



University of Pennsylvania
ScholarlyCommons

Technical Reports (CIS)

Department of Computer & Information Science

6-2010

LNgen: Tool Support for Locally Nameless Representations

Brian Aydemir
University of Pennsylvania

Stephanie Weirich
University of Pennsylvania, sweirich@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Recommended Citation

Brian Aydemir and Stephanie Weirich, "LNgen: Tool Support for Locally Nameless Representations", . June 2010.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-10-24.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/933
For more information, please contact repository@pobox.upenn.edu.

LNgen: Tool Support for Locally Nameless Representations

Abstract

Given the complexity of the metatheoretic reasoning about current programming languages and their type systems, techniques for mechanical formalization and checking of such metatheory have received much recent attention. In previous work, we advocated a combination of locally nameless representation and cofinite quantification as a lightweight style for carrying out such formalizations in the Coq proof assistant. As part of the presentation of that methodology, we described a number of operations associated with variable binding and listed a number of properties, called “infrastructure lemmas”, about those operations that needed to be shown. The proofs of these infrastructure lemmas are straightforward but tedious.

In this work, we present LNgen, a prototype tool for automatically generating statements and proofs of infrastructure lemmas from Ott language specifications. Furthermore, the tool also generates a recursion scheme for defining functions over syntax, which was not available in our previous work. LNgen works in concert with Ott to effectively alleviate much of the tedium of working with locally nameless syntax. For the case of untyped lambda terms, we show that the combined output from the two tools is sound and complete, with LNgen automatically proving many of the key lemmas. We prove the soundness of our representation with respect to a fully concrete representation, and we argue that the representation is complete—that we generate the right set of lemmas—with respect to Gordon and Melham’s “Five Axioms of Alpha-Conversion.”

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-10-24.

LNgen: Tool Support for Locally Nameless Representations

MS-CIS-10-24

Brian Aydemir
University of Pennsylvania
baydemir@cis.upenn.edu

Stephanie Weirich
University of Pennsylvania
sweirich@cis.upenn.edu

June 2010

Abstract

Given the complexity of the metatheoretic reasoning about current programming languages and their type systems, techniques for mechanical formalization and checking of such metatheory have received much recent attention. In previous work, we advocated a combination of locally nameless representation and cofinite quantification as a lightweight style for carrying out such formalizations in the Coq proof assistant. As part of the presentation of that methodology, we described a number of operations associated with variable binding and listed a number of properties, called “infrastructure lemmas”, about those operations that needed to be shown. The proofs of these infrastructure lemmas are straightforward but tedious.

In this work, we present LNgen, a prototype tool for automatically generating statements and proofs of infrastructure lemmas from Ott language specifications. Furthermore, the tool also generates a recursion scheme for defining functions over syntax, which was not available in our previous work. LNgen works in concert with Ott to effectively alleviate much of the tedium of working with locally nameless syntax. For the case of untyped lambda terms, we show that the combined output from the two tools is sound and complete, with LNgen automatically proving many of the key lemmas. We prove the soundness of our representation with respect to a fully concrete representation, and we argue that the representation is complete—that we generate the right set of lemmas—with respect to Gordon and Melham’s “Five Axioms of Alpha-Conversion.”

1 Introduction

Mechanical formalizations of programming languages have received much recent attention. One question that is foremost in any mechanization is the treatment of binding. Many tools exist to aid in this practice—Abella [1], Hybrid [2], Lambda Tamer [3], Nominal Isabelle [4], Twelf [5]—as well as many representation techniques—de Bruijn indices [6], higher-order abstract syntax (HOAS) [7], locally named [8], locally nameless [9], weak HOAS [10], etc.

As a programming language designer, how should we compare these methodologies? What tools should we use? The POPLMARK challenge [11] laid out a number of criteria, which we have come to interpret with respect to existing technologies, for evaluating potential answers:

1. *Transparency.* Reasoning should be similar to that done with pencil and paper. For example, de Bruijn indices are not transparent. Metatheory involving them often includes many lemmas about shifting—lemmas that have no correspondence to pencil and paper proofs.
2. *Logical expressivity.* There should be minimal restriction on the logic that we use for formal developments. For example, the models of Nominal Logic [12] require that all definable relations be equivariant. To allow similar reasoning in higher-order logic, where this is not the case, Nominal Isabelle must require equivariance proofs (many of which can be provided automatically).

3. *Traction.* The strategy should draw on the strengths of the proof assistant. For example, in previous work [13], we explored nominal reasoning in Coq by defining an interface which specified the constructors of a nominal datatype, as well as an induction principal and recursion scheme for that datatype. We chose not to pursue that line of work because the interface, while usable, prevented users from taking advantage of Coq’s built-in features. Utilizing distinctness and injectivity of datatype constructors, reasoning by induction, and defining functions by recursion all required the explicit use of special theorems and combinators. Furthermore, functions defined by the recursion combinator would not reduce by Coq’s definitional equality—we had to use explicit rewriting.

From these criteria, we draw the following conclusions: We want a representation that involves reasoning about variable names, not indices, because that is the most transparent. We want to use this representation in a general purpose logic, such as higher-order logic or the Calculus of Inductive Constructions (CIC), but we want to automate as much of the tedious machinery as possible. And we want our representation of syntax to use what proof assistants are good at: specifying inductive datatypes and generating their associated induction principles and recursion schemes.

In previous work [14], we proposed a completely manual scheme for reasoning about binding structure based on locally nameless representations and defining inference rules with cofinite quantification. We described a number of operations associated with variable binding (free variable calculation, index substitution, free variable substitution, and free variable closing) and listed a number of properties, called “infrastructure lemmas”, about those operations that needed to be shown. This strategy is lightweight in that the definitions of the operations are simple structural recursions, so proofs of their properties are straightforward. We have successfully used this strategy in our own developments and know of its use by others—for example, by Jia et al. [15], Pratikakis et al. [16], Rossberg et al. [17], and many more.

However, our previous work did not fully explain its own success. Why were the “infrastructure lemmas” the right set of lemmas to show? Would future formalizations require still more lemmas? Furthermore, if the proofs of the infrastructure lemmas are so straightforward and mechanical, should it not be possible to automatically generate those lemma statements and their proofs?

In this paper, we describe a prototype tool, LNgen, that we have developed for exactly this last purpose. LNgen uses the same input language as Ott [18], a tool for translating language specifications written in an intuitive syntax into output for L^AT_EX and proof assistants. While Ott generates locally nameless definitions—datatypes for syntax and relations, functions to calculate free variables and substitutions—from the specification, LNgen provides recursion schemes for defining functions over syntax and a large collection of infrastructure lemmas. LNgen automates much of the tedium associated with the locally nameless style, even in our streamlined style, by allowing users to focus on the more interesting aspects of their developments instead of on infrastructure lemmas. In Sec. 2, we describe in additional detail the input to and output from LNgen, highlighting the important properties that are automatically proved.

Following the overview of LNgen, we discuss soundness (Sec. 3) and completeness (Sec. 4) in the particular case of the untyped lambda calculus. For soundness, we prove that the locally nameless definition generated by Ott is adequate with respect to fully concrete terms identified up to alpha equivalence. The lemmas proved by LNgen provide many of the key lemmas required in this proof. For completeness, we prove that even though we use a locally nameless representation, the lemmas generated by LNgen are enough to shield users from the de Bruijn indices used to represent bound variables. Specifically, we give a model of Gordon and Melham’s “Five Axioms of Alpha-Conversion” [19]. Although we think that the output of Ott and LNgen is more convenient to work with than the five axioms, we can implement these five axioms in an extremely straightforward manner, by using the lemmas proved by LNgen and without reasoning about de Bruijn indices or by induction on syntax.

We and others have experience with using LNgen in significant developments. Section 5 gives an overview of the case studies. Our experiences suggest that this tool has the advantages of code generators without the drawbacks of generating executable code. In particular, the output of LNgen is straightforward for programmers to effectively understand (definitions and lemma statements must be comprehended, but proofs do not) and robust to change (lemma statements do not change significantly as the language is modified).

We conclude the paper with related work (Sec. 6), and our conclusions and future work (Sec. 7).

<pre> metavar expvar, x, y, z ::= {{ repr-locally-nameless }} grammar exp, e, f, g :: '' ::= x :: :: var e1 e2 :: :: app \ x . e :: :: abs (+ bind x in e +) substitutions single e x :: subst freevars e x :: fv </pre>	<pre> Definition expvar := var. Inductive exp : Set := var_b : nat -> exp var_f : expvar -> exp app : exp -> exp -> exp. abs : exp -> exp </pre>
--	--

Figure 1: Input file (left) and output Coq datatype (right) for lambda terms

2 The LNgen Tool

LNgen is a prototype tool for generating locally nameless definitions and infrastructure for the Coq proof assistant. While LNgen is still under active development, the current version is available and has been used for significant developments.¹ LNgen relies on Ott [18] to generate the core locally nameless definitions for a language. It then generates additional definitions and lemmas that are often needed in developments—the main benefit that it provides to users over using Ott alone.

The input language for LNgen is a proper subset of the Ott specification language. Figure 1 shows an example input file for untyped lambda terms. The syntax is intended to mimic what one might write informally. Ott is specifically designed for specifying programming languages in a manner that is both convenient for people and machines, e.g., proof assistants. Thus, Ott is a natural starting point for the input language to LNgen. We can take advantage of the work that has gone into the design of Ott, not require users to learn a new specification language, and allow our tool to work in parallel with Ott, relying on Ott for the generation of some of the Coq definitions as well as \LaTeX output.

Below, we use the example to give a brief overview of the subset of Ott that LNgen supports; a detailed description of the Ott language can be found elsewhere [20]. The first part of an input file for LNgen consists of a list of **metavar** declarations. Each declaration defines a new type for object language variables—LNgen and Ott define binding and substitution for these variables. In Fig. 1, the text **repr-locally-nameless** indicates that binding should be represented using a locally nameless encoding. (Ott can also output definitions using a concrete representation of binding.) The second part, the **grammar**, consists of a list of context-free grammar definitions for nonterminals. Each declaration defines a new, inductively defined type for object-language abstract syntax trees. Binding specifications may be attached to each constructor. For example in the **abs** constructor, the metavariable **x** is a binding occurrence in the nonterminal **e**. The third part follows the **substitutions** keyword and indicates that functions for substituting for free variables should be generated. The final part follows the **freevars** keyword and indicates that functions for calculating free variables should be generated. Anything else in the file is ignored by LNgen but may be processed by Ott, e.g., specifications of inductively defined relations.

2.1 Generated Definitions

Figure 1 also shows the output representation that Ott produces for the untyped lambda terms. Metavariables are implemented by the type **var**, which is provided by our metatheory library.² The constructor names

¹LNgen is available from <http://www.cis.upenn.edu/~baydemir/>.

²The metatheory library is included with LNgen and also available from <http://www.plclub.org/metelib/>.

for the syntactic forms are determined by the input file, except the constructors for free and bound variables, where `_f` and `_b` are appended to the specified name. The variable in the `abs` constructor disappears because the binding specification indicates that this is a binding constructor.

Figure 2 lists the basic operations and predicates generated from the input in Fig. 1. For accessibility and brevity, we use mathematical notation instead of listing the Coq output directly. Ott generated the definitions of `fv` and `subst`; LNgen generated everything else. In general, the output follows our previously described style for working with locally nameless representations [14]. The operations include calculating the free variables of an expression (`fv`), substituting for an index (`openi`), replacing a free variable with an index (`closei`), and substituting for a free variable (`subst`). Note that `close` allows us to construct a concrete expression without explicitly referring to indices. Using `lam x` as an abbreviation for `abs ∘ close x`, we can transparently write $\lambda x.\lambda y.\lambda z.z(xy)$ as

$$\text{lam } x \text{ (lam } y \text{ (lam } z \text{ (app (var.f } z) \text{ (app (var.f } x) \text{ (var.f } y))))).$$

Note also that the versions of `openi` and `closei` presented here are derived from those of Pollack [9] and are slightly more general than that of our previous work—they may initially be called with an index other than zero. Previously, we promoted the absolute simplest definitions to make working by hand easy. Here, we have tool support, so it makes little difference if these definitions are more complicated. If anything, they are actually *easier* for LNgen to work with because they require tracking fewer invariants.

The final definitions in Fig. 2 give the constructors for the inductively defined `lc` and `lc.set` predicates, which hold for *locally closed* lambda terms—those with no unresolved de Bruijn indices. Only expressions that satisfy these predicates correspond to lambda calculus terms. The only difference between the two predicates is that the former is in `Prop` and the latter is in `Set`; their definitions are otherwise identical. Because of Coq’s distinction between `Prop` and `Set`, their uses are not. An object of type `lc e` is treated as a proof and may be analyzed only to produce another proof; an object of type `lc.set e` may be analyzed freely. The inductive definition of `lc` provides an induction principle for reasoning about expressions, while the inductive definition of `lc.set` provides a recursion scheme for defining functions over expressions. The induction principle and recursion scheme are both shown in Fig. 3.

Our treatment of local closure departs from our previous work in that we previously did not provide `lc.set` and the recursion scheme that comes with it. We can use the scheme, for example, to define a function to perform parallel β -reduction on lambda terms:

$$\begin{aligned} \text{beta} &= \text{lc.set.rec}(\lambda _ . \text{exp}) \text{fvar fapp fabs} \text{ where} \\ \text{fvar } x &= \text{var.f } x \\ \text{fapp } _ _ _ (\text{abs } e'_1) _ e'_2 &= \text{open } e'_2 \ e'_1 \\ \text{fapp } _ _ _ e'_1 _ e'_2 &= \text{app } e'_1 \ e'_2 \\ \text{fabs } e_1 _ f' &= \text{abs}(\text{close } x \ (f' \ x)) \text{ for some } x \notin \text{fv } e_1 \end{aligned}$$

(In Coq, one would use this recursion scheme via `Fixpoint`, writing the function more naturally using explicit pattern matching on the local closure proof, and explicit recursive calls.) In the variable case, `beta` simply returns that variable. In the application case, the result of reducing the first component is examined: if it is an abstraction (`abs e'_1`), `beta` substitutes the reduced second component `e'_2` for the first index in the body of the abstraction. Otherwise, reduction continues into both components of the application. In the abstraction case, `beta` reduces the body of the abstraction by picking a fresh variable to give to `f'`. This argument to `fabs` is a function that, when given a name for the variable bound at this location, computes the result of `beta` for the body of the abstraction using that name. After this recursive call, the branch removes that fresh variable from the result with `close` and creates a new abstraction.

In another departure from our previous work, neither `lc` nor `lc.set` uses cofinite quantification. Instead, both use “universal” quantification in the `abs` case, by requiring that the premise hold for all names. This choice results in the strongest possible induction principle and recursion scheme. For `lc`, LNgen generates as a lemma an “existential” form of the `lc.abs` constructor (lemma `lc-abs-exists` in Fig. 4) that requires showing the premise for only one name. This lemma provides the easiest to use introduction principle for proving `lc(abs e)`. This style of using a “universal” and an “existential” rule is based on the style of McKinna and Pollack [8]. While cofinite quantification is a good compromise between these two extremes when doing everything by hand, with tool support, it makes sense to provide these stronger principles. Using universal quantification also allows us to prove the uniqueness of `lc` proofs (lemma `lc-unique` in Fig. 4).

fv	$:$	$\text{exp} \rightarrow \text{expvarset}$	
$\text{fv}(\text{var_f } x)$	$=$	$\{x\}$	
$\text{fv}(\text{var_b } i)$	$=$	\emptyset	
$\text{fv}(\text{abs } e_1)$	$=$	$\text{fv } e_1$	
$\text{fv}(\text{app } e_1 e_2)$	$=$	$\text{fv } e_1 \cup \text{fv } e_2$	
open_i	$:$	$\text{nat} \rightarrow \text{exp} \rightarrow \text{exp} \rightarrow \text{exp}$	
$\text{open}_i e (\text{var_b } i_1)$	$=$	$(\text{var_b } i_1)$	when $i_1 < i$
$\text{open}_i e (\text{var_b } i_1)$	$=$	e	when $i_1 = i$
$\text{open}_i e (\text{var_b } i_1)$	$=$	$(\text{var_b } (i_1 - 1))$	when $i_1 > i$
$\text{open}_i e (\text{var_f } x)$	$=$	$\text{var_f } x$	
$\text{open}_i e (\text{abs } e_1)$	$=$	$\text{abs}(\text{open}_{(i+1)} e e_1)$	
$\text{open}_i e (\text{app } e_1 e_2)$	$=$	$\text{app}(\text{open}_i e e_1) (\text{open}_i e e_2)$	
$\text{open } e_1 e_2$	$=$	$\text{open}_0 e_1 e_2$	
close_i	$:$	$\text{nat} \rightarrow \text{expvar} \rightarrow \text{exp} \rightarrow \text{exp}$	
$\text{close}_i x (\text{var_b } i_1)$	$=$	$\text{var_b } i_1$	when $i_1 < i$
$\text{close}_i x (\text{var_b } i_1)$	$=$	$\text{var_b } (1 + i_1)$	when $i_1 \geq i$
$\text{close}_i x (\text{var_f } y)$	$=$	$\text{var_b } i$	when $x = y$
$\text{close}_i x (\text{var_f } y)$	$=$	$\text{var_f } y$	when $x \neq y$
$\text{close}_i x (\text{abs } e_1)$	$=$	$\text{abs}(\text{close}_{(1+i)} x e_1)$	
$\text{close}_i x (\text{app } e_1 e_2)$	$=$	$\text{app}(\text{close}_i x e_1) (\text{close}_i x e_2)$	
$\text{close } x e$	$=$	$\text{close}_0 x e$	
subst	$:$	$\text{exp} \rightarrow \text{expvar} \rightarrow \text{exp} \rightarrow \text{exp}$	
$\text{subst } e x (\text{var_b } i_1)$	$=$	$\text{var_b } i_1$	
$\text{subst } e x (\text{var_f } y)$	$=$	e	when $x = y$
$\text{subst } e x (\text{var_f } y)$	$=$	$\text{var_f } y$	when $x \neq y$
$\text{subst } e x (\text{abs } e_1)$	$=$	$\text{abs}(\text{subst } e x e_1)$	
$\text{subst } e x (\text{app } e_1 e_2)$	$=$	$\text{app}(\text{subst } e x e_1) (\text{subst } e x e_2)$	
lc	$:$	$\text{exp} \rightarrow \text{Prop}$	
lc_var	$:$	$\forall x, \text{lc}(\text{var_f } x)$	
lc_app	$:$	$\forall e_1 e_2, \text{lc } e_1 \rightarrow \text{lc } e_2 \rightarrow \text{lc}(\text{app } e_1 e_2)$	
lc_abs	$:$	$\forall e_1, (\forall x, \text{lc}(\text{open}(\text{var_f } x) e_1)) \rightarrow \text{lc}(\text{abs } e_1)$	
lc_set	$:$	$\text{exp} \rightarrow \text{Set}$	
lc_set_var	$:$	$\forall x, \text{lc_set}(\text{var_f } x)$	
lc_set_app	$:$	$\forall e_1 e_2, \text{lc_set } e_1 \rightarrow \text{lc_set } e_2 \rightarrow \text{lc_set}(\text{app } e_1 e_2)$	
lc_set_abs	$:$	$\forall e_1, (\forall x, \text{lc_set}(\text{open}(\text{var_f } x) e_1)) \rightarrow \text{lc_set}(\text{abs } e_1)$	

Convention: The first two arguments to lc_app and lc_set_app are implicit, as are the first arguments to lc_abs and lc_set_abs .

Figure 2: Definitions generated by Ott and LNgen

Induction principle (`lc.ind`)

$$\begin{aligned}
& \forall (P : \text{exp} \rightarrow \text{Prop}), \\
& (\forall x, P(\text{var_f } x)) \rightarrow \\
& (\forall e_1 e_2, \text{lc } e_1 \rightarrow P e_1 \rightarrow \text{lc } e_2 \rightarrow P e_2 \rightarrow P(\text{app } e_1 e_2)) \rightarrow \\
& (\forall e_1, \\
& \quad (\forall x, \text{lc}(\text{open}(\text{var_f } x) e_1)) \rightarrow (\forall x, P(\text{open}(\text{var_f } x) e_1)) \rightarrow P(\text{abs } e_1)) \rightarrow \\
& \forall e, \text{lc } e \rightarrow P e
\end{aligned}$$

Recursion scheme (`lc.set_rec`)

`lc.set_rec` has the same type as `lc.ind`, except with `Set` instead of `Prop`, and `lc.set` instead of `lc`. It behaves as follows: If $f = \text{lc.set_rec } P \text{ fvar } \text{fapp } \text{fabs}$, then

$$\begin{aligned}
f \quad (\text{var_f } x) \quad (\text{lc_var } x) &= \text{fvar } x \\
f \quad (\text{app } e_1 e_2) \quad (\text{lc_app } lcp_1 lcp_2) &= \text{fapp } e_1 e_2 lcp_1 (f e_1 lcp_1) lcp_2 (f e_2 lcp_2) \\
f \quad (\text{abs } e_1) \quad (\text{lc_abs } lcp) &= \text{fabs } e_1 lcp (\lambda x. f(\text{open}(\text{var_f } x) e_1) (lcp x)).
\end{aligned}$$

Figure 3: Induction principal and recursion scheme

2.2 Generated Lemmas

The main benefit to using LNgen is that it automatically generates a collection of lemmas (with their proofs) about expressions that are useful in metatheoretic reasoning. We highlight the most important of these in Fig. 4. The collection shown includes all of the lemmas that we discussed in our previous work [14]. For convenience, LNgen also generates several variants of the lemmas shown and others besides. Our goal in picking the set of lemmas to generate was not to determine some minimal “complete” set for working with metatheory but to generate a set that, from our experience, we know to be useful in formalizations.

Many of the lemmas in Fig. 4 describe the interaction between the various operations. For example, the first group of lemmas (1–6) describe what happens when `fv` is applied to expressions built from `open`, `close` and `subst`.

The next eight lemmas (7–14) are primarily about `subst`. Lemma *subst-spec* decomposes substitution into `open` composed with `close`, which was Gordon’s definition of substitution [21]. We prefer our version because it commutes directly with constructors. (A definition in terms of `open` and `close` would need to use `openi` and `closei` once it went under a binder.) Lemma *subst-abs* lets us reason about how substitution interacts with abstractions, while making sure that we call `subst` only on locally closed terms. (The definition of `subst` just pushes through an abstraction, calling itself recursively on the body, which may have an unresolved index.)

The remaining lemmas (15–23) describe properties of `open`, `close`, and `lc`. Lemma *lc-abs-exists* asserts the existence of an operation that constructs a local closure proof for an abstraction from a proof about a single variable. (Recall that the definition of local closure required that the body be closed for any name for the free variable; this one requires only a single name.) Lemma *lc-subst* asserts the existence of an operation that shows that local closure proofs are preserved by substitution. Lemma *lc-unique* shows that all local closure proofs about the same expression are equivalent.³ Finally, lemmas *lc-of-lc-set* and *lc-set-of-lc* show the equivalence between `lc` and `lc.set`.

2.3 Generated Proofs

LNgen is able to automatically generate the proofs of each of the lemmas in Fig. 4 because, in general, they are “boring” infrastructure lemmas whose proofs are straightforward inductions. At any given point in a proof, there is little choice about what step to take next. Thus, most of the proof scripts start by applying an induction tactic and then use a “power tactic” to apply a default set of simplifications to the resulting

³The proof of this lemma requires extensional equality on functions, which may safely be asserted in Coq as an axiom.

1. *fv-open-upper*:
 $\text{fv}(\text{open } e_1 e_2) \subseteq \text{fv } e_1 \cup \text{fv } e_2$
2. *fv-open-lower*:
 $\text{fv } e_2 \subseteq \text{fv}(\text{open } e_1 e_2)$
3. *fv-close*:
 $\text{fv}(\text{close } x e) = \text{fv } e \setminus \{x\}$
4. *fv-subst-upper*:
 $\text{fv}(\text{subst } e_1 x e_2) \subseteq \text{fv } e_1 \cup (\text{fv } e_2 \setminus \{x\})$
5. *fv-subst-lower*:
 $(\text{fv } e_2 \setminus \{x\}) \subseteq \text{fv}(\text{subst } e_1 x e_2)$
6. *fv-subst-fresh*:
 $\text{fv}(\text{subst } e_1 x e_2) = \text{fv } e_2$
when $x \notin \text{fv } e_2$
7. *subst-fresh-eq*:
 $\text{subst } e_1 x e_2 = e_2$
when $x \notin \text{fv } e_2$.
8. *subst-subst*:
 $\text{subst } e_1 x (\text{subst } e_2 y e) =$
 $\text{subst}(\text{subst } e_1 x e_2) y (\text{subst } e_1 x e)$
when $y \notin \text{fv } e_1$ and $y \neq x$
9. *subst-spec*:
 $\text{subst } e_1 x e_2 = \text{open } e_1 (\text{close } x e_2)$
10. *subst-open*:
 $\text{subst } e_1 x (\text{open } e_2 e_3) =$
 $\text{open}(\text{subst } e_1 x e_2) (\text{subst } e_1 x e_3)$
when $\text{lc } e_1$
11. *subst-open-var*:
 $\text{subst } e_1 x (\text{open}(\text{var_f } y) e_2) = \text{open}(\text{var_f } y) (\text{subst } e_1 x e_2)$
when $x \neq y$ and $\text{lc } e_1$
12. *subst-abs*:
 $\text{subst } e_1 x (\text{abs } e_2) = \text{abs}(\text{close } z (\text{subst } e_1 x (\text{open}(\text{var_f } z) e_2)))$
when $z \notin \text{fv } e_1 \cup \text{fv } e_2 \cup \{x\}$ and $\text{lc } e_1$
13. *subst-close*:
 $\text{subst } e_1 x (\text{close } y e_2) = \text{close } y (\text{subst } e_1 x e_2)$
when $x \neq y$ and $y \notin \text{fv } e_1$ and $\text{lc } e_1$
14. *subst-intro*:
 $\text{open } e_1 e_2 = \text{subst } e_1 x (\text{open}(\text{var_f } x) e_2)$ when $x \notin \text{fv } e_2$
15. *open-close*:
 $\text{open}(\text{var_f } x) (\text{close } x e) = e$
16. *close-open*:
 $\text{close } x (\text{open}(\text{var_f } x) e) = e$
when $x \notin \text{fv } e$
17. *open-inj*:
 $\text{open}(\text{var_f } x) e_1 = \text{open}(\text{var_f } x) e_2$
implies $e_1 = e_2$
when $x \notin \text{fv } e_1 \cup \text{fv } e_2$
18. *close-inj*:
 $\text{close } x e_1 = \text{close } x e_2$
implies $e_1 = e_2$
19. *lc-abs-exists*:
 $\text{lc_abs_exists } x \text{ lcp} : \text{lc}(\text{abs } e)$
when $\text{lcp} : \text{lc}(\text{open}(\text{var_f } x) e)$
20. *lc-subst*:
 $\text{lc_subst } \text{lcp}_1 x \text{ lcp}_2 : \text{lc}(\text{subst } e_1 x e_2)$
when $\text{lcp}_1 : \text{lc } e_1$ and $\text{lcp}_2 : \text{lc } e_2$
21. *lc-unique*:
If $(\text{lcp}_1 : \text{lc } e)$ and $(\text{lcp}_2 : \text{lc } e)$,
then $\text{lcp}_1 = \text{lcp}_2$
22. *lc-of-lc-set*:
 $\text{lc_set } e$ implies $\text{lc } e$
23. *lc-set-of-lc*:
 $\text{lc } e$ implies $\text{lc_set } e$

Figure 4: Some of the lemmas generated by LNgen

subgoals. In cases where this is not sufficient, LNgen generates more complex scripts based on our knowledge of how such proofs normally proceed. There is no worry about the soundness of our reasoning: the scripts generated by LNgen must be run by Coq to generate proof terms that are then checked.

We favor generating proof scripts over proof terms because it keeps the implementation of LNgen simple. Proof terms are specific to individual lemmas and vary from language to language. By contrast, our tactics—which are useful in their own right—apply to multiple lemmas and do not need to vary from language to language. Unfortunately, because Coq’s tactic language is incompletely specified, it is impossible for us to guarantee that our scripts will always succeed. These scripts have never failed on any of our case studies. However, if some proof should fail, the effect is localized. The user may have to do that proof by hand (if they would like to use that lemma) but other generated definitions, lemmas, and proofs will still be available.

2.4 Input Restrictions

LNgen supports only a subset of the Ott language. List forms (for specifying constructors of variable arity) and subgrammars (for indicating that, for example, values are a subset of expressions) are both unsupported. The only binding specifications accepted by LNgen are those where a single metavariable binds in a single nonterminal. This excludes Ott’s auxiliary functions for computing the set of binders in an object, e.g., those introduced by nested record patterns. We see no reason why some future version of LNgen could not be extended with these forms.

3 Soundness

Since everything generated by Ott and LNgen must be run through Coq, there is no need to worry that one is building a development on top of an inconsistent foundation—Coq will complain if a definition is ill-formed or if a proof is incomplete. However, this is not the same as saying that their outputs faithfully reflect the language that the user specified. Binding specifications in Ott use names (i.e., metavariables) to indicate binding occurrences of variables, as is common in informal practice. Intuitively, terms in the specification use a fully concrete encoding of binding: all variables are named, and terms are identified up to alpha equivalence. On the other hand, we use Ott and LNgen to generate output that uses a locally nameless representation for binding, where bound variables are represented as de Bruijn indices and where syntactic equality corresponds to alpha equivalence.

In this section, we prove that the user need not worry about this difference in representations: the locally nameless representations generated by Ott and LNgen are adequate representations of the fully concrete ones. Informally, this means that there is a bijection between the terms of the two representations and that substitution is compositional with respect to this bijection [22]. Terms representable in one representation are representable in the other, and substitution means the same thing for both representations. Below, we make these notions precise and carry out the proofs for the specific case of untyped lambda terms (Fig. 1). By considering adequacy for a particular (and small) language, we keep the proofs below relatively simple, while still demonstrating the utility of the lemmas generated by LNgen. A language-independent account of adequacy would require a precise semantics for Ott specifications and a precise specification of how Ott and LNgen generate their output. We leave developing these for future work. We also leave as future work formalization in Coq of the proofs below. Ott does not generate a definition of capture avoiding substitution or of alpha equivalence. Furthermore, *mechanized* reasoning about these notions is difficult and extremely tedious—precisely the reasons why we prefer to represent binding in some other way! Without tool support, we must work out ourselves properties of capture avoiding substitution and alpha equivalence that are ordinarily taken for granted when writing out proofs by hand.

Fully concrete lambda terms are defined in Fig. 5, along with free variables, capture-avoiding substitution, and alpha equivalence. Note that capture-avoiding substitution is defined by induction on the height of terms simultaneously with a proof that substituting a variable preserves the height of terms. (In the second case for lambda abstractions, the recursive call is not on an immediate subterm.) By assuming that picking a variable fresh for a finite set is deterministic, we obviate the need to show that the definition of substitution actually defines a function—this is trivially the case. We find it convenient to work with a definition of

Expressions

$$M, N ::= x \mid M_1 M_2 \mid \lambda x. M_1$$

Free variables

$$\begin{aligned} \text{fv}(x) &\stackrel{\text{def}}{=} \{x\} \\ \text{fv}(M_1 M_2) &\stackrel{\text{def}}{=} (\text{fv } M_1) \cup (\text{fv } M_2) \\ \text{fv}(\lambda x. M_1) &\stackrel{\text{def}}{=} (\text{fv } M_1) \setminus \{x\} \end{aligned}$$

Capture avoiding substitution

$$\begin{aligned} [N/x](x) &\stackrel{\text{def}}{=} N \\ [N/x](y) &\stackrel{\text{def}}{=} y \text{ when } y \neq x \\ [N/x](M_1 M_2) &\stackrel{\text{def}}{=} ([N/x] M_1) ([N/x] M_2) \\ [N/x](\lambda x. M_1) &\stackrel{\text{def}}{=} \lambda x. M_1 \\ [N/x](\lambda y. M_1) &\stackrel{\text{def}}{=} (\lambda z. [N/x][z/y] M_1) \\ &\text{for some } z \notin \text{fv } N \cup \text{fv } M_1 \text{ and when } y \neq x \end{aligned}$$

Alpha equivalence

The binary relation $=_\alpha$ on expressions is the least congruence closed under

$$\lambda x. M_1 =_\alpha \lambda y. [y/x] M_1 \quad \text{when } y \notin \text{fv } M_1$$

Figure 5: Fully concrete lambda terms

capture-avoiding substitution that is total, so the abstraction case always renames the bound variable to avoid capture.

To show the adequacy of our locally nameless representation, we prove that there is an alpha-equivalence respecting bijection between concrete terms and locally nameless terms that are locally closed. We give this bijection by defining functions between the two sets and then proving that they are inverses of each other. We define the function $\llbracket - \rrbracket$ from concrete terms to locally nameless ones as follows:

$$\begin{aligned} \llbracket x \rrbracket &\stackrel{\text{def}}{=} \text{var_f } x \\ \llbracket M_1 M_2 \rrbracket &\stackrel{\text{def}}{=} \text{app } \llbracket M_1 \rrbracket \llbracket M_2 \rrbracket \\ \llbracket \lambda x. M_1 \rrbracket &\stackrel{\text{def}}{=} \text{abs } (\text{close } x \llbracket M_1 \rrbracket). \end{aligned}$$

The fact that this function yields only locally closed terms follows by structural induction on its argument, using lemmas *lc-abs-exists* and *open-close* in the case for abstractions. We define the function $\llbracket - \rrbracket$ from locally nameless terms that are locally closed to concrete terms using the recursion principle in Fig. 3. Note that this definition is also a function, again because we assume that picking a fresh variable not in a particular set is deterministic.

$$\begin{aligned} \llbracket \text{var_f } x \rrbracket &\stackrel{\text{def}}{=} x \\ \llbracket \text{app } e_1 e_2 \rrbracket &\stackrel{\text{def}}{=} \llbracket e_1 \rrbracket \llbracket e_2 \rrbracket \\ \llbracket \text{abs } e_1 \rrbracket &\stackrel{\text{def}}{=} \lambda x. \llbracket \text{open } (\text{var_f } x) e_1 \rrbracket \quad \text{for some } x \notin \text{fv } e_1 \end{aligned}$$

In the remainder of this section, we sketch out the proof of adequacy; additional details can be found in the appendix. The proofs below are straightforward given the lemmas generated by LNgen. We need only to be careful about ordering properly the lemmas and theorems.

We first need to show that both $\llbracket - \rrbracket$ and $\llbracket - \rrbracket$ preserve free variables. These proofs also serve as basic sanity checks: it would be odd for corresponding terms in the two representations to have different sets of free variables.

Lemma 1 $\text{fv}(M) = \text{fv}(\lceil M \rceil)$ for any M .

Proof

By induction on the structure of M . In the case for abstractions, we need lemma *fv-close*. \square \square

Lemma 2 $\text{fv}(\lfloor e \rfloor) = \text{fv}(e)$ for any locally closed e .

Proof

By induction on the proof that e is locally closed. In the case for abstractions, we need lemmas *fv-close* and *close-open*. \square \square

Next, we prove simultaneously that $\lceil - \rceil$ commutes with substitution and that it preserves alpha equivalence. For $\lfloor - \rfloor$, we prove that it commutes with substitution; it trivially preserves alpha equivalence.

Theorem 3 For all M ,

1. Substitution commutes with $\lceil - \rceil$. That is, for any N and x ,

$$\lceil [N / x] M \rceil = \text{subst } \lceil N \rceil x \lceil M \rceil.$$

2. $\lceil - \rceil$ respects alpha-equivalence. That is, for any N such that $N =_\alpha M$,

$$\lceil M \rceil = \lceil N \rceil.$$

Proof

We prove these two results simultaneously by induction on the height of M , observing that substituting a variable does not change the height of a term. We need lemmas *fv-close*, *subst-fresh-eq*, *subst-spec*, *subst-close*, and *close-open*. \square

Theorem 4 $\lfloor - \rfloor$ commutes with substitution. That is,

$$\lfloor \text{subst } g x e \rfloor =_\alpha \lfloor [g] / x \rfloor \lfloor e \rfloor$$

for all locally closed e and g , and for all x .

Proof

By induction on the proof that e is locally closed. In the case for abstractions, we need lemmas *subst-fresh-eq*, *subst-spec*, *subst-abs*, *open-close*, and *close-open*. \square \square

Finally, we prove that $\lceil - \rceil$ and $\lfloor - \rfloor$ are inverses of each other. It follows that each function defines a bijection.

Theorem 5 $\lfloor \lceil M \rceil \rfloor =_\alpha M$ for any M .

Proof

By induction on the structure of M . In the case for abstractions, we need theorem 4, and lemmas *fv-close* and *subst-spec*. \square \square

Theorem 6 $\lceil \lfloor e \rfloor \rceil = e$ for any locally closed e .

Proof

By induction on the proof that e is locally closed. In the case for abstractions, we need lemma *close-open*. \square \square

Taken together, theorems 3–6 suffice to prove that the locally nameless representation generated by Ott and LNgen is adequate with respect to the fully concrete interpretation of the original Ott specification.

4 Completeness

Does LNgen generate enough definitions and properties to get work done? Of course, this is an impossible question to answer because the tool cannot possibly generate proofs of every property that one could need or want. However, we can limit the scope of the question by showing that LNgen *trivially* models some specification of binding. By choosing a specification that makes no mention of de Bruijn indices, this result implies that the user need only work with locally-closed terms and never reason explicitly about de Bruijn indices.

We make our claim by showing that the output of Ott and LNgen for untyped lambda terms (Fig. 1) is not very far from Gordon and Melham’s “Five Axioms of Alpha-Conversion” [19]. In fact, we can derive these axioms with only currying, uncurrying, and applications of lemmas generated by LNgen. This work is a bit tedious, but none of it includes reasoning about de Bruijn indices, doing induction on raw expressions, or doing induction on local closure derivations. Thus, it substantiates our claim that the output of our tool provides users with enough machinery to reason about binding. The LNgen distribution includes a straightforward, mechanical formalization in Coq of the results of this section.

Gordon and Melham’s five axioms are defined in terms of a type `Term`, three constructors for that type,

$$\begin{aligned} \text{Var} & : \text{expvar} \rightarrow \text{Term} \\ \text{App} & : \text{Term} \rightarrow \text{Term} \rightarrow \text{Term} \\ \text{Lam} & : \text{expvar} \rightarrow \text{Term} \rightarrow \text{Term}, \end{aligned}$$

and three operations for that type,

$$\begin{aligned} \text{Fv} & : \text{Term} \rightarrow \text{expvarset} \\ \text{Subst} & : \text{Term} \rightarrow (\text{Term} \times \text{expvar}) \rightarrow \text{Term} \\ \text{Abs} & : (\text{expvar} \rightarrow \text{Term}) \rightarrow \text{Term}. \end{aligned}$$

Our implementation starts by defining `Term` as a dependent pair of a raw expression and a proof that it is locally closed.

Definition 7 (Term)

$$\text{Term} \stackrel{\text{def}}{=} \Sigma e : \text{exp}. \text{lc } e.$$

The definitions of the three constructors simply construct and propagate local closure proofs. In the definition of `Lam`, we explicitly use the “existential” version of `lc_abs` (i.e., `lc_abs_exists`) and implicitly use lemma *open-close* to show that the local closure proof applies to first component of the tuple.

Definition 8 (Gordon-Melham Constructors)

$$\begin{aligned} \text{Var } x & \stackrel{\text{def}}{=} (\text{var_f } x, \text{lc_var } x) \\ \text{App } (e_1, \text{lc}_1) (e_2, \text{lc}_2) & \stackrel{\text{def}}{=} (\text{app } e_1 e_2, \text{lc_app } \text{lc}_1 \text{lc}_2) \\ \text{Lam } x (e_1, \text{lc}_1) & \stackrel{\text{def}}{=} (\text{abs } (\text{close } x e_1), \text{lc_abs_exists } x \text{lc}_1) \end{aligned}$$

The definitions for free variables (`Fv`) and substitution (`Subst`) simply push the operations on raw terms through the dependent pair. For substitution, we rely on the fact that substitution preserves local closure.

Definition 9 (Fv and Subst)

$$\begin{aligned} \text{Fv } (e_1, \text{lc}_1) & \stackrel{\text{def}}{=} \text{fv } e_1 \\ \text{Subst } (e_1, \text{lc}_1) ((e_2, \text{lc}_2), x) & \stackrel{\text{def}}{=} (\text{subst } e_2 x e_1, \text{lc_subst } \text{lc}_2 x \text{lc}_1) \end{aligned}$$

The final operation, `Abs`, reifies a function from variable names to terms into a lambda term. We defer its definition until later, when we discuss the last of the five axioms.

With the model above, we can derive Gordon and Melham’s five axioms. The proofs of their five axioms involve little more than projecting out components of dependent pairs and applying lemmas generated by

LNgen to construct local closure derivations. In fact, the only interesting aspect of these proofs is that they are so uninteresting. Below, we only mention the lemmas that the proofs depend on; additional details can be found in the appendix.

The first three axioms are basic facts about free variables, capture-avoiding substitution, and alpha conversion.

Theorem 10 (Axiom 1: Free variables)

1. $\text{Fv}(\text{Var } x) = \{ x \}$
2. $\text{Fv}(\text{App } t_1 t_2) = \text{Fv } t_1 \cup \text{Fv } t_2$
3. $\text{Fv}(\text{Lam } x t_1) = \text{Fv } t_1 \setminus \{ x \}$

Proof

By unfolding definitions. Part 3 requires lemma *fv-close*. □ □

Theorem 11 (Axiom 2: Substitution)

1. $\text{Subst}(\text{Var } x)(u, x) = u$
2. $x \neq y$ implies $\text{Subst}(\text{Var } y)(u, x) = \text{Var } y$
3. $\text{Subst}(\text{App } t_1 t_2)(u, x) = \text{App}(\text{Subst } t_1(u, x))(\text{Subst } t_2(u, x))$
4. $\text{Subst}(\text{Lam } x t)(u, x) = \text{Lam } x t$
5. $x \neq y$ and $y \notin (\text{Fv } u)$ imply $\text{Subst}(\text{Lam } y t)(u, x) = \text{Lam } y(\text{Subst } t(u, x))$

Proof

By unfolding definitions. All parts require lemma *lc-unique*. Part 4 also requires lemmas *fv-close* and *subst-fresh-eq*. Part 5 also requires lemma *subst-close*. □

□

Theorem 12 (Axiom 3: Alpha conversion)

$$y \notin \text{Fv}(\text{Lam } x t) \text{ implies } \text{Lam } x t = \text{Lam } y(\text{Subst } t(\text{Var } y, x))$$

Proof

By unfolding definitions. The proof requires lemmas *fv-close*, *subst-spec*, *close-open*, and *lc-unique*. □ □

To support the definition of functions over lambda-calculus expressions, Gordon and Melham's work states an iteration axiom and uses it to derive a recursion scheme through pairing. However, because Coq produces recursion schemes already, we define the recursion scheme directly. The iterative version follows as a simple corollary.

Theorem 13 (Axiom 4: Recursion scheme) *For all result types R and all*

$$\begin{aligned} & (fvar : \text{expvar} \rightarrow R) \\ & (fapp : R \rightarrow R \rightarrow \text{Term} \rightarrow \text{Term} \rightarrow R) \\ & (fabs : (\text{expvar} \rightarrow R) \rightarrow (\text{expvar} \rightarrow \text{Term}) \rightarrow R), \end{aligned}$$

there exists a unique f of type $\text{Term} \rightarrow R$ such that

1. $f(\text{Var } x) = fvar x$
2. $f(\text{App } t_1 t_2) = fapp(f t_1)(f t_2) t_1 t_2$

3. $f(\text{Lam } x \ t) = f_{\text{abs}}(\lambda y. f(\text{Subst } t(\text{Var } y, x))) (\lambda y. \text{Subst } t(\text{Var } y, x))$.

Proof

By unfolding definitions. All parts require lemma *lc-unique*. Part 3 also requires lemma *subst-spec*. \square \square

The final axiom concerns **Abs**, an operation for turning functions from `expvars` to `Terms` into lambda abstractions. This operation allows the Gordon-Melham recursion combinator to create a new term in the lambda case. The trickiest part of the definition of **Abs** is picking a variable name to use for the binder that is fresh for the body of the abstraction. We do this in two stages: We first access the body with an arbitrary variable x_0 (which may already appear in the body), and then we use the resulting term to pick a variable certain to be fresh for body. We use `lc_abs_exists` and lemma *open-close* similarly to how we did in the definition of `Lam`.

Definition 14 (Abs)

$$\begin{aligned} \text{Abs } f &= (\text{abs}(\text{close } y \ e_2), \text{lc_abs_exists } y \ \text{lc } p_2) \\ \text{where } (e_1, -) &= f \ x_0 \\ y &\notin (\text{fv } e_1) \\ (e_2, \text{lc } p_2) &= f \ y \end{aligned}$$

With **Abs** defined, we can now state and derive the final axiom.

Theorem 15 (Axiom 5: Abstraction)

$$\text{Abs}(\lambda y. \text{Subst } t(\text{Var } y, x)) = \text{Lam } x \ t$$

Proof

By unfolding definitions. The proof requires lemmas *fv-close*, *fv-subst-lower*, *subst-spec*, *close-open*, and *lc-unique*. \square

\square

The abstraction operation is the only definition that is not trivial in that it first must calculate a fresh variable for the term. The advantage of axiom 5 is that it lets one have a lambda expression without naming its binder. However, in some sense, **Abs** is not necessary for our style of reasoning. Certainly, all of this effort is not required to define functions with `lc_set_rec`, e.g., *beta* in Sec. 2.1.

5 Case Studies

We have used `LNgen` to streamline proofs of type safety for the simply-typed lambda calculus and for System F with subtyping, i.e., parts 1A and 2A of the `POPLMARK` challenge. In both cases, the only proofs that needed to be mechanized by hand were lemmas about the relations of their respective systems. (Because `LNgen` works only with syntax, it cannot be expected to generate these proofs.) Every necessary lemma concerning only the calculation of free variables, substitution, and local closure was automatically proved by `LNgen`.

Others have used `LNgen` for far more substantial developments than the two above. Greenberg et al. [23] used `LNgen` to help formalize a proof of confluence for parallel reduction in dependent λ^h , a language with manifest contracts. Greenberg reports⁴ that, “All in all, `LNgen` was great—it covered most of the stupid facts I needed.” The tool failed to generate only one set of lemmas, which concerned how substitution maintains invariants about the free variables of terms. Jia et al. [24] used `LNgen` when they proved type soundness for a dependently-typed language with strong eliminators and an abstract definition of program equivalence. The authors report⁵ that without the 9000 lines of lemmas and proofs that `LNgen` generated for their language, they would have been unable to complete their formalization in a timely fashion. Because the tool provided every infrastructure lemma they needed, they were able to focus their efforts on the novel

⁴By personal communication.

⁵Again, by personal communication.

aspects of their language’s design and complete their formalization in about nine days—an impressive feat given the complexity of their design and the fact that they were tweaking the design in the process. Taken together, these two non-trivial developments provide a compelling story about the effectiveness of LNgen in eliminating the tedium associated with locally nameless encodings.

6 Related Work

Much work has been done in the area of representing binding. For example, we have already discussed the “Five Axioms of Alpha-Conversion.” In previous work [14], we also gave an extensive survey of first-order representation techniques. Thus, we focus this section on work that is specifically related to the issues described in this paper.

Logical frameworks—such as Abella [1], Hybrid [2], and Twelf [5]—are specifically designed to represent and reason about logics and programming languages. Their specialized meta-logics encourage the use of higher-order abstract syntax (HOAS), which represents binding in an object language using binding in the framework’s meta-logic. Thus, when reasoning about an object language, one gets facts about alpha equivalence, substitution, and free variables “for free.” Unfortunately, the generality of Coq’s logic precludes traditional HOAS encodings, and first-order representations (e.g., locally nameless) require that one explicitly deal with free variable calculation and substitution. LNgen steps in here to recover the benefits to working in a traditional logical framework by automatically proving properties about syntax that one expects to have “for free.”

The Lambda Tamer project [3] also automatically proves a variety of facts about programming languages encoded in Coq. Compared to LNgen, Lambda Tamer favors the use of dependent types when representing syntax, ensuring that only well-typed syntax, according to the type system of the *object* language, can be represented. It uses generic programming techniques to ensure that generated proofs are correct by construction. As mentioned previously (Sec. 2.3), we prefer to generate proof scripts because of the approach’s simplicity—writing generic proofs directly is a non-trivial exercise and would have slowed the development of LNgen.

Parametric higher-order abstract syntax (PHOAS) [25] is a representation technique that allows one to use HOAS-like approach to represent binding, thus obtaining “for free” facts about syntax that LNgen has to prove about locally nameless encodings. The key idea is to represent the body of an abstraction not as a function from expression to expressions, as with HOAS, but as a function from *variables* to expressions, an approach reminiscent of weak HOAS [10]. Ill-formed terms are ruled out by universally quantifying over the type of variables and appealing to parametricity to ensure that the type for variables is treated abstractly. Without a general proof of parametricity for Coq, one must assert that parametricity holds for particular terms as needed or as an axiom.

7 Conclusions and Future Work

Since LNgen is currently only a prototype, there are a number of promising avenues for future development and research. We developed LNgen independently from Ott in order to make it easier to experiment with its output: which definitions to generate, which lemmas to generate, how to generate proofs, etc. But, it might be beneficial to add such support to Ott directly. The ideas we have presented here are not particular to Coq, and we expect that they can be generalized to the full spectrum of Ott’s binding forms. We also believe that it is possible to automatically generate theorems about some judgements: equivariance (invariance under swappings of variables), weakening, and substitution, for example. Support for defining functions directly in Ott specifications and having them translated into locally nameless definitions, using schemes such as `lc_set_rec`, would also be useful. In particular, one would like to know that something similar to the “freshness condition for binders” from Nominal Isabelle holds whenever a function is defined. In the case of binding constructors, this would allow one to conclude that the behavior of the function does not depend on the particular choice of name for the bound variable (recall the definition of *beta* in Sec. 2.1). On a more theoretical note, we envision giving a general account of how to transform a fully concrete representation into a locally nameless one, thus making it possible to give a general account of soundness for LNgen.

In the end, what we provide *now* is a usable prototype tool for taking our locally nameless style—already a lightweight representation technique—and making it even lighter weight. We have shown that Ott and LNgen are sound and complete in the specific case of untyped lambda terms. Compared to our previous work, we now provide a recursion scheme for defining functions, and it comes “for free” from our definitions. On a day to day basis, the benefit of our work is simple: no more boring infrastructure proofs.

Acknowledgements

This work was funded by DARPA, *CSSG Phase II: Machine Checked Metatheory for Security-Oriented Languages*.

References

- [1] Gacek, A.: The Abella interactive theorem prover (system description). In Armando, A., Baumgartner, P., Dowek, G., eds.: *Automated Reasoning: Fourth International Joint Conference, IJCAR 2008*. Volume 5195 of *Lecture Notes in Artificial Intelligence*. Springer (2008) 154–161
- [2] Momigliano, A., Martin, A.J., Felty, A.P.: Two-level Hybrid: A system for reasoning using higher-order abstract syntax. In Abel, A., Urban, C., eds.: *Proceedings of the International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2008)*. Volume 228 of *Electronic Notes in Theoretical Computer Science*. Elsevier (2009) 85–93
- [3] Chlipala, A.: Generic programming and proving for programming language metatheory. Technical Report UCB/EECS-2007-147, University of California, Berkeley (2007)
- [4] Urban, C.: Nominal techniques in Isabelle/HOL. *Journal of Automated Reasoning* **40**(4) (2008) 327–356
- [5] Pfenning, F., Schürmann, C.: System description: Twelf — A meta-logical framework for deductive systems. In Ganzinger, H., ed.: *Automated Deduction, CADE 16: 16th International Conference on Automated Deduction*. Volume 1632 of *Lecture Notes in Artificial Intelligence*. Springer (1999) 202–206
- [6] de Bruijn, N.G.: Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae* **34**(5) (1972) 381–392
- [7] Pfenning, F., Elliot, C.: Higher-order abstract syntax. In: *PLDI '88: Proceedings of the ACM SIGPLAN 1988 Conference on Programming Language Design and Implementation*. ACM (1988) 199–208
- [8] McKinna, J., Pollack, R.: Some lambda calculus and type theory formalized. *Journal of Automated Reasoning* **23**(3–4) (1999) 373–409
- [9] Pollack, R.: Closure under alpha-conversion. In Barendregt, H., Nipkow, T., eds.: *Types for Proofs and Programs: International Workshop, TYPES 1993*. Volume 806 of *Lecture Notes in Computer Science*. Springer (1994) 313–332
- [10] Despeyroux, J., Felty, A., Hirschowitz, A.: Higher-order abstract syntax in Coq. In: *Typed Lambda Calculi and Applications, Second International Conference on Typed Lambda Calculi and Applications, TLCA '95*. Volume 902 of *Lecture Notes in Computer Science*. Springer (1995) 124–138. Also available as INRIA Research report 2556
- [11] Aydemir, B.E., Bohannon, A., Fairbairn, M., Foster, J.N., Pierce, B.C., Sewell, P., Vytiniotis, D., Washburn, G., Weirich, S., Zdancewic, S.: Mechanized metatheory for the masses: The POPLMARK challenge. In Hurd, J., Melham, T., eds.: *Theorem Proving in Higher Order Logics: 18th International Conference, TPHOLs 2005*. Volume 3603 of *Lecture Notes in Computer Science*. Springer (2005) 50–65
- [12] Pitts, A.M.: Nominal logic, a first order theory of names and binding. *Information and Computation* **186** (2003) 165–193

- [13] Aydemir, B., Bohannon, A., Weirich, S.: Nominal reasoning techniques in Coq (extended abstract). In Momigliano, A., Pientka, B., eds.: *Proceedings of the First International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2006)*. Volume 174 of *Electronic Notes in Theoretical Computer Science*. Elsevier (2007) 69–77
- [14] Aydemir, B., Charguéraud, A., Pierce, B.C., Pollack, R., Weirich, S.: Engineering formal metatheory. In: *POPL '08: Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM (2008) 3–15
- [15] Jia, L., Vaughan, J.A., Mazurak, K., Zhao, J., Zarko, L., Schorr, J., Zdancewic, S.: AURA: A programming language for authorization and audit. In: *ICFP '08: Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming*. ACM (2008) 27–38
- [16] Pratikakis, P., Foster, J.S., Hicks, M., Neamtiu, I.: Formalizing soundness of contextual effects. In Ait Mohamed, O., Muñoz, C., Tahar, S., eds.: *Theorem Proving in Higher Order Logics: 21st International Conference, TPHOLs 2008*. Volume 5170 of *Lecture Notes in Computer Science*. Springer (2008) 262–277
- [17] Rossberg, A., Russo, C., Dreyer, D.: F-ing modules. Submitted for publication (October 2010)
- [18] Sewell, P., Zappa Nardelli, F., Owens, S., Peskine, G., Ridge, T., Sarkar, S., Strniša, R.: Ott: Effective tool support for the working semanticist. In: *ICFP '07: Proceedings of the 2007 ACM SIGPLAN International Conference on Functional Programming*. ACM (2007) 1–12
- [19] Gordon, A.D., Melham, T.: Five axioms of alpha-conversion. In von Wright, J., Grundy, J., Harrison, J., eds.: *Theorem Proving in Higher Order Logics: 9th International Conference, TPHOLs '96*. Volume 1125 of *Lecture Notes in Computer Science*. Springer (1996) 173–190
- [20] Sewell, P., Zappa Nardelli, F.: Ott. Available from <http://www.cl.cam.ac.uk/~pes20/ott/> (2009)
- [21] Gordon, A.D.: A mechanisation of name-carrying syntax up to alpha-conversion. In Joyce, J.J., Seger, C.J.H., eds.: *Higher-order Logic Theorem Proving And Its Applications, Proceedings, 1993*. Volume 780 of *Lecture Notes in Computer Science*. Springer (1994) 414–426
- [22] Harper, R., Honsell, F., Plotkin, G.: A framework for defining logics. *Journal of the ACM* **40**(1) (1993) 143–184
- [23] Greenberg, M., Pierce, B., Weirich, S.: Contracts made manifest. In: *POPL '10: Proceedings of the 37th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, Madrid, Spain, ACM (January 2010)*. To appear.
- [24] Jia, L., Zhao, J., Sjöberg, V., Weirich, S.: Dependent types and program equivalence. In: *POPL '10: Proceedings of the 37th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, Madrid, Spain, ACM (January 2010)*. To appear.
- [25] Chlipala, A.: Parametric higher-order abstract syntax for mechanized semantics. In: *ICFP '08: Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming*. ACM (2008) 143–156

A Proofs

A.1 Proof of Theorem 3

We prove these two results simultaneously by induction on the height of M , observing that substituting a variable does not change the height of a term. We need lemmas *fv-close*, *subst-fresh-eq*, *subst-spec*, *subst-close*, and *close-open*. For example, the abstraction case of the second part is shown below, where $M = (\lambda x . M_1)$ and $N = \lambda y . [y/x] M_1$, with $y \notin \text{fv } M_1$.

$$\begin{aligned}
& \lceil \lambda x . M_1 \rceil \\
& \quad \text{by definition of } \lceil - \rceil \\
& = \text{abs}(\text{close } x \lceil M_1 \rceil) \\
& \quad \text{by lemma } \textit{close-open} \\
& = \text{abs}(\text{close } y (\text{open } (\text{var}_f y) (\text{close } x \lceil M_1 \rceil))) \\
& \quad \text{by lemma } \textit{subst-spec} \\
& = \text{abs}(\text{close } y (\text{subst } (\text{var}_f y) x \lceil M_1 \rceil)) \\
& \quad \text{by IH(1)} \\
& = \text{abs}(\text{close } y \lceil [y/x] M_1 \rceil) \\
& \quad \text{by definition of } \lceil - \rceil \\
& = \lceil \lambda y . [y/x] M_1 \rceil
\end{aligned}$$

A.2 Proof of Theorem 4

By induction on the proof that e is locally closed. The only interesting case is when $e = \text{abs } e_1$, for some e_1 . Let $y \notin \text{fv } e_1$ and $w \notin \text{fv } g \cup \{x\} \cup \text{fv } e_1 \cup \{y\}$. We consider two cases for x and y . First, suppose that $x \neq y$.

$$\begin{aligned}
& \lceil [g] / x \rceil \lceil \text{abs } e_1 \rceil \\
& \quad \text{for some } y \notin \text{fv } e_1 \\
& = \lceil [g] / x \rceil (\lambda y . \lceil \text{open } (\text{var}_f y) e_1 \rceil) \\
& \quad \text{by alpha conversion} \\
& =_\alpha \lceil [g] / x \rceil (\lambda w . \lceil [w/y] \lceil \text{open } (\text{var}_f y) e_1 \rceil \rceil) \\
& \quad \text{by property of substitution} \\
& =_\alpha \lambda w . \lceil [g] / x \rceil \lceil [w/y] \lceil \text{open } (\text{var}_f y) e_1 \rceil \rceil \\
& \quad \text{by IH} \\
& =_\alpha \lambda w . \lceil [g] / x \rceil (\lceil \text{subst } (\text{var}_f w) y (\text{open } (\text{var}_f y) e_1) \rceil) \\
& \quad \text{by lemmas } \textit{subst-spec} \text{ and } \textit{close-open} \\
& = \lambda w . \lceil [g] / x \rceil (\lceil \text{open } (\text{var}_f w) e_1 \rceil) \\
& \quad \text{by IH} \\
& =_\alpha \lambda w . \lceil \text{subst } g x (\text{open } (\text{var}_f w) e_1) \rceil \\
& \quad \text{by lemma } \textit{open-close} \\
& = \lambda w . \lceil \text{open } (\text{var}_f w) (\text{close } w (\text{subst } g x (\text{open } (\text{var}_f w) e_1))) \rceil \\
& \quad \text{by the definition of } \lceil - \rceil \\
& = \lceil \text{abs}(\text{close } w (\text{subst } g x (\text{open } (\text{var}_f w) e_1))) \rceil \\
& \quad \text{by lemma } \textit{subst-abs} \\
& = \lceil \text{subst } g x (\text{abs } e_1) \rceil
\end{aligned}$$

Second, suppose that $x = y$. Since $y \notin \text{fv } e_1$, it is also the case that $x \notin \text{fv } e_1$.

$$\begin{aligned}
& \llbracket [g] / x \rrbracket \llbracket \text{abs } e_1 \rrbracket \\
& \quad \text{for some } y \notin \text{fv } e_1 \\
& = \llbracket [g] / x \rrbracket (\lambda y . \llbracket \text{open } (\text{var.f } y) e_1 \rrbracket) \\
& \quad \text{by definition of substitution} \\
& = \lambda y . \llbracket \text{open } (\text{var.f } y) e_1 \rrbracket \\
& \quad \text{by definition of } \llbracket - \rrbracket \\
& = \llbracket \text{abs } e_1 \rrbracket \\
& \quad \text{by lemma } \textit{subst-fresh-eq} \\
& = \llbracket \text{subst } g \ x \ (\text{abs } e_1) \rrbracket
\end{aligned}$$

A.3 Proof of Theorem 5

By induction on the structure of M . The only interesting case is when $M = \lambda x . M_1$, for some M_1 .

$$\begin{aligned}
& \llbracket \llbracket \lambda x . M_1 \rrbracket \rrbracket \\
& \quad \text{by definition of } \llbracket - \rrbracket \\
& = \llbracket \text{abs } (\text{close } x \llbracket M_1 \rrbracket) \rrbracket \\
& \quad \text{for some } y \notin \text{fv } (\text{close } x \llbracket M_1 \rrbracket) \\
& = \lambda y . \llbracket \text{open } (\text{var.f } y) (\text{close } x \llbracket M_1 \rrbracket) \rrbracket \\
& \quad \text{by lemma } \textit{subst-spec} \\
& = \lambda y . \llbracket \text{subst } (\text{var.f } y) \ x \ \llbracket M_1 \rrbracket \rrbracket \\
& \quad \text{by theorem 4} \\
& = \lambda y . \llbracket y / x \rrbracket \llbracket \llbracket M_1 \rrbracket \rrbracket \\
& \quad \text{by alpha conversion and lemma } \textit{fv-close} \\
& =_\alpha \lambda x . \llbracket \llbracket M_1 \rrbracket \rrbracket \\
& \quad \text{by IH} \\
& =_\alpha \lambda x . M_1
\end{aligned}$$

Note that in the alpha conversion step, we assumed that $x \neq y$. When $x = y$, the result follows trivially.

A.4 Proof of Theorem 6

By induction on the proof that e is locally closed. The only interesting case is when $e = \text{abs } e_1$, for some e_1 .

$$\begin{aligned}
& \llbracket \llbracket \text{abs } e_1 \rrbracket \rrbracket \\
& \quad \text{for some } y \notin \text{fv } e_1 \\
& = \llbracket \lambda y . \llbracket \text{open } (\text{var.f } y) e_1 \rrbracket \rrbracket \\
& \quad \text{by definition of } \llbracket - \rrbracket \\
& = \text{abs } (\text{close } y \llbracket \llbracket \text{open } (\text{var.f } y) e_1 \rrbracket \rrbracket) \\
& \quad \text{by IH} \\
& = \text{abs } (\text{close } y (\text{open } (\text{var.f } y) e_1)) \\
& \quad \text{by lemma } \textit{close-open} \\
& = \text{abs } e_1
\end{aligned}$$

A.5 Proof of Theorem 11

We first observe that for any e , any two derivations of $\text{lc } e$ are equal by *lc-unique*. Therefore, to show that each equality holds, it suffices to show that the first components of each side of the equality are equal. In the proofs below, we use $_$ as a place holder for the second components.

After unfolding definitions, parts 1, 2, and 3 are trivial.

proofs, much as we did in the proof of theorem 11, we have the following:

$$\begin{aligned}
& f(\text{Lam } x (e_1, -)) \\
& \quad \text{by definition} \\
& = f'(\text{abs}(\text{close } x e_1)) - \\
& \quad \text{by property of } \text{lc_set_rec} \\
& = \text{fabs}'(\text{close } x e_1) - (\lambda y. f(\text{open}(\text{var_f } y)(\text{close } x e_1), -)) \\
& \quad \text{by definition of } \text{fabs}' \\
& = \text{fabs}(\lambda y. f(\text{open}(\text{var_f } y)(\text{close } x e_1), -)) (\lambda y. (\text{open}(\text{var_f } y)(\text{close } x e_1), -)) \\
& \quad \text{by lemma } \text{subst-spec} \\
& = \text{fabs}(\lambda y. f(\text{subst}(\text{var_f } y) x e_1, -)) (\lambda y. (\text{subst}(\text{var_f } y) x e_1, -)) \\
& \quad \text{by definition of } \text{Subst} \\
& = \text{fabs}(\lambda y. f(\text{Subst } t(\text{Var } y, x))) (\lambda y. (\text{Subst } t(\text{Var } y, x)))
\end{aligned}$$

A.8 Proof of Theorem 15

We decompose t as $(e_1, -)$ and make use of lemma *lc-unique* in the same way we did as in the proof of theorem 11.

$$\begin{aligned}
& \text{Abs}(\lambda y. \text{Subst } t(\text{Var } y, x)) \\
& \quad \text{by definition of } \text{Abs} \text{ and } \text{Subst} \\
& \quad \text{for some } y \notin \text{Fv}(\text{Subst } t(\text{Var } x_0, x)) \\
& = (\text{abs}(\text{close } y(\text{subst}(\text{var_f } y) x e_1)), -) \\
& \quad \text{by lemma } \text{subst-spec} \\
& = (\text{abs}(\text{close } y(\text{open}(\text{var_f } y)(\text{close } x e_1))), -) \\
& \quad \text{by lemma } \text{close-open}, \\
& \quad \text{discharging the side condition by lemmas } \text{fv-close} \text{ and } \text{fv-subst-lower} \\
& = (\text{abs}(\text{close } x e_1), -) \\
& \quad \text{by definition of } \text{Lam} \text{ and lemma } \text{lc-unique} \\
& = \text{Lam } x t
\end{aligned}$$