Functional programming with λ-tree syntax

Ulysse Gérard, Dale Miller, and Gabriel Scherer

Inria Saclay & LIX, Ecole Polytechnique

Abstract. We present the design of a new functional programming language, MLTS, that uses the λ-*tree* syntax approach to encoding bindings that appear within data structures. In this setting, bindings never become free nor escape their scope: instead, binders in data structures are permitted to *move* into binders within programs. The design of MLTS—whose concrete syntax is based on that of OCaml—includes additional sites within programs that directly support this movement of bindings. We illustrate the features of MLTS by presenting several collections of examples. We also present a typing discipline that naturally extends the typing of OCaml programs. In order to formally define the language's operational semantics, we present an abstract syntax for MLTS and a natural semantics for its evaluation. We shall view such natural semantics as a logical theory with a rich logic that includes both nominal abstraction and the ∇-quantifier: as a result, the natural semantic specification of MLTS can be given a succinct and elegant presentation.

Keywords: λ-tree syntax; binder mobility; functional programming; higher-order abstract syntax

1 Introduction

Even from the earliest days of high-level programming, functional programming languages were used to build systems that manipulated the syntax of various programming languages and logics. For example, Lisp was a common language for building theorem provers, interpreters, compilers, and parsers, and the ML programming language was designed as a "meta-language" for a proof checker [Gordon, Milner, and Wadsworth, 1979]. While these various tasks involve the manipulation of syntax, none of these earliest functional programming languages provided support for a key feature of almost all programming languages and logics: variable binding.

Bindings in syntactic expressions have been given, of course, a range of different treatments within the functional programming setting. Common approaches are to implement bindings by using variable names or, in a more abstract way, by using de Bruijn's nameless dummies [de Bruijn, 1979]. Since such techniques are quite complex to get right and since bindings are so pervasive, a great deal of energy has gone into making libraries of procedures that can help deal with binders: for example, there is the *locally nameless* approach [Charguéraud, 2011, Gordon, 1994, McBride and McKinna, 2004] and the *parametric higher-order abstract syntax* approach [Chlipala, 2008].

Extending a functional programming language with features that support bindings in data has been considered before: for example, there have been the FreshML [Shinwell, Pitts, and Gabbay, 2003, Pottier, 2007] and CαML [Pottier, 2006] extensions to ML-style functional programming languages. Also, entirely new functional programming languages, such as the dependently typed Beluga [Pientka and Dunfield, 2010] language, have been designed and implemented with the goal to support bindings in syntax. In the domain of logic programming and theorem prover conception, several designs and implemented systems exist that incorporate approaches to binding: such systems include Isabelle's generic reasoning core [Paulson, 1989], λProlog [Nadathur and Miller, 1988, Miller and Nadathur, 2012], Qu-Prolog [Cheng, Robinson, and Staples, 1991], Twelf [Pfenning and Schürmann, 1999], αProlog [Cheney and Urban, 2004], the Minlog prover [Schwichtenberg, 2006], and the Abella theorem prover [Baelde, Chaudhuri, Gacek, Miller, Nadathur, Tiu, and Wang, 2014].

In this paper we present MLTS, a new language that extend (the core of) ML and incorporates the λ-*tree syntax* approach to encoding the abstract syntax of data structures containing binders. Briefly, we can define the λ-tree syntax approach to syntax as following the three tenets: (1) Syntax is encoded as simply typed λ-terms in which the primitive types are identified with syntactic categories. (2) Equality of syntax must include $\alpha\beta_0\eta$ conversion (defined in Section 7.2). (3) Bound variables never become free: instead, their binding scope can move. This latter tenet introduces the most characteristic aspect of λ-tree syntax which is often called *binder mobility*. MLTS is, in fact, an acronym for *mobility and λ-tree syntax*.

2 The new features of MLTS

We chose the concrete syntax of MLTS to be an extension of that of the OCaml programming language (a program in MLTS not using the new language features should be accepted by the `ocamlc` compiler). We shall assume that the reader is familiar with basic syntactic conventions of OCaml [OCaml, 2018], many of which are shared with most ML-like programming languages. MLTS contains the following five new language features.

1. Datatypes can be extended to contain new *nominal* constants and the `(new X in body)` program phrase provides a binding that declares that the nominal `X` is new within the lexical scope given by `body`.
2. A new typing constructor `=>` is used to type bindings within term structures. This constructor is an addition to the already familiar constructor `->` used for the typing of functional expressions.
3. The *backslash* (`\` as an infix symbol that associates to the right) is used to form an abstraction of a nominal over its scope. For example, `(X\body)` is a syntactic expression that hides the nominal `X` in the scope `body`. Thus the backslash *introduces* an abstraction.

4. The `@` *eliminates* an abstraction: for example, the expression `((X\body) @ t)` denotes the result of substituting the abstracted nominal `X` with the term `t` in `body`.
5. Clauses within match-expressions can also contain the `(nab X in rule)` binder: in the scope of this binder, the symbol `X` can match existing nominals introduced by the `new` binder and the `\` operator. Note that `X` is bound over the entire rule (including both the left and right-side of the rule).

All three bindings expressions—`(X\body)`, `(new X in body)` and `(nab X in rule)`—are subject to α-renaming of bound variables, just as the names of variables bound in `let` declarations and function definitions. As we shall see, nominals are best thought of as constructors: as a consequence, we follow the OCaml convention of capitalizing the name of their binders. We are assuming that, in all parts of MLTS, the names of nominals (of bound variables in general) are not available to programs since α-conversion (the alphabetic change of bound variables) is always applicable. Thus, compilers are free to implement nominals in any number of ways, even ways in which they do not have, say, print names.

Expressions involving `@` are greatly restricted within patterns of match expressions: in particular the expression `(m @ X1 ... Xj)` is restricted so that `m` is a pattern variable and `X1`, ..., `Xj` are distinct nominals bound within the scope of the pattern binding on `m`. This restriction is essentially the same as those required by *higher-order pattern unification* [Miller, 1991]: as a result, pattern matching in this setting is a simple generalization of usual first-order pattern matching.

We note that the expression `(X\ r @ X)` is interchangeable with the simple expression `r`: that is, when `r` is of `=>` type, an η-equality holds.

We now present several sets of examples of MLTS programs in the next sections: the Appendix contains an additional example. We hope that the informal semantics given above plus the simplicity of the examples will give a working understanding of the semantics of MLTS. We delay the formal definition of the operational semantics of MLTS until Section 7.

3 MLTS examples: the untyped λ-calculus

The untyped λ-terms can be defined in MLTS as the following datatype:

```
type tm =
     | App of tm * tm
     | Abs of tm => tm;;
```

The use of the `=>` type constructor here indicates that the argument of `Abs` is a *binding abstraction* of a `tm` over a `tm`. Just as the type `tm` denotes a syntactic category of untyped λ-terms, the type `tm => tm` denotes the syntactic category of terms abstracted over such terms.

Following usual conventions, expressions whose concrete syntax have nested binders using the same name are disambiguated by the parser by linking the named variable with the closest binder. Thus, the concrete syntax `(Abs(X\ Abs(X\ X)))`

is parsed as a term α-equivalent to `(Abs(Y\ Abs(X\ X)))`. Similarly, the expression `(let n = 2 in let n = 3 in n)` is parsed as an expression α-equivalent to `(let m = 2 in let n = 3 in n)`: this expression has value `3`.

The following MLTS program computes the size of an untyped λ-term.

```
let rec size term =
        match term with
        | App(n, m)  -> 1 + size n + size m
        | Abs(r)     -> 1 + new X in size (r @ X)
        | nab X in X -> 1;;
```

For example, `(size (App(Abs(X\X), Abs(X\X))))` evaluates to 5. In the second match rule, the match-variable `r` will be bound to an expression built using the backslash. On the right of that rule, `r` is applied to a single argument which is a newly provided nominal constructor of type `tm`. The third match rule contains the `nab` binder that allows the token `X` to match any nominal: alternatively, that last clause could have matched any non-`App` and non-`Abs` term by using the clause `| _ -> 1`. (Note that as written, the three match rules used to define `size` could have been listed in any order.) The following sequence of expressions shows the evolution of a computation involving the `size` function.

```
size (Abs (X\ Abs (Y\ App(X,Y))));;
1 + new X in size (Abs (Y\ App(X,Y)));;
1 + new X in 1 + new Y in size (App(X,Y));;
1 + new X in 1 + new Y in 1 + size X + size Y;;
1 + new X in 1 + new Y in 1 + 1 + 1;;
```

The first call to `size` will bind the pattern variable `r` to `X\ Abs(Y\ App(X,Y))`. It is important to note that the names of bound variables within MLTS programs and data structures are fictions: in the expressions above, binding names are chosen for readability.

Figure 1 defines the function `(subst t u)` that takes an abstraction over terms `t` and a term `u` and returns the result of substituting the (top-level) bound variable of `t` with `u`. This function works by first introducing a new nominal `X` and then defining an auxiliary function that replaces that nominal in a term with the term `u`. Finally, that auxiliary function is called on the expression `(t @ X)` which is the result of "moving" the top-level bound variable in `t` to the binding occurrence of the expression `new X in`. (As we note at the end of Section 10.2, such binder movement can be implemented in constant time and does not need to involve an actual substitution of a nominal for a bound variable.) This substitution function has the type `(tm => tm) -> (tm -> tm)`: that is, it is used to inject the abstraction type `=>` into the function type `->`. Substitution is then used by the second function of Figure 1, `beta`, to compute the β-normal form of a given term of type `tm`. This figure also contains the Church numeral for 2 and operations for addition and multiplication on Church numerals. In the resulting evaluation context, the values computed by `(beta (App(App(plus, two), two)))` and `(beta (App(App(times, two), two)))` are both the Church numeral for 4.

```
let subst t u = new X in
    let rec aux t = match t with
      | X -> u
      | nab Y in Y -> Y
      | App(u, v) -> App(aux u, aux v)
      | Abs r -> Abs(Y\ aux (r @ Y))
    in aux (t @ X);;

let rec beta t = match t with
  | Abs r      -> Abs(Y\ beta (r @ Y))
  | nab X in X -> X
  | App(m, n)  ->
    let m = beta m in let n = beta n in
    begin
        match m with
        | Abs r -> beta (subst r n)
        | _ -> App(m, n)
    end ;;

let two   = Abs(F\ Abs(X\ App(F, App(F, X))));;
let plus  = Abs(M\ Abs(N\ Abs(F\ Abs(X\
                    App(App(M, F), App(App(N, F), X))))));;
let times = Abs(M\ Abs(N\ Abs(F\ Abs(X\
                           App(App(M, App(N, F)), X)))));;
```

Fig. 1. The function for computing the substitution $[t/x]u$ and the (partial) function that returns the β-normal form of its argument.

For another example, consider a program that returns `true` if and only if its argument, of type `tm => tm`, is such that its top-level bound variable is a "vacuous" binding. Figure 2 contains three implementations of this boolean-valued function. The first implementation proceeds by matching patterns with the prefix `X\`, thereby, matching expressions of type `tm => tm`. The second implementation uses a different style: it creates a new nominal `X` and proceeds to work on the term `t @ X`, in the same fashion as the `size` example. The internal `aux` function is then defined to search for occurrences of `X` in that term. The third implementation, `vacp3`, is not (overtly) recursive since the entire effort of checking for the vacuous binding is done during pattern matching. The first match rule of this third implementation is essentially asking the question: is there an instantiation for the (pattern) variable s so that the $\lambda x.s$ equals t? This question can be posed as asking if the logical formula $\exists s.(\lambda x.s) = t$ can be proved. In this latter form, it should be clear that since substitution is intended as a logical operation, the result of substituting for s never allows for variable capture. Hence, every instance of the existential quantifier yields an equation with a left-hand side that is a vacuous abstraction. Of course, this kind of pattern matching requires a recursive analysis of the term t.

```
let rec vacp1 t = match t with
    | X\ X                  -> false
    | nab Y in X\ Y         -> true
    | X\ App(m @ X, n @ X) -> vacp1 m && vacp1 n
    | X\ Abs(Y\ r @ X Y)    -> new Y in vacp1 (X\ r @ X Y);;

let rec vacp2 t =
    new X in
    let rec aux term = match term with
        | X           -> false
        | nab Y in Y -> true
        | App(m, n)   -> aux m && aux n
        | Abs(u)      -> new Y in aux (u @ Y)
    in aux (t @ X);;

let vacp3 t = match t with
    | X\ s  -> true
```

Fig. 2. Three implementations for determining if an abstraction is vacuous.

```
let rec assoc x alist = match alist with
  | (u,y)::alst -> if (u = x) then y else assoc x alst;;

type tm' =
     | App' of tm' * tm'
     | Abs' of tm' => tm';;

let rec id g term = match term with
  | App(m,n)    -> App'(id g m, id g n)
  | Abs(r)      -> new X in Abs'(Y\ id ((X, Y)::g) (r @ X))
  | nab X in X -> assoc X g;;
```

Fig. 3. Translating from `tm` to its mirror version `tm'`.

For a simple example of computing on the untyped λ-calculus, consider introducing a mirror version of `tm`, as is done in Figure 3, and writing the function that constructs the mirror term in `tm'` from an input term `tm`. This computation is achieved by adding a context (an association list) as an extra argument that maintains the association of bound variables of type `tm` and those of type `tm'`. The value of `id [] (Abs(X\ Abs(Y\ App(X,Y))))` is `(Abs'(X\ Abs'(Y\ App'(X,Y))))` (the types of `X` and `Y` in these two expressions are, of course, different).

Figure 4 presents a datatype for the untyped λ-calculus in De Bruijn's style nameless dummies [de Bruijn, 1972] as well as the functions that can convert between that syntax and the one with explicit bindings. The auxiliary functions `nth` and `index` take a list of nominals as their second argument: `nth` takes also

```
type deb =
     | Dapp of deb * deb
     | Dabs of deb
     | Dvar of int;;

let rec nth n l = match (n, l) with
  | (0, x::k) -> x
  | (c, x::k) -> nth (c - 1) k;;

let index x l =
    let rec aux c x k = match (x, k) with
          | nab X   in (X, X::(l @ X))   -> c
          | nab X Y in (X, Y::(l @ X Y)) ->
                                    aux (c + 1) x (l @ X Y)
     in aux 0 x l;;

let rec trans prefix term = match term with
  | App(m, n)  -> Dapp(trans prefix m, trans prefix n)
  | Abs r      -> new X in Dabs(trans (X::prefix) (r @ X))
  | nab Y in Y -> Dvar (index Y prefix);;

let rec dtrans prefix term = match term with
  | Dapp(m, n) -> App(dtrans prefix m, dtrans prefix n)
  | Dabs r     -> Abs(X\ dtrans (X::prefix) r)
  | Dvar c     -> nth c prefix;;
```

Fig. 4. De Bruijn's nameless dummy syntax and its conversions with type `tm`.

an integer `n` and returns the n^{th} nominal in that list while `index` takes a nominal and returns its ordinal position in that list. For example, the value of

```
trans [] (Abs(X\ Abs(Y\ Abs(Z\ App(X, Abs(W\ Z))))));;
```

is the term `DAbs(DAbs(DAbs(DApp(Dvar 2, DAbs(Dvar 1)))))` of type `deb`. If `dtrans []` is applied to this second term, the former term is returned (modulo α-renaming, of course).

4 Higher-order programming examples

Recall the familiar "fold-right" higher-order function.

```
let rec foldr f a lst = match lst with
  | [] -> a
  | x :: xs -> f x (foldr f a xs);;
```

This function can be viewed as replacing all occurrences of `::` with the binary function `f` and all occurrences of `[]` with `a`. The higher-order program `maptm` in Figure 5 does the analogous operation on the datatype of untyped λ-terms

```
let rec maptm fapp fabs fvar t = match t with
    | App(m,n) -> fapp (maptm fapp fabs fvar m)
                       (maptm fapp fabs fvar n)
    | Abs r  -> fabs (fun x -> maptm fapp fabs fvar (r @ x))
    | nab X in X -> fvar X;;

let lookup sub var = match var with
  | nab X in X ->
    let rec aux s = match s with
                    | []          -> X
                    | (X,t)::sub -> t
                    | (y,t)::sub -> aux sub
     in aux sub;;

let mapvar = maptm (fun m -> fun n -> App (m, n))
                   (fun r -> Abs (X \ r X));;

let subst_tm sub = mapvar (lookup sub);;

let fv term = maptm union (fun r -> new X in remove X (r X))
                          (fun x -> x::[]) term;;

let size term = maptm (fun x -> fun y -> 1 + x + y)
                      (fun r -> new X in 1 + (r X))
                      (fun x -> 1) term;;

let terminals term = maptm (fun x -> fun y -> x + y)
                           (fun r -> new X in (r X))
                           (fun x -> 1) term;;
```

Fig. 5. Various computations on untyped λ-terms using higher-order programs.

`tm`. In particular, the constructors `App` and `Abs` are replaced by functions `fapp` and `fabs` respectively. In addition, the function `fvar` is applied to all nominals encountered in the term. This higher-order function can be used to define a number of other useful and familiar functions. For example, `mapvar` function is a specialization of the `maptm` function that just applies a given function to all nominals in an untyped λ-term. The application of a substitution (an expression of type `(tm * tm) list`) to a term of type `tm` can then be seen as the result of applying the `lookup` function to every variable in the term (using `mapvar`). Using the functions in Figure 5, the three expressions

```
Abs(X\
 mapvar (fun x -> X) (Abs(U\ Abs(V\ App(U, V))))));;
new X in new Y in
 lookup ((X, Abs(U\U))::(Y, Abs(U\ App(U,U)))::[]) X;;
new X in new Y in
```

```
lookup ((X, Abs(U\U))::(Y, Abs(U\ App(U,U)))::[]) Y;;
```

evaluate to the following three λ-terms.

```
Abs(X\ Abs(Y\ Abs(Z\ App(X, X))))
Abs(X\ X)
Abs(X\ App(X, X))
```

Three additional functions are defined in Figure 5: `fv` constructs the list of free variables in a term; `size` is a re-implementation of the `size` function presented in Section 3; and `terminals` counts the number of variable occurrences (terminal nodes) in its argument.

5 Typing

Given that MLTS is a rather mild extension to OCaml at the syntax level, a typing system for MLTS is simple to present and follows standard practices. Figure 6 contains the rules for typing the new features of MLTS: additional rules for encoding `let` and `let rec` constructions (as well as for built-in types such as integers) must also be added, but these follow the usual pattern. The inference rules in this figure involve the following typing judgments.

$$\Gamma \vdash \texttt{M} : \texttt{A} \qquad \Gamma \vdash \texttt{A} : \texttt{R} : \texttt{B} \qquad \Gamma \vdash \texttt{M} : \texttt{A} \vdash \Delta \qquad \text{open } \texttt{A}$$

In all of these rules, Γ is the usual association between bound variables and a type: in our situation, Γ will associate both variables and nominals to type expressions. (We also assume that the order of pairs in Γ is not important.) The first of these judgments is the usual typing judgment between a program expression `M` and `A`. The second of these judgments is used to type a rewriting rule `R` that has a left-hand side of type `A` and a right-hand side of type `B`. For example, the following typing judgment should be provable.

$$\Gamma \vdash \texttt{tm} : \texttt{Abs(r) -> 1 + (new X in size (r @ X))} : \texttt{int}$$

Since this rule expression is intended to be closed (that is, the variable `r` is quantified implicitly around this rule), the actual value of Γ will not impact this particular typing judgment. The third typing judgment above is used to analyze the left-hand-side of a match rule: in particular, $\Gamma \vdash \texttt{M} : \texttt{A} \vdash \Delta$ holds if during the process of analyzing the pattern `M`, pattern variables are produced (since these are implicitly quantified) and placed into the typing context Δ. For example, the following should be provable.

$$\Gamma \vdash \texttt{Abs(r)} : \texttt{tm} \vdash \{\texttt{r} : \texttt{tm => tm}\}$$

Some of the inference rules in Figure 6 contain premises of the form (open `A`) where `A` is a primitive type. Types for which this judgment holds are called *open types* and are the types of bindings in the `new` and backslash expressions: equivalently, open types can contain nominals. For our purposes here, we can

$$\frac{}{\Gamma, x : C \vdash \texttt{x} : \texttt{C}} \qquad \frac{\Gamma \vdash \texttt{M} : \texttt{A -> B} \qquad \Gamma \vdash \texttt{N} : \texttt{A}}{\Gamma \vdash \texttt{(M N)} : \texttt{B}} \qquad \frac{\Gamma, \texttt{x} : \texttt{A} \vdash \texttt{M} : \texttt{B}}{\Gamma \vdash \texttt{(fun x -> M)} : \texttt{A -> B}}$$

$$\frac{\Gamma, \texttt{X} : \texttt{A} \vdash \texttt{M} : \texttt{B} \qquad \text{open A}}{\Gamma \vdash \texttt{(new X in M)} : \texttt{B}} \qquad \frac{\Gamma, \texttt{X} : \texttt{A} \vdash \texttt{M} : \texttt{B} \qquad \text{open A}}{\Gamma \vdash \texttt{(X \ M)} : \texttt{A => B}}$$

$$\frac{\Gamma \vdash \texttt{r} : \texttt{A1 => ... => An => A} \quad \Gamma \vdash \texttt{t1} : \texttt{A1} \quad \ldots \quad \Gamma \vdash \texttt{tn} : \texttt{An}}{\Gamma \vdash \texttt{(r @ t1 ... tn)} : \texttt{A}}$$

$$\frac{\Gamma \vdash \texttt{term} : \texttt{B} \qquad \Gamma \vdash \texttt{B} : \texttt{R1} : \texttt{A} \qquad \ldots \qquad \Gamma \vdash \texttt{B} : \texttt{Rn} : \texttt{A}}{\Gamma \vdash \texttt{match term with R1 | ... | Rn} : \texttt{A}}$$

$$\frac{\Gamma, \texttt{X} : \texttt{C} \vdash \texttt{A} : \texttt{R} : \texttt{B} \qquad \text{open C}}{\Gamma \vdash \texttt{A} : \texttt{nab X in R} : \texttt{B}} \qquad \frac{\Gamma \vdash \texttt{L} : \texttt{A} \vdash \Delta \qquad \Gamma, \Delta \vdash \texttt{R} : \texttt{B}}{\Gamma \vdash \texttt{A} : \texttt{L -> R} : \texttt{B}}$$

$$\frac{}{\Gamma, \texttt{x} : \texttt{A} \vdash \texttt{x} : \texttt{A} \vdash \cdot} \qquad \frac{\Gamma \vdash \texttt{X1} : \texttt{A1} \ \ldots \ \Gamma \vdash \texttt{Xn} : \texttt{An} \quad \text{open A1} \ldots \text{open An}}{\Gamma \vdash \texttt{(r @ X1 ... Xn)} : \texttt{A} \vdash \texttt{r} : \texttt{A1 => ... => An => A}}$$

$$\frac{\Gamma \vdash \texttt{p} : \texttt{A} \vdash \Delta_1 \qquad \Gamma \vdash \texttt{q} : \texttt{B} \vdash \Delta_2}{\Gamma \vdash \texttt{(p,q)} : \texttt{A * B} \vdash \Delta_1, \Delta_2}$$

$$\frac{\Gamma \vdash \texttt{t1} : \texttt{A1} \vdash \Delta_1 \quad \ldots \quad \Gamma \vdash \texttt{tn} : \texttt{An} \vdash \Delta_n}{\Gamma \vdash \texttt{C(t1,...,tn)} : \texttt{A} \vdash \Delta_1, \ldots, \Delta_n} \ \text{provided } C \text{ is a constructor of type } \texttt{A1 * ... * An -> A}$$

Fig. 6. Typing rules based on the concrete syntax for the new features of MLTS.

assume that every type that is defined in a program (using the `type` command) is presumed to be open. For example, the judgment (open `tm`) needs to be true so that the type `tm => tm` can be formed in the various typing rules. On the other hand, the built-in type for integers `int` should not be considered open in this sense. Clearly a keyword must be added to datatype declarations to indicate if a type is intended as open in this sense.

In the inference rules in Figure 6, whenever we extend the typing context Γ to, say, $\Gamma, \texttt{X} : \texttt{A}$, we always assume that `X` is not declared a type in Γ already. Since α-conversion is always possible within terms, this assumption can always be satisfied. Note that since pattern variables are restricted (as is usual) so that they have at most one occurrence in a given pattern, the union of contexts, in the form $\Delta_1, \ldots, \Delta_n$ never attributes more than one type to the same variable.

The prototype implementation TryMLTS [Gérard, Miller, and Scherer, 2018] of MLTS contains a type inference engine that runs on top of λProlog: given the hypothetical judgments available in λProlog, the implemented typing system is structured differently (but equivalently) to the one given in Figure 6.

6 Abstract syntax, untyped λ-calculus, and arity typing

Although MLTS is designed as a strongly typed functional programming language, evaluation for this language is fundamentally untyped. The *abstract syntax* for MLTS is based on the untyped λ-calculus along with a few extensions to capture the new features of MLTS.

Recall the semantic description of the untyped λ-calculus given by Scott in [Scott, 1970]. Scott was able to present a semantic domain D that was isomorphic to its own function space: that is, $D \equiv [D \to D]$. This equivalence is witnessed by the two continuous mappings $\Phi\colon D \to (D \to D)$ (encoding application) and $\Psi\colon (D \to D) \to D$ (encoding abstraction). For example, the untyped λ-term $\lambda x \lambda y((xy)y)$ is encoded as a value in domain D using the expression $(\Psi(\lambda x(\Psi(\lambda y(\Phi(\Phi\ X\ Y)\ X))))))$.

Note that syntactically, application in the untyped λ-calculus is captured by two domain-level features: function application and the mapping Φ. Similarly, abstraction is captured by two domain-level features: function abstraction (the creation of an element of $[D \to D]$) and the mapping Ψ. We can thus identify two different *syntactic categories* in this encoding: those denoted by the domain D and those identified by the domain of (continuous) functions $D \to D$. In what follows, we need to make a similar distinction between $(\lambda x.T)$ of type $D \to D$ and $(\Psi(\lambda x.T))$ of type D.

To capture this distinction in a more general setting, we employ the notion of *arity typing* that has been used by Martin-Löf [Nordstrom, Petersson, and Smith, 1990]. In particular, we inductively define arity types as follows.

- There is one primitive arity type, written as $\mathbf{0}$.
- If ρ_1 and ρ_2 are arity types then so is $(\rho_1 \to \rho_2)$.
- If $\rho_1, \ldots, \rho_n$ $(n \geq 2)$ are arity types then so is $\rho_1 \otimes \cdots \otimes \rho_n$.

Here, $\mathbf{0}$ formally plays a role in the syntax of expressions that is played by domain D in denotational semantics. As is common practice, the infix arrow $\to$ associates to the right. In the encoding of the untyped λ-calculus, the operator Φ takes two arguments of arity type $\mathbf{0}$ while the operator Ψ takes one argument of arity $\mathbf{0} \to \mathbf{0}$. The arity type constructor $\otimes$ will not be used in our setting except for the possible convenience of writing an arity expressions in an *uncurried* form. In particular, we follow the usual OCaml convention that constructors must have arity $(\rho_1 \otimes \ldots \otimes \rho_n) \to \rho_0$ where ρ_0 is a primitive arity. The abstract syntax of such constructors could well have the curried arity type $\rho_1 \to \ldots \to \rho_n \to \rho_0$.

In most formalizations of ML-style programming languages, expressions of non-zero arity generally only arise in the application of a function to its argument: all other features of the language only take arguments of arity type $\mathbf{0}$. In MLTS, expressions of non-zero arity play extended roles: for example, in MLTS, pattern matching variables can have non-primitive arity while in most ML-languages, pattern variables are always of primitive arity. It is important to keep arity typing and ML-style typing separated. For example, the type of `subst` in Section 4 can be inferred to be `(tm => tm) -> tm -> tm`. The arity typing of `subst` is, however, the simple expression $(\mathbf{0} \to \mathbf{0}) \to \mathbf{0} \to \mathbf{0}$: that is, the first argument

given to `subst` must be a binding at the level of the abstract syntax. As we shall see in the following sections, the arity typing is used in the specification of the operational semantics of MLTS.

7 Formalizing the design of MLTS

Bindings are such an intimate part of the nature of syntax that we should expect that our high-level programming languages accounts for them directly in, for example, any built-in notion of equality or matching. (The paper [Miller, 2018] contains an extended argument of this point in the setting of logic programming and proof assistants.) Another reason to include binders as a primitive within a functional programming languages is that their semantics have a well understood declarative and operational treatment. For example, Church's higher-order logic STT [Church, 1940] contains an elegant integration of bindings in both terms and formulas. His logic also identifies equality for both terms and formulas with $\alpha\beta\eta$-conversion. Church's integration is also a popular one in theorem proving—being the core logic of the Isabelle [Paulson, 1994], HOL [Harrison, 2009, Gordon, 1991], and Abella [Baelde, Chaudhuri, Gacek, Miller, Nadathur, Tiu, and Wang, 2014] theorem provers—as well as the logic programming language λProlog [Miller and Nadathur, 2012]. Given the existence of these provers, a good literature now exists that describes how to effectively implement STT and closely related logics. Below, we describe what that literature can tell us about the meaning and implementation of the novel features of MLTS.

7.1 Equality modulo α, β, η conversion

The abstract syntax behind MLTS is essentially a simply typed λ-term where the types are identified with arity types over the primitive arity $\mathbf{0}$ and the binary arity constructor $\to$. Furthermore, the equality theory of such terms is given by the familiar α, β, η conversion rules. As a result, a programming language that adopts this notion of equality cannot take an abstraction and return, say, the name of its bound variable: since that name can be changed via the α-conversion, such an operation would not be a proper function. Thus, it is not possible to decompose the untyped λ-term $\lambda x.t$ into the two components x and t. Not being able to retrieve a bound variables name might appear as a serious deficiency but, in fact, it can be a valuable feature of the language: for example, a compiler does not need to maintain such names and can choose any number of different, low-level representations of bindings to exploit during execution. Since the names of bindings seldom have semantically meaningful value, dropping them entirely is an interesting design choice. That choice is similar to one taken in ML-style languages in which the location in memory of a reference cell is not maintained as a value in the language.

The relation of λ-conversion is invoked when evaluating the expression `(t @ s1 ... sn)`. If we assume that expressions $\mathtt{s1}, \ldots, \mathtt{sn}$ have arities $\rho_1, \ldots, \rho_n$, respectively, then `t` must have arity $\rho_1 \to \cdots \to \rho_n \to \mathbf{0}$. Thus, `t` is η-equivalent

to a term with n abstractions, for example, `X1\...Xn\ t'` and the value of the expression `(t @ s1 ... sn)` is the result of performing λ-normalization of `(X1\...Xn\ t')` to the arguments `s1`, ..., `sn`.

As we illustrated in Section 3, it is possible to implement both substitution and λ-conversion in MLTS. Thus, it is possible to limit the occurrences of `@` to appear only within the scope of match clauses and only then with a pattern variable as the first argument of `@`. For the sake of the rest of this paper, we will not enforce that restriction.

7.2 Pattern unification and matching

Since we are not able to decompose bindings into their bound variable and body, we need to find alternative means for analyzing the structure of terms containing bindings. As our earlier examples illustrated, matching within patterns can be used to probe terms and their bindings. If we do not place restrictions on the use of pattern variables, then patterns can have complex behaviors.

No repeated pattern variable occurrences. We impose a familiar restriction on the match rules: a pattern variable cannot have more than one occurrence within a match pattern. The main reason this is done in ML-style languages is that it relieves pattern matching from the need to check equality of terms. Since terms can be large, pattern matching could involve a costly recursive descent of terms. It is far more common to forbid repeated occurrences of pattern variables and force the programmer to insert equality checking outside the pattern matching operation. Thus, instead of defining `memb : tm -> tm list -> bool` with the following code using a repeated match variable

```
let rec memb x l = match (x,l) with
    | (x,[])      -> false
    | (x,(x::l)) -> true
    | (y,(x::l)) -> memb x l;;
```

we can require the programmer to write an equality predicate for type `tm` and then rewrite the program above as follows.

```
let rec eqtm t s = match (t,s) with
    | (App(m1,m2),
       App(n1,n2))     -> eqtm m1 n1 && eqtm m2 n2
    | (Abs r, Abs s)   -> new X in eqtm (r @ X) (s @ X)
    | nab X in (X, X) -> true
    | _                -> false;;

let rec memb x l = match (x,l) with
    | (x,[])      -> false
    | (x,(y::l)) -> if (eqtm x y) then true
                                  else (memb x l);;
```

Given the definition of the `tm` datatype, it is clear that a compiler for MLTS could define its own equality predicate for this type. In that case, repeated variable occurrences in patterns could be allowed since resolving such patterns could be done using these equality predicates.

Restricted use of higher-order pattern variables. Since pattern variables within match rules can have higher-order arity (and higher-order types), occurrences of those variables within patterns need to be restricted: otherwise, undesirable features of higher-order matching could appear. Fortunately, there is a natural restriction on occurrences of pattern variables that guarantees that a match either fails or succeeds with at most one solution. That restriction is the following: every occurrence of an expression of the form `(r @ X1 ... Xn)` in the left-hand side of a match rule must be such that the pattern variable `r` is applied to $n \geq 0$ *distinct* nominals `X1 ... Xn` and those nominals are bound *within* the scope of the binding for `r`. For example, the following expression is not well formed

```
Abs(X\ (match Abs(Y\ App(X,Y)) with
            | Abs(Z\ r @ Z X) -> Abs(Z\ r @ X Z)))
```

since the scope of the nominal `X` contains the (implicit) scope of the pattern variable `r`, which is around the rule `(Abs(Z\ r @ Z X) -> Abs(Z\ r @ X Z))`.

This restriction can be motivated within a purely logical setting as follows. Let j be a primitive type and let $f : j \to j \to j$ be a simply typed constant. The formula $\exists G : j \to j\ \forall x : j\ [G\ x = (f\ x\ x)]$ has a unique proof in which G is instantiated by the term $\lambda w.(f\ w\ w)$. Note that the binding scope of the variable x is inside the binding scope of the variable G. If, however, one switches the order of the quantifiers, yielding $\forall x : j\ \exists G : j \to j\ [G\ x = (f\ x\ x)]$, then there are four different proofs of this equation: if one replaces the outermost universal quantifier with an eigenvariable, say a, then there are four different solutions for G, namely, $\lambda w.(f\ a\ a)$, $\lambda w.(f\ a\ w)$, $\lambda w.(f\ w\ a)$, and $\lambda w.(f\ w\ w)$.

The subset of higher-order unification in which unification variables (a.k.a., logic variables, meta-variables, pattern variables) are applied to distinct bound variables restricted as described above, is called *higher-order pattern unification* or L_λ *unification* [Miller, 1991]. (Nipkow provides a functional programming implementation of such unification in [Nipkow, 1993].) This particular subset of higher-order unification is commonly implemented in theorem provers such as Abella [Baelde, Chaudhuri, Gacek, Miller, Nadathur, Tiu, and Wang, 2014], Minlog [Schwichtenberg, 2006], and Twelf [Pfenning and Schürmann, 1999] as well as recent implementations of λProlog [Dunchev, Guidi, Coen, and Tassi, 2015, Qi, Gacek, Holte, Nadathur, and Snow, 2015].

The following results about about higher-order pattern unification are proved in [Miller, 1991].

1. It is decidable and unitary, meaning that if there is a unifier then there exists a most general unifier.

2. It does not depend on typing (or on arity). As a result, it is possible to add it to the evaluator for MLTS based on untyped terms.
3. The only form of β-conversion that is needed to solve such unification problems is what is called β_0-conversion which is a form of the β rule that equates $(\lambda x.t)x$ with t.

An equivalent way to write the β_0-conversion rule (assuming the presence of α-conversion) is that $(\lambda x.t)y$ converts to $t[y/x]$ *provided* that y is not free in $\lambda x.t$. Notice that applying β_0 reduction actually makes a term smaller and does not introduce new β redexes: as a result it is not a surprise that such unification (and, hence, matching) has low computational complexity (the paper [Qian, 1996] claims that such unification is, in fact, solvable in linear time).

All `nab` bound variables must have a rigid occurrence. There is an additional restriction on match rules that is associated to the `nab` quantifiers that appear in such rules. We say that an occurrence of a `nab`-quantified nominal is *flexible* if it is in the scope of an `@`. For example, in the code

```
Abs(X\ (match Abs(Y\ App(X,Y)) with
            | nab W in Abs(Z\ r @ Z W) ->
                       Abs(Z\ r @ W Z)));;
```

the nominal binding `W` has two occurrences that are flexible: one each within `(r @ Z W)` and `(r @ W Z)`. All other occurrences of a `nab` quantified nominal is *rigid*. For example, in the match rule `| nab X in X -> 1`, `X` has a binding occurrence and a rigid occurrence. In the auxiliary function used by the `insert` function in Figure 4, namely,

```
let rec aux c x k = match (x, k) with
    | nab X   in (X, X::(l @ X))   -> c
    | nab X Y in (X, Y::(l @ X Y)) ->
                                aux (c + 1) x (l @ X Y)
```

the nominals `X` and `Y` have both rigid and flexible occurrences within their scope.

The one additional restriction that we need is the following: every `nab` quantified variable must have at least one rigid occurrence in the left part of the match rule (the pattern) that falls within the scope of its binder. For example, the code listed above (for an expression of type `tm`) does not satisfy this restriction since every occurrence of of `W` in the pattern is flexible (there is just one such occurrence). The necessity of this restriction can be seen when we consider a pattern of the form

```
| nat X Y in  (r @ X Y) -> term
```

In the event that a nominal, say `U`, is matched with the pattern in this rule, there are two possible instantiations for `r` that could succeed, namely, the terms `X\Y\X` and `X\Y\Y`: we wish to avoid multiple successful matches of the same rule. The following clause is also ruled out by this restriction

```
| nat X in  1 -> X
```

since `X` has no rigid occurrence in the expression `1`. Discarding this match rule makes sense since the nominal that is returned as the result of this match is not constrained by the input to the match.

7.3 β_0 versus β

In order to ensure that matching a rule either fails or has a unique, most general solution, we will insist that in the left side of a match rule, all subexpressions of the form `(r @ X1 ... Xn)` are such that the scope of the binding for `r` contains the scope of the bindings for the distinct variables in `X1`, ..., `Xn`. On the right-hand side of a match rule, however, it seems that one has an interesting choice. If on the right, we have an expression of the form `(r @ t1 ... tn)` then clearly, the terms `t1`, ..., `tn` are intended to be substituted into the abstraction that is instantiated for the pattern variable `r`: that is, we need to use β-conversion on this redex. One choice is that we restrict the terms `t1`, ..., `tn` to be distinct nominals just as on the left-hand-side: in this case, β-reduction of the expression `(r @ t1 ... tn)` requires only β_0 reductions. A second choice is that we allow the terms `t1`, ..., `tn` to be unrestricted: in this case, β-reduction of the expression `(r @ t1 ... tn)` requires more general (and costly) β-reductions.

A similar trade-off between allowing β-conversion or just β_0 conversion has also been studied within the theory and design of the π-calculus. In particular, the full π-calculus allows the substitution of arbitrary names into input prefixes (modeled by β-conversion) while the π_I-calculus (π-calculus with internal mobility [Sangiorgi, 1996]) is restricted in such a way that the only instances of β-conversions are, in fact, β_0-conversions.

Another reason to identify the β_0 fragment of β-conversion is that β_0 reduction provides support for binder mobility and it can be given effective implementations, sometime involving only constant time (see Section 10.2).

7.4 Match rule quantification

Match rules in MLTS contain two kinds of quantification. The familiar quantification of pattern variables can be interpreted as being universal quantifiers. For example, the first rule defining the `size` function in Section 3, namely,

```
| App(n, m)  -> 1 + size n + size m
```

can be encoded as the logical statement

$$\forall \texttt{m} \forall \texttt{n}[\texttt{(size (App(n, m))) = 1 + size n + size m}].$$

On the other hand, the third match rule for `size` contains the binder `nab`

```
| nab X in X -> 1
```

which corresponds approximately to the *generic* ∇-quantifier (pronounced nabla) that is found in various efforts to formalize the metatheory of computational systems [Miller and Tiu, 2005, Baelde, Chaudhuri, Gacek, Miller, Nadathur, Tiu, and Wang, 2014]. That is, this rule can be encoded as the quantified equation ∇`x`.(`size x` = `1`): that is, the size of a nominal constant is 1.

Although there are two kinds of quantifiers around such match rules, the ones corresponding to the universal quantifiers are implicit while the ones corresponding to the ∇-quantifiers are explicit. Our design for MLTS places the implicit quantifiers at outermost scope: that is, the quantification over a match rule is of the form $\forall\nabla$. Another choice might be to allow some (all) universal quantifiers to be explicitly written and placed among any `nab` bindings. While this is a sensible choice, the $\forall\nabla$-prefixes is, in fact, a reduction class in the sense that if one has a $\forall$ quantifier inside a ∇-quantifier, it is possible to rotate that ∇-quantifier inside using a technique called *raising* [Miller, 1991, Miller and Tiu, 2005]. That is, the formula $\nabla x : \gamma \forall y : \tau(Bxy)$ is logically equivalent to the formula $\forall h : (\gamma \to \tau)\nabla x : \gamma(Bx(hx))$: note that as the ∇-quantifier of type γ is moved to the right over a universal quantifier, the type of that quantifier is raised from τ to $\gamma \to \tau$. Thus, it is possible for an arbitrary mixing of $\forall$ and ∇ quantifiers to be simplified to be of the form $\forall\nabla$.

7.5 Nominal abstraction

Before we can present the formal operational semantics of MLTS, we need to introduce one final logical concept: *nominal abstraction* which allows implicit bindings represented by nominals to be moved into explicit abstractions over terms [Gacek, Miller, and Nadathur, 2011]. The following notation is useful for defining this relationship.

Let t be a term, let $c_1, \ldots, c_n$ be distinct nominals that possibly occur in t, and let $y_1, \ldots, y_n$ be distinct variables not occurring in t and such that, for $1 \leq i \leq n$, y_i and c_i have the same type. Then we write $\lambda c_1 \ldots \lambda c_n.t$ to denote the term $\lambda y_1 \ldots \lambda y_n.t'$ where t' is the term obtained from t by replacing c_i by y_i for $1 \leq i \leq n$. There is an ambiguity in this notation in that the choice of variables $y_1, \ldots, y_n$ is not fixed. However, this ambiguity is harmless: the terms that are produced by acceptable choices are all equivalent under a renaming of bound variables.

Let $n \geq 0$ and let s and t be terms of type $\tau_1 \to \cdots \to \tau_n \to \tau$ and τ, respectively; notice, in particular, that s takes n arguments to yield a term of the same type as t. The formula $s \trianglerighteq t$ is a *nominal abstraction of degree* n (or, simply, a *nominal abstraction*). The symbol $\trianglerighteq$ is used here in an overloaded way in that the degree of the nominal abstraction it participates in can vary. The nominal abstraction $s \trianglerighteq t$ of degree n is said to hold just in the case that s is λ-convertible to $\lambda c_1 \ldots c_n.t$ for some distinct nominals $c_1, \ldots, c_n$.

Clearly, nominal abstraction of degree 0 is the same as equality between terms based on λ-conversion, and we will use $=$ to denote this relation in that case. In the more general case, the term on the left of the operator serves as a pattern for isolating occurrences of nominals. For example, if p is a binary constructor

$$\frac{\vdash val\ V}{\vdash V \Downarrow V} \qquad \frac{\vdash M \Downarrow F \quad \vdash N \Downarrow U \quad \vdash apply\ F\ U\ V}{\vdash M@N \Downarrow V} \qquad \frac{\vdash (R\ (fixpt\ R)) \Downarrow V}{\vdash (fixpt\ R) \Downarrow V}$$

$$\frac{\vdash C \Downarrow tt \quad \vdash L \Downarrow V}{\vdash cond\ C\ L\ M \Downarrow V} \qquad \frac{\vdash C \Downarrow ff \quad \vdash M \Downarrow V}{\vdash cond\ C\ L\ M \Downarrow V}$$

$$\frac{\vdash M \Downarrow U \quad \vdash (R\ U) \Downarrow V}{\vdash (let\ M\ R) \Downarrow V} \qquad \frac{\vdash (R\ U) \Downarrow V}{\vdash apply\ (\text{lam}\ R)\ U\ V}$$

$$\frac{\vdash \nabla x.(E\ x) \Downarrow (V\ x)}{\vdash \lambda x.Ex \Downarrow \lambda x.Vx} \qquad \frac{\vdash \nabla x.(E\ x) \Downarrow V}{\vdash new\ E \Downarrow V}$$

$$\frac{\vdash \text{pattern}\ T\ Rule\ U \quad \vdash U \Downarrow V}{\vdash (\text{match}\ T\ (Rule :: Rules)) \Downarrow V} \qquad \frac{\vdash (\text{match}\ T\ Rules) \Downarrow V}{\vdash (\text{match}\ T\ (Rule :: Rules)) \Downarrow V}$$

$$\frac{\vdash \exists x.\text{pattern}\ T\ (P\ x)\ U}{\vdash \text{pattern}\ T\ (all\ \lambda x.P\ x)\ U} \qquad \frac{\vdash (\lambda z_1 \ldots \lambda z_m.(t \Longrightarrow s)) \trianglerighteq (T \Longrightarrow U)}{\vdash \text{pattern}\ T\ (nab\ z_1 \ldots nab\ z_m.(t \Longrightarrow s))\ U}$$

Fig. 7. A natural semantic specification of evaluation.

and c_1 and c_2 are nominals, then the nominal abstractions of the first row below hold while those in the second row do not.

$\lambda x.x \trianglerighteq c_1$	$\lambda x.p\ x\ c_2 \trianglerighteq p\ c_1\ c_2$	$\lambda x.\lambda y.p\ x\ y \trianglerighteq p\ c_1\ c_2$
$\lambda x.x \not\trianglerighteq p\ c_1\ c_2$	$\lambda x.p\ x\ c_2 \not\trianglerighteq p\ c_2\ c_1$	$\lambda x.\lambda y.p\ x\ y \not\trianglerighteq p\ c_1\ c_1$

A logic with equality generalized to nominal abstraction has been studied in [Gacek, 2009, Gacek, Miller, and Nadathur, 2011] where a logic, named $\mathcal{G}$, that contains fixed points, induction, coinduction, ∇-quantification, and nominal abstraction is given a sequent calculus presentation. Cut-elimination for $\mathcal{G}$ is proved in [Gacek, Miller, and Nadathur, 2011] and algorithms and implementations for nominal abstraction are presented in [Gacek, 2009, Wang, Chaudhuri, Gacek, and Nadathur, 2013]. An important feature of the Abella prover—∇ in the head of a definition—can be explained and encoded using nominal abstraction [Gacek, Miller, and Nadathur, 2008].

8 Natural semantic specification of MLTS

We can now define the operational semantics of MLTS by giving inference rules in the style of natural semantic (a.k.a. big-step semantic) following Kahn [Kahn, 1987]. The semantic definition for the core of MLTS is defined in Figure 7. Since those inference rules are written using an abstract syntax for MLTS, we need to describe briefly how that abstract syntax is derived from the concrete syntax we have presented for our several examples.

```
prog "size" (fixpt size\ lam term\ match term
   [(all n\ all m\ ((app n m) ==>
       (sum @ (i 1) @ (sum @ (size @ n) @ (size @ m)))))),
    (all' r\        ((abs r)    ==>
       (sum @ (i 1) @ (new x\ size @ (r x))))),
    (nab x\         (x          ==> (i 1)))]).
```

Fig. 8. The abstract syntax of the `size` program.

Instead of detailing the translation from concrete to abstract syntax, we illustrate this translation with an example. There is an implementation of MLTS that includes a parser and a transpiler into λProlog code: this system is available for online use and for downloading at https://trymlts.github.io [Gérard, Miller, and Scherer, 2018]. The abstract syntax used in this section and in that implementation are different in detail. For presentation purposes, we use a simplified version of the abstract syntax of MLTS. For example, the λProlog code in Figure 8 is the abstract syntax for the MLTS program for `size` given in Section 3.

The backslash (as infix notation) is also used in λProlog to denote binders; the `@` denotes the untyped λ-calculus application; `lam` denotes the untyped λ-calculus abstraction; recursive function definitions are encoded using the `fixpt` operator; the infix symbol `==>` denotes a match rule; `nab` denotes the `nab`-quantifier; and `all` and `all'` are explicit universal quantification bindings for variables of arity $\mathbf{0}$ and $\mathbf{0} \to \mathbf{0}$, respectively.[1] Finally, the concrete syntax `(let x = t in s)` is translated to the abstract syntax `(let (x\ s) t)`.

It is intended that the inference rules given in Figure 7 are, in fact, notations for formulas in the logic $\mathcal{G}$. For example, schema variables of the inference figure are universally quantified around the intended formula; the horizontal line is an implication; the list of premises is a conjunction; and $\Downarrow$ is a binary (infix) predicate, etc. Some features of $\mathcal{G}$ are exploited by some of those inference rules: those features are enumerated below.

In the rules for `fixpt`, `let`, and `apply`, a variable of arity type $\mathbf{0} \to \mathbf{0}$ (namely, R) is applied to a term of arity type $\mathbf{0}$. These rules make use of the underlying equality theory of simply typed λ-terms in $\mathcal{G}$ to perform a substitution. In the rule for apply, for example, if R is instantiated by the term $\lambda w.t$ and U is instantiated by the term s, then the expression written as $(R\ U)$ is equal (in $\mathcal{G}$) to the result of substituting s for the free occurrences of w in t: that is, to the result of a β-reduction on the expression $((\lambda w.t)\ s)$.

Existential quantification is written explicitly into the first rule for patterns. It is possible (as is done in other rules) to drop the explicit existential quantifier and instead have the quantification be implicitly universally quantified around the entire rule. We write it explicitly here to highlight the fact that solving the problem of finding instances of pattern variables in matching rules is lifted to

[1] The abstract syntax used in the `tryMLTS` implementation contains only one such quantifier: pattern variables of any arity can be encoded using `all`.

the general problem of finding substitution terms in $\mathcal{G}$. Also, the arity of the existentially quantified variable in that inference rule can range over $\mathbf{0}$, $\mathbf{0} \to \mathbf{0}$, $\mathbf{0} \to \mathbf{0} \to \mathbf{0}$, ...

The proof rules for natural semantics are nondeterministic in principle. Consider attempting to prove that t, a term of arity type $\mathbf{0}$, has a value: that is, $\exists V, t \Downarrow V$. It can be the case that no proof exists or that there might be several proofs with different values for V. No proofs are possible if, for example, the condition in a conditional phrase does not evaluate to a boolean or if there are insufficient match rules provided to cover all the possible values given to a match expression. Similarly, if there are overlapping matching rules, it is possible to have multiple values computed. Ultimately, we will want to provide a static check that could issue a warning if the rules listed in a match expression are not exhaustive. Making sure that only the first successful match in a list of matching rules is selected is easily achieved in the implementation of natural semantics: for example, we added a single Prolog-cut operator into our λProlog implementation of the natural semantics of MLTS to satisfy that restriction.

The *nominal abstraction* of $\mathcal{G}$ is directly invoked to solve pattern matching in which nominals are explicitly abstracted using the `nab` binding construction. When attempting to prove the judgment $\vdash$ pattern T *Rule* U, the inference rules in Figure 7 eventually lead to an attempt to prove in $\mathcal{G}$ an existentially quantified nominal abstraction of the form

$$\exists x_1 \ldots \exists x_n [(\lambda z_1 \ldots \lambda z_m.(t \Longrightarrow s)) \trianglerighteq (T \Longrightarrow U)].$$

Here, the arrow $\Longrightarrow$ is simply a formal (syntactic) pairing operator. The schema variables $x_1, \ldots, x_n$ can be of arity $\mathbf{0}$, $\mathbf{0} \to \mathbf{0}$, $\mathbf{0} \to \mathbf{0} \to \mathbf{0}$, ... and can appear free only in t and s: furthermore, if any of these variables are free in s they must be free in t. Also, if any of the variables $z_1, \ldots, z_m$ (all of which are assumed to have arity $\mathbf{0}$) are free in s they are also free in t. While the variables $x_1, \ldots, x_n$ cannot appear more than once in t, the variables $z_1, \ldots, z_m$ are not restricted in this fashion. In order to prove the formula $\exists \bar{x}(\lambda \bar{z}.t) \trianglerighteq s$, one must find a collection of distinct nominals $\bar{c}$ and witness terms $\bar{t}$ that do not contain any of the elements of $\bar{c}$ such that $[\bar{t}/\bar{x}, \bar{c}/\bar{z}]t = s$ [Wang, Chaudhuri, Gacek, and Nadathur, 2013].

It is worth pointing out that given the way we have defined the operational semantics of MLTS, it is immediate that "nominals cannot escape their scopes." For example, the expression `(new X in X)` does not have a value (in abstract syntax, this expression translates to `(new X\ X)`). More precisely, there is no proof of $\vdash \exists v.(new \lambda x.x) \Downarrow v$ using the inference rules in Figure 7. To understand why this is an immediate consequence of the specification of evaluation, consider the formula (which encodes the inference rule in Figure 7 defining *new*)

$$\forall E \forall V [(\nabla x.(E\ x) \Downarrow V) \supset (new\ E \Downarrow V)].$$

Given that the scope of the ∇x is inside the scope of $\forall V$, it is not possible for any instance of this formula to allow the x binder to appear as the second argument of the $\Downarrow$ predicate. While such escaping is easily ruled out using this logical specification, a direct implementation of this logic must incur a cost, however,

to constantly ensure that no escaping is permitted. (See Section 10 for more discussion on this point.)

9 Binder mobility

We started this programming language project with the desire to treat binders in syntax as directly and naturally as possible. We approached this project by designing the MLTS language with more binders than, say, OCaml: it has not only the usual binders for building functions and for refactoring computation (via the `let` construction) but also new binders that are directly linked to binders in data (via the `new X in`, `nab X in`, and `X\` operators). Finally, the natural semantics of MLTS in $\mathcal{G}$ and its implementation in λProlog are all based on using logics that contain rich binding operators that go beyond the usual universal and existential quantifiers. It is worth noting that if one were to write MLTS programs that do not need to manipulate data structures containing bindings, then the new binding features of MLTS would not be needed and neither would the novel features of both $\mathcal{G}$ and λProlog. Thus, in a sense, binders have not been formally implemented in this story: instead, binders of one kind have been implemented and specified using binders in another system. We were able to complete a prototype implementation of MLTS since we know how to implement the high-level logics, and those techniques can be applied directly to the natural semantic specification.

One way to view the processing a binder is that one needs to first *open* the abstraction, process the result (by "freshening" the newly freed names), and then *close* the abstraction [Pottier, 2006]. In the setting of MLTS, it is better to view such processing as the *movement* of a binder: that is, the binder in a data structure actually gets re-identified with an actual binder in the programming language. As we illustrated in Section 3 with the following step-by-step evaluation

```
size (Abs (X\ (Abs (Y\ (App(X,Y))))));;
new X in 1 + (size (Abs (Y\ (App(X,Y)))));;
new X in 1 + new Y in 1 + (size (App(X,Y)));;
new X in 1 + new Y in 1 + 1 + (size X) + (size Y);;
new X in 1 + new Y in 1 + 1 + 1 + 1;;
```

the bound variable occurrences for `X` and `Y` simply move. It is never the case that a bound variable actually becomes free: instead, it just becomes bound elsewhere. Thus, our strategy for strengthening the expressiveness of MLTS over other ML-style languages has been to add to the language more binding sites to which bindings can move.

10 Interpreters for MLTS

We have a prototype implementation of MLTS. A parser from our extended OCaml syntax and a transpiler that generates λProlog code are implemented in OCaml. A simple evaluator and type checker written in λProlog can then

be used to type check and execute MLTS code. The implementation of the evaluator in λProlog is rather compact but not completely trivial since neither ∇-quantification nor nominal abstraction are native to λProlog: they needed to be implemented. Both the Teyjus [Qi, Gacek, Holte, Nadathur, and Snow, 2015] and the Elpi [Dunchev, Guidi, Coen, and Tassi, 2015] implementations of λProlog can be used to execute the MLTS interpreter.

The `TryMLTS` web site [Gérard, Miller, and Scherer, 2018] provides a means for anyone with a recent web browser to create and execute MLTS programs online without needing to install any software. Since Elpi, the parser, and the transpiler are written in OCaml, web-based execution was made possible by compiling the OCaml bytecode to a Javascript client library with `js_of_ocaml` [js-of-ocaml].

There is little about this prototype implementation that is focused on providing an efficient implementation of MLTS. Instead, the prototype is a useful device for exploring the exact meaning and possible uses of the new program features. Never-the-less, we can comment here briefly on some costs of the underlying system that will likely appear in any implementation of MLTS.

10.1 Nominal-escape checking

As we have mentioned in Section 8, nominals are not allowed to escape their scope during evaluation and quantifier alternation can be used to enforce this restriction at the logic level. When one implements the logic, one needs to implement (parts of) the unification of simply typed λ-terms [Huet, 1975] and such unification is constantly checking that bound variable scopes are properly restricted. There are times, however, when the expensive check for escaping nominals are not, in fact, needed. In particular, it is possible to rewrite the inference rule in Figure 7 for the `new` binding operator as the following rule.

$$\frac{\vdash \nabla x.(E\ x) \Downarrow (U\ x) \qquad U = \lambda x.V}{\vdash new\ E \Downarrow V}$$

Here, both U and V are quantified universally around the inference rule. Attempting a proof of the first premise can result in the construction of some (possibly large) value, say t such that $\vdash (E\ x) \Downarrow t$ holds. We can immediately form the binding of $U \mapsto \lambda x.t$ without checking the structure of t. The second premise is where the examination of t may need to take place: if x is free in t, then there is no substitution for V that makes $\lambda x.t$ equal to $\lambda x.V$. This check can be expensive, of course, since one might in principle need to examine the entire structure of t to solve this second premise. There are many situations, however, where such an examination is not needed and they can be revealed by the typing system. For example, if the type of U is, say, `tm => int`, there should not be any possible way for an untyped λ-term to have an occurrence inside an integer. Furthermore, there are static methods for examining type declarations in order to describe if a type $\tau_1 \to \tau_2$ (for primitive types τ_1 and τ_2) can be inhabited by only vacuous λ-terms (see, for example, [Miller, 1992]). Of course, if the types of τ_1 and τ_2 are the same (say, `tm`), then type information is not useful here and

a check of the entire structure t might be necessary. Other static checks and program analyses might be possible as a way to reduce the costs of checking for escaping nominals: the paper [Pottier, 2007] includes such static checks albeit for a technically different functional programming language, namely FreshML [Shinwell, Pitts, and Gabbay, 2003].

10.2 Costs of moving binders

As we have mentioned before, binders are able to move from, say, a term-level binding to a program-level binding by the use of β_0. In particular, if y is a binder that does not appear free in the abstraction $\lambda x.B$ then the β_0 reduction of $(\lambda x.B)y$ causes the x binding in B to move and to be identified with the y binder in $B[y/x]$. If one must actually do the substitution of y for x in B, a possibly large term (at least its spine) must be copied. However, there are some situations where this movement of a binding can be inexpensive. For example, consider again the following match rule for `size`.

```
| Abs(r) -> 1 + (new X in size (r @ X))
```

If we assume that the underlying implementation of terms use De Bruijn's nameless dummies, it is possible to understand the rewriting needed in applying this match clause to be a constant time operation. In particular, if `r` is instantiated with an abstraction then it's top-level constructor would indicate where a binder of value 0 points. If we were to compile the syntax `(r @ X)` as simply meaning that that top-level constant is stripped away, then a binder of value 0 in the resulting term would automatically point (move) to being bound by the `new X` binder. While such a treatment of binder mobility without doing substitution is possible in many of our examples, it does not cover all cases. In general, a more involved scheme for implementing binder mobility must be considered. This kind of analysis and implementation of binder mobility is used in the ELPI implementation of λProlog [Dunchev, Guidi, Coen, and Tassi, 2015].

11 Future work

There is clearly much more work to do. While the examples presented in this paper illustrate that the new features in MLTS can provide elegant and direct support for computing with binding structures, we plan to develop many more examples. The general area of theorem proving implementation and compiler construction is an early target for us. A more effective implementation is also something we wish to target soon. It seems likely that we will need to consider extensions to the usual abstract machine models for functional programming in order to get such a direct implementation.

The cost of basic operations in MLTS must also be understood better. As we noted in Section 3, we could design pattern matching in clauses in such a way that they might require the recursive descent of entire terms in order to know if a match was successful. The language could also be designed so that such a

costly check is never performed during pattern matching: for example, one could insist that every pattern variable is `@`-applied to a list of *all* nominal abstractions that are in the scope of the binding for that pattern variable. In that case, a recursive descent of terms is not needed.

Given the additional expressivity of MLTS, the usual static checks used to produce warnings for non-exhaustive matchings are missing cases that we should add. As mentioned in Section 10, still other static checks are needed to help a future compiler avoid making costly checks.

It would also be interesting to see to what extent binders might interact with a range of non-functional features like references found in languages such as OCaml. A natural starting point to explore the possible interaction of effectful features would be to use a natural semantic treatment based on linear logic (see, for example, [Chirimar, 1995]): the logical features of $\mathcal{G}$ should also work well in a linear logic setting.

Finally, the treatment of syntax with bindings generally leads to the need to manipulate contexts and association lists that relate bindings to other bindings, to types, or to bits of code. We have already seen association lists used in Figure 3. It seems likely that more sophisticated MLTS examples will require singling out contexts for special treatment. Although the current design of MLTS does not commit to any special treatment of context, we are interested to see what kind of treatment will actually prove useful in a range of applications.

12 Related work

The term *higher-order abstract syntax (HOAS)* was introduced in [Pfenning and Elliott, 1988] to describe an encoding technique available in λProlog. A subsequent paper identified HOAS as a technique "whereby variables of an object language are mapped to variables in the metalanguage" [Pfenning and Schürmann, 1999]. When applied to functional programming, this latter description of HOAS describes the mapping of bindings in syntax to the bindings which create functions. Unfortunately, this encoding technique often lacks adequacy (since "exotic terms" can appear [Despeyroux, Felty, and Hirschowitz, 1995]), and structural recursion can slip away [Gabbay and Pitts, 1999]. The term *λ-tree syntax* was introduced in [Miller and Palamidessi, 1999] and the term *binder mobility* was introduced in [Miller, 2004] to describe the different approach that we have used here.

The ML_λ [Miller, 1990] extension to ML is similar to MLTS in that it also contained two different arrow type constructors (`->` and `=>`) and pattern matching was extended to allow for pattern variables to be applied to a list of distinct bound variables. The `new` operator of MLTS could be emulated by using the backslash operator and a "discharge" function. Critically missing from that language was anything similar to the `nab` binding of MLTS. Also, no formal specification and no implementation were ever offered.

Nominals and nominal abstraction, in the sense used in this paper, were first conceived, studied, and implemented as part of the Abella theorem prover [Baelde, Chaudhuri, Gacek, Miller, Nadathur, Tiu, and Wang, 2014]. Although

the design of Abella does not use the $\trianglerighteq$ relation directly, the notion of "∇ in the head" of definitions is essentially equivalent to having the $\trianglerighteq$ relation in the logic.

The Delphin [Poswolsky and Schürmann, 2008] and Beluga [Pientka and Dunfield, 2010] computer systems provide functional programming support for object-level terms that are taken from the dependently typed λ-calculus LF. Both of these systems are rather ambitious and make many extensions to the core of the ML family of programming languages. For example, these languages have well established notions of contexts and proofs (as dependently typed λ-terms) that are part of their programming language's design and evaluation. Our approach here has been much more minimal and incremental.

The FreshML [Shinwell, Pitts, and Gabbay, 2003] and CαML [Pottier, 2006] functional programming languages provide an approach to names based on nominal logic [Pitts, 2003]. In a sense, these two programming languages provide for an abstract treatment of names and naming. Once naming is available, binding structures can also be implemented. In a sense, the design of these two ML-variants are also more ambitious than the design goal intended for MLTS: in the latter, we were not focused on naming but just bindings.

The recent paper [Ferreira and Pientka, 2017] introduces a syntactic framework that treats bindings as primitives. That framework is then integrated with various tools and with the framework of contextual types (similar to that found in Beluga) in order to provide a programmer of, say, OCaml with sophisticated tools for the manipulation of syntax and binders. A possible future target for MLTS could be to provide some aspects of such tools more directly in the language itself.

13 Conclusion

While the λ-tree syntax approach to computing with syntax containing bindings has been successfully developed within the logic programming setting (in particular, in λProlog and Twelf), we provide in this paper another example of how binding can be captured in a functional programming language. Most of the expressiveness of MLTS arises from its increased use of program-level binding. The sophistication needed to correctly exploit binders and quantifiers in MLTS is a skill most people have learned from using quantification in, for example, predicate logic.

We have presented a number of MLTS programs and we note that they are both natural and unencumbered by concerns about managing bound variable names. We have also presented a typing discipline for MLTS as well as a formal specification of its natural semantics: this latter task was aided by being able to directly exploit a rich logic, called $\mathcal{G}$, that incorporates λ-tree syntax principles within quantificational logic. Finally, this natural semantic specification was directly implementable in λProlog. As a consequence, a prototype implementation is available for helping to judge the expressiveness of MLTS programs.

Acknowledgments. We thank Kaustuv Chaudhuri, François Pottier, Enrico Tassi, the HOPE Workshop 2018 audience, and the reviewers on an earlier draft of this paper for their helpful comments and observations.

Bibliography

David Baelde, Kaustuv Chaudhuri, Andrew Gacek, Dale Miller, Gopalan Nadathur, Alwen Tiu, and Yuting Wang. Abella: A system for reasoning about relational specifications. *Journal of Formalized Reasoning*, 7(2), 2014. https://doi.org/10.6092/issn.1972-5787/4650. URL http://jfr.unibo.it/article/download/4650/4137.

Arthur Charguéraud. The locally nameless representation. *Journal of Automated Reasoning*, pages 1–46, May 2011. https://doi.org/10.1007/s10817-011-9225-2.

James Cheney and Christian Urban. Alpha-Prolog: A logic programming language with names, binding, and alpha-equivalence. In Bart Demoen and Vladimir Lifschitz, editors, *Logic Programming, 20th International Conference*, volume 3132 of *LNCS*, pages 269–283. Springer, 2004.

Anthony S. K. Cheng, Peter J. Robinson, and John Staples. Higher level meta programming in qu-prolog 3: 0. In Koichi Furukawa, editor, *Logic Programming, Proceedings of the Eigth International Conference, Paris, France, June 24-28, 1991*, pages 285–298. MIT Press, 1991.

Jawahar Chirimar. *Proof Theoretic Approach to Specification Languages*. PhD thesis, University of Pennsylvania, February 1995. URL http://www.lix.polytechnique.fr/Labo/Dale.Miller/chirimar/phd.ps.

Adam Chlipala. Parametric higher-order abstract syntax for mechanized semantics. In James Hook and Peter Thiemann, editors, *Proceeding of the 13th ACM SIGPLAN international conference on Functional programming, ICFP 2008, Victoria, BC, Canada, September 20-28, 2008*, pages 143–156. ACM, 2008. https://doi.org/10.1145/1411204.1411226.

Alonzo Church. A formulation of the Simple Theory of Types. *J. of Symbolic Logic*, 5:56–68, 1940. https://doi.org/10.2307/2266170.

N. G. de Bruijn. Lambda calculus notation with namefree formulas involving symbols that represent reference transforming mappings. *Indag. Math.*, 40(3): 348–356, 1979.

Nicolaas Govert de Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with an application to the Church-Rosser theorem. *Indagationes Mathematicae*, 34(5):381–392, 1972.

Joëlle Despeyroux, Amy Felty, and Andre Hirschowitz. Higher-order abstract syntax in Coq. In *Second International Conference on Typed Lambda Calculi and Applications*, pages 124–138, April 1995.

Cvetan Dunchev, Ferruccio Guidi, Claudio Sacerdoti Coen, and Enrico Tassi. ELPI: fast, embeddable, λProlog interpreter. In Martin Davis, Ansgar Fehnker, Annabelle McIver, and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning - 20th International Conference, LPAR-20 2015, Suva, Fiji, November 24-28, 2015, Proceedings*, volume 9450 of *LNCS*, pages 460–468. Springer, 2015. https://doi.org/10.1007/978-3-662-48899-7_32. URL http://dx.doi.org/10.1007/978-3-662-48899-7_32.

Francisco Ferreira and Brigitte Pientka. Programs using syntax with first-class binders. In Hongseok Yang, editor, *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, volume 10201 of *Lecture Notes in Computer Science*, pages 504–529. Springer, 2017. ISBN 978-3-662-54433-4; 978-3-662-54434-1.

M. J. Gabbay and A. M. Pitts. A new approach to abstract syntax involving binders. In *14th Symp. on Logic in Computer Science*, pages 214–224. IEEE Computer Society Press, 1999.

Andrew Gacek. *A Framework for Specifying, Prototyping, and Reasoning about Computational Systems*. PhD thesis, University of Minnesota, 2009.

Andrew Gacek, Dale Miller, and Gopalan Nadathur. Combining generic judgments with recursive definitions. In F. Pfenning, editor, *23th Symp. on Logic in Computer Science*, pages 33–44. IEEE Computer Society Press, 2008. URL http://www.lix.polytechnique.fr/Labo/Dale.Miller/papers/lics08a.pdf.

Andrew Gacek, Dale Miller, and Gopalan Nadathur. Nominal abstraction. *Information and Computation*, 209(1):48–73, 2011. https://doi.org/10.1016/j.ic.2010.09.004.

Ulysse Gérard, Dale Miller, and Gabriel Scherer. Try MLTS online, March 2018. https://trymlts.github.io/.

A. Gordon. A mechanisation of name-carrying syntax up to alpha-conversion. In *International Workshop on Higher Order Logic Theorem Proving and its Applications*, volume 780 of *Lecture Notes in Computer Science*, pages 414–426, 1994.

Michael J. Gordon, Arthur J. Milner, and Christopher P. Wadsworth. *Edinburgh LCF: A Mechanised Logic of Computation*, volume 78 of *LNCS*. Springer, 1979.

Michael J. C. Gordon. Introduction to the HOL system. In Myla Archer, Jeffrey J. Joyce, Karl N. Levitt, and Phillip J. Windley, editors, *Proceedings of the International Workshop on the HOL Theorem Proving System and its Applications*, pages 2–3. IEEE Computer Society, 1991.

John Harrison. HOL light: an overview. In *International Conference on Theorem Proving in Higher Order Logics*, pages 60–66. Springer, 2009.

Gérard Huet. A unification algorithm for typed λ-calculus. *Theoretical Computer Science*, 1:27–57, 1975.

js-of-ocaml. Js_of_ocaml, 2018. http://ocsigen.org/js_of_ocaml/.

Gilles Kahn. Natural semantics. In Franz-Josef Brandenburg, Guy Vidal-Naquet, and Martin Wirsing, editors, *Proceedings of the Symposium on Theoretical Aspects of Computer Science*, volume 247 of *LNCS*, pages 22–39. Springer, March 1987.

Conor McBride and James McKinna. Functional pearl: I am not a number - I am a free variable. In Henrik Nilsson, editor, *Proceedings of the ACM SIGPLAN Workshop on Haskell, Haskell 2004, Snowbird, UT, USA, September 22-22, 2004*, pages 1–9. ACM, 2004. URL http://doi.acm.org/10.1145/1017472.1017477.

Dale Miller. An extension to ML to handle bound variables in data structures: Preliminary report. In *Proceedings of the Logical Frameworks BRA Workshop*, pages 323–335, Antibes, France, June 1990. URL http://www.lix.polytechnique.fr/Labo/Dale.Miller/papers/mll.pdf. Available as UPenn CIS technical report MS-CIS-90-59.

Dale Miller. A logic programming language with lambda-abstraction, function variables, and simple unification. *J. of Logic and Computation*, 1(4):497–536, 1991.

Dale Miller. Unification under a mixed prefix. *Journal of Symbolic Computation*, 14(4):321–358, 1992.

Dale Miller. Bindings, mobility of bindings, and the ∇-quantifier. In Jerzy Marcinkowski and Andrzej Tarlecki, editors, *18th International Conference on Computer Science Logic (CSL) 2004*, volume 3210 of *LNCS*, page 24, 2004. https://doi.org/10.1007/978-3-540-30124-0_4.

Dale Miller. Mechanized metatheory revisited. *Journal of Automated Reasoning*, October 2018. ISSN 1573-0670. https://doi.org/10.1007/s10817-018-9483-3.

Dale Miller and Gopalan Nadathur. *Programming with Higher-Order Logic*. Cambridge University Press, June 2012. https://doi.org/10.1017/CBO9781139021326.

Dale Miller and Catuscia Palamidessi. Foundational aspects of syntax. *ACM Computing Surveys*, 31, September 1999.

Dale Miller and Alwen Tiu. A proof theory for generic judgments. *ACM Trans. on Computational Logic*, 6(4):749–783, October 2005. https://doi.org/10.1145/1094622.1094628. URL http://www.lix.polytechnique.fr/Labo/Dale.Miller/papers/tocl-nabla.pdf.

Robin Milner. Functions as processes. In *Automata, Languages and Programming 17th Int. Coll.*, volume 443 of *LNCS*, pages 167–180. Springer, July 1990.

Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, Part I. *Information and Computation*, 100(1):1–40, September 1992.

Gopalan Nadathur and Dale Miller. An Overview of λProlog. In *Fifth International Logic Programming Conference*, pages 810–827, Seattle, August 1988. MIT Press. URL http://www.lix.polytechnique.fr/Labo/Dale.Miller/papers/iclp88.pdf.

Tobias Nipkow. Functional unification of higher-order patterns. In M. Vardi, editor, *8th Symp. on Logic in Computer Science*, pages 64–74. IEEE, June 1993.

Bengt Nordstrom, Kent Petersson, and Jan M. Smith. *Programming in Martin-Löf's type theory : an introduction*. International Series of Monographs on Computer Science. Oxford: Clarendon, 1990.

OCaml. http://ocaml.org/, 2018.

Lawrence C. Paulson. The foundation of a generic theorem prover. *Journal of Automated Reasoning*, 5:363–397, September 1989.

Lawrence C. Paulson. *Isabelle: A Generic Theorem Prover*. Number 828 in LNCS. Springer Verlag, 1994.

Frank Pfenning and Conal Elliott. Higher-order abstract syntax. In *Proceedings of the ACM-SIGPLAN Conference on Programming Language Design and Implementation*, pages 199–208. ACM Press, June 1988.

Frank Pfenning and Carsten Schürmann. System description: Twelf — A meta-logical framework for deductive systems. In H. Ganzinger, editor, *16th Conf. on Automated Deduction (CADE)*, number 1632 in LNAI, pages 202–206, Trento, 1999. Springer. https://doi.org/10.1007/3-540-48660-7_14.

Brigitte Pientka and Joshua Dunfield. Beluga: A framework for programming and reasoning with deductive systems (system description). In J. Giesl and R. Hähnle, editors, *Fifth International Joint Conference on Automated Reasoning*, number 6173 in LNCS, pages 15–21, 2010.

Andrew M. Pitts. Nominal logic, A first order theory of names and binding. *Information and Computation*, 186(2):165–193, 2003.

Adam Poswolsky and Carsten Schürmann. System description: Delphin - A functional programming language for deductive systems. In A. Abel and C. Urban, editors, *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2008)*, volume 228, pages 113–120, 2008.

François Pottier. An overview of Cαml. In *Proceedings of the ACM-SIGPLAN Workshop on ML (ML 2005)*, volume 148 of *Electr. Notes Theor. Comput. Sci.*, pages 27–52, 2006. https://doi.org/10.1016/j.entcs.2005.11.039.

François Pottier. Static name control for FreshML. In *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*, pages 356–365. IEEE, 2007.

Xiaochu Qi, Andrew Gacek, Steven Holte, Gopalan Nadathur, and Zach Snow. The Teyjus system – version 2, 2015. URL http://teyjus.cs.umn.edu/. http://teyjus.cs.umn.edu/.

Zhenyu Qian. Unification of higher-order patterns in linear time and space. *J. of Logic and Computation*, 6(3):315–341, 1996.

Davide Sangiorgi. π-calculus, internal mobility and agent-passing calculi. *Theoretical Computer Science*, 167(2):235–274, 1996.

Helmut Schwichtenberg. Minlog. In Freek Wiedijk, editor, *The Seventeen Provers of the World*, volume 3600 of *LNCS*, pages 151–157. Springer, 2006. https://doi.org/10.1007/11542384_19.

Dana Scott. Outline of a mathematical theory of computation. In *Proceedings, Fourth Annual Princeton Conference on Information Sciences and Systems*, pages 169–176. Princeton University, 1970. Also, Programming Research Group Technical Monograph PRG–2, Oxford University.

M. R. Shinwell, A. M. Pitts, and M. J. Gabbay. FreshML: Programming with binders made simple. In *Eighth ACM SIGPLAN International Conference on Functional Programming (ICFP 2003), Uppsala, Sweden*, pages 263–274. ACM Press, August 2003.

Yuting Wang, Kaustuv Chaudhuri, Andrew Gacek, and Gopalan Nadathur. Reasoning about higher-order relational specifications. In Tom Schrijvers, editor, *Proceedings of the 15th International Symposium on Princples and Practice of Declarative Programming (PPDP)*, pages 157–168, Madrid, Spain, September 2013. https://doi.org/10.1145/2505879.2505889. URL http://chaudhuri.info/papers/draft13hhw.pdf.

A Another MLTS example: the π-calculus

The π-calculus [Milner, Parrow, and Walker, 1992, Milner, 1990] is a language for modeling processes in which interactions are name-based. In particular, this calculus permits communication via named channels, including the communication of the names of the channels themselves. The basic calculus has two syntactic categories: names and processes.

Process expressions are defined by the following syntax rule.

$$P := 0 \quad | \quad P \mid P \quad | \quad P{+}P \quad | \quad x(y).P \quad | \quad \bar{x}y.P \quad | \quad [x = y].P \quad | \quad \tau.P \quad | \quad (y)P \quad | \quad !\, P.$$

Here, x and y range over names. The process 0 cannot perform any actions. The expressions $P \mid P$ and $P + P$ denote, respectively, the parallel composition and the choice of two processes. The next four expressions are *prefixed* processes:

- $x(y).P$ represents a process that can accept a name on the channel x and will then become P with y bound to the input name;
- $\bar{x}y.P$ is a process that can output the name y on the channel x;
- $[x = y].P$ is a process that can become P provided that the names x and y are equal;
- $\tau.P$ is a process that can evolve through a silent action.

The expression $(y)P$ represents the restriction of the name y to P: interactions can take place internally to P through this name but the process cannot communicate externally along the channels $\bar{y}$ or y. Finally, $!\, P$ denotes the parallel composition of any number of copies of P.

To represent expressions of the π-calculus in MLTS, we define the two datatypes `name` and `proc` for names and processes that are given in Figure 9. Note that the two process expressions $x(y).P$ and $(y)P$ embody a binding notion. The λ-tree syntax for these expressions will accordingly include an explicit abstraction. For example, the two π-calculus expressions

$$(y)\bar{a}y.((y(w).0) \mid (\bar{b}b.0)) \qquad \text{and} \qquad (y)\bar{a}y.((y(w).\bar{b}b.0) + (\bar{b}b.y(w).0))$$

are encoded in MLTS with the terms, respectively.

```
Nu(Y\ Out(A,Y,Par(In(Y, W\ Null),Out(B,B,Null))))
Nu(Y\ Out(A,Y,Plus(In(Y, W\ Out(B,B,Null)),
                   Out(B,B, In(Y, W\ Null)))))
```

In this encoding of the π-calculus (Figure 9), the type `name` must be considered open (in the sense described in Section 5) while the type `proc` is not open. The operational semantics of the π-calculus is generally described using a non-deterministic, labeled transition systems. That semantics is easily specified in λProlog [Miller and Nadathur, 2012] and reasoned with in Abella [Baelde, Chaudhuri, Gacek, Miller, Nadathur, Tiu, and Wang, 2014].

One way to demonstrate the expressiveness of the π-calculus is to encode within it the call-by-name evaluation in the untyped λ-calculus. Such a translation

```
type name = | A | B | C;;

type proc =
  | Null
  | Plus of proc * proc
  | Par  of proc * proc
  | In   of name * (name => proc)
  | Out  of name * name * proc
  | Eqn  of name * name * proc
  | Taup of proc
  | Bang of proc
  | Nu   of name => proc;;
```

Fig. 9. Two data types for encoding the π-calculus.

```
let rec trans gamma term = match term with
  | App(m, n) ->
    let p = trans gamma m in
    let q = trans gamma n in
    (U\ Nu(V\ Par(
      p @ V,
      Nu(X\ Out(V, X, Out(V, U, Bang(In(X, q))))))))
  | Abs(m) ->
    new X in (U\ In(U, Y\
      let p = trans ((X,Y)::gamma) (m @ X) in
      In(U, V\ p @ V)))
  | nab X in X -> (U\ Out(assoc X gamma, U, Null));;
```

Fig. 10. Encoding of the call-by-name evaluation of untyped λ-terms into the π-calculus.

function was given by Milner in [Milner, 1990] and it can be written as follows.

$$\begin{aligned}
[\![x]\!](u) &= \bar{x}u.0 \\
[\![\lambda x\ M]\!](u) &= u(x).u(v).[\![M]\!](v) \\
[\![(M\ N)]\!](u) &= (v).([\![M]\!](v) \mid (x).(\bar{v}x.\bar{v}u.!x(w).[\![N]\!](w)))
\end{aligned}$$

Here, the translation function $[\![M]\!](u)$ takes an untyped λ-term M and a name u and returns a process that encodes the λ-term M in such a way that it expects to receive its arguments on channel u. In Figure 10, we provide an MLTS implementation of this translation: in particular, if $[\![M]\!](u)$ is the process calculus expression P, then the function `trans`, when applied to (the encoding of) M would yield (the encoding of) $\lambda u.P$. (The function `assoc` used here is defined in Figure 3.) For example, the value of `(transf [] Abs(X\X))` is

```
(U\ In(U,X\ In(U,(Y\ Out(X,Y, Null))))).
```