

Resolving Inductive Definitions with Binders in Higher-Order Typed Functional Programming*

Matthew R. Lakin and Andrew M. Pitts

University of Cambridge Computer Laboratory, Cambridge CB3 0FD, UK
{Matthew.Lakin,Andrew.Pitts}@cl.cam.ac.uk

Abstract. This paper studies inductive definitions involving binders, in which aliasing between free and bound names is permitted. Such aliasing occurs in informal specifications of operational semantics, but is excluded by the common representation of binding as meta-level λ -abstraction. Drawing upon ideas from functional logic programming, we represent such definitions with aliasing as recursively defined functions in a higher-order typed functional programming language that extends core ML with types for name-binding, a type of “semi-decidable propositions” and existential quantification for types with decidable equality. We show that the representation is sound and complete with respect to the language’s operational semantics, which combines the use of evaluation contexts with constraint programming. We also give a new and simple proof that the associated constraint problem is NP-complete.

1 Introduction

Perhaps the single most important technique in the study of programming language semantics is the use of inductive definitions. This is especially the case for operational semantics, which in broad terms consists of one or more inductively defined relations between data structures involving programming language syntax. The inductive definition commonly takes the form of finitely many “schematic” rules containing parameters that can be instantiated, usually in infinitely many different ways, to get concrete rules for inductively generating instances of the relations. Schematic rules are necessarily written in some meta-language whose definition is often left implicit in published research. Having to be completely precise about the meta-language of rule schemes is an inescapable part of the current trend toward mechanization of semantics, whether it be machine-assisted proof construction/checking, or executable semantic specifications. In this paper we are concerned with the latter, but in either case it is clear that the ubiquitous presence of binding constructs in the “object-language” (that is, the programming language whose semantics is being formalized) creates difficulties for mechanized meta-languages (see the *POPLmark Challenge* wiki, for example). Ideally one would like the executable meta-language for rule schemes to provide a fully automatic treatment of α -conversion of bound names

* Research supported by UK EPSRC grant EP/D000459/1.

in object-languages. Here we investigate a way of doing that within the context of higher-order typed functional programming. In doing so we take a “nominal” approach to object-level binders for the following reasons to do with *name aliasing*.

One way of dealing with issues of α -conversion is to make the representation of object-level binders completely anonymous. This can be achieved through a case-by-case use of de Bruijn indices [8], or more systematically by use of Higher-Order Abstract Syntax (HOAS) [21] to enforce that object-level bound names are only represented by meta-level bound names. In either case the conceptually simple operation of instantiating parameters in a rule scheme, which may involve capture of a free name by a binder of the same name, has to be replaced by something more complicated—simply because “binder of the same name” makes no sense if binders have been anonymized. But there is a more serious problem with anonymous representations of object-level binding: they rule out the common practice of *name-aliasing* involving binders, be it the use of the same name at two different binding occurrences, or the use of the same name for both a free and a binding occurrence. For example, in an inductive definition of β -reduction of λ -terms it is natural to use the rule scheme

$$\frac{t \rightarrow t'}{\lambda x. t \rightarrow \lambda x. t'}$$

where the two different binding occurrences within the conclusion are both named x . In this case the name-aliasing does not cause a problem for a formalization using HOAS, which might render the above rule as

$$\frac{f(x) \rightarrow f'(x)}{\lambda(f) \rightarrow \lambda(f')}$$

where f and f' are meta-variables of function type. The top half of Fig. 1 contains another example. (The bottom half of the figure will be explained in Sect. 2.) In this case the conclusion of the third rule contains both free (the first x) and binding (the x within $L(\langle x \rangle t)$) occurrences with the same name. We leave the reader to ponder how to convert these rules into an extensionally equivalent HOAS formalization (probably by ignoring the third rule completely). In fact we do not know any definitive results comparing the class of relations (on tuples of α -equivalence classes of λ -terms, say) defined by HOAS rule schemes with the class defined by first-order rule schemes with conventional, named binders. In any case, the phenomenon of name-aliasing seems too convenient to give up unless we really have to.

So we advocate the study of executable meta-languages for rule schemes that allow object-level binders to be named. More specifically we study such an executable meta-language which is a higher-order, typed functional programming language, drawing upon the ideas of *functional logic programming* [14]. Our motivation for favouring this paradigm over the relational paradigm of first-order logic programming has to do with the expressiveness and modularity afforded by

$$\begin{array}{c}
\text{nfv}(x, t): \text{“}x : \mathbf{vr} \text{ is not a free variable of the } \lambda\text{-term } t : \mathbf{tm}\text{”} \\
\frac{x \# x'}{\text{nfv}(x, \mathbf{V} x')} \quad \frac{\text{nfv}(x, t) \ \& \ \text{nfv}(x, t')}{\text{nfv}(x, \mathbf{A}(t, t'))} \quad \frac{}{\text{nfv}(x, \mathbf{L}(\langle x \rangle t))} \quad \frac{\text{nfv}(x, t) \ \& \ x \# x'}{\text{nfv}(x, \mathbf{L}(\langle x' \rangle t))} \\
\text{Equivalent standard form } \frac{\varphi}{\text{nfv } y} \text{ has} \\
\varphi \triangleq \exists x, x' (y = (x, \mathbf{V} x') \ \& \ x \# x') \vee \exists x, t, t' (y = (x, \mathbf{A}(t, t')) \ \& \ \text{nfv}(x, t) \ \& \ \text{nfv}(x, t')) \\
\vee \exists x, t (y = \text{nfv}(x, \mathbf{L}(\langle x \rangle t))) \vee \exists x, x', t (y = \text{nfv}(x, \mathbf{L}(\langle x' \rangle t)) \ \& \ \text{nfv}(x, t) \ \& \ x \# x')
\end{array}$$

Fig. 1. Example α -inductive definition

higher-order functions. For example, having higher-order functions allows one to encode definitions that are parameterised by other definitions (such as the various operations on relations that occur in the relational approach to contextual equivalence of programs [12, 17, 25]).

Contributions of this paper. We begin by fixing a simple, yet expressive class of inductive definitions permitting name-aliasing in binders, where binding is handled generically through the existing notion of a “nominal signature”. These *α -inductive definitions* (Sect. 2) may involve side-conditions asserting constraints in terms of α -equivalence and the “not-a-free-variable-of” relation. In Sect. 3 we make the apparently new observation that such constraints can express membership in finite sets; consequently the associated constraint satisfaction problem is NP-complete (Theorem 3.3). Section 4 introduces the main contribution of the paper, a typed higher-order functional programming language that extends core ML with name-binding types, a type of “semi-decidable propositions” and existential quantification for types in a class of equality types coinciding with the arities of a user-declared nominal signature. This language, which we call α ML, draws upon the ideas of α Prolog [7], extending them to higher-order functional programming. α ML is a simplification of both our first attempt to do this [16] and of α Prolog itself, in that *it avoids the use of concrete names and name-permutations in programs* (see Remark 2.2 and Sect. 6 for the significance of this). α ML has a remarkably simple operational semantics that combines the use of Felleisen-style evaluation contexts with constraint programming; we show that it restricts to the usual operational semantics on the purely functional part of α ML (Theorem 4.1). By design, α ML represents α -inductive definitions as certain recursively defined functions; we prove that this representation is sound and complete (Corollary 5.3). Finally, Sect. 6 discusses related and future work.

2 α -Inductive Definitions

In this section we give a simple, yet expressive class of *inductively defined relations between α -equivalence classes of expressions*, or *α -inductive definitions* for short. Alpha-equivalence arises from the presence of binding constructs in

the expressions and we will deal with this in a generic way by using *nominal signatures* [28], Σ . These generalize the usual notion of many-sorted algebraic signature to encompass constructors that bind names of various sorts. Such a Σ is given by a finite set of *name sorts* N , a disjoint finite set of *data sorts* S , and a finite set of *constructors* $K : A \rightarrow S$, each with a specified result sort S and argument arity A —where the *nominal arities* of the signature are given by:

$$\begin{aligned}
A ::= S & \quad (\text{data sort}) \\
& A * \dots * A \quad (\text{tuples}) \\
& N & \quad (\text{name sort}) \\
& [N]A & \quad (\text{name-abstractions}).
\end{aligned} \tag{1}$$

For example, the nominal signature for untyped λ -calculus has a name sort \mathbf{vr} for variables, a data sort \mathbf{tm} for λ -terms and constructors $\mathbf{V} : \mathbf{vr} \rightarrow \mathbf{tm}$, $\mathbf{A} : \mathbf{tm} * \mathbf{tm} \rightarrow \mathbf{tm}$ and $\mathbf{L} : [\mathbf{vr}] \mathbf{tm} \rightarrow \mathbf{tm}$. For other examples, see [24, Sect. 2]. As in that paper, we associate with each arity A of a nominal signature Σ a set $\alpha\text{-Tree}_\Sigma(A)$ of α -equivalence classes of abstract syntax trees, or α -trees for short. The elements of each $\alpha\text{-Tree}_\Sigma(A)$ are equivalence classes $[g]_\alpha$ of syntax trees $g \in \text{Tree}_\Sigma(A)$ built up from countably many *names* $n \in \text{Name}(N)$ (for each name sort N of Σ) by repeatedly applying the following three operations.

Constructor application: $K g \in \text{Tree}_\Sigma(S)$, if $g \in \text{Tree}_\Sigma(A)$ and $K : A \rightarrow S$.

Tupling: $(g_1, \dots, g_n) \in \text{Tree}_\Sigma(A_1 * \dots * A_n)$, if $g_i \in \text{Tree}_\Sigma(A_i)$ for $i = 1..n$.

Name-abstraction: $\langle n \rangle g \in \text{Tree}_\Sigma([N]A)$, if $n \in \text{Name}(N)$ and $g \in \text{Tree}_\Sigma(A)$.

(These trees are the *ground* nominal terms from [28], that is, the ones not involving variables.) The third operation, name-abstraction, is the generic binding form provided by nominal signatures: renaming $\langle n \rangle$ -bound names in trees gives an equivalence relation $=_\alpha$ [24, Fig. 1] and $\alpha\text{-Tree}_\Sigma(A)$ is the quotient $\text{Tree}_\Sigma(A)/=_\alpha$. To specify inductively defined relations between α -trees we make use of a simple meta-language of patterns.

Definition 2.1 (patterns and valuations). The *patterns* $p \in \text{Pat}_\Sigma(A)$ for describing α -trees of each arity A of Σ are built up from countably many *variables* $x \in \text{Var}(A)$ (for each A) by repeatedly applying the three tree-forming operations mentioned above:

Constructor application: $K p \in \text{Pat}_\Sigma(S)$, if $p \in \text{Pat}_\Sigma(A)$ and $K : A \rightarrow S$.

Tupling: $(p_1, \dots, p_n) \in \text{Pat}_\Sigma(A_1 * \dots * A_n)$, if $p_i \in \text{Pat}_\Sigma(A_i)$ for $i = 1..n$.

Name-abstraction: $\langle x \rangle p \in \text{Pat}_\Sigma([N]A)$, if $x \in \text{Var}(N)$ and $p \in \text{Pat}_\Sigma(A)$.

A *valuation* V is a finite function mapping variables to α -trees (of the same arity). If the variables occurring in a pattern $p \in \text{Pat}_\Sigma(A)$ are in $\text{dom}(V)$ (the domain of definition of V), then $\llbracket p \rrbracket_V \in \alpha\text{-Tree}_\Sigma(A)$ denotes the α -tree resulting from p by replacing each $x \in \text{dom}(V)$ with $V(x)$.

Remark 2.2. The following points about patterns should be noted.

(a) *Variables stand for unknown α -trees, not unknown trees, and (hence) a pattern $p \in \text{Pat}_\Sigma(A)$ describes an α -tree rather than a tree* (just which one depends upon how its variables are instantiated by a valuation). This reflects the common

practice of leaving α -equivalence implicit and referring to a class via a representative, signalled by a phrase like “we identify expressions up to α -equivalence”. (Our own Figs. 2 and 3 in Sect. 4 provide examples of this!)

(b) *No concrete names $n \in \text{Name}(N)$ occur in patterns.* In particular, although the meta-language allows us to name object-level binding occurrences, $\langle x \rangle p$, we use variables x of name sort rather than names themselves to do so. Again, this reflects common practice. For example, in Barendregt’s classic text [1], Definition 2.1.1 says that λ -terms are words over an alphabet containing, among other things, “variables v_0, v_1, \dots ”; then Notation 2.1.2 says that “ $x, y, z \dots$ denote arbitrary variables”; the concrete variables v_0, v_1, \dots are never mentioned again and only the meta-variables $x, y, z \dots$ are used throughout the rest of the book.

(c) *There are no meta-level variable-binding constructs in patterns—all variables in a pattern are free.* In particular x occurs free in the name-abstraction pattern $\langle x \rangle p$. This allows patterns to support the phenomenon of name-aliasing discussed in the Introduction.

(d) *Valuation of patterns is a form of “possibly-capturing” substitution.* Once again, this reflects common practice when instantiating the meta-variables of schematic rules in operational semantics. Note that, in addition to the previous point, this is another reason why it makes no sense to try to identify patterns up to renaming $\langle x \rangle (-)$ -scoped variables, since valuations do not respect such a notion of α -equivalence. For example, we cannot regard $\langle x \rangle z$ and $\langle y \rangle z$ as equivalent (where x, y and z are distinct), since the valuation $V = \{x \mapsto [n]_\alpha, y \mapsto [n']_\alpha, z \mapsto [n]_\alpha\}$ (with $n \neq n'$) has $\llbracket \langle x \rangle z \rrbracket_V = \llbracket \langle n \rangle n \rrbracket_\alpha \neq \llbracket \langle n' \rangle n \rrbracket_\alpha = \llbracket \langle y \rangle z \rrbracket_V$.

Fix a finite set of *relation symbols* $r <: A$, each with a specified arity A . Such an r is intended to denote a subset of $\alpha\text{-Tree}_\Sigma(A)$. Schematic rules for inductively defining such subsets take the form

$$\frac{r_1(p_1) \ \& \ \cdots \ \& \ r_m(p_m) \ \& \ c_1 \ \& \ \cdots \ \& \ c_n}{r(p)} \quad (2)$$

(see Fig. 1 for an example). The conclusion of (2) is an *atomic formula* $r(p)$ with $r <: A$ and $p \in \text{Pat}_\Sigma(A)$ for some arity A ; the hypothesis is a finite (possibly empty) conjunction of such atomic formulas and “side-conditions” c_i , that is, constraints on how the rule may be instantiated by a valuation. What form of constraints should we use? At the very least we need *name-inequality* constraints $x \neq x'$, where $x, x' \in \text{Var}(N)$ with N a name sort. Experience with nominal logic [23, 11] and nominal logic programming [7] shows that it is useful to generalize name-inequality to *name freshness* constraints, $x \# p$ where $x \in \text{Var}(N)$ and $p \in \text{Pat}_\Sigma(A)$, even though these can be inductively defined in terms of name-inequality (cf. Fig. 1). The intended meaning of $x \# p$ on α -trees is as follows.

Definition 2.3 (free names and freshness). If $t \in \alpha\text{-Tree}_\Sigma(A)$ we write $FN t$ for the finite set of names that occur freely in some (indeed, any) tree g with $[g]_\alpha = t$; in other words, $n \in FN[g]_\alpha$ iff n occurs in g , but not within the

scope of any name-abstraction $\langle n \rangle(-)$. Note that each α -tree $t \in \alpha\text{-Tree}_\Sigma(N)$ of name sort N is of the form $t = [n]_\alpha = \{n\}$ for some $n \in \text{Name}(N)$ (because the constructors of a nominal signature only produce results of data sort, rather than of more general arities, and these are disjoint from the name sorts). Given $t \in \alpha\text{-Tree}_\Sigma(N)$ and $t' \in \alpha\text{-Tree}_\Sigma(A)$, we write $t \# t'$ and say t is fresh for t' if $t = [n]_\alpha$ and $n \notin \text{FN } t'$.

Note that the case of several mutually inductively defined relation symbols $r_1 <: A_1, \dots, r_k <: A_k$ reduces to the case of a single one at the expense of extending the signature: we add a new data sort S and new constructors $I_i : A_i \rightarrow S$ for $i \in \{1..k\}$ and use the fact that subsets $R \subseteq \alpha\text{-Tree}_\Sigma(S)$ are in bijection with n -tuples of subsets, $R_1 \subseteq \alpha\text{-Tree}_\Sigma(A_1), \dots, R_n \subseteq \alpha\text{-Tree}_\Sigma(A_n)$. So from now on we will fix on a single arity A_r of a nominal signature Σ and a single relation symbol $r <: A_r$.

Definition 2.4 (formulas and satisfaction). Let Form_Σ be the set of first-order formulas built up from atomic formulas $r(p)$ (where $p \in \text{Pat}_\Sigma(A_r)$), equalities $p = p'$ ($p, p' \in \text{Pat}_\Sigma(A)$, some A) and freshnesses $x \# p$ ($x \in \text{Var}(N)$, $p \in \text{Pat}_\Sigma(A)$, some N, A) just using finite conjunctions $\&$, finite disjunctions \vee and existential quantification $\exists x(-)$ ($x \in \text{Var}(A)$, some A). Given an interpretation $R \subseteq \alpha\text{-Tree}_\Sigma(A_r)$ for the relation symbol r and a valuation V (Definition 2.1) for the free variables of $\varphi \in \text{Form}_\Sigma$ (i.e. those not within the scope of an $\exists x$), let $(R, V) \models \varphi$ denote the associated *satisfaction relation*. Thus $(R, V) \models r(p)$ holds if $\llbracket p \rrbracket_V \in R$; $(R, V) \models p = p'$ holds if $\llbracket p \rrbracket_V = \llbracket p' \rrbracket_V$; $(R, V) \models x \# p$ holds if $V(x) \# \llbracket p \rrbracket_V$ (Definition 2.3); and satisfaction is extended to compound formulas in the usual way.

Generalizing (2), we will allow the hypothesis of a schematic rule to be a formula in Form_Σ . Allowing equality, disjunction and existential quantification in addition to freshness and conjunction does not increase the expressive power of inductive definitions; but it does allow us to write inductive definitions in a “standard form”, illustrated in Fig. 1.

Definition 2.5 (standard α -inductive definitions). An α -inductive definition \mathcal{D} in standard form of a set of α -trees of arity A_r is given by

$$\frac{\varphi}{r(x)} \quad (3)$$

where $x \in \text{Var}(A_r)$ and $\varphi \in \text{Form}_\Sigma$ is a formula with at most x free. The meaning $\llbracket \mathcal{D} \rrbracket \subseteq \alpha\text{-Tree}_\Sigma(A_r)$ of \mathcal{D} is by definition the least fixed point of the monotone function $\Phi_{\mathcal{D}}$ on subsets of α -trees that maps each $R \subseteq \alpha\text{-Tree}_\Sigma(A_r)$ to

$$\Phi_{\mathcal{D}}(R) \triangleq \{t \in \alpha\text{-Tree}_\Sigma(A_r) \mid (R, \{x \mapsto t\}) \models \varphi\}. \quad (4)$$

The definition of $\llbracket \mathcal{D} \rrbracket$ via (4) is a fancy way of stating the usual meaning of a rule-based inductive definition: the rule (3) is schematic in the sense that it has the variable x as parameter; we instantiate x to get many concrete rules (this

is the effect of the valuations $\{x \mapsto t\}$ in the definition of $\Phi_{\mathcal{D}}$) and take the least set of α -trees closed under these rules, in other words, the least R such that $\Phi_{\mathcal{D}}(R) \subseteq R$. The existence of $\llbracket \mathcal{D} \rrbracket$ is an application of the usual Tarski fixed point theorem ($\Phi_{\mathcal{D}}$ is monotone because the relation symbol r only occurs positively in φ). Indeed $\Phi_{\mathcal{D}}$ is finitary and we can construct $\llbracket \mathcal{D} \rrbracket$ as the union of the countable chain of subsets $\emptyset \subseteq \Phi_{\mathcal{D}}(\emptyset) \subseteq \Phi_{\mathcal{D}}(\Phi_{\mathcal{D}}(\emptyset)) \subseteq \dots$ of α -Tree $_{\Sigma}(A_r)$.

3 α -Tree Constraint Problems

The hypothesis φ of an α -inductive definition (3) contains non-inductive equality and freshness constraints. When instantiated by a particular valuation, the validity of such constraints amounts to α -equivalence of trees $g \in \text{Tree}_{\Sigma}(A)$ and to non-membership of the set of free names of such trees. These are properties that can be decided in linear time; see [3] for example. However, the problem of checking whether or not there is some valuation that validates a collection of equality and freshness constraints is surprisingly more complicated, mainly because of the presence of variables x in binding position in name-abstraction patterns $\langle x \rangle p$ (see points (c) and (d) in Remark 2.2).

Definition 3.1 (constraints and their satisfaction). A formula $\varphi \in \text{Form}_{\Sigma}$ is an α -tree *constraint* if it is of the form $\exists x_1 \dots \exists x_m (c_1 \& \dots \& c_n)$ with each c_i either an equality ($p = p'$) or a freshness ($x \# p$). Since such formulas do not involve the relation symbol $r <: A_r$, the satisfaction relation of Definition 2.4 restricts to a relation $V \models \varphi$ between valuations and constraints. A *constraint problem* is a closed constraint formula and it is *satisfiable* if $\emptyset \models \varphi$ holds, where \emptyset denotes the valuation with empty domain.

That satisfaction of α -tree constraint problems is decidable and in NP can be deduced from results about nominal unification [28]: see [5, Theorem 7.1.2]. One can show that it is also NP-hard via the following simple observation, which seems to be new.¹ In stating it we use the abbreviation $\langle x_1, \dots, x_n \rangle (-)$ to stand for iterated name-abstraction $\langle x_1 \rangle (\dots \langle x_n \rangle (-) \dots)$.

Lemma 3.2 (set membership as an α -tree constraint). *Given distinct variables $x, x_1, \dots, x_k, x', x'_1, \dots, x'_k \in \text{Var}(N)$ (for some name sort N), define $\text{mem}(x, x_1, \dots, x_k) \triangleq \exists x' \exists x'_1 \dots \exists x'_k (x \# x' \& \langle x_1, \dots, x_k \rangle x = \langle x'_1, \dots, x'_k \rangle x')$. Then a valuation V on $\{x, x_1, \dots, x_k\}$ satisfies $\text{mem}(x, x_1, \dots, x_k)$ iff $V(x)$ is a member of the finite set $\{V(x_i) \mid i = 1..k\}$. \square*

We can use this lemma to show NP-hardness by reduction of GRAPH 3-COLOURABILITY. Given a finite graph with vertices v_1, \dots, v_n (which we can take to be variables of some name sort), edges e_1, \dots, e_m and source/target functions $s, t : \{e_1, \dots, e_m\} \rightrightarrows \{v_1, \dots, v_n\}$, then the formula

¹ Cheney's proof of NP-hardness [4] for his constraint problems is not applicable here, because it relies upon the use of concrete names and name-permutations.

$\exists r, g, b, v_1, \dots, v_n (r \# g \& g \# b \& b \# r \& \&_{i=1}^n \text{mem}(v_i, r, g, b) \& \&_{j=1}^m s(e_j) \# t(e_j))$
 is logically equivalent to an α -tree constraint problem which is satisfiable iff the graph's vertices can be coloured with one of three colours (r, g, b) so that no edge connects vertices of the same colour. So altogether we have:

Theorem 3.3. *Satisfiability of α -tree constraint problems is NP-complete.* \square

4 α ML

We are going to make the α -inductive definitions of Sect. 2 executable by embedding the simple meta-language language of patterns and formulas in which they are expressed within an ML-like functional programming language, called α ML. The embedding has two attractive features:

- (a) *nominal signatures Σ are subsumed within recursive data type declarations;*
- (b) *α -inductive definitions \mathcal{D} become instances of recursively defined functions.*

To achieve point (a) we mimic FreshML [26] and extend ML's type system with types of name N and name-abstraction types $[N]T$. However, unlike FreshML and for the reasons given below, we will restrict the use of $[N]T$ to the case when T is an *equality type* in the sense of Standard ML [20, Sect. 4.4]. To achieve point (b) we note that the meaning $\llbracket \mathcal{D} \rrbracket$ of \mathcal{D} is the fixed point of a higher-order function (4). We represent subsets of $\alpha\text{-Tree}_\Sigma(A)$ by α ML functions of type $A \rightarrow \mathbf{prop}$, where \mathbf{prop} is a new type of "semi-decidable propositions"; and then $\Phi_{\mathcal{D}}$ is represented by a function of type $(A \rightarrow \mathbf{prop}) \rightarrow (A \rightarrow \mathbf{prop})$. In order to be able to write this function, α ML extends the pure functional core of ML with name-binding patterns, with equality and freshness constraints, and with existential quantification over values of equality types. The syntax of α ML types and expressions is given in Fig. 2. For simplicity's sake α ML has a monomorphic type system with a single, top-level data type declaration of some name sorts (N), of some data sorts (S , including a distinguished one `bool` with constructors `T ()` and `F ()`) whose recursive definitions may only involve equality types (E), and of some general data types (D) whose recursive definitions may involve function types and \mathbf{prop} . Note that in accord with point (a) above, such a declaration subsumes the notion of nominal signature [28] that we used in Sect. 2; in particular, the signature's nominal arities (1) coincide with equality types.

Turning to α ML's operational semantics, the behaviour of the pure functional constructs is completely straightforward and could be formalized in any of the standard ways. We use Felleisen-style evaluation contexts [9], formalized using *frame stacks* [22], because this makes the combination with α ML's impure features smoother. These impure features are α -tree equality and freshness constraints, and existentially quantified variables of equality type. We describe their behaviour by combining the use of frame stacks with the techniques of *constraint logic programming* (CLP) [15] applied to the α -tree constraint problems of the previous section. A constraint-based approach gives a clean, abstract presentation that avoids the use of unifying substitutions; this is especially useful here

<p><i>Equality types</i> $E ::=$</p> <p>S (data sort)</p> <p>$E * \dots * E$ (tuples, including nullary case, unit)</p> <p>N (name sort)</p> <p>$[N]E$ (name-abstraction)</p> <p><i>Types</i> $T ::=$</p> <p>E (equality type)</p> <p>D (data type)</p> <p>$T * \dots * T$ (tuples)</p> <p>$T \rightarrow T$ (functions)</p> <p>prop (semi-decidable propositions)</p> <p><i>Data type declaration:</i></p> <p>names $N \dots$</p> <p>data</p> <p>bool = T of unit F of unit</p> <p>$S = K_1$ of E_1 \dots K_n of E_n</p> <p>\vdots</p> <p>$D = K'_1$ of T_1 \dots K'_n of T_n</p> <p>\vdots</p>	<p><i>Expressions</i> $e, v ::=$</p> <p>x, f (variables)</p> <p>let $x = e$ in e (sequencing)</p> <p>$K v$ (constructor application)</p> <p>case v of $K x \Rightarrow e$ (case analysis)</p> <p> \dots</p> <p> $K x \Rightarrow e$</p> <p>(v, \dots, v) (tuple)</p> <p>$e.i$ (projections, $i \in \mathbb{N}$)</p> <p>fun $f x = e$ (recursive function)</p> <p>$v v$ (function application) <i>pure</i></p> <p>\dots <i>impure</i></p> <p>$\langle v \rangle v$ (name-abstraction)</p> <p>T (empty constraint)</p> <p>c (atomic constraint)</p> <p>$\exists x e$ (existential)</p> <p><i>Atomic constraints</i> $c ::=$</p> <p>$v = v$ (equality)</p> <p>$v \# v$ (freshness)</p> <hr/> <p><i>Frame stacks</i> $s ::=$</p> <p>id (empty)</p> <p>$s \circ (x.e)$ (non-empty)</p> <hr/> <p>Expressions and frame stacks in <i>A-normal form</i> are obtained by restricting v to range over</p> <p><i>Values</i> $v ::=$</p> <p>x, f $K v$ (v, \dots, v) fun $f x = e$ $\langle v \rangle v$ T</p>
--	---

• We only consider well-typed expressions and frame stacks. We use explicitly typed variables $x \in \text{Var}(T)$ as T ranges over types. The typing of the pure functional part of αML is entirely standard; the types of αML 's impure features are:

- Name-abstraction* $\langle e \rangle e' : [N]E$ if $e : N$ and $e' : E$
- Empty constraint* $\mathbf{T} : \mathbf{prop}$
- Equality constraint* $e = e' : \mathbf{prop}$ if $e : E$ and $e' : E$
- Freshness constraint* $e \# e' : \mathbf{prop}$ if $e : N$ and $e' : E$
- Existential* $\exists x e : T$ if $x \in \text{Var}(E)$ and $e : T$.

• We identify αML expressions up to renaming bound variables. Despite the fact that αML is a meta-language for object-level languages with binding, there is no reason not to adopt the usual conventions (see Remark 2.2(a)) for αML 's own variable-binding constructs. In the pure part, variable-binding occurs in the usual way, in **let**, **case** and **fun** expressions and in frame stacks $s \circ (x. -)$; and in the impure part, $\exists x(-)$ is a binder. We write $FV(e)$ for the finite set of free variables of an expression e and say e is *closed* if this set is empty. We write $e[e'/x]$ for the capture-avoiding substitution of e' for all free occurrences of x in e (well-defined up to renaming bound variables).

Fig. 2. αML syntax

<i>Pure transitions</i> $s, e \rightarrow s', e'$	
(P1)	$s \circ (x.e), v \rightarrow s, e[v/x]$
(P2)	$s, (\text{let } x = e \text{ in } e') \rightarrow s \circ (x.e'), e$
(P3)	$s, (v_1, \dots, v_n).i \rightarrow s, v_i \text{ if } i \in \{1..n\}$
(P4)	$s, (\text{case } K_i \text{ v of } K_1 x_1 \Rightarrow e_1 \mid \dots \mid K_n x_n \Rightarrow e_n) \rightarrow s, e_i[v/x_i] \text{ if } i \in \{1..n\}$
(P5)	$s, v v' \rightarrow s, e[v/f, v'/x] \text{ if } v \text{ is } \text{fun } f x = e$
<i>Impure transitions</i> $\exists \vec{x}(\vec{c}; s; e) \rightarrow \exists \vec{x}'(\vec{c}'; s'; e')$	
(I1)	$\exists \vec{x}(\vec{c}; s; e) \rightarrow \exists \vec{x}'(\vec{c}'; s'; e') \text{ if } s, e \rightarrow s', e'$
(I2)	$\exists \vec{x}(\vec{c}; s; x.i) \rightarrow \exists \vec{x}, x_1, \dots, x_n(\vec{c} \ \& \ x = (x_1, \dots, x_n); s; x_i)$
(I3)	$\exists \vec{x}(\vec{c}; s; \text{case } x \text{ of } K_1 x_1 \Rightarrow e_1 \mid \dots \mid K_n x_n \Rightarrow e_n) \rightarrow \exists \vec{x}, x_i(\vec{c} \ \& \ x = K_i x_i; s; e_i)$ if $i \in \{1..n\}$ and $\emptyset \models \exists \vec{x}, x_i(\vec{c} \ \& \ x = K_i x_i)$
(I4)	$\exists \vec{x}(\vec{c}; s; c) \rightarrow \exists \vec{x}(\vec{c} \ \& \ c; s; \top) \text{ if } \emptyset \models \exists \vec{x}(\vec{c} \ \& \ c)$
(I5)	$\exists \vec{x}(\vec{c}; s; \exists x e) \rightarrow \exists \vec{x}, x(\vec{c}; s; e) \text{ if } \emptyset \models \exists x(\top)$

- e, s, e', s', e_i, \dots range over expressions and frame stacks in A-normal form (Fig. 2).
- Impure transitions are between *configurations* $\exists \vec{x}(\vec{c}; s; e)$ where \vec{c} is a finite conjunction of atomic constraints and \vec{x} is a finite list of distinct variables of equality type containing the free variables of \vec{c} , s and e . As for expressions, we identify configurations up to renaming of \exists -bound variables. The initial configuration is $\exists \emptyset(\top; id; e)$.
- In (I2) $x \in \text{Var}(E_1 * \dots * E_n)$ and $x_i \in \text{Var}(E_i) - \vec{x}$ for $i = 1..n$.
- In (I3) $x \in \text{Var}(S)$, $S = K_1 \text{ of } E_1 \mid \dots \mid K_n \text{ of } E_n$ and $x_i \in \text{Var}(E_i) - \vec{x}$ for $i = 1..n$.
- In (I5) $x \notin \vec{x}$. If $x \in \text{Var}(E)$ say, then constraint $\exists x(\top)$ is satisfiable iff E is non-empty, in the sense that there is an α -tree of arity E . We allow empty data sorts, e.g. that given by the declaration $\text{es} = K$ of es .

Fig. 3. α ML operational semantics

because the “possibly-capturing” nature of substitution (cf. Remark 2.2(d)) complicates unification—see for example the use of terms involving explicit name-permutations in nominal [28] and equivariant [6] unification algorithms. α ML’s operational semantics is specified in Fig. 3. To simplify the presentation we have restricted to the *A-normal forms* [10] from Fig. 2; transitions for general expressions can be derived by reducing them to A-normal form. The α ML transition relation is non-deterministic, because of the “narrowing” that occurs when evaluating a **case**-expression whose subject is an existentially quantified variable (transition (I3) in Fig. 3). There is a considerable literature about this specific source of non-determinism, centred around the semantics of the functional logic programming language Curry; see [14] for a survey. Since α ML features non-trivial computational effects and we do not wish to impose a monadic programming style, we prefer a strict evaluation strategy, rather than the call-by-need strategy that is more common in the functional logic programming literature; and for simplicity’s sake we wish to avoid residuation and concurrent execution [14, Sect. 2.4]. So we use a simple-minded design where the “rigid/flexible” behaviour of **case**-analysis is part of the dynamics (pure transition (P4) versus impure transition (I3)), rather than user-specified.

The following theorem shows that we do achieve the design goal of embedding within α ML the usual operational behaviour of pure functional programming

with recursive data types and call-by-value higher-order functions. It depends on a notion of configurations being well-typed, $\exists \vec{x}(\vec{c}; s; e) : T$, whose straightforward definition we omit here.

Theorem 4.1 (embedded pure functional language). *An α ML expression or frame stack is pure if it does not contain sub-expressions of the form $\langle e \rangle e'$, \top , $e = e'$, $e \# e'$, or $\exists x e$. Suppose s and e are pure and that $\exists \emptyset(\top; s; e) : T$ holds for some type T . Then $\exists \emptyset(\top; s; e) \rightarrow \exists \vec{x}(\vec{c}; s'; e')$ holds iff $\vec{x} = \emptyset$, $\vec{c} = \top$, s' and e' are pure, and there is a pure transition $(s; e) \rightarrow (s'; e')$. \square*

α ML restricts the name-abstraction type-former $[N](-)$ to apply only to equality types E (that is, to nominal arities) rather than to general types T , for which equality constraints are in general uncomputable. This allows expressions of name-abstraction type to be deconstructed by unification with name-abstraction patterns $\langle x \rangle p$ in the presence of freshness constraints on x , rather than using the “generative name unbinding” [25] mechanism of FreshML, which is based on a supply of dynamically allocated fresh names (“gensym”). Here is an example.

Example 4.2. If the nominal signature for λ -terms mentioned in Sect. 2 is part of the data type declaration, then the λ -term substitution operation $(t, t', x') \mapsto t[t'/x']$ can be encoded in α ML by the following function of type $\mathbf{tm} * \mathbf{tm} * \mathbf{vr} \rightarrow \mathbf{tm}$

$$\begin{aligned} \mathbf{fun} \text{subst}(t, t', x') = \mathbf{case} \ t \ \mathbf{of} \\ \quad \mathbf{V} \ x \Rightarrow \mathbf{if} \ x = x' \ \mathbf{then} \ t' \ \mathbf{else} \ t \\ \quad \mathbf{I} \ \mathbf{A} \ x \Rightarrow \mathbf{A}(\text{subst}(x.1, t', x'), \text{subst}(x.2, t', x')) \\ \quad \mathbf{I} \ \mathbf{L} \ x \Rightarrow \exists x_1 \exists t_1 \ \underline{x_1 \# (t', x')} \ \& \ \langle x_1 \rangle t_1 = x \ \& \ \mathbf{L}(\langle x_1 \rangle \text{subst}(t_1, t', x')) \end{aligned} \quad (5)$$

where we have used some syntactic sugar for tuple-pattern matching, together with the following abbreviations:

$$e_1 \ \& \ e_2 \triangleq \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 \quad (6)$$

$$\mathbf{if} \ e \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2 \triangleq \mathbf{case} \ e \ \mathbf{of} \ \mathbf{T} \ x_1 \Rightarrow e_1 \ \mathbf{I} \ \mathbf{F} \ x_2 \Rightarrow e_2 \quad (7)$$

(where $x \notin FV(e_2)$ in (6) and $x_i \notin FV(e_i)$ in (7)). The underlined freshness constraint in (5) enforces the usual “capture-avoiding” property when substituting under a λ -binder (cf. [7, Example 2.3]).

Remark 4.3 (dynamically allocated names). We can add dynamically allocated names to α ML without breaking the “names as meta-variables” aspect of its design (Remark 2.2(b)): extend its syntax with expressions \mathbf{fresh}_N of type N (for each name sort N) and its operational semantics with the impure transition:

$$(I6) \quad \exists \vec{x}(\vec{c}; s; \mathbf{fresh}_N) \rightarrow \exists \vec{x}, x(\vec{c} \ \& \ x \ \# \ \vec{x}; s; x)$$

where $x \in Var(N)$ is not in \vec{x} and $x \ \# \ \vec{x}$ is the constraint $x \ \# \ x_1 \ \& \ \dots \ \& \ x \ \# \ x_n$ when $\vec{x} = x_1, \dots, x_n$. Using this we can define a uniform operation of *generative name-unbinding*

$$\mathbf{unbinde} \ \mathbf{as} \ \langle x \rangle x' \ \mathbf{in} \ e' \triangleq \mathbf{let} \ x = \mathbf{fresh}_N \ \mathbf{in} \ \exists x' \langle x \rangle x' = e \ \& \ e' \quad (8)$$

where $e : [N]E$ and $x \in \text{Var}(N), x' \in \text{Var}(E)$ are distinct variables not occurring in e . Then, for example, we can replace the last branch of the **case** expression in (5) with $\text{L } x \Rightarrow \text{unbind } x \text{ as } \langle x_1 \rangle t_1 \text{ in } \text{L}(\langle x_1 \rangle \text{subst}(t_1, t', x'))$. Rule (I6) and definition (8) together give a version of generative unbinding that is like the one used by MLSOS [16]. It is operationally different from FreshML's version [26], which pushes a swap of x with a fresh name into the body e' . Are these two forms of generative unbinding behaviourally equivalent? To determine this requires developing the properties of *contextual equivalence* for αML , which we defer to a future paper. Is fresh_N definable up to contextual equivalence in terms of the language presented in Figs 2 and 3? It seems unlikely, but we do not have a proof.

5 α -Inductive Definitions as αML Recursive Functions

We remarked in the previous section that nominal arities are the same thing as αML equality types. Regarding the relation symbol $r <: A_r$ of nominal arity A_r as a variable of type $A_r \rightarrow \text{prop}$, we identify α -inductive definitions in standard form \mathcal{D} (Definition 2.5) with certain αML recursive function values $v_{\mathcal{D}}$ of type $A_r \rightarrow \text{prop}$:

$$v_{\mathcal{D}} \triangleq (\text{fun } r \ x = \varphi) \quad \text{where } \mathcal{D} \text{ is } \frac{\varphi}{r(x)}. \quad (9)$$

For this to make sense we have to embed formulas $\varphi \in \text{Form}_{\Sigma}$ over a nominal signature Σ (Definition 2.4) as αML expressions of type **prop**. Clearly the patterns p of each arity A (Definition 2.1) coincide with values v (Fig. 2) of equality type A . So αML syntax has all the necessary constituents for expressing formulas except possibly for conjunction and disjunction. We define conjunction as in (6) and express disjunction using a flexible **case**-expression (cf. [27, Sect. 3.1]):

$$e_1 \vee e_2 \triangleq \exists x (\text{case } x \text{ of } \text{T } x_1 \Rightarrow e_1 \mid \text{F } x_2 \Rightarrow e_2) \quad (10)$$

(where $x, x_1, x_2 \notin \text{FV}(e_1, e_2)$).

So given an α -inductive definition \mathcal{D} in standard form with associated αML function $v_{\mathcal{D}} : A_r \rightarrow \text{prop}$ as in (9), for each formula $\varphi' \in \text{Form}_{\Sigma}$ we get an αML expression $\varphi'[v_{\mathcal{D}}/r]$ of type **prop**. The following theorem characterizes satisfaction of φ' in terms of the operational behaviour of this expression. It uses the *solution set* of φ' (with respect to a set of variables \vec{x} containing those free in φ'): this is defined to be the set $\text{solns}(\varphi')$ of constraints $\exists \vec{x}'(\vec{c})$ such that $\exists \vec{x}(\text{T}; \text{id}; \varphi'[v_{\mathcal{D}}/r]) \rightarrow \dots \rightarrow \exists \vec{x}, \vec{x}'(\vec{c}; \text{id}; \text{T})$ and $\emptyset \models \exists \vec{x}', \vec{c}(\vec{c})$.

Theorem 5.1. *For any formula φ' and valuation V we have:*

Soundness: *if $\exists \vec{x}'(\vec{c}) \in \text{solns}(\varphi')$ and $V \models \exists \vec{x}'(\vec{c})$ then $(\llbracket \mathcal{D} \rrbracket, V) \models \varphi'$.*

Completeness: *if $(\llbracket \mathcal{D} \rrbracket, V) \models \varphi'$ then there is $\exists \vec{x}'(\vec{c}) \in \text{solns}(\varphi')$ with $V \models \exists \vec{x}'(\vec{c})$.*

Proof. The theorem can be deduced using standard techniques from constraint logic programming (CLP) [15, Sect. 4.4 and 4.5]. This is because one can prove that α ML's operational semantics agrees with the semantics of CLP goal states $\langle \varphi, \varphi_1, \dots, \varphi_n \mid \bar{c} \rangle$ if the latter are encoded as α ML configurations of the form $\exists \vec{x}(\bar{c}; s_{\varphi_n, \dots, \varphi_1}; \varphi)$, where the frame stack $s_{\varphi_n, \dots, \varphi_1}$ is defined by: $s_\emptyset \triangleq id$ and $s_{\bar{\varphi}, \varphi} \triangleq s_{\bar{\varphi}} \circ (x. \varphi)$ (where $x \notin FV(\varphi)$). \square

Definition 5.2 (success). We say that a configuration $\exists \vec{x}(\bar{c}; s; e) : \mathbf{prop}$ may succeed and write $\exists \vec{x}(\bar{c}; s; e) \downarrow$ if there is a finite sequence of transitions from $\exists \vec{x}(\bar{c}; s; e)$ to a configuration of the form $\exists \vec{x}, \vec{x}'(\bar{c}'; id; \top)$, for some \vec{x}' and \bar{c}' with $\emptyset \models \exists \vec{x}, \vec{x}'(\bar{c}')$.

We can use Theorem 5.1 to deduce that the operational semantics of $v_{\mathcal{D}} : A_r \rightarrow \mathbf{prop}$ in α ML detects, through the above notion of success, all and only the α -trees $t \in \alpha\text{-Tree}(A_r)$ lying in the inductively defined subset $\llbracket \mathcal{D} \rrbracket$ (Definition 2.5). To do so, we first have to discuss how α ML represents α -trees, since they involve concrete names $n \in \text{Name}(N)$ whereas α ML follows the common practice (Remark 2.2(b)) of only using variables of name sort, $x \in \text{Var}(N)$. What matters about names when they are used to describe binding structure is not their particular identity, but rather the *distinctions* between them—and those can be expressed using constraints asserting that all the variables in a list $\vec{x} = x_1, \dots, x_k$ are distinct:

$$\#_{\vec{x}} \triangleq \big\&_{1 \leq i < j \leq k} x_i \# x_j.$$

A valuation V with domain \vec{x} satisfies $\#_{\vec{x}}$ iff $V(x_1), \dots, V(x_k)$ are (α -equivalence classes of) mutually distinct names. We can represent a particular α -tree in α ML by a pattern in the presence of such a constraint: if $t \in \alpha\text{-Tree}(A_r)$ is the α -equivalence class of a tree involving k distinct names (bound or free), we can find a pattern $p : A_r$ with k variables \vec{x} of name sort and a valuation V with $t = \llbracket p \rrbracket_V$ and $V \models \#_{\vec{x}}$. Then taking φ' to be $\#_{\vec{x}} \& r(p)$ in Theorem 5.1 we get:

Corollary 5.3. *Let \mathcal{D} be an α -inductive definition in standard form with associated α ML function $v_{\mathcal{D}} : A_r \rightarrow \mathbf{prop}$. If $t \in \alpha\text{-Tree}(A_r)$ is represented as above by a pattern $p : A_r$ and a valuation V (with $\text{dom}(V) = \vec{x}$, the variables of p), then $t \in \llbracket \mathcal{D} \rrbracket$ iff $\exists \vec{x}(\#_{\vec{x}}; id; v_{\mathcal{D}} p) \downarrow$. \square*

6 Related and Future Work

A popular approach to executable operational semantics is to use *higher-order logic programming*, where binders in inductive definitions are represented via higher-order abstract syntax (HOAS): see Miller [19] for an overview. We think it is both useful and interesting to study executable operational semantics also using *functional logic programming* [14]. It has proved harder to integrate HOAS representations with functional programming: see [18] for a recent view on this. In any case, in the Introduction we advocated leaving the HOAS mainstream and pursuing a nominal approach, for reasons to do with name-aliasing. As far as we know the

first such approach was Cheney and Urban’s first-order logic programming language α Prolog [7]. Our first attempt to combine α Prolog’s computational mechanism (resolution based on nominal unification [28]) with higher-order typed functional programming was influenced by the work on FreshML [26] and produced MLSOS [16]. Byrd and Friedman’s α Kanren [2] combines it with the untyped functional language Scheme. α Prolog, MLSOS and α Kanren allow the use of constants to name object-level bound entities. We argued in Remark 2.2(b) that such concrete names are never used in practice when specifying inductively defined relations. Moreover their use naturally leads to “equivariance” (that is, invariance under permutations of concrete names) becoming an explicit part of the meta-language’s operational semantics [6], rather than just a useful meta-theoretic property of the semantics. By contrast, α ML only uses *meta-variables* rather than constants to name object-level bound entities, plus freshness constraints on meta-variables when distinctions between names are needed. As well as being closer to informal practice, this approach leads both to a pleasingly simple design for α ML’s operational semantics (Fig. 3) and a correctness result (Corollary 5.3) that was lacking for MLSOS. Our design also avoids the use of explicit name-swapping; although this is a characteristic feature of nominal logic [23], nominal unification and α Prolog, it is not needed from the point of view of a user specifying operational semantics in an executable meta-language.

The presence of unifiable meta-variables in binding position in α ML patterns does mean that, as for Cheney’s equivariant unification [4], our constraint satisfaction problem is NP-complete (Theorem 3.3). There are at least two different approaches to obtaining a practically useful implementation of α ML that should be investigated. One approach is to identify restrictions on α -inductive definitions that do not limit their applicability for specifying operational semantics too much, but for which the associated α -tree constraint problems are in P rather than NP; cf. Cheney and Urban [7, Sect. 5.3]. Since degrees of “applicability for specifying operational semantics” are hard to pin down, perhaps a more attractive alternative is to stick with the general and conceptually simple form of α -inductive definitions used in this paper, but investigate transformations on α -tree constraint problems that allow the highly developed technology of SAT-solvers to be applied.

References

- [1] Barendregt, H.P.: The Lambda Calculus: Its Syntax and Semantics. North-Holland, Amsterdam (1984) (revised edition)
- [2] Byrd, W.E., Friedman, D.P.: α Kanren: A fresh name in nominal logic programming. In: Proc. 2007 Workshop on Scheme and Functional Programming, number DIUL-RT-0701 in Université Laval Technical Reports, pp. 79–90 (2007)
- [3] Calvès, C., Fernández, M.: Nominal matching and α -equivalence. In: Hodges, W., de Queiroz, R. (eds.) Logic, Language, Information and Computation. LNCS, vol. 5110, pp. 111–122. Springer, Heidelberg (2008)
- [4] Cheney, J.: The complexity of equivariant unification. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 332–344. Springer, Heidelberg (2004)
- [5] Cheney, J.: Nominal Logic Programming. PhD thesis, Cornell Univ. (2004)

- [6] Cheney, J.: Equivariant unification. In: Giesl, J. (ed.) RTA 2005. LNCS, vol. 3467, pp. 74–89. Springer, Heidelberg (2005)
- [7] Cheney, J., Urban, C.: Nominal logic programming. *Trans. Prog. Lang. and Systems* 30(5), 1–47 (2008)
- [8] de Bruijn, N.G.: Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indag. Math.* 34, 381–392 (1972)
- [9] Felleisen, M., Hieb, R.: The revised report on the syntactic theories of sequential control and state. *Theoret. Comput. Science* 103, 235–271 (1992)
- [10] Flanagan, C., Sabry, A., Duba, B.F., Felleisen, M.: The essence of compiling with continuations. *SIGPLAN Not.* 28, 237–247 (1993)
- [11] Gabbay, M.J.: Fresh logic: Proof-theory and semantics for FM and nominal techniques. *J. Appl. Logic* 5, 356–387 (2007)
- [12] Gordon, A.D.: Operational equivalences for untyped and polymorphic object calculi. In: Gordon and Pitts [13], pp. 9–54
- [13] Gordon, A.D., Pitts, A.M. (eds.): *Higher Order Operational Techniques in Semantics*. Cambridge University Press, Cambridge (1998)
- [14] Hanus, M.: Multi-paradigm declarative languages. In: Dahl, V., Niemelä, I. (eds.) ICLP 2007. LNCS, vol. 4670, pp. 45–75. Springer, Heidelberg (2007)
- [15] Jaffar, J., Maher, M., Marriott, K., Stuckey, P.: Semantics of constraint logic programming. *J. Logic Programming* 37, 1–46 (1998)
- [16] Lakin, M.R., Pitts, A.M.: A metalanguage for structural operational semantics. In: *Trends in Functional Programming*, vol. 8, pp. 19–35. Intellect (2008)
- [17] Lassen, S.B.: Relational reasoning about contexts. In: Gordon and Pitts [13], pp. 91–135
- [18] Licata, D.R., Zeilberger, N., Harper, R.: Focusing on binding and computation. In: *LICS 2008 Proceedings*, pp. 241–252. IEEE Computer Society, Los Alamitos (2008)
- [19] Miller, D.A.: Abstract syntax for variable binders: An overview. In: Palamidessi, C., Moniz Pereira, L., Lloyd, J.W., Dahl, V., Furbach, U., Kerber, M., Lau, K.-K., Sagiv, Y., Stuckey, P.J. (eds.) CL 2000. LNCS, vol. 1861, pp. 239–253. Springer, Heidelberg (2000)
- [20] Milner, R., Tofte, M., Harper, R., MacQueen, D.: *The Definition of Standard ML (Revised)*. MIT Press, Cambridge (1997)
- [21] Pfenning, F., Elliott, C.: Higher-order abstract syntax. In: *PLDI 1988 Proceedings*, pp. 199–208. ACM Press, New York (1988)
- [22] Pitts, A.M.: Operational semantics and program equivalence. In: Barthe, G., Dybjer, P., Pinto, L., Saraiva, J. (eds.) APPSEM 2000. LNCS, vol. 2395, pp. 378–412. Springer, Heidelberg (2002)
- [23] Pitts, A.M.: Nominal logic, a first order theory of names and binding. *Inf. and Comput.* 186, 165–193 (2003)
- [24] Pitts, A.M.: Alpha-structural recursion and induction. *J. ACM* 53, 459–506 (2006)
- [25] Pitts, A.M., Shinwell, M.R.: Generative unbinding of names. *Logical Methods in Comput. Science* 4, 1–33 (2008)
- [26] Shinwell, M.R., Pitts, A.M., Gabbay, M.J.: FreshML: Programming with binders made simple. In: *ICFP 2003 Proceedings*, pp. 263–274. ACM Press, New York (2003)
- [27] Tolmach, A., Antoy, S.: A monadic semantics for core Curry. In: *WFLP 2003 Proceedings*. *Electr. Notes in Theoret. Comp. Science*, vol. 86(3), pp. 16–34. Elsevier, Amsterdam (2003)
- [28] Urban, C., Pitts, A.M., Gabbay, M.J.: Nominal unification. *Theoret. Comp. Science* 323, 473–497 (2004)