January 2008

Engineering Formal Metatheory

Brian Aydemir
University of Pennsylvania

Arthur Charguéraud
INRIA

Benjamin C. Pierce
University of Pennsylvania, bcpierce@cis.upenn.edu

Randy Pollack
University of Edinburgh

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

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

Recommended Citation

Brian Aydemir, Arthur Charguéraud, Benjamin C. Pierce, Randy Pollack, and Stephanie Weirich, "Engineering Formal Metatheory", . January 2008.

Postprint version. Published in *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, January 2008, pages 3-15.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_papers/369
For more information, please contact repository@pobox.upenn.edu.

Engineering Formal Metatheory

Abstract

Machine-checked proofs of properties of programming languages have become a critical need, both for increased confidence in large and complex designs and as a foundation for technologies such as proof-carrying code. However, constructing these proofs remains a black art, involving many choices in the formulation of definitions and theorems that make a huge cumulative difference in the difficulty of carrying out large formal developments. The representation and manipulation of terms with variable binding is a key issue. We propose a novel style for formalizing metatheory, combining *locally nameless* representation of terms and cofinite quantification of free variable names in inductive definitions of relations on terms (typing, reduction, ...). The key technical insight is that our use of cofinite quantification obviates the need for reasoning about equivariance (the fact that free names can be renamed in derivations); in particular, the structural induction principles of relations defined using cofinite quantification are strong enough for metatheoretic reasoning, and need not be explicitly strengthened. Strong inversion principles follow (automatically, in Coq) from the induction principles. Although many of the underlying ingredients of our technique have been used before, their combination here yields a significant improvement over other methodologies using first-order representations, leading to developments that are faithful to informal practice, yet require no external tool support and little infrastructure within the proof assistant. We have carried out several large developments in this style using the Coq proof assistant and have made them publicly available. Our developments include type soundness for System *F* and core *ML* (with references, exceptions, datatypes, recursion, and patterns) and subject reduction for the Calculus of Constructions. *Not only do these developments demonstrate the comprehensiveness of our approach; they have also been optimized for clarity and robustness, making them good templates for future extension.*

Keywords

binding, Coq, locally nameless

Comments

Postprint version. Published in Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), January 2008, pages 3-15.

Engineering Formal Metatheory

Brian Aydemir

University of Pennsylvania
baydemir@cis.upenn.edu

Arthur Charguéraud

INRIA
arthur.chargueraud@inria.fr

Benjamin C. Pierce

University of Pennsylvania
bcpierce@cis.upenn.edu

Randy Pollack

University of Edinburgh
rpollack@inf.ed.ac.uk

Stephanie Weirich

University of Pennsylvania
sweirich@cis.upenn.edu

Abstract

Machine-checked proofs of properties of programming languages have become a critical need, both for increased confidence in large and complex designs and as a foundation for technologies such as proof-carrying code. However, constructing these proofs remains a black art, involving many choices in the formulation of definitions and theorems that make a huge cumulative difference in the difficulty of carrying out large formal developments. The representation and manipulation of terms with variable binding is a key issue.

We propose a novel style for formalizing metatheory, combining *locally nameless* representation of terms and cofinite quantification of free variable names in inductive definitions of relations on terms (typing, reduction, ...). The key technical insight is that our use of cofinite quantification obviates the need for reasoning about equivariance (the fact that free names can be renamed in derivations); in particular, the structural induction principles of relations defined using cofinite quantification are strong enough for metatheoretic reasoning, and need not be explicitly strengthened. Strong inversion principles follow (automatically, in Coq) from the induction principles. Although many of the underlying ingredients of our technique have been used before, their combination here yields a significant improvement over other methodologies using first-order representations, leading to developments that are faithful to informal practice, yet require no external tool support and little infrastructure within the proof assistant.

We have carried out several large developments in this style using the Coq proof assistant and have made them publicly available. Our developments include type soundness for System F_{λ} and core ML (with references, exceptions, datatypes, recursion, and patterns) and subject reduction for the Calculus of Constructions. Not only do these developments demonstrate the comprehensiveness of our approach; they have also been optimized for clarity and robustness, making them good templates for future extension.

Categories and Subject Descriptors F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Mechanical theo-

rem proving; D.3.1 [Programming Languages]: Formal Definitions and Theory—Syntax

General Terms Design, Documentation, Languages, Theory, Verification

Keywords binding, Coq, locally nameless

1. Introduction

Recent years have seen burgeoning interest in the use of proof assistants for formalizing definitions of programming languages and checking proofs of their properties. However, despite several successful *tours de force* (Appel 2001; Crary 2003; Klein and Nipkow 2006; Leroy 2006; Lee et al. 2007, etc.), the community remains fragmented, with little synergy between groups and, for newcomers wanting to join the game, a perplexing array of choices between different logics, proof assistants, and representation techniques.

To stimulate progress on this problem, Aydemir et al. (2005) proposed the POPLMARK challenge, a set of tasks designed to stress many of the critical issues in formalizing programming language metatheory. They laid out three criteria for evaluating proposed formalization techniques: *infrastructure overhead*—formalization overheads, such as additional operations and their associated proof obligations, should not be prohibitive for large developments; *transparency*—formal definitions and theorems should not depart radically from the usual informal conventions familiar to a technical audience; and *cost of entry*—the infrastructure should be usable by someone who is not an expert in theorem prover technology.

In this paper, we propose a new style for formalizing programming language metatheory that performs well on all these criteria. The style builds upon two key ingredients: locally nameless representation of syntax involving binders and a cofinite style of quantification for introducing free names in rules dealing with binders.

In locally nameless representation, bound variables are represented by de Bruijn indices while free variables are represented by names. This mixed representation combines the benefits of both approaches, avoiding difficulties associated with alpha-conversion by ensuring that each alpha-equivalence class of terms has a unique representation, while supporting formal reasoning that closely follows informal practice.

The second ingredient is the use of a cofinite quantification of free names introduced by rules dealing with binders, instead of the standard “exists-fresh” quantification—these styles are illustrated by the following variants of the typing rule for abstraction in the

simply-typed lambda calculus ($\lambda_{\rightarrow}$):

$$\frac{\text{EXISTS-FRESH} \quad x \notin \text{FV}(t) \quad E, x:S \vdash t^x : T}{E \vdash \text{abs } t : S \rightarrow T} \quad \frac{\text{COFINITE} \quad \forall x \notin L. E, x:S \vdash t^x : T}{E \vdash \text{abs } t : S \rightarrow T}$$

Here, t^x is the body of abstraction ($\text{abs } t$) *opened up* with a free variable named x . In the second rule, L is some finite set of names that is chosen when we apply the rule. Just as EXISTS-FRESH is applicable if there exists some $x \notin \text{FV}(t)$ such that $E, x:S \vdash t^x : T$, so COFINITE is applicable if there exists some L such that $\forall x \notin L. E, x:S \vdash t^x : T$. The induction hypothesis for rule EXISTS-FRESH holds for one particular name x , which (notionally) comes from the derivation being eliminated by the induction; if that x is not fresh enough, the only way to proceed is to reason about equivariance of the typing judgement. The cofinite rule, on the other hand, asserts from the start that for $E \vdash \text{abs } t : S \rightarrow T$ to hold there must be infinitely many x s with $E, x:S \vdash t^x : T$; surely one of them is fresh enough since only finitely many names could have been used so far. In rare cases, renaming is necessary for proofs, but it is straightforwardly supported by the cofinite presentation; the required lemmas may be easily derived from standard properties such as substitution and weakening, which are themselves derivable using only the cofinite definition. These renaming lemmas also provide the core arguments for the equivalence between the exists-fresh and cofinite presentations of the relations.

Neither of these ingredients is entirely new. The locally nameless representation dates back to the introduction of de Bruijn indices. Several strengthened induction principles that avoid manual renaming have been proposed (Gordon 1994; McKinna and Pollack 1993, 1999; Urban et al. 2007b). The idea of reasoning about the freshness of names by considering all but those in some finite set is at the heart of nominal logic (Pitts 2003) and also appears in definitions of alpha-equivalence by Krivine (1990) and Ford and Mason (2001). Our contribution lies in the precise way that we combine and apply these ingredients. In particular, the cofinitely quantified typing rule for abstraction (presented earlier), which reflects both key aspects of our design, does not appear in prior work. The observation that we can use this formulation to derive renaming lemmas with little infrastructure is a key point, of both theoretical and practical interest.

To demonstrate the comprehensiveness of our approach we have developed several significant examples in the Coq proof assistant (Coq Development Team 2007): proofs of type soundness for $\lambda_{\rightarrow}$, System F_{ω} (the core of the POPLMARK challenge—parts 1A and 2A), and ML extended with references, exceptions, datatypes, recursion and patterns, as well as a proof of subject reduction for the Calculus of Constructions (including a proof of the Church-Rosser property). These developments require no external tool support and could be reproduced with another general purpose theorem prover. Little infrastructure is required between the statement of the system formalized and the core of the formal reasoning. Furthermore, our formal reasoning faithfully follows informal practice, preserving the skeletons of proofs and the kinds of arguments involved, unpolluted by details about alpha-conversion or freshness. We have carefully organized and commented these proof scripts and made them freely available as starting points for comparison and extension.¹

The paper is organized as follows. Section 2 surveys possible approaches to representing binding structures, locates the locally nameless approach in this space, and sketches its history. Section 3 presents a complete specification of $\lambda_{\rightarrow}$ using locally nameless representation and the standard exists-fresh presentation of the se-

mantics. Section 4, the technical meat of the paper, describes the problem faced by the exists-fresh specification, shows how cofinite quantification solves this problem, and discusses why our solution is significantly better than other concrete approaches. Finally, Section 5 gives a practical overview of our larger formalizations. This section also quantitatively compares our developments to each other and to other publicly available Coq solutions to the POPLMARK challenge.

2. A Survey of Binding Representations

There are two main categories of approaches to representing and reasoning about syntax with binding, depending on whether variables and binders are represented as concrete nodes in first-order algebraic structures or whether these aspects are “lifted to the metalanguage” by representing the bodies of binding constructs as metalanguage functions. In each category there are many ingenious variations, and we do not attempt to be exhaustive in this survey. Outside the main categories, there is a plethora of hybrid representations, multi-level representations, . . . , which we do not consider.

2.1 Concrete Approaches

With concrete (or first-order) approaches, variables are typically encoded using names or natural numbers. Capture-avoiding substitution must then be defined explicitly as a function on terms, and in the case where bound variables are named, alpha-equivalence must also be defined explicitly. Concrete approaches can be subdivided roughly in three categories: those using names to represent variables, those using de Bruijn indices, and those distinguishing bound and free variables (which may use either names or indices for each kind of variable).

The most standard representation on paper uses names to represent variables, generally quotienting raw expressions by some notion of alpha-equivalence. Although used since the first days of the lambda calculus, this representation has not proved convenient for formal developments. For example, naive substitution is not structurally recursive in this setting but requires well-founded recursion. Nevertheless, the Church-Rosser property in pure lambda calculus has been proved using named representations by Homeier (2001) using HOL, Ford and Mason (2001) using PVS, and Vestergaard and Brotherston (2003) using Isabelle/HOL (unusually, Vestergaard and Brotherston reason about unquotiented raw terms). Other attempts to tame this representation include work by Stoughton (1988), where simultaneous substitution is defined structurally and puts terms into a canonical alpha-normal form, and by Hendriks and van Oostrom (2003).

To overcome the clumsiness of explicit alpha-conversion or quotienting, the first large formalizations of languages with binders were based on de Bruijn indices (1972) rather than names. In this approach, each alpha-equivalence class of lambda-terms has a unique syntactic representation in which variables are encoded using natural numbers giving the depth of the variable relative to its binder. Indices greater than the number of enclosing local binders may indicate a position in some enclosing context such as a typing context. There is abundant evidence of the effectiveness of this representation, from proofs of the Church-Rosser property for lambda calculus—for example, by Shankar (1988) using the Boyer-Moore prover, by Huet (1994) in Coq, by Rasmussen (1995) in Isabelle/ZF, and by Nipkow (2001) in Isabelle/HOL—to harder results such as strong normalization for System F in LEGO (Altenkirch 1993) and formalizing Coq in Coq (Barras and Werner 1997). In de Bruijn representation, the treatment of bound variables incurs minor technical annoyances (lifting over binders, etc.), while the treatment of free variables (e.g., variables bound by a typing context) requires reasoning far from natural informal presentations. For example, consider the statement of a lemma that says thinning the typing

¹<http://arthur.chargueraud.org/research/2007/binders/>.

context (weakening and permuting) preserves a judgement: indices occurring in the judgement must be updated to match the permutation of their binding points in the context.

Nominal logic (Pitts 2003) provides another way to address the problem of alpha-conversion inherent in the named representation. Urban (2007) and Urban et al. (2007a,b) have adapted these ideas to standard higher-order logic, and are developing a *nominal package* in Isabelle/HOL that provides facilities for formal reasoning with binders over names. The package helps by automating language definitions that equate alpha-equivalent terms, by providing lemmas about renaming and freshness of names, by deriving strengthened induction principles that provide name freshness facts for relations over these languages, and by supporting recursion over these languages. Some syntactic restrictions must be observed to use this package. See Section 4.6 for further discussion of the nominal package.

Besides these “homogeneous” concrete approaches—where there is a single notion of “variable,” either a name or a de Bruijn index—there is a third category of concrete representations that uses two distinct syntactic classes, called variables (for locally bound variables) and parameters (for free, or globally bound variables). The idea that variables and parameters should be distinguished goes back at least to Gentzen (1969) and Prawitz (1965). Several combinations can be considered, using names, indices, or levels to represent bound and free variables; of these, two have been used in the literature.

The first is the *locally named* representation, which uses distinct species of names to represent variables and parameters respectively. McKinna and Pollack introduced this technique to formalize Pure Type Systems (PTS) (1993; 1999), following a suggestion by Coquand (1991). This representation avoids the difficulties of reasoning about capture-avoiding substitutions: since variables and parameters are syntactically distinguished, no parameter can ever be captured by a variable binder during substitution. McKinna and Pollack also introduced a technique for handling the requirement of choosing fresh global variables that often occurs in reasoning about binding (weakening lemmas are a prototypical example of the problem). By careful choice of definitions, they avoid the need to reason about alpha-equivalence of bound variables throughout their large formal body of PTS metatheory. With these techniques, reasoning about lambda calculus and type theory is straightforward, if heavy.

Nonetheless, the use of names for bound variables is not a perfect fit to the intuitive notion of binding: eventually one needs to reason about alpha-conversion. For example, parallel reduction has the diamond property concretely in locally named representation, but beta reduction has this property only up to alpha-conversion (Pollack 1994b, Section 3.3.6). So Pollack (1994a) suggested that the McKinna–Pollack approach to reasoning with two species of variables also works well with a representation that uses names for parameters and de Bruijn indices for bound variables. This *locally nameless* representation, in which alpha-equivalence classes have canonical representation, was already mentioned in the conclusion of de Bruijn’s famous paper (1972). It had been used for implementation in Huet’s *Constructive Engine* (1989), and later in the implementations of the Coq, LEGO (Luo and Pollack 1992), HOL 4 (Norris and Slind 2007), Isabelle, and EPIGRAM (McBride and McKinna 2004) proof assistants.

In the context of formal proofs, Gordon (1994) appears to be the first to have used *locally nameless* representation. Rather than reason directly with locally nameless terms, he builds a representation of named lambda terms on top of locally nameless terms and demonstrates how the named terms may be used as a basis for representing syntax. Using this technique, Gordon formalized Abramsky’s lazy lambda calculus. Later work by Gordon and Melham (1996) also uses locally nameless terms as a model for an ab-

stract “axiomatic” representation of named terms. Pollack (2006) has more recently emphasized the benefits of locally nameless representation in the context of the POPLMARK challenge. Locally nameless representation with (variants of) McKinna–Pollack style reasoning has been used by several researchers (with several proof tools) for solutions to the POPLMARK challenge, first by Leroy (2007) and later by Chlipala (2006) and Ricciotti (2007).

2.2 Higher-Order Approaches

Higher-order representations use the function space of the meta-logic to encode binding in an object language, allowing issues like capture-avoidance and alpha-equivalence to be handled once and for all by the meta-logic, rather than facing them anew for each object language. There is a bewildering variety of higher-order approaches, which we survey only superficially.

In *higher-order abstract syntax* (HOAS) (Harper et al. 1993; Pfenning and Elliot 1988), the introduction form for lambda-abstractions has type $(\text{term} \rightarrow \text{term}) \rightarrow \text{term}$, i.e., the lambda constructor packages a function of type $(\text{term} \rightarrow \text{term})$, which should be thought of as the function substituting its argument into the body of the lambda. Not only are alpha-equivalent terms equal, but substitution is handled by the meta-language.

HOAS was introduced by Church and developed by de Bruijn, Martin-Löf, Plotkin, Pfenning, and their co-workers. The modern form, using dependent types and *hypothetical* and *schematic* judgements as well as HOAS term representation, appeared in the Edinburgh Logical Framework (LF) (Harper et al. 1993). By design, LF is a weak type theory to be used as a metalanguage, not supporting inductive definition or primitive recursion. It allows *faithful* representation of object languages, where LF itself adds no mathematical strength to the object language representation. The use of LF methodology for metatheory (as opposed to just representation) has been highly developed by Pfenning and his co-workers; the implementation of this approach is the Twelf system (Pfenning and Schürmann 1999), which is widely used and very successful for formalizing the metatheory of a wide variety of programming languages (Ashley-Rollman et al. 2005; Lee et al. 2007). The approach continues to be developed foundationally as well as in practice (Harper and Licata 2007).

A second main stream of work, *weak* higher-order encodings, gives a type such as $(\text{name} \rightarrow \text{term}) \rightarrow \text{term}$ to the lambda constructor of terms.² This constructor packages a function of type $(\text{name} \rightarrow \text{term})$, which should be thought of as a function that takes a name and returns the instance of the lambda’s body instantiated with that name. Here, alpha-equivalent terms are equated, but substitution must be defined as a relation between terms. Despeyroux et al. (1995) demonstrate an early use of this approach. The approach has been most developed by Honsell, Miculan, and their co-workers as the Theory of Contexts (Honsell et al. 2002; Bucalo et al. 2006).

2.3 Why Locally Nameless?

We have chosen to concentrate on concrete approaches to represent binding structure. Despite the attractiveness of higher-order approaches and their demonstrated success in many large developments, they are, at present, not completely general-purpose. Twelf is currently the only industrial-strength proof assistant for devel-

²The technical advantage of the weak variant is that the type of terms can be given as an inductive datatype. For strong HOAS, due to the negative occurrence of term in $(\text{term} \rightarrow \text{term}) \rightarrow \text{term}$, there is no inductive datatype representing the set of terms. In the weak version, on the other hand, the type name replaces the negative occurrence of term , so there is a datatype of weak higher-order terms in, e.g., Coq and HOL. Nonetheless technical problems arise unless name is a sufficiently abstract type (Honsell et al. 2002).

oping metatheory based on HOAS representations, and it is (in its current incarnation) strictly limited in its proof-theoretic strength. The Theory of Contexts is also interesting, but hasn't yet been taken up beyond its originators. These are the subject of active research and may become more generally applicable in the future.

Among the concrete approaches, locally nameless representation offers the best combination of features. Because free variables are represented using names, lemmas are stated and proofs structured just as they are with the standard representation using names, making locally nameless developments intuitive. In particular, the shifting operations that clutter both statements and proofs in a pure de Bruijn setting are avoided. Because bound variables are represented using indices, each alpha-equivalence class of lambda terms has a unique representation, thus avoiding the difficulties associated with alpha-equivalence. Finally, because bound and free variables are syntactically distinguished there is no issue of variable capture, and substitution can be defined by simple structural recursion.

These observations motivate the choice of the locally nameless representation. The next section presents this style in detail.

3. Locally Nameless Representation

In this section, we describe locally nameless representation, using $\lambda_{\rightarrow}$ as a running example. All the material presented is formal, in the sense that we have implemented it in Coq.³ We write $\Rightarrow$ for logical implication.

3.1 Definitions

The fundamental idea of the locally nameless approach is to distinguish bound from free variables, using de Bruijn indices for the former and names for the latter. Figure 1 presents a locally nameless specification of $\lambda_{\rightarrow}$. An occurrence of a bound variable is represented as $(\text{bvar } i)$ where i is a natural number index denoting the number of enclosing lambda abstractions that must be traversed before reaching the abstraction binding that variable. Abstractions do not bind names; for example, $(\text{abs } (\text{bvar } 0))$ is the identity function.

In its raw form, this representation is not isomorphic to ordinary named lambda terms because a bound variable may not resolve to any binder. For example, $(\text{abs } (\text{bvar } 2))$ does not represent any lambda-term. We therefore define a predicate, $(\text{term } t)$, on *pre-terms* (the syntactic objects defined by the grammar in Figure 1) that selects the *locally closed* pre-terms—or just *terms*—in which all indices resolve to binders.

The key operation on pre-terms is *opening* an abstraction by instantiating a bound variable with a term. If $(\text{abs } t)$ is an abstraction and u another term, then the result of opening the abstraction with u , written t^u , is formed by replacing in the body t all bound variables that point to the outermost lambda with u . Figure 1 gives the structurally recursive definition of the open function; the idea is to explore the term, keeping track of the current number of binders passed through.⁴ One use of the open operation is in the conclusion of the beta reduction rule, RED-BETA, near the bottom of the figure.

The open operation is also frequently applied to a free variable; we call this specialized version *variable opening*, which we write with abuse of notation as (t^x) instead of $(t^{\text{fvar } x})$. Variable opening is used when passing through a binder to turn the bound variable into a free variable. In a named representation, we may have our hands on an abstraction $(\text{abs } x \ t)$ and wish to talk about its body t . With the locally nameless representation, the abstraction would

Syntax:

$$\begin{aligned} S, T &\equiv A \mid T_1 \rightarrow T_2 \\ t, u, w &\equiv \text{bvar } i \mid \text{fvar } x \mid \text{app } t_1 \ t_2 \mid \text{abs } t \\ E, F, G &\equiv \emptyset \mid E, x : T \end{aligned}$$

Open: $t^u \equiv \{0 \rightarrow u\} t$, with

$$\begin{aligned} \{k \rightarrow u\} (\text{bvar } k) &= u \\ \{k \rightarrow u\} (\text{bvar } i) &= \text{bvar } i \quad \text{when } i \neq k \\ \{k \rightarrow u\} (\text{fvar } x) &= \text{fvar } x \\ \{k \rightarrow u\} (\text{app } t_1 \ t_2) &= \text{app } (\{k \rightarrow u\} t_1) (\{k \rightarrow u\} t_2) \\ \{k \rightarrow u\} (\text{abs } t) &= \text{abs } (\{(k+1) \rightarrow u\} t) \end{aligned}$$

Free variables:

$$\begin{aligned} \text{FV}(\text{bvar } i) &= \emptyset \\ \text{FV}(\text{fvar } x) &= \{x\} \\ \text{FV}(\text{app } t_1 \ t_2) &= \text{FV}(t_1) \cup \text{FV}(t_2) \\ \text{FV}(\text{abs } t) &= \text{FV}(t) \end{aligned}$$

Locally closed terms:

$$\begin{aligned} \frac{}{\text{term } (\text{fvar } x)} \text{TERM-VAR} \quad & \frac{\text{term } t_1 \quad \text{term } t_2}{\text{term } (\text{app } t_1 \ t_2)} \text{TERM-APP} \\ \frac{x \notin \text{FV}(t) \quad \text{term } (t^x)}{\text{term } (\text{abs } t)} \text{TERM-ABS} \end{aligned}$$

Well-formed environments (no duplicate names):

$$\frac{}{\text{ok } \emptyset} \text{OK-NIL} \quad \frac{\text{ok } E \quad x \notin \text{dom}(E)}{\text{ok } (E, x : T)} \text{OK-CONS}$$

Typing:

$$\begin{aligned} \frac{\text{ok } E \quad (x : T) \in E}{E \vdash \text{fvar } x : T} \text{TYPING-VAR} \\ \frac{E \vdash t_1 : S \rightarrow T \quad E \vdash t_2 : S}{E \vdash \text{app } t_1 \ t_2 : T} \text{TYPING-APP} \\ \frac{x \notin \text{FV}(t) \quad E, x : T_1 \vdash t^x : T_2}{E \vdash \text{abs } t : T_1 \rightarrow T_2} \text{TYPING-ABS} \end{aligned}$$

Call-by-value evaluation:

$$\begin{aligned} \frac{\text{term } (\text{abs } t)}{\text{value } (\text{abs } t)} \text{VALUE-ABS} \\ \frac{\text{term } (\text{abs } t) \quad \text{value } u}{\text{app } (\text{abs } t) \ u \mapsto t^u} \text{RED-BETA} \\ \frac{t_1 \mapsto t'_1 \quad \text{term } t_2}{\text{app } t_1 \ t_2 \mapsto \text{app } t'_1 \ t_2} \text{RED-APP-1} \\ \frac{\text{value } t_1 \quad t_2 \mapsto t'_2}{\text{app } t_1 \ t_2 \mapsto \text{app } t_1 \ t'_2} \text{RED-APP-2} \end{aligned}$$

Type soundness lemmas (preservation and progress):

$$\begin{aligned} E \vdash t : T &\Rightarrow t \mapsto t' \Rightarrow E \vdash t' : T \\ \emptyset \vdash t : T &\Rightarrow (\text{value } t \vee \exists t', t \mapsto t') \end{aligned}$$

Figure 1. Locally nameless presentation of $\lambda_{\rightarrow}$.

³The corresponding development is contained in files `STLC.Core.*.v`.

⁴Experts will note that this definition of open works correctly only when zero is the only unbound index. The definition can be generalized to allow multiple binders to be opened simultaneously as explained in Section 5 in the paragraph about our ML development.

be written $(\text{abs } t)$; to talk about its body as a term (rather than a pre-term with an unbound index), we open t with some fresh name x . A typical example of variable opening appears in the typing rule for abstractions, TYPING-ABS.

In variable opening, the chosen variable should not already appear free inside the term. Otherwise, the operation would be a form of variable capture: the name being given to the index, a bound variable, would be the same as a free variable. Therefore, the rule TYPING-ABS also includes the premise $x \notin \text{FV}(t)$,⁵ which uses the *free variable* function to collect the set of free variables in a pre-term. Like open, this function is defined by simple structural recursion and need not worry about confusing free and bound variables—all variable names in the term are free, since bound variables are represented by indices.

To state the typing judgement of $\lambda_{\rightarrow}$, we need to formalize typing environments. In Figure 1, environments are concretely represented as association lists, for which we assume some standard operations. The *concatenation* of two environments is written (E, F) . The *lookup* of a variable in an environment is expressed by the predicate $(x:T) \in E$, which holds when (x, T) is the rightmost pair of E whose first component is x . Finally, a variable is *fresh* for an environment, written $x \notin \text{dom}(E)$, when that variable does not appear in the first component of any pair in E . In principle, the association list representation allows multiple assumptions for the same variable, but our typing rules ensure that environments participating in derivable judgements are well formed (ok) in the sense that a given variable is bound at most once in a given environment.

The final portion of Figure 1 completes the definition of $\lambda_{\rightarrow}$, by stating the properties we wish to prove: the standard preservation and progress lemmas. The figure, as a whole, forms our “trusted base” of definitions: one must believe in the correctness of everything stated in the figure in order to believe that a formal proof of type soundness is really about one’s informal notion of $\lambda_{\rightarrow}$. We discuss this issue in additional detail in Section 3.4.

3.2 Substitution

The definition of $\lambda_{\rightarrow}$ depends only on the open and FV operations. In proofs, however, the operation of *free variable substitution*, replacing free names with terms, is also needed. Written $[x \rightarrow u] t$, it is defined by structural recursion in Figure 2. As in the definition of open, there is no arbitrary choice of name for the bound variable in the abstraction case, so we need not worry that this choice will affect the result. Therefore, the behavior of substitution is natural and easy to reason about.

The properties of the substitution function, shown in the bottom portion of Figure 2, are fundamental when working with a locally nameless representation. Lemma `subst_fresh` states that substitution for an unused variable has no effect, and lemma `subst_open` states that substitution distributes over the opening operation. The remaining two lemmas are immediate corollaries of the former two. Lemma `subst_open_var` permutes a substitution with variable opening, and lemma `subst_intro` decomposes an open operation into a variable opening operation and a substitution. Examples of when these lemmas are used can be found in the proof of the substitution lemma for $\lambda_{\rightarrow}$ (Section 4.1) and in the proof of a renaming lemma (Section 4.2).

We occasionally need a related operation—intuitively, a reciprocal of variable opening—that abstracts the name x from the term t , replacing it with a bound variable. We call this operation *variable closing* and define it in Figure 3, again by structural recursion.

⁵ The rule TYPING-ABS could also explicitly include a premise to ensure that x does not appear in the domain of the environment E . However, that restriction is already implied by the second hypothesis.

Substitution of a term for a free name:

$$\begin{aligned} [z \rightarrow u] (\text{bvar } i) &= \text{bvar } i \\ [z \rightarrow u] (\text{fvar } z) &= u \\ [z \rightarrow u] (\text{fvar } x) &= \text{fvar } x \quad \text{when } x \neq z \\ [z \rightarrow u] (\text{app } t_1 \ t_2) &= \text{app } ([z \rightarrow u] t_1) ([z \rightarrow u] t_2) \\ [z \rightarrow u] (\text{abs } t) &= \text{abs } ([z \rightarrow u] t) \end{aligned}$$

Properties of substitution:

$$\begin{aligned} \text{subst_fresh} \quad x \notin \text{FV}(t) &\Rightarrow [x \rightarrow u] t = t \\ \text{subst_open} \quad [x \rightarrow u] (t^w) &= ([x \rightarrow u] t)^{([x \rightarrow u] w)} \\ &\quad \text{when term } u \\ \text{subst_open_var} \quad [x \rightarrow u] (t^y) &= ([x \rightarrow u] t)^y \\ &\quad \text{when } x \neq y \text{ and term } u \\ \text{subst_intro} \quad [x \rightarrow u] (t^x) &= t^u \\ &\quad \text{when } x \notin \text{FV}(t) \end{aligned}$$

Figure 2. Substitution

Close: $\backslash^x t \equiv \{0 \leftarrow x\} t$, with

$$\begin{aligned} \{k \leftarrow x\} (\text{bvar } i) &= \text{bvar } i \\ \{k \leftarrow x\} (\text{fvar } x) &= \text{bvar } k \\ \{k \leftarrow x\} (\text{fvar } y) &= \text{fvar } y \quad \text{when } x \neq y \\ \{k \leftarrow x\} (\text{app } t_1 \ t_2) &= \text{app } (\{k \leftarrow x\} t_1) (\{k \leftarrow x\} t_2) \\ \{k \leftarrow x\} (\text{abs } t) &= \text{abs } (\{(k+1) \leftarrow x\} t) \end{aligned}$$

Figure 3. Variable Closing

3.3 Working With Local Closure

We use an inductive definition of local closure, given in Figure 1. A bound variable on its own is never locally closed: there is no case for $(\text{bvar } i)$. An abstraction is locally closed if the body of that abstraction, once opened up with a fresh name, is locally closed.⁶ There are alternative ways to define local closure, but we favor this style which is similar to the `VClosed` relation of McKinna and Pollack (1993). As they point out, this definition provides a structural induction principle for locally closed terms:

$$\frac{\begin{aligned} \forall x. P(x) \\ \forall t_1 \ t_2. P(t_1) \Rightarrow P(t_2) \Rightarrow P(\text{app } t_1 \ t_2) \\ \forall x \ t. x \notin \text{FV}(t) \Rightarrow P(t^x) \Rightarrow P(\text{abs } t) \end{aligned}}{\forall t. \text{term } t \Rightarrow P(t)}$$

This definition can be used in many different logics: intensional or extensional, dependent types or simply typed HOL.⁷

Although the relations defining the semantics of $\lambda_{\rightarrow}$ in Figure 1 are formally defined over pre-terms, these relations actually include only locally closed terms and the basic operations preserve local

⁶ The freshness assumption is not actually required for this judgement; we include it to maintain the invariant that a term is never opened with a name that occurs free in that term.

⁷ In extensional higher order logic the (non-empty) predicate term can be made into a *type* of locally closed terms. In intensional dependent type systems (including Coq), this construction isn’t available. However, in such systems we can define the inductive family of pre-terms indexed by the set of bvvars occurring free (e.g. McBride and McKinna (2004); for discussion of inductive families see Dybjer (1994)). Then the section of this family indexed by the empty set of bvvars is exactly the locally closed terms. Improvements to our style may be possible along these lines, but for the moment we stick with the representation given above, for its simplicity of use and uniformity of presentation across different logics.

closure.

$$\begin{array}{lll}
\text{term } (\text{abs } t) \wedge \text{term } u \Rightarrow \text{term } (t^u) & & \\
\text{term } u \wedge \text{term } t \Rightarrow \text{term } ([x \rightarrow u] t) & & \\
E \vdash t : T \Rightarrow \text{term } t & & \\
t \mapsto t' \Rightarrow \text{term } t \wedge \text{term } t' & &
\end{array}$$

We formulate the definitions in this way for two reasons. First, it is required for the adequacy of our formalization: only locally closed terms correspond to alpha-equivalence classes of ordinary named terms. Second, some useful properties of the definitions are true only for locally closed terms.⁸ The typing relation ensures local closure of its subject without any local closure premises, but not every relation of interest is this clean. For example, to ensure this property of small-step evaluation, the definition includes several local closure premises. In general, such premises are only required at the leaves of the definition: e.g., we do not need a premise (term t) if there is already a premise (value t), since the latter implies the former.

Similarly, because our definitions guarantee local closure, we usually do not need to add local closure hypotheses to the statements of lemmas and theorems. For example, the “substitution preserves typing” lemma

$$\begin{array}{lcl}
E, z:S, F \vdash t : T & \Rightarrow & E \vdash u : S \\
E, F \vdash [z \rightarrow u] t : T & &
\end{array}$$

requires none because the typing relation ensures that t , u , and $[z \rightarrow u] t$ are all locally closed. As a result, even though the use of the nameless representation means that we must be careful about local closure, the management of this predicate does not require a significant departure from common informal practice.

3.4 Adequacy

Why should you believe this mechanized, formal representation of $\lambda_{\rightarrow}$ is “correct,” i.e., expresses your understanding of the rigorous, informal system? This is the question of *adequacy* of the representation, set out by Harper et al. (1993) in their context of HOAS. (See also (Harper and Licata 2007) for a more recent take.) First, note that it is an informal question, because it involves the relationship between an informal thing and a formal thing. No matter how much faith you put in Coq, no Coq proof will completely settle this question.

Nonetheless, formal proofs are useful for this question in three ways. First, you may prove that the formalization has the properties you expect from the informal system and does not have unexpected properties. E.g., our Coq representation of $\lambda_{\rightarrow}$ is shown to have Church–Rosser, weakening, subject reduction, etc., some of which is discussed below. (You must also be convinced that our formal statements of these properties are correct.) If you think our rule TYPING-ABS requires a stronger side condition, prove your expected rule is admissible in our system. Second, you may prove some formal relationship with another formalization of the same informal system that you, and other readers, believe to be adequate, e.g., that our formalization is related to a pure de Bruijn representation in some way (but we haven’t done that). Finally, if intensional, impredicative, constructive type theory is just too weird, you may carry out the same representation and its development in some other proof system. We have avoided using special properties of Coq or its logic, and our style could be carried out in HOL, NuPrl, PVS, etc.

Harper et al. (1993) ask that there be a compositional bijection between the informal syntactic entities (terms, formulas, ...) and some collection of entities of the formal representation. (They also note this may be taken to have different meanings: e.g., adequacy

for theorems versus adequacy for derivations.) In our setting, you can carry this out (on paper, not in Coq) with a fairly simple bijection.

4. Cofinite Quantification

Although the definitions in the previous section adequately represent $\lambda_{\rightarrow}$, they yield insufficiently strong induction principles: in many situations we must rename the variable used to open an abstraction to complete a proof. The need to reason about renaming motivates our use of “cofinite quantification” for inductively defined relations.

In this section, we present a complete style of proof development based on cofinite quantification. We first explain cofinite quantification and give a new definition of $\lambda_{\rightarrow}$ using it. Definitions in this style are sufficient, by themselves, to develop significant amounts of metatheory—e.g., to prove weakening and substitution results—since they naturally balance strong induction and inversion principles with useful introduction forms. In a few situations, however, *renaming lemmas* are required to build derivations using the cofinitely quantified rules. We show how these lemmas can be derived directly from weakening and substitution. Finally, to be confident that relations defined with cofinite quantification are adequate encodings of $\lambda_{\rightarrow}$, we use renaming to show that they are equivalent to their counterparts shown in Figure 1.

Having described our complete style of development, we then observe that some significant streamlining is possible, notably in cases where one believes in the adequacy of the cofinite presentation of a relation without requiring a formal proof of equivalence with the “exists-fresh” variant. We conclude the section with comparisons to previous work, arguing that the added value of our approach is significant in each case.

4.1 Stronger Induction Principles

The definitions of the local closure and typing relations in Figure 1 use an “exists-fresh” style of quantification. When an abstraction is opened in the premise of a rule (TERM-ABS and TYPING-ABS), the name used in the premise is required to be fresh for the body of the abstraction. Thus, to build a derivation using one of these rules, it suffices to show that there exists a single sufficiently fresh name for which the premise holds. The premises of rules written in this way are easy to satisfy, making them ideal introduction forms.

However, such rules give rise to weak elimination forms (inversion and induction principles). The fact that the premises of a rule defined using exists-fresh quantification need only hold for one particular name means that the corresponding induction hypothesis also holds for one particular name. That name will only be as fresh as the premise of the rule requires, which may not be enough.

For example, consider the proof of the weakening lemma for the typing relation, which allows one to insert additional typing assumptions anywhere into a typing environment:

$$\begin{array}{lcl}
\text{typing_weaken} : & & \\
E, G \vdash t : T & \Rightarrow & \text{ok } (E, F, G) \Rightarrow \\
E, F, G \vdash t : T & &
\end{array}$$

When proving this by induction on the given typing derivation, in the case for TYPING-ABS we are given a derivation ending with

$$\frac{x \notin \text{FV}(t_1) \quad E, G, x:T_1 \vdash t_1^x : T_2}{E, G \vdash \text{abs } t_1 : T_1 \rightarrow T_2}$$

and an induction hypothesis

$$\text{ok } (E, F, G, x:T_1) \Rightarrow (E, F, G, x:T_1 \vdash t_1^x : T_2)$$

and we must derive

$$E, F, G \vdash \text{abs } t_1 : T_1 \rightarrow T_2$$

⁸These include both low-level technical properties such as `subst.open` and more interesting properties such as reflexivity of subtyping in System $F_{<}$.

$$\begin{array}{c}
\overline{\text{term}_c(\text{fvar } x)} \quad \text{C-TERM-VAR} \\
\\
\frac{\text{term}_c t_1 \quad \text{term}_c t_2}{\text{term}_c(\text{app } t_1 t_2)} \quad \text{C-TERM-APP} \\
\\
\frac{\forall x \notin L. \text{term}_c(t^x)}{\text{term}_c(\text{abs } t)} \quad \text{C-TERM-ABS} \\
\\
\frac{\text{ok } E \quad (x:T) \in E}{E \vdash_c \text{fvar } x : T} \quad \text{C-TYPING-VAR} \\
\\
\frac{E \vdash_c t_1 : S \rightarrow T \quad E \vdash_c t_2 : S}{E \vdash_c \text{app } t_1 t_2 : T} \quad \text{C-TYPING-APP} \\
\\
\frac{\forall x \notin L. (E, x:T_1 \vdash_c t^x : T_2)}{E \vdash_c \text{abs } t : T_1 \rightarrow T_2} \quad \text{C-TYPING-ABS}
\end{array}$$

Figure 4. Local closure and typing using cofinite quantification

for an arbitrary F . By TYPING-ABS, it suffices to derive

$$E, F, G, y:T_1 \vdash t_1^y : T_2,$$

for some name $y \notin \text{FV}(t_1)$. Intuitively, we want to use $y = x$, since the induction hypothesis applies to x . However, the given typing derivation shows that t_1^x is well typed in the environment $(E, G, x:T_1)$, so we can only derive that x is fresh for (E, G) . We need to type t_1^x in the extended environment $(E, F, G, x:T_1)$, but x is not guaranteed to be sufficiently fresh for that environment as it could appear in F .

The particular name x appearing in the induction hypothesis is (notionally) the name that occurs in that position of the particular derivation of $E, G \vdash t : T$ being eliminated in the proof. In an informal proof, we avoid the problem by assuming this derivation uses sufficiently fresh names. (Depending on our informal view, this is usually justified by appeal to the Barendregt Variable Convention (1984), which allows us to assume that the name of the bound variable in the abstraction case is sufficiently fresh.)

Our solution is to strengthen the induction principle for the typing relation by changing its definition to the one at the bottom of Figure 4. This definition differs only in the rule for abstractions, C-TYPING-ABS. This new rule intuitively states: “to show that $(\text{abs } t)$ is well typed it suffices to choose some finite set of names, L , and show that t^x is well-typed for every x not in L .” We say that this definition of the typing relation uses “cofinite quantification” since we must show that the premise of C-TYPING-ABS (“C” for “cofinite”) holds for all but finitely many names. The set of names L may be different for each use of the rule. To similarly strengthen the structural induction principle for terms, we redefine term in the cofinite style at the top of Figure 4, and replace all references to term in the reduction relation with this one. Although this new definition of $\lambda \rightarrow$ differs from the one given in Figure 1, it is equivalent, as we show in Section 4.3.

Using cofinite quantification in the rules provides a stronger induction principle because, in the cases that open abstractions, the induction hypotheses will hold for all names except those in some finite set L , rather than just for a single name. On the other hand, a rule defined using cofinite quantification is nearly as easy to use for introduction as the exists-fresh version because L can include the finitely many names that could potentially lead to clashes.

Consider again the proof of the weakening lemma—this time, for the cofinite typing relation $\vdash_c$. In the case for C-TYPING-ABS,

we are given the assumption

$$\forall x \notin L'. (E, G, x:T_1 \vdash_c t_1^x : T_2)$$

for some finite set of names L' , and an induction hypothesis

$$\forall x \notin L'. \text{ok } (E, F, G, x:T_1) \Rightarrow E, F, G, x:T_1 \vdash_c t_1^x : T_2,$$

and we must derive

$$E, F, G \vdash_c \text{abs } t_1 : T_1 \rightarrow T_2.$$

By C-TYPING-ABS, it suffices to find an L for which we can show

$$\forall x \notin L. (E, F, G, x:T_1 \vdash_c t_1^x : T_2).$$

Choosing $L = L' \cup \text{dom}(F)$ does the trick: any $x \notin L$ is also not in L' , so the induction hypothesis directly gives us the result.

As an additional example of the usefulness of the stronger induction hypothesis, we prove the standard “substitution preserves typing” lemma for the cofinite typing relation:

$$\begin{array}{l}
\text{typing_subst :} \\
E, z:S, F \vdash_c t : T \Rightarrow E \vdash_c u : S \Rightarrow \\
E, F \vdash_c [z \rightarrow u] t : T.
\end{array}$$

The proof of this lemma also demonstrates the ease of working with substitution in a locally nameless setting. It proceeds by induction on the given typing derivation for t . In the case for C-TYPING-ABS, we have a derivation that ends in

$$\frac{\forall x \notin L'. (E, z:S, F, x:T_1 \vdash_c t_1^x : T_2)}{E, z:S, F \vdash_c \text{abs } t_1 : T_1 \rightarrow T_2}$$

for some finite set of names L' , and an induction hypothesis

$$\forall x \notin L'. (E, F, x:T_1 \vdash_c [z \rightarrow u](t_1^x) : T_2),$$

and we must derive

$$E, F \vdash_c [z \rightarrow u](\text{abs } t_1) : T_1 \rightarrow T_2.$$

By the definition of substitution, we simplify the substitution to $\text{abs } ([z \rightarrow u] t_1)$. Then to apply C-TYPING-ABS, we must find an L for which we can show

$$\forall x \notin L. (E, F, x:T_1 \vdash_c ([z \rightarrow u] t_1)^x : T_2).$$

We use `subst_open_var` to write $([z \rightarrow u] t_1)^x$ as $[z \rightarrow u](t_1^x)$ (to match the induction hypothesis). To do this rewriting, we need x and z to be distinct. Since the induction hypothesis holds only for $x \notin L'$, we choose $L = L' \cup \{z\}$.

In this proof and in the proof of weakening, we carefully chose L when applying C-TYPING-ABS to create a new typing derivation. In practice, we use a tactic to instantiate L with the set of all names appearing in the current proof context: intuitively, showing that a judgement holds for fewer names is easier than showing that it holds for more.

4.2 Renaming

There are a few situations in which we need to apply a derivation rule with a cofinitely quantified premise but only know that this premise holds for one particular fresh name. Alternatively, knowing that an abstraction type checks, we may wish to use inversion to show that the body of an abstraction type checks with a particular name, but we do not know that that particular name is excluded from the set L used with that invocation of C-TYPING-ABS. In such situations, we need to explicitly invoke the fact that our judgments are equivariant through a *renaming lemma*. We present now an example of such a situation and explain how to prove the renaming lemma by reusing the corresponding substitution lemma.

Suppose that we are trying to show that type checking is decidable.^{9 10}

typing_decidable :

$$\text{term } t \Rightarrow \text{ok } E \Rightarrow (E \vdash_c t : T) \vee \neg(E \vdash_c t : T)$$

We prove this result by induction on term t . Consider the case for abstractions (where $t = \text{abs } t_1$) with an arbitrary environment E and an arrow type (where $T = T_1 \rightarrow T_2$). Here we are given an induction hypothesis that states that type checking the body is decidable (for some finite set L').

$$\begin{aligned} \forall x \notin L'. \forall E' T', \\ \text{ok } E' \Rightarrow (E' \vdash_c t_1^x : T') \vee \neg(E' \vdash_c t_1^x : T') \end{aligned}$$

At this point we must determine whether the body type checks to know whether the abstraction should type check. So we pick an arbitrary fresh variable, $x \notin L' \cup \text{FV}(t_1) \cup \text{dom}(E)$, and consider the two cases that arise from instantiating E' with $(E, x:T_1)$ and T' with T_2 .¹¹ If the body does type check, we have a derivation of

$$E, x:T_1 \vdash_c t_1^x : T_2.$$

However, to show that the abstraction type checks, using rule C-TYPING-ABS, we must show that there is some L such that

$$\forall y \notin L. E, y:T_1 \vdash_c t_1^y : T_2.$$

Although x was arbitrary, we do not know whether the particular variable that we chose influenced whether the body type checked. Therefore, we complete the proof with the help of a renaming lemma that asserts that typing judgments are stable under certain renamings:¹²

typing_rename :

$$\begin{aligned} x \notin (\text{dom}(E) \cup \text{FV}(t_1)) \Rightarrow E, x:T_1 \vdash_c t_1^x : T_2 \Rightarrow \\ \forall y \notin (\text{dom}(E) \cup \text{FV}(t_1)). E, y:T_1 \vdash_c t_1^y : T_2 \end{aligned}$$

This renaming lemma allows us to go from a typing derivation for a single variable x , to a typing derivation for a cofinite set of variables. Using this lemma and choosing $L' = \text{dom}(E) \cup \text{FV}(t_1)$ lets us complete the decidability proof.

The renaming lemma is simply a consequence of properties (substitution and weakening) that we have already proven about the cofinite typing judgement. Importantly, showing these properties does not require renaming. We can derive typing_rename as follows. If $x = y$, then there is nothing to do. Otherwise, we rewrite the conclusion using subst_intro as

$$E, y:T_1 \vdash_c [x \rightarrow \text{fvar } y](t_1^x) : T_2.$$

By the substitution lemma, typing_subst, it suffices to show that

1. $E, y:T_1 \vdash_c \text{fvar } y : T_1$ and
2. $E, y:T_1, x:T_1 \vdash_c t_1^x : T_2$.

⁹ Coq is a constructive logic, so a proof of this lemma is an algorithm that produces either a typing derivation, or a proof that one isn't possible.

¹⁰ The simplicity of the example requires a slight change to the formalization of $\lambda_{\rightarrow}$, to annotate applications (app) with the argument type. This change permits the proof to go through in the application case, but is otherwise independent of binding.

¹¹ The variable x must be fresh for L' for the induction hypothesis to apply and must be fresh for $\text{dom}(E)$ so that $\text{ok } (E, x:T_1)$. We also include $\text{FV}(t_1)$ to maintain the invariant that terms are only opened with fresh variables.

¹² While the freshness condition on $y \notin \text{dom}(E)$ ensures that $(E, y:T_1)$ is well-formed, the freshness condition $x \notin \text{dom}(E)$ could be derived with some work from the second hypothesis; nevertheless we choose to stick to a simpler and symmetric presentation.

To show (1), we use C-TYPING-VAR. To show (2), we use the weakening lemma, by which it suffices to show that $E, x:T_1 \vdash_c t^x : T_2$. This result follows from the assumptions of typing_rename.

In our experience, renaming lemmas are rarely needed. Aside from typing_decidable, the only situations in our developments where they are required are in the proofs of confluence for beta reduction in the lambda calculus and the Calculus of Constructions. In these proofs, the conclusion of the theorem is an existential statement: “there exists a common reduct...” The cases which require renaming involve rules with binding, where we have an induction hypothesis of the form “for all x not in L' , there exists t such that ...” which we need to use to prove a statement of the form “there exists t such that for all y not in $L, \dots$ ” The induction hypothesis is strictly weaker than the statement we are trying to prove. For each name x , it gives us an object t for which some relation holds. We must use the renaming lemma for that relation to show that the relation holds for all sufficiently fresh names y .

4.3 Equivalence of Exists-Fresh and Cofinite Definitions

At this point we have defined two sets of rules for $\lambda_{\rightarrow}$: one consisting of the original exists-fresh definitions, which corresponds directly to the “paper presentation,” and one that uses cofinite quantification, which yields stronger induction principles. To be confident that the cofinite definitions are adequate, we verify (formally) that they define the same relations as the exists-fresh variants.¹³ Note that these proofs of equivalence are straightforward consequences of renaming lemmas.

Let us take the typing relation as an example. The theorem to prove is the following:

$$E \vdash_c t : T \Leftrightarrow E \vdash t : T.$$

Both directions are straightforward inductions. The TYPING-ABS case of the $\Leftarrow$ direction is the only interesting case, where we must show $(E \vdash_c \text{abs } t_1 : T_1 \rightarrow T_2)$ from the induction hypothesis $(E, x:T_1 \vdash_c t_1^x : T_2)$ and an assumption $(x \notin \text{FV}(t_1))$. This case follows by applying C-TYPING-ABS, instantiating L to be $(\text{FV}(t_1) \cup \text{dom}(E) \cup \{x\})$. We are left with the subgoal

$$\forall y \notin L. (E, y:T_1 \vdash_c t_1^y : T_2)$$

and conclude using typing_rename.

4.4 Streamlining Developments

In general, then, developing metatheory in the locally nameless style involves the following steps.

1. State definitions using exists-fresh quantification.
2. Restate the definitions using cofinite quantification.
3. Prove basic lemmas, such as subst_intro.
4. Prove substitution and weakening lemmas.
5. Prove renaming lemmas.
6. Prove the equivalence of the cofinite and exists-fresh presentations.

However, we can sometimes get by with even less. To begin with, observe that the exists-fresh definitions are needed only in the first and last steps above. The main reason for stating the exists-fresh definitions is to check (informally) that the cofinite definitions adequately encode the relations of interest. Thus, it is only necessary to state the exists-fresh definitions for relations whose adequacy is crucial—those relations that appear hereditarily in the statements of theorems of interest. The style can be streamlined by omitting

¹³ The formal proofs can be found in file STLC_Core_Adequacy.v.

steps 1 and 6 for definitions of technical relations that are introduced only to help with proofs, as long as the equivalence of the existential and universal forms of definition is not actually needed in proofs (see Section 4.2).

In fact, the style can be streamlined further, according to need and personal taste. One may (after gaining some familiarity) feel convinced that the cofinite definitions are adequate encodings of one’s informal ideas, without the need of a formal proof of equivalence with their “more obviously adequate” exists-fresh versions.¹⁴ In this case, steps 1 and 6 can be omitted. By doing this, one may be risking adequacy of representation, but one is not risking consistency: the soundness of cofinite-quantified relations checked by Coq is no more in question than the soundness of Coq itself.

Finally, if we omit step 6, we can then prove renaming lemmas—step 5—only as needed, rather than going to the trouble of stating and proving them for all relations. Indeed, as we described in Section 4.2, our experience has been that they are rarely needed.

In short, the core of our style is to define relations using cofinite quantification and to reason about those particular definitions, i.e., steps 2 through 4. The remaining steps need not be carried out for every relation.

4.5 Comparison

As we mentioned in Section 2, we are not the first to propose a locally nameless representation for languages with binders. Nevertheless, our style of formal reasoning using this representation differs from previous work in important ways, requiring much less infrastructure as a result. In this subsection, we give a comparison to the styles used by Gordon (1994), McKinna and Pollack (1993; 1999), and Leroy (2007). To obtain a meaningful comparison, we sometimes apply their styles to our definitions.

Gordon was the first to use a locally nameless representation for formal reasoning about binders: he used them as a concrete datatype upon which to erect a type of named terms up to alpha equivalence. In contrast to our judgements (both the exists-fresh and cofinite versions), his are stated using the close operation instead of open. For example, Gordon’s local closure rule for abstraction (in our notation) is:

$$\frac{\text{term } t}{\text{term } (\text{abs } (\lambda^x t))}$$

Gordon observed that the rule induction principle arising from this form is “too weak” (Gordon 1994, Section 4) in that renaming a variable with a fresh name is sometimes required. He derives a strengthened induction principle by well-founded induction on the length of a term (as is standard). Then, to avoid further reasoning about length, he uses this strengthened induction principle to derive another strengthened induction principle that picks a name fresh for some finite set. However, Gordon’s strengthened induction principle is not the same as that arising from our form of cofinite quantification, since the induction hypothesis only applies to a single variable in the abstraction case. In this respect, Gordon’s strengthened principle is more similar to that of Urban et al. (2007b) than to ours. We get a stronger induction principle and stronger inversion principles automatically from the cofinite quantified forms.

McKinna and Pollack (1993, 1999) proposed a style of development for working with locally named terms. The most significant difference from our style is that it obtains strengthened induction principles by stating *universal* definitions of inductive relations,

which quantify bound variables over as many names as possible. For example, they would state the typing rule for abstractions as follows:

$$\frac{\forall x \notin \text{dom}(E). (E, x:S \vdash_a t^x : T)}{E \vdash_a \text{abs } t : S \rightarrow T} \text{A-TYPING-ABS}$$

This form of quantification is natural, as there is no arbitrary choice of fresh name in rule A-TYPING-ABS, hence there is at most one derivation of a judgement $E \vdash_a t : T$, and the associated induction principle is as strong as possible. The problem with relation $\vdash_a$ is that rule A-TYPING-ABS is too weak as an introduction rule, since it requires its premise to hold for every sufficiently fresh name, where cofinite quantification (rule C-TYPING-ABS) allows a finite number of problematic names to be excluded. In McKinna-Pollack style one must prove the equivalence between $\vdash$ and $\vdash_a$ (at significant cost in infrastructure) before metatheoretic results—step 4 in our style—can be developed, since $\vdash$ is too weak in elimination and $\vdash_a$ is too weak in introduction for either to prove both weakening and substitution lemmas for themselves. Thus our strategy for proving equivalence between definitions, which piggybacks on standard metatheoretic results, fails in the McKinna-Pollack style.

Leroy’s locally nameless solution to the POPLMARK Challenge uses McKinna and Pollack’s ideas. He takes universal definitions as the official ones, obtaining strong induction principles “for free”, and does not explicitly define the exist-fresh definitions at all. However, to make his proofs go through he needs to show that exists-fresh forms are admissible for the universal forms. E.g., Leroy proves the following lemma for his universally quantified typing relation:

$$\begin{aligned} \text{t.abs}' : \\ x \notin \text{FV}(t) \Rightarrow E, x:T_1 \vdash_a t^x : T_2 \Rightarrow \\ E \vdash_a \text{abs } t : T_1 \rightarrow T_2. \end{aligned}$$

Such lemmas, which are essentially equivalent to showing that $\vdash$ and $\vdash_a$ define the same relation, must be stated and proved for every rule with binding in the development. The proof of such lemmas requires a significant amount of infrastructure related to renaming, so Leroy’s style is a minor streamlining of the McKinna-Pollack style.

4.6 Nominal Isabelle and Local Representation

The nominal Isabelle package (Urban 2007; Urban et al. 2007a,b) supports representation in Isabelle/HOL of languages with binding (as an extended notion of datatype) using a single species of name for both bound and free variables. Using the HOL type definition facility, the package arranges that HOL equality of terms in these represented languages is up-to alpha equivalence, while users need not define alpha equivalence or see any quotient structure. The package supports multiple classes of names (e.g., type variables and term variables) in mutually inductive datatypes (e.g., terms and dependent types) and multiple relations on them (e.g., typing and reduction). Using the Isabelle typeclass facility, the package automatically provides definitions of name freshness and name permutation (renaming) polymorphic over these multiple aspects of a formalization. Importantly, given some syntactic restrictions on the represented languages, the package automatically derives equivariance lemmas and strengthened induction principles for formalized languages. The nominal package involves no axiomatic extension to Isabelle/HOL, hence is as safe as Isabelle/HOL itself. This nominal representation is a serious alternative to our style for users who want to do formal metatheory. However it is based on a great deal of bespoke infrastructure only available in Isabelle/HOL, there are some syntactic restrictions, and certain aspects of formalization (such as rule inversion and recursive definition over terms) are not

¹⁴ It is worth emphasizing that the exists-fresh versions *are* more obviously adequate, in a technical sense: they can be formalized in a very weak theory—Primitive Recursive Arithmetic—while cofinite presentations involve rules with “infinitely many hypotheses” that cannot be formalized in any conservative extension of PRA.

well supported (at time of writing) because of the need to respect alpha equivalence.

A locally nameless formalization (as we propose) or a locally named formalization (*à la* McKinna-Pollack) can be formalized in many logics and proof tools, depending only on basic notions like datatypes and inductively defined relations.¹⁵ In particular, the local representation can be seen as a nominal representation that doesn’t happen to use any nominal binding; in this case the nominal relation of freshness of names for language structures (terms, types, contexts, ...) is the same as the local representation’s notion of non-occurrence. This leads to an interesting possibility: local representations can be defined and reasoned about using the nominal Isabelle package, getting the definitions of freshness and renaming and the proofs of equivariance for free from the package. Interestingly, the nominal package can also automatically infer a strengthened induction principle for inductive relations over local representations (with some syntactic restrictions); see (Urban and Pollack 2007) for more details.

A detailed comparison between nominal representation and locally nameless representation is beyond the scope of this paper, but, as we are focusing on strengthened induction principles, we can remark that the locally nameless strengthened induction principle (e.g., from rule C-TYPING-ABS or rule A-TYPING-ABS) is stronger than the strengthened induction principle automatically derived by the nominal Isabelle package. For a full explanation of the nominal strengthened induction, see Urban et al. (2007b); here it is enough to observe that, when using such an induction principle, one specifies not only the derivation being eliminated (“induction on ...”), but also an arbitrary finitely supported structure to be used as a *freshness context*. Whenever a chosen name appears in an induction hypothesis (e.g., the name that is bound in an abstraction typing rule), that name is known to be fresh for the specified freshness context. This solves many problems, as in the example above of weakening for the typing judgement. However, you must specify the freshness context when *invoking* induction. In contrast, the locally nameless strengthened induction principles, with their universally quantified fresh names, allow fresh instantiation of these names later, at the point where you have seen the particular cases that must be proved. Since particular cases may involve existential quantifiers, you may not know what the necessary freshness context is until after seeing a case and eliminating an existential quantifier. In such rare examples, the automatically derived nominal strengthened induction principle is not adequate (a detailed hand proof is necessary) while the locally nameless strengthened induction principle works perfectly.

5. Practical Formal Metatheory

We close with a concrete discussion of our Coq realization of working with cofinite quantification. We present the major developments we have carried out—System $F_{<}$, the Calculus of Constructions, and extensions of ML—to demonstrate that our techniques scale to larger languages and are expressive enough for wide use. We discuss issues specific to each development (e.g., dealing with multi-binders) as well as properties common to all of them (e.g., the conciseness and robustness on change of our scripts). Finally, we give an overview of the structure of our developments for readers who may wish to use them as a basis for their own formalizations.

¹⁵ Some specialized tactics, as in the accompanying Coq code <http://arthur.chargueraud.org/research/2007/binders/>, are useful, but not essential. In any case these tactics are very much simpler than the nominal Isabelle package.

5.1 System $F_{<}$

The heart of the POPLMARK Challenge is a proof of type soundness of System $F_{<}$. This task was designed as a stress test for formalized metatheory. Our solution closely follows the structure proposed in the appendix of the challenge’s description.

To get a sense of how our solution compares to other solutions, we measured the number of lemmas and *steps of reasoning* involved in several developments;¹⁶ the results are summarized in Figure 5. Our study is restricted to Coq developments since there are no meaningful metrics across different proof assistants. In this comparison, a step of reasoning is defined as the invocation of one tactic, excluding trivial ones: *intros* (introduces hypotheses into the context), *auto* (performs automated proof search), and simple variations of these two. We also exclude from the counts material from libraries that are reusable across metatheory developments. We isolate part 1A of the challenge (technical properties of the subtyping relation) from parts 1A+2A (preservation and progress) because some submissions cover only part 1A.

Even when attention is restricted to Coq developments, the comparison is necessarily rough, since the developments were carried out by different people with different preferences in the amount of automation they used, etc. Nevertheless, a few interesting points emerge. The first line of the table shows that our style compares well against Vouillon’s development, which uses pure de Bruijn indices—a representation commonly thought to be quite effective. The solutions that do not use a standard de Bruijn or locally nameless representation (Stump and Hirschowitz & Maggesi) require significantly more steps. The last line shows that our style requires significantly less infrastructure than Leroy’s solution, a locally nameless solution faithful to McKinna and Pollack’s way of dealing with the quantification of names. Chlipala’s is a variation on an earlier version of Leroy’s development and is the only one that is shorter than ours, mainly because Chlipala’s proof script relies heavily on development-specific tactics.

Although clever automation and development-specific tactics can make scripts extremely small (Nipkow [2001] gives an impressively automated proof of confluence, for example), they can also make them slow (hard to work with) and brittle (hard to reuse). Brittleness emerges when automation fails after some changes to earlier definitions or lemmas: it can be hard to know exactly what has been broken and how to fix it. Of course, the other extreme of no automation at all is also unmanageable, especially for metatheory proofs, which involve many related arguments. Moreover, too little automation can also make scripts brittle: even a tiny change in definitions or statements of lemmas will require update. Aiming for maximal robustness, our convention is to rely as much as possible on automation for low-level details (e.g., proving that a term is locally closed) while giving all the “important” arguments (the ones that would appear in a rigorous paper proof) manually.

5.2 Calculus of Constructions (CoC)

This formalization contains three main components. The first is a proof of confluence for parallel beta-reduction. Its use of the variable closing operation requires slightly more advanced reasoning about binders. We have also used this development to evaluate our scripts for robustness: we first completed a proof of confluence for the pure lambda calculus and then migrated it to CoC, preserving the same flow of arguments and just adding the extra cases and updating the names of some hypotheses.

The second part consists of a set of inversion results for the typing and the conversion judgments, which are quite technical in themselves but do not involve much difficulty with respect to

¹⁶ All are publicly available from <http://alliance.seas.upenn.edu/~plclub/cgi-bin/poplmark/>.

Author (chronological order)	Representation used	Lemmas 1A	Proof steps 1A	Lemmas 1A+2A	Proof steps 1A+2A
Vouillon	de Bruijn	30	402	72	1175
Leroy	locally nameless	49	495	128	1364
Stump	levels/names	56	938	-	-
Hirschowitz & Maggesi	de Bruijn (nested datatype)	49	1574	-	-
Chlipala	locally nameless	23	75	-	-
Our development	locally nameless	22	101	65	576

Figure 5. Comparison of Coq submissions to the POPLMARK Challenge

Language	Results formalized	Number of trusted definitions	Infrastructure lemmas	Core lemmas	Infrastructure proof steps	Core proof steps
$\lambda \rightarrow$	Preservation and progress	13	13	4	49	45
System F_{λ}	Preservation and progress	20	48	17	335	241
ML core only	Preservation and progress	23	37	8	181	98
ML with features	Preservation and progress	42	55	18	367	396
Lambda calculus	Church-Rosser	11	21	25	112	160
CoC	Church-Rosser, preservation	17	36	55	214	475

Figure 6. Complexity of our developments

binders. This part is representative of the effort required for the CoC formalization: we felt that the largest share of the work was to understand and be able to state clearly the arguments involved. The fact that the proofs were machine-checked appeared to be only a way to force oneself to understand all the details of the proofs formalized, not as an extra burden to be dealt with.

The third and last part is the core of the preservation proof, which shares a significant amount of structure with the corresponding proof for System F_{λ} . (only the steps of reasoning are more complex to follow). Compared to the System F_{λ} development, the fact that terms and types are represented in a uniform way in CoC and that proofs are more complex reduces the binding infrastructure from about 60% of the total development down to 35% percent.

Compared to Barras’s formalization of CoC in Coq (1997) using a pure de Bruijn representation, our style saves us from numerous issues associated with shifting indices—particularly bothersome when working with CoC’s dependent types—and so, unlike Barras, we did not need clever engineering of the statements of lemmas to make them fit the representation used.

5.3 Extensions of ML

We also investigated how to extend the basic $\lambda \rightarrow$ development with features such as datatypes, fixpoints, references, exceptions, pattern matching, and ML-style polymorphic types. Then, building on the individual successes of each of these extensions, we set up the formalization of a language containing all of these features together. This experiment demonstrates that our style can be applied to full-featured programming languages, not just tiny lambda-calculi.

Pleasingly, most of the arguments that would be omitted as “trivial” on paper and which do not involve binders are solved by the `auto` tactic. In particular, the infrastructure lemmas do not need any change when we add features like pairs or references. The key proofs—preservation and progress—contain respectively in 60 and 80 lines, which seems reasonable for a system with 24 typing rules.

A key idea involved in this development is a treatment of *multi-binders*, used to encode ML type schemes, patterns, and recursion. The basic operation on a multi-binder is to open it with a *list* of terms. A bound variable pointing to the j th variable bound by the i th binder above the current position is represented as `(trm.bvar  $i$   $j$ )`. Free variables are still represented as one single name and substitution lemmas are still stated in terms of an atomic

substitution from one name to one term. Then, to prove subject reduction, for example, the substitution lemma is applied iteratively in the cases involving multi-binders.

5.4 Organization of Our Developments

Each of our developments is divided into three main parts: *trusted definitions*, *infrastructure*, and *core lemmas*.

The *trusted definitions* part contains the description of the language formalized—the syntax of the language, its semantics, and its type system, if any—and the statements of the main theorems which are to be proven about the language. This part is the only material that needs to be trusted, in the following sense: if one is convinced about the adequacy of these definitions and trusts that Coq correctly type-checks all proofs in the development, then one can have confidence that the system has been formally certified.

The *infrastructure* part sets up the machinery required for the *core lemmas* and consists of several components:

1. Language-specific specializations of tactics for working with cofinite quantification, e.g., to automatically choose a set L when applying a rule that uses cofinite quantification.
2. Proofs about properties of substitution (Figure 2).
3. Proofs that local closure is preserved by various operations, e.g., substitution (Section 3.3).
4. *Regularity lemmas* which state that relations contain only locally closed terms (Section 3.3).
5. Hints to enable Coq’s automation to use regularity lemmas.

Note that the lemmas about substitution are similar from one language to the next. When setting up a new development, those who wish to get to “interesting” proofs quickly may state as axioms properties of substitution and regularity lemmas, proving them only after they believe that their main proofs will go through. In this manner, one can “test-drive” a language definition without having to revise infrastructure proofs as the definition is modified.

Finally, the *core lemmas* part contains the lemmas that would normally be stated in an informal presentation. The statements closely match their informal counterparts, and the proofs contain few uninteresting details—facts about local closure and freshness side conditions are almost always handled automatically.

We conclude this section with a breakdown of our developments by the number of definitions and lemmas in each part, as well as the number of proof steps required to prove the lemmas; the results are summarized in Figure 6. The number of trusted definitions gives an idea of the size of the language formalized. The number of core lemmas gives an idea of the theoretical complexity of the development; recall that these lemmas correspond closely with results that would be proven informally. The amount of infrastructure, in terms of both lemmas and steps, is proportional to the number of binding constructs and relations in the language. A comparison between figures from the last two columns suggests that the amount of infrastructure work is reasonable compared to the core proof work—especially since infrastructure proofs follow a standard pattern and are easily set up. More precisely, the ratio “infrastructure over core proofs” increases with the number of binding constructs involved but decreases with the inherent complexity of the development.

6. Conclusion

We have argued for a novel style for formalizing programming-language metatheory, based on a combination of locally nameless representation with a cofinite idiom for quantifying free names in inductive definitions.

This design satisfies the evaluation criteria of POPLMARK. First, our presentation is transparent: the proofs closely follow their informal equivalents. Second the overheads of the approach are low: we do not need manual proofs of freshness side-conditions nor reasoning about alpha-equivalence, and only on rare occasions do we need to justify well-formedness of objects (e.g. local closure) or explicitly rename variables. When we do, the required proofs of renaming lemmas follow with virtually no infrastructure. At the same time, there is no need for external tools, and the style works in any general purpose theorem prover (although we found Coq to be well-suited to the task). In our scripts, definitions and results about variables, freshness, and environments are factored into a reusable library. Finally, experience with a number of large developments suggests that the approach is “complete” in the informal sense that any language definition and accompanying reasoning techniques that would be accepted as informally correct can be carried out in this style. Our formalizations provide a good starting point for new work—reports from early adopters confirm that modifying them is much easier than starting a new development from scratch.

In the future, we would like to integrate our approach with the Ott tool (Sewell et al. 2007), which automatically generates Coq (and $\text{\LaTeX}$) definitions from compact, high-level language descriptions. A more ambitious next step is to study the possibility of automatically generating the required infrastructure proofs, which follow a fairly simple pattern. Finally, we would like to explore certified programming of tools such as type checkers and compilers that must deal with binders. In the long term, we hope that our techniques will help with the verification of both specifications and implementations of programming languages.

Acknowledgments Many thanks to the members of the Penn PLClub for extensive feedback and discussion, and to Michael Norrish and Christian Urban for important technical insights and to the POPL reviewers, Bob Harper, Peter Sewell, and Karl Crary for helping improve the final version. This work was supported by NSF grant 545886, *CRI: Machine Assistance for Programming Language Research*. Pollack thanks the International Centre for Mathematical Sciences (Edinburgh) Workshop on Abstraction, Substitution and Naming and the EU Framework 6 Coordination Action 510996, TYPES.

References

- Thorsten Altenkirch. A formalization of the strong normalization proof for System F in LEGO. In Bezem and Groote (1993), pages 13–28.
- Andrew W. Appel. Foundational proof-carrying code. In *IEEE Symposium on Logic in Computer Science (LICS)*, Boston, Massachusetts, pages 247–258, June 2001.
- Michael Ashley-Rollman, Karl Crary, and Robert Harper. Submission to the POPLMARK challenge. Available from <http://www.cis.upenn.edu/~plclub/mmm/>, 2005.
- Brian E. Aydemir, Aaron Bohannon, Matthew Fairbairn, J. Nathan Foster, Benjamin C. Pierce, Peter Sewell, Dimitrios Vytiniotis, Geoffrey Washburn, Stephanie Weirich, and Steve Zdancewic. Mechanized metatheory for the masses: The POPLMARK challenge. In Joe Hurd and Tom Melham, editors, *Theorem Proving in Higher Order Logics: 18th International Conference, TPHOLs 2005*, volume 3603 of *Lecture Notes in Computer Science*, pages 50–65. Springer, 2005.
- Henk P. Barendregt. *The Lambda Calculus*. North Holland, revised edition, 1984.
- Bruno Barras and Benjamin Werner. Coq in coq. Available from http://pauillac.inria.fr/~barras/coq_work-eng.html, 1997.
- M. Bezem and J. F. Groote, editors. *Typed Lambda Calculi and Applications: International Conference on Typed Lambda Calculi and Applications, TLCA '93*, volume 664 of *Lecture Notes in Computer Science*, 1993. Springer.
- Anna Bucalo, Furio Honsell, Marino Miculan, Ivan Scagnetto, and Martin Hoffman. Consistency of the theory of contexts. *J. Funct. Program*, 16 (3), 2006.
- Adam Chlipala. Submission to the POPLMARK challenge, part 1a. Available from <http://www.cs.berkeley.edu/~adamc/poplmark/>, 2006.
- The Coq Development Team. The Coq proof assistant reference manual, version 8.1. Available from <http://coq.inria.fr/>, 2007.
- 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.
- Karl Crary. Toward a foundational typed assembly language. In *POPL '03: Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 198–212. ACM Press, 2003.
- N. G. de Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae*, 34(5):381–392, 1972.
- Joëlle Despeyroux, Amy Felty, and André Hirschowitz. 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*, pages 124–138. Springer, 1995. Also available as INRIA Research Report 2556.
- Peter Dybjer. Inductive families. *Formal Aspects of Computing*, 6:1–26, 1994.
- Jonathan M. Ford and Ian A. Mason. Operational techniques in PVS — A preliminary evaluation. *Electronic Notes in Theoretical Computer Science*, 42, 2001.
- Gerhard Gentzen. *The Collected Papers of Gerhard Gentzen*. North-Holland, 1969. Edited by Mandred Szabo.
- Andrew D. Gordon. A mechanisation of name-carrying syntax up to alpha-conversion. In J. J. Joyce and C.-J. H. Seger, editors, *Higher-order Logic Theorem Proving And Its Applications, Proceedings, 1993*, volume 780 of *Lecture Notes in Computer Science*, pages 414–426. Springer, 1994.
- Andrew D. Gordon and Tom Melham. Five axioms of alpha-conversion. In J. von Wright, J. Grundy, and J. Harrison, editors, *Theorem Proving in Higher Order Logics: 9th International Conference, TPHOLs '96*, volume 1125 of *Lecture Notes in Computer Science*, pages 173–190. Springer, 1996.

- Robert Harper and Daniel R. Licata. Mechanizing metatheory in a logical framework. *Journal of Functional Programming*, 17(4–5):613–673, 2007.
- Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
- Dimitri Hendriks and Vincent van Oostrom. Adbmal. In F. Baader, editor, *Automated Deduction – CADE-19*, volume 2741 of *Lecture Notes in Artificial Intelligence*, pages 136–150. Springer-Verlag, 2003.
- Peter Homeier. A proof of the Church-Rosser theorem for the lambda calculus in higher order logic. In Richard J. Boulton and Paul B. Jackson, editors, *TPHOLs 2001: Supplemental Proceedings*, pages 207–222. Division of Informatics, University of Edinburgh, September 2001. Available as Informatics Research Report EDI-INF-RR-0046.
- Furio Honsell, Marino Miculan, and Ivan Scagnetto. The theory of contexts for first order and higher order abstract syntax. *Electronic Notes in Theoretical Computer Science*, 62, 2002.
- Grard Huet. The constructive engine. In Raghavan Narasimhan, editor, *A Perspective in Theoretical Computer Science: Commemorative Volume for Gift Siromoney*. World Scientific Publishing, 1989. Also available as INRIA Technical Report 110.
- Grard Huet. Residual theory in λ -calculus: A formal development. *Journal of Functional Programming*, 4(3):371–394, July 1994. Also available as INRIA Research Report 2009 (August 1993).
- Gerwin Klein and Tobias Nipkow. A machine-checked model for a Java-like language, virtual machine, and compiler. *ACM Transactions on Programming Languages and Systems*, 28(4):619–695, 2006.
- J. L. Krivine. *Lambda-Calculus, Types and Models*. Ellis Horwood, 1990.
- Daniel K. Lee, Karl Crary, and Robert Harper. Towards a mechanized metatheory of Standard ML. In *POPL ’07: Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 173–184. ACM Press, 2007.
- Xavier Leroy. Formal certification of a compiler back-end, or: programming a compiler with a proof assistant. In *Proc. of the 33rd Symposium on Principles of Programming Languages*, pages 42–54. ACM Press, 2006.
- Xavier Leroy. A locally nameless solution to the POPLmark challenge. Research report 6098, INRIA, January 2007.
- Zhaohui Luo and Robert Pollack. The LEGO proof development system: A user’s manual. Technical Report ECS-LFCS-92-211, University of Edinburgh, May 1992.
- Conor McBride and James McKinna. Functional pearl: I am not a number—I am a free variable. In *Haskell ’04: Proceedings of the 2004 ACM SIGPLAN Workshop on Haskell*, pages 1–9. ACM Press, 2004.
- James McKinna and Robert Pollack. Pure Type Systems formalized. In Bezem and Groote (1993), pages 289–305.
- James McKinna and Robert Pollack. Some lambda calculus and type theory formalized. *Journal of Automated Reasoning*, 23(3–4):373–409, 1999.
- Tobias Nipkow. More Church-Rosser proofs (in Isabelle/HOL). *Journal of Automated Reasoning*, 26(1):51–66, January 2001.
- Michael Norrish and Konrad Slind. HOL 4. Available from <http://hol.sourceforge.net/>, 2007.
- Frank Pfenning and Conal Elliot. Higher-order abstract syntax. In *PLDI ’88: Proceedings of the ACM SIGPLAN 1988 Conference on Programming Language Design and Implementation*, pages 199–208. ACM Press, 1988.
- Frank Pfenning and Carsten Schrmann. System description: Twelf — A meta-logical framework for deductive systems. In Harald Ganzinger, editor, *Automated Deduction, CADE 16: 16th International Conference on Automated Deduction*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 202–206. Springer, 1999.
- Andrew M. Pitts. Nominal logic, a first order theory of names and binding. *Information and Computation*, 186:165–193, 2003.
- Randy Pollack. Reasoning about languages with binding: Can we do it yet?, February 2006. Presentation, slides available from <http://homepages.inf.ed.ac.uk/rpollack/>.
- Robert Pollack. Closure under alpha-conversion. In H. Barendregt and T. Nipkow, editors, *TYPES’93: Workshop on Types for Proofs and Programs, Nijmegen, May 1993, Selected Papers*, volume 806 of *Lecture Notes in Computer Science*, pages 313–332. Springer, 1994a.
- Robert Pollack. *The Theory of LEGO: A Proof Checker for the Extended Calculus of Constructions*. PhD thesis, Univ. of Edinburgh, 1994b.
- Dag Prawitz. *Natural Deduction: Proof Theoretical Study*. Almqvist and Wiksell, Stockholm, 1965.
- Ole Rasmussen. The Church-Rosser theorem in Isabelle: A proof porting experiment. Technical Report 364, University of Cambridge, Computer Laboratory, March 1995.
- Wilmer Ricciotti. Submission to the POPLMARK challenge, part 1a. Available from <http://ricciott.web.cs.unibo.it/>, 2007.
- Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Gilles Peskine, Thomas Ridge, Susmit Sarkar, and Rok Strnia. Ott: Effective tool support for the working semanticist. In *ICFP ’07: Proceedings of the 2007 ACM SIGPLAN International Conference on Functional Programming*, pages 1–12. ACM, 2007.
- Natarajan Shankar. A mechanical proof of the Church-Rosser theorem. *Journal of the Association for Computing Machinery*, 35(3):475–522, 1988.
- Allen Stoughton. Substitution revisited. *Theoretical Computer Science*, 59(3):317–325, 1988.
- Christian Urban. Nominal techniques in Isabelle/HOL. *Journal of Automatic Reasoning*, 2007. To appear; available from <http://www4.in.tum.de/~urbanc/publications.html>.
- Christian Urban and Randy Pollack. Locally nameless representation in Nominal Isabelle. Talk at Workshop on Mechanizing Metatheory. Available from www4.in.tum.de/~urbanc/Publications/ln.pdf, 2007.
- Christian Urban, Stefan Berghofer, and Julien Narboux. Nominal datatype package for Isabelle/HOL. Available from <http://isabelle.in.tum.de/nominal/>, 2007a.
- Christian Urban, Stefan Berghofer, and Michael Norrish. Barendregt’s variable convention in rule inductions. In *Proceedings of the 21th Conference on Automated Deduction (CADE 2007)*, volume 4603 of *Lecture Notes in Computer Science*, pages 35–50. Springer, 2007b.
- Ren Vestergaard and James Brotherston. A formalised first-order confluence proof for the λ -calculus using one-sorted variable names. *Information and Computation*, 183(2):212–244, 2003.