ML-Like Inference for Classifiers

Cristiano Calcagno[1], Eugenio Moggi[2*], and Walid Taha[3**]

[1] Imperial College, London, UK ccris@doc.ic.ac.uk
[2] DISI, Univ. of Genova, Italy moggi@disi.unige.it
[3] Rice University, TX, USA taha@cs.rice.edu

Abstract. Environment classifiers were proposed as a new approach to typing multi-stage languages. Safety was established in the simply-typed and let-polymorphic settings. While the motivation for classifiers was the feasibility of inference, this was in fact not established. This paper starts with the observation that inference for the full classifier-based system fails. We then identify a subset of the original system for which inference is possible. This subset, which uses *implicit classifiers*, retains significant expressivity (e.g. it can embed the calculi of Davies and Pfenning) and eliminates the need for classifier names in terms. Implicit classifiers were implemented in MetaOCaml, and no changes were needed to make an existing test suite acceptable by the new type checker.

1 Introduction

Introducing explicit staging constructs into programming languages is the goal of research projects including 'C [9], Popcorn [25], MetaML [30,20], MetaOCaml [4, 19], and Template Haskell [23]. Staging is an essential ingredient of macros [10], partial evaluation [15], program generation [16], and run-time code generation [12]. In the untyped setting, the behavior of staging constructs resembles the quasi-quotation mechanisms of LISP and Scheme [2]. But in the statically-typed setting, such quotation mechanisms may prohibit static type-checking of the quoted expression. Some language designs, such as that of 'C, consider this acceptable. In Template Haskell, this is considered a feature; namely, a form of staged type inference [24]. But in the design of MetaML and MetaOCaml, it is seen as a departure from the commitment of ML and OCaml to static prevention of runtime errors.[1]

Multi-stage Basics. The use of staging constructs can be illustrated in a multi-stage language such as MetaOCaml [19] with a classic example[2]:

[1] Dynamic typing can be introduced as orthogonal and non-pervasive feature [1].
[2] Dots are used around brackets and escapes to disambiguate the syntax in the implementation, but they are dropped when we consider the underlying calculus.

```
let rec power n  x = (* : int -> int code -> int code *)
  if n=0 then .<1>. else .<.~x * .~(power (n-1) x)>.
let power72 : int -> int = .! .<fun x -> .~(power 72 .<x>.)>.
```

Ignoring the type constructor t `code` and the three staging annotations *brackets* `.<e>.`, *escapes* `.~e`, and *run* `.!`, the above code is a standard definition of a function that computes x^n, which is then used to define the specialized function x^{72}. Without staging, however, the last step just produces a closure that invokes the power function every time it gets a value for x. To understand the effect of the staging annotations, it is best to start from the end of the example. Whereas a term `fun x -> ` e `x` is a value, an annotated term `.<fun x -> .~(` e `.<x>.)>.` is not. Brackets indicate that we are constructing a future stage computation, and an escape indicates that we must perform an immediate computation *while* building the enclosing bracketed computation. The application e `.<x>.` has to be performed first, even though `x` is still an uninstantiated *symbol*. In the `power` example, `power 72 .<x>.` is performed immediately, once and for all, and not repeated every time we have a new value for `x`. In the body of the definition of the `power` function, the recursive application of `power` is escaped to make sure that it is performed immediately. The run construct (`.!`) on the last line invokes the compiler on the generated code fragment, and incorporates the result of compilation into the runtime system.

Background. Starting with the earliest statically typed languages supporting staging (including those of Gomard and Jones [11] and Nielson and Nielson [22]), most proposals to date fall under two distinct approaches: one treating code as always open, the other treating code as always closed. The two approaches are best exemplified by two type systems corresponding to well-known logics:

$\lambda^\bigcirc$ Motivated by the next modality $\bigcirc$ of linear time temporal logic, this system provides a sound framework for typing constructs that have the same operational semantics as bracket and escape [6]. As illustrated above, brackets and escapes can be used to annotate λ-abstractions so as to force evaluation under lambda. This type system supports code generation but does not provide a construct for code execution.

$\lambda^\square$ Motivated by the necessity modality $\square$ of S4 modal logic, this system provides constructs for generating and executing closed code [7]. The exact correspondence between the constructs of $\lambda^\square$ and LISP-style quotation mechanism is less immediate than for $\lambda^\bigcirc$.

Combining the two approaches to realize a language that allows evaluation under lambda *and* a run construct is challenging [26]. In particular, evaluation under lambda gives rise to code fragments that contain free variables that are not yet "linked" to any fixed value. Running such open code fragments can produce a runtime error. Several type systems [3,29,21] have been proposed for safely combining the key features of $\lambda^\bigcirc$ (the ability to manipulate open code) and $\lambda^\square$ (the ability to execute closed code). But a practical solution to the problem requires meeting a number of demanding criteria simultaneously:

- **Safety:** the extension should retain static safety;
- **Conservativity:** the extension should not affect programs that do not use multi-stage facilities;
- **Inference:** the extension should support type inference;
- **Light annotations:** the extension should minimize the amount of programmer annotations required to make type inference possible.

All the above proposals were primarily concerned with the safety criterion, and were rarely able to address the others. Because previous proposals seemed notationally heavy, implementations of multi-stage languages (such MetaML and MetaOCaml) often chose to sacrifice safety. For example, in MetaOCaml .! e raises an exception, when the evaluation of e produces open code.

The type system for environment classifiers λ^α of [29] appears to be the most promising starting point towards fulfilling all criteria. The key feature of λ^α is providing a code type $\langle\tau\rangle^\alpha$ decorated with a classifier α that constrains the unresolved variables that may occur free in code. Intuitively, in the type system of λ^α, variables are declared at levels annotated by classifiers, and code of type $\langle\tau\rangle^\alpha$ may contain only unresolved variables declared at a level annotated with α. Classifiers are also used explicitly in terms. Type safety for λ^α was established [29], but type inference was only conjectured.

Contributions and Organization of this Paper. The starting point for this work is the observation that inference for full λ^α fails. To address this problem, a subset of the original system is identified for which inference is not only possible but is in fact *easy*. This subset uses *implicit classifiers*, thus eliminates the need for classifier names in terms, and retains significant expressivity (e.g. it embeds the paradigmatic calculi $\lambda^\bigcirc$ and $\lambda^\square$). Implicit classifiers have been implemented in MetaOCaml, and no changes were needed to make an existing test suite acceptable by the type checker. The paper proceeds as follows:

- Section 2 extends a core subset of ML with environment classifiers. The new calculus, called λ^i_{let}, corresponds to a proper subset of λ^α but eliminates classifier names in terms. This is an improvement on λ^α in making annotations lighter, and the proof of type safety for λ^α adapts easily to λ^i_{let}.
- Section 3 gives two inference algorithms:
    1. a principal typing algorithm for λ^i, the simply-typed subset of λ^i_{let} (i.e. no type schema and let-binding), which extends Hindley's principal typing algorithm for the λ-calculus.
    2. a principal type algorithm for λ^i_{let}, which extends Damas and Milner's algorithm W.

    Therefore classifiers are a natural extension to well-established type systems.
- Section 4 relates λ^i to λ^α and exhibits some terms typable in λ^α that fail to have a principal type (thus, λ^α fails to meet the inference criterion). It also shows that λ^i retains significant expressivity, namely there are *typability-preserving* embeddings of $\lambda^\bigcirc$ and a variant of $\lambda^\square$ into λ^i (similar to the embeddings into λ^α given in [29]). However, if one restricts λ^i_{let} further, by considering a runClosed construct similar to Haskell's runST [17], then the embedding of $\lambda^\square$ is lost (but term annotations disappear completely).

Variables	$x \in \mathsf{X}$
Classifiers	$\alpha \in \mathsf{A}$
Named Levels	$A \in \mathsf{A}^*$
Terms	$e, v \in \mathsf{E} ::= x \mid \lambda x.e \mid e\,e \mid \langle e \rangle \mid {\tilde{}}e \mid$
	$\mathsf{run}\ e \mid \%e \mid \mathsf{open}\ e \mid \mathsf{close}\ e \mid \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2$
Type Variables	$\beta \in \mathsf{B}$
Types	$\tau \in \mathsf{T} ::= \beta \mid \tau_1 \to \tau_2 \mid \langle \tau \rangle^\alpha \mid \langle \tau \rangle$
Type Schema	$\sigma \in \mathsf{S} ::= \tau \mid \forall \alpha.\sigma \mid \forall \beta.\sigma$ or equivalently
	$\forall \bar{\kappa}.\tau$ with $\bar{\kappa}$ sequence of distinct α and β
Assignments	$\Gamma \in \mathsf{X} \overset{fin}{\mapsto} (\mathsf{T} \times \mathsf{A}^*)$ of types and named levels
Assignments	$\Delta \in \mathsf{X} \overset{fin}{\mapsto} (\mathsf{S} \times \mathsf{A}^*)$ of type schema and named levels

Fig. 1. Syntax of λ_{let}^i

Notation. Throughout the paper we use the following notation and conventions:

- We write m to range over the set N of natural numbers. Furthermore, $m \in \mathsf{N}$ is identified with the set of its predecessors $\{i \in \mathsf{N} | i < m\}$.
- We write $\bar{a}$ to range over the set A^* of finite sequences $(a_i | i \in m)$ with $a_i \in \mathsf{A}$, and $|\bar{a}|$ denotes its length m. We write $\emptyset$ for the empty sequence and $\bar{a}_1, \bar{a}_2$ for the concatenation of $\bar{a}_1$ and $\bar{a}_2$.
- We write $f : A \overset{fin}{\mapsto} B$ to say that f is a partial function from A to B with a finite domain, written $\mathrm{dom}(f)$. We write $A \to B$ to denote the set of total functions from A to B. We use the following operations on partial functions:
  - $\{a_i : b_i | i \in m\}$ is the partial function mapping a_i to b_i (where the a_i are distinct, i.e. $a_i = a_j$ implies $i = j$); in particular, $\emptyset$ is the everywhere undefined partial function;
  - $f \backslash a$ denotes the partial function g s.t. $g(a') = b$ iff $f(a') = b$ when $a' \neq a$, and undefined otherwise;
  - $f\{a : b\}$ denotes the (possibly) partial function g s.t. $g(a) = b$ and $g(a') = f(a')$ when $a' \neq a$;
  - f, g denotes the union of two partial functions with disjoint domains.
  - $f = g \bmod X$ means that $\forall x \in X.f(x) = g(x)$.
- We write $X \# X'$ to mean that X and X' are disjoint sets, and $X \uplus X'$ for their disjoint union.

2 λ_{let}^i: A Calculus with Implicit Classifiers

This section defines λ_{let}^i, an extension of the functional subset of ML with environment classifiers. In comparison to λ^α [29], classifier names do not appear in terms. In particular, the constructs of λ^α for explicit abstraction $(\alpha)e$ and instantiation $e[\alpha]$ of classifiers are replaced by the constructs close e and open e, but with more restrictive typing rules. As we show in this paper, this makes it possible to support ML-style inference of types and classifiers in a straightforward manner.

Figure 1 gives the syntax of λ_{let}^i. Intuitively, *classifiers* allow us to name parts of the environment in which a term is typed. Classifiers are described as *implicit* in λ_{let}^i because they do not appear in terms. *Named levels* are sequences of environment classifiers. They are used to keep track of the environments used as we build nested code. Named levels are thus an enrichment of the traditional notion of levels in multi-stage languages [6,30,28], the latter being a natural number which keeps track only of the depth of nesting of brackets. Terms include:

- the standard λ-terms, i.e. variables, λ-abstraction and application;
- the staging constructs of MetaML [30], i.e. Brackets $\langle e \rangle$, escape ~e, and run run e, and a construct %e for cross-stage persistence (CSP) [3,29];
- the constructs close e and open e are the implicit versions of the λ^α constructs for classifiers abstraction $(\alpha)e$ and instantiation $e[\alpha]$ respectively;
- the standard let-binding for supporting Hindley-Milner polymorphism.

Types include type variables, functional types, and code types $\langle \tau \rangle^\alpha$ annotated with a classifier (exactly as in λ^α of [29], thus refining open code types). The last type $\langle \tau \rangle$ is for executable code (it is used for typing run e) and basically corresponds to the type $(\alpha)\langle \tau \rangle^\alpha$ of λ^α (as explained in more detail in Section 4.1).

As in other Hindley-Milner type systems, type schema restrict quantification at the outermost level of types. Since the types of λ_{let}^i may contain not only type variables but also classifiers, type schema allow quantification of both.

Notation. The remainder of the paper makes use of the following definitions:

- We write $\{x_i : \tau_i^{A_i} | i \in m\}$ and $\{x_i : \sigma_i^{A_i} | i \in m\}$ to describe type and schema assignments, and in this context τ^A denotes the pair (τ, A).
- FV(_) denotes the set of variables free in _. In λ_{let}^i, there are three kinds of variables: term variables x, classifiers α, and type variables β. The definition of FV(_) for terms, types, and type schema is standard, and it extends in the obvious way to A, Γ, and Δ, e.g. $\text{FV}(\Delta) = \cup\{\text{FV}(\sigma) \cup \text{FV}(A) \mid \Delta(x) = \sigma^A\}$.
- We write $\equiv$ for equivalence up to α-conversion on terms, types, and type schema. _$[x : e]$ denotes substitution of x with e in _ modulo $\equiv$, i.e. the bound variables in _ are automatically renamed to avoid clashes with FV(e). Similarly we write _$[\alpha : \alpha']$ and _$[\beta : \tau]$ for classifiers and type variables.
- We write $\rho \in \mathsf{Sub}$ for the set of *substitutions*, i.e. functions (with domain $\mathsf{A} \cup \mathsf{B}$) mapping classifiers α to classifiers and type variables β to types τ, and having a finite *support* defined as $\{\alpha | \rho(\alpha) \neq \alpha\} \cup \{\beta | \rho(\beta) \neq \beta\}$. Then _$[\rho]$ denotes *parallel substitution*, where each free occurrence in _ of a classifier α and type variable β is replaced by its ρ-image. With some abuse of notation, we write $e[\rho]$ also when ρ is a partial function with finite domain, by extending it as the identity outside dom(ρ).
- $\sigma \succ \tau$ means that type τ is an *instance* of σ, i.e. $\forall \overline{\kappa}.\tau \succ \tau' \overset{\text{def}}{\Longleftrightarrow} \tau[\rho] = \tau'$ for some $\rho \in \mathsf{Sub}$ with support $\overline{\kappa}$ (the order in $\overline{\kappa}$ is irrelevant).
  $\succ$ extends to a relation on type schemas and type schema assignments:

$$\text{var}\ \dfrac{\sigma \succ \tau \quad \Delta(x) = \sigma^A}{\Delta \overset{A}{\vdash} x : \tau} \qquad \text{lam}\ \dfrac{\Delta\{x : \tau_1^A\} \overset{A}{\vdash} e : \tau_2}{\Delta \overset{A}{\vdash} \lambda x.e : \tau_1 \to \tau_2} \qquad \text{app}\ \dfrac{\Delta \overset{A}{\vdash} e_1 : \tau_1 \to \tau_2 \quad \Delta \overset{A}{\vdash} e_2 : \tau_1}{\Delta \overset{A}{\vdash} e_1 e_2 : \tau_2}$$

$$\text{brck}\ \dfrac{\Delta \overset{A,\alpha}{\vdash} e : \tau}{\Delta \overset{A}{\vdash} \langle e \rangle : \langle \tau \rangle^\alpha} \qquad \text{esc}\ \dfrac{\Delta \overset{A}{\vdash} e : \langle \tau \rangle^\alpha}{\Delta \overset{A,\alpha}{\vdash} {\sim}e : \tau} \qquad \text{run}\ \dfrac{\Delta \overset{A}{\vdash} e : \langle \tau \rangle}{\Delta \overset{A}{\vdash} \mathsf{run}\ e : \tau} \qquad \text{csp}\ \dfrac{\Delta \overset{A}{\vdash} e : \tau}{\Delta \overset{A,\alpha}{\vdash} \%e : \tau}$$

$$\text{open}\ \dfrac{\Delta \overset{A}{\vdash} e : \langle \tau \rangle}{\Delta \overset{A}{\vdash} \mathsf{open}\ e : \langle \tau \rangle^\alpha} \qquad \text{close}\ \dfrac{\Delta \overset{A}{\vdash} e : \langle \tau \rangle^\alpha}{\Delta \overset{A}{\vdash} \mathsf{close}\ e : \langle \tau \rangle} \quad \alpha \notin \mathrm{FV}(\Delta, A, \tau)$$

$$\text{let}\ \dfrac{\Delta \overset{A}{\vdash} e_1 : \tau_1 \quad \Delta\{x : (\forall\overline{\kappa}.\tau_1)^A\} \overset{A}{\vdash} e_2 : \tau_2}{\Delta \overset{A}{\vdash} \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2 : \tau_2} \quad \mathrm{FV}(\Delta, A)\#\overline{\kappa}$$

Fig. 2. Type System for λ_{let}^i

- $\forall\overline{\kappa}.\tau \succ \forall\overline{\kappa}'.\tau' \overset{\text{def}}{\iff} \forall\overline{\kappa}.\tau \succ \tau'$, where we assume $\overline{\kappa}'\#\mathrm{FV}(\forall\overline{\kappa}.\tau)$ by α-conversion
- $\Delta_1 \succ \Delta_2 \overset{\text{def}}{\iff} \mathrm{dom}(\Delta_1) = \mathrm{dom}(\Delta_2)$ and, for all x, whenever $\Delta_1(x) = \sigma_1^{A_1}$ and $\Delta_2(x) = \sigma_2^{A_2}$ then it is also the case that $\sigma_1 \succ \sigma_2$ and $A_1 = A_2$.

2.1   Type System

Figure 2 gives the type system for λ_{let}^i. The first three rules are mostly standard. As in ML, a polymorphic variable x whose type schema is σ can be assigned any type which is an instance of σ. In these constructs the named level A is propagated without alteration to the sub-terms. In the variable rule, the named level associated with the variable being typed-checked is required to be the *same* as the current level. In the lambda abstraction rule, the named level of the abstraction is recorded in the environment.

The rule for brackets is almost the same as in previous type systems. First, for every code type a classifier must be assigned. Second, while typing the body of the code fragment inside brackets, the named level of the typing judgment is extended by the name of the "current" classifier. This information is used in both the variable and the escape rules to make sure that only variables and code fragments with the same classification are ever incorporated into this code fragment. The escape rule at named level A, α only allows the incorporation of code fragments of type $\langle \tau \rangle^\alpha$. The rule for CSP itself is standard: It allows us to incorporate a term e at a "higher" level. The rule for run allows to execute a code fragment that has type $\langle \tau \rangle$.

The *close* and *open* rules are introduction and elimination for the runnable code type $\langle \tau \rangle$ respectively. One rule says that close e is runnable code when e can be classified with *any* α, conversely the other rule says that code open e can be classified by any α provided e is runnable code. The rule for let is standard and allows the introduction of variables of polymorphic type.

The following proposition summarizes the key properties of the type system relevant for type safety as well as type inference.

Proposition 1 (Basic Properties). *The following rules are admissible:*

$$- \;\alpha\tau\text{-subst}\quad \frac{\Delta \overset{A}{\vdash} e : \tau}{(\Delta \overset{A}{\vdash} e : \tau)[\rho]}\;\; \rho \in \mathsf{Sub} \qquad \Delta\text{-sub}\quad \frac{\Delta_2 \overset{A}{\vdash} e : \tau}{\Delta_1 \overset{A}{\vdash} e : \tau}\;\; \Delta_1 \succ \Delta_2$$

$$- \;\text{strength}\quad \frac{\Delta \overset{A}{\vdash} e : \tau}{\Delta \backslash x \overset{A}{\vdash} e : \tau}\;\; x \notin \mathrm{FV}(e) \qquad \text{weaken}\quad \frac{\Delta \overset{A}{\vdash} e : \tau}{\Delta, x : \sigma_1^{A_1} \overset{A}{\vdash} e : \tau}\;\; x \notin dom(\Delta)$$

$$- \;\text{e-subst}\quad \frac{\Delta \overset{A_1}{\vdash} e_1 : \tau_1 \qquad \Delta, x : (\forall\overline{\kappa}.\tau_1)^{A_1} \overset{A_2}{\vdash} e_2 : \tau_2}{\Delta \overset{A_2}{\vdash} e_2[x : e_1] : \tau_2}\;\; \overline{\kappa}\#\mathrm{FV}(\Delta, A_1)$$

3 Inference Algorithms

This section describes two inference algorithms. The first algorithm extends Hindley's principal typing algorithm [13] for the simply typed λ-calculus with type variables to λ^i (the simply-typed subset of λ^i_{let}, i.e. without type schema and let-binding). Existence of principal typings is very desirable but hard to get (see [14,31]). Thus it is reassuring that it is retained after the addition of classifiers. The second algorithm extends Damas and Milner's algorithm W [5] to λ^i_{let} and proves that it is sound and complete for deriving principal types. Damas and Milner's algorithm is at the core of type inference for languages such as ML, OCaml, and Haskell. That this extension is possible (and easy) is of paramount importance to the practical use of the proposed type system.

Both algorithms make essential use of a function $mgu(T)$ computing a most general unifier $\rho \in \mathsf{Sub}$ for a finite set T of equations between types or between classifiers. For convenience, we also introduce the following derived notation for sets of equations (used in side-conditions to the rules describing the algorithms):

- (A_1, A_2) denotes $\{(\alpha_{1,j}, \alpha_{2,j}) \mid j \in n\}$ when $n = |A_1| = |A_2|$ and $\alpha_{i,j}$ is the j-th element of A_i, and is undefined when $|A_1| \neq |A_2|$
- (Γ_1, Γ_2) denotes $\cup\{(\tau_{1,x}, \tau_{2,x}), (A_{1,x}, A_{2,x}) | x \in dom(\Gamma_1) \cap dom(\Gamma_2)\}$ where $\Gamma_i(x) = \tau_{i,x}^{A_{i,x}}$.

3.1 Principal Typing

We extend Hindley's principal typing algorithm for the simply typed λ-calculus with type variables to λ^i. Wells [31] gives a general definition of principal typing and related notions, but we need to adapt his definition of Hindley's principal typing to our setting, mainly to take into account levels.

Definition 1 (Typing). *A triple (Γ, τ, A) is a typing of $(e, n) \overset{def}{\Longleftrightarrow} \Gamma \overset{A}{\vdash} e : \tau$ is derivable and $n = |A|$. A Hindley principal typing of (e, n) is a typing (Γ, τ, A) of (e, n) s.t. for every other typing (Γ', τ', A') of (e, n)*

- $\Gamma[\rho] \subseteq \Gamma'$ *and* $\tau' = \tau[\rho]$ *and* $A' = A[\rho]$ *for some* $\rho \in \mathsf{Sub}$.

Remark 1. Usually one assigns typings to terms. We have chosen to assign typings to a pair (e, n), because the operational semantics of a term is level-dependent. However, one can easily assign typings to terms (and retain the existence of principal typings). First, we introduce an infinite set of variables $\phi \in \Phi$ ranging over annotated levels. Then, we modify the BNF for annotated levels to become $A ::= \phi \mid A, \alpha$. Unification will also have to deal with equations for annotated levels, e.g. $\phi_1 = \phi_2, \alpha$. A posteriori, one can show that a principal typing will contain exactly one variable ϕ.

Figure 3 defines the algorithm by giving a set of rules (directed by the structure of e) for deriving judgments of the form $\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau, A)$. $\mathcal{K} \subseteq_{fin} \mathsf{A} \uplus \mathsf{B}$ is an auxiliary parameter (instrumental to the algorithm), which is threaded in recursive calls for recording the classifiers and type variables used so far. The algorithm computes a typing (and updates $\mathcal{K}$) or fails, and it enjoys the following properties (which imply that every (e, n) with a typing has a principal typing).

Theorem 1 (Soundness). *If $\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma', \tau', A')$, then $\Gamma' \overset{A'}{\vdash} e : \tau'$ and $n = |A'|$, moreover $dom(\Gamma') = \mathrm{FV}(e)$, $\mathcal{K} \subseteq \mathcal{K}'$ and $\mathrm{FV}(\Gamma', \tau', A') \subseteq \mathcal{K}' \setminus \mathcal{K}$.*

Theorem 2 (Completeness). *If $\Gamma' \overset{A'}{\vdash} e : \tau'$, then $\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau, A)$ is derivable (for any choice of $\mathcal{K}$) and exists $\rho' \in \mathsf{Sub}$ s.t. $\Gamma[\rho'] \subseteq \Gamma'$, $\tau' = \tau[\rho']$ and $A' = A[\rho']$.*

Moreover, from general properties of the most general unifier and the similarity of our principal typing algorithm with that for the λ-calculus, one can also show

Theorem 3 (Conservative Extension). *If $e ::= x \mid \lambda x.e \mid e\ e$ is a λ-term, then (Γ, τ) is a principal typing of e in $\lambda \iff (\Gamma, \tau, \emptyset)$ is a principal typing of $(e, 0)$ in λ^i, where we identify $x : \tau$ with $x : \tau^\emptyset$.*

3.2   Principal Type Inference

In this section, we extend Damas and Milner's [5] principal type algorithm to λ^i_{let} and prove that it is sound and complete. Also in this case we have to adapt to our setting the definition of Damas-Milner principal type in [31].

Definition 2 (Principal Type). *A Damas-Milner principal type of (Δ, A, e) is a type τ s.t. $\Delta \overset{A}{\vdash} e : \tau$ and for every $\Delta \overset{A}{\vdash} e : \tau'$*

- $\tau' = \tau[\rho]$ *for some $\rho \in \mathsf{Sub}$ with support $\mathrm{FV}(\tau) - \mathrm{FV}(\Delta, A)$*

We define a principal type algorithm $W(\Delta, A, e, \mathcal{K})$, where $\mathcal{K} \subseteq_{fin} \mathsf{A} \uplus \mathsf{B}$ is an auxiliary parameter that is threaded in recursive calls for recording the classifiers and type variables used so far. The algorithm either computes a type and a substitution for Δ and A (and updates $\mathcal{K}$) or fails. Figure 4 derives judgments of the form $\mathcal{K}, (\Delta, e, A) \Rightarrow \mathcal{K}', (\rho, \tau)$. When the judgment is derivable, it means that $W(\Delta, A, e, \mathcal{K}) = (\rho, \tau, \mathcal{K}')$. The rules use the following notation:

- $\text{close}(\tau, \Delta, A) \overset{\text{def}}{\equiv} \forall \overline{\kappa}.\tau$, where $\overline{\kappa} = \text{FV}(\tau) - \text{FV}(\Delta, A)$
- $\rho'\rho$ denotes composition of substitutions, i.e. $e[\rho'\rho] = (e[\rho])[\rho']$

The algorithm enjoys the following soundness and completeness properties.

Theorem 4 (Soundness). *If* $W(\Delta, A, e, \mathcal{K}) = (\rho, \tau, \mathcal{K}')$ *and* $\text{FV}(\Delta, A) \subseteq \mathcal{K}$ *then* $\Delta[\rho] \overset{A[\rho]}{\vdash} e : \tau$, *moreover* $\mathcal{K} \subseteq \mathcal{K}'$ *and* $\text{FV}(\tau, \Delta[\rho], A[\rho]) \subseteq \mathcal{K}'$.

Theorem 5 (Completeness). *If* $\Delta[\rho'] \overset{A[\rho']}{\vdash} e : \tau'$ *and* $\text{FV}(\Delta, A) \subseteq \mathcal{K}$ *and* $\mathcal{K} \subseteq_{fin} \mathsf{A} \cup \mathsf{B}$ *then* $W(\Delta, A, e, \mathcal{K}) = (\rho, \tau, \mathcal{K}')$ *is defined and exists* $\rho'' \in \mathsf{Sub}$ *s.t.* $\tau' = \tau[\rho'']$ *and* $\Delta[\rho'] \equiv \Delta[\rho''\rho]$ *and* $A[\rho'] = A[\rho''\rho]$.

Remark 2. In practice, one is interested in typing a complete program e, i.e. in computing $W(\emptyset, \emptyset, e, \emptyset)$. If the algorithm returns $(\rho, \tau, \mathcal{K}')$, then τ is the principal type and ρ and $\mathcal{K}'$ can be ignored. Even when the program uses a library, one can ignore the substitution ρ, since $\text{FV}(\Delta) = \emptyset$.

$$\frac{\beta \text{ and } A = (\alpha_i | i \in n) \text{ distinct and } \notin \mathcal{K}}{\mathcal{K}, (x, n) \Rightarrow \mathcal{K} \uplus \{\beta, A\}, (x : \beta^A, \beta, A)}$$

$$\frac{\mathcal{K}, (e_1, n) \Rightarrow \mathcal{K}', (\Gamma_1, \tau_1, A_1) \quad \mathcal{K}', (e_2, n) \Rightarrow \mathcal{K}'', (\Gamma_2, \tau_2, A_2) \quad \rho = mgu((\tau_1, \tau_2 \to \beta), (\Gamma_1, \Gamma_2), (A_1, A_2))}{\mathcal{K}, (e_1\ e_2, n) \Rightarrow \mathcal{K}'' \uplus \{\beta\}, (\Gamma_1 \cup \Gamma_2, \beta, A_1)[\rho]}$$

$$\frac{\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau, A) \quad x \notin \text{FV}(e) \text{ and } \beta \notin \mathcal{K}'}{\mathcal{K}, (\lambda x.e, n) \Rightarrow \mathcal{K}' \uplus \{\beta\}, (\Gamma, \beta \to \tau, A)}$$

$$\frac{\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau_2, A_2) \quad x \in \text{FV}(e), \ \Gamma(x) = \tau_1^{A_1} \text{ and } \rho = mgu(A_1, A_2)}{\mathcal{K}, (\lambda x.e, n) \Rightarrow \mathcal{K}', (\Gamma \setminus x, \tau_1 \to \tau_2, A_2)[\rho]}$$

$$\frac{\mathcal{K}, (e, n+) \Rightarrow \mathcal{K}', (\Gamma, \tau, (A, \alpha))}{\mathcal{K}, (\langle e \rangle, n) \Rightarrow \mathcal{K}', (\Gamma, \langle \tau \rangle^\alpha, A)}$$

$$\frac{\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau, A) \quad \rho = mgu(\tau, \langle \beta \rangle^\alpha) \text{ and } \beta, \alpha \notin \mathcal{K}'}{\mathcal{K}, (\tilde{}e, n+) \Rightarrow \mathcal{K}' \uplus \{\beta, \alpha\}, (\Gamma, \beta, (A, \alpha))[\rho]}$$

$$\frac{\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau, A) \quad \alpha \notin \mathcal{K}'}{\mathcal{K}, (\%e, n+) \Rightarrow \mathcal{K}' \uplus \{\alpha\}, (\Gamma, \tau, (A, \alpha))}$$

$$\frac{\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau, A) \quad \rho = mgu(\tau, \langle \beta \rangle) \text{ and } \beta \notin \mathcal{K}'}{\mathcal{K}, (\text{run } e, n) \Rightarrow \mathcal{K}' \uplus \{\beta\}, (\Gamma, \beta, A)[\rho]}$$

$$\frac{\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau, A) \quad \rho = mgu(\tau, \langle \beta \rangle^\alpha) \text{ and } \beta, \alpha \notin \mathcal{K}'}{\mathcal{K}, (\text{open } e, n) \Rightarrow \mathcal{K}' \uplus \{\beta, \alpha\}, (\Gamma, \langle \beta \rangle^\alpha, A)[\rho]}$$

$$\frac{\mathcal{K}, (e, n) \Rightarrow \mathcal{K}', (\Gamma, \tau, A) \quad \rho = mgu(\tau, \langle \beta \rangle^\alpha), \ \beta, \alpha \notin \mathcal{K}' \text{ and } \alpha[\rho] \notin \text{FV}((\Gamma, \beta, A)[\rho])}{\mathcal{K}, (\text{close } e, n) \Rightarrow \mathcal{K}'' \uplus \{\beta, \alpha\}, (\Gamma, \langle \beta \rangle, A)[\rho]}$$

Fig. 3. Principal Typing Algorithm

4 Relation to Other Calculi

This section studies the expressivity of the type system for λ^i, the simply-typed subset of λ^i_{let} (i.e. no let-binding and no quantification in type schema). The typing judgment for λ^i takes the form $\Gamma \overset{A}{\vdash} e : \tau$, since type schema collapse into types, and the typing rules are restricted accordingly. In summary, we have the following results:

- λ^i is a proper subset of λ^α, but the additional expressivity of λ^α comes at a price: the type system has *no principal types*.
- λ^i retains significant expressivity, namely, the embeddings given in [29] for two paradigmatic calculi $\lambda^\bigcirc$ and λ^{S4} (a variant of $\lambda^\square$) factor through λ^i.
- λ^i can be simplified further, by replacing run e with a construct runClosed e similar to Haskell's runST, and then removing $\langle\tau\rangle$, close e and open e, but doing so implies that the embedding of λ^{S4} no longer holds.

$$\frac{\Delta(x) \equiv (\forall\overline{\kappa}.\tau)^{A_1} \quad \rho = mgu(A, A_1) \quad \overline{\kappa}\#\mathcal{K}}{\mathcal{K}, (\Delta, x, A) \;\Rightarrow\; \mathcal{K} \uplus \{\overline{\kappa}\}, (\rho, \tau[\rho])}$$

$$\frac{\mathcal{K} \uplus \{\beta\}, (\Delta\{x : \beta^A\}, e, A) \;\Rightarrow\; \mathcal{K}', (\rho, \tau) \quad \beta \notin \mathcal{K}}{\mathcal{K}, (\Delta, \lambda x.e, A) \;\Rightarrow\; \mathcal{K}', (\rho, \beta[\rho] \to \tau)}$$

$$\frac{\mathcal{K}, (\Delta, e_1, A) \;\Rightarrow\; \mathcal{K}', (\rho_1, \tau_1) \quad \mathcal{K}', (\Delta[\rho_1], e_2, A[\rho_1]) \;\Rightarrow\; \mathcal{K}'', (\rho_2, \tau_2) \quad \rho = mgu(\tau_1[\rho_2], \tau_2 \to \beta) \quad \beta \notin \mathcal{K}''}{\mathcal{K}, (\Delta, e_1 e_2, A) \;\Rightarrow\; \mathcal{K}'' \uplus \{\beta\}, (\rho\rho_2\rho_1, \beta[\rho])}$$

$$\frac{\mathcal{K} \uplus \{\alpha\}, (\Delta, e, (A, \alpha)) \;\Rightarrow\; \mathcal{K}', (\rho, \tau) \quad \alpha \notin \mathcal{K}}{\mathcal{K}, (\Delta, \langle e\rangle, A) \;\Rightarrow\; \mathcal{K}', (\rho, \langle\tau\rangle^{\alpha[\rho]})}$$

$$\frac{\mathcal{K}, (\Delta, e, A) \;\Rightarrow\; \mathcal{K}', (\rho, \tau) \quad \rho' = mgu(\tau, \langle\beta\rangle^\alpha) \quad \beta \notin \mathcal{K}'}{\mathcal{K}, (\Delta, {\sim}e, (A, \alpha)) \;\Rightarrow\; \mathcal{K}' \uplus \{\beta\}, (\rho'\rho, \beta[\rho'])}$$

$$\frac{\mathcal{K}, (\Delta, e, A) \;\Rightarrow\; \mathcal{K}', (\rho, \tau)}{\mathcal{K}, (\Delta, \%e, (A, \alpha)) \;\Rightarrow\; \mathcal{K}', (\rho, \tau)}$$

$$\frac{\mathcal{K}, (\Delta, e, A) \;\Rightarrow\; \mathcal{K}', (\rho, \tau) \quad \rho' = mgu(\tau, \langle\beta\rangle) \quad \beta \notin \mathcal{K}'}{\mathcal{K}, (\Delta, \mathsf{run}\ e, A) \;\Rightarrow\; \mathcal{K}' \uplus \{\beta\}, (\rho'\rho, \beta[\rho'])}$$

$$\frac{\mathcal{K}, (\Delta, e, A) \;\Rightarrow\; \mathcal{K}', (\rho, \tau) \quad \rho' = mgu(\tau, \langle\beta\rangle) \quad \alpha, \beta \notin \mathcal{K}'}{\mathcal{K}, (\Delta, \mathsf{open}\ e, A) \;\Rightarrow\; \mathcal{K}' \uplus \{\alpha, \beta\}, (\rho'\rho, \langle\beta[\rho']\rangle^\alpha)}$$

$$\frac{\mathcal{K}, (\Delta, e, A) \;\Rightarrow\; \mathcal{K}', (\rho, \tau) \quad \rho' = mgu(\tau, \langle\beta\rangle^\alpha) \quad \alpha, \beta \notin \mathcal{K}' \quad \alpha[\rho'] \notin \mathrm{FV}(\Delta[\rho'\rho], A[\rho'\rho], \beta[\rho'])}{\mathcal{K}, (\Delta, \mathsf{close}\ e, A) \;\Rightarrow\; \mathcal{K}' \uplus \{\alpha, \beta\}, (\rho'\rho, \langle\beta[\rho']\rangle)}$$

$$\frac{\mathcal{K}, (\Delta, e_1, A) \;\Rightarrow\; \mathcal{K}', (\rho_1, \tau_1) \quad \mathcal{K}', (\Delta[\rho_1]\{x : \mathsf{close}(\tau_1, \Delta[\rho_1], A[\rho_1])^{A[\rho_1]}\}, e_2, A[\rho_1]) \;\Rightarrow\; \mathcal{K}'', (\rho_2, \tau_2)}{\mathcal{K}, (\Delta, \mathsf{let}\ x = e_1\ \mathsf{in}\ e_2, A) \;\Rightarrow\; \mathcal{K}'', (\rho_2\rho_1, \tau_2)}$$

Fig. 4. Principal Type Algorithm

4.1   Relation to λ^α

The key feature of λ^α is the inclusion of a special quantifier $(\alpha)\tau$ in the language of types, representing universal quantification over classifiers. Figure 5 recalls the BNF for terms and types, and the most relevant typing rules [29]. In λ^i the main difference is that the quantifier $(\alpha)\tau$ of λ^α is replaced by the runnable code type $\langle\tau\rangle$. In fact, $\langle\tau\rangle$ corresponds to a restricted form of quantification, namely $(\alpha)\langle\tau\rangle^\alpha$ with $\alpha \notin \mathrm{FV}(\tau)$. It is non-trivial to define formally a typability-preserving embedding of λ^i into λ^α, since we need to recover classifier names in terms. Therefore, we justify the correspondence at the level of terms only informally:

- The terms close e and open e of λ^i correspond to $(\alpha)e$ and $e[\alpha]$ of λ^α. Since λ^i has no classifier names in terms, these constructs record that a classi-

Terms $e \in \mathsf{E} ::= x \mid \lambda x.e \mid e\ e \mid \langle e \rangle^{\alpha} \mid\ \tilde{}\ e \mid \%e \mid \mathsf{run}\ e \mid (\alpha)e \mid e[\alpha]$
Types $\tau \in \mathsf{T} ::= \beta \mid \tau_1 \rightarrow \tau_2 \mid \langle \tau \rangle^{\alpha} \mid (\alpha)\tau$

$$\text{brck}\ \dfrac{\Gamma \overset{A,\alpha}{\vdash} e : \tau}{\Gamma \overset{A}{\vdash} \langle e \rangle^{\alpha} : \langle \tau \rangle^{\alpha}} \qquad \text{esc}\ \dfrac{\Gamma \overset{A}{\vdash} e : \langle \tau \rangle^{\alpha}}{\Gamma \overset{A,\alpha}{\vdash}\ \tilde{}\ e : \tau} \qquad \text{csp}\ \dfrac{\Gamma \overset{A}{\vdash} e : \tau}{\Gamma \overset{A,\alpha}{\vdash} \%e : \tau} \qquad \text{all-run}\ \dfrac{\Gamma \overset{A}{\vdash} e : (\alpha)\langle \tau \rangle^{\alpha}}{\Gamma \overset{A}{\vdash} \mathsf{run}\ e : (\alpha)\tau}$$

$$\text{all-close}\ \dfrac{\Gamma \overset{A}{\vdash} e : \tau}{\Gamma \overset{A}{\vdash} (\alpha)e : (\alpha)\tau}\ \alpha \notin \mathrm{FV}(\Gamma, A) \qquad \text{all-open}\ \dfrac{\Gamma \overset{A}{\vdash} e : (\alpha)\tau}{\Gamma \overset{A}{\vdash} e[\alpha'] : \tau[\alpha : \alpha']}$$

Fig. 5. Type System for λ^{α} (adapted from [29])

fier abstraction and instantiation has occurred *without* naming the classifier involved. (Similarly, the term $\langle e \rangle$ in λ^i corresponds to $\langle e \rangle^{\alpha}$ in λ^{α}.)

- The term $\%e$ has exactly the same syntax and meaning in the two calculi.
- The term $\mathsf{run}\ e$ of λ^i corresponds to $(\mathsf{run}\ e)[\alpha']$ of λ^{α}, where α' can be chosen arbitrarily without changing the result type. In fact, the type of run in λ^{α} is $((\alpha)\langle \tau \rangle^{\alpha}) \rightarrow (\alpha)\tau$, while in λ^i it is $\langle \tau \rangle \rightarrow \tau$, which corresponds to $((\alpha)\langle \tau \rangle^{\alpha}) \rightarrow \tau$ with $\alpha \notin \mathrm{FV}(\tau)$.

We conclude the comparison between λ^i and λ^{α} by showing that type inference in λ^{α} is problematic.

Lack of principal types in λ^{α}. Consider the closed term $e \equiv (\lambda x.\mathsf{run}\ x)$. We can assign to e exactly the types of the form $((\alpha)\langle \tau \rangle^{\alpha}) \rightarrow (\alpha)\tau$ with an arbitrary type τ, including ones with $\alpha \in \mathrm{FV}(\tau)$. However, e does not have a *principal type*, i.e. one from which one can recover all other types (modulo α-conversion of bound classifiers) by applying a substitution $\rho \in \mathsf{Sub}$ for classifiers and type variables. In fact, the obvious candidate for the principal type, i.e. $((\alpha)\langle \beta \rangle^{\alpha}) \rightarrow (\alpha)\beta$, allows us to recover only the types of the form $((\alpha)\langle \tau \rangle^{\alpha}) \rightarrow (\alpha)\tau$ with $\alpha \notin \mathrm{FV}(\tau)$, since substitution should be capture avoiding.

Lack of principal types in previous polymorphic extensions of λ^{α}. A more expressive type system for λ^{α} was previously proposed [29], where type variables β are replaced by variables β^n ranging over types parameterized w.r.t. n classifiers. Thus, the BNF for types becomes:

$$\tau \in \mathsf{T} ::= \beta^n[A] \mid \tau_1 \rightarrow \tau_2 \mid \langle \tau \rangle^{\alpha} \mid (\alpha)\tau \quad \text{with } |A| = n$$

In this way, there is a better candidate for the principal type of e, namely $((\alpha)\langle \beta^1[\alpha] \rangle^{\alpha}) \rightarrow (\alpha)\beta^1[\alpha]$.

In this extension, standard unification techniques are no longer applicable, and some form of higher-order unification is needed. However, even in this system, there are typable terms that do not have a principal type. For instance, the term $e = (x(x_1[\alpha]), f(x_2[\alpha]))$ (for simplicity, we assume that we have pairing and product types) has no principal typing, in fact

$$
\begin{array}{ll}
\text{Terms} & e \in \mathsf{E} ::= x \mid \lambda x.e \mid e\ e \mid \mathsf{box}\ e \mid \mathsf{unbox}_n\ e \\
\text{Types} & \tau \in \mathsf{T} ::= \beta \mid \tau_1 \to \tau_2 \mid \Box\tau \\
\text{Assignments} & \Gamma \in \mathsf{X} \xrightarrow{fin} \mathsf{T}\ \text{of types} \\
\text{Stacks} & \Psi \in (\mathsf{X} \xrightarrow{fin} \mathsf{T})^*\ \text{of type assignments}
\end{array}
$$

$$
\dfrac{}{\Psi; \Gamma \vdash x : \tau}\ \Gamma(x) = \tau
\qquad
\dfrac{\Psi; \Gamma, x : \tau_1 \vdash e : \tau_2}{\Psi; \Gamma \vdash \lambda x.e : \tau_1 \to \tau_2}
\qquad
\dfrac{\Psi; \Gamma \vdash e_1 : \tau_1 \to \tau_2 \qquad \Psi; \Gamma \vdash e_2 : \tau_1}{\Psi; \Gamma \vdash e_1 e_2 : \tau_2}
$$

$$
\dfrac{\Psi; \Gamma; () \vdash e : \tau}{\Psi; \Gamma \vdash \mathsf{box}\ e : \Box\tau}
\qquad
\dfrac{\Psi; \Gamma \vdash e : \Box\tau}{\Psi; \Gamma; \Gamma_1; \dots; \Gamma_n \vdash \mathsf{unbox}_n\ e : \tau}
$$

Fig. 6. Type system for λ^{S4} [8, Section 4.3]

- x must be a function, say of type $\tau \to \tau'$
- x_i must be of type $(\alpha_i)\tau_i$, among them the most general is $(\alpha_i)\beta_i^1[\alpha_i]$
- τ and $\tau_i[\alpha_i : \alpha]$ must be the same, but there is no most general unifier for $\beta_1^1[\alpha] = \beta_2^1[\alpha]$.

4.2   Embedding of $\lambda^{\bigcirc}$

The embedding of $\lambda^{\bigcirc}$ [6] into λ^i is direct. We pick one arbitrary classifier α and define the embedding as follows:

$$
[\![\beta]\!] \equiv \beta, \quad [\![\bigcirc\tau]\!] \equiv \langle [\![\tau]\!] \rangle^\alpha, \quad [\![\tau_1 \to \tau_2]\!] \equiv [\![\tau_1]\!] \to [\![\tau_2]\!]
$$
$$
[\![n]\!] \equiv \alpha^n, \quad [\![x_i : \tau_i^{n_i}]\!] \equiv x_i : [\![\tau_i]\!]^{[\![n_i]\!]}
$$
$$
[\![x]\!] \equiv x, \quad [\![\lambda x.e]\!] \equiv \lambda x.[\![e]\!], \quad [\![e_1\ e_2]\!] \equiv [\![e_1]\!]\ [\![e_2]\!]
$$
$$
[\![\mathsf{next}\ e]\!] \equiv \langle [\![e]\!] \rangle \quad [\![\mathsf{prev}\ e]\!] \equiv {}^\sim[\![e]\!]
$$

The translation preserves the typing, i.e.

Theorem 6. *If* $\Gamma \overset{n}{\vdash} e : \tau$ *is derivable in* $\lambda^{\bigcirc}$, *then* $[\![\Gamma]\!] \overset{[\![n]\!]}{\vdash} [\![e]\!] : [\![\tau]\!]$ *is derivable in* λ^i.

The translation preserves also the big-step operational semantics.

4.3   Embedding of λ^{S4}

Figure 6 recalls the type system of λ^{S4} [8, Section 4.3]. This calculus is *equivalent* to $\lambda^{\Box}$ [7], but makes explicit use of levels in typing judgments. The operational semantics of λ^{S4} is given indirectly [8, Section 4.3] via the translation into $\lambda^{\Box}$. The embedding of this calculus into λ^i is as follows: the embedding maps types to types:

$$
[\![\Box\tau]\!] \equiv \langle [\![\tau]\!] \rangle \quad [\![\tau_1 \to \tau_2]\!] \equiv [\![\tau_1]\!] \to [\![\tau_2]\!] \quad [\![\beta]\!] \equiv \beta
$$

The embedding on terms is parameterized by a level m:

$$
[\![x]\!]_m \equiv x \quad [\![\lambda x.e]\!]_m \equiv \lambda x.[\![e]\!]_m \quad [\![e_1 e_2]\!]_m \equiv [\![e_1]\!]_m [\![e_2]\!]_m
$$
$$
[\![\mathsf{box}\ e]\!]_m \equiv \mathsf{close}\ \langle [\![e]\!]_{m+1} \rangle
$$
$$
[\![\mathsf{unbox}_0\ e]\!]_m \equiv \mathsf{run}\ [\![e]\!]_m \quad [\![\mathsf{unbox}_{n+1}\ e]\!]_{m+n+1} \equiv \%^n({}^\sim(\mathsf{open}\ [\![e]\!]_m))
$$
$$
\%^0(e) \equiv e \quad \%^{n+1}(e) \equiv \%^n(\%e)
$$

The translation of unbox_m depends on the subscript m. unbox_0 corresponds to running code. When $m > 0$ the term unbox_m corresponds to $\tilde{\ }-$, but if $m > 1$ it also digs into the environment stack to get code from previous stages, and thus the need for the sequence of %s. To define the translation of typing judgments, we must fix a sequence of distinct classifiers $\alpha_1, \alpha_2, \ldots$, and we write A_i for the prefix of the first i classifiers, i.e. $A_i = \alpha_1, \ldots, \alpha_i$:

$$[\![\Gamma_0; \ldots; \Gamma_n \vdash e : \tau]\!] \equiv [\![\Gamma_0]\!]^{A_0}, \ldots, [\![\Gamma_n]\!]^{A_n} \stackrel{A_n}{\vdash} [\![e]\!]_n : [\![\tau]\!]$$

where $[\![x_1 : \tau_1, \ldots, x_n : \tau_n]\!]^A \equiv x_1 : [\![\tau_1]\!]^A, \ldots, x_n : [\![\tau_n]\!]^A$. The translation preserves the typing, i.e.

Theorem 7. *If $\Gamma_0; \ldots; \Gamma_n \vdash e : \tau$ is derivable in λ^{S4}, then $[\![\Gamma_0; \ldots; \Gamma_n \vdash e : \tau]\!]$ is derivable in λ^i.*

4.4   Relation to Haskell's runST

The typing rules for close and run can be combined into one rule analogous to that for the Haskell's runST [17,18], namely,

$$\text{runClosed } \frac{\Delta \stackrel{A}{\vdash} e : \langle\tau\rangle^\alpha}{\Delta \stackrel{A}{\vdash} \mathsf{runClosed}\ e : \tau} \quad \alpha \notin \mathrm{FV}(\Delta, A, \tau)$$

With this rule in place, there is no need to retain the type $\langle\tau\rangle$ and the terms close e and open e, thus resulting in a proper fragment of λ^i. There is a loss in expressivity, because the embedding of λ^{S4} does not factor through this fragment. In fact, the term $\lambda x.\mathsf{runClosed}\ x$ is not typable, while $\lambda x.\mathsf{run}\ x$ is typable in λ^i_{let} (but $\lambda x.\mathsf{close}\ x$ is still not typable). This variant was implemented in the MetaOCaml system (giving `.!` the typing of runClosed), and it was found that it required no change to the existing code base of multi-stage programs.

References

1. M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin. Dynamic typing in a statically typed language. *ACM Transactions on Programming Languages and Systems*, 13(2):237–268, April 1991.
2. Alan Bawden. Quasiquotation in LISP. In O. Danvy, editor, *Proceedings of the Workshop on Partial Evaluation and Semantics-Based Program Manipulation*, pages 88–99, San Antonio, 1999. University of Aarhus, Dept. of Computer Science. Invited talk.
3. Cristiano Calcagno, Eugenio Moggi, and Tim Sheard. Closed types for a safe imperative MetaML. *Journal of Functional Programming*, 2003. To appear.
4. Cristiano Calcagno, Walid Taha, Liwen Huang, and Xavier Leroy. Implementing multi-stage languages using asts, gensym, and reflection. In Krzysztof Czarnecki, Frank Pfenning, and Yannis Smaragdakis, editors, *Generative Programming and Component Engineering (GPCE)*, Lecture Notes in Computer Science. Springer-Verlag, 2003.

5. Luís Damas and Robin Milner. Principal type schemes for functional languages. In *9th ACM Symposium on Principles of Programming Languages*. ACM, August 1982.
6. Rowan Davies. A temporal-logic approach to binding-time analysis. In *the Symposium on Logic in Computer Science (LICS '96)*, pages 184–195, New Brunswick, 1996. IEEE Computer Society Press.
7. Rowan Davies and Frank Pfenning. A modal analysis of staged computation. In *the Symposium on Principles of Programming Languages (POPL '96)*, pages 258–270, St. Petersburg Beach, 1996.
8. Rowan Davies and Frank Pfenning. A modal analysis of staged computation. *Journal of the ACM*, 48(3):555–604, 2001.
9. Dawson R. Engler, Wilson C. Hsieh, and M. Frans Kaashoek. 'C: A language for high-level, efficient, and machine-independent dynaic code generation. In *In proceedings of the ACM Symposium on Principles of Programming Languages (POPL)*, pages 131–144, St. Petersburg Beach, 1996.
10. Steven Ganz, Amr Sabry, and Walid Taha. Macros as multi-stage computations: Type-safe, generative, binding macros in MacroML. In *the International Conference on Functional Programming (ICFP '01)*, Florence, Italy, September 2001. ACM.
11. Carsten K. Gomard and Neil D. Jones. A partial evaluator for untyped lambda calculus. *Journal of Functional Programming*, 1(1):21–69, 1991.
12. Brian Grant, Matthai Philipose, Markus Mock, Craig Chambers, and Susan J. Eggers. An evaluation of staged run-time optimizations in DyC. In *Proceedings of the Conference on Programming Language Design and Implementation*, pages 293–304, 1999.
13. J. Roger Hindley. *Basic Simple Type Theory*, volume 42 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, 1997.
14. Trevor Jim. What are principal typings and what are they good for? In *Conf. Rec. POPL '96: 23rd ACM Symp. Princ. of Prog. Langs.*, 1996.
15. Neil D. Jones, Carsten K. Gomard, and Peter Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice-Hall, 1993.
16. Sam Kamin, Miranda Callahan, and Lars Clausen. Lightweight and generative components II: Binary-level components. In *[27]*, pages 28–50, 2000.
17. John Launchbury and Simon L. Peyton Jones. State in haskell. *LISP and Symbolic Computation*, 8(4):293–342, 1995.
18. John Launchbury and Amr Sabry. Monadic state: Axiomatization and type safety. In *Proceedings of the International Conference on Functional Programming*, Amsterdam, 1997.
19. MetaOCaml: A compiled, type-safe multi-stage programming language. Available online from http://www.cs.rice.edu/ taha/MetaOCaml/, 2001.
20. The MetaML Home Page, 2000. Provides source code and documentation online at http://www.cse.ogi.edu/PacSoft/projects/metaml/index.html.
21. A. Nanevski and F. Pfenning. Meta-programming with names and necessity. submitted, 2003.
22. Flemming Nielson and Hanne Riis Nielson. Two-level semantics and code generation. *Theoretical Computer Science*, 56(1):59–133, 1988.
23. Tim Sheard and Simon Peyton-Jones. Template meta-programming for haskell. In *Proc. of the workshop on Haskell*, pages 1–16. ACM, 2002.
24. Mark Shields, Tim Sheard, and Simon L. Peyton Jones. Dynamic typing through staged type inference. In *In proceedings of the ACM Symposium on Principles of Programming Languages (POPL)*, pages 289–302, 1998.

25. Frederick Smith, Dan Grossman, Greg Morrisett, Luke Hornof, and Trevor Jim. Compiling for run-time code generation. *Journal of Functional Programming*, 2003.
26. Walid Taha. *Multi-Stage Programming: Its Theory and Applications.* PhD thesis, Oregon Graduate Institute of Science and Technology, 1999.
27. Walid Taha, editor. *Semantics, Applications, and Implementation of Program Generation*, volume 1924 of *Lecture Notes in Computer Science*, Montréal, 2000. Springer-Verlag.
28. Walid Taha. A sound reduction semantics for untyped CBN multi-stage computation. Or, the theory of MetaML is non-trivial. In *Proceedings of the Workshop on Partial Evaluation and Semantics-Based Program Maniplation (PEPM)*, Boston, 2000. ACM Press.
29. Walid Taha and Michael Florentin Nielsen. Environment classifiers. In *The Symposium on Principles of Programming Languages (POPL '03)*, New Orleans, 2003.
30. Walid Taha and Tim Sheard. Multi-stage programming with explicit annotations. In *Proceedings of the Symposium on Partial Evaluation and Semantic-Based Program Manipulation (PEPM)*, pages 203–217, Amsterdam, 1997. ACM Press.
31. Joe Wells. The essence of principal typings. In *Proc. 29th Int'l Coll. Automata, Languages, and Programming*, volume 2380 of *Lecture Notes in Computer Science*, pages 913–925. Springer-Verlag, 2002.