# The ∇-Calculus. Functional Programming with Higher-order Encodings.

Carsten Schürmann, Adam Poswolsky, Jeffrey Sarnat

Yale University
New Haven, CT, 06511
USA

**Abstract.** Higher-order encodings use functions provided by one language to represent variable binders of another. They lead to concise and elegant representations, which historically have been difficult to analyze and manipulate.
In this paper we present the ∇-calculus, a calculus for defining general recursive functions over higher-order encodings. To avoid problems commonly associated with using the same function space for representations and computations, we separate one from the other. The simply-typed $\lambda$-calculus plays the role of the representation-level. The computation-level contains not only the usual computational primitives but also an embedding of the representation-level. It distinguishes itself from similar systems by allowing recursion under representation-level $\lambda$-binders while permitting a natural style of programming which we believe scales to other logical frameworks. Sample programs include bracket abstraction, parallel reduction, and an evaluator for a simple language with first-class continuations.

## 1 Introduction

Higher-order abstract syntax refers to the technique of using a meta-language, or logical framework, to encode an object language in such a way that variables of the object language are represented by the variables of the logical framework. This deceptively simple idea has far reaching consequences for the design of languages that aim to manipulate these encodings. On one hand, higher-order encodings are often very concise and elegant since they take advantage of common concepts and operations automatically provided by the logical framework, including variable renaming, capture avoiding substitutions, and hypothetical judgments. On the other hand, higher-order encodings are not inductive in the usual sense, which means that they are difficult to analyze and manipulate.

Many attempts have been made to integrate advanced encoding techniques into functional programming languages. FreshML [GP99] supports implicit variable renaming for first-order encodings. The modal $\lambda$-calculus supports primitive recursion over higher-order encodings via an iterator. However, function definition via iteration is naturally limited [SDP01].

In this paper, we present the $\nabla$-calculus, a step towards integrating logical frameworks into functional programming. It supports *general* recursive functions over higher-order encodings without burdening the representational expressiveness of the logical framework. The $\nabla$-calculus distinguishes itself from similar systems by allowing recursion under representation-level $\lambda$-binders while permitting a natural style of programming, which we believe scales to other logical frameworks.

To avoid problems commonly associated with using the same function space for representations and computations, we separate one from the other. The simply-typed $\lambda$-calculus plays the role of the representation-level and provides a function space enabling higher-order encodings. A second simply-typed language plays the role of the computation-level. It provides embeddings of the higher-order encodings, function definition by cases, and insurances for safe returns from computation under representation-level $\lambda$-binders.

The resulting system allows us, for example, to write computation-level functions that recurse over the usual higher-order encoding of the untyped $\lambda$-calculus (see Example 3). It is general enough to permit case analysis over any representation-level object of any representation-level type. In the accompanying Technical Report [SPS04] the reader may find a wide collection of examples, such as translation to de Bruijn indices, parallel reduction, and an evaluator for a simple language with first-class continuations. A prototype implementation [PS04] of the $\nabla$-calculus, including a type-checker, an interactive runtime-system, and a collection of examples is available from the website `http://www.cs.yale.edu/~delphin`.

This paper is organized as follows. We explain the use of the simply-typed $\lambda$-calculus as a logical framework in Section 2. We introduce the $\nabla$-calculus in Section 3. It is divided into several subsections describing the conventional features of the $\nabla$-calculus and those constructs that facilitate programming with higher-order encodings. The static and operational semantics of the $\nabla$-calculus are given in Section 4, while the meta-theoretic properties of the calculus are discussed and analyzed in Section 5. We assess results and discuss related and future work in Section 6.

## 2 The Simply-Typed Logical Framework

We choose the simply-typed $\lambda$-calculus as our logical framework. It is not as expressive as dependently-typed frameworks, such as LF [HHP93], but is expressive enough to permit interesting higher-order encodings.

$$
\begin{array}{lll}
\text{Types:} & A, B & ::= a \mid A \to B \\
\text{Objects:} & M, N & ::= x \mid c \mid \lambda x : A.\, M \mid M\ N \\
\text{Signatures:} & \Sigma & ::= \cdot \mid \Sigma, a : \mathit{type} \mid \Sigma, c : A \\
\text{Contexts:} & \Gamma & ::= \cdot \mid \Gamma, x : A
\end{array}
$$

We use $a$ for type constants, $c$ for object constants, and $x$ for variables. We assume that constants and variables are declared at most once in a signature

and context, respectively. To maintain this invariant, we tacitly rename bound variables and use capture-avoiding substitutions. The typing judgments for objects and signatures are standard. Type-level and term-level constants must be declared in the signature.

**Definition 1 (Typing judgment).** *$\Gamma \vdash M : A$ is defined by the following rules:*

$$\frac{\Gamma(x) = A}{\Gamma \vdash x : A}\ \mathsf{ofvar} \qquad \frac{\Sigma(c) = A}{\Gamma \vdash c : A}\ \mathsf{ofconst}$$

$$\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x : A.\, M : A \to B}\ \mathsf{oflam} \qquad \frac{\Gamma \vdash M : A \to B \quad \Gamma \vdash N : A}{\Gamma \vdash M\ N : B}\ \mathsf{ofapp}$$

Our notion of definitional equality is obtained by taking the reflexive, transitive, and symmetric closure of $\beta$- and $\eta$-conversion [Coq91]. We write $\Gamma \vdash M \equiv N : A$ if and only if $M$ is $\beta\eta$-equivalent to $N$ and both have type $A$. For every well-typed object $M$ of type $A$, there exists a unique $\beta$-normal, $\eta$-long term $M'$ such that $\Gamma \vdash M \equiv M' : A$ [Pfe92]. We refer to $M'$ as being canonical, which we denote as $\Gamma \vdash M' \Uparrow A$.

Throughout this paper, our examples will use encodings of natural numbers, first-order logic, and the untyped $\lambda$-calculus. An encoding consists of a signature and a representation function, which maps elements from our domain of discourse into canonical forms in our logical framework. We say that an encoding is *adequate* if the representation function is an isomorphism.

In all of the examples below, the signatures for our encoding are listed in italics and our translation functions $\ulcorner - \urcorner$ are defined by the given sets of equations.

*Example 1 (Natural numbers).*

$$\begin{array}{rll} & \mathit{nat} & : \mathit{type} \\ \ulcorner 0 \urcorner = z & z & : \mathit{nat} \\ \ulcorner n+1 \urcorner = s\, \ulcorner n \urcorner & s & : \mathit{nat} \to \mathit{nat} \end{array}$$

□

Example 1 is a first-order encoding because none of the constants take arguments of functional types.

*Example 2 (First order logic with equality).* Terms $t ::= x$ and first order formulas $F ::= \forall x.\, F \mid F_1 \supset F_2 \mid \neg F \mid t_1 = t_2$ are represented as objects of type $i$ and type $o$, respectively, in a signature that also includes the following declarations:

$$\begin{array}{rll} \ulcorner \forall x.\, F \urcorner = \mathit{forall}\ (\lambda x : i. \ulcorner F \urcorner) & \mathit{forall} & : (i \to o) \to o \\ \ulcorner \neg F \urcorner = \mathit{neg}\ \ulcorner F \urcorner & \mathit{neg} & : o \to o \\ \ulcorner F_1 \supset F_2 \urcorner = \mathit{impl}\ \ulcorner F_1 \urcorner\ \ulcorner F_2 \urcorner & \mathit{impl} & : o \to o \to o \\ \ulcorner t_1 = t_2 \urcorner = \mathit{eq}\ \ulcorner t_1 \urcorner\ \ulcorner t_2 \urcorner & \mathit{eq} & : i \to i \to o \\ \ulcorner x \urcorner = x & & \end{array}$$

*Example 3 (Untyped λ-expressions).* Untyped λ-expressions $e ::= x \mid \mathbf{lam}\ x.\, e \mid e_1\ e_2$ are encoded as follows:

$$\begin{array}{rll} & & \mathit{exp} : \mathit{type} \\ \ulcorner \mathbf{lam}\ x.\, e \urcorner = \mathit{lam}\ (\lambda x : \mathit{exp}. \ulcorner e \urcorner) & & \mathit{lam} : (\mathit{exp} \to \mathit{exp}) \to \mathit{exp} \\ \ulcorner e_1\ e_2 \urcorner = \mathit{app}\ \ulcorner e_1 \urcorner\ \ulcorner e_2 \urcorner & & \mathit{app} : \mathit{exp} \to \mathit{exp} \to \mathit{exp} \\ \ulcorner x \urcorner = x & & \end{array}$$

The encodings of first-order formulas and λ-calculus expressions illustrate the use of higher-order abstract syntax since object-language variable-binders use logical-framework functions. Because little meaningful analysis can be done on variables in our logical framework, the only interesting operation that can be performed on a variable is substitution. Thus, it is most helpful to think of a term of type $A \to B$ not as representing a computation, but as representing a term of type $B$ that has a hole of type $A$.

We demonstrate the formulation of an adequacy theorem. Each case can be proven by a straightforward induction.

**Theorem 1 (Adequacy of *exp*).** *Adequacy holds for our representation of untyped λ-expressions.*

1. *If $e$ is an expression with free variables among $x_1, \ldots, x_n$, then $x_1 : \mathit{exp}, \ldots, x_n : \mathit{exp} \vdash \ulcorner e \urcorner \Uparrow \mathit{exp}$.*
2. *If $x_1 : \mathit{exp}, \ldots, x_n : \mathit{exp} \vdash M \Uparrow \mathit{exp}$ then $M = \ulcorner e \urcorner$ for some expression $e$ with free variables among $x_1, \ldots, x_n$.*
3. *$\ulcorner - \urcorner$ is a bijection between expressions and canonical forms where $\ulcorner [e'/x]e \urcorner = [\ulcorner e' \urcorner / x] \ulcorner e \urcorner$.* □

## 3 The ∇-Calculus

The logical-framework type *exp* is not inductive because the constructor $\mathit{lam} : (\mathit{exp} \to \mathit{exp}) \to \mathit{exp}$ has a negative occurrence [PM93] of *exp*. This is not just a formal observation, since this property has deep consequences for the design of the ∇-calculus, which needs to provide a notion of computation general enough to handle higher-order datatypes of this kind. We offer the ability to recurse under λ-binders and consider cases over functions of type $\mathit{exp} \to \mathit{exp}$ while continuing to guarantee the adequacy of the encoding. Allowing for this, as well as general recursive computation, can be seen as the main contribution of this work.

In the ∇-calculus, expressions permit function definition by cases and alternations instead of providing explicit λ-binders on the computation-level. Computation-level expressions and types are summarized in Figure 1 and explained in the remainder of this section.

### 3.1 Function Definition by Cases and Recursion

In the ∇-calculus, we draw a separating line between the levels of representation and computation. Representation-level types, such as *nat* and *exp* are

$$
\begin{array}{lll}
\text{Types:} & \tau, \sigma ::= & \langle A \rangle \mid \tau \Rightarrow \sigma \mid \Box\tau \\
\text{Expressions:} & e, f ::= & u \mid \langle M \rangle \mid e_1 \mapsto_\tau e_2 \mid \epsilon x : A.\, e \mid \epsilon u \in \tau.\, e \\
& & \mid e_1 \cdot e_2 \mid (e_1 \mid e_2) \mid \text{rec } u \in \tau.\, e \\
& & \mid \nu x : A.\, e \mid \text{pop } e \mid \nabla x : A.\, e
\end{array}
$$

**Fig. 1.** Syntactic categories of the $\nabla$-calculus.

*injected* into computation-level types $\langle nat \rangle$ and $\langle exp \rangle$. Likewise, representation-level constants, such as $(s\ z)$ and $lam\ (\lambda x : exp.\, x)$, are injected into computation-level terms $\langle s\ z \rangle$ and $\langle lam\ (\lambda x : exp.\, x) \rangle$. There are no user defined datatypes on the computation-level; all type and constant declarations must be done at the representation-level.

*Example 4 (Addition).* We informally define the function *plus* over the representation of natural numbers from Example 1 in the following manner:

$$
\begin{array}{l}
plus\ z\ y \quad\ = y \\
plus\ (s\ x)\ y = s\ (plus\ x\ y)
\end{array}
$$

We represent this formally in the $\nabla$-calculus as follows:

$$
\begin{array}{l}
\text{rec } plus \in \langle \text{nat} \rangle \Rightarrow \langle \text{nat} \rangle \Rightarrow \langle \text{nat} \rangle. \\
\quad \epsilon y : \text{nat}.\, \langle \text{z} \rangle \mapsto \langle y \rangle \mapsto \langle y \rangle \\
\quad \mid\ \epsilon x : \text{nat}.\, \langle \text{s}\ x \rangle \mapsto \epsilon y : \text{nat}.\, \langle y \rangle \mapsto \langle \text{s} \rangle \circ (plus \cdot \langle x \rangle \cdot \langle y \rangle)
\end{array}
$$

□

The recursion operator is conventional. In later examples we will omit it for the sake of readability. Alternation, "|", separates cases that may be chosen for evaluation non-deterministically. It binds more tightly than the recursion operator rec $u \in \tau.\, e$, but not as tight as any of the other operators. Individual cases are of the form $e_1 \mapsto_\tau e_2$, where $e_1$ can be thought of as a guard. Only when such a case is applied to an object equivalent to $e_1$ (as defined in Section 4.2) is $e_2$ evaluated. In particular, if $e_1$ is a value of type $\langle A \rangle$, then our notion of equality is given by our logical framework's notion of definitional equality. We refer to $e_1$ as the *pattern* and $e_2$ as the *body* of the case. The index $\tau$ states the type of the pattern, but is usually omitted when the type of the pattern can be easily inferred. In conventional programming languages, variables that occur in patterns are implicitly declared, whereas in the $\nabla$-calculus they must be declared explicitly by $\epsilon x : A.\, e$ for reasons explained in Section 3.2. A similar declaration for the computation-level $\epsilon u \in \tau.\, e$ permits higher-order functions and is discussed in detail in the accompanying Technical Report [SPS04]. Application in the $\nabla$-calculus is written as $e_1 \cdot e_2$ in order to avoid confusion with

representation-level application, which is expressed via juxtaposition. The notation $e_1 \circ e_2$ is syntactic sugar that lifts representation-level application to the computation-level.

$$e_1 \circ_{A,B} e_2 = \epsilon x : A \to B.\, \langle x \rangle \mapsto_{\langle A \to B \rangle} \epsilon y : A.\, \langle y \rangle \mapsto_{\langle A \rangle} \langle x\ y \rangle$$

We refer to $\circ$ without type annotations because they are easily inferable.

### 3.2 Traversal of λ-Binders

Next, we explain the operators $\nu$ and pop from Figure 1. Recall the encoding of first-order logic from Example 2.

As a running example, we consider Kolmogorov's double-negation interpretation, which transforms formulas from classical logic into intuitionistic logic in the following way:

$$\begin{array}{ll} \mathit{dneg}\,(\mathit{eq}\,t_1\ t_2) & = \mathit{neg}\,(\mathit{neg}\,(\mathit{eq}\,t_1\ t_2)) \\ \mathit{dneg}\,(\mathit{impl}\,F_1\ F_2) & = \mathit{neg}\,(\mathit{neg}\,(\mathit{impl}\,(\mathit{dneg}\,F_1)\ (\mathit{dneg}\,F_2))) \\ \mathit{dneg}\,(\mathit{neg}\,F) & = \mathit{neg}\,(\mathit{neg}\,(\mathit{neg}\,(\mathit{dneg}\,F))) \\ \mathit{dneg}\,(\mathit{forall}\,F) & = \mathit{neg}\,(\mathit{neg}\,(\mathit{forall}\,F')) \\ & \qquad \text{where } F'\ x = \mathit{dneg}\,(F\ x) \\ & \qquad \text{for some new parameter } x : i \end{array}$$

In the last case *dneg* must recurse on the body $F$ of the *forall* term, which is a representation-level function of type $i \to o$. Since $F$ is definitionally equivalent to a canonical term that starts with a $\lambda$-binder, we strip away the $\lambda$-binder by applying $F$ to some new parameter $x$ before invoking *dneg*. The result of the computation depends on $x$ and is hence written as $F'\ x$, where $F'$ is a representation-level function of type $i \to o$.

The first three cases of *dneg* can be implemented in the $\nabla$-calculus with constructs we have already introduced. As for the *forall* case, we need to add new constructs to our language. We feel that there are several interesting possibilities worth considering. One possibility would be to introduce a computation-level operator $\hat{\lambda}$, which lifts representation-level abstraction to the computation-level in much the same way that the syntactic-sugar $\circ$ lifts representation-level application. In this case, we could write the *forall* case as

$$\epsilon F : \mathrm{i} \to \mathrm{o}.\, \langle \mathrm{forall}\ F \rangle \mapsto \langle \mathrm{neg} \rangle \circ (\langle \mathrm{neg} \rangle \circ (\langle \mathrm{forall} \rangle \circ (\hat{\lambda} x : \mathrm{i}.\, \mathit{dneg} \cdot \langle F\ x \rangle)))$$

where the subterm $(\hat{\lambda} x : \mathrm{i}.\, \mathit{dneg} \cdot \langle F\ x \rangle)$ has type $\langle i \to o \rangle$. In principle this is a possible solution. Adequacy is preserved because although the body of $\hat{\lambda}$ may diverge or get stuck, any value it computes must be of the form $\langle M \rangle$. However, $\hat{\lambda}$ is too limited for our purposes because it always returns a representation-level function, even if the expected result is of a base type (see Example 5). Meta-ML [TS00] employs a construct similar to $\hat{\lambda}$.

Another possibility is to add an explicit parameter introduction operator $\bar{\lambda}$

$$\begin{array}{l}\epsilon F : \mathrm{i} \to \mathrm{o}.\, \langle \mathrm{forall}\ F \rangle \mapsto \bar{\lambda} x : \mathrm{i}. \\ \qquad \mathrm{case}\ \mathit{dneg} \cdot \langle F\ x \rangle \\ \qquad\quad \mathrm{of}\ \epsilon F' : \mathrm{i} \to \mathrm{o}.\, \langle F'\ x \rangle \mapsto \langle \mathrm{neg}\ (\mathrm{neg}\ (\mathrm{forall}\ F')) \rangle\end{array}$$

where we write "case $e_1$ of $e_2$" as syntactic sugar for "$e_2 \cdot e_1$". In contrast to $\hat{\lambda}$, the type of the subterm starting with $\bar{\lambda}$ is $\langle o \rangle$. Since the recursive call results in a value of type $\langle o \rangle$, and *forall* requires a value of type $i \to o$, we need a way to turn the result into a value of type $\langle i \to o \rangle$. Furthermore, because this value escapes $x$'s declaration, it should not contain any free occurrences of $x$. Ideally, higher-order pattern matching would yield $F'$, which is the result of abstracting all occurrences of $x$ from the result of the recursive call. But there is no guarantee that this will succeed, because $F'$ is declared within the scope of $x$. For example, if $\mathit{dneg} \cdot \langle F\ x \rangle$ returns $\langle \mathrm{eq}\ x\ x \rangle$, then $F' = (\lambda y : i.\, \mathrm{eq}\ x\ x)$ and $F' = (\lambda y : i.\, \mathrm{eq}\ y\ y)$ are among the possible solutions to this matching problem. To remedy this, $F'$ can be declared outside of the scope of $x$, and thus could not possibly be instantiated with a term containing $x$:

$$\begin{array}{l}\epsilon F : \mathrm{i} \to \mathrm{o}.\, \langle \mathrm{forall}\ F \rangle \mapsto \epsilon F' : \mathrm{i} \to \mathrm{o}. \\ \qquad \bar{\lambda} x : \mathrm{i}.\mathrm{case}\ \mathit{dneg} \cdot \langle F\ x \rangle\ \mathrm{of}\ \langle F'\ x \rangle \mapsto \langle \mathrm{neg}\ (\mathrm{neg}\ (\mathrm{forall}\ F')) \rangle\end{array}$$

In this case, the only solution to the matching problem is $F' = (\lambda y : i.\, \mathrm{eq}\ y\ y)$, which illustrates the necessity of explicit $\epsilon$-declarations. However, we do not include $\bar{\lambda}$ in the $\nabla$-calculus since, as we have seen, it allows us to write functions that let parameters escape their scope.

Instead, we do include two operators and one new type constructor that can be found in Figure 1. The operator $\nu$ is similar to $\bar{\lambda}$ in that it introduces new parameters, but different because it statically requires that these parameters cannot extrude their scope. The operator "pop" provides such guarantees. These guarantees are communicated through the type $\Box\tau$, which pop introduces and $\nu$ eliminates. The complete function *dneg* is given below.

$$\begin{array}{l}\mathit{dneg}\ :\ \langle \mathrm{o} \rangle \Rightarrow \langle \mathrm{o} \rangle \\ \quad = \epsilon t_1 : \mathrm{i}.\, \epsilon t_2 : \mathrm{i}.\, \langle \mathrm{eq}\ t_1\ t_2 \rangle \mapsto \langle \mathrm{neg}\ (\mathrm{neg}\ (\mathrm{eq}\ t_1\ t_2)) \rangle \\ \qquad \mid\ \epsilon F_1 : \mathrm{o}.\, \epsilon F_2 : \mathrm{o}. \\ \qquad\quad \langle \mathrm{imp}\ F_1\ F_2 \rangle \mapsto \langle \mathrm{neg} \rangle \circ (\langle \mathrm{neg} \rangle \circ (\langle \mathrm{imp} \rangle \circ (\mathit{dneg} \cdot \langle F_1 \rangle) \circ (\mathit{dneg} \cdot \langle F_2 \rangle))) \\ \qquad \mid\ \epsilon F : \mathrm{o}.\, \langle \mathrm{neg}\ F \rangle \mapsto \langle \mathrm{neg} \rangle \circ (\langle \mathrm{neg} \rangle \circ (\langle \mathrm{neg} \rangle \circ (\mathit{dneg} \cdot \langle F \rangle))) \\ \qquad \mid\ \epsilon F : \mathrm{i} \to \mathrm{o}.\, \langle \mathrm{forall}\ F \rangle \mapsto \epsilon F' : \mathrm{i} \to \mathrm{o}. \\ \qquad\quad \nu x : \mathrm{i}.\, \mathrm{case}\ \mathit{dneg} \cdot \langle F\ x \rangle\ \mathrm{of}\ \langle F'\ x \rangle \mapsto \mathrm{pop}\ \langle \mathrm{neg}\ (\mathrm{neg}\ (\mathrm{forall}\ F')) \rangle\end{array}$$

The body of the $\nu$ is of type $\Box\langle o \rangle$; the $\Box$ ensures that whatever value this expression evaluates to does not contain $x$. The body of pop has type $\langle o \rangle$ only because it neither contains $x$ nor any $\epsilon$-quantified variable the may depend on $x$. Thus, the subexpression "pop $\langle \mathrm{forall}\ F' \rangle$" introduces type $\Box\langle o \rangle$. A precise type theoretic definition and analysis of the $\Box$ type will be given in Section 4.

### 3.3 Pattern-matching Parameters

Finally, we turn to the last unexplained operator from Figure 1, the $\nabla$-operator, which is used to match parameters introduced by $\nu$.

*Example 5 (Counting variable occurrences).* Consider a function that counts the number of occurrences of bound variables in an untyped $\lambda$-expression from Example 3.

$$\begin{array}{ll} \mathit{cntvar}\,(x) & = (s\ z) \text{ where } x : \mathit{exp} \text{ is a parameter} \\ \mathit{cntvar}\,(\mathit{app}\ e_1\ e_2) & = \mathit{plus}\,(\mathit{cntvar}\,e_1)\,(\mathit{cntvar}\,e_2) \\ \mathit{cntvar}\,(\mathit{lam}\,e) & = \mathit{cntvar}\,(e\ x) \text{ for some new parameter } x : \mathit{exp} \end{array}$$

The first of the three cases corresponds to the parameter case that matches *any* parameter of type *exp* regardless of where and when it was introduced. Formally, we use the $\nabla$-operator to implement this case.

$$\begin{array}{l} \mathit{cntvar}\ :\ \langle \mathrm{exp} \rangle \Rightarrow \langle \mathrm{nat} \rangle \\ \quad = \nabla x : \mathrm{exp}.\, \langle x \rangle \mapsto \langle \mathrm{s\ z} \rangle \\ \quad\ \mid\ \epsilon e_1 : \mathrm{exp}.\, \epsilon e_2 : \mathrm{exp}. \\ \quad\quad \langle \mathrm{app}\ e_1\ e_2 \rangle \mapsto \mathit{plus} \cdot (\mathit{cntvar} \cdot \langle e_1 \rangle) \cdot (\mathit{cntvar} \cdot \langle e_2 \rangle) \\ \quad\ \mid\ \epsilon e : \mathrm{exp} \to \mathrm{exp}. \\ \quad\quad \langle \mathrm{lam}\ e \rangle \mapsto \epsilon n : \mathrm{nat}. \\ \quad\quad\ \nu x : \mathrm{exp}. \\ \quad\quad\quad (\langle n \rangle \mapsto \mathrm{pop}\ \langle n \rangle) \cdot (\mathit{cntvar} \cdot \langle e\ x \rangle) \end{array}$$

□

Notice that, in the above example, if we were to replace the $\nabla$ with $\epsilon$, it would still be possible for *cntvar* to return correct answers, since $\epsilon x : \mathit{exp}$ can match any expression of type *exp* including parameters; however, it would also be possible for *cntvar* to always return $\langle s\ z \rangle$ for the same reason.

*Example 6 (Combinators).* The combinators $c ::= \mathbf{S} \mid \mathbf{K} \mid \mathbf{MP}\ c_1\ c_2$ are represented as objects of type *comb* as follows:

$$\begin{array}{rll} \ulcorner \mathbf{K} \urcorner = K & & K\ : \mathit{comb} \\ \ulcorner \mathbf{S} \urcorner = S & & S\ : \mathit{comb} \\ \ulcorner \mathbf{MP}\ c_1\ c_2 \urcorner = \mathit{MP} \ulcorner c_1 \urcorner \ulcorner c_2 \urcorner & & \mathit{MP} : \mathit{comb} \to \mathit{comb} \to \mathit{comb} \end{array}$$

Any simply-typed $\lambda$-expression from Example 3 can be converted into a combinator in a two-step algorithm. The first step is called bracket abstraction, or *ba*, which converts a parametric combinator (a representation-level function of type $\mathit{comb} \to \mathit{comb}$) into a combinator with one less parameter (of type *comb*). If $M$ has type $\mathit{comb} \to \mathit{comb}$ and $N$ has type *comb* then $\langle \mathit{MP} \rangle \circ (\mathit{ba} \cdot \langle M \rangle) \circ \langle N \rangle$ results in a term that is equivalent to $\langle MN \rangle$ in combinator logic.

$$\begin{array}{l} \mathit{ba}\,(\lambda x : \mathit{comb}.\, x) = \mathit{MP}\ (\mathit{MP}\ S\ K)\ K \\ \mathit{ba}\,(\lambda x : \mathit{comb}.\, z) = \mathit{MP}\ K\ z \text{ where } z : \mathit{comb} \text{ is a parameter} \\ \mathit{ba}\,(\lambda x : \mathit{comb}.\, K) = \mathit{MP}\ K\ K \\ \mathit{ba}\,(\lambda x : \mathit{comb}.\, S) = \mathit{MP}\ K\ S \\ \mathit{ba}\,(\lambda x : \mathit{comb}.\, \mathit{MP}\,(c_1\ x)\,(c_2\ x)) = \mathit{MP}\,(\mathit{MP}\ S\,(\mathit{ba}\ c_1))\,(\mathit{ba}\ c_2) \end{array}$$

$$
\begin{array}{l}
ba \;:\; \langle \text{comb} \to \text{comb} \rangle \Rightarrow \langle \text{comb} \rangle \\
\quad = \langle \lambda x : \text{comb}.\, x \rangle \mapsto \langle \text{MP (MP S K) K} \rangle \\
\qquad \mid\; \nabla z : \text{comb}.\, \langle \lambda x : \text{comb}.\, z \rangle \mapsto \langle \text{MP K } z \rangle \\
\qquad \mid\; \langle \lambda x : \text{comb}.\, \text{K} \rangle \mapsto \langle \text{MP K K} \rangle \\
\qquad \mid\; \langle \lambda x : \text{comb}.\, \text{S} \rangle \mapsto \langle \text{MP K S} \rangle \\
\qquad \mid\; \epsilon c_1 : \text{comb} \to \text{comb}.\, \epsilon c_2 : \text{comb} \to \text{comb}. \\
\qquad\quad \langle \lambda x : \text{comb}.\, \text{MP } (c_1\ x)\ (c_2\ x) \rangle \mapsto \\
\qquad\qquad \langle \text{MP} \rangle \circ (\langle \text{MP} \rangle \circ \langle \text{S} \rangle \circ (ba \cdot \langle c_1 \rangle)) \circ (ba \cdot \langle c_2 \rangle)
\end{array}
$$

The first two cases of *ba* illustrate how to distinguish $x$, which is to be abstracted, from parameters that are introduced in the function *convert*, which we discuss next. The function *convert* traverses $\lambda$-expressions and uses *ba* to convert them into combinators.

$$
\begin{array}{l}
\mathit{convert}\ (y\ z) = z \text{ where } y : \mathit{comb} \to \mathit{exp} \text{ and } z : \mathit{comb} \text{ are parameters} \\
\mathit{convert}\ (\mathit{app}\ e_1\ e_2) = \mathit{MP}\ (\mathit{convert}\ e_1)\ (\mathit{convert}\ e_2) \\
\mathit{convert}\ (\mathit{lam}\ e) = \mathit{ba}\ c \text{ where } c\ z = \mathit{convert}\ (e\ (y\ z)) \\
\qquad\qquad\qquad\qquad\qquad \text{and } y : \mathit{comb} \to \mathit{exp} \\
\qquad\qquad\qquad\qquad\qquad \text{and } z : \mathit{comb} \text{ are parameters}
\end{array}
$$

The last case illustrates how a parameter of functional type may introduce information to be used when the parameter is matched. Rather than introduce a parameter $x$ of type *exp*, we introduce a parameter of type $\mathit{comb} \to \mathit{exp}$ that carries a combinator as "payload." In our example, the payload is another parameter $z : \mathit{comb}$, the image of $x$ under *convert*. This technique is applicable to a wide range of examples (see the Technical Report [SPS04] for details). We formalize *convert* below:

$$
\begin{array}{l}
\mathit{convert} \;:\; \langle \text{exp} \rangle \Rightarrow \langle \text{comb} \rangle \\
\qquad = \nabla y : \text{comb} \to \text{exp}.\, \nabla z : \text{comb}.\, \langle y\ z \rangle \mapsto \langle z \rangle \\
\qquad\quad \mid\; \epsilon e_1 : \text{exp}.\, \epsilon e_2 : \text{exp}. \\
\qquad\qquad\quad \langle \text{app } e_1\ e_2 \rangle \mapsto \langle \text{MP} \rangle \circ (\mathit{convert} \cdot \langle e_1 \rangle) \circ (\mathit{convert} \cdot \langle e_2 \rangle) \\
\qquad\quad \mid\; \epsilon e : \text{exp} \to \text{exp}.\, \langle \text{lam } e \rangle \mapsto \epsilon c : \text{comb} \to \text{comb}. \\
\qquad\qquad\quad \nu y : \text{comb} \to \text{exp}.\, \nu z : \text{comb}. \\
\qquad\qquad\qquad \text{case } \mathit{convert} \cdot \langle e\ (y\ z) \rangle \text{ of } \langle c\ z \rangle \mapsto \text{pop (pop } (\mathit{ba} \cdot \langle c \rangle))
\end{array}
$$

□

We summarize a few of the most important properties of the $\nabla$-operator. First, it is intuitively appealing to have one base case (the $\nabla$-case) for each class of parameters, because what happens in these cases is uniquely defined in one place. Second, payload carrying parameters permit sophisticated base cases, which simplify the reading of a program because all information shared between the introduction and matching of parameters must be made explicit.

$$\frac{\Phi(u) = \tau}{\Omega, (\Gamma; \Phi) \vdash u \in \tau}\ \textsf{tpvar} \qquad \frac{\Gamma \vdash M : A}{\Omega, (\Gamma; \Phi) \vdash \langle M \rangle \in \langle A \rangle}\ \textsf{tpinj}$$

$$\frac{\Omega \vdash e_1 \in \tau \qquad \Omega \vdash e_2 \in \sigma}{\Omega \vdash e_1 \mapsto_\tau e_2 \in \tau \to \sigma}\ \textsf{tpfun} \qquad \frac{\Omega, (\Gamma; \Phi, u \in \tau) \vdash e \in \tau}{\Omega, (\Gamma; \Phi) \vdash \text{fix}\ u \in \tau.\, e \in \tau}\ \textsf{tpfix}$$

$$\frac{\Omega, (\Gamma, x : A; \Phi) \vdash e \in \tau}{\Omega, (\Gamma; \Phi) \vdash \epsilon x : A.\, e \in \tau}\ \textsf{tptheobj} \qquad \frac{\Omega, (\Gamma; \Phi, u \in \tau) \vdash e \in \tau}{\Omega, (\Gamma; \Phi) \vdash \epsilon u \in \tau.\, e \in \tau}\ \textsf{tpthemeta}$$

$$\frac{\Omega \vdash e_1 \in \sigma \to \tau \qquad \Omega \vdash e_2 \in \sigma}{\Omega \vdash e_1 \cdot e_2 \in \tau}\ \textsf{tpapp} \qquad \frac{\Omega \vdash e_1 \in \tau \qquad \Omega \vdash e_2 \in \tau}{\Omega \vdash (e_1 \mid e_2) \in \tau}\ \textsf{tpalt}$$

$$\frac{\Omega \vdash e \in \tau}{\Omega, (\Gamma; \Phi) \vdash \text{pop}\ e \in \Box\tau}\ \textsf{tpbox} \qquad \frac{\Omega, (\Gamma; \Phi), (\Gamma, x : A; \Phi) \vdash e \in \Box\tau}{\Omega, (\Gamma; \Phi) \vdash \nu x : A.\, e \in \tau}\ \textsf{tpnew}$$

$$\frac{\Omega, (\Gamma, x : A; \Phi) \vdash e \in \tau}{\Omega, (\Gamma; \Phi) \vdash \nabla x : A.\, e \in \tau}\ \textsf{tpnabla}$$

**Fig. 2.** The static semantics of the $\nabla$-calculus

## 4 Semantics

The operators $\nu$ and pop have guided the design of the static and operational semantics of the $\nabla$-calculus. To reiterate, once a parameter is introduced by a $\nu$, all other declarations that take place within its scope may depend on the new parameter. As we will see, pop statically ensures that an expression is valid outside $\nu$'s scope by discarding all declarations since the last parameter introduction in a manner reminiscent of popping elements off a stack. The ambient environment is therefore formally captured in form of *scope stacks*. A scope consists of two parts: The context $\Gamma$ (defined in Section 2), which summarizes all object-level declarations $x : A$, and the context $\Phi$, which summarizes all meta-level declarations $u \in \tau$.

$$\text{Meta Contexts: } \Phi ::= \cdot \mid \Phi, u \in \tau$$
$$\text{Scope Stacks: } \quad \Omega ::= \cdot \mid \Omega, (\Gamma; \Phi)$$

We refer to the top and second-from-top elements of $\Omega$ as the *current* and *previous* scopes, respectively. The scope stack $\Omega$ grows monotonically, which means that the current scope always extends the previous scope.

### 4.1 Static semantics

We define the typing judgment $\Omega \vdash e \in \tau$ by the rules depicted in Figure 2. Many of the rules are self-explanatory. All rules except for tpnew and tppop touch

only the current scope. For example, tpvar relates variables and types, whereas tpinj enforces that only representation-level objects valid in the current scope can be lifted to the computation-level. For functions, the pattern must be of the argument type, whereas the body be of the result type. Variables that may occur in patterns must be declared by a preceding $\epsilon x : A$ or $\epsilon u \in \tau$ declaration, which will be recorded in the current scope by tptheobj and tpthemeta, respectively. The rules tpapp, tpalt, and tpfix are standard. The tpbox rule is the introduction rule for $\Box\tau$. The expression pop $e$ is valid if $e$ is valid in the previous scope. The corresponding elimination rule is tpnew. The expression $\nu x : A.\, e$ has type $\tau$ when $e$ is of type $\Box\tau$ in the properly extended scope stack.

### 4.2 Operational Semantics

Computation level function application in the $\nabla$-calculus is more demanding than the usual substitution of an argument for a free variable. It relies on the proper instantiation of all $\epsilon$- and $\nabla$-bound variables that occur in the function's pattern. Perhaps not surprisingly, the behavior of our calculus depends on when these instantiations are committed. For example,

$$(\epsilon f \in \langle nat \rangle \to \langle nat \rangle.\, f \mapsto plus \cdot (f \cdot \langle z \rangle) \cdot (f \cdot \langle s\ z \rangle)) \cdot (\epsilon n : nat.\, \langle n \rangle \mapsto \langle n \rangle)$$

may either return $s\ z$ under a call-by-name semantics, or no solution at all under a call-by-value semantics because $n : nat$ may be instantiated either by $z$ or $s\ z$ but not both. Consequently, our calculus adopts a call-by-name evaluation strategy. We can define computational-level $\lambda$-abstraction "lambda $u \in \tau.\, e$" as syntactic sugar for $(\epsilon u \in \tau.\, u \mapsto e)$ and "let $u \in \tau = e_1$ in $e_2$ end" as syntactic sugar for $((\epsilon u \in \tau.\, u \mapsto e_2)\ e_1)$.

**Definition 2 (Values).** *The set of values of the $\nabla$-calculus is defined as follows.*

$$\textit{Values: } v ::= \langle M \rangle \mid \text{pop } e \mid e_1 \mapsto_\tau e_2$$

The operational semantics of the $\nabla$-calculus combines a system of reduction rules of the form $\Omega \vdash e \to e'$ with an equivalence relation on meta-level expressions $\Omega \vdash e \equiv e' \in \tau$. We give the reduction rules in Figure 3 and the equality rules in Figure 4. During runtime, all $\epsilon$-quantified variables are instantiated with concrete objects, so evaluation always takes place in a scope stack of the form $\Omega ::= \cdot \mid \Omega, (\Gamma; \cdot)$, where $\Gamma$ contains only $\nu$-quantified parameter declarations.

The rules in Figure 3 are organized into three parts. The top part shows the essential reduction rules redbeta and rednupop. The rule rednupop states that it is unnecessary to traverse into a new scope to return an expression that is valid in the previous scope.

Among the second block of rules, $\mathsf{redalt}_1$ and $\mathsf{redalt}_2$ express a non-deterministic choice in the control flow. Similarly, redsome and redsomeM express a non-deterministic choice of instantiations. The abbreviations $f/u$ and $M/x$ stand for single-point substitutions that can easily be expanded into simultaneous substitutions given in Definition 3. During evaluation, the current scope only contains

$$\frac{\Omega \vdash e_1 \equiv e_1' \in \tau}{\Omega \vdash (e_1 \mapsto_\tau e_2) \cdot e_1' \to e_2}\ \mathsf{redbeta} \qquad \frac{}{\Omega \vdash \nu x : A.\, \mathrm{pop}\ e \to e}\ \mathsf{rednupop}$$

$$\frac{}{\Omega \vdash (e_1 \mid e_2) \to e_1}\ \mathsf{redalt_1} \qquad \frac{}{\Omega \vdash (e_1 \mid e_2) \to e_2}\ \mathsf{redalt_2}$$

$$\frac{\Gamma \vdash M : A}{\Omega, (\Gamma; \cdot) \vdash \epsilon x : A.\, e \to [M/x]e}\ \mathsf{redsome} \qquad \frac{\Omega \vdash f \in \tau}{\Omega \vdash \epsilon u \in \tau.\, e \to [f/u]e}\ \mathsf{redsomeM}$$

$$\frac{\Gamma(y) = A}{\Omega, (\Gamma; \cdot) \vdash \nabla x : A.\, e \to [y/x]e}\ \mathsf{rednabla} \qquad \frac{}{\Omega \vdash \mathrm{fix}\ u \in \tau.\, e \to [\mathrm{fix}\ u \in \tau.\, e/u]e}\ \mathsf{redfix}$$

$$\frac{\Omega \vdash e_1 \to e_1'}{\Omega \vdash e_1 \cdot e_2 \to e_1' \cdot e_2}\ \mathsf{redfun} \qquad \frac{\Omega, (\Gamma; \cdot), (\Gamma, x : A; \cdot) \vdash e \to e'}{\Omega, (\Gamma; \cdot) \vdash \nu x : A.\, e \to \nu x : A.\, e'}\ \mathsf{rednew}$$

**Fig. 3.** Small-step semantics (Reductions)

$$\frac{\Gamma \vdash M \equiv N : A}{\Omega, (\Gamma; \cdot) \vdash \langle M \rangle \equiv \langle N \rangle \in \langle A \rangle}\ \mathsf{eqinjV} \qquad \frac{\Omega \vdash e_1 \to^* e_1' \quad \Omega \vdash e_2 \to^* e_2' \quad \Omega \vdash e_1' \equiv e_2' \in \langle A \rangle}{\Omega \vdash e_1 \equiv e_2 \in \langle A \rangle}\ \mathsf{eqinjR}$$

$$\frac{\Omega \vdash e_1 \equiv e_2 \in \tau}{\Omega, (\Gamma; \cdot) \vdash \mathrm{pop}\ e_1 \equiv \mathrm{pop}\ e_2 \in \Box\tau}\ \mathsf{eqpopV} \qquad \frac{\Omega \vdash e_1 \to^* e_1' \quad \Omega \vdash e_2 \to^* e_2' \quad \Omega \vdash e_1' \equiv e_2' \in \Box\tau}{\Omega \vdash e_1 \equiv e_2 \in \Box\tau}\ \mathsf{eqpopR}$$

$$\frac{}{\Omega \vdash e \equiv e \in \tau_1 \Rightarrow \tau_2}\ \mathsf{eqfun}$$

**Fig. 4.** Small-step semantics (Equality)

parameters introduced by $\nu$, and thus rednabla expresses a non-deterministic choice of parameters. Finally, redfix implements the unrolling of the recursion operator.

The bottom two rules are necessary to give us a congruence closure for reductions on $\nabla$-expressions. Because the $\nabla$-calculus is call-by-name, we do not evaluate $e_2$ in the rule redfun. Finally, rednew reduces under the $\nu$ after appropriately copying and extending the current scope.

Thus, equivalence on functions is decided only by syntactic equality, as shown by rule eqfun in Figure 4. For all other types, we give two rules: the rule ending in V refers to the case where the left and right hand side are already values, while the rule ending in R is used when further reduction steps are required on either side.

$$\frac{}{\Gamma \vdash \cdot : \cdot}\ \mathsf{tpEObjS} \qquad \frac{\Gamma \vdash M : A \quad \Gamma \vdash \gamma : \Gamma'}{\Gamma \vdash (\gamma, M/x) : (\Gamma', x : A)}\ \mathsf{tpIObS}$$

$$\frac{}{\Omega \vdash \cdot : \cdot}\ \mathsf{tpEMetaS} \qquad \frac{\Omega \vdash f \in \tau \quad \Omega \vdash \varphi : \Phi}{\Omega \vdash (\varphi, f/u) : (\Phi, u \in \tau)}\ \mathsf{tpIMetaS}$$

$$\frac{}{\Omega \vdash \cdot : \cdot}\ \mathsf{tpEStackS} \qquad \frac{\Gamma \vdash \gamma : \Gamma' \quad \Omega, (\Gamma; \Phi) \vdash \varphi : \Phi' \quad \Omega \vdash \omega : \Omega'}{\Omega, (\Gamma; \Phi) \vdash \omega, (\gamma; \varphi) : \Omega', (\Gamma', \Phi')}\ \mathsf{tpIStackS}$$

**Fig. 5.** The static semantics of substitutions

## 5 Meta Theory

We study the meta-theory of the $\nabla$-calculus culminating in the type-preservation theorem, which entails that parameters cannot escape their scope.

Substituting for $\epsilon$ and $\nabla$-bound variables is essential for defining the operational meaning of our expressions. In this section we elaborate on object-level and meta-level substitutions, as well as substitution stacks, which are defined on scope stacks. As is standard, we make our substitutions capture avoiding by tacitly renaming variable names.

**Definition 3 (Substitutions).**

$$\begin{array}{ll} \textit{Object-Level Substitutions:} & \gamma ::= \cdot \mid \gamma, M/x \\ \textit{Meta-Level Substitutions:} & \varphi ::= \cdot \mid \varphi, f/u \\ \textit{Substitution Stacks:} & \omega ::= \cdot \mid \omega, (\gamma; \varphi) \end{array}$$

We define the meaning of the three typing judgments for substitutions $\Gamma \vdash \gamma : \Gamma'$, $\Omega \vdash \phi : \Phi$, and $\Omega \vdash \omega \in \Omega'$ in Figure 5. The domains of the substitutions are $\Gamma'$, $\Phi$, and $\Omega'$, respectively, and the codomains of the substitutions are $\Gamma$, $\Omega$, and $\Omega$, respectively. The definition of substitution application is given in Figure 6.

We write $\Gamma < \Gamma'$ if $\Gamma'$ strictly extends $\Gamma$, and $\Phi \leq \Phi'$, if $\Phi = \Phi'$ or $\Phi < \Phi'$.

**Definition 4 (Well-Formed Context Stacks).** *We say that a context stack $\Omega$ is well-formed if the proposition $\vdash \Omega\ \mathsf{ok}$ can be proved using the following judgments:*

$$\frac{}{\vdash \cdot\ \mathsf{ok}}\ \mathsf{okempty} \qquad \frac{}{\vdash \cdot, (\Gamma; \Phi)\ \mathsf{ok}}\ \mathsf{okinit} \qquad \frac{\vdash \Omega, (\Gamma; \Phi)\ \mathsf{ok} \quad \Gamma < \Gamma' \quad \Phi \leq \Phi'}{\vdash \Omega, (\Gamma; \Phi), (\Gamma'; \Phi')\ \mathsf{ok}}\ \mathsf{oknew}$$

The following substitution lemma is the key lemma for proving type preservation.

**Lemma 1 (Substitution).** *If $\Omega \vdash e \in \tau$, $\Omega\ \mathsf{ok}$, $\Omega'\ \mathsf{ok}$ and $\Omega' \vdash \omega \in \Omega$ then $\Omega' \vdash [\omega]e \in \tau$.*

$$
\begin{array}{rl}
[\gamma, M/x]x = M & [\omega, (\gamma;\varphi)]u = [\varphi]u \\
[\gamma, M/x]y = [\gamma]y & [\omega, (\gamma;\varphi)]\langle N\rangle = \langle[\gamma]N\rangle \\
[\gamma]c = c & [\omega, (\gamma;\varphi)](\text{pop } e) = \text{pop } [\omega]e \\
[\gamma](N_1\ N_2) = ([\gamma]N_1)\ ([\gamma]N_2) & [\omega](e_1 \mapsto_\tau e_2) = ([\omega]e_1) \mapsto_\tau ([\omega]e_2) \\
[\gamma](\lambda x : A.\, N) = \lambda x : A.\, [\gamma, x/x]N & [\omega](e_1 \cdot e_2) = ([\omega]e_1) \cdot ([\omega]e_2) \\
 & [\omega](e_1 \mid e_2) = ([\omega]e_1 \mid [\omega]e_2) \\
 & [\omega, (\gamma;\varphi)](\text{fix } u \in \tau.\, e) = \text{fix } u \in \tau.\, [\omega, (\gamma;\varphi, u/u)]e \\
 & [\omega, (\gamma;\varphi)](\epsilon x : A.\, e) = \epsilon x : A.\, [\omega, (\gamma, x/x;\varphi)]e \\
[\varphi, e/u]u = e & [\omega, (\gamma;\varphi)](\epsilon u \in \tau.\, e) = \epsilon u \in \tau.\, [\omega, (\gamma;\varphi, u/u)]e \\
[\varphi, e/u]v = [\varphi]v & [\omega, (\gamma;\varphi)](\nabla x : A.\, e) = \nabla x : A.\, [\omega, (\gamma, x/x;\varphi)]e \\
 & [\omega, (\gamma;\varphi)](\nu x : A.\, e) = \nu x : A.\, [\omega, (\gamma;\varphi), (\gamma, x/x;\varphi)]e
\end{array}
$$

**Fig. 6.** Substitution Application

*Proof.* By induction on the structure of $\Omega \vdash e \in \tau$. See the Technical Report [SPS04] for details.

We are now ready to prove the type preservation theorem.

**Theorem 2 (Type Preservation).** *If* $\vdash \Omega$ ok *and* $\Omega \vdash e \in \tau$ *and* $\Omega \vdash e \to e'$ *then* $\Omega \vdash e' \in \tau$.

*Proof.* By induction on the structure of $\Omega \vdash e \to e'$. See the Technical Report [SPS04] for details.

As a corollary we obtain the property that parameters cannot escape their scope.

**Corollary 1 (Scope Preservation).** *If* $\vdash \Omega, (\Gamma;\cdot)$ ok *and* $\Omega, (\Gamma;\cdot) \vdash e \in \Box\tau$ *and* $\Omega, (\Gamma;\cdot) \vdash e \to^* v$ *and* $v$ *is a value then* $v = \text{pop } e'$ *and* $\Omega \vdash e' \in \tau$.

In future work, we will investigate further meta-theoretical properties of the $\nabla$-calculus, such as progress and termination. Neither of these two properties is satisfied without additional side conditions on the typing rules.

## 6 Conclusion

In this paper we have presented the $\nabla$-calculus. We allow for evaluation under $\lambda$-binders, pattern matching against parameters, and programming with higher-order encodings. The $\nabla$-calculus has been implemented as a stand-alone programming language, called Elphin [PS04]. The $\nabla$-calculus solves many problems associated with programming with higher-order abstract syntax. We allow for, and can usefully manipulate, datatype declarations whose constructor types make reference to themselves in negative positions while maintaining a closed description of the functions. Many examples, such as parallel reduction and an evaluator for a simple language with first-class continuations can be found in the Technical Report [SPS04].

The $\nabla$-calculus is the result of many years of design, originally inspired by an extension to ML proposed by Dale Miller [Mil90]. Other influencing works include pattern-matching calculi as employed in ALF [CNSvS94] or proposed by Jouannaud and Okada [JO91], the type theory $\mathcal{T}_\omega^+$ [Sch01], and Hofmann's work on higher-order abstract syntax [Hof99]. A direct predecessor to the $\nabla$-calculus is the modal $\lambda$-calculus with iterators [SDP01]. We conjecture that any function written in the modal $\lambda$-calculus with iterators can also be expressed in the $\nabla$-calculus.

Closely related to our work are programming languages with freshness [PG00,GP99], which provide a built-in $\alpha$-equivalence relation for first-order encodings but provide neither $\beta\eta$ nor any support for higher-order encodings. Also closely related to the $\nabla$-calculus are meta-programming languages, such as MetaML [TS00,Nan02], which provide hierarchies of computation levels, but do not single out a particular level for representation. Many other attempts have been made to combine higher-order encodings and functional programming, in particular Honsell, Miculan, and Scagnetto's embedding of the $\pi$-calculus in Coq[HMS01], and Momgliano, Amber, and Crole's Hybrid system [MAC03].

In future work, we plan to extend the $\nabla$-calculus to a dependently-typed logical framework, add polymorphism to the computation-level, and study termination and progression.

*Acknowledgments.* We would like to thank Henrik Nilsson, Simon Peyton-Jones, and Valery Trifonov for comments on earlier drafts of this paper.

## References

[CNSvS94] Thierry Coquand, Bengt Nordström, Jan M. Smith, and Björn von Sydow. Type theory and programming. *Bulletin of the European Association for Theoretical Computer Science*, 52:203–228, February 1994.

[Coq91] Thierry Coquand. An algorithm for testing conversion in type theory. In Gérard Huet and Gordon Plotkin, editors, *Logical Frameworks*, pages 255–279. Cambridge University Press, 1991.

[DPS97] Joëlle Despeyroux, Frank Pfenning, and Carsten Schürmann. Primitive recursion for higher-order abstract syntax. In R. Hindley, editor, *Proceedings of the Third International Conference on Typed Lambda Calculus and Applications (TLCA'97)*, pages 147–163, Nancy, France, April 1997. Springer-Verlag LNCS. An extended version is available as Technical Report CMU-CS-96-172, Carnegie Mellon University.

[GP99] Murdoch Gabbay and Andrew Pitts. A new approach to abstract syntax involving binders. In G. Longo, editor, *Proceedings of the 14th Annual Symposium on Logic in Computer Science (LICS'99)*, pages 214–224, Trento, Italy, July 1999. IEEE Computer Society Press.

[HHP93] Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the Association for Computing Machinery*, 40(1):143–184, January 1993.

[HMS01] Furio Honsell, Marino Miculan, and Ivan Scagnetto. pi-calculus in (Co)inductive-type theory. *Theoretical Computer Science*, 253(2):239–285, 2001.

[Hof99] Martin Hofmann. Semantical analysis for higher-order abstract syntax. In G. Longo, editor, *Proceedings of the 14th Annual Symposium on Logic in Computer Science (LICS'99)*, pages 204–213, Trento, Italy, July 1999. IEEE Computer Society Press.

[JO91] Jean-Pierre Jouannaud and Mitsuhiro Okada. A computation model for executable higher-order algebraic specification languages. In Gilles Kahn, editor, *Proceedings of the 6th Annual Symposium on Logic in Computer Science*, pages 350–361, Amsterdam, The Netherlands, July 1991. IEEE Computer Society Press.

[MAC03] Alberto Momgliano, Simon Ambler, and Roy Crole. A definitional approach to primitive recursion over higher order abstract syntax. In Alberto Momgliano and Marino Miculan, editors, *Proceedings of the Merlin Workshop*, Uppsala, Sweden, June 2003. ACM Press.

[Mil90] Dale Miller. An extension to ML to handle bound variables in data structures: Preliminary report. In *Proceedings of the Logical Frameworks BRA Workshop*, Nice, France, May 1990.

[Nan02] Aleksander Nanevski. Meta-programming with names and necessity. In Cindy Norris and Jr. James B. Fenwick, editors, *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP-02)*, Pittsburgh, PA, October 2002. ACM Press.

[Pfe92] Frank Pfenning. Computation and deduction. Unpublished lecture notes, 277 pp. Revised May 1994, April 1996, May 1992.

[PG00] A. M. Pitts and M. J. Gabbay. A metalanguage for programming with bound names modulo renaming. In R. Backhouse and J. N. Oliveira, editors, *Mathematics of Program Construction, MPC2000, Proceedings, Ponte de Lima, Portugal, July 2000*, volume 1837 of *Lecture Notes in Computer Science*, pages 230–255. Springer-Verlag, Heidelberg, 2000.

[PM93] Christine Paulin-Mohring. Inductive definitions in the system Coq: Rules and properties. In M. Bezem and J.F. Groote, editors, *Proceedings of the International Conference on Typed Lambda Calculi and Applications*, pages 328–345, Utrecht, The Netherlands, March 1993. Springer-Verlag LNCS 664.

[PS04] Adam Poswolsky and Carsten Schürmann. Elphin: Functional programming with higher-order encodings. Technical report, Yale University, 2004. to appear.

[Sch01] Carsten Schürmann. Recursion for higher-order encodings. In Laurent Fribourg, editor, *Proceedings of the Conference on Computer Science Logic (CSL 2001)*, pages 585–599, Paris, France, August 2001. Springer Verlag LNCS 2142.

[SDP01] Carsten Schürmann, Joëlle Despeyroux, and Frank Pfenning. Primitive recursion for higher-order abstract syntax. *Theoretical Computer Science*, (266):1–57, 2001. Journal version of [DPS97].

[SPS04] Carsten Schürmann, Adam Poswolsky, and Jeffrey Sarnat. The $\nabla$-calculus. functional programming with logical frameworks. Technical Report YALEU/DCS/TR-1272, Yale University, October 2004.

[TS00] Walid Taha and Tim Sheard. MetaML: Multi-stage programming with explicit annotations. *Theoretical Computer Science*, 248(1-2), 2000.