el, (list_for_n value (1 + Const.arity c) el) → trm. (* δ-rules *)
  ... (* 3 more properties *)

Module MkJudge (Delta : DeltaIntf).
  Inductive ⊢ : kenv → env → trm → typ → Prop := ... (* the typing judgment *)
  Inductive → : trm → trm → Prop := ... (* the reduction relation *)
  Inductive value : trm → Prop := ... (* values *)
  ...

Module Type SndHypIntf.
  Parameter delta_typed : ∀ c el vl K Γ τ,
    (K; Γ ⊢ const_app c el : τ) → (K; Γ ⊢ Delta.reduce c el vl : τ).

Module MkSound (SH : SndHypIntf).
  Theorem preservation : ∀ K Γ e e' τ, (K; Γ ⊢ e : τ) → (e → e') → (K; Γ ⊢ e' : τ).
  Theorem progress : ∀ K e τ, (K; ∅ ⊢ e : τ) → (value e ∨ ∃ e', e → e').

```

Fig. 5. Module structure

alternate domain-specific definitions, and applications of framework functors to those definitions, each new definition using the result of the previous functor application. The problem appears not so much with constraints themselves, but rather with constants and δ -rules. In order to obtain the definitions for typing judgments, one has to provide implementations for constraints and constants, extract the definition of types and terms, and use them to provide constant types and δ -rules. We enforce the completeness of δ -rules by requiring a function `reduce` which will be applied to a list of values of length $(1 + \text{Const.arity } c)$; through well-typedness they will be only used if `Const.arity c` is smaller than the arity of type `c`. Type soundness itself is another functor, that requires some lemmas whose proofs may use infrastructure lemmas on type judgments, and returns proofs of preservation and progress. The real structure is even more complex, because the proofs span several files, and each file must mimick this structure. The same problem is known to occur in programs using heavily ML functors, so this is not specific to Coq. But the level of stratification of definitions we see in this proof rarely occurs in programs.

This instantiation has been done for a constraint domain containing both polymorphic variants and records, and a fixpoint operator. We show the constraint domain in Fig. 6; we write $\langle \rangle$ for `None`, which denotes here the set of all possible labels. Constants and δ -rules are in Fig. 7, using the nameful syntax for types. You can see the duality between variants and records, at least for `tag` and `get`.

Both in the framework and domain proofs, cofinite quantification demonstrated its power, as no renaming of type or term variables was needed at all. It helped also in an indirect way: in the original rule for `GENERALIZE`, one has to close the set of free

Module Cstr.

Definition $\text{attr} := \text{nat}$.

Inductive $\text{ksort} : \text{Set} := \text{Ksum} \mid \text{Kprod} \mid \text{Kbot}$.

Record $\text{cstr} : \text{Set} := \text{C} \{ \text{sort} : \text{ksort}; \text{low} : \text{list nat}; \text{high} : \text{option}(\text{list nat}) \}$.

Definition $\text{valid } c := \text{sort } c \neq \text{Kbot} \wedge (\text{high } c = \langle \rangle \vee \text{low } c \subset \text{high } c)$.

Definition $s_1 \leq s_2 := s_1 = \text{Kbot} \vee s_1 = s_2$.

Definition $c_1 \models c_2 :=$

$\text{sort } s_2 \leq \text{sort } s_1 \wedge \text{low } c_2 \subset \text{low } c_1 \wedge (\text{high } c_2 = \langle \rangle \vee \text{high } c_1 \subset \text{high } c_2)$.

...

Fig. 6. Constraint domain for polymorphic variants and records

$$\begin{aligned}
 \text{type}(\text{tag}_l) &= \alpha :: (\langle \text{Ksum}, \{l\}, \langle \rangle \rangle, \{l \mapsto \beta\}) \triangleright \beta \rightarrow \alpha \\
 \text{type}(\text{match}_{l_1 \dots l_n}) &= \alpha :: (\langle \text{Ksum}, \emptyset, \{l_1, \dots, l_n\} \rangle, \{l_1 \mapsto \alpha_1, \dots, l_n \mapsto \alpha_n\}) \\
 &\quad \triangleright (\alpha_1 \rightarrow \beta) \rightarrow \dots \rightarrow (\alpha_n \rightarrow \beta) \rightarrow \alpha \rightarrow \beta \\
 \text{type}(\text{record}_{l_1 \dots l_n}) &= \alpha :: (\langle \text{Kprod}, \emptyset, \{l_1, \dots, l_n\} \rangle, \{l_1 \mapsto \alpha_1, \dots, l_n \mapsto \alpha_n\}) \\
 &\quad \triangleright \alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow \alpha \\
 \text{type}(\text{get}_l) &= \alpha :: (\langle \text{Kprod}, \{l\}, \langle \rangle \rangle, \{l \mapsto \beta\}) \triangleright \alpha \rightarrow \beta \\
 \text{type}(\text{recf}) &= ((\alpha \rightarrow \beta) \rightarrow (\alpha \rightarrow \beta)) \rightarrow (\alpha \rightarrow \beta) \\
 \\
 \text{match}_{l_1 \dots l_n} f_1 \dots f_n (\text{tag}_{l_i} e) &\longrightarrow f_i e \\
 \text{get}_{l_i} (\text{record}_{l_1 \dots l_n} e_1 \dots e_n) &\longrightarrow e_i \\
 \text{recf } f e &\longrightarrow f (\text{recf } f) e
 \end{aligned}$$

Fig. 7. Types and δ -rules for constants

variables of a type with the free variables of their kinds; but the cofinite quantification takes care of that implicitly, without any extra definitions.

While cofinite quantification may seem perfect, there is a pitfall in this perfection itself. One forgets that some proof transformations intrinsically require variable renaming. Concretely, to make typing more modular, I added a rule that discards irrelevant kinds from the kinding environment. Fig. 8 shows both the normal and cofinite forms. Again one can see the elegance of the cofinite version, where there is no need to specify which kinds are irrelevant: just the ones whose names have no impact on typability. Proofs went on smoothly, until I realized that I needed the following inversion lemma, relating derivations using KIND GC , and those without it.

$$\forall \text{K} \Gamma e \tau, (\text{K}; \Gamma \vdash_{\text{GC}} e : \tau) \rightarrow \exists \text{K}', (\text{K}, \text{K}'; \Gamma \vdash e : \tau)$$

Namely, by putting back the kinds we discarded, we shall be able to obtain a derivation that does not rely on KIND GC . This is very intuitive, but since this requires making KIND GC commute with GENERALIZE , we end up commuting quantifiers. And this is just impossible without a true renaming lemma. I got stuck there for a while, unable to see what was going wrong³. Even more confusing, the same problem occurs when we try to make KIND GC commute with ABSTRACTION , whereas intuitively the choice of names for term variables is independent of the choice of names for type variables.

³ Thanks to Arthur Charguéraud for opening my eyes.

$$\frac{\text{KIND GC} \quad \text{K}, \text{K}'; \Gamma \vdash e : \tau \quad \text{FV}_K(\Gamma, \tau) \cap \text{dom}(\text{K}') = \emptyset}{\text{K}; \Gamma \vdash e : \tau} \qquad \frac{\text{CO-FINITE KIND GC} \quad \forall \bar{\alpha} \notin L \quad \text{K}, \bar{\alpha} :: \bar{\kappa}^{\bar{\alpha}}; \Gamma \vdash e : \tau}{\text{K}; \Gamma \vdash e : \tau}$$

Fig. 8. Kind discarding

Finally this lemma required about 1000 lines to prove it, including renaming lemmas for both term and type variables.

$$\begin{aligned} \text{Lemma typing_rename} &: \forall \text{K} \Gamma x y \sigma \Gamma' e \tau, \\ &\text{K}; \Gamma, x : \sigma, \Gamma' \vdash e : \tau \rightarrow y \notin \text{dom}(\Gamma, \Gamma') \cup \{x\} \cup \text{FV}(e) \rightarrow \text{K}; \Gamma, y : \sigma, \Gamma' \vdash [y/x]e : \tau. \\ \text{Lemma typing_rename_typ} &: \forall \text{K} \Gamma \bar{\kappa} \tau \bar{\alpha} \bar{\alpha}' e, \\ &\bar{\alpha} \notin \text{FV}(\Gamma) \cup \text{FV}(\bar{\kappa} \triangleright \tau) \cup \text{dom} \text{K} \cup \text{FV}(\text{K}) \rightarrow \bar{\alpha}' \notin \text{dom} \text{K} \cup \{\bar{\alpha}\} \rightarrow \\ &\text{K}, \bar{\alpha} :: \bar{\kappa}^{\bar{\alpha}}; \Gamma \vdash e : \tau^{\bar{\alpha}} \rightarrow \text{K}, \bar{\alpha}' :: \bar{\kappa}^{\bar{\alpha}'}; \Gamma \vdash e : \tau^{\bar{\alpha}'}. \end{aligned}$$

The renaming lemmas were harder to prove than expected (100 lines each). Contrary to what was suggested in [1], we found it rather difficult to prove these lemmas starting from the substitution lemmas of the soundness proof; while renaming for types used this approach, renaming for terms was proved by a direct induction, and they ended up being of the same length. On the other hand, one could argue that the direct proof was easy precisely thanks to cofinite quantification, which eschews the need for extra machinery.

Once the essence of the problem (*i.e.* commutation of quantifiers) becomes clear, one can see a much simpler solution: in most situations, it is actually sufficient to have KIND GC occur only just above ABSTRACTION and GENERALIZE, and the canonicalization lemma is just 100 lines, as it doesn't change the quantifier structure of the proof. This also raises the issue of how to handle several variants of a type system in the same proof. Here this was done by parameterizing the predicate \vdash with the canonicity of the derivation, and whether KIND GC is allowed at this point. This gives 4 cases for the availability of KIND GC: allowed nowhere, allowed everywhere, or inside a canonical derivation where it is allowed or not at the current point. Functions `gc_ok`, `gc_raise` and `gc_lower`, which are used by the definitions themselves, allow to manipulate this state transparently.

4 Type Inference

The main goal of using local constraints was to keep the simplicity of unification-based type inference. Of course, unification has to be extended in order to handle kinding, but the algorithms for unification and type inference stay reasonably simple.

4.1 Unification

Unification has been a target of formal verification for a long time, with formal proofs as early as 1985 [16]. Here we just wrote down the algorithm in Coq, and proved

$[\tilde{\alpha}] \tau = \tau_*$ <p style="text-align: center;">such that $\tau_*^{\tilde{\alpha}} = \tau$ and $FV(\tau_*) \cap \tilde{\alpha} = \emptyset$</p> $[\tilde{\alpha}](\bar{\kappa} \triangleright \tau) = ([\tilde{\alpha}]\bar{\kappa} \triangleright [\tilde{\alpha}]\tau)$ <p>Definition $\text{generalize}(K, \Gamma, L, \tau) :=$ let $A = FV_K(\Gamma)$ and $B = FV_K(\tau)$ in let $K' = K _{K \setminus A}$ in let $\tilde{\alpha} :: \bar{\kappa} = K' _B$ in let $\tilde{\alpha}' = B \setminus (A \cup \tilde{\alpha})$ in let $\bar{\kappa}' = \text{map}(\lambda _ . \bullet) \tilde{\alpha}'$ in $\langle (K _A, K' _L), [\tilde{\alpha}'\tilde{\alpha}](\bar{\kappa}' \triangleright \tau) \rangle$.</p>	<p>Definition $\text{typinf}(K, \Gamma, \text{let } e_1 \text{ in } e_2, \tau, \theta, L) :=$ let $\alpha = \text{fresh}(L)$ in match $\text{typinf}(K, \Gamma, e_1, \alpha, \theta, L \cup \{\alpha\})$ with $\langle K', \theta', L' \rangle \Rightarrow$ let $K_1 = \theta'(K')$ and $\Gamma_1 = \theta'(\Gamma)$ in let $L_1 = \theta'(\text{dom}(K))$ and $\tau_1 = \theta'(\tau)$ in let $\langle K_A, \sigma \rangle = \text{generalize}(K_1, \Gamma_1, L_1, \tau_1)$ in let $x = \text{fresh}(\text{dom}(\Gamma) \cup FV(e_1) \cup FV(e_2))$ in $\text{typinf}(K_A, (\Gamma, x : \sigma), e_2^x, \tau, \theta', L')$ $\langle \rangle \Rightarrow \langle \rangle$ end.</p>
--	---

Fig. 9. Type inference algorithm

both partial-correctness and completeness. A rule-based version of the algorithm can be found in [8]. The following statements were proved:

Definition $\text{unifies } \theta \ l := \forall \tau_1 \tau_2, \text{In } (\tau_1, \tau_2) \ l \rightarrow \theta(\tau_1) = \theta(\tau_2)$.

Theorem $\text{unify_types} : \forall h \ l \ K \ \theta, \text{unify } h \ l \ K \ \theta = \langle K', \theta' \rangle \rightarrow \text{unifies } \theta' \ l$.

Theorem $\text{unify_kinds} : \forall h \ l \ K \ \theta,$
 $\text{unify } h \ l \ K \ \theta = \langle K', \theta' \rangle \rightarrow \text{dom}(\theta) \cap \text{dom}(K) = \emptyset \rightarrow$
 $K \vdash \theta' : \theta'(K') \wedge \text{dom}(\theta') \cap \text{dom}(K') = \emptyset$.

Theorem $\text{unify_mgu} : \forall h \ l \ K_0 \ K \ \theta,$
 $\text{unify } h \ l \ K_0 \ \text{id} = \langle K, \theta \rangle \rightarrow \text{unifies } \theta' \ l \rightarrow K_0 \vdash \theta' : K' \rightarrow \theta' \sqsupseteq \theta \wedge K \vdash \theta : K'$.

Theorem $\text{unify_complete} : \forall K \ \theta \ K_0 \ l \ h,$
 $\text{unifies } \theta \ l \rightarrow K_0 \vdash \theta : K \rightarrow \text{size_pairs } \text{id } K_0 \ l < h \rightarrow \text{unify } h \ l \ K_0 \ \text{id} \neq \langle \rangle$.

The first argument to unify is the number of type variables, which is used to enforce termination. Then come a list of type pairs to unify and the original kinding environment. Last is a starting substitution, so that the algorithm is tail-recursive. To keep the statement clear, well-formedness conditions are omitted here. The proof is rather long, as kinds need particular treatment, but there was no major stumbling block. The proof basically follows the algorithms, but there are two useful tricks. One concerns substitutions. Rather than using the relation “ θ is more general than θ' ” ($\exists \theta_1, \theta' = \theta_1 \circ \theta$), we used the more direct “ θ' extends θ ” ($\forall \alpha, \theta'(\theta(\alpha)) = \theta'(\alpha)$). In the above theorem it is noted $\theta' \sqsupseteq \theta$. When θ is idempotent, the two definitions are equivalent, but the latter can be used directly through rewriting. The other idea was to define a special induction lemma for successful unification, which uses symmetries to reduce the number of cases to check. Unification being done on first-order terms, the types we are unifying shall contain no de Bruijn indices, but only global variables. Since we started with a representation allowing both kinds of variables, there was no need to change it.

4.2 Inference

The next step is type inference itself. Again, correctness has been proved before for Core ML [13,6,19], but to our knowledge never for a system containing equi-recursive types.

$$\begin{array}{ll}
\text{Theorem soundness : } \forall K \Gamma e \tau \theta L K' \theta' L', & \text{Theorem principality : } \forall K \Gamma e \tau \theta K_1 \theta_1 L, \\
\text{typinf}(K, \Gamma, e, \tau, \theta, L) = \langle K', \theta', L' \rangle \rightarrow & K; \theta(\Gamma) \vdash e : \theta(\tau) \rightarrow K_1 \vdash \theta : K \rightarrow \\
\text{dom}(\theta) \cap \text{dom}(K) = \emptyset \rightarrow & \theta \sqsupseteq \theta_1 \rightarrow \text{dom}(\theta_1) \cap \text{dom}(K_1) = \emptyset \rightarrow \\
\text{FV}(\theta, K, \Gamma, \tau) \subset L \rightarrow & \text{dom}(\theta) \cup \text{FV}(\theta_1, K_1, \Gamma, \tau) \subset L \rightarrow \\
\theta'(K'); \theta'(\Gamma) \vdash e : \theta'(\tau) \wedge & \exists K' \theta' L', \\
K \vdash \theta' : \theta'(K') \wedge \theta' \sqsupseteq \theta \wedge & \text{typinf}(K_1, \Gamma, e, \tau, \theta_1, L) = \langle K', \theta', L' \rangle \wedge \\
\text{FV}(\theta', K', \Gamma) \cup L \subset L' \wedge & \exists \theta'', K' \vdash \theta \theta'' : K \wedge \theta \theta'' \sqsupseteq \theta' \wedge \\
\text{dom}(\theta') \cap \text{dom}(K') = \emptyset. & \text{dom}(\theta'') \subset L' \setminus L.
\end{array}$$

Fig. 10. Properties of type inference

Proving both soundness and principality was rather painful. This time one problem was the complexity of the algorithm itself, in particular the behaviour of type generalization. The usual behaviour for ML is just to find the variables that are not free in the typing environment and generalize them, but with a kinding environment several extra steps are required. First, the free variables should be closed transitively using the kinding environment. Then, the kinding environment also should be split into generalizable and non-generalizable parts. Last, some generalizable parts of the kinding environment need to be duplicated, as they might be used independently in some other parts of the typing derivation. The definitions for `generalize` and the `let` case of `typinf` are shown in Fig. 9. $[\bar{\alpha}] \tau$ stands for the generalization of τ with respect to $\bar{\alpha}$, obtained by replacing the occurrences of variables of $\bar{\alpha}$ in τ by their indices.

Due to the large number of side-conditions required, the statements for the inductive versions of soundness of principality become very long. In Fig. 10 we show slightly simplified versions, omitting well-formedness properties. These statements can be proved directly by induction. From those, we can derive the following corollaries for a simplified version of `typinf`, taking only a term and a closed environment as arguments.

$$\begin{array}{l}
\text{Corollary soundness' : } \forall K \Gamma e \tau, \text{FV}(\Gamma) = \emptyset \rightarrow \text{typinf}' \Gamma e = \langle K, \tau \rangle \rightarrow K; \Gamma \vdash e : \tau. \\
\text{Corollary principality' : } \forall K \Gamma e \tau, \text{FV}(\Gamma) = \emptyset \rightarrow K; \Gamma \vdash e : \tau \rightarrow \\
\exists K', \exists T', \text{typinf}' \Gamma e = \langle K', T' \rangle \wedge \exists \theta, K' \vdash \theta : K \wedge \tau = \theta(T').
\end{array}$$

As usual, the proof of principality requires the following lemma, which states that if a term e has a type τ under an environment Γ , then we can give it the same type under a more general environment Γ_1 .

$$\text{Lemma typing_moregen : } \forall K \Gamma \Gamma_1 e \tau, K; \Gamma \vdash e : \tau \rightarrow K \vdash \Gamma_1 \leq \Gamma \rightarrow K; \Gamma_1 \vdash e : \tau.$$

$K \vdash \Gamma_1 \leq \Gamma$ means that the polytypes of Γ are instances of those in Γ_1 . Due to the presence of kinds, the definition of the instantiation order gets a bit complicated.

$$K \vdash \bar{\kappa}_1 \triangleright \tau_1 \leq \bar{\kappa} \triangleright \tau \stackrel{\text{def}}{=} \forall \bar{\alpha}, \text{dom}(K) \cap \bar{\alpha} = \emptyset \rightarrow \exists \bar{\tau}, K, \bar{\alpha} :: \bar{\kappa}^{\bar{\alpha}} \vdash \bar{\tau} :: \bar{\kappa}_1^{\bar{\tau}} \wedge \tau_1^{\bar{\tau}} = \tau^{\bar{\alpha}}.$$

It may be easier to consider the version without de Bruijn indices.

$$K \vdash \forall \bar{\alpha}_1. K_1 \triangleright \tau_1 \leq \forall \bar{\alpha}_2. K_2 \triangleright \tau_2 \stackrel{\text{def}}{=} \exists \theta, \text{dom}(\theta) \subset \bar{\alpha}_1 \wedge K, K_1 \vdash \theta : K, K_2 \wedge \theta(\tau_1) = \tau_2.$$

Another difficulty is that, since we are building a derivation, cofinite quantification appears as a requirement rather than a given, and we need renaming for both terms and

```

Inductive clos : Set :=
| clos_abs   : trm → list clos → clos
| clos_const : Const.const → list clos → clos.

Fixpoint clos2trm(c : clos) : trm :=
match c with
| clos_abs e l ⇒ trm_inst (λe) (map clos2trm l)
| clos_const c l ⇒ const_app c (map clos2trm l)
end.

Record frame : Set := Frame {frm_benv : list clos; frm_app : list clos; frm_trm : trm}.

Inductive eval_res : Set :=
| Result : nat → clos → eval_res
| Inter  : list frame → eval_res.

Fixpoint eval (h : nat) (benv : list clos) (app : list clos) (e : trm) (stack : list frame)
{struct h} : eval_res := ...

Theorem eval_sound : ∀hKeτ,
(K; Γ ⊢ e : τ) → (K; Γ ⊢ res2trm (eval h nil nil t nil) : τ).

Theorem eval_complete : ∀Ke'e'τ,
(K; Γ ⊢ e : τ) → (e  $\xrightarrow{*}$  e') → value e' →
∃h, ∃cl, eval h nil nil t nil = Result 0 cl ∧ e' = clos2trm cl.

```

Fig. 11. Definitions and theorems for stack-based evaluation

types in many places. This is true both for soundness and principality, since in the latter the type variables of the inferred derivation and of the provided derivation are different. As a result, while we could finally avoid using the renaming lemmas for type soundness, they were ultimately needed for type inference.

5 Interpreter

Type soundness ensures that evaluation according to a set of source code rewriting rules cannot go wrong. However, programming languages do not evaluate a program by rewriting it, but rather interpreting it with a virtual machine. We defined a stack-based abstract machine, and proved that at every step the state of the abstract machine could be converted back to a term whose typability was a direct consequence of the typability of the reduced program. This ensures that evaluation cannot go wrong, and the final result, if reached, shall be either a constant or a function closure. Once the relation between program and state was properly specified, the proof was mostly straightforward.

The basic definitions and the statements for soundness and completeness are in Fig. 11. We omit here the concrete definition of `eval` for lack of space. A closure is either a function body paired with its environment, or a partially applied constant. `clos2trm` converts back a closure to an equivalent term, `trm_inst` instantiating all de Bruijn indices at once with a list of terms, and `const_app` building the curried application of `c` to a list of terms. Since evaluation may not terminate, `eval` takes as argument the number `h` of reduction steps to compute. The remaining arguments are the environment `benv`, accessed through de Bruijn indices, the application stack `app` which contains the arguments to

the term being evaluated, the term e itself, which provides an efficient representation of code thanks to de Bruijn indices, and the control stack $stack$. Here the nameless representation of terms was handy, as it maps naturally to a stack machine. The result of `eval` is either a closure, with the number of evaluation steps remaining, or the current state of the machine.

We also proved completeness with respect to the rewriting rules, *i.e.* if the rewriting based evaluation reaches a normal form, then evaluation by the abstract machine terminates with the same normal form. This required building a bisimulation between the two evaluations, and was trickier than expected. Namely we need to prove the following lemma:

Definition `inst t benv := trm_inst t (map clos2trm benv)`.

Lemma `complete_rec : ∀ args args' fl fl' e e' benv benv' τ,`
`args ≡ args' → fl ≡ fl' → (inst e benv → inst e' benv') →`
`K; Γ ⊢ stack2trm (app2trm (inst e benv) args) fl : τ →`
`∃ h, ∃ h', eval h benv args e fl ≡ eval h' benv' args' e' fl'.`

where \equiv denotes the equality of closures after substitution by their environment, *i.e.* `clos_abs e benv ≡ clos_abs e' benv'` iff `inst (λe) benv = inst (λe') benv'`. Proving this by case analysis on e and e' ended up being very time consuming. The proofs being rather repetitive, they may profit from better lemmas.

6 Dependent Types

As we pointed in section 4, the statements of many lemmas and theorems include lots of well-formedness properties, which are expected to be true of any value of a given type. For instance, substitutions should be idempotent, environments should not bind the same variable twice, de Bruijn indices should not escape, kinds should be valid, *etc.* . . . A natural impulse is to use dependent types to encode these properties. Yet proofs from [1] only use dependent types for the generation of fresh variables. The reason is simple enough: as soon as a value is defined as a dependent sum, using rewriting on it becomes much more cumbersome. I attempted using it for the well-formedness of polytypes, but had to abandon the idea because there were too many things to prove upfront. On the other hand, using dependent types to make sure that kinds are valid and coherent was not so hard, and helped to streamline the proofs. This is probably due to the abstract nature of constraint domains, which limits interactions between kinds and other features. The definition of kinds becomes:

Definition `coherent kc kr := ∀ x (τ τ' : typ),`

`Cstr.unique kc x = true → ln (x, τ) kr → ln (x, τ') kr → τ = τ'.`

Record `ckind : Set := Kind{`

`kind_cstr : Cstr.cstr; kind_valid : Cstr.valid kind_cstr;`

`kind_rel : list (Cstr.attr × typ); kind_coherent : coherent kind_cstr kind_rel}.`

We still need to apply substitutions to kinds, but this is not a problem as substitutions do not change the constraint, and preserve the coherence. We just need the following function.

Definition `ckind_map_spec` : $\forall (f : \text{typ} \rightarrow \text{typ})(k : \text{ckind}),$
 $\{k' : \text{ckind} \mid \text{kind_cstr } k = \text{kind_cstr } k' \wedge \text{kind_rel } k' = \text{map_snd } f (\text{kind_rel } k)\}.$

We also sometimes have to prove the equality of two kinds obtained independently. This requires the following lemma, which can be proved using proof irrelevance⁴.

Lemma `ckind_pi` : $\forall k k' : \text{ckind},$
 $\text{kind_cstr } k = \text{kind_cstr } k' \rightarrow \text{kind_rel } k = \text{kind_rel } k' \rightarrow k = k'.$

Another application of dependent types is ensuring termination for the unification and type inference algorithms. In Coq all functions must be complete. Originally, this was ensured by adding a step counter, and proving separately that one can choose a number of steps sufficient to obtain a result. This is the style used in section 4.1. This approach is simple, but this extra parameter stays in the extracted code. In a first version of the proof, the parameter was so big that the unification algorithm would just take forever trying to compute the number of steps it needed. I later came up with a smaller value, but it would be better to have it disappear completely. This is supported in Coq through well-founded recursion. In practice this works by moving the extra parameter to the universe of proofs (Prop), so that it will disappear during extraction. The Function command automates this, but there is a pitfall: while it generates dependent types, it doesn't support them in its input. The termination argument for unification being rather complex, this limitation proved problematic. Attempts with Program Fixpoint didn't succeed either. Finally I built the dependently typed function by hand. While this requires a rather intensive use of dependent types, the basic principle is straightforward, and it makes the proof of completeness simpler. As a result the overall size of the proof for unification didn't change. However, since the type inference algorithm calls unification, it had to be modified too, and its size grew by about 10%. An advantage of building our functions by hand is that we control exactly the term produced; since rewriting on dependently typed terms is particularly fragile, this full control proves useful.

7 Program Extraction

Both the type checker and interpreter can be extracted to Objective Caml code. This lets us build a fully certified⁵ implementation for a fragment of Objective Caml's type system. Note that there is no parser or read-eval-print loop yet, making it just a one-shot interpreter for programs written directly in abstract syntax. Moreover, since Coq requires all programs to terminate, one has to indicate the number of steps to be evaluated explicitly. Well-founded recursion cannot be used here, as our language is Turing-complete. (Actually, Objective Caml allows one to define cyclic constants, so that we can build a value representing infinity, and remove the need for an explicit number of steps. However, this is going around the soundness of Coq.)

Here is an example of program written in abstract syntax (with a few abbreviations), and its inferred type (using lots of pretty printing).

⁴ Since both validity and coherence are decidable, proof irrelevance could be avoided here by slightly changing definitions.

⁵ The validity of our certification relies on the correctness of Coq and Objective Caml, which are rather strong assumptions.

```

# let rev_append =
  recf (abs (abs (abs
    (matches [0;1]
      [abs (bvar 1);
        abs (apps (bvar 3) [sub 1 (bvar 0); cons (sub 0 (bvar 0)) (bvar 1)]));
      bvar 1]))) ;
val rev_append : trm = ...
# typinf2 Nil rev_append;;
- : (var * kind) list * typ =
  [(10, <Ksum, {}, {0; 1}, {0 => tv 15; 1 => tv 34}>);
   (29, <Ksum, {1}, any, {1 => tv 26}>);
   (34, <Kprod, {1; 0}, any, {0 => tv 30; 1 => tv 10}>);
   (30, any);
   (26, <Kprod, {}, {0; 1}, {0 => tv 30; 1 => tv 29}>);
   (15, any)],
  tv 10 @> tv 29 @> tv 29)

```

Here `recf` is an extra constant which implements the fixpoint operator. Our encoding of lists uses 0 and 1 as labels for both variants and records, but we could have used any other natural numbers: their meaning is not positional, but associative. Since de Bruijn indices can be rather confusing, here is a version translated to a syntax closer to Objective Caml, with meaningful variable names and labels.

```

let rec rev_append l1 l2 =
  match l1 with
  | 'Nil _ -> l2
  | 'Cons c ->
    rev_append c.tl ('Cons {hd=c.hd; tl=l2})
val rev_append :
  ([< 'Nil of '15 | 'Cons of {hd:'30; tl:'10; ..}] as '10) ->
  ([> 'Cons of {hd:'30; tl:'29}] as '29) -> '29

```

8 Related Works

The mechanization of type safety proofs for programming languages has been extensively studied. Existing works include Core ML using Coq [5], Java using Isabelle/HOL [14], and more recently full specification of OCaml light using HOL-4 [15] and Standard ML using Twelf [11,4]. The main difference in our system is the presence of structural polymorphism and recursion. In particular, among the above works, only [11] handles inclusion problems for iso-recursive types (in a simpler setting than ours, since when checking signature subtyping no structural polymorphism is allowed). It is also the work closest to our goal of handling advanced type features (it already handles fully Standard ML). OCaml-light rather focuses on subtle points in the dynamic semantics of the language. Typed Scheme [18] has a type system remarkably similar to ours, and part of the soundness proof was mechanized in Isabelle/HOL, but the mechanized part does not contain recursive types.

Concerning unification and type inference, we have already mentioned the works of Paulson in LCF [16], Dubois and Ménéssier-Morain in Coq [6], and Naraschewski and Nipkow in Isabelle [13], and the more recent Isabelle/Nominal proof by Urban and Nipkow [19]. The main difference is the introduction of structural polymorphism,

Table 1. Components of the proof

File	Lines	Contents
Lib_*	1706	Auxiliary lemmas and tactics from [1]
Metatheory	1376	Metatheory lemmas and tactics from [1]
Metatheory_SP	1304	Additional lemmas and tactics
Definitions	458	Definition of the type system
Infrastructure	1152	Common lemmas
Soundness	633	Soundness proof
Rename	985	Renaming and inversion lemmas
Eval	2935	Stack-based evaluation
Unify	1832	Unification
Inference	3159	Type inference
Domain	1085	Constraint domain specific proofs
Unify_wf	1827	Unification using dependent measure
Inference_wf	3443	Inference using dependent measure

which results in much extended statements to handle admissible substitutions. Even in the absence of structural polymorphism, just handling equi-recursive types makes type inference more complex, and we are aware of no proof of principality including them. It might be interesting to compare these different proofs of W in more detail, as the first two use de Bruijn indices [6,13], the latter nominal datatypes [19], and ours cofinite quantification. However, as Urban and Nipkow already observed, while there are clear differences between the different approaches, in the case of type inference lots of low-level handling of type variables has to be done, and as a result clever encodings do not seem to be that helpful.

More generally, all the literature concerning the PoplMark challenge [2] can be seen as relevant here, at least for the type soundness part. In particular, one could argue that structural polymorphism being related to structural subtyping, challenges 1B and 2B (transitivity of subtyping with records, and type safety with records and pattern matching) should be relevant. However, in the case of structural polymorphism, the presence of recursive types requires the use of a graph structure to represent types, which does not seem to be necessary for those challenges, where trees are sufficient. We believe that this changes the complexity of the proof.

9 Conclusion

We have reached our first goal, providing a fully certified type checker and interpreter. We show the size and contents of the various components of the proof in table 1. While this is a good start, it currently handles only a very small subset of Objective Caml. The next goal is of course to add new features. A natural next target would be the addition of side-effects, with the relaxed value restriction. Note that since the value restriction relies on subtyping, it would be natural to also add type constructors, with variance annotations, at this point. Considering the difficulties we have met up to now, we do not expect it to be an easy task.

Acknowledgments. I wish to thank the anonymous reviewers for their detailed and helpful comments.

References

1. Aydemir, B., Charguéraud, A., Pierce, B.C., Pollack, R., Weirich, S.: Engineering formal metatheory. In: Proc. ACM Symposium on Principles of Programming Languages, pp. 3–15 (2008)
2. Aydemir, B.E., Bohannon, A., Fairbairn, M., Foster, J.N., Pierce, B.C., Sewell, P., Vytiniotis, D., Washburn, G., Weirich, S., Zdancewic, S.: Mechanized metatheory for the masses: The PoplMark challenge. In: Hurd, J., Melham, T. (eds.) TPHOLs 2005. LNCS, vol. 3603, pp. 50–65. Springer, Heidelberg (2005)
3. Barras, B.: Auto-validation d'un système de preuves avec familles inductives. Thèse de doctorat, Université Paris 7 (November 1999)
4. Crary, K., Harper, B.: Mechanized definition of Standard ML alpha release. Twelf proof scripts (August 2009)
5. Dubois, C.: Proving ML type soundness within Coq. In: Aagaard, M.D., Harrison, J. (eds.) TPHOLs 2000. LNCS, vol. 1869, pp. 126–144. Springer, Heidelberg (2000)
6. Dubois, C., Ménéssier-Morain, V.: Certification of a type inference tool for ML: Damas-Milner within Coq. *Journal of Automated Reasoning* 23(3), 319–346 (1999)
7. Furuse, J.P., Garrigue, J.: A label-selective lambda-calculus with optional arguments and its compilation method. RIMS Preprint 1041, Research Institute for Mathematical Sciences, Kyoto University (October 1995)
8. Garrigue, J.: Simple type inference for structural polymorphism. In: The Ninth International Workshop on Foundations of Object-Oriented Languages, Portland, Oregon (2002)
9. Garrigue, J.: Relaxing the value restriction. In: Kameyama, Y., Stuckey, P.J. (eds.) FLOPS 2004. LNCS, vol. 2998, Springer, Heidelberg (2004)
10. Garrigue, J., Rémy, D.: Extending ML with semi-explicit higher order polymorphism. *Information and Computation* 155, 134–171 (1999)
11. Lee, D.K., Crary, K., Harper, R.: Towards a mechanized metatheory of standard ML. In: Proc. ACM Symposium on Principles of Programming Languages, pp. 173–184 (January 2007)
12. Leroy, X., Doligez, D., Garrigue, J., Rémy, D., Vouillon, J.: The Objective Caml system release 3.11, Documentation and user's manual. Projet Gallium, INRIA (November 2008)
13. Naraschewski, W., Nipkow, T.: Type inference verified: Algorithm W in Isabelle/HOL. *Journal of Automated Reasoning* 23, 299–318 (1999)
14. Oheimb, D.v., Nipkow, T.: Machine-checking the Java specification: Proving type-safety. In: Alves-Foss, J. (ed.) Formal Syntax and Semantics of Java. LNCS, vol. 1523, pp. 119–156. Springer, Heidelberg (1999)
15. Owens, S.: A sound semantics for OCaml light. In: Drossopoulou, S. (ed.) ESOP 2008. LNCS, vol. 4960, pp. 1–15. Springer, Heidelberg (2008)
16. Paulson, L.: Verifying the unification algorithm in LCF. *Science of Computer Programming* 5, 143–169 (1985)
17. The Coq Team. The Coq Proof Assistant, Version 8.2. INRIA (2009)
18. Tobin-Hochstadt, S., Felleisen, M.: The design and implementation of typed scheme. In: Proc. ACM Symposium on Principles of Programming Languages (2008)
19. Urban, C., Nipkow, T.: Nominal verification of algorithm W. In: Huet, G., Lévy, J.-J., Plotkin, G. (eds.) From Semantics to Computer Science. Essays in Honour of Gilles Kahn, pp. 363–382. Cambridge University Press, Cambridge (2009)