

Mechanizing Metatheory in a Logical Framework

Robert Harper and Daniel R. Licata

Carnegie Mellon University
(e-mail: {rwh, drl}@cs.cmu.edu)

Abstract

The LF logical framework codifies a methodology for representing deductive systems, such as programming languages and logics, within a dependently typed λ -calculus. In this methodology, the syntactic and deductive apparatus of a system is encoded as the canonical forms of associated LF types; an encoding is correct (*adequate*) if and only if it defines a *compositional bijection* between the apparatus of the deductive system and the associated canonical forms. Given an adequate encoding, one may establish metatheoretic properties of a deductive system by reasoning about the associated LF representation. The Twelf implementation of the LF logical framework is a convenient and powerful tool for putting this methodology into practice. Twelf supports both the representation of a deductive system and the mechanical verification of proofs of metatheorems about it.

The purpose of this paper is to provide an up-to-date overview of the LF λ -calculus, the LF methodology for adequate representation, and the Twelf methodology for mechanizing metatheory. We begin by defining a variant of the original LF language, called *Canonical LF*, in which only canonical forms (long $\beta\eta$ -normal forms) are permitted. This variant is parameterized by a *subordination relation*, which enables modular reasoning about LF representations. We then give an adequate representation of a simply typed λ -calculus in Canonical LF, both to illustrate adequacy and to serve as an object of analysis. Using this representation, we formalize and verify the proofs of some metatheoretic results, including preservation, determinacy, and strengthening. Each example illustrates a significant aspect of using LF and Twelf for formalized metatheory.

1 Introduction

A formal definition of a programming language provides a precise specification for programmers, ensures compatibility among compilers, and enables rigorous analysis of its properties. However, a language definition is an intricate artifact, and the proofs of its properties are often complex and subtle. Consequently, it can be difficult to define a language correctly and prove the appropriate theorems about it, let alone to maintain the definition and associated proofs as a language evolves. Fortunately, some of this burden can be alleviated by mechanizing the definition and metatheory of a language.

In this article, we present a technical introduction to mechanizing languages using LF (Harper *et al.*, 1993) and Twelf (Pfenning & Schürmann, 1999). The literature on formalizing and verifying languages in LF and Twelf is dispersed among numerous research papers, dissertations, course notes, and on-line repositories. The

goal of this paper is to consolidate these ideas and make them more accessible to a broader audience. We present an up-to-date overview of the LF λ -calculus, the LF methodology for representing languages, and the Twelf methodology for mechanically verifying the metatheory of a language. We hope to suggest that this methodology for mechanizing languages is effective and feasible; however, the successful Twelf mechanization efforts, some of which are listed in Section 5, support this claim better than this article can.

Mechanizing Definitions in LF. Language formalization efforts frequently start from an informal, on-paper definition of the language to be mechanized—in, for example, a textbook or a research paper. What does it mean to correctly represent such a language definition in a formal framework? In the methodology of the the LF logical framework (Harper *et al.*, 1993), an LF representation of a language is *adequate* iff it is isomorphic to the informal definition of the language. Adequacy is a useful correctness criterion because any reasoning in or about an adequate LF representation applies equally well to the informal definition, and vice versa. For example, metatheoretic proofs about an adequate LF representation establish properties of the original definition as well.

LF is a minimal dependent type theory. The syntax and judgements of an *object language* (the object of study) are represented in LF as the canonical forms (essentially long $\beta\eta$ -normal forms) of associated LF types. These canonical forms are specified by an LF *signature*, which declares type and term constants, and by a *world*, which specifies the LF contexts under consideration. A representation is adequate iff there is an isomorphism between the informal presentation of the object language and the associated canonical forms. The LF methodology has been successfully used to derive representations of a wide variety of logical systems, often leading to new insights about their structure.

One can often use variable binding in LF to represent object-language binding forms and hypothetical judgements. With such higher-order representations of syntax and judgements, LF binding provides object-language operations such as α -conversion and capture-avoiding substitution and object-language judgement properties such as weakening and substitution. Because binding and substitution are present in both the informal presentation of a language and its LF representation, an isomorphism between the two must respect this structure. Thus, we say that a higher-order LF representation is adequate iff there is a *compositional bijection* between the informal presentation of the language and the associated canonical forms, where a bijection is compositional iff it commutes with substitution. This property ensures that the object language’s variables and hypothetical judgements are correctly modeled.

Establishing the adequacy of a representation requires two main technical tools. First, adequacy proofs proceed by *induction on the canonical forms of LF*. The representation of an object language is not an inductive definition inside the LF type theory; indeed, higher-order encodings rely on negative occurrences of types. However, externally, the representation of a language as the canonical forms of par-

ticular types in a particular world *is* an inductive definition because the canonical forms of LF are inductively defined.

Second, modular adequacy proofs require *subordination-based transport of adequacy*. To modularly encode languages in LF, it is necessary to consider each fragment of the object language in only the LF signature and contexts relevant to its encoding. Combining these fragments requires understanding whether an encoding that is adequate in one signature and world remains adequate in another. To understand when adequacy may be transported, it suffices to consider conditions under which the canonical forms of a type remain invariant when transported from one signature and world to another. The judgements of LF are parametrized by a *subordination relation* (Virga, 1999), which determines whether canonical forms of one type can appear in canonical forms of another. Using subordination, it is possible to establish general conditions under which the canonical forms of a type, and therefore the adequacy theorems proved about them, remain invariant when transported from one signature and world to another.

Though adequacy underlies the LF methodology, it is not an obstacle to mechanization in practice. Experienced users of LF often define a language solely by its LF representation, never even stating an informal description. However, even such a definition will rely on an understanding of how to represent a language as the canonical forms of particular LF types in contexts of particular forms, which is exactly what an adequacy proof clarifies. Thus, we claim that all users of LF should understand adequacy, even if they do not carry out a proof for every object language.

Mechanizing Metatheory in Twelf. One application of adequate LF encodings is the formalization and verification of a language’s metatheory. Since the entire deductive apparatus of a language is represented by certain canonical forms in LF, it is possible to reason about its properties by reasoning about the associated canonical forms. Specifically, an informal proof by induction on the structure of derivations can be re-cast, via an adequate encoding, as a structural induction on the canonical forms of the associated types.

This capability is exploited by the Twelf implementation of LF (Pfenning & Schürmann, 1999), which supports the mechanized proofs of $\forall\exists$ -propositions over the canonical forms of specified types. This is sufficient to capture many useful properties of programming languages and logics. The Twelf methodology has proved to be very robust: it has been used to develop the metatheory of a wide variety of systems, including Mini-ML (Michaylov & Pfenning, 1991), the typed assembly language TALT (Crary, 2003), the POPLmark Challenge problem (Aydemir *et al.*, 2005), and the Harper-Stone internal language for Standard ML (Lee *et al.*, 2007). The same technical tools required for adequacy proofs are necessary for Twelf metatheory: Twelf proofs proceed by induction on the canonical forms of LF, and modular Twelf proofs require subordination-based transport of canonical forms.

Outline. In Section 2, we give a presentation of the LF type theory. Our presentation, called *Canonical LF*, follows Watkins *et al.* (2002, 2004a) in defining the

type theory so that only canonical forms are permitted, giving a direct inductive definition of the canonical forms. Additionally, we follow Virga (1999) in supporting subordination in the type theory. Next, we use a simply typed λ -calculus as a running example of mechanizing a language and its metatheory in LF and Twelf. In Section 3, we give an adequate LF encoding of the syntax and semantics of the simply typed λ -calculus. In Section 4, we show how to mechanize some of the metatheory of the simply typed λ -calculus in Twelf. We give mechanized proofs of several of its properties, including type preservation, determinacy of the operational semantics, and strengthening.

An electronic copy of the Twelf code presented in this article is available as supplementary material from the Journal of Functional Programming Web site (<http://www.cambridge.org/jfp/>).

2 Canonical LF

As we stated above, the LF methodology consists of representing object languages as canonical forms in a dependent type theory. In a dependent type theory, terms of the type theory are allowed to appear in types. In LF, dependency arises from considering families of types indexed by LF terms; these type families are used to represent object-language judgements, as we discuss in Section 3. In the presence of dependency, it is necessary to generalize the simple function type $A_1 \rightarrow A_2$ to a dependent function type $\Pi x:A_1. A_2$. In such a type, the variable x , which stands for the argument to the function, may appear in the result type A_2 ; this permits the result type of the function to vary with the argument provided to it.

The canonical forms of LF are the long $\beta\eta$ -normal forms: a term is a canonical form iff it is not β -reducible and it is η -expanded as much as possible without introducing β -redices. Traditional presentations of LF (Harper *et al.*, 1993; Salvesen, 1990; Geuvers, 1992; Harper & Pfenning, 2005) first define the type theory so that both canonical and non-canonical forms are well-typed and then give an inductive characterization of which terms are canonical forms. However, because the LF representation methodology requires only canonical forms, the non-canonical terms can be seen as a technical device, rather than as an essential part of the type theory. Technically, the reason to admit non-canonical forms is that they arise from the substitution of one canonical form into another—i.e., canonical forms are not closed under substitution. Such a substitution occurs in the typing rule for the application of a dependent function to an argument:

$$\frac{M_1 : \Pi x:A_2. A \quad M_2 : A_2}{M_1 M_2 : [M_2/x]A}$$

Even if M_2 and all the terms embedded in the type $\Pi x:A_2. A$ are in canonical form, the substitution $[M_2/x]A$ may introduce non-canonical forms if A_2 is itself a function type. To equate non-canonical terms with an associated canonical form, traditional presentations of LF consider terms modulo β - or $\beta\eta$ -equivalence.

Canonical LF, which is a fragment of Concurrent LF (Watkins *et al.*, 2002, 2004a), is a presentation of LF in which *only* the canonical forms are well-typed. The key

Kinds	K	$::=$	type	$ \ \Pi x:A. K$
Canonical Type Families	A	$::=$	P	$ \ \Pi x:A_2. A$
Atomic Type Families	P	$::=$	a	$ \ P M$
Canonical Terms	M	$::=$	R	$ \ \lambda x. M$
Atomic Terms	R	$::=$	x	$ \ c \mid R M$
Signatures	Σ	$::=$	\cdot	$ \ \Sigma, c : A \mid \Sigma, a : K$
Contexts	Γ	$::=$	\cdot	$ \ \Gamma, x : A$

Fig. 1. LF Syntax

technical contribution of the canonical-forms approach is a notion of *hereditary substitution*, which directly computes the canonical form of the result of an ordinary substitution. In Canonical LF’s application typing rule, hereditary substitution is used to directly compute the canonical form resulting from an ordinary substitution $[M_2/x]A$; the intermediate non-canonical form is never considered. When a traditional substitution would result in a redex $(\lambda x. M_1) M_2$, the corresponding hereditary substitution continues by hereditarily substituting M_2 for x in M_1 . The key insight made in Watkins *et al.* (2004a), which enables this approach, is that hereditary substitution can be defined in a structurally recursive manner, using the same algorithm as structural cut elimination (Pfenning, 1994). The corresponding structural induction principle is used to establish the metatheoretic properties of Canonical LF. Consequently, the metatheory of Canonical LF is simpler than presentations that admit non-canonical forms—there is no need to consider $\beta\eta$ -equality.

Following Virga (1999), the judgements of Canonical LF are parametrized by a subordination relation, which determines when canonical forms of one type are relevant to canonical forms of another. Under appropriate subordination conditions, the canonical forms of a type remain invariant when considered in a different signature and world. Specifically, the addition or removal of canonical forms of types that are *not* subordinate to a given type does not change the canonical forms of that type.

In Sections 2.1, 2.2, and 2.3, we overview the syntax, judgements, and properties of Canonical LF. In Section 2.4, we give a formal definition of subordination and prove the transport of canonical forms theorem.

2.1 Syntax

In Figure 1, we present the syntax of LF. The term level includes functions, application, variables x (which are bound in contexts Γ), and constants c (which are declared in signatures Σ). Functions are given dependent function types $\Pi x:A_2. A$; we write $A_2 \rightarrow A$ as an abbreviation when x is not free in A . The base types are family-level constants a and their applications to terms. The original presentation of LF (Harper *et al.*, 1993) included family-level λ -abstractions; because these cannot

$$\boxed{\vdash_{\preceq} \Sigma \text{ sig}}$$

$$\frac{}{\vdash_{\preceq} \cdot \text{sig}} \text{SIG_EMPTY} \qquad \frac{\vdash_{\preceq} \Sigma \text{ sig} \cdot \vdash_{\Sigma, \preceq} \mathbf{A} \text{ type} \quad c \# \Sigma}{\vdash_{\preceq} \Sigma, c : \mathbf{A} \text{ sig}} \text{SIG_TERM}$$

$$\frac{\vdash_{\preceq} \Sigma \text{ sig} \cdot \vdash_{\Sigma, \preceq} \mathbf{K} \text{ kind} \quad a \# \Sigma}{\vdash_{\preceq} \Sigma, a : \mathbf{K} \text{ sig}} \text{SIG_FAM}$$

$$\boxed{\vdash_{\Sigma, \preceq} \Gamma \text{ ctx}}$$

$$\frac{}{\vdash_{\Sigma, \preceq} \cdot \text{ctx}} \text{CTX_EMPTY} \qquad \frac{\vdash_{\Sigma, \preceq} \Gamma \text{ ctx} \quad \Gamma \vdash_{\Sigma, \preceq} \mathbf{A} \text{ type} \quad x \# \Gamma}{\vdash_{\Sigma, \preceq} \Gamma, x : \mathbf{A} \text{ ctx}} \text{CTX_TERM}$$

Fig. 2. LF Signature and Context Formation

appear in canonical types, we omit them from the current presentation. The kind **type** classifies types; the inhabitants of dependent function kinds are family-level constants and their applications to terms.

To admit only canonical forms, the syntax of terms is stratified into two categories, the canonical terms \mathbf{M} and the atomic terms \mathbf{R} ; this syntactically precludes β -redices. For example, the term $(\lambda x. x) c$ is not syntactically correct, let alone well-typed, because the λ -abstraction is not an atomic term \mathbf{R} . The type family level is similarly stratified into \mathbf{A} and \mathbf{P} to differentiate between Π -types and base types; this stratification will be used in the formation rules to ensure that well-typed terms are η -long.

We implicitly consider all expressions \mathbf{M} , \mathbf{R} , \mathbf{A} , \mathbf{P} , and \mathbf{K} up to α -equivalence. By convention, when we write a binding form $(\lambda x. M, \Pi x : \mathbf{A}_2. A, \text{ or } \Pi x : \mathbf{A}. K)$, the bound variable is chosen to be fresh.

2.2 Judgements

The judgements defining LF are presented in Figures 2, 3, 4, and 5. For each judgement, we first identify the judgement form in a box, (e.g., $\boxed{\Sigma \text{ sig}}$), and then we give the inference rules defining that judgement.

Parameters to the Formation Judgements. The signature formation judgement is parametrized by a subordination relation \preceq , which is a binary relation between type families. The context, kind, type family, and term formation judgements are parametrized by a signature Σ and a subordination relation \preceq . Subordination is discussed in detail in Section 2.4. Taking \preceq to be the complete relation imposes no restrictions on which judgements are derivable, recovering a presentation of Canonical LF without subordination as a special case.

Signature and Context Formation. The judgements in Figure 2 define signature and context formation. These judgements ensure that all variables or constants

$\Gamma \vdash_{\Sigma, \preceq} K \text{ kind}$

$$\frac{}{\Gamma \vdash_{\Sigma, \preceq} \text{type kind}} \text{CANON_KIND_TYPE}$$

$$\frac{\Gamma \vdash_{\Sigma, \preceq} A \text{ type} \quad \Gamma, x:A \vdash_{\Sigma, \preceq} K \text{ kind}}{\Gamma \vdash_{\Sigma, \preceq} \Pi x:A. K \text{ kind}} \text{CANON_KIND_PI}$$

$\Gamma \vdash_{\Sigma, \preceq} A \text{ type}$

$$\frac{\Gamma \vdash_{\Sigma, \preceq} P \Rightarrow \text{type}}{\Gamma \vdash_{\Sigma, \preceq} P \text{ type}} \text{CANON_FAM_ATOM}$$

$$\frac{\Gamma \vdash_{\Sigma, \preceq} A_2 \text{ type} \quad \Gamma, x:A_2 \vdash_{\Sigma, \preceq} A \text{ type} \quad A_2 \preceq A}{\Gamma \vdash_{\Sigma, \preceq} \Pi x:A_2. A \text{ type}} \text{CANON_FAM_PI}$$

$\Gamma \vdash_{\Sigma, \preceq} P \Rightarrow K$

$$\frac{a : K \text{ in } \Sigma}{\Gamma \vdash_{\Sigma, \preceq} a \Rightarrow K} \text{ATOM_FAM_CONST}$$

$$\frac{\Gamma \vdash_{\Sigma, \preceq} P_1 \Rightarrow \Pi x:A_2. K_1 \quad \Gamma \vdash_{\Sigma, \preceq} M_2 \Leftarrow A_2 \quad [M_2/x]_{A_2}^k K_1 = K}{\Gamma \vdash_{\Sigma, \preceq} P_1 M_2 \Rightarrow K} \text{ATOM_FAM_APP}$$

$\Gamma \vdash_{\Sigma, \preceq} M \Leftarrow A$

$$\frac{\Gamma \vdash_{\Sigma, \preceq} R \Rightarrow P}{\Gamma \vdash_{\Sigma, \preceq} R \Leftarrow P} \text{CANON_TERM_ATOM}$$

$$\frac{\Gamma, x:A_2 \vdash_{\Sigma, \preceq} M \Leftarrow A}{\Gamma \vdash_{\Sigma, \preceq} \lambda x. M \Leftarrow \Pi x:A_2. A} \text{CANON_TERM_LAM}$$

$\Gamma \vdash_{\Sigma, \preceq} R \Rightarrow A$

$$\frac{x : A \text{ in } \Gamma}{\Gamma \vdash_{\Sigma, \preceq} x \Rightarrow A} \text{ATOM_TERM_VAR} \quad \frac{c : A \text{ in } \Sigma}{\Gamma \vdash_{\Sigma, \preceq} c \Rightarrow A} \text{ATOM_TERM_CONST}$$

$$\frac{\Gamma \vdash_{\Sigma, \preceq} R_1 \Rightarrow \Pi x:A_2. A_1 \quad \Gamma \vdash_{\Sigma, \preceq} M_2 \Leftarrow A_2 \quad [M_2/x]_{A_2}^a A_1 = A}{\Gamma \vdash_{\Sigma, \preceq} R_1 M_2 \Rightarrow A} \text{ATOM_TERM_APP}$$

Fig. 3. LF Formation Judgements

Simple Types $\alpha ::= a \mid \alpha_1 \rightarrow \alpha_2$

$(A)^- = \alpha$

$$\frac{}{(a)^- = a} \quad \frac{(P)^- = \alpha}{(PM)^- = \alpha} \quad \frac{(A_2)^- = \alpha_2 \quad (A)^- = \alpha}{(\Pi x:A_2. A)^- = \alpha_2 \rightarrow \alpha}$$

Fig. 4. Erasure to Simple Types

$$\boxed{[M_0/x_0]_{\alpha_0}^k K = K'}$$

$$\frac{}{[M_0/x_0]_{\alpha_0}^k \mathbf{type} = \mathbf{type}} \text{SUBST_K_TYPE} \quad \frac{[M_0/x_0]_{\alpha_0}^a A = A' \quad [M_0/x_0]_{\alpha_0}^k K = K'}{[M_0/x_0]_{\alpha_0}^k \Pi x:A. K = \Pi x:A'. K'} \text{SUBST_K_PI}$$

$$\boxed{[M_0/x_0]_{\alpha_0}^a A = A'}$$

$$\frac{[M_0/x_0]_{\alpha_0}^p P = P'}{[M_0/x_0]_{\alpha_0}^a P = P'} \text{SUBST_A_P} \quad \frac{[M_0/x_0]_{\alpha_0}^a A_2 = A'_2 \quad [M_0/x_0]_{\alpha_0}^a A = A'}{[M_0/x_0]_{\alpha_0}^a \Pi x:A_2. A = \Pi x:A'_2. A'} \text{SUBST_A_PI}$$

$$\boxed{[M_0/x_0]_{\alpha_0}^p P = P'}$$

$$\frac{}{[M_0/x_0]_{\alpha_0}^p \mathbf{a} = \mathbf{a}} \text{SUBST_P_CONST} \quad \frac{[M_0/x_0]_{\alpha_0}^p P = P' \quad [M_0/x_0]_{\alpha_0}^m M = M'}{[M_0/x_0]_{\alpha_0}^p P M = P' M'} \text{SUBST_P_APP}$$

$$\boxed{[M_0/x_0]_{\alpha_0}^m M = M'}$$

$$\frac{[M_0/x_0]_{\alpha_0}^r R = M' : \alpha'}{[M_0/x_0]_{\alpha_0}^m R = M'} \text{SUBST_M_RH} \quad \frac{[M_0/x_0]_{\alpha_0}^r R = R'}{[M_0/x_0]_{\alpha_0}^m R = R'} \text{SUBST_M_R}$$

$$\frac{[M_0/x_0]_{\alpha_0}^m M = M'}{[M_0/x_0]_{\alpha_0}^m \lambda x. M = \lambda x. M'} \text{SUBST_M_LAM}$$

$$\boxed{[M_0/x_0]_{\alpha_0}^r R = M' : \alpha}$$

$$\frac{}{[M_0/x_0]_{\alpha_0}^r x_0 = M_0 : \alpha_0} \text{SUBST_RH_VAR}$$

$$\frac{[M_0/x_0]_{\alpha_0}^r R_1 = \lambda x. M'_1 : \alpha_2 \rightarrow \alpha \quad [M_0/x_0]_{\alpha_0}^m M_2 = M'_2 \quad [M'_2/x]_{\alpha_2}^m M'_1 = M'}{[M_0/x_0]_{\alpha_0}^r R_1 M_2 = M' : \alpha} \text{SUBST_RH_APP}$$

$$\boxed{[M_0/x_0]_{\alpha_0}^r R = R'}$$

$$\frac{x \# x_0}{[M_0/x_0]_{\alpha_0}^r x = x} \text{SUBST_R_VAR} \quad \frac{}{[M_0/x_0]_{\alpha_0}^r c = c} \text{SUBST_R_CONST}$$

$$\frac{[M_0/x_0]_{\alpha_0}^r R_1 = R'_1 \quad [M_0/x_0]_{\alpha_0}^m M_2 = M'_2}{[M_0/x_0]_{\alpha_0}^r R_1 M_2 = R'_1 M'_2} \text{SUBST_R_APP}$$

$$\boxed{[M_0/x_0]_{\alpha_0}^c \Gamma = \Gamma'}$$

$$\frac{}{[M_0/x_0]_{\alpha_0}^c \cdot = \cdot} \text{SUBST_C_EMPTY}$$

$$\frac{x_0 \# x \quad x \# M_0 \quad [M_0/x_0]_{\alpha_0}^c \Gamma = \Gamma' \quad [M_0/x_0]_{\alpha_0}^a A = A'}{[M_0/x_0]_{\alpha_0}^c \Gamma, x : A = \Gamma', x : A'} \text{SUBST_C_TERM}$$

Fig. 5. LF Hereditary Substitution

declared in a context or signature are distinct and that all classifiers are well-formed in the preceding declarations. The judgements $x\#\Gamma$, $c\#\Sigma$, and $a\#\Sigma$ assert that the variable or constant is not declared in the context or signature; we omit their inductive definitions. Note that the context formation judgement presupposes that the signature is well-formed, rather than checking signature formation at the leaves.

We use the term *world* to refer to a set of well-formed contexts:

Definition 2.1 (World)

Given Σ and \preceq such that $\vdash_{\preceq} \Sigma \text{ sig}$, a world \mathcal{W} is a set containing contexts Γ for which the judgement $\vdash_{\Sigma, \preceq} \Gamma \text{ ctx}$ is derivable.

Kind, Family, and Term Formation. The judgements in Figure 3 define the formation of kinds, type families, and terms. We now call attention to several aspects of these rules:

- All of the formation judgements presuppose that $\vdash_{\preceq} \Sigma \text{ sig}$ and $\vdash_{\Sigma, \preceq} \Gamma \text{ ctx}$. In the rules for binding forms, the presuppositions of the judgements and the premises of the rules always entail that the extended context in the premise of the rule is well-formed.
- The judgement $\Gamma \vdash M \Leftarrow A$ presupposes that the type A is well-formed in Γ . In contrast, the judgement $\Gamma \vdash R \Rightarrow A$ ensures that the type A is well-formed. These modes correspond to a bidirectional operational interpretation: a canonical term is checked against a type, whereas a type is synthesized from an atomic term. The atomic type family judgement also synthesizes its kind, while the remaining judgements check that a type or kind is well-formed. The direction of the arrow in the judgement form serves as a mnemonic for the flow of information: in $M \Leftarrow A$, information in the type is used to check the term; in $R \Rightarrow A$, information in the term is used to synthesize the type. Because λ -abstractions are always checked against a known type, they do not require a type annotation for the bound variable.
- Variables and constants must be declared in the context or signature to be well-formed. As the variable rule `ATOM_TERM_VAR` explicates, the notation $x:A$ in contexts Γ stands for hypothetical assumptions of $\vdash_{\Sigma, \preceq} x \Rightarrow A$. Consequently, the usual substitution principle for the hypothetical judgement permits the substitution of a derivation of $R \Rightarrow A$ for an assumption $x \Rightarrow A$. In contrast, it requires further justification to substitute a derivation of $M \Leftarrow A$ for an assumption $x \Rightarrow A$; this justification is provided by hereditary substitution. Note that we do not consider substitutions for constants.
- The subordination relation parameter to the judgements is used only in the premise of the rule `CANON_FAM_PI`; this premise ensures that for any well-formed type $\Pi x:A_2. A$, the relation $A_2 \preceq A$ holds. When \preceq is taken to be the complete relation, this premise is always satisfiable. The rationale for this premise is described in Section 2.4.
- The rules `ATOM_TERM_APP` and `ATOM_FAM_APP` have hereditary substitution premises that compute the result type and kind of an application.

- In the rule `CANON_TERM_ATOM`, the syntactic restriction of the classifier to a \mathbf{P} , rather than any \mathbf{A} , ensures that canonical forms are η -long. For example, the judgement $\mathbf{f} : (\mathbf{a} \rightarrow \mathbf{a}) \vdash \mathbf{f} \Leftarrow (\mathbf{a} \rightarrow \mathbf{a})$ is not derivable: $(\mathbf{a} \rightarrow \mathbf{a})$ is a Π -type, so `CANON_TERM_ATOM` cannot be applied, and the variable \mathbf{f} is not a λ -abstraction, so neither can `CANON_TERM_LAM`. However, in a signature containing the declaration $\mathbf{a} : \mathbf{type}$, the canonicity of the η -expansion of \mathbf{f} can be derived as follows (let Γ abbreviate the context $\mathbf{f} : (\mathbf{a} \rightarrow \mathbf{a}), \mathbf{x} : \mathbf{a}$):

$$\frac{\frac{\frac{\Gamma \vdash \mathbf{f} \Rightarrow (\mathbf{a} \rightarrow \mathbf{a})}{\Gamma \vdash \mathbf{f} \Rightarrow (\mathbf{a} \rightarrow \mathbf{a})} \quad \frac{\frac{\Gamma \vdash \mathbf{x} \Rightarrow \mathbf{a}}{\Gamma \vdash \mathbf{x} \Leftarrow \mathbf{a}} \text{CTA} \quad \frac{\vdots}{[\mathbf{x}/_]\mathbf{a} = \mathbf{a}}}{\frac{\Gamma \vdash \mathbf{f} \mathbf{x} \Rightarrow \mathbf{a}}{\Gamma \vdash \mathbf{f} \mathbf{x} \Leftarrow \mathbf{a}} \text{CTA}}{\mathbf{f} : (\mathbf{a} \rightarrow \mathbf{a}) \vdash \lambda \mathbf{x}. \mathbf{f} \mathbf{x} \Leftarrow (\mathbf{a} \rightarrow \mathbf{a})} .$$

Both inferences labeled CTA for `CANON_TERM_ATOM` occur at base type, as required.

These formation judgements give a direct inductive definition of the canonical forms of LF. Consequently, rule induction for these judgements may be used to reason about canonical forms.

Hereditary Substitution. Next, we define hereditary substitution, which computes the canonical result of substituting one canonical form into another. The hereditary substitution judgements are defined in Figure 5. Hereditary substitution is defined on α -equivalence classes of expressions; by our notational convention, a bound variable \mathbf{x} is always distinct from both the variable being substituted for and the term being substituted for it (\mathbf{x}_0 and \mathbf{M}_0 in the rules). Written in this relational style, a standard substitution judgement $[\mathbf{M}_0/\mathbf{x}]\mathbf{M} = \mathbf{M}'$ would have four parameters: the term being substituted (\mathbf{M}_0), the variable being substituted for (\mathbf{x}), the term being substituted into (\mathbf{M}), and the result of the substitution (\mathbf{M}'). The hereditary substitution judgement $[\mathbf{M}_0/\mathbf{x}]_{\alpha_0}^m \mathbf{M} = \mathbf{M}'$ adds an extra parameter α_0 , which is the simple type of the substituted term \mathbf{M}_0 (the superscript m is simply notation for differentiating the judgements for the various syntactic classes from one another). The syntax of simple types α is defined in Figure 4. The simple type guides the process of hereditary substitution: because of the simple type parameter to the judgement, it is decidable whether or not a hereditary substitution exists even when the terms involved are ill-formed (see *Theorem 2.4* below).

The key hereditary substitution judgement is $[\mathbf{M}_0/\mathbf{x}]_{\alpha_0}^r \mathbf{R} = \mathbf{M}' : \alpha'$. This judgement computes the canonical result of substituting a canonical form \mathbf{M}_0 into an atomic term \mathbf{R} whose head is the variable \mathbf{x} ; it computes a new canonical term \mathbf{M}' and its simple type α' . The key rule is `SUBST_RH_APP`: when the substitution into the function position of an application yields a λ -abstraction, this rule continues by hereditarily substituting the argument \mathbf{M}'_2 into the body of the function. This is the rule where the simple type is used to guide the process of hereditary substitution: the simple type resulting from the first premise must have the form $\alpha_2 \rightarrow \alpha$, and the simple-type input to the third premise is α_2 (which, on well-typed terms, is

the simple type of M'_2) rather than α_0 . We prove below that the simple type α_2 in this final premise is always smaller than the input simple type α_0 ; this fact is used to establish decidability of hereditary substitution and to prove that hereditary substitutions exist under the appropriate typing premises.

The other hereditary substitution judgement on atomic terms, $[M_0/x]_{\alpha_0}^r R = R'$, is derivable only when the variable x is not the head variable of R ; it computes another atomic term compositionally. The remaining hereditary substitution judgements are defined compositionally as well. Because there is no single scope in which a variable declared in Γ can be renamed, the premises of the rule `SUBST_C_TERM`, which defines hereditary substitution into a context, insist that the variable in the context does not interfere with the substitution. The auxiliary judgement $x \# M$ holds when x is not free in the term M ; we elide its standard inductive definition. Intuitively, the premises of this rule ensure that the hereditary substitution into a context Γ is not defined when either the variable x_0 is declared in Γ or when a free variable in M_0 would be captured by a declaration in Γ .

The judgement $(A)^- = \alpha$ in Figure 4 defines an erasure relation that computes a simple type α from a dependent type A . Observe (by induction on the structure of type families) that for all A , there exists a unique α such that $(A)^- = \alpha$; this justifies using function notation for $(A)^-$. As a convenience, we write $[M_0/x]_{A_0}^m M = M'$ to mean $[M_0/x]_{\alpha_0}^m M = M'$ where $(A_0)^- = \alpha_0$; we adopt the analogous convention for the other syntactic categories. This notational convention is used by the hereditary substitution premises of the formation rules.

Notation. In the remainder of this article, we adopt several additional notational conveniences. We elide the signature Σ and the subordination relation \preceq parameters to the formation judgements when they are clear from context, as they are invariant throughout a derivation. To make the default instantiation of the parameters clear in a particular context, we will say that we work in $\text{LF}[\Sigma, \preceq]$. Also, we abbreviate a hypothetical judgement in the empty context by eliding the turnstile, writing, for example, $R \Rightarrow A$ for $\cdot \vdash R \Rightarrow A$.

2.3 Metatheory of Canonical LF

In this section, we prove two major theorems about Canonical LF. First, we prove decidability results for hereditary substitution and type checking. Next, we prove that hereditary substitutions exist and preserve types when the terms involved are well-formed and the types align in the appropriate way. Watkins *et al.* (2002) discuss these theorems and their proofs in more detail.

The remainder of this article does not require understanding the proofs of these two theorems about Canonical LF, so we provide proof sketches only as a reference for interested readers. However, the adequacy proofs below do make use of these results. A reader who is not interested in the metatheoretic development of LF may skip this section on first reading and refer back to the theorem statements as necessary to understand the subsequent adequacy proofs.

As a notational convenience in the following theorem statements, we use the word

“expression” to refer generically to any of the syntactic classes of Canonical LF. We also write E as a general metavariable for any syntactic category K, A, P, M, R, Γ , and we use e in $\{k, a, p, m, r, c\}$ for the corresponding tag on the hereditary substitution judgement. We write $x \# E$ to mean that the variable x does not occur free in the expression E .

The staging of lemmas leading up to the main results is somewhat intricate, and in some cases it is desirable to strengthen a theorem statement so that it holds even when certain subjects are not well-formed. In particular, we will state some theorems about formation judgements whose contexts are not known to be well-formed. However, even in these cases, we tacitly assume that all expressions have correct variable scoping. Specifically, when we write a context $x_1 : A_1, \dots, x_n : A_n$, we assume that all variables x_i are distinct and that the free variables of A_i are a subset of $\{x_1, \dots, x_{i-1}\}$. Additionally, we only write a judgement form such as $\Gamma \vdash M \Leftarrow A$ when all free variables of M and A are declared in Γ . Observe that, given a hereditary substitution $[M_0/x_0]_{A_0}^e E = E'$, the free variables of E' , written $FV(E')$, are $(FV(E) - \{x_0\}) \cup FV(M_0)$.

In the following theorem statements, when we leave the signature and subordination relation parameters to a theorem statement implicit, we tacitly assume a subordination relation \preceq and signature Σ such that $\vdash_{\preceq} \Sigma$ **sig**.

2.3.1 Decidability

Lemma 2.2 (Head Substitution Size)

If $[M_0/x_0]_{\alpha_0}^r R = M' : \alpha$ then α is a subexpression of α_0 .

Proof

By induction on the derivation of $[M_0/x_0]_{\alpha_0}^r R = M' : \alpha$. In the case for `SUBST_RH_VAR`, α_0 and α are identical. In the case for `SUBST_RH_APP`, the inductive hypothesis gives that $\alpha_2 \rightarrow \alpha$ is a subexpression of α_0 , so α is as well. \square

Lemma 2.3 (Uniqueness of Substitution and Synthesis)

1. If $[M_0/x_0]_{\alpha_0}^r R = R'$ and $[M_0/x_0]_{\alpha_0}^r R = M' : \alpha'$ then false.
2. For any E in $\{K, A, P, M, R\}$, if $[M_0/x_0]_{\alpha_0}^e E = E'$ and $[M_0/x_0]_{\alpha_0}^e E = E''$ then $E' = E''$.
3. If $\Gamma \vdash R \Rightarrow A$ and $\Gamma \vdash R \Rightarrow A'$ then $A = A'$.
4. If $\Gamma \vdash P \Rightarrow K$ and $\Gamma \vdash P \Rightarrow K'$ then $K = K'$.

Because we develop the metatheory of Canonical LF in an informal constructive logic, the following statements of decidability are sensible.

Theorem 2.4 (Decidability of Substitution)

1. For any E in $\{K, A, P, M, \Gamma\}$, given M_0, α_0, x_0 , either there exists an E' such that $[M_0/x_0]_{\alpha_0}^e E = E'$ or there is no such E' .
2. Given M_0, α_0, x_0 , and R , either there exists an R' such that $[M_0/x_0]_{\alpha_0}^r R = R'$ or there exist M' and α' such that $[M_0/x_0]_{\alpha_0}^r R = M' : \alpha'$ or there are no such R' and M' .

Proof

First, we prove part 2 and the clause of part 1 for canonical terms M by mutual lexicographic induction on $(A_0)^-$, the terms M and R being substituted into, and an order permitting inductive calls to the clause for atomic terms from the clause for canonical terms on the same simple type and term (to account for the silent injection from R to M). Then, we prove the clauses of part 1 for P , A , K , and Γ in that order, each by induction on the expression being substituted into, and each using the previous parts. \square

Theorem 2.5 (Decidability of Formation)

Assume that the subordination relation $A \preceq B$ is decidable.

1. For all Σ and \preceq , either $\vdash_{\preceq} \Sigma$ **sig** or not.

Assume \preceq and Σ such that that $\vdash_{\preceq} \Sigma$ **sig**.

2. For all Γ and K , either $\Gamma \vdash K$ **kind** or not.
3. For all Γ and A , either $\Gamma \vdash A$ **type** or not.
4. For all Γ and P , either there exists a K such that $\Gamma \vdash P \Rightarrow K$ or there is no such K .
5. For all Γ , M , and A , either $\Gamma \vdash M \Leftarrow A$ or not.
6. For all Γ and R , either there exists an A such that $\Gamma \vdash R \Rightarrow A$ or there is no such A .
7. For all Γ , either Γ **ctx** or not.

Proof

First, we prove parts 5 and 6 by mutual lexicographic induction on first the term and second an order permitting inductive calls to the clause for atomic terms from the clause for canonical terms on the same term. Then, we prove parts 4, 3, 2, 7, and 1, each by induction on the expression being judged well-formed. \square

2.3.2 Substitution Theorem

We now prove that under the appropriate typing conditions hereditary substitutions exist and preserve types. We write $\Gamma \subseteq \Gamma'$ iff Γ can be obtained from Γ' by removing zero or more declarations; we use the notation $\Sigma \subseteq \Sigma'$ analogously. Viewing subordination relations as sets of pairs of types, we write $\preceq \subseteq \preceq'$ for the usual subset order.

Lemma 2.6 (Weakening of Signature and Context)

Assume $\Gamma \subseteq \Gamma'$, $\Sigma \subseteq \Sigma'$ and $\preceq \subseteq \preceq'$.

1. For all five formation judgements \mathcal{J} , if $\Gamma \vdash_{\Sigma, \preceq} \mathcal{J}$ then $\Gamma' \vdash_{\Sigma', \preceq'} \mathcal{J}$.
2. If $\vdash_{\Sigma, \preceq} \Gamma$ **ctx** then $\vdash_{\Sigma', \preceq'} \Gamma$ **ctx**.
3. If $\vdash_{\preceq} \Sigma$ **sig** then $\vdash_{\preceq'} \Sigma$ **sig**.

Lemma 2.7 (Exchange)

For all five formation judgements, if $\Gamma, x:A, y:B, \Gamma' \vdash \mathcal{J}$ then $\Gamma, y:B, x:A, \Gamma' \vdash \mathcal{J}$.

Exchange for the signature Σ also holds, but we do not require it in this article.

Lemma 2.8 (Vacuous Substitutions)

For all M_0 , x_0 , A_0 , and E among $\{K, A, P, M, R\}$, if $x_0 \# E$ then $[M_0/x_0]_{A_0}^e E = E$.

Lemma 2.9 (Erasure is Invariant Under Substitution)

For all E in $\{A, P\}$, if $[M_0/x_0]_{\alpha_0}^e E = E'$ then $(E)^- = (E')^-$.

Ordinary substitutions enjoy the following composition property:

$$[e_0/x_0][e_2/x]e_1 = [[e_0/x_0]e_2/x][e_0/x_0]e_1$$

assuming that $x \# x_0$ and $x \# e_0$. The following lemma establishes an analogous property for hereditary substitutions. In addition to showing that the results of the two substitutions are equal, it shows that the outer two substitutions are defined if the inner three are.

Lemma 2.10 (Composition of Substitution)

Assume $x \# x_0$ and $x \# M_0$.

1. For all E in $\{K, A, P, M, R\}$, if $[M_0/x_0]_{\alpha_0}^m M_2 = M'_2$ and $[M_2/x]_{\alpha_2}^e E_1 = E$ and $[M_0/x_0]_{\alpha_0}^e E_1 = E'_1$ then there exists an E' such that $[M_0/x_0]_{\alpha_0}^e E = E'$ and $[M'_2/x]_{\alpha_2}^e E'_1 = E'$.
2. If $[M_0/x_0]_{\alpha_0}^m M_2 = M'_2$ and $[M_2/x]_{\alpha_2}^r R_1 = M : \alpha$ and $[M_0/x_0]_{\alpha_0}^r R_1 = R'_1$ then there exists an M' such that $[M_0/x_0]_{\alpha_0}^m M = M'$ and $[M'_2/x]_{\alpha_2}^r R'_1 = M' : \alpha$.
3. If $[M_0/x_0]_{\alpha_0}^m M_2 = M'_2$ and $[M_2/x]_{\alpha_2}^r R_1 = R$ and $[M_0/x_0]_{\alpha_0}^r R_1 = M'_1 : \alpha$ then there exists an M' such that $[M_0/x_0]_{\alpha_0}^r R = M' : \alpha$ and $[M'_2/x]_{\alpha_2}^m M'_1 = M'$.

The additional two composition properties for atomic terms R apply when the head variable of R is x or x_0 .

Proof

Define $size(a) = 1$ and $size(\alpha_1 \rightarrow \alpha_2) = 1 + size(\alpha_1) + size(\alpha_2)$. The proof is by mutual lexicographic induction on first $size(\alpha_0) + size(\alpha_2)$ and then the derivation of the substitution of M_2 . The size metric is necessary to justify an inductive call in which α_0 and α_2 are swapped. \square

Theorem 2.11 (Substitution)

Assume Σ and \preceq such that $\vdash_{\preceq} \Sigma \text{ sig}$. Assume $\Gamma_L, x_0 : A_0, \Gamma_R \text{ ctx}$ and $\Gamma_L \vdash M_0 \Leftarrow A_0$. Further, assume that subordination relationships $A \preceq B$ are preserved by hereditary substitution into A and B . Then:

1. There exists a Γ'_R such that $[M_0/x_0]_{A_0}^c \Gamma_R = \Gamma'_R$ and $\Gamma_L, \Gamma'_R \text{ ctx}$.
2. If $\Gamma_L, x_0 : A_0, \Gamma_R \vdash K \text{ kind}$ then there exists a K' such that $[M_0/x_0]_{A_0}^k K = K'$ and $\Gamma_L, \Gamma'_R \vdash K' \text{ kind}$.
3. If $\Gamma_L, x_0 : A_0, \Gamma_R \vdash A \text{ type}$ then there exists an A' such that $[M_0/x_0]_{A_0}^a A = A'$ and $\Gamma_L, \Gamma'_R \vdash A' \text{ type}$.
4. If $\Gamma_L, x_0 : A_0, \Gamma_R \vdash A \text{ type}$ and $\Gamma_L, x_0 : A_0, \Gamma_R \vdash M \Leftarrow A$ then there exist A' and M' such that $[M_0/x_0]_{A_0}^a A = A'$ and $[M_0/x_0]_{A_0}^m M = M'$ and $\Gamma_L, \Gamma'_R \vdash M' \Leftarrow A'$.

Proof

To prove this substitution theorem, we strengthen the theorem statements so that they work over contexts and types that are not necessarily well-formed. By strengthening the theorem in this way, we may prove the clause for terms without yet knowing that substitutions into type families preserve well-formedness. For conciseness, in the following theorem statements we leave implicit the existential quantification of the results of hereditary substitution. We first prove the following two clauses:

- If $\Gamma_L, \mathbf{x}_0 : \mathbf{A}_0, \Gamma_R \vdash \mathbf{M} \Leftarrow \mathbf{A}$ and $\Gamma_L \vdash \mathbf{M}_0 \Leftarrow \mathbf{A}_0$ and $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^c \Gamma_R = \Gamma'_R$ and $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^a \mathbf{A} = \mathbf{A}'$ then $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^m \mathbf{M} = \mathbf{M}'$ and $\Gamma_L, \Gamma'_R \vdash \mathbf{M}' \Leftarrow \mathbf{A}'$.
- If $\Gamma_L, \mathbf{x}_0 : \mathbf{A}_0, \Gamma_R \vdash \mathbf{R} \Rightarrow \mathbf{A}$ and $\Gamma_L \vdash \mathbf{M}_0 \Leftarrow \mathbf{A}_0$ and $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^c \Gamma_R = \Gamma'_R$ then $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^a \mathbf{A} = \mathbf{A}'$ and either $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^r \mathbf{R} = \mathbf{R}'$ and $\Gamma_L, \Gamma'_R \vdash \mathbf{R}' \Rightarrow \mathbf{A}'$ or $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^r \mathbf{R} = \mathbf{M}' : (\mathbf{A}')^-$ and $\Gamma_L, \Gamma'_R \vdash \mathbf{M}' \Leftarrow \mathbf{A}'$.

The proof is by mutual lexicographic induction on first the simple type $(\mathbf{A}_0)^-$ and then the derivations of $\Gamma_L, \mathbf{x}_0 : \mathbf{A}_0, \Gamma_R \vdash \mathbf{M} \Leftarrow \mathbf{A}$ and $\Gamma_L, \mathbf{x}_0 : \mathbf{A}_0, \Gamma_R \vdash \mathbf{R} \Rightarrow \mathbf{A}$.

Next, we prove the analogous statements for the remaining syntactic categories:

- If $\Gamma_L, \mathbf{x}_0 : \mathbf{A}_0, \Gamma_R \vdash \mathbf{P} \Rightarrow \mathbf{K}$ and $\Gamma_L \vdash \mathbf{M}_0 \Leftarrow \mathbf{A}_0$ and $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^c \Gamma_R = \Gamma'_R$ then $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^k \mathbf{K} = \mathbf{K}'$ and $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^p \mathbf{P} = \mathbf{P}'$ and $\Gamma_L, \Gamma'_R \vdash \mathbf{P}' \Rightarrow \mathbf{K}'$.
- If $\Gamma_L, \mathbf{x}_0 : \mathbf{A}_0, \Gamma_R \vdash \mathbf{A}$ type and $\Gamma_L \vdash \mathbf{M}_0 \Leftarrow \mathbf{A}_0$ and $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^c \Gamma_R = \Gamma'_R$ then $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^a \mathbf{A} = \mathbf{A}'$ and $\Gamma_L, \Gamma'_R \vdash \mathbf{A}'$ type.
- If $\Gamma_L, \mathbf{x}_0 : \mathbf{A}_0, \Gamma_R \vdash \mathbf{K}$ kind and $\Gamma_L \vdash \mathbf{M}_0 \Leftarrow \mathbf{A}_0$ and $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^c \Gamma_R = \Gamma'_R$ then $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^k \mathbf{K} = \mathbf{K}'$ and $\Gamma_L, \Gamma'_R \vdash \mathbf{K}'$ kind.
- If $\Gamma_L, \mathbf{x}_0 : \mathbf{A}_0, \Gamma_R$ ctx and $\Gamma_L \vdash \mathbf{M}_0 \Leftarrow \mathbf{A}_0$ then $[\mathbf{M}_0/\mathbf{x}_0]_{\mathbf{A}_0}^c \Gamma_R = \Gamma'_R$ and Γ_L, Γ'_R ctx.

Each part is proved in sequence by induction on the derivation of the expression being substituted into, using the previous parts. Then the clauses of the main theorem may be obtained as simple corollaries. \square

Lemma 2.12 (Regularity)

1. If Γ ctx and $\Gamma \vdash \mathbf{P} \Rightarrow \mathbf{K}$ then $\Gamma \vdash \mathbf{K}$ kind.
2. If Γ ctx and $\Gamma \vdash \mathbf{R} \Rightarrow \mathbf{A}$ then $\Gamma \vdash \mathbf{A}$ type.

This final lemma will be convenient in the adequacy proofs below; it asserts an n -ary hereditary substitution inversion principle for iterated Π -types and applications.

Lemma 2.13 (Iterated Hereditary Substitution Inversion)

1. If $[\mathbf{M}/\mathbf{x}]_{\mathbf{A}}^a \mathbf{B} = \Pi \mathbf{x}_1 : \mathbf{A}_1 \dots \Pi \mathbf{x}_n : \mathbf{A}_n. \mathbf{A}_{n+1}$ then there exist $\mathbf{A}'_1, \dots, \mathbf{A}'_{n+1}$ such that $\mathbf{B} = \Pi \mathbf{x}_1 : \mathbf{A}'_1 \dots \Pi \mathbf{x}_n : \mathbf{A}'_n. \mathbf{A}'_{n+1}$ and $[\mathbf{M}/\mathbf{x}]_{\mathbf{A}}^a \mathbf{A}'_i = \mathbf{A}_i$ for $1 \leq i \leq n+1$.
2. If $[\mathbf{M}/\mathbf{x}]_{\mathbf{A}}^a \Pi \mathbf{x}_1 : \mathbf{A}_1 \dots \Pi \mathbf{x}_n : \mathbf{A}_n. \mathbf{A}_{n+1} = \mathbf{B}$ then there exist $\mathbf{A}'_1, \dots, \mathbf{A}'_{n+1}$ such that $\mathbf{B} = \Pi \mathbf{x}_1 : \mathbf{A}'_1 \dots \Pi \mathbf{x}_n : \mathbf{A}'_n. \mathbf{A}'_{n+1}$ where $[\mathbf{M}/\mathbf{x}]_{\mathbf{A}}^a \mathbf{A}'_i = \mathbf{A}_i$ for $1 \leq i \leq n+1$.
3. If $[\mathbf{M}/\mathbf{x}]_{\mathbf{A}}^a \mathbf{B} = (\mathbf{a} \mathbf{M}_1 \dots \mathbf{M}_n)$ then there exist $\mathbf{M}'_1, \dots, \mathbf{M}'_n$ such that $\mathbf{B} = \mathbf{a} \mathbf{M}'_1 \dots \mathbf{M}'_n$ where $[\mathbf{M}/\mathbf{x}]_{\mathbf{A}}^m \mathbf{M}'_i = \mathbf{M}_i$ for $1 \leq i \leq n$.
4. If $[\mathbf{M}/\mathbf{x}]_{\mathbf{A}}^a (\mathbf{a} \mathbf{M}_1 \dots \mathbf{M}_n) = \mathbf{B}$ then there exist $\mathbf{M}'_1, \dots, \mathbf{M}'_n$ such that $\mathbf{B} = \mathbf{a} \mathbf{M}'_1 \dots \mathbf{M}'_n$ where $[\mathbf{M}/\mathbf{x}]_{\mathbf{A}}^m \mathbf{M}'_i = \mathbf{M}_i$ for $1 \leq i \leq n$.

The complete metatheory of Canonical LF includes an additional theorem witnessing that a variable $\mathbf{x} : \mathbf{A}$ can be η -expanded into a canonical form of type \mathbf{A} .

Whereas hereditary substitution corresponds to cut admissibility for a sequent calculus, this theorem corresponds to an identity principle (i.e., that $A \vdash A$ for any A). Because we do not require the identity theorem in this article, we refer the interested reader to Watkins *et al.* (2002) for details.

2.4 Subordination

Intuitively, a type family \mathbf{a} is subordinate to a type family \mathbf{b} if terms of type \mathbf{a} can appear in either terms of type \mathbf{b} or indices of the type family \mathbf{b} (Virga, 1999). The definition of subordination requires an auxiliary judgement identifying the *head* constant of a type family A .

$$\boxed{|A| = \mathbf{a}}$$

$$\frac{}{|\mathbf{a}| = \mathbf{a}} \quad \frac{|P| = \mathbf{a}}{|PM| = \mathbf{a}} \quad \frac{|A| = \mathbf{a}}{|\Pi x:A_2. A| = \mathbf{a}}$$

This judgement identifies the family-level constant at the head of the base type to which terms of type A contribute, where a term of base type contributes to that base type and a term of function type contributes to the base type that results from fully applying it to arguments. Observe (by induction over the structure of type families) that for all A there exists a unique family-level constant \mathbf{a} such that $|A| = \mathbf{a}$; this justifies using $|A|$ in a functional notation.

We now define the conditions under which a subordination relation is well-formed.

Definition 2.14 (Subordination Relation)

A *subordination relation* for a signature Σ is a binary relation \preceq between family-level constants declared in Σ , presented as a list of tuples, that satisfies the following properties:

1. Well-formedness: The judgement $\vdash_{\preceq} \Sigma \text{ sig}$ is derivable.
2. Index subordination: For all declarations $\mathbf{a} : \Pi x_1:A_1 \dots \Pi x_n:A_n. \text{type}$ in Σ , $|A_i| \preceq \mathbf{a}$ for all $1 \leq i \leq n$.
3. Reflexivity: For all \mathbf{a} declared in Σ , $\mathbf{a} \preceq \mathbf{a}$.
4. Transitivity: if $\mathbf{a}_1 \preceq \mathbf{a}_2$ and $\mathbf{a}_2 \preceq \mathbf{a}_3$ then $\mathbf{a}_1 \preceq \mathbf{a}_3$.

The first condition, $\vdash_{\preceq} \Sigma \text{ sig}$, implies that the subordination relation is permissive enough that the signature itself is well-formed. The second condition ensures that the types of the indices to a type family are subordinate to that family. The third and fourth ensure that subordination is a pre-order.

We tacitly extend the notation for subordination from constants to arbitrary type families by writing $A_1 \preceq A_2$ to mean $|A_1| \preceq |A_2|$. We write $\mathbf{a}_1 \not\preceq \mathbf{a}_2$ to mean that \mathbf{a}_1 is not subordinate to \mathbf{a}_2 ; we extend this notation to $A_1 \not\preceq A_2$ analogously to subordination. We also write $K \preceq A$, where $K = \Pi x_1:A_1 \dots \Pi x_n:A_n. \text{type}$, to mean $A_i \preceq A$ for all $1 \leq i \leq n$.

Theorem 2.5 assumes that the subordination relation $A \preceq B$ is decidable. This assumption is true of any subordination relation: $|A|$ maps every type family A to a unique constant \mathbf{a} ; and subordination on constants $\mathbf{a} \preceq \mathbf{b}$ is decidable because

we require a subordination relation to be presented as a list of pairs. *Theorem 2.11* assumes that $A \preceq B$ is preserved by substitution into A and B . This assumption is true of any subordination relation by the following lemma:

Lemma 2.15 (Head is Invariant Under Substitution)

For all M_0 , x_0 , α_0 , and E in $\{A, P\}$, if $[M_0/x_0]_{\alpha_0}^e E = E'$ then $|E| = |E'|$.

2.4.1 Transport of Canonical Forms

In this section, we show that the canonical forms of a type are unchanged by adding or removing canonical forms of other types that are not subordinate to it. We begin by defining the restriction of a context, signature, or subordination relation to a type, which removes all declarations or relationships that are not subordinate to that type.

Definition 2.16 (Context, Signature, and Subordination Relation Restriction)

$$\boxed{\Gamma|_{\mathbf{a}}^{\preceq} = \Gamma'}$$

$$\frac{}{\cdot|_{\mathbf{a}}^{\preceq} = \cdot} \quad \frac{\Gamma|_{\mathbf{a}}^{\preceq} = \Gamma' \quad \mathbf{A}_2 \preceq \mathbf{a}}{(\Gamma, \mathbf{x} : \mathbf{A}_2)|_{\mathbf{a}}^{\preceq} = \Gamma', \mathbf{x} : \mathbf{A}_2} \quad \frac{\Gamma|_{\mathbf{a}}^{\preceq} = \Gamma' \quad \mathbf{A}_2 \not\preceq \mathbf{a}}{(\Gamma, \mathbf{x} : \mathbf{A}_2)|_{\mathbf{a}}^{\preceq} = \Gamma'}$$

$$\boxed{\Sigma|_{\mathbf{a}}^{\preceq} = \Sigma'}$$

$$\frac{}{\cdot|_{\mathbf{a}}^{\preceq} = \cdot} \quad \frac{\Sigma|_{\mathbf{a}}^{\preceq} = \Sigma' \quad \mathbf{A}_2 \preceq \mathbf{a}}{(\Sigma, \mathbf{c} : \mathbf{A}_2)|_{\mathbf{a}}^{\preceq} = \Sigma', \mathbf{c} : \mathbf{A}_2} \quad \frac{\Sigma|_{\mathbf{a}}^{\preceq} = \Sigma' \quad \mathbf{A}_2 \not\preceq \mathbf{a}}{(\Sigma, \mathbf{c} : \mathbf{A}_2)|_{\mathbf{a}}^{\preceq} = \Sigma'}$$

$$\frac{\Sigma|_{\mathbf{a}}^{\preceq} = \Sigma' \quad \mathbf{a}' \preceq \mathbf{a}}{(\Sigma, \mathbf{a}' : \mathbf{K})|_{\mathbf{a}}^{\preceq} = \Sigma', \mathbf{a}' : \mathbf{K}} \quad \frac{\Sigma|_{\mathbf{a}}^{\preceq} = \Sigma' \quad \mathbf{a}' \not\preceq \mathbf{a}}{(\Sigma, \mathbf{a}' : \mathbf{K})|_{\mathbf{a}}^{\preceq} = \Sigma'}$$

Given a subordination relation \preceq , we define its restriction to a constant \mathbf{c} , written $\preceq|_{\mathbf{c}}$, by

$$\mathbf{a} \preceq|_{\mathbf{c}} \mathbf{b} \quad \text{iff} \quad \mathbf{a} \preceq \mathbf{b} \text{ and } \mathbf{b} \preceq \mathbf{c}.$$

Observe by induction on the structure of Γ that for all Γ , \preceq , \mathbf{a} , there exists a unique Γ' such that $\Gamma|_{\mathbf{a}}^{\preceq} = \Gamma'$; this justifies using $\Gamma|_{\mathbf{a}}^{\preceq}$ in function notation. A similar result holds for signature restriction. As a convenience, we write $\Gamma|_{\mathbf{A}}^{\preceq}$ for $\Gamma|_{\mathbf{a}}^{\preceq}$ where $|\mathbf{A}| = \mathbf{a}$; we adopt the analogous convention for signature restriction $\Sigma|_{\mathbf{A}}^{\preceq}$ and subordination restriction $\preceq|_{\mathbf{A}}$.

We now prove that canonical forms are invariant under subordination-based context, signature, and subordination relation restriction and extension:

Theorem 2.17 (Transport of Canonical Forms)

Assume that \preceq is a subordination relation for Σ .

1. For all \mathbf{A} , if $\preceq|_{\mathbf{A}} = \preceq'$ and $\Sigma|_{\mathbf{A}}^{\preceq} = \Sigma'$ then $\vdash_{\preceq'} \Sigma'$ sig.

Assume a type \mathbf{B} and let $\Sigma|_{\mathbf{B}}^{\preceq} = \Sigma'$ and $\preceq|_{\mathbf{B}} = \preceq'$.

2. If $A \preceq B$ and $\vdash_{\Sigma, \preceq} \Gamma \text{ ctx}$ and $\Gamma \vdash_{\Sigma, \preceq} A \text{ type}$ and $\Gamma|_B^{\preceq} = \Gamma'$ then $\Gamma \vdash_{\Sigma, \preceq} M \Leftarrow A$ iff $\Gamma' \vdash_{\Sigma', \preceq'} M \Leftarrow A$. The analogous statement for $\Gamma \vdash_{\Sigma, \preceq} R \Rightarrow A$ also holds.
3. If $A \preceq B$ and $\vdash_{\Sigma, \preceq} \Gamma \text{ ctx}$ and $\Gamma|_B^{\preceq} = \Gamma'$ then $\Gamma \vdash_{\Sigma, \preceq} A \text{ type}$ iff $\Gamma' \vdash_{\Sigma', \preceq'} A \text{ type}$. The analogous statement for $\Gamma \vdash_{\Sigma, \preceq} P \Rightarrow K$ also holds.
4. If $K \preceq B$ and $\vdash_{\Sigma, \preceq} \Gamma \text{ ctx}$ and $\Gamma|_B^{\preceq} = \Gamma'$ then $\Gamma \vdash_{\Sigma, \preceq} K \text{ kind}$ iff $\Gamma' \vdash_{\Sigma', \preceq'} K \text{ kind}$.
5. If $\vdash_{\Sigma, \preceq} \Gamma \text{ ctx}$ and $\Gamma|_B^{\preceq} = \Gamma'$ then $\vdash_{\Sigma', \preceq'} \Gamma' \text{ ctx}$.

Proof

The “if” directions of parts 2, 3, and 4 are consequences of weakening (*Lemma 2.6*). This leaves the “only if” directions for each of these three parts. The two clauses in part 2 can be proved by mutual induction on the derivations of $\Gamma \vdash_{\Sigma, \preceq} M \Leftarrow A$ and $\Gamma \vdash_{\Sigma, \preceq} R \Rightarrow A$. Then, to prove part 3, we prove the clause for atomic families by induction on the derivation of $\Gamma \vdash_{\Sigma, \preceq} P \Rightarrow K$, and then we prove the clause for types by induction on the derivation of $\Gamma \vdash_{\Sigma, \preceq} A \text{ type}$. Next, we prove part 4 by induction on the derivation of $\Gamma \vdash_{\Sigma, \preceq} K \text{ kind}$. The proofs of these parts use *Lemma 2.12* and *Lemma 2.15*. These parts also require the following lemma: if \preceq is a subordination relation for Σ and $\Gamma \vdash_{\Sigma, \preceq} P \Rightarrow K$ then $K \preceq P$; this lemma can be proved by induction on the formation derivation for P .

Next, we prove part 5 by induction on the derivation of $\vdash_{\Sigma, \preceq} \Gamma \text{ ctx}$. For part 1, we prove that if \preceq is a subordination relation for Σ and $\Sigma|_A^{\preceq} = \Sigma'$ and $\preceq|_A = \preceq'$ then $\vdash_{\preceq'} \Sigma' \text{ sig}$. The proof is by induction on the restriction derivation; it requires a lemma stating that a subordination relation for a signature $(\Sigma, c : A)$ is also a subordination relation for Σ (and similarly for $\Sigma, a : K$). \square

The proof of this theorem uses the fact that the relationship $A_2 \preceq A$ holds whenever a type $\Pi x:A_2. A$ is well-formed. Moreover, it is possible to construct a counterexample to this transport theorem using a type that does not satisfy this condition.

2.4.2 Strongest Subordination Relation

We define the *strongest subordination relation* for a signature Σ , written \preceq_Σ , to be the intersection of all subordination relations for Σ . If a signature Σ is well-formed without regard to subordination (i.e., it is well-formed in the complete relation), then there exists a unique strongest subordination relation for Σ . Intuitively, the strongest subordination relation establishes as few subordination relationships as possible, subject to the constraint that the signature itself be well-formed. Working with any subordination relation other than the strongest one causes *Theorem 2.17* to produce overapproximate results: extra subordination relationships prevent context and signature restriction from dropping assumptions that are, in fact, irrelevant. For example, if this theorem is applied with the complete subordination relation, the restrictions are the identity and the theorem yields no information.

In the remainder of this paper, we consider only the strongest subordination relation for each signature. By convention, when we instantiate a judgement with an LF signature Σ , we also implicitly take the subordination relation parameter to be \preceq_Σ . For example, we will write $\text{LF}[\Sigma]$ to mean $\text{LF}[\Sigma, \preceq_\Sigma]$. Additionally, we write $\Sigma|_A$ to mean $\Sigma|_A^{\preceq_\Sigma}$ for the strongest subordination relation \preceq_Σ . We will also

$$\begin{array}{l}
\tau ::= \text{unit} \mid \tau_1 \rightarrow \tau_2 \\
e ::= x \mid \langle \rangle \mid \lambda x:\tau. e \mid e_1 e_2
\end{array}$$

$\mathcal{X} \vdash e \text{ term}$

$$\begin{array}{c}
\frac{}{\mathcal{X}, x, \mathcal{X}' \vdash x \text{ term}} \text{TERM_VAR} \quad \frac{}{\mathcal{X} \vdash \langle \rangle \text{ term}} \text{TERM_EMPTY} \\
\frac{\mathcal{X}, x \vdash e \text{ term}}{\mathcal{X} \vdash \lambda x:\tau. e \text{ term}} \text{TERM_LAM} \quad \frac{\mathcal{X} \vdash e_1 \text{ term} \quad \mathcal{X} \vdash e_2 \text{ term}}{\mathcal{X} \vdash e_1 e_2 \text{ term}} \text{TERM_APP}
\end{array}$$

$[e_0/x]e = e'$

$$\begin{array}{c}
\frac{}{[e_0/x]x = e_0} \text{SUBST_VAR_MATCH} \quad \frac{y \# x}{[e_0/x]y = y} \text{SUBST_VAR_MISMATCH} \\
\frac{}{[e_0/x]\langle \rangle = \langle \rangle} \text{SUBST_EMPTY} \quad \frac{[e_0/x]e_1 = e'_1 \quad [e_0/x]e_2 = e'_2}{[e_0/x]e_1 e_2 = e'_1 e'_2} \text{SUBST_APP} \\
\frac{x \# y \quad y \# e_0 \quad [e_0/x]e = e'}{[e_0/x]\lambda y:\tau. e = \lambda y:\tau. e'} \text{SUBST_LAM}
\end{array}$$

Fig. 6. Syntax of the STLC

write $\Gamma|_{\mathbf{a}}$ when the signature and its strongest subordination relation are clear from context.

All of the signatures Σ we consider in this article have the property that

$$(\preceq_{\Sigma})|_{\mathbf{a}} = \preceq_{\Sigma'} \text{ where } \Sigma|_{\mathbf{a}}^{\preceq_{\Sigma}} = \Sigma'.$$

for any family \mathbf{a} declared in the signatures. That is, the restriction to a type of the strongest subordination relation for the signature is the strongest subordination relation for the restriction of the signature. We sometimes tacitly pass between these two subordination relations.

3 Mechanizing the Definition of the STLC in LF

3.1 Encoding of Syntax

We now begin mechanizing the simply typed λ -calculus (STLC) in LF. In this section, we encode the language's syntax. For reference, in Figure 6, we present the syntax of the STLC in informal mathematical notation. The metavariable \mathcal{X} ranges over comma-separated lists of distinct variables, which stand for assumptions of $x \text{ term}$; the hypothetical judgement $\mathcal{X} \vdash e \text{ term}$ is derivable when the free variables of the term e are contained in \mathcal{X} . This judgement will be used to state the adequacy theorems below. The judgement $[e_0/x]e = e'$ defines the standard notion of capture-avoiding substitution; we omit the standard definition of $x \# e$.

The LF signature defined in Figure 7 represents the syntax of the STLC. It

```

    tp : type
arrow : tp → tp → tp
unit  : tp

    tm : type
empty : tm
app   : tm → tm → tm
lam   : tp → (tm → tm) → tm

```

Fig. 7. LF Signature for the STLC Syntax

$\tau \text{ type} \gg M \Leftarrow \text{tp}$

$$\frac{}{\text{unit type} \gg \text{unit} \Leftarrow \text{tp}} \text{ENC_TP_UNIT}$$

$$\frac{\tau_1 \text{ type} \gg M_1 \Leftarrow \text{tp} \quad \tau_2 \text{ type} \gg M_2 \Leftarrow \text{tp}}{\tau_1 \rightarrow \tau_2 \text{ type} \gg \text{arrow } M_1 M_2 \Leftarrow \text{tp}} \text{ENC_TP_ARROW}$$

$\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}$

$$\frac{}{\cdot \text{ terms} \gg \cdot \text{ ctx}} \text{ENC_TERMS_NIL} \quad \frac{\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}}{\mathcal{X}, x \text{ terms} \gg \Gamma, x : \text{tm} \text{ ctx}} \text{ENC_TERMS_TERM}$$

$\mathcal{X} \vdash e \text{ term} \gg \Gamma \vdash M \Leftarrow \text{tm}$

$$\frac{}{\mathcal{X}, x, \mathcal{X}' \vdash x \text{ term} \gg \Gamma, x : \text{tm}, \Gamma' \vdash x \Leftarrow \text{tm}} \text{ENC_TM_VAR}$$

$$\frac{}{\mathcal{X} \vdash \langle \rangle \text{ term} \gg \Gamma \vdash \text{empty} \Leftarrow \text{tm}} \text{ENC_TM_EMPTY}$$

$$\frac{\tau \text{ type} \gg M_t \Leftarrow \text{tp} \quad \mathcal{X}, x \vdash e \text{ term} \gg \Gamma, x : \text{tm} \vdash M_e \Leftarrow \text{tm}}{\mathcal{X} \vdash (\lambda x : \tau. e) \text{ term} \gg \Gamma \vdash \text{lam } M_t (\lambda x. M_e) \Leftarrow \text{tm}} \text{ENC_TM_LAM}$$

$$\frac{\mathcal{X} \vdash e_1 \text{ term} \gg \Gamma \vdash M_1 \Leftarrow \text{tm} \quad \mathcal{X} \vdash e_2 \text{ term} \gg \Gamma \vdash M_2 \Leftarrow \text{tm}}{\mathcal{X} \vdash e_1 e_2 \text{ term} \gg \Gamma \vdash \text{app } M_1 M_2 \Leftarrow \text{tm}} \text{ENC_TM_APP}$$

Fig. 8. Encoding of STLC Syntax

declares an LF type for each syntactic category (**tp** for τ and **tm** for **e**) along with constants inhabiting those types with the representations of the language's types and terms. The informal syntax **unit** has no subexpressions in the informal grammar, so it is represented by an LF constant **unit** of type **tp**; the informal syntax $\tau_1 \rightarrow \tau_2$ has two type subexpressions, so it is represented by an LF constant **arrow** of type **tp** \rightarrow **tp** \rightarrow **tp**. The representation of terms uses a technique called *higher-order abstract syntax*: object-language variables are represented by LF variables; with this representation, the framework provides α -conversion and substitution for the object language. To build intuition, consider the following expressions and their intended representations:

<u>Expression</u>		<u>LF Representation</u>
<code>unit → unit</code>	\gg	<code>(arrow unit unit)</code>
<code>x y</code>	\gg	<code>(app x y)</code>
<code>λx:unit. x</code>	\gg	<code>(lam unit (λ x. x))</code>

The second and third examples illustrate higher-order abstract syntax. In the second example, object-language variables are translated to LF variables. The third example illustrates the representation of binding forms: in the informal syntax $\lambda x:\tau. e$, the variable x is bound in the body e ; in the LF representation, the body is represented by an LF function of type $(\mathbf{tm} \rightarrow \mathbf{tm})$ that binds the equivalent variable. This representation strategy is the reason for the higher-order type of the constant `lam`.

The judgements in Figure 8 formally define the encoding of the STLC into LF. We consider both object-language terms and LF terms up to α -equivalence, and by convention all bound variables are chosen fresh. The intended reading of these judgements is as follows:

- $\tau \text{ type} \gg M \Leftarrow \mathbf{tp}$ Encode an object-language type τ to an LF term M of LF type \mathbf{tp} in the empty context.
- $\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}$ Encode an object-language list of variables \mathcal{X} to an LF context declaring each of those variables to have type \mathbf{tm} .
- $\mathcal{X} \vdash e \text{ term} \gg \Gamma \vdash M \Leftarrow \mathbf{tm}$ Assuming $\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}$, encode an object-language term e with free variables in \mathcal{X} to an LF term M of type \mathbf{tm} in LF context Γ . This judgement has four parameters (e , \mathcal{X} , M , and Γ). The notation is meant to suggest the intended invariant that whenever the judgement $\mathcal{X} \vdash e \text{ term} \gg \Gamma \vdash M \Leftarrow \mathbf{tm}$ is derivable, so are $\mathcal{X} \vdash e \text{ term}$ and $\Gamma \vdash M \Leftarrow \mathbf{tm}$. This invariant is verified in the adequacy proof below.

The term encoding rule `ENC_TM_LAM` illustrates higher-order abstract syntax: the encoding of the object language term e with free variables (\mathcal{X}, x) is an LF term M_e with free variables in the context $(\Gamma, x:\mathbf{tm})$; the variable x is then bound in $\lambda x. M_e$, which creates an LF term of type $\mathbf{tm} \rightarrow \mathbf{tm}$.

These encoding judgements clarify the fact that the LF representation of the object-language syntax is specified not just by the LF signature, but also by the LF contexts in which that signature is considered. For example, to know that \mathbf{tm} adequately represents object-language terms with free variables, it is necessary to know not just that the constants inhabiting \mathbf{tm} are `empty`, `app`, and `lam`, but also that contexts containing LF variables of type \mathbf{tm} are considered.

The strongest subordination relation for the signature Σ in Figure 7, written \preceq_Σ , is $\{\mathbf{tp} \preceq \mathbf{tp}, \mathbf{tp} \preceq \mathbf{tm}, \mathbf{tm} \preceq \mathbf{tm}\}$. Intuitively, STLC types appear in the syntax of STLC terms, but STLC terms do not appear in the syntax of STLC types.

3.2 Adequacy of Syntax Encodings

Adequacy is the correctness criterion for an encoding judgement; it establishes that the encoding is an isomorphism between the informal object-language entities and

their LF representation. We break the statement of adequacy into four parts: First, we show that the subjects of the encoding judgement are well-formed, which establishes that the encoding relates the stated object-language entities to LF terms of the appropriate type. Second, we show that the encoding judgement relates every informal entity to a unique LF term; third, we show that the encoding judgement relates every canonical form of the appropriate LF type to a unique informal entity. Because the single encoding judgement is functional in both directions, these functions are mutually inverse, yielding a bijection. Finally, we show that the encoding commutes with substitution, establishing compositionality.

For the remainder of this section, we work in $\text{LF}[\Sigma]$, where Σ stands for the LF signature defined in Figure 7.

3.2.1 Types

For parallelism with later adequacy statements, we use the notation τ **type** to mean that τ is a syntactically well-formed type. The adequacy theorem for types is degenerate: because types do not involve binding, no compositionality condition is necessary.

Theorem 3.1 (Adequacy for Types)

The encoding relation τ **type** \gg $M \Leftarrow \mathbf{tp}$ defines a bijection between the types τ and the LF terms M such that $M \Leftarrow \mathbf{tp}$:

1. If τ **type** \gg $M \Leftarrow \mathbf{tp}$ then τ **type** and $M \Leftarrow \mathbf{tp}$.
2. If τ **type**, there exists a unique LF term M such that τ **type** \gg $M \Leftarrow \mathbf{tp}$.
3. If $M \Leftarrow \mathbf{tp}$ then there exists a unique τ such that τ **type** \gg $M \Leftarrow \mathbf{tp}$.

The third part of this theorem is proved by induction on canonical forms, using a specialized inversion principle for the canonical forms of type \mathbf{tp} in the empty LF context:

Lemma 3.2 (Inversion of Canonical Forms of Type \mathbf{tp})

If \mathcal{D} derives $M \Leftarrow \mathbf{tp}$ then either

- $M = \mathbf{unit}$, or
- $M = \mathbf{arrow} M_1 M_2$, and $M_1 \Leftarrow \mathbf{tp}$ and $M_2 \Leftarrow \mathbf{tp}$ were derived as strict subderivations of \mathcal{D} .

This lemma permits a proof by induction on the derivation of $M \Leftarrow \mathbf{tp}$ to consider only these two cases and to appeal to induction on the strict subderivations—it is equivalent to a specialized induction principle for the canonical forms of type \mathbf{tp} in the empty context.

Proof

Characterizing the canonical forms of type \mathbf{tp} requires first characterizing the canonical forms of certain higher types:

1. There is no R such that $R \Rightarrow \Pi x_1:A_1 \dots \Pi x_n:A_n. \mathbf{tp}$ for $n \geq 3$.
2. If $R \Rightarrow \Pi x_1:A_1. \Pi x_2:A_2. \mathbf{tp}$ then $A_1 = A_2 = \mathbf{tp}$ and $R = \mathbf{arrow}$.

3. If $R \Rightarrow \Pi x:A_2. \mathbf{tp}$ then $A_2 = \mathbf{tp}$, $R = \mathbf{arrow} M_1$, and $M_1 \Leftarrow \mathbf{tp}$ was derived as a strict subderivation.

To prove the first part, assume there is such a term and then obtain a contradiction by induction on the given derivation. The proofs of the next two parts proceed by case analysis on the derivation, each using the previous part. Each proof uses *Lemma 2.13*. Then the overall lemma is proved by inverting the derivation of $M \Leftarrow \mathbf{tp}$ and showing, in each case, that M has the required form. The proof uses *Lemma 2.13* and part 3 above. \square

Using this lemma, we prove adequacy:

Proof of Theorem 3.1

1. To show: If $\tau \mathbf{type} \gg M \Leftarrow \mathbf{tp}$ then $\tau \mathbf{type}$ and $M \Leftarrow \mathbf{tp}$. We show one case.
In the case for `ENC_TP_ARROW`, assume as the inductive hypothesis that $\tau_1 \mathbf{type}$, $\tau_2 \mathbf{type}$, $M_1 \Leftarrow \mathbf{tp}$, and $M_2 \Leftarrow \mathbf{tp}$. Then $\tau_1 \rightarrow \tau_2 \mathbf{type}$ and the judgement $\mathbf{arrow} M_1 M_2 \Leftarrow \mathbf{tp}$ can be derived using the inductive hypotheses and the following rules: `ATOM_TERM_APP`, `ATOM_TERM_CONST`, `SUBST_A_PI`, `SUBST_A_P`, `SUBST_P_CONST`, and `CANON_TERM_ATOM`.
2. To show: For all τ , there exists a unique term M such that $\tau \mathbf{type} \gg M \Leftarrow \mathbf{tp}$.
The proof is by structural induction on τ . We show one case.
In the case for $\tau_1 \rightarrow \tau_2$, we assume that there exist unique M_1 and M_2 such that $\tau_1 \mathbf{type} \gg M_1 \Leftarrow \mathbf{tp}$ and $\tau_2 \mathbf{type} \gg M_2 \Leftarrow \mathbf{tp}$. We must show that there exists a unique M such that $\tau_1 \rightarrow \tau_2 \mathbf{type} \gg M \Leftarrow \mathbf{tp}$. To establish existence, take M to be $\mathbf{arrow} M_1 M_2$. Then using `ENC_TP_ARROW` on the inductive hypotheses proves that $\tau_1 \rightarrow \tau_2 \mathbf{type} \gg M \Leftarrow \mathbf{tp}$. To show uniqueness, assume some other M' such that $\tau_1 \rightarrow \tau_2 \mathbf{type} \gg M' \Leftarrow \mathbf{tp}$. By inversion, the only rule that can have applied to $\tau_1 \rightarrow \tau_2$ is `ENC_TP_ARROW`, so $M' = \mathbf{arrow} M'_1 M'_2$ where $\tau_1 \mathbf{type} \gg M'_1 \Leftarrow \mathbf{tp}$ and $\tau_2 \mathbf{type} \gg M'_2 \Leftarrow \mathbf{tp}$. But then the inductive hypotheses that M_1 and M_2 are unique show that $M_1 = M'_1$ and $M_2 = M'_2$. Therefore $\mathbf{arrow} M'_1 M'_2 = \mathbf{arrow} M_1 M_2$, establishing uniqueness.
3. To show: For all M such that $M \Leftarrow \mathbf{tp}$, there exists a unique τ such that $\tau \mathbf{type} \gg M \Leftarrow \mathbf{tp}$. The proof is by induction on the derivation of $M \Leftarrow \mathbf{tp}$.
Lemma 3.2 gives two cases to consider; we show the case for $M = \mathbf{arrow} M_1 M_2$, where $M_1 \Leftarrow \mathbf{tp}$ and $M_2 \Leftarrow \mathbf{tp}$ were derived as strict subderivations. To show: there exists a unique τ such that $\tau \mathbf{type} \gg \mathbf{arrow} M_1 M_2 \Leftarrow \mathbf{tp}$. By the inductive hypothesis applied to the subderivations, there exist unique τ_1 and τ_2 such that $\tau_1 \mathbf{type} \gg M_1 \Leftarrow \mathbf{tp}$ and $\tau_2 \mathbf{type} \gg M_2 \Leftarrow \mathbf{tp}$. To establish existence, take $\tau = \tau_1 \rightarrow \tau_2$; then `ENC_TP_ARROW` applied to the derivations from the inductive hypothesis shows that $\tau \mathbf{type} \gg \mathbf{arrow} M_1 M_2 \Leftarrow \mathbf{tp}$. To establish uniqueness, assume some other τ' such that $\tau' \mathbf{type} \gg \mathbf{arrow} M_1 M_2 \Leftarrow \mathbf{tp}$. By inversion, only the rule `ENC_TP_ARROW` could have applied, so $\tau' = \tau'_1 \rightarrow \tau'_2$ where $\tau'_1 \mathbf{type} \gg M_1 \Leftarrow \mathbf{tp}$ and $\tau'_2 \mathbf{type} \gg M_2 \Leftarrow \mathbf{tp}$. The inductive hypothesis that τ_1 and τ_2 are unique shows that $\tau_1 = \tau'_1$ and $\tau_2 = \tau'_2$, so $\tau_1 \rightarrow \tau_2 = \tau'_1 \rightarrow \tau'_2$, as we needed to show.

\square

3.2.2 Terms

The following lemma, analogous to the one for \mathbf{tp} above, gives an inversion principle for the canonical forms of type \mathbf{tm} :

Lemma 3.3 (Inversion of Canonical Forms of Type \mathbf{tm})

If \mathcal{X} terms $\gg \Gamma$ \mathbf{ctx} and $\Gamma \vdash M \Leftarrow \mathbf{tm}$ then one of the following holds:

- $M = x$ and $\Gamma = \Gamma_1, x : \mathbf{tm}, \Gamma_2$.
- $M = \mathbf{empty}$.
- $M = \mathbf{lam} M_t \lambda x. M_e$, where $\Gamma \vdash M_t \Leftarrow \mathbf{tp}$ and $\Gamma, x : \mathbf{tm} \vdash M_e \Leftarrow \mathbf{tm}$ were derived as strict subderivations.
- $M = \mathbf{app} M_1 M_2$, where $\Gamma \vdash M_1 \Leftarrow \mathbf{tm}$ and $\Gamma \vdash M_2 \Leftarrow \mathbf{tm}$ were derived as strict subderivations.

This lemma states that an LF term of type \mathbf{tm} in a context containing variables of type \mathbf{tm} is either a variable or a constant applied to arguments which may contain those variables—i.e., it states that the LF term is parametric in the variables in the context. If, hypothetically, LF had an unrestricted case-analysis construct for analyzing terms of type \mathbf{tm} , this lemma would not be true: a case-analysis of a variable would be an “exotic” canonical form of type \mathbf{tm} , violating adequacy. Extensions of LF with other types must take care to preserve this property; in Concurrent LF, certain connectives are confined to a monad so that they do not interfere with this style of higher-order representation (Watkins *et al.*, 2002).

Because types appear in the syntax of terms, adequacy for terms requires adequacy for types. However, types are adequately represented in the empty LF context, whereas the judgements in Figure 8 encode terms in LF contexts of the form $x_1 : \mathbf{tm}, \dots, x_n : \mathbf{tm}$. Thus, *Theorem 3.1* as stated in the previous section is not strong enough, as adequacy for terms requires that types remain adequate in these extended contexts. This requirement could fail—for example, if term contexts were of the form $x : \mathbf{tm}, \dots, u : \mathbf{tp}, \dots$, these contexts would induce additional canonical forms of type \mathbf{tp} that have no informal counterpart. Consequently, it is necessary to prove a lemma showing that types remain adequate in the contexts Γ such that \mathcal{X} terms $\gg \Gamma$ \mathbf{ctx} . This lemma is our first application of the transport of canonical forms theorem, *Theorem 2.17*.

Lemma 3.4 (Transport of Adequacy for Terms)

1. If \mathcal{X} terms $\gg \Gamma$ \mathbf{ctx} then Γ \mathbf{ctx} .
2. If \mathcal{X} terms $\gg \Gamma$ \mathbf{ctx} then $\cdot \vdash M_t \Leftarrow \mathbf{tp}$ iff $\Gamma \vdash M_t \Leftarrow \mathbf{tp}$.

Proof

The proof of the first part is a straightforward induction on the premise. The second part uses transport of canonical forms (*Theorem 2.17*). First, one application of *Theorem 2.17* shows that $\cdot \vdash_{\Sigma} M \Leftarrow \mathbf{tp}$ iff $\cdot \vdash_{\Sigma|\mathbf{tp}} M \Leftarrow \mathbf{tp}$. Next, because $\mathbf{tm} \not\leq \mathbf{tp}$, a simple induction shows that if \mathcal{X} terms $\gg \Gamma$ \mathbf{ctx} then $\Gamma|_{\mathbf{tp}}^{\leq \Sigma} = \cdot$. Then the proof is direct using *Theorem 2.17* and the first part. \square

We now prove the main adequacy result. The judgement $\mathcal{X} \vdash e$ term is used here to state the invariant on the encoding. When we claim uniqueness for an object-language or LF entity that involves binding, we mean uniqueness up to α -conversion.

Theorem 3.5 (Adequacy for Terms)

1. If $\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}$ and $\mathcal{X} \vdash e \text{ term} \gg \Gamma \vdash M \Leftarrow \text{tm}$ then $\mathcal{X} \vdash e \text{ term}$ and $\Gamma \vdash M \Leftarrow \text{tm}$.
2. If $\mathcal{X} \vdash e \text{ term}$ and $\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}$ then there exists a unique LF term M such that $\mathcal{X} \vdash e \text{ term} \gg \Gamma \vdash M \Leftarrow \text{tm}$.
3. If $\Gamma \vdash M \Leftarrow \text{tm}$ and $\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}$ then there exists a unique e such that $\mathcal{X} \vdash e \text{ term} \gg \Gamma \vdash M \Leftarrow \text{tm}$.

Proof

1. By rule induction on the second premise. The case for `ENC_TM_LAM` is the only one in which the conclusion does not follow immediately from the inductive hypotheses. In this case, by the inductive hypothesis, $\mathcal{X}, x \vdash e \text{ term}$ and $\Gamma, x : \text{tm} \vdash M_e \Leftarrow \text{tm}$. Then the first conclusion is immediate by `TERM_LAM`. To derive the second conclusion, we also need to know that $\Gamma \vdash M_t \Leftarrow \text{tp}$. By *Theorem 3.1* applied to the premise of the rule, $\cdot \vdash M_t \Leftarrow \text{tp}$. Then by *Lemma 3.4*, $\Gamma \vdash M_t \Leftarrow \text{tp}$.
2. By rule induction on $\mathcal{X} \vdash e \text{ term}$, we show that if $\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}$ then there exists a unique M such that $\mathcal{X} \vdash e \text{ term} \gg \Gamma \vdash M \Leftarrow \text{tm}$. We give the cases that involve binding:
 - Case for `TERM_VAR`. To show: if $\mathcal{X}_1, x, \mathcal{X}_2 \text{ terms} \gg \Gamma \text{ ctx}$ then there exists a unique LF term M such that $\mathcal{X} \vdash x \text{ term} \gg \Gamma \vdash M \Leftarrow \text{tm}$. By inversion, Γ must have the form $\Gamma_1, x : \text{tm}, \Gamma_2$. To establish existence, use `ENC_TM_VAR` to derive $\mathcal{X}_1, x, \mathcal{X}_2 \vdash x \text{ term} \gg \Gamma \vdash x \Leftarrow \text{tm}$. To establish uniqueness, assume some other M' such that $\mathcal{X}_1, x, \mathcal{X}_2 \vdash x \text{ term} \gg \Gamma \vdash M' \Leftarrow \text{tm}$; inversion on this derivation gives the result because only `ENC_TM_VAR` can have applied.
 - Case for `TERM_LAM`. Assume a derivation of $\mathcal{X} \text{ terms} \gg \Gamma \text{ ctx}$. Using the rule `ENC_TERMS_TERM`, we can derive $(\mathcal{X}, x) \text{ terms} \gg (\Gamma, x : \text{tm}) \text{ ctx}$. Then we can appeal to the inductive hypothesis to conclude that there exists a unique canonical form M_e such that $\mathcal{X}, x \vdash e \text{ term} \gg \Gamma, x : \text{tm} \vdash M_e \Leftarrow \text{tm}$. By *Theorem 3.1*, there exists a unique M_t such that $\tau \text{ type} \gg M_t \Leftarrow \text{tp}$. To show existence, take M to be `lam` $M_t \lambda x. M_e$; then we can derive the desired property by applying `ENC_TM_LAM` to the above encoding derivations. Uniqueness is immediate by inversion and the results above.
3. By induction on the derivation of $\Gamma \vdash M \Leftarrow \text{tm}$. *Lemma 3.3* gives four cases; we show those involving binding:
 - $M = x$ and $\Gamma = (\Gamma_1, x : \text{tm}, \Gamma_2)$.
By inversion on the derivation of $\mathcal{X} \text{ terms} \gg (\Gamma_1, x : \text{tm}, \Gamma_2) \text{ ctx}$, \mathcal{X} is $\mathcal{X}_1, x, \mathcal{X}_2$. Thus, we can take e to be x , which has the desired property by `ENC_TM_VAR`. Uniqueness is immediate by inversion.
 - $M = \text{lam } M_t \lambda x. M_e$, where $\Gamma \vdash M_t \Leftarrow \text{tp}$ and $\Gamma, x : \text{tm} \vdash M_e \Leftarrow \text{tm}$ were derived as subderivations.
We would like to appeal to *Theorem 3.1* on the derivation of $\Gamma \vdash M_t \Leftarrow \text{tp}$ to conclude that there exists a unique τ such that $\tau \text{ type} \gg M_t \Leftarrow \text{tp}$. Unfortunately, *Theorem 3.1* is only stated for typing derivations in the

empty context. Thus, we first apply *Lemma 3.4* to conclude $\cdot \vdash M_t \Leftarrow \mathfrak{tp}$; then we can appeal to *Theorem 3.1*.

Next, using the assumed derivation of \mathcal{X} terms $\gg \Gamma \text{ ctx}$, we can create a derivation of (\mathcal{X}, x) terms $\gg (\Gamma, x : \mathfrak{tm}) \text{ ctx}$ by `ENC_TERMS_TERM`. Then, by the inductive hypothesis on the second subderivation, there exists a unique e' such that $\mathcal{X}, x \vdash e' \text{ term} \gg \Gamma, x : \mathfrak{tm} \vdash M'_e \Leftarrow \mathfrak{tm}$. To show existence, we take e to be $\lambda x : \tau. e'$; `ENC_TM_LAM` applied to the derivations produced above shows that $\mathcal{X} \vdash \lambda x : \tau. e' \text{ term} \gg \Gamma \vdash \mathfrak{lam} M_t \lambda x. M'_e \Leftarrow \mathfrak{tm}$. Uniqueness is immediate by inversion and the uniqueness of e' and τ .

□

Next, we prove compositionality. The proof requires a lemma stating that a property similar to weakening holds for the encoding:

Lemma 3.6 (Weakening of the Encoding)

Assume $\mathcal{X}, \mathcal{X}'$ terms $\gg \Gamma, \Gamma' \text{ ctx}$ and $\mathcal{X}, x, \mathcal{X}'$ terms $\gg \Gamma, x : \mathfrak{tm}, \Gamma' \text{ ctx}$. Then if $\mathcal{X}, \mathcal{X}' \vdash e \text{ term} \gg \Gamma, \Gamma' \vdash M \Leftarrow \mathfrak{tm}$ then $\mathcal{X}, x, \mathcal{X}' \vdash e \text{ term} \gg \Gamma, x : \mathfrak{tm}, \Gamma' \vdash M \Leftarrow \mathfrak{tm}$.

Exchange and contraction also hold, but we do not require them in this development. Compositionality is essentially a substitution principle:

Theorem 3.7 (Compositionality for Terms)

Assume \mathcal{X} terms $\gg \Gamma \text{ ctx}$ and $\mathcal{X} \vdash e_2 \text{ term} \gg \Gamma \vdash M_2 \Leftarrow \mathfrak{tm}$.

If $\mathcal{X}, x, \mathcal{X}'$ terms $\gg \Gamma, x : \mathfrak{tm}, \Gamma' \text{ ctx}$ and $\mathcal{X}, x, \mathcal{X}' \vdash e \text{ term} \gg \Gamma, x : \mathfrak{tm}, \Gamma' \vdash M \Leftarrow \mathfrak{tm}$ and $[e_2/x]e = e'$ and $[M_2/x]_{\mathfrak{tm}}^m M = M'$ then $\mathcal{X}, \mathcal{X}' \vdash e' \text{ term} \gg \Gamma, \Gamma' \vdash M' \Leftarrow \mathfrak{tm}$.

Proof

Using *Lemma 3.6*, *Theorem 3.1*, and *Lemma 2.8*, we prove this theorem by induction on the derivation of $\mathcal{X}, x, \mathcal{X}' \vdash e \text{ term} \gg \Gamma, x : \mathfrak{tm}, \Gamma' \vdash M \Leftarrow \mathfrak{tm}$. This theorem assumes the hereditary substitution $[M_2/x]_{\mathfrak{tm}}^m M = M'$ only to make the inductive argument more convenient; applying *Lemma 3.4*, *Theorem 3.5*, and *Theorem 2.11* to the other assumptions shows that the substitution must exist. □

3.2.3 Discussion

Transport of adequacy lemmas like *Lemma 3.4* demonstrate the convenience of the general subordination-based transport of canonical forms theorem (*Theorem 2.17*). Absent this general theorem, we could prove each individual transport of adequacy lemma inductively. Alternatively, we could state and prove every adequacy theorem for the largest context necessary for the encoding of the entire language. Fortunately, subordination saves us from the tedium of the first alternative and the inherent non-modularity of the second, which requires knowing, in advance, all future uses of any piece of the object language.

3.3 Encoding of Judgements

The judgements in Figure 9 define the static and dynamic semantics of the STLC in informal notation. In the type system, we use γ to notate object-language contexts

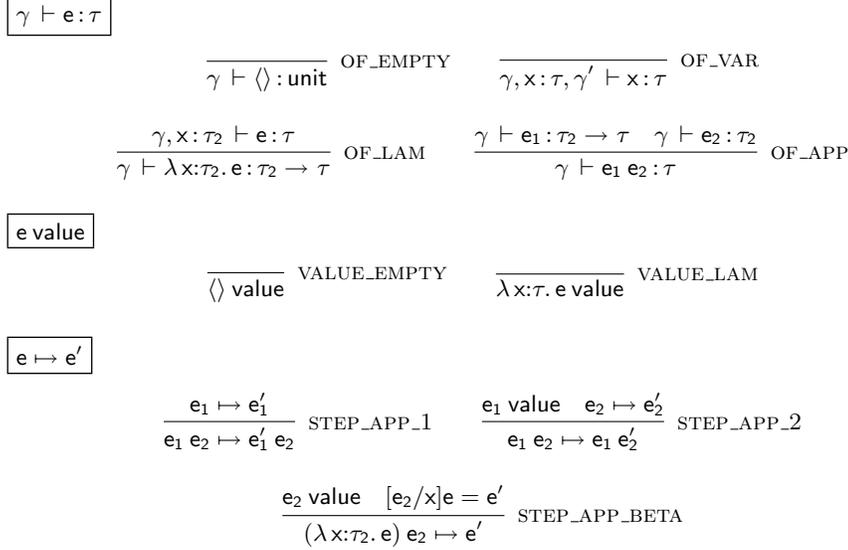


Fig. 9. Semantics of the STLC

containing assumptions of the form $x_i : \tau_i$; such a context is well-formed when all variables in it are distinct. Whenever we write γ , we tacitly presuppose that this context is well-formed. The dynamic semantics are a standard call-by-value structural operational semantics on closed terms. The following property of the STLC is necessary for the adequacy proofs below:

Lemma 3.8 (Uniqueness of Substitution Derivations)

If \mathcal{D} and \mathcal{D}' both derive $[e_2/x]e = e'$ then $\mathcal{D} = \mathcal{D}'$.

The LF signature in Figure 10 extends the signature from Figure 7 to represent these judgements. The representation of the static and dynamic semantics is guided by the *judgements-as-types* principle: an object-language judgement is represented by a family of LF types indexed by the subjects of the judgement; derivations of a judgement are represented as LF terms inhabiting this type. Such a representation is adequate iff there is an isomorphism between the informal derivations of the judgement and the canonical forms of the associated type family.

For instance, the operational semantics judgement $e \mapsto e'$, whose subject is two terms, is represented by the LF type family $\text{step} : \mathbf{tm} \rightarrow \mathbf{tm} \rightarrow \mathbf{type}$. Each constant with head **step** corresponds to the object-language inference rule with the same name. A derivation of the judgement $e \mapsto e'$ is represented by an LF term of type $\text{step } M_e M'_e$, where M_e and M'_e are the encodings of e and e' . For example, assuming terms M_1, M_2 , and M'_1 of LF type \mathbf{tm} and an LF term $M \Leftarrow \text{step } M_1 M'_1$, the LF term $(\text{step.app.1 } M_1 M_2 M'_1 M)$ represents a derivation of $\text{step } (\text{app } M_1 M_2) (\text{app } M'_1 M_2)$.

The object-language hypothetical judgement $\gamma \vdash e : \tau$ is represented by the LF type family of , which is indexed by the LF types \mathbf{tm} and \mathbf{tp} . One might expect the family to also be indexed by the encoding of a context γ . However, rather

```

      of      :  tm → tp → type
of_empty    :  of empty unit
of_app      :  Π E1, E2:tm. Π T2, T:tp.
               (of E1 (arrow T2 T)) → (of E2 T2) → of (app E1 E2) T
of_lam      :  Π T2, T:tp. Π E:tm → tm.
               (Π x:tm. (of x T2) → (of (E x) T)) → of (lam T2 λ x. E x) (arrow T2 T)

      value   :  tm → type
value_empty  :  value empty
value_lam    :  Π T:tp. Π E:tm → tm. value (lam T λ x. E x)

      step    :  tm → tm → type
step_app_1   :  Π E1, E2, E'1:tm. (step E1 E'1) → step (app E1 E2) (app E'1 E2)
step_app_2   :  Π E1, E2, E'2:tm.
               (value E1) → (step E2 E'2) → step (app E1 E2) (app E1 E'2)
step_app_beta : Π E2:tm. Π E:tm → tm. Π T2:tp.
               (value E2) → step (app (lam T2 λ x. E x) E2) (E E2)

```

Fig. 10. LF Signature for the STLC Judgements

than representing the context as an explicit argument, we represent the object-language hypotheses as LF hypotheses, identifying the object-language hypothetical judgement with a hypothetical judgement in LF. Specifically, an object-language context

$$x_1 : \tau_1, \dots, x_n : \tau_n$$

is represented by an LF context

$$x_1 : \text{tm}, dx_1 : \text{of } x_1 \text{ } T_1, \dots, x_n : \text{tm}, dx_n : \text{of } x_n \text{ } T_n$$

where each dx_i is a fresh variable that is distinct from all x_i and from the other dx_j . Each object-language hypothesis $x_i : \tau_i$ is represented by an LF variable dx_i . Correspondingly, context extension is represented by LF terms of higher type. Consider the object-language rule OF_LAM:

$$\frac{\gamma, x : \tau_2 \vdash e : \tau}{\gamma \vdash \lambda x : \tau_2. e : \tau_2 \rightarrow \tau} \text{ OF_LAM}$$

In the premise of the rule, the context is extended with the assumption $x : \tau_2$ for a fresh variable x . This rule is represented by the LF constant `of_lam`. The premise of `of_lam` is represented by an LF term of type of the following dependent function type:

$$\Pi x : \text{tm}. (\text{of } x \text{ } T_2) \rightarrow (\text{of } (E \ x) \ T).$$

$$\begin{array}{c}
 \boxed{\text{vars } \gamma \ \mathcal{X}} \\
 \\
 \boxed{\gamma \text{ ctx} \gg \Gamma \text{ ctx}} \\
 \\
 \frac{\text{vars } \gamma \ \mathcal{X}}{\text{vars } \cdot \cdot \quad \text{vars } (\gamma, x : \tau) \ (\mathcal{X}, x)} \\
 \\
 \frac{\gamma \text{ ctx} \gg \Gamma \text{ ctx} \quad \tau \text{ type} \quad \tau \gg M_t \Leftarrow \text{tp}}{\gamma, x : \tau \text{ ctx} \gg \Gamma, x : \text{tm}, dx : \text{of } x \ M_t \text{ ctx}} \text{ENC_CTX_CONS} \\
 \\
 \frac{\cdot \text{ctx} \gg \cdot \text{ctx}}{\text{ENC_CTX_EMPTY}} \\
 \\
 \boxed{\mathcal{D} :: \gamma \vdash e : \tau \gg \Gamma \vdash M \Leftarrow \text{of } M_e \ M_t} \\
 \\
 \frac{\gamma, x : \tau, \gamma' \vdash x : \tau \gg \Gamma, x : \text{tm}, dx : \text{of } x \ M_t, \Gamma' \vdash dx \Leftarrow \text{of } x \ M_t}{\text{ENC_OF_VAR}} \\
 \\
 \frac{\gamma \vdash \text{empty} : \text{unit} \gg \Gamma \vdash \text{of_empty} \Leftarrow \text{of empty unit}}{\text{ENC_OF_EMPTY}} \\
 \\
 \frac{\mathcal{D}_1 :: \gamma \vdash e_1 : \tau_2 \rightarrow \tau \gg \Gamma \vdash M_1 \Leftarrow \text{of } M_{e_1} \ (\text{arrow } M_{t_2} \ M_t) \quad \mathcal{D}_2 :: \gamma \vdash e_2 : \tau_2 \gg \Gamma \vdash M_2 \Leftarrow \text{of } M_{e_2} \ M_{t_2}}{\mathcal{D}_1 :: \gamma \vdash e_1 : \tau_2 \rightarrow \tau \quad \mathcal{D}_2 :: \gamma \vdash e_2 : \tau_2 \quad \gg \Gamma \vdash \text{of_app } M_{e_1} \ M_{e_2} \ M_t \ M_1 \ M_2 \Leftarrow \text{of } (\text{app } M_{e_1} \ M_{e_2}) \ M_t} \text{ENC_OF_APP} \\
 \\
 \frac{\tau_2 \text{ type} \gg M_{t_2} \Leftarrow \text{tp} \quad \mathcal{D}_1 :: \gamma, x : \tau_2 \vdash e : \tau \gg \Gamma, x : \text{tm}, dx : \text{of } x \ M_{t_2} \vdash M_2 \Leftarrow \text{of } M_e \ M_t}{\mathcal{D}_1 :: \gamma, x : \tau_2 \vdash e : \tau} \text{ENC_OF_LAM} \\
 \\
 \frac{\gamma \vdash \lambda x : \tau_2. e : \tau_2 \rightarrow \tau \gg \Gamma \vdash \text{of_lam } M_{t_2} \ M_t \ (\lambda x. \lambda dx. M_2) \Leftarrow \text{of } (\text{lam } M_t \ (\lambda x. M_e)) \ (\text{arrow } M_{t_2} \ M_t)}{\gamma \vdash \lambda x : \tau_2. e : \tau_2 \rightarrow \tau \gg \Gamma \vdash \text{of_lam } M_{t_2} \ M_t \ (\lambda x. \lambda dx. M_2) \Leftarrow \text{of } (\text{lam } M_t \ (\lambda x. M_e)) \ (\text{arrow } M_{t_2} \ M_t)}
 \end{array}$$

Fig. 11. Encoding of STLC Static Semantics

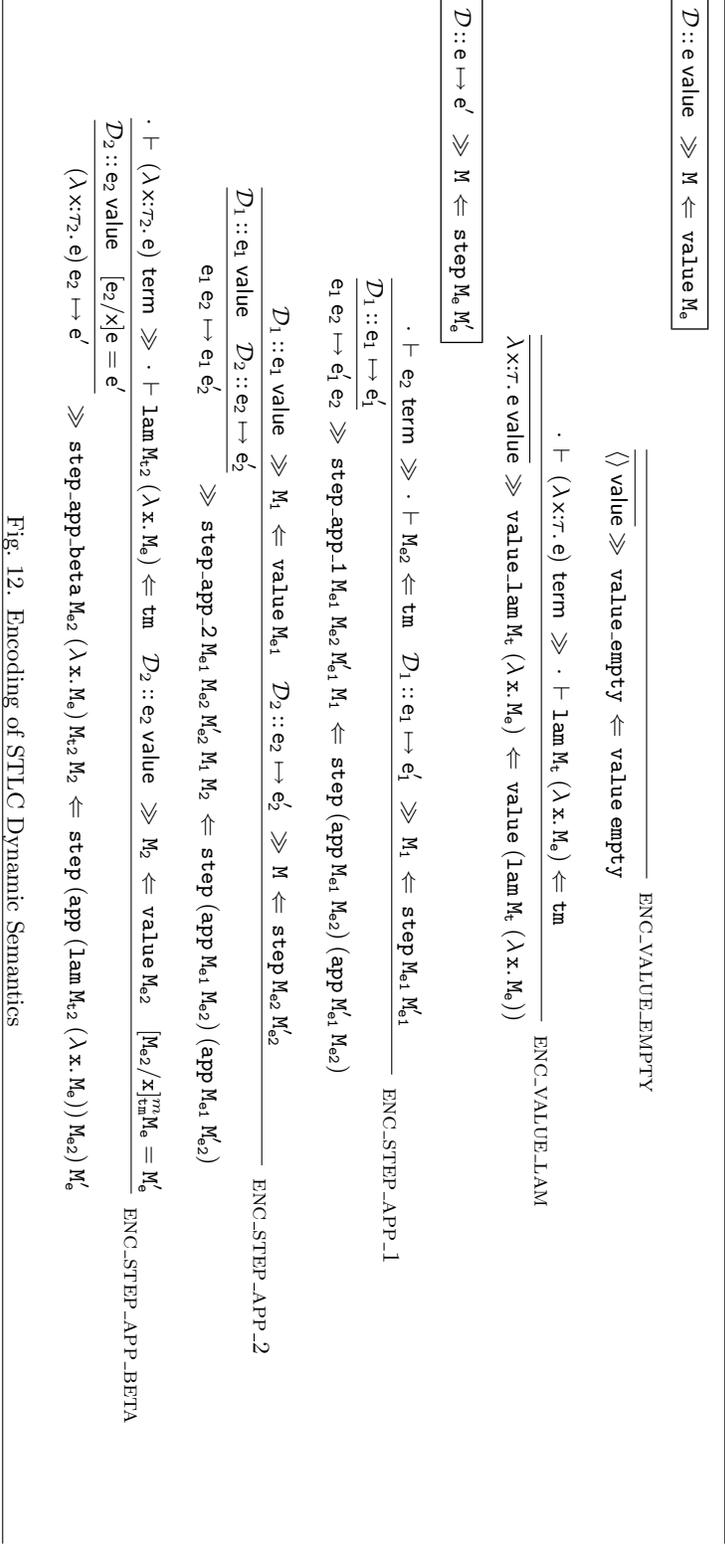


Fig. 12. Encoding of STLC Dynamic Semantics

The canonical forms of this type are terms $\lambda x. \lambda dx. M$, where M has LF type $\text{of } (\mathbf{E} \ x) \ T$ in an LF context including $x : \mathbf{tm}, dx : \text{of } x \ T_2$. These LF assumptions correspond exactly to a derivation under the hypothesis $x : \tau_2$ for a fresh variable x . This higher-order representation of hypotheses is natural and advantageous: there is no need to explicitly represent contexts and operations on them, and, moreover, LF provides structural properties such as weakening and substitution when hypothetical judgements are represented in such a fashion.

The correspondence between object-language derivations and their LF representations is made precise by the judgements in Figures 11 and 12. The intended reading of these judgements is as follows:

- $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ Encode an object-language context γ to an LF context Γ .
- $\mathcal{D} :: \gamma \vdash e : \tau \gg \Gamma \vdash M \Leftarrow \text{of } M_e M_t$ Assuming $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$, encode an object-language derivation \mathcal{D} of $\gamma \vdash e : \tau$ to an LF term M of type $\text{of } M_e M_t$ in Γ . The notation $\mathcal{D} :: \gamma \vdash e : \tau$ is simply linear notation for

$$\begin{array}{c} \mathcal{D} \\ \gamma \vdash e : \tau. \end{array}$$

- $\mathcal{D} :: e \text{ value} \gg M \Leftarrow \text{value } M_e$ Encode an object-language derivation \mathcal{D} of $e \text{ value}$ for a closed term e to an LF term M of type $\text{value } M_e$ in the empty LF context.
- $\mathcal{D} :: e \mapsto e' \gg M \Leftarrow \text{step } M_e M'_e$ Encode an object-language derivation \mathcal{D} of $e \mapsto e'$ for closed terms e and e' to an LF term M of type $\text{step } M_e M'_e$ in the empty LF context.

The rule `ENC_OF_LAM` expresses the higher-order representation strategy described above. The rule `ENC_STEP_APP_BETA` includes a hereditary substitution premise because the right-hand term of the informal rule `STEP_APP_BETA` is $[e_2/x]e$; compositionality of the encoding (*Theorem 3.7*) shows that this substitution can be represented by writing $(\mathbf{E} \ E_2)$ as the right-hand term in `step_app_beta`.

These encoding judgements clarify the fact that the LF representation of the object language judgements is specified not just by the LF signature, but also by the LF contexts in which that signature is considered. For example, the object language operational semantics judgement $e \mapsto e'$ presupposes that the terms e and e' are closed. In the LF representation, this presupposition is reflected in the fact that we consider only the inhabitants of `step` in the empty LF context. Note that the same LF signature for `step`, considered in other LF contexts, could adequately represent an object-language transition system that does not presuppose closed terms.

3.4 Adequacy of the Judgement Encodings

Let Σ stand for the signature in Figure 10. When proving adequacy, we consider each judgement of the object language in the fragment of Σ relevant to it. (This is in contrast to Section 3.2, where for simplicity we considered adequacy of both types and terms in the full signature.) Observe by the definition of signature restriction that

- $\Sigma|_{\text{tp}}$ contains the family declaration **tp** and the constants **unit** and **arrow**.
- $\Sigma|_{\text{tm}}$ contains the family declarations **tp**, **tm** and their associated constants.
- $\Sigma|_{\text{of}}$ contains the family declarations **tp**, **tm**, **of** and their associated constants.
- $\Sigma|_{\text{value}}$ contains the family declarations **tp**, **tm**, **value** and their associated constants.
- $\Sigma|_{\text{step}}$ contains the family declarations **tp**, **tm**, **value**, **step** and their associated constants.

The strongest subordination relation \preceq_Σ is the reflexive-transitive closure of the following relation: $\text{tp} \preceq \text{tm}$, $\text{tm} \preceq \text{of}$, $\text{tp} \preceq \text{of}$, $\text{tm} \preceq \text{value}$, $\text{tp} \preceq \text{value}$, $\text{tm} \preceq \text{step}$, $\text{tp} \preceq \text{step}$, and $\text{value} \preceq \text{step}$. Additionally, for each **a** of these type families, $\preceq_\Sigma|_{\mathbf{a}} = \preceq_{\Sigma|_{\mathbf{a}}}$. Thus, *Theorem 2.17* shows that each of these signature restrictions is well-formed in the strongest subordination relation for it. Our previous adequacy theorems show that the encodings of terms and types are adequate in $\Sigma|_{\text{tm}}$. An easy application of *Theorem 2.17* shows that adequacy for STLC types, *Theorem 3.1*, holds in $\Sigma|_{\text{tp}}$ as well.

3.4.1 Static Semantics

In this section, we work in $\text{LF}[\Sigma|_{\text{of}}]$. Just as it was necessary to ensure that the encoding of types remained adequate in the signature and contexts for terms, it is now necessary to show that the encodings of both types and terms remain adequate in the signature and contexts we consider for typing derivations:

Lemma 3.9 (Transport of Adequacy for Typing)

1. $(\Sigma|_{\text{of}})|_{\text{tp}} = \Sigma|_{\text{tp}}$ and $(\Sigma|_{\text{of}})|_{\text{tm}} = \Sigma|_{\text{tm}}$.
2. If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ then $\Gamma \text{ ctx}$.
3. If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ then $\Gamma|_{\text{tp}} = \cdot$.
4. If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$, $\text{vars } \gamma \mathcal{X}$, and $\Gamma|_{\text{tm}} = \Gamma'$ then $\mathcal{X} \text{ terms} \gg \Gamma' \text{ ctx}$.
5. If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ then $\Gamma \vdash_{\Sigma|_{\text{of}}} M_t \Leftarrow \text{tp}$ iff $\cdot \vdash_{\Sigma|_{\text{tp}}} M_t \Leftarrow \text{tp}$.
6. If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ then $\Gamma \vdash_{\Sigma|_{\text{of}}} M_e \Leftarrow \text{tm}$ iff $\Gamma|_{\text{tm}} \vdash_{\Sigma|_{\text{tm}}} M_e \Leftarrow \text{tm}$.

Next, we state the inversion lemma:

Lemma 3.10 (Inversion of Canonical Forms of Type of)

If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ and $\Gamma \vdash M \Leftarrow \text{of } M_e M_t$ then one of the following hold:

- $M = \text{dx}$, $M_e = \mathbf{x}$, and $\Gamma = \Gamma_1, \mathbf{x} : \text{tm}, \text{dx} : \text{of } \mathbf{x} M_t, \Gamma_2$ for some variables dx and \mathbf{x} .
- $M = \text{of_empty}$, $M_e = \text{empty}$, and $M_t = \text{unit}$.
- $M = \text{of_app } M_{e1} M_{e2} M_{t2} M_t M_1 M_2$, $M_e = (\text{app } M_{e1} M_{e2})$, and the following were derived as strict subderivations: $\Gamma \vdash M_{e1} \Leftarrow \text{tm}$ and $\Gamma \vdash M_{e2} \Leftarrow \text{tm}$ and $\Gamma \vdash M_{t2} \Leftarrow \text{tp}$ and $\Gamma \vdash M_t \Leftarrow \text{tp}$ and $\Gamma \vdash M_1 \Leftarrow \text{of } M_{e1} (\text{arrow } M_{t2} M_t)$ and $\Gamma \vdash M_2 \Leftarrow \text{of } M_{e2} M_{t2}$.
- $M = \text{of_lam } M_{t2} M'_t (\lambda \mathbf{x}. M'_e) (\lambda \mathbf{x}. \lambda \text{dx}. M_2)$, $M_e = (\text{lam } M_t (\lambda \mathbf{x}. M'_e))$, $M_t = (\text{arrow } M_{t2} M'_t)$, and the following were derived as strict subderivations: $\Gamma \vdash M_{t2} \Leftarrow \text{tp}$ and $\Gamma \vdash M'_t \Leftarrow \text{tp}$ and $\Gamma, \mathbf{x} : \text{tm} \vdash M'_e \Leftarrow \text{tm}$ and $\Gamma, \mathbf{x} : \text{tm}, \text{dx} : \text{of } \mathbf{x} M_{t2} \vdash M_2 \Leftarrow \text{of } M'_e M'_t$.

In the statement and proof of adequacy, it will be convenient to abuse notation by writing $\gamma \vdash e \text{ term} \gg \Gamma \vdash M_e \Leftarrow \text{tm}$ to mean the following: $\text{vars } \gamma \mathcal{X}$ and $\Gamma|_{\text{tm}} = \Gamma'$ and $\mathcal{X} \vdash e \text{ term} \gg \Gamma' \vdash M_e \Leftarrow \text{tm}$. In *Definition 2.16*, we observed that given Γ and A , there exists a unique Γ' such that $\Gamma|_A = \Gamma'$. Observe (by induction on γ) that given a well-formed context γ , there exists a unique \mathcal{X} such that $\text{vars } \gamma \mathcal{X}$. Thus, both \mathcal{X} and Γ' are uniquely determined when we write $\gamma \vdash e \text{ term} \gg \Gamma \vdash M_e \Leftarrow \text{tm}$.

We can now prove adequacy. The first condition states not only that the subjects are well-formed, as above, but also that the indices to the judgement are encoded in the appropriate manner.

Theorem 3.11 (Adequacy for Typing Derivations)

1. If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ and $\mathcal{D} :: \gamma \vdash e : \tau \gg \Gamma \vdash M \Leftarrow \text{of } M_e M_t$ then
 - $\tau \text{ type} \gg M_t \Leftarrow \text{tp}$ and $\gamma \vdash e \text{ term} \gg \Gamma \vdash M_e \Leftarrow \text{tm}$,
 - \mathcal{D} derives $\gamma \vdash e : \tau$, and
 - $\Gamma \vdash M \Leftarrow \text{of } M_e M_t$.
2. If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ and \mathcal{D} derives $\gamma \vdash e : \tau$ then there exist unique M_e, M_t and M such that $\mathcal{D} :: \gamma \vdash e : \tau \gg \Gamma \vdash M \Leftarrow \text{of } M_e M_t$.
3. If $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ and $\Gamma \vdash M \Leftarrow \text{of } M_e M_t$ then there exist unique τ, e , and \mathcal{D} such that $\mathcal{D} :: \gamma \vdash e : \tau \gg \Gamma \vdash M \Leftarrow \text{of } M_e M_t$.

Proof

The proof follows the same general pattern as adequacy for syntax. It uses *Lemma 3.9*, *Theorem 3.5*, *Theorem 3.1*, *Lemma 2.12*, and *Lemma 3.10*. \square

In informal descriptions of programming languages, we do not usually consider the operation of substituting one typing derivation into another (i.e., the computational content of the proof of the substitution theorem). However, it is possible to define such an operation and then prove a compositionality theorem, analogous to *Theorem 3.7*, for the judgement $\mathcal{D} :: \gamma \vdash e : \tau \gg \Gamma \vdash M \Leftarrow \text{of } M_e M_t$. However, because this compositionality result is not necessary for the remainder of this article, we elide the details.

3.4.2 Dynamic Semantics

The adequacy proof for the value judgement is simple; we elide the transport of adequacy lemma and the canonical forms inversion lemma, which are analogous to those given above.

Theorem 3.12 (Adequacy for Values)

1. If $\mathcal{D} :: e \text{ value} \gg M \Leftarrow \text{value } M_e$ then
 - $\cdot \vdash e \text{ term} \gg \cdot \vdash M_e \Leftarrow \text{tm}$,
 - \mathcal{D} derives $e \text{ value}$, and
 - $\cdot \vdash_{\Sigma|_{\text{value}}} M \Leftarrow \text{value } M_e$.
2. If $\cdot \vdash e \text{ term}$ and $\mathcal{D} :: e \text{ value}$ then there exist unique M_e and M such that $\mathcal{D} :: e \text{ value} \gg M \Leftarrow \text{value } M_e$.

3. If $\cdot \vdash_{\Sigma|_{\text{value}}} M \Leftarrow \text{value } M_e$ then there exist unique e and \mathcal{D} such that $\mathcal{D} :: e \text{ value} \gg M \Leftarrow \text{value } M_e$.

In the remainder of this section, we work in $\text{LF}[\Sigma|_{\text{step}}]$.

Lemma 3.13 (Transport of Adequacy for Operational Semantics)

1. For $a \in \{\text{tm}, \text{tp}, \text{value}\}$, $(\Sigma|_{\text{step}})|_a = \Sigma|_a$.
2. $\cdot \vdash_{\Sigma|_{\text{step}}} M_e \Leftarrow \text{tm}$ iff $\cdot \vdash_{\Sigma|_{\text{tm}}} M_e \Leftarrow \text{tm}$.
3. $\cdot \vdash_{\Sigma|_{\text{step}}} M_t \Leftarrow \text{tp}$ iff $\cdot \vdash_{\Sigma|_{\text{tp}}} M_t \Leftarrow \text{tp}$.
4. $\cdot \vdash_{\Sigma|_{\text{step}}} M \Leftarrow \text{value } M_e$ iff $\cdot \vdash_{\Sigma|_{\text{value}}} M \Leftarrow \text{value } M_e$.

Next, we state the inversion principle:

Lemma 3.14 (Inversion of Canonical Forms of Type step)

If $M \Leftarrow \text{step } M_e M'_e$ then one of the following hold:

- $M = \text{step_app_1 } M_{e1} M_{e2} M'_{e1} M_1$, $M_e = (\text{app } M_{e1} M_{e2})$, $M'_e = (\text{app } M'_{e1} M_{e2})$, and the following were derived as strict subderivations: $M_{e1} \Leftarrow \text{tm}$ and $M_{e2} \Leftarrow \text{tm}$ and $M'_{e1} \Leftarrow \text{tm}$ and $M_1 \Leftarrow \text{step } M_{e1} M'_{e1}$.
- $M = \text{step_app_2 } M_{e1} M_{e2} M'_{e2} M_1 M_2$, $M_e = (\text{app } M_{e1} M_{e2})$, $M'_e = (\text{app } M_{e1} M'_{e2})$, and the following were derived as strict subderivations: $M_{e1} \Leftarrow \text{tm}$ and $M_{e2} \Leftarrow \text{tm}$ and $M'_{e2} \Leftarrow \text{tm}$ and $M_1 \Leftarrow \text{value } M_{e1}$ and $M_2 \Leftarrow \text{step } M_{e2} M'_{e2}$.
- $M = \text{step_app_beta } M_{e2} (\lambda x. M_b) M_{t2} M_2$, $M_e = (\text{app } (\text{lam } M_{t2} (\lambda x. M_b)) M_{e2})$, and the following were derived as strict subderivations: $M_{e2} \Leftarrow \text{tm}$ and $x : \text{tm} \vdash M_b \Leftarrow \text{tm}$ and $M_{t2} \Leftarrow \text{tp}$ and $M_2 \Leftarrow \text{value } M_{e2}$ and $[M_{e2}/x]_{\text{tm}}^m M_b = M'_e$.

Finally, we state adequacy.

Theorem 3.15 (Adequacy for Operational Semantics)

1. If $\mathcal{D} :: e \mapsto e' \gg M \Leftarrow \text{step } M_e M'_e$ then
 - $\cdot \vdash e \text{ term} \gg \cdot \vdash M_e \Leftarrow \text{tm}$ and $\cdot \vdash e' \text{ term} \gg \cdot \vdash M'_e \Leftarrow \text{tm}$,
 - \mathcal{D} derives $e \mapsto e'$, and
 - $M \Leftarrow \text{step } M_e M'_e$.
2. If $\cdot \vdash e \text{ term}$ and $\cdot \vdash e' \text{ term}$ and \mathcal{D} derives $e \mapsto e'$ then there exist unique M_e, M'_e and M such that $\mathcal{D} :: e \mapsto e' \gg M \Leftarrow \text{step } M_e M'_e$.
3. If $M \Leftarrow \text{step } M_e M'_e$ then there exist unique e, e' , and \mathcal{D} such that $\mathcal{D} :: e \mapsto e' \gg M \Leftarrow \text{step } M_e M'_e$.

Proof

An object-language derivation using the rule `STEP_APP_BETA` contains a subderivation deriving $[e_2/x]e$. However, because substitution in the object-language is represented as substitution in LF, this substitution derivation is not encoded as an explicit LF object. Consequently, it is necessary to show that there is at most one object-language substitution derivation; *Lemma 3.8* establishes this fact. Other than this complication, the proof is similar to adequacy of typing. The proof uses *Theorem 2.11*, *Theorem 3.5*, *Theorem 3.7*, *Theorem 3.12*, *Lemma 3.13*, and *Lemma 3.14*. \square

4 Mechanizing the Metatheory of the STLC

A *metatheorem* is a statement about an object language. The Twelf implementation of LF (Pfenning & Schürmann, 1999) permits the mechanization of metatheorems about languages encoded in LF.

Type Families as Relations. In describing the metatheoretic capabilities of Twelf, it is useful to regard an LF type family as defining a relation on the family’s indices, where indices are related iff their instance of the type family is inhabited. In the simplest case, when a type family’s kind is not dependent and when only closed LF terms are considered, a type family defines a single relation on the family’s indices, where indices are related iff the corresponding type is inhabited. For example, the type family `step`: `tm` \rightarrow `tm` \rightarrow `type` defines a relation between two LF terms of type `tm`, where `E` and `E'` are related iff there exists an LF term `D` of type `step E1 E2`. More formally, we say that a type family `a`: `A1` \rightarrow \dots \rightarrow `An` \rightarrow `type` defines a relation R on terms $\cdot \vdash M_1 \Leftarrow A_1, \dots, \cdot \vdash M_n \Leftarrow A_n$, where $R(M_1, \dots, M_n)$ iff there exists an LF term `D` such that `D` \Leftarrow `a M1 ... Mn`.

This simple case generalizes in two ways. First, when LF terms in non-empty contexts are considered, a type family constitutes a simultaneous definition of a context-indexed family of relations on the type family’s indices, where indices are related by the relation for a particular context iff their instance of the family is inhabited in that context. For example, the type family `of`: `tm` \rightarrow `tp` \rightarrow `type` defines a context-indexed family of relations R , where each relation R_Γ relates terms `E` and `T` such that $\Gamma \vdash E \Leftarrow \text{tm}$ and $\Gamma \vdash T \Leftarrow \text{tp}$, and $R_\Gamma(E, T)$ holds iff there exists an LF term `D` such that $\Gamma \vdash D \Leftarrow \text{of } E\ T$. Note that the type family constitutes a simultaneous inductive definition of all the relations R_Γ , as the definition of R_Γ refers to other relations $R_{\Gamma'}$ (e.g., when the premise of `of.lam` extends the LF context). More formally, we say that a type family `a`: `A1` \rightarrow \dots \rightarrow `An` \rightarrow `type` defines a context-indexed family of relations R , where each R_Γ is a relation on terms $\Gamma \vdash M_1 \Leftarrow A_1, \dots, \Gamma \vdash M_n \Leftarrow A_n$ and $R_\Gamma(M_1, \dots, M_n)$ iff there exists a canonical form `D` such that $\Gamma \vdash D \Leftarrow$ `a M1 ... Mn`.

Second, the kind of a type family may in general be a dependent kind of the form $\Pi x_1:A_1. \dots \Pi x_n:A_n. \text{type}$, in which case the type family defines a dependent relation—the types of the related terms may vary with the other indices to the relation. In full generality, a type family `a`: $\Pi x_1:A_1. \dots \Pi x_n:A_n. \text{type}$ defines a context-indexed family of relations R , where R_Γ relates terms `M1` such that $\Gamma \vdash M_1 \Leftarrow A_1$, `M2` such that $[M_1/x_1]_{A_1}^a A_2 = A'_2$ and $\Gamma \vdash M_2 \Leftarrow A'_2$, `M3` such that $[M_2/x_2]_{A_2}^a [M_1/x_1]_{A_1}^a A_3 = A'_3$, and so on, and $R_\Gamma(M_1, \dots, M_n)$ iff there exists a canonical form `D` such that $\Gamma \vdash D \Leftarrow$ `a M1 ... Mn`. Though we have not used any dependent kinds in the previous sections of this article, we discuss many examples of them below.

Totality Assertions. Twelf has the ability to mechanically verify *totality assertions* about LF type families. To a first approximation, a totality assertion for a type family corresponds with the standard notion of totality for the relation defined by the type family: a totality assertion is specified by designating some indices as

inputs and the remaining indices as outputs; then the totality assertion asserts that for all inputs, there exist outputs that stand in the relation. However, because a type family defines a family of relations R_Γ , the specification of a totality assertion must also clarify for which contexts Γ the relations R_Γ are asserted to be total. In general, we may wish to state a totality assertion only for contexts of a particular form—because, for example, the totality assertion may be true for contexts of one form but false for contexts of another. To allow this, we parametrize the specification of a totality assertion by a world (set of contexts) \mathcal{W} , and the totality assertion asserts that for all contexts $\Gamma \in \mathcal{W}$, the relation R_Γ is total.

Thus, the specification of a totality assertion for a given type family consists of two declarations: a *mode* declaration, which identifies some of the family’s indices as inputs (we notate these A_i) and the others as outputs (we notate these B_j), and a *world* declaration, which restricts the totality assertion to LF contexts in a particular world \mathcal{W} . Given these declarations, the totality assertion for a type family is defined as follows.

Definition 4.1 (Totality Assertions)

The totality assertion for a type family $\mathbf{a} : \Pi x_1 : A_1 \dots \Pi y_1 : B_1 \dots \mathbf{type}$ with inputs A_1, \dots, A_m and outputs B_1, \dots, B_n in world \mathcal{W} is the proposition

$$\begin{aligned} &\text{For all } \Gamma \in \mathcal{W}, \text{ for all } M_1 \text{ such that } \Gamma \vdash M_1 \Leftarrow A'_1, \dots, \text{ and } M_m \text{ such that} \\ &\Gamma \vdash M_m \Leftarrow A'_m, \text{ there exist } N_1 \text{ such that } \Gamma \vdash N_1 \Leftarrow B'_1, \dots, \text{ and } N_n \text{ such that} \\ &\Gamma \vdash N_n \Leftarrow B'_n, \text{ and } D \text{ such that } \Gamma \vdash D \Leftarrow \mathbf{a} M_1 \dots M_m N_1 \dots N_n. \end{aligned}$$

where we write A'_i and B'_j for the appropriate substitutions into A_i and B_j to account for the dependent nature of the relation:

$$\begin{aligned} &[M_{i-1}/x_{i-1}]_{A_{i-1}} \dots [M_1/x_1]_{A_1} A_i = A'_i \\ &[N_{j-1}/y_{j-1}]_{B_{j-1}} \dots [N_1/y_1]_{B_1} [M_m/x_m]_{A_m} \dots [M_1/x_1]_{A_1} B_j = B'_j. \end{aligned}$$

For example, we may specify a totality assertion for the type family `step E E'` by declaring `E` an input and `E'` an output and by restricting consideration to the empty LF context. These declarations specify the following totality assertion: for all `E` such that `E` \Leftarrow `tm`, there exist `E'` and `D` such that `E'` \Leftarrow `tm` and `D` \Leftarrow `step E E'`. Of course, this totality assertion for the STLC is false, as there are terms that cannot take a step. Other judgements of an object language, such as a compiler transformation, might in fact define total relations.

The Twelf implementation permits a totality assertion to be specified by a mode declaration and a *regular worlds* declaration, which defines a world \mathcal{W} by a regular expression over contexts. Moreover, Twelf provides a totality declaration which instructs Twelf to (attempt to) prove a totality assertion using a specified induction metric. Twelf proves a totality assertion for a type family by interpreting the type family as a higher-order logic program and proving that the logic program is total; see Schürmann & Pfenning (2003) for details. As with any theorem prover for sufficiently rich statements, there are many true totality assertions that Twelf is unable to prove. In the remainder of this section, we present many examples of type

families that Twelf can prove total, and we discuss how totality assertions are used in mechanizing metatheory.

Mechanizing Metatheory. The machinery of totality assertions can be deployed in a particular way to permit the mechanical verification of a much wider class of metatheorems than totality assertions themselves. General $\forall\exists$ -statements of the form

$$\begin{array}{l} \text{For all } \Gamma \in \mathcal{W}, \text{ for all } M_1 \text{ such that } \Gamma \vdash M_1 \Leftarrow A_1, \dots, \text{ and } M_m \text{ such that} \\ \Gamma \vdash M_m \Leftarrow A_m, \text{ there exist } N_1 \text{ such that } \Gamma \vdash N_1 \Leftarrow B_1, \dots, \text{ and } N_n \text{ such that} \\ \Gamma \vdash N_n \Leftarrow B_n. \end{array}$$

can be mechanized in Twelf. A proof of a $\forall\exists$ -statement consists of a transformation from the universally quantified terms (inputs) into the existentially quantified terms (outputs), typically defined by induction on the inputs. One way of presenting such a transformation is as a total relation. Consequently, a $\forall\exists$ -statement and its inductive proof can be represented as an LF type family and verified by Twelf's totality checker. Specifically, a statement of this form can be translated into the type family, mode, and worlds declarations specifying the totality of a type family $\mathbf{a} : \Pi x_1:A_1 \dots \Pi y_n:B_n$. **type**. An inductive proof of such a statement can be translated into LF constants that inductively define the type family \mathbf{a} in such a way that Twelf can verify the its totality. The original $\forall\exists$ -statement is a corollary of the totality assertion for this particular relation. That is, we deploy Twelf's ability to *prove* totality assertions in order to *check* proofs of general $\forall\exists$ -statements by representing these proofs explicitly as LF terms. A successful proof constitutes a sufficiently detailed definition of a relation that Twelf can verify its totality.

This methodology is useful because a wide class of metatheorems can expressed as $\forall\exists$ -statements about the canonical forms of LF. Because of the judgements-as-types representation strategy, the quantifiers of $\forall\exists$ -statements range over not just the representations of object-language individuals such as terms and types but also the representations of derivations of object-language judgements. For example, the standard informal statement of preservation is the following:

Theorem 4.2 (Informal Statement of Preservation)

For all types τ and all closed terms e and e' , if $e \mapsto e'$ and $\cdot \vdash e : \tau$ then $\cdot \vdash e' : \tau$.

By adequacy (*Theorem 3.1*, *Theorem 3.5*, *Theorem 3.11*, and *Theorem 3.15*), this statement can be recast as the following $\forall\exists$ -statement:

Theorem 4.3 (Statement of Preservation via the Encoding)

For all E such that $E \Leftarrow \mathbf{tm}$, E' such that $E' \Leftarrow \mathbf{tm}$, T such that $T \Leftarrow \mathbf{tp}$, D_{step} such that $D_{\text{step}} \Leftarrow \mathbf{step} E E'$, and D_{of} such that $D_{\text{of}} \Leftarrow \mathbf{of} E T$, there exists a D'_{of} such that $D'_{\text{of}} \Leftarrow \mathbf{of} E' T$.

To prove this statement in Twelf, we recast the informal proof of preservation as constants inhabiting an LF type family

$$\text{preserv} \quad : \quad \Pi E:\mathbf{tm}. \Pi E':\mathbf{tm}. \Pi T:\mathbf{tp}. \mathbf{step} E E' \rightarrow \mathbf{of} E T \rightarrow \mathbf{of} E' T \rightarrow \mathbf{type}.$$

The type family `preserv` has a dependent kind, accounting for the occurrences of the metavariables in the theorem statement. Then, we ask Twelf to verify the appropriate totality assertion:

Theorem 4.4 (Totality of Preservation)

For all E such that $E \Leftarrow \text{tm}$, E' such that $E' \Leftarrow \text{tm}$, T such that $T \Leftarrow \text{tp}$, D_{step} such that $D_{\text{step}} \Leftarrow \text{step } E E'$, and D_{of} such that $D_{\text{of}} \Leftarrow \text{of } E T$, there exist D'_{of} such that $D'_{\text{of}} \Leftarrow \text{of } E' T$ and D such that $D \Leftarrow \text{preserv } E E' T D_{\text{of}} D_{\text{step}} D'_{\text{of}}$.

To state and prove preservation, it suffices to consider a $\forall\exists$ -statement for the world containing only the empty LF context. In general, $\forall\exists$ -statements about non-empty LF contexts are useful for stating properties of object languages encoded using higher-order representations. For example, we may recast a statement about STLC typing derivations in non-empty contexts γ as a statement about LF terms of type `of` in contexts Γ such that $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$. We prove such a statement in Twelf by verifying the totality assertion for a corresponding type family in an appropriate world.

In summary, proving an informal $\forall\exists$ -statement about an object language in Twelf requires four steps. First, via adequacy, recast the metatheorem as a $\forall\exists$ -statement about canonical forms of particular types in a particular world. Second, define a corresponding LF type family \mathbf{r} and annotate it with mode and worlds declarations specifying the corresponding totality assertion. Third, extend the LF signature with constants inhabiting \mathbf{r} for all canonical inputs. Fourth, ask Twelf to verify that \mathbf{r} satisfies the totality assertion by executing a totality declaration. In the remainder of this section, we show several examples of this methodology for mechanizing metatheorems in Twelf.

Twelf Concrete Syntax. In this section, we will use Twelf’s concrete syntax for LF. The type $\Pi x:A_2. A$ is written as $\{x:A_2\} A$, the kind $\Pi x:A_2. K$ is written as $\{x:A_2\} K$, and the term $\lambda x.M$ is written as $[x] M$. Parentheses are used for grouping, and the usual λ -calculus associativities and precedences apply: application associates to the left; \rightarrow associates to the right; the scope of a binder extends as far to the right as possible. Twelf’s concrete syntax also includes several conveniences that we will exploit:

- When declaring a constant in a signature, if an identifier beginning with a lower-case letter is not bound, Twelf reports an error. If an identifier beginning with an upper-case letter is not bound, Twelf implicitly binds it in a Π at the front of the constant’s type or kind. The application of a constant to these implicit arguments is then inferred.
- Twelf allows a programmer to write an arrow type $A \rightarrow B$ with a backward arrow— $B \leftarrow A$. This makes it easier to see the head family of a constant.
- Twelf allows arbitrary type annotations by writing $M : A$ in place of any M .
- Twelf permits non-canonical forms, which are treated as syntactic sugar for the corresponding canonical form.

In Figure 13, we present an example of the first two conveniences: we show Twelf concrete syntax for the signature from Figure 10 defining the static and dynamic

```

of      : tm -> tp -> type.
of_empty : of empty unit.
of_lam  : of (lam T2 ([x] E x)) (arrow T2 T)
        <- ({x:tm} (of x T2) -> (of (E x) T)).
of_app  : of (app E1 E2) T
        <- of E2 T2
        <- of E1 (arrow T2 T).

value   : tm -> type.
value_empty : value empty.
value_lam : value (lam T2 ([x] E x)).

step      : tm -> tm -> type.
step_app_1 : step (app E1 E2) (app E1' E2)
          <- step E1 E1'.
step_app_2 : step (app E1 E2) (app E1 E2')
          <- step E2 E2'
          <- value E1.
step_app_beta : step (app (lam T2 ([x] E x)) E2) (E E2)
          <- value E2.

```

Fig. 13. Static and Dynamic Semantics for the STLC in Twelf Concrete Syntax

semantics of the STLC. In idiomatic Twelf, it is common to reverse the order of the premises so that they read, from top to bottom, in the same order as they would read from left to right in an inference rule; however, to keep the order of arguments consistent with the signature in Figure 10, we have not done this here.

4.1 Type Preservation

Type preservation is our first example of a Twelf metatheorem.

4.1.1 Twelf Proof

Figure 14 contains a complete Twelf proof of preservation for the STLC. The type family `preserv` defines a relation among the appropriate types. The `%mode` and `%worlds` declarations state the appropriate totality assertion: the mode declaration declares that the first two indices of `preserv` are universally quantified (+), whereas the last one is existentially quantified (-). Implicitly quantified variables are universally quantified if they appear in any inputs and existentially quantified if they appear only in outputs, so in this case `E`, `E'`, and `T` are all inputs. The `%worlds` declaration declares that these quantifiers range over canonical forms in the world containing only the empty LF context.

The term-level constants `preserv_app_1`, `preserv_app_2`, and `preserv_app_beta` constitute an inductive definition of the `preserv` type family, codifying the cases of the proof of the $\forall\exists$ -theorem. We prove preservation by induction on the operational semantics, writing one LF constant (i.e., case of the proof) for each operational semantics rule. For example, consider `preserv_app_1`. The intuitive reading of this

```

preserv : step E E' -> of E T -> of E' T -> type.
%mode preserv +Dstep +Dof -Dof'.

preserv_app_1  : preserv
  (step_app_1 (DstepE1 : step E1 E1'))
  (of_app (DofE1 : of E1 (arrow T2 T))
    (DofE2 : of E2 T2))
  (of_app DofE1' DofE2)
  <- preserv DstepE1 DofE1 (DofE1' : of E1' (arrow T2 T)).

preserv_app_2  : preserv
  (step_app_2 (DvalE1 : value E1) (DstepE2 : step E2 E2'))
  (of_app (DofE1 : of E1 (arrow T2 T))
    (DofE2 : of E2 T2))
  (of_app DofE1 DofE2')
  <- preserv DstepE2 DofE2 (DofE2' : of E2' T2).

preserv_app_beta : preserv
  (step_app_beta (Dval : value E2))
  (of_app (of_lam (([x] [dx] DofE x dx)
    : {x : tm} (of x T2) -> (of (E x) T)))
    (DofE2 : of E2 T2))
  (DofE E2 DofE2).

%worlds () (preserv _ _ _).
%total D (preserv D _ _).

```

Fig. 14. Twelf Proof of Preservation for the STLC

case is as follows: in the case for `step_app_1`, we invert the typing derivation because `of_app` is the only rule that could have applied, yielding typing derivations for `E1` and `E2`; we then appeal inductively to `preserv` on the smaller step derivation `DstepE1` and the typing derivation `DofE1` for `E1`, which yields a typing derivation `DofE1'` for `E1'`; then we apply the rule `of_app` to this derivation and the derivation for `E2` to obtain the result. More formally, `preserv_app_1` is an LF constant inhabiting the family

$$\text{preserv } (\text{step_app_1 } D\text{step1}) (\text{of_app } D\text{ofE1 } D\text{ofE2}) (\text{of_app } D\text{ofE1}' D\text{ofE2})$$

as long as the family `preserv Dstep1 DofE1 DofE2` is inhabited. This extends the canonical forms that are related by `preserv`. The fact that the premise of the constant is on a smaller step derivation `Dstep1` is used in showing the totality of `preserv`. The constant `preserv_app_2` is similar. In `preserv_app_beta`, the intuitive reading is that we invert twice based on the syntactic form of the term on the left-hand side of the step derivation; then we apply the hypothetical typing derivation for the body of the function `DofE` to the argument `E2` and its typing derivation `DofE2`—because of the higher-order encoding, there is no need for a substitution lemma. Of course, `preserv_app_beta` is in fact just an LF constant inhabiting the type `preserv` for particular indices. The type annotations on each constant are

included only for readability; if we omitted their types, Twelf would infer them. The `%total` declaration asks Twelf to verify, by induction on the first argument `D`, that `preserv` satisfies the totality assertion in *Theorem 4.4*.

4.1.2 Correctness of the Statement of Preservation

In the introduction to this section, we stated the equivalence of the informal statement of preservation (*Theorem 4.2*) and the statement via the encoding (*Theorem 4.3*). We discuss the proof of this equivalence in more detail here. The quantifiers and their types clearly correspond, so adequacy should imply the equivalence of these two statements. However, there are two ways in which the previous adequacy statements could fail to give the necessary result:

1. Though we did not make this explicit in the introduction to this section, the Twelf proof proves *Theorem 4.3* in a different signature than the one in which the relevant types are adequate: the type family `preserv` and the terms inhabiting it extend the LF signature. It is necessary to check that the previous adequacy results can be transferred to this extended signature.
2. The `%worlds` declaration specifies the class of LF contexts for the metatheorem. It is necessary to check that the previous adequacy results give the desired correspondence in these contexts.

To address the first concern, let Σ be the original signature in Figure 10, and call the extension of this signature with `preserv` and its constants Σ' . The strongest subordination relation $\preceq_{\Sigma'}$ extends \preceq_{Σ} with the conditions that `a` $\preceq_{\Sigma'}$ `preserv` for all `a` declared in Σ , but in $\preceq_{\Sigma'}$ no additional families are subordinate to any family declared in Σ . With this subordination relation, we can calculate that $\Sigma'|_{\text{of}} = \Sigma|_{\text{of}}$ and $\Sigma'|_{\text{step}} = \Sigma|_{\text{step}}$. Therefore, *Theorem 2.17* immediately implies the following lemma:

Lemma 4.5 (Transport of Adequacy for Preservation)

1. $\cdot \vdash_{\Sigma'} M \Leftarrow \text{of } M_e M_t$ iff $\cdot \vdash_{\Sigma|_{\text{of}}} M \Leftarrow \text{of } M_e M_t$.
2. $\cdot \vdash_{\Sigma'} M \Leftarrow \text{step } M_e M'_e$ iff $\cdot \vdash_{\Sigma|_{\text{step}}} M \Leftarrow \text{step } M_e M'_e$.

To address the second concern, observe that the `%worlds` declaration asserts the theorem only for canonical forms in the empty LF context. The previous adequacy results (*Theorem 3.1*, *Theorem 3.5*, *Theorem 3.11*, and *Theorem 3.15*) show that the empty LF context adequately represents types, closed terms, typing derivations in the empty object-language context, and operational semantics derivations.

4.1.3 Correctness of the Proof of Preservation

In processing the `%total` declaration, Twelf automatically verifies that `preserv` satisfies the totality specification in *Theorem 4.4*. However, we can also prove the totality of this relation by hand.

Proof of Theorem 4.4

```

%% identity (alpha-equivalence) of terms
id  : tm -> tm -> type.
refl : id E E.

%% application congruence lemma for identity
id_app_cong : id E1 E1'
              -> id E2 E2'
              -> id (app E1 E2) (app E1' E2')
              -> type.
%mode id_app_cong +X1 +X2 -X3.

- : id_app_cong refl refl refl.

%worlds () (id_app_cong _ _ _).
%total {} (id_app_cong _ _ _).

%% determinacy of evaluation
det : step E E' -> step E E'' -> id E' E'' -> type.
%mode det +X1 +X2 -X3.

det-1 : det (step_app_1 DstepE1') (step_app_1 DstepE1'') DidApp
       <- det DstepE1' DstepE1'' DidE1
       <- id_app_cong DidE1 refl DidApp.

det-2 : det (step_app_2 _ DstepE2') (step_app_2 _ DstepE2'') DidApp
       <- det DstepE2' DstepE2'' DidE2
       <- id_app_cong refl DidE2 DidApp.

det-b : det (step_app_beta _) (step_app_beta _) refl.

%worlds () (det _ _ _).
%total D (det D _ _).

```

Fig. 15. Twelf Proof of Determinacy of the STLC Operational Semantics

In what follows, we work in $\text{LF}[\Sigma']$. Take Γ to be \cdot . *Lemma 4.5* above justifies appealing to *Lemma 3.10* and *Lemma 3.14* on derivations of $\cdot \vdash_{\Sigma'} M \Leftarrow \text{of } M_e M_t$ and $\cdot \vdash_{\Sigma'} M \Leftarrow \text{step } M_e M'_e$. Because preservation is proved by induction on the dynamic semantics derivation, we can use *Lemma 3.14* on D_{step} . This gives three cases to consider. In each case, we use *Lemma 3.10* to invert D_{of} . In the case for **step_app_1**, the inductive hypothesis gives an inhabitant of the premise of **preserv_app_1**. The case for **step_app_2** is similar. In the case for **step_app_beta**, we use *Theorem 2.11* to instantiate the hypothetical typing derivation for the body of the function—because of the higher-order encoding, substitution for LF is used where one would expect to use substitution for the object language. \square

4.2 Determinacy

Next, we show that the STLC’s operational semantics are deterministic. This example illustrates several additional aspects of Twelf metatheory: we show how to give a simple representation of object-language α -equivalence as an LF type family; we show how to use a lemma in a Twelf proof; and we illustrate a circumstance in which Twelf’s metatheorem checker allows vacuously-true proof cases to be elided. Determinacy is stated as follows:

Theorem 4.6 (Informal Statement of Determinacy)

For closed terms e , e' , and e'' , if $e \mapsto e'$ and $e \mapsto e''$ then $e' =_{\alpha} e''$.

Translating this informal statement into Twelf requires representing α -equivalence of STLC terms as an LF type family. Because of our higher-order representation strategy, α -equivalence is an intrinsic property of the representation: two object-language terms are α -equivalent iff their representations are α -equivalent canonical forms in LF (by *Theorem 3.5*). Consequently, we can represent object-language α -equivalence by internalizing α -equivalence of LF terms as a type family. The type family `id` at the top of Figure 15 does just this. For each E of type `tm`, there is one canonical form inhabiting `id`, `refl E`, expressing reflexivity (recall Twelf’s convention that the variable E in the type of `refl` is implicitly quantified). With this definition, `id E E'` is inhabited exactly when the canonical forms E and E' are α -equivalent in LF. There is no need to give an inductive definition of α -equivalence, as one would give for raw abstract syntax trees.

Using `id`, the informal statement of determinacy can be rephrased as follows:

Theorem 4.7 (Statement of Determinacy via the Encoding)

For all E, E', E'' such that $E, E', E'' \Leftarrow \text{tm}$, D' such that $D' \Leftarrow \text{step } E E'$, and D'' such that $D'' \Leftarrow \text{step } E E''$, there exists a D such that $D \Leftarrow \text{id } E' E''$.

Adequacy (*Theorem 3.5*, *Theorem 3.15*, and the fact that `id` corresponds to object-language α -equivalence) implies that this theorem statement is equivalent to the informal statement in *Theorem 4.6*.

Figure 15 contains a complete Twelf proof of determinacy. The theorem statement in *Theorem 4.7* is represented by the type family `det` and its mode and world declarations in Figure 15, which state the corresponding totality assertion for the relation defined by the constants in the figure.

This proof uses a lemma, represented by the type family `id.app.cong`, its mode and world declarations, and its inhabitants. This lemma states a congruence property of α -equivalence: two applications are identical if their subterms are identical. The proof of this lemma has exactly one case, in which the two subterm equalities were both derived using reflexivity; in this case, the applications are identical, so reflexivity gives the result. (In the Twelf concrete syntax, the application of `refl` to its implicit term argument is suppressed, so all three derivations are written simply as `refl`.) The `%total` declaration for `id.app.cong` uses the empty lexicographic termination metric `{}` because the proof is not inductive. The totality of this lemma justifies using it as a premise in any future proof to produce a derivation

of `id (app E1 E2) (app E1' E2')` from two input derivations of the appropriate type.

The three constants with head `det`, named `det-1`, `det-2`, and `det-b`, define the relation that proves determinacy. In each of these three cases, the two `step` derivations conclude with the same final rule. In the case labeled `det-1` for `step_app_1`, we are given a derivation of the judgement `step (app E1 E2) (app E1' E2)` and a derivation of `step (app E1 E2) (app E1'' E2)` with subderivations of `step E1 E1'` and `step E1 E1''`. The `det` premise represents an appeal to the inductive hypothesis on the two subderivations, concluding that `id E1' E1''`. The `id_app_cong` premise appeals to the lemma on this identity derivation to conclude that the applications are equal. The case `det-2` is similar. When a `step` derivation ends with `step_app_beta`, the left-hand term completely determines the right-hand term, so in the case `det-b` the result is immediate by reflexivity. The `%total` declaration instructs Twelf to check that this type family represents a total relation by induction on the first `step` derivation—though, because of the symmetry, using the second derivation would also suffice. The totality of `id_app_cong` is used as a lemma in showing the totality of `det`—for example, it is used to show the existence of an inhabitant of `id_app_cong DidE1 refl DidApp` given `DidE1` and `refl`.

Twelf successfully proves the totality of `det`, even though the relation contains only cases where the final rule used in both `step` derivations is the same. This is because the off-diagonal cases are all vacuously true—and moreover, Twelf’s metatheorem checker rules out these contradictory cases automatically. For example, if one derivation concluded with `step_app_1` and the other with `step_app_2`, then there would be subderivations concluding both `value E1` and `step E1 E1'`. These two types can never be simultaneously inhabited: `step` is only inhabited when `E1` is an application, and there is no rule inhabiting `value` for an application. Similarly, if one derivation concluded with `step_app_1` and the other with `step_app_beta`, subderivations would give a `step` derivation whose left-hand side is a `lam`, which cannot exist. The other off-diagonal cases can be contradicted in a similar manner. Twelf’s coverage checker rules out cases like these where the inputs to a metatheorem result in an uninhabited instance of some type family (Schürmann & Pfenning, 2003). This feature both eases the development of Twelf proofs and keeps these details from cluttering the final product. As another example, in the companion Twelf code, we give a proof of progress for the STLC that exploits this feature to avoid an explicit canonical forms lemma.

4.3 Strengthening

The strengthening property of the STLC is stated as follows:

Theorem 4.8 (Informal Statement of Strengthening)

If $\gamma, y : \tau_0 \vdash e : \tau$ and $y \# e$ then $\gamma \vdash e : \tau$.

Strengthening is easy to prove if all types are inhabited (simply substitute any term of the appropriate type for `y`), but the proof we give in this section handles uninhabited types as well. This theorem might, naïvely, seem difficult to prove for

our higher-order representation of the object language, as it requires a detailed statement about variables that would seem to require their representation as data. In this section, we give a simple Twelf proof of this property without sacrificing the higher-order representation. Strengthening illustrates three additional features of Twelf metatheory: it is a theorem stated for a non-trivial world; its proof appeals to induction in an extended LF context; and its proof illustrates a common Twelf device for handling the variable case of a theorem.

To formalize the statement of strengthening in Twelf, we will need to capture the condition $y \# e$, that the variable y is not free in the term e . However, under the representation strategy defined in Figure 8, the variable y is free in the term e exactly when the LF variable y is free in the term M_e representing e . Therefore, adequacy (*Theorem 3.1*, *Theorem 3.5*, *Theorem 3.11*) implies that this informal statement of strengthening is equivalent to the following, where we let \mathcal{W}_{of} consist of all contexts Γ such that $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ and we work in $\text{LF}[\Sigma]$, where Σ is the signature for the STLC in Section 3.

Theorem 4.9 (Statement of Strengthening via the Encoding)

For all $\Gamma \in \mathcal{W}_{\text{of}}$, for all E such that $\Gamma \vdash E \Leftarrow \text{tm}$, T and T_0 such that $\Gamma \vdash T, T_0 \Leftarrow \text{tp}$, and D such that $\Gamma, y : \text{tm}, dy : \text{of } y T_0 \vdash D \Leftarrow \text{of } ET$, there exists a D' such that $\Gamma \vdash D' \Leftarrow \text{of } ET$.

Unlike preservation, which was stated for the world $\{\cdot\}$, this theorem is stated for a world \mathcal{W}_{of} containing the representations of object-language typing hypotheses. The world declaration is critical for equivalence with the informal statement in *Theorem 4.8*—for example, if this theorem were stated for the world $\{\cdot\}$ like preservation, the LF statement would only adequately correspond to a weaker theorem about object-language typing derivations in the empty context. Moreover, the proof we give below requires the more general inductive hypothesis, so it would not suffice to prove the weaker theorem, even if that were the informal theorem of interest.

The statement of *Theorem 4.9* is almost in the form supported by Twelf, but the condition $\Gamma, y : \text{tm}, dy : \text{of } y T_0 \vdash D \Leftarrow \text{of } ET$ is in a context other than Γ , which is not permitted. However, it is simple to put it in the correct form by abstracting over the extra variables y and dy , premising the theorem on a term of higher LF type. The revised theorem is as follows:

Theorem 4.10 (Revised Statement of Strengthening)

For all $\Gamma \in \mathcal{W}_{\text{of}}$, for all E such that $\Gamma \vdash E \Leftarrow \text{tm}$, T and T_0 such that $\Gamma \vdash T, T_0 \Leftarrow \text{tp}$, and D such that $\boxed{\Gamma \vdash D \Leftarrow \Pi y : \text{tm}. \text{of } y T_0 \rightarrow \text{of } ET}$, there exists a D' such that $\Gamma \vdash D' \Leftarrow \text{of } ET$.

4.3.1 Twelf Proof

This $\forall\exists$ -statement is translated into the Twelf type family, mode, and worlds declaration in Figure 16. The only subtlety in the type family and mode declarations is that T_0 , E , and T are tacitly universally quantified at the outside by Twelf's implicit

```

strengthen : ({y : tm} of y T0 -> of E T)
             -> of E T
             -> type.
%mode strengthen +X1 -X2.

str-e : strengthen
      ([y] [dy : of y T0] of_empty)
      of_empty.

str-a : strengthen
      ([y] [dy : of y T0]
       (of_app
        ((Dof1 : {y} of y T0 -> of E1 (arrow T2 T)) y dy)
        ((Dof2 : {y} of y T0 -> of E2 T2) y dy)))
      (of_app Dof1' Dof2')
      <- strengthen Dof1 (Dof1' : of E1 (arrow T2 T))
      <- strengthen Dof2 (Dof2' : of E2 T2).

str-l : strengthen
      ([y] [dy : of y T0]
       (of_lam ([x] [dx : of x T2]
                (Dof : {y} of y T0 -> {x} of x T2 -> of (E x) T)
                y dy x dx)))
      (of_lam Dof')
      <- ({x} {dx : of x T2}
          strengthen ([y] [dy : of y T0] Dof y dy x dx)
                    ((Dof' : {x} of x T2 -> of (E x) T) x dx)).

%block ofblock : some {T : tp} block {x : tm} {dx : of x T}.
%worlds (ofblock) (strengthen _ _).
%total D (strengthen D _).

```

Fig. 16. Incomplete First Attempt to Prove Strengthening for the STLC

arguments convention. The `%block` and `%worlds` declaration are Twelf's concrete syntax for defining the world \mathcal{W}_{of} . The block declaration defines a context block called `ofblock` containing the declarations `x : tm, dx : of x T` for some term $T \Leftarrow \text{tp}$. The `%worlds` declaration states `strengthen` for the world of LF contexts containing any number of `ofblocks`. This set of contexts corresponds exactly to those for which the judgement $\gamma \text{ ctx} \gg \Gamma \text{ ctx}$ is derivable.

Figure 16 contains an incomplete first proof attempt; it is instructive to see where this proof fails, and the modification necessary to correct the proof is small. The constants `str-e`, `str-a`, and `str-l` give an inductive definition of the relation `strengthen`. The case `str-e` is simple: when the open typing derivation was derived by the rule `of_empty`, the conclusion can be derived by `of_empty` because this constant does not mention the variables (this corresponds to the fact that the on-paper rule `OF_EMPTY` is stated for an arbitrary context γ). The case `str-a` applies when the open derivation was constructed with the rule `of_app`, in which case the subderivations `Dof1` and `Dof2` can potentially mention `y` and `dy`. However, because

y is not free in (**app** E1 E2), it is not free in the subterms, so by induction on each subderivation (two **strengthen** premises to the constant), we can create derivations of the same facts that do not mention y and dy ; then reapplying the constant of **of_app** gives the result.

The case **str-1** is slightly more involved because it passes under a binder. It is helpful to first consider how this case would proceed on paper. The input is a derivation of

$$\frac{\gamma, y : \tau_0, x : \tau_2 \vdash e : \tau}{\gamma, y : \tau_0 \vdash \lambda x : \tau_2. e : \tau_2 \rightarrow \tau} \text{OF_LAM}$$

where $y \# \lambda x : \tau_2. e$. Under this condition, $y \# e$ as well, so we can use exchange for the object language typing judgement to permute the assumptions $y : \tau_0$ and $x : \tau_2$ and then appeal to the inductive hypothesis on a derivation in the extended context $\Gamma, x : \tau_2, y : \tau_0$ to give a derivation of $\gamma, x : \tau_2 \vdash e : \tau$. Then reapplying **OF_LAM** gives the result.

The Twelf case is precisely analogous. When the last rule applied was **OF_LAM**, the subderivation **Dof** in an extended context is represented by an LF term of function type $\{y\} \text{ of } y \text{ T0} \rightarrow \{x\} \text{ of } x \text{ T2} \rightarrow \text{of } (E \ x) \ T$. Observe that the arguments to the function are the encoding of the assumptions (other than γ) in the premise of **OF_LAM**. Because the term **lam** ($[x] \ E \ x$) is well-formed without y in the context, **E** can only mention the free variable x . The inductive call to **strengthen** takes place in a context extended with $x : \text{tm}$, $dx : \text{of } x \ \text{T2}$; this is represented by giving the constant **str-1** a higher-order premise that abstracts over these two variables. Observe that this context extension stays in \mathcal{W}_{of} ; if this were not the case, Twelf would report an error, signaling that the inductive hypothesis of the totality assertion for this relation, which is stated only for contexts in \mathcal{W}_{of} , would not apply. The use of exchange is implicit in the fact that the strengthening variables for the inductive call are bound inside this dependent function type. The inductive call outputs a derivation **Dof'** that can mention only the bound variables x and dx , and then applying the constant **of_lam** gives the result.

Though each of these three cases is correct, they do not satisfy the totality assertion defined by the mode and worlds declarations. Specifically, this relation does not cover the typing derivation given by a variable dx in the context—there is no inhabitant of the type **strengthen** ($[y] \ [dy] \ dx$) **D** for variables dx . In an informal proof, this is the case for **OF_VAR**: assume $y \# x$ and $\gamma, y : \tau_0 \vdash x : \tau$, then the variable x must be bound in γ , so **OF_VAR** derives the conclusion.

It may not be immediately obvious how to formalize a case for a variable—where can one even mention an arbitrary LF variable dx from the context? In Twelf, such cases can be covered by putting the case for theorem *in the context itself*. Proofs of metatheorems are simply constants of particular types, so the standard apparatus of hypotheses can be deployed to cover cases for variables. For this theorem, whenever we assume a typing derivation dx , we also assume a case of **strengthen** for it. Instead of proving the theorem for contexts in \mathcal{W}_{of} , we prove it for a world \mathcal{W}_{str} , which contains contexts of the form

$$x : \text{tm}, dx : \text{of } x \ \text{T2}, \text{dstr} : \Pi \text{T0} : \text{tp}. \text{strengthen } (\lambda y. \lambda dy. dx) \ dx, \dots$$

```

strengthen : ({y : tm} of y T0 -> of E T)
             -> of E T
             -> type.
%mode strengthen +X1 -X2.

%% str-e and str-a are unchanged

str-l : strengthen
      ([y] [dy] (of_lam ([x] [dx] Dof x dx y dy)))
      (of_lam Dof')
      <- ({x} {dx : of x T2}
          {_ : {T0:tp} strengthen ([y] [dy : of y T0] dx) dx}
          strengthen ([y] [dy] Dof x dx y dy) (Dof' x dx)).

%block strblock : some {T : tp}
                  block {x : tm} {dx : of x T}
                  {_ : {T0:tp} strengthen ([y] [dy : of y T0] dx) dx}.

%worlds (strblock) (strengthen _ _).
%total D (strengthen D _).

```

Fig. 17. Completed Proof of Strengthening for the STLC

As in the informal proof, the variable typing derivation dx is the necessary output in this case. Another way to understand these contexts is to recall that an LF variable stands for all possible substitution instances—by insisting that contexts have this particular form, we are restricting the substitution instances for x and dx to those for which strengthening holds.

Figure 17 contains the completed proof. The theorem statement has changed according to the above discussion: the block `strblock` declares blocks of the above form, and the worlds declaration states the theorem for the world generated by this block. The theorem case in the block defines the relation `strengthen` for the canonical forms that were not covered in the previous attempt. Only one other change to the previous proof is necessary: the inductive call in `str-l` adds the `strengthen` case to the context to stay in the world for which the theorem is stated.

4.3.2 Correctness of the Statement of Strengthening

Let Σ' stand for the extension of the signature Σ from Section 3 with the constants for `strengthen`. The above Twelf proof implies the following statement of strengthening, which is stated in $\text{LF}[\Sigma']$.

Theorem 4.11 (Statement of Strengthening in \mathcal{W}_{str})

For all $\Gamma \in \mathcal{W}_{\text{str}}$, for all E such that $\Gamma \vdash E \Leftarrow \text{tm}$, T and T_0 such that $\Gamma \vdash T, T_0 \Leftarrow \text{tp}$, and D such that $\Gamma \vdash D \Leftarrow \Pi y:\text{tm}. \text{of } y T_0 \rightarrow \text{of } E T$, there exists a D' such that $\Gamma \vdash D' \Leftarrow \text{of } E T$.

```

weaken : {T0} of E T -> ({x} (of x T0) -> of E T) -> type.
%mode weaken +X1 +X2 -X3.

- : weaken T0 D ([x] [dx] D).

%worlds (ofblock) (weaken _ _ _).
%total {} (weaken _ _ _).

```

Fig. 18. Direct Twelf Proof of Weakening for the STLC

This theorem differs from our desired theorem statement (*Theorem 4.10*) in two ways: it is stated for an extended signature Σ' and for a different world \mathcal{W}_{str} . Thus, it is possible that we successfully proved the wrong theorem. For example, *Theorem 4.11* might not cover all object-language typing derivations, or it might assume some additional typing derivations that have no informal counterpart. To assuage these concerns, we should check that *Theorem 4.11* implies *Theorem 4.10*. Because the only differences between these two theorems are the LF signature and world, it should not be surprising that the content of this theorem is a transport of adequacy result:

Lemma 4.12 (Transport of Adequacy for strengthen)

1. $\Sigma'|_{\text{of}} = \Sigma|_{\text{of}}$.
2. For all $\Gamma' \in \mathcal{W}_{\text{of}}$, there exists a $\Gamma \in \mathcal{W}_{\text{str}}$ such that $\Gamma|_{\text{of}}^{\preceq_{\Sigma'}} = \Gamma'$.

Proof

Observe that $\preceq_{\Sigma'}$ adds the edges $\text{tm} \preceq_{\Sigma'} \text{strengthen}$, $\text{tp} \preceq_{\Sigma'} \text{strengthen}$, and $\text{of} \preceq_{\Sigma'} \text{strengthen}$ to \preceq_{Σ} . The first part is true by calculation. The second part can be proved by a simple induction on the number of blocks in Γ . For any ofblock of the form $x : \text{tm}, dx : \text{of } x \text{ T}$ there is a strblock

$$x : \text{tm}, dx : \text{of } x \text{ T}, \text{strx} : \Pi T_0 : \text{tp. strengthen } (\lambda y. \lambda dy. dx) dx$$

because the indices to the theorem case for strengthen mention only terms bound in the ofblock ; the restriction of this block to of is the original block. \square

Using this lemma and *Theorem 2.17*, the proof that *Theorem 4.11* implies *Theorem 4.10* is direct. Thus, the Twelf proof does indeed prove the informal statement that we wanted to show.

4.3.3 Contrasting Strengthening with Weakening

Whereas strengthening requires an inductive proof, weakening can be proved directly, as we show in Figure 18. The type argument $T0$ is made explicit so that it can be universally quantified; otherwise, it would be existentially quantified because it would appear only in an output. The proof of weaken includes only one case, in which, given a derivation of $\text{of } E \text{ T}$, we introduce functions to create a derivation that is abstracted over x and dx . The fact that this LF constant is well-typed

```

%% previously proved congruence lemma for identity
id_app_cong : id E1 E1'
             -> id E2 E2'
             -> id (app E1 E2) (app E1' E2')
             -> type.
%mode id_app_cong +X1 +X2 -X3.
%worlds () (id_app_cong _ _ _).

%% all free variables of a well-typed term must have a typing assumption
%% in the context
all_declared : ({y : tm} of (E y) T)
              -> ({y : tm} id (E y) E')
              -> type.
%mode all_declared +X1 -X2.

- : all_declared ([y] of_app (D1 y) (D2 y)) DidApp
  <- all_declared D1 (Did1 : {y:tm} id (E1 y) E1')
  <- all_declared D2 (Did2 : {y:tm} id (E2 y) E2')
  <- ({y : tm} id_app_cong (Did1 y) (Did2 y) (DidApp y)).

...

%worlds (ofblock) (all_declared _ _).
%total D (all_declared D _).

```

Fig. 19. Excerpt of Twelf Proof of Motivating Example for World Subsumption

corresponds with weakening being admissible in LF: the use of the variable D is insensitive to which other variables are in the LF context, so we may derive

$$\dots, D : \text{of } E \ T, x : \text{tm}, dx : \text{of } x \ T_0 \vdash D \Rightarrow \text{of } E \ T$$

by `ATOM_TERM_VAR`, and then use `CANON_TERM_ATOM` and `CANON_TERM_LAM` to introduce the functions.

It may at first seem curious that there is such a direct proof of weakening but not of strengthening—why does strengthening for LF not give the result? In this instance, strengthening for LF would prove the following:

If $\Gamma, y : \text{tm}, dy : \text{of } y \ T_0 \vdash D \Leftarrow \text{of } E \ T$ and neither y nor dy are free in E or D , then $\Gamma \vdash D \Leftarrow \text{of } E \ T$.

That is, strengthening in LF only allows the variables to be dropped when they are known not to be free *in the derivation* D . In the statement of strengthening in *Theorem 4.9*, the variable y is not free in the term E , but both variables may appear in the derivation D . This stronger statement of strengthening requires an inductive proof even when the weaker version does not; indeed, one can imagine object languages where the weaker statement is true but the stronger one is not.

4.4 World Subsumption

4.4.1 Motivating Example

When proving theorems that are stated for a non-trivial LF world, it is common to want to use a lemma that is stated for one world in the proof of a theorem stated for another. We illustrate this with the following (somewhat contrived) theorem:

Theorem 4.13

If $\gamma \vdash e : \tau$, where $\mathcal{X}, y, \mathcal{X}' \vdash e$ term but y is not in γ , then $y \# e$.

That is, all free variables in a well-typed term must be declared in the context.

Figure 19 contains an excerpt of a Twelf proof of this theorem (see the companion Twelf code for the remaining cases, which are simple but not necessary for the discussion in this section). The premise that $\gamma \vdash e : \tau$, where y is potentially free in e but is not declared in γ , is represented by the input LF term of type $(\{y : \mathbf{tm}\} \text{ of } (\mathbf{E} \ y) \ \mathbf{T})$ —the term \mathbf{E} may potentially mention the variable y , but there is no typing derivation for y given as a hypothesis. The conclusion, that $y \# e$ is represented by the output LF term of type $(\{y : \mathbf{tm}\} \text{ id } (\mathbf{E} \ y) \ \mathbf{E}')$ —the term \mathbf{E} , which may mention the variable y , is α -equivalent to a term \mathbf{E}' in which y does not occur (recall the definition of id in Section 4.2). To adequately represent the informal statement, which is stated for arbitrary object-language contexts γ , the Twelf statement is stated for the world consisting of `ofblocks`.

The excerpt of the proof shows the case when the typing derivation was derived using `of_app`. Appealing to the inductive hypothesis on the two subderivations shows that y is not free in either $\mathbf{E1}$ (with evidence `Did1`) or $\mathbf{E2}$ (with evidence `Did2`), and we must use these facts to show that it is not free in `app E1 E2`. To do so, the case appeals to the congruence lemma `id_app_cong` proved in Figure 15 (for reference, the statement of this lemma is reiterated at the top of Figure 19). This call is in an extended context binding the variable y .

Should this call be allowed? The current theorem, `all_declared`, is stated for the world \mathcal{W}_{of} , whereas the lemma, `id_app_cong`, was proved in the world $\{\cdot\}$. Unfortunately, a lemma proved for one world may not be true in another, and, in this case, there is indeed a problem. The world \mathcal{W}_{of} contains additional canonical forms of type \mathbf{tm} that are not present in $\{\cdot\}$, so it is possible that there are input derivations of `id EE'` in \mathcal{W}_{of} that `id_app_cong` is not prepared to handle. More formally, the totality assertion for `id_app_cong` is the following:

For all $\Gamma \in \{\cdot\}$, if $\Gamma \vdash D1 \Leftarrow \text{id } \mathbf{E1} \ \mathbf{E1}'$ and $\Gamma \vdash D2 \Leftarrow \text{id } \mathbf{E2} \ \mathbf{E2}'$ then there exist $D3$ and D such that $\Gamma \vdash D3 \Leftarrow \text{id } (\text{app } \mathbf{E1} \ \mathbf{E2}) \ (\text{app } \mathbf{E1}' \ \mathbf{E2}')$ and $\Gamma \vdash D \Leftarrow \text{id_app_cong } D1 \ D2 \ D3$.

Appealing to this theorem from \mathcal{W}_{of} instantiates Γ with a context other than \cdot . This is impermissible, absent some additional reasoning justifying that the theorem still holds in the extended world.

To solve this problem, we can prove `id_app_cong` in the larger world. Intuitively, we need to show that the application congruence rule for `id` is admissible not just on closed terms, but also on open terms. Fortunately, our original proof does

```

id_app_cong : id E1 E1'
  -> id E2 E2'
  -> id (app E1 E2) (app E1' E2')
  -> type.
%mode id_app_cong +X1 +X2 -X3.

- : id_app_cong refl refl refl.

%worlds (ofblock) (id_app_cong _ _ _).
%total {} (id_app_cong _ _ _).

```

Fig. 20. Twelf Proof of Lemma `id_app_cong` in \mathcal{W}_{of}

indeed prove the stronger theorem—in Section 4.2 we were concerned only with closed terms, so we did not consider the extended theorem statement. This can be expressed with the Twelf code in Figure 20. The only difference from the proof in Figure 15 is the worlds declaration, which states the theorem for \mathcal{W}_{of} . The proof goes through as before.

Unfortunately, this revised theorem is still insufficient for two reasons. The first is that the call to it from the proof of `all_declared` is not in \mathcal{W}_{of} , as the extended context contains a variable `y` but no typing derivation. This problem can be solved by revising the case of the theorem to add an unused typing hypothesis for `y` to the context. The more serious issue is one of modularity: the statement and proof of `id_app_cong` have nothing to do with typing derivations, but yet the lemma needed to be stated for a world containing typing derivations to support its later use. This issue is analogous to modularity issues we encountered with adequacy proofs, and we can deploy the same machinery to solve it.

4.4.2 Definition of World Subsumption

Calling a lemma stated for a world \mathcal{W} from a theorem stated for a world \mathcal{W}' is not always permissible: it is possible that \mathcal{W}' contains additional canonical forms that the lemma is not prepared to handle (as in the above example), or that the lemma outputs canonical forms in \mathcal{W} that do not exist in \mathcal{W}' . In either case, the lemma proved in \mathcal{W} may no longer be true in \mathcal{W}' . Fortunately, the same techniques that we used to transfer adequacy from one world to another can be used to transfer proofs of metatheorems from one world to another—the key issue in both cases is the characterization of the relevant canonical forms. In particular, we define an order on worlds as follows, where the relation $\Gamma \sim \Gamma'$ is identity modulo permutation of independent hypotheses:

Definition 4.14 (World Subsumption)

$\mathcal{W} \lesssim_a \mathcal{W}'$ iff for all $\Gamma' \in \mathcal{W}'$, there exists a $\Gamma \in \mathcal{W}$ such that $\Gamma|_a \sim \Gamma'|_a$.

Because the proof that one world subsumes another may be non-trivial, Twelf implements a conservative but useful approximation of the relation $\mathcal{W} \lesssim_a \mathcal{W}'$.

The intuitive justification for choosing this orientation of the relation $\lesssim_{\mathbf{a}}$ is that world subsumption is a sufficient condition for weakening the world of a totality assertion (i.e., for transporting a theorem proved in a smaller world to a larger one):

Theorem 4.15 (Totality Assertions Remain True in Larger Worlds)

Assume an LF constant $\mathbf{a} \Rightarrow \Pi \mathbf{x}_1:\mathbf{A}_1 \dots \Pi \mathbf{y}_1:\mathbf{B}_1 \dots \mathbf{type}$ that satisfies the following totality assertion:

For all $\Gamma \in \mathcal{W}$, for all \mathbf{M}_1 such that $\Gamma \vdash \mathbf{M}_1 \Leftarrow \mathbf{A}'_1, \dots$, and \mathbf{M}_m such that $\Gamma \vdash \mathbf{M}_m \Leftarrow \mathbf{A}'_m$, there exist \mathbf{N}_1 such that $\Gamma \vdash \mathbf{N}_1 \Leftarrow \mathbf{B}'_1, \dots$, and \mathbf{N}_n such that $\Gamma \vdash \mathbf{N}_n \Leftarrow \mathbf{B}'_n$, and \mathbf{D} such that $\Gamma \vdash \mathbf{D} \Leftarrow \mathbf{a} \mathbf{M}_1 \dots \mathbf{M}_m \mathbf{N}_1 \dots \mathbf{N}_n$.

where

$$[\mathbf{M}_{i-1}/\mathbf{x}_{i-1}]_{\mathbf{A}_{i-1}} \dots [\mathbf{M}_1/\mathbf{x}_1]_{\mathbf{A}_1} \mathbf{A}_i = \mathbf{A}'_i \\ [\mathbf{N}_{j-1}/\mathbf{y}_{j-1}]_{\mathbf{B}_{j-1}} \dots [\mathbf{N}_1/\mathbf{y}_1]_{\mathbf{B}_1} [\mathbf{M}_m/\mathbf{x}_m]_{\mathbf{A}_m} \dots [\mathbf{M}_1/\mathbf{x}_1]_{\mathbf{A}_1} \mathbf{B}_j = \mathbf{B}'_j.$$

Then for all \mathcal{W}' such that $\mathcal{W} \lesssim_{\text{thm}} \mathcal{W}'$, the totality assertion for \mathcal{W}' is also true.

Proof

The proof is direct using *Theorem 2.17*, *Theorem 2.11*, and two additional lemmas:

1. If $\Gamma \sim \Gamma'$ then $\Gamma|_{\mathbf{a}} \sim \Gamma'|_{\mathbf{a}}$.
2. If $\Gamma \vdash \mathbf{M} \Leftarrow \mathbf{A}$ and $\Gamma \sim \Gamma'$ then $\Gamma' \vdash \mathbf{M} \Leftarrow \mathbf{A}$.

The proof of the second part uses exchange (*Lemma 2.7*). \square

Intuitively, the condition $\mathcal{W} \lesssim_{\mathbf{a}} \mathcal{W}'$ ensures that all canonical forms of type \mathbf{a} in \mathcal{W}' also exist in \mathcal{W} . This condition is necessary to use a lemma proved for \mathcal{W} from a theorem stated for \mathcal{W}' , as it ensures that \mathcal{W}' does not add additional inputs that the lemma is not prepared to handle. However, the condition $\mathcal{W} \lesssim_{\mathbf{a}} \mathcal{W}'$ does not imply that all of the canonical forms of type \mathbf{a} in \mathcal{W} exist in \mathcal{W}' . This may intuitively seem problematic for the above theorem—what if the lemma returns a canonical form in \mathcal{W} that does not exist in \mathcal{W}' ? The lemma cannot return such a derivation because totality assertions quantify over the context Γ at the outside: when the lemma is given canonical inputs in a particular Γ , it returns canonical outputs in the same Γ .

To understand *Theorem 4.15*, it may be helpful to consider an example in which it does not apply. Recall the transport of adequacy theorem for strengthening, *Lemma 4.12* above. The second part of this lemma is equivalent to stating that $\mathcal{W}_{\text{str}} \lesssim_{\text{of}} \mathcal{W}_{\text{of}}$. Moreover, the proof that the statement of strengthening in \mathcal{W}_{str} (*Theorem 4.11*) implies the statement of strengthening in \mathcal{W}_{of} (*Theorem 4.10*) is similar to the proof of *Theorem 4.15*. However, we cannot use *Theorem 4.15* to give an alternate proof of this theorem. *Lemma 4.12* proves $\mathcal{W}_{\text{str}} \lesssim_{\text{of}} \mathcal{W}_{\text{of}}$, but it does not prove $\mathcal{W}_{\text{str}} \lesssim_{\text{strengthen}} \mathcal{W}_{\text{of}}$, which would be necessary to satisfy the premise of *Theorem 4.15* (in fact, it is not true that $\mathcal{W}_{\text{str}} \lesssim_{\text{strengthen}} \mathcal{W}_{\text{of}}$). In this example, the $\forall\exists$ -statement of strengthening transfers from \mathcal{W}_{str} to \mathcal{W}_{of} , as we proved above, but the totality assertion for **strengthen** does not—precisely because the contexts in \mathcal{W}_{str} contribute proof cases to **strengthen**.

```

id_app_cong : id E1 E1'
             -> id E2 E2'
             -> id (app E1 E2) (app E1' E2')
             -> type.
%mode id_app_cong +X1 +X2 -X3.

- : id_app_cong refl refl refl.

%block tmblock : block {x : tm}.
%worlds (tmblock) (id_app_cong _ _ _).
%total {} (id_app_cong _ _ _).

```

Fig. 21. Twelf Proof of Lemma `id_app_cong` in \mathcal{W}_{tm}

4.4.3 Using World Subsumption

The proof of `all_declared` at the beginning of this section required a lemma asserting that identity is a congruence on open terms. The correct statement and proof of this lemma is presented in Figure 21. The theorem must be stated for the world containing term variables (call this world \mathcal{W}_{tm}) to account for open terms. However, unlike the previous statement, this theorem does not mention typing derivations, which are irrelevant to this particular lemma and its proof. It is easy to show that $\mathcal{W}_{\text{tm}} \lesssim_{\text{id_app_cong}} \mathcal{W}_{\text{of}}$ of $\not\leq \text{id_app_cong}$, so for any context Γ of the form $x : \text{tm}, dx : \text{of } x \text{ T}, \dots$, it is true that $\Gamma|_{\text{id_app_cong}} \in \mathcal{W}_{\text{tm}}$. Thus, *Theorem 4.15* implies that `id_app_cong` holds in \mathcal{W}_{of} as well. This authorizes the call from the proof of `all_declared` in Figure 19.

Observe that $\{\cdot\} \not\lesssim_{\text{id_app_cong}} \mathcal{W}_{\text{of}}$ because $\text{tm} \not\leq \text{id_app_cong}$, so *Theorem 4.15* does not permit the theorem to call the weaker version of the lemma, which was stated only for the empty context. Of course, had we anticipated needing the open-term congruence lemma, we could have proved that lemma in the first place, as it includes closed terms as a special case. It is worth noting that, though these examples have to do with worlds, the general principle has nothing to do with Twelf: in a paper proof, one might choose to prove a weaker version of a lemma for use at one point in a proof, and then, later in the proof, discover that a stronger version is necessary.

The world subordination criterion emphasizes that transport of canonical forms is important not just for modular adequacy proofs, but also for modular Twelf proofs.

4.5 Discussion

Though Twelf *proofs* are mechanically checked, a programmer still must verify that a Twelf theorem *statement* corresponds to the informal statement of the theorem he wants to prove. As we have seen in this section, the adequacy methodology permits a precise account of the correspondence between informal and mechanized theorem statements. That said, a Twelf programmer does not need to make this correspondence for every theorem he proves. In a large Twelf proof, there are usually

a few top-level theorems stating the major results about a system (e.g., progress and preservation for a language) that are proved using many lemmas. In this case, it is necessary only to verify that the top-level theorem statements correspond to the desired informal theorems; there is no need to consider the informal analogues of lemmas that are not of independent interest.

5 Related Work

A reader who is interested in learning more about Twelf should visit the Twelf Wiki (<http://twelf.plparty.org/>), which contains numerous tutorials and case studies.

LF and Twelf. Harper *et al.* (1993) introduce the LF type theory and representation methodology. The canonical-forms presentation of LF in this article is due to the treatment of Concurrent LF by Watkins *et al.* (2002, 2004a); it has since been applied to several other type theories (for example, see Nanevski *et al.* (2007); Nanevski & Morrisett (2006); Lee *et al.* (2007)). The general method of defining a type theory algorithmically goes back to the AUTOMATH languages (de Bruijn, 1993; van Daalen, 1980). Previous treatments of the metatheory of LF include the following: Harper *et al.* (1993) discuss the metatheory of LF with β -conversion but not η -conversion. Salvesen (1990), Geuvers (1992), and Goguen (1999) discuss the extension to $\beta\eta$ -conversion. Coquand (1991) gives a shape-directed equality algorithm for LF. Harper & Pfenning (2005) present a type-directed equality algorithm. Felty (1991) presents LF with formation judgements that admit only canonical forms, but defines equality using β -reduction on non-canonical forms.

Several logical frameworks extend LF with additional features. Linear LF (LLF) extends LF with linear connectives (Cervesato & Pfenning, 2002). Concurrent LF (CLF) extends LLF with more linear connectives and with an intrinsic notion of concurrency (Watkins *et al.*, 2002, 2004b; Cervesato *et al.*, 2002). These extensions ease the representation of language features for which LF provides no native support—e.g., a language with state can conveniently be represented using the linear connectives in LLF.

The Twelf User’s Guide (Pfenning & Schürmann, 2002) describes the features of Twelf. The metatheory of Twelf is discussed in several sources. Pfenning (1991) introduces the logic-programming operational semantics for LF. Pfenning & Rohwedder (1992) give an early overview of the Twelf methodology. Rohwedder & Pfenning (1996) discuss mode and termination checking. Pientka & Pfenning (2000) discuss an extended termination checker. Schürmann (2000) and Schürmann & Pfenning (2003) discuss coverage and world checking.

Several extensions of Twelf have been studied but not yet fully implemented. Schürmann & Pfenning (1998) discuss a metatheorem prover. Anderson & Pfenning (2004) discuss uniqueness checking. Schürmann *et al.* (2005) describe an extended metatheorem language that would permit more general theorems than $\forall\exists$ -statements. Reed (2006) describes an approach to extending the metatheoretic capabilities of Twelf to Linear LF.

Applications of LF and Twelf. Several papers discuss applications of LF and Twelf. Lee *et al.* (2007) discuss a Twelf proof of type safety for the Harper-Stone internal language for Standard ML. Avron *et al.* (1989) discuss LF representation and present several examples. Michaylov & Pfenning (1991) present some of the metatheory of Mini-ML. Pfenning (1992) gives a proof of the Church-Rosser theorem for the simply typed λ -calculus. Pfenning (1994) gives a proof of cut elimination for intuitionistic logic. Schürmann *et al.* (2001) present some of the metatheory of F^ω . Schürmann & Stehr (2005) present a formalization of Howe’s embedding of HOL into NuPRL. Appel (2001) presents a foundational proof-carrying code system. Appel & Felty (2002) use Twelf to implement tactical theorem provers. Crary & Sarkar (2003) present a typed assembly language and applications to proof-carrying code. Simmons (2005) presents the metatheory of a functional language with references. Pfenning (1999) surveys various methods of representing deductive systems in logical frameworks. The `examples` directory of the Twelf distribution contains additional examples of deductive systems and their metatheory.

Though they do not focus on mechanization, several additional papers are accompanied by Twelf appendices mechanizing the metatheory of the presented languages. Crary (2003) formalizes a foundational typed assembly language. Murphy VII *et al.* (2004) and Murphy VII *et al.* (2005) formalize modal type systems for distributed computing. Licata & Harper (2005) formalize a language extending ML with a form of dependent types. Garg & Pfenning (2006) formalize an authorization logic. Acar *et al.* (2007) prove consistency and correctness of a semantics for self-adjusting computation. Fluett *et al.* (2006) formalize a substructural type system for region-based memory management.

Other Mechanization Tools. The LF logical framework is similar in spirit to the AUTOMATH languages (Nederpelt *et al.*, 1994)—both provide frameworks for representing machine-checkable mathematical arguments, while neither make any foundational commitment about what logics can be represented.

Other proof assistants that have been used to formalize mathematics and computer science include Coq (Coq Development Team, 2007; Bertot & Castéran, 2004), Isabelle/HOL (Nipkow *et al.*, 2002), NuPRL (Constable *et al.*, 1986), HOL Light (Harrison, 1996), and ACL2 (Kaufmann *et al.*, 2000). Many of these proof assistants have been applied to the domain of programming languages and logics. For example, Miculan (1997) shows how to encode several logical systems in Coq; Leroy (2006) reports on implementing a certified compiler in Coq. Klein & Nipkow (2006) formalize a semantics for a Java-like language in Isabelle/HOL.

Twelf, Coq, Isabelle/HOL, and other proof assistants differ significantly in the representation techniques they support, the classes of metatheorems that can be proved, the style in which proofs are written, and the degree of automation they provide. A detailed analysis of the trade-offs between these tools is beyond the scope of this article. In the literature, several efforts have been made to compare and contrast proofs of the same theorem formalized in different proof assistants. For example, the POPLmark Challenge (Aydemir *et al.*, 2005) is one benchmark on which proof assistants can be compared; solutions have been presented in Twelf,

Coq, Isabelle/HOL, and other systems. Appel & Leroy (2006) compare Coq and Twelf proofs of a first-order list-machine benchmark.

6 Conclusion

In this article, we have surveyed the LF λ -calculus, the LF methodology for representing languages, and the Twelf methodology for mechanizing metatheory. Following Watkins *et al.* (2002), we have shown how LF, a minimal dependent type theory, can be defined so that only canonical forms are well-typed, giving a direct inductive definition of the canonical forms. We have shown how the minimality of LF enables higher-order representations of syntax and judgements, whereby an object language inherits α -equivalence, capture-avoiding substitution, and properties of hypothetical judgements from the logical framework. We have shown how these representations are proved adequate by simple structural inductions on the informal object-language entities and on the canonical forms of LF; and we have shown how subordination-based transport of adequacy facilitates modular adequacy proofs. Finally, we have shown how the proofs of $\forall\exists$ -statements over LF types can be represented relationally within LF and mechanically verified by Twelf, again using subordination and induction on canonical forms.

There are several important and useful features of Twelf that we have not discussed in this article. Chiefly, Twelf implements a logic-programming operational semantics for LF (Pfenning, 1991). Using this operational semantics, one can run the LF specification of a programming language directly. For example, the type family `step` defined in Figure 13 can be run as an interpreter for the STLC, and the type family `of` can be run as a type checker. The interested reader should consult the Twelf User’s Guide (Pfenning & Schürmann, 2002) or the Twelf Wiki to learn about the additional features of Twelf.

In future work, we expect that the methodologies that we have surveyed in this paper will be scaled to richer logical frameworks, to tools that provide more automated theorem proving, and to more-general metatheorem languages. Research on Linear LF and Concurrent LF has shown how these frameworks permit facile representations of systems that are cumbersome to represent in LF (Cervesato & Pfenning, 2002; Cervesato *et al.*, 2002). The LF methodology for adequate representations scales to these richer frameworks; the programmer simply has a richer collection of types available for generating canonical forms. The extension of Twelf’s logic-programming operational semantics and totality checker to these logical frameworks is current and future work (Reed, 2006).

At present, Twelf is most often used as a metatheorem checker rather than as a metatheorem prover, in the sense that most often a proof of a $\forall\exists$ -statement is represented explicitly as an LF type family and then verified by a totality assertion. A metatheorem prover (i.e., an extension of Twelf with the ability to automatically generate an LF type family satisfying a given totality assertion) has been studied and implemented in a previous release of Twelf (Schürmann & Pfenning, 1998). However, the current Twelf implementation of this metatheorem prover does not output the LF type family that it discovers, limiting the feature’s utility. This

metatheorem prover’s implementation may eventually be completed, and the development of more-sophisticated metatheorem provers is a subject of active research.

Because Twelf metatheorems are proved using totality assertions about LF type families, the class of metatheorems that can be mechanized is restricted to $\forall\exists$ -statements over LF types. On the one hand, as the successful Twelf formalizations cited in Section 5 demonstrate, these $\forall\exists$ -statements have proved to be sufficient for formalizing a wide variety of metatheorems about programming languages and logics. On the other hand, we have no way to quantify when metatheorems of this form will be sufficient, and there are some well-known examples of proofs that cannot be formalized directly using Twelf’s metatheorem language. For example, proofs by logical relations often require more quantifier complexity than $\forall\exists$ -statements afford. In future work, we expect that this restriction may be lifted by moving to a different metatheorem language that permits more general statements (see Schürmann *et al.* (2005) for some preliminary work). Though such a metatheorem language is a departure from the specific methodology outlined in this paper, it is in essence just a different way of expressing proofs by induction on the canonical forms of LF.

In summary, we believe that the LF and Twelf methodology for mechanizing languages and their metatheory is both useful in practice today and the foundation of interesting future research. We hope that the interested reader will use LF and Twelf and assess their utility himself.

Acknowledgements. We thank Todd Wilson, Rob Simmons, Steve Brookes, William Lovas, Tom Murphy, Jason Reed, Ruy Ley-Wild, Daniel Lee, Daniel Spoonhower, Akiva Leffert, Neel Krishnaswami, Jake Donham, Sean McLaughlin, Karl Crary, and Frank Pfenning for discussions about, and feedback on, previous drafts of this article. We thank the anonymous referees for their helpful feedback on previous drafts.

References

- Acar, Umut A., Blume, Matthias, & Donham, Jacob. (2007). A consistent semantics of self-adjusting computation. *European Symposium on Programming*. Springer-Verlag.
- Anderson, Penny, & Pfenning, Frank. (2004). Verifying uniqueness in a logical framework. *Pages 18–33 of: Slind, K., Bunker, A., & Gopalakrishnan, G. (eds), International Conference on Theorem Proving in Higher-Order Logics*. Lecture Notes in Computer Science, vol. 3223. Springer.
- Appel, Andrew, & Leroy, Xavier. (2006). A list-machine benchmark for mechanized metatheory. *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice*. Electronic Notes in Theoretical Computer Science.
- Appel, Andrew W. (2001). Foundational proof-carrying code. *Page 247 of: IEEE Symposium on Logic in Computer Science*. IEEE Computer Society.
- Appel, Andrew W., & Felty, Amy P. (2002). Dependent types ensure partial correctness of theorem provers. *Journal of Functional Programming*, **14**(1), 3–19.

- Avron, Arnon, Honsell, Furio, & Mason, Ian A. (1989). An overview of the Edinburgh Logical Framework. Birtwistle, Graham, & Subrahmanyam, P.A. (eds), *Current trends in hardware verification*. Springer Verlag.
- Aydemir, Brian E., Bohannon, Aaron, Fairbairn, Matthew, Foster, J. Nathan, Pierce, Benjamin C., Sewell, Peter, Vytiniotis, Dimitrios, Washburn, Geoffrey, Weirich, Stephanie, & Zdancewic, Steve. (2005). Mechanized metatheory for the masses: The POPLmark challenge. *International Conference on Theorem Proving in Higher-Order Logics*. Springer-Verlag.
- Bertot, Yves, & Castéran, Pierre. (2004). *Interactive theorem proving and program development: Coq'art: The calculus of inductive constructions*. Texts in Theoretical Computer Science. Springer.
- Cervesato, Iliano, & Pfenning, Frank. (2002). A linear logical framework. *Information and Computation*, **179**(1), 19–75.
- Cervesato, Iliano, Pfenning, Frank, Walker, David, & Watkins, Kevin. (2002). *A concurrent logical framework II: Examples and applications*. Tech. rept. CMU-CS-02-102. Department of Computer Science, Carnegie Mellon University. Revised May 2003.
- Constable, Robert L., Allen, Stuart F., Bromley, H. M., Cleaveland, W. R., Cremer, J. F., Harper, R. W., Howe, Douglas J., Knoblock, T. B., Mendler, N. P., Panangaden, P., Sasaki, James T., & Smith, Scott F. (1986). *Implementing mathematics with the NuPRL proof development system*. Prentice Hall.
- Coq Development Team. (2007). *The Coq proof assistant reference manual*. INRIA. Available from <http://coq.inria.fr/>.
- Coquand, Thierry. (1991). An algorithm for testing conversion in type theory. *Pages 255–279 of: Huet, Gérard, & Plotkin, Gordon D. (eds), Logical frameworks*. New York, NY, USA: Cambridge University Press.
- Crary, Karl. (2003). Toward a foundational typed assembly language. *Pages 198–212 of: ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press.
- Crary, Karl, & Sarkar, Susmit. (2003). Foundational certified code in a metalogical framework. *Pages 106–120 of: International Conference on Automated Deduction*. Springer-Verlag.
- de Bruijn, Nicolaas Govert. (1993). Algorithmic definition of lambda-typed lambda calculus. *Pages 131–145 of: Huet, Gérard, & Plotkin, Gordon D. (eds), Logical environment*. Cambridge University Press.
- Felty, Amy. (1991). Encoding dependent types in an intuitionistic logic. *Pages 214–251 of: Huet, Gérard, & Plotkin, Gordon D. (eds), Logical frameworks*. Cambridge University Press.
- Fluet, Matthew, Morrisett, Greg, & Ahmed, Amal. (2006). Linear regions are all you need. *Pages 7–21 of: European Symposium on Programming*. Springer-Verlag.
- Garg, Deepak, & Pfenning, Frank. (2006). Non-interference in constructive authorization logic. *Pages 183–293 of: Computer security foundations workshop*.
- Geuvers, Herman. (1992). The Church-Rosser property for $\beta\eta$ -reduction in typed λ -calculi. *Pages 453–460 of: Scedrov, A. (ed), IEEE Symposium on Logic in Computer Science*. IEEE Press.

- Goguen, Healfdene. (1999). Soundness of the logical framework for its typed operational semantics. *International Conference on Typed Lambda Calculi and Applications*. Lecture Notes in Computer Science, vol. 1581. Springer-Verlag.
- Harper, Robert, & Pfenning, Frank. (2005). On equivalence and canonical forms in the LF type theory. *ACM Transactions on Computational Logic*, **6**, 61–101.
- Harper, Robert, Honsell, Furio, & Plotkin, Gordon. (1993). A framework for defining logics. *Journal of the Association for Computing Machinery*, **40**(1).
- Harrison, John. (1996). HOL Light: A tutorial introduction. *Pages 265–269 of: Formal Methods in Computer-Aided Design*. Lecture Notes in Computer Science, vol. 1166. Springer-Verlag.
- Kaufmann, Matt, Manolios, Panagiotis, & Moore, J Strother. (2000). *Computer-aided reasoning: An approach*. Kluwer Academic Publishers.
- Klein, Gerwin, & Nipkow, Tobias. (2006). A machine-checked model for a Java-like language, virtual machine and compiler. *ACM Transactions on Programming Languages and Systems*, **28**(4), 619–695.
- Lee, Daniel K., Crary, Karl, & Harper, Robert. (2007). Towards a mechanized metatheory of Standard ML. *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press.
- Leroy, Xavier. (2006). Formal certification of a compiler back-end, or: programming a compiler with a proof assistant. *Pages 42–54 of: ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press.
- Licata, Daniel R., & Harper, Robert. (2005). *A formulation of Dependent ML with explicit equality proofs*. Tech. rept. CMU-CS-05-178. Department of Computer Science, Carnegie Mellon University.
- Michaylov, Spiro, & Pfenning, Frank. (1991). Natural semantics and some of its meta-theory in Elf. *Pages 299–344 of: Eriksson, L.H., Hallnäs, L., & Schroeder-Heister, P. (eds), International Workshop on Extensions of Logic Programming*. Lecture Notes in Artificial Intelligence, vol. 596. Springer-Verlag.
- Miculan, Marino. (1997). *Encoding logical theories of programs*. Ph.D. thesis, Dipartimento di Informatica, Università di Pisa.
- Murphy VII, Tom, Crary, Karl, Harper, Robert, & Pfenning, Frank. (2004). A symmetric modal lambda calculus for distributed computing. *Pages 286–295 of: Ganzinger, H. (ed), IEEE Symposium on Logic in Computer Science*. IEEE Press.
- Murphy VII, Tom, Crary, Karl, & Harper, Robert. (2005). Distributed control flow with classical modal logic. *Pages 51–69 of: Computer Science Logic*. Lecture Notes in Computer Science, vol. 3634. Springer-Verlag.
- Nanevski, Aleksandar, & Morrisett, Greg. (2006). *Dependent type theory of stateful higher-order functions*. Tech. rept. TR-24-05. Harvard Computer Science.
- Nanevski, Aleksandar, Pfenning, Frank, & Pientka, Brigitte. (2007). Contextual modal type theory. *ACM Transactions on Computational Logic*. To appear.
- Nederpelt, R.P., Geuvers, J.H., & de Vrijer, R.C. (eds). (1994). *Selected papers on AUTOMATH*. Studies in Logic and the Foundations of Mathematics, vol. 133. North-Holland.
- Nipkow, Tobias, Paulson, Lawrence C., & Wenzel, Markus. (2002). *Isabelle/HOL*

- *a proof assistant for higher-order logic*. Lecture Notes in Computer Science, vol. 2283. Springer-Verlag.
- Pfenning, Frank. (1991). Logic programming in the LF logical framework. *Pages 149–181 of: Huet, Gérard, & Plotkin, Gordon D. (eds), Logical frameworks*. Cambridge University Press.
- Pfenning, Frank. (1992). *A proof of the Church-Rosser theorem and its representation in a logical framework*. Tech. rept. CMU-CS-92-186. Department of Computer Science, Carnegie Mellon University.
- Pfenning, Frank. (1994). *A structural proof of cut elimination and its representation in a logical framework*. Tech. rept. CMU-CS-94-218. Department of Computer Science, Carnegie Mellon University.
- Pfenning, Frank. (1999). Logical Frameworks. Robinson, Alan, & Voronkov, Andrei (eds), *Handbook of Automated Reasoning*. Elsevier Science and MIT Press.
- Pfenning, Frank, & Rohwedder, Ekkehard. (1992). Implementing the meta-theory of deductive systems. *Pages 537–551 of: Kapur, D. (ed), International Conference on Automated Deduction*. Lecture Notes in Artificial Intelligence, vol. 607. Springer-Verlag.
- Pfenning, Frank, & Schürmann, Carsten. (1999). System description: Twelf - a meta-logical framework for deductive systems. *Pages 202–206 of: Ganzinger, Harald (ed), International Conference on Automated Deduction*.
- Pfenning, Frank, & Schürmann, Carsten. (2002). *Twelf user's guide, version 1.4*.
- Pientka, Brigitte, & Pfenning, Frank. (2000). Termination and reduction checking in the logical framework. Schürmann, Carsten (ed), *Workshop on Automation of Proofs by Mathematical Induction*.
- Reed, Jason. (2006). Hybridizing a logical framework. *International workshop on hybrid logic*. Electronic Notes in Theoretical Computer Science. Elsevier.
- Rohwedder, Ekkehard, & Pfenning, Frank. (1996). Mode and termination checking for higher-order logic programs. *Pages 296–310 of: Nielson, Hanne Riis (ed), European Symposium on Programming*. Lecture Notes in Computer Science, vol. 1058. Springer-Verlag.
- Salvesen, Anne. (1990). The Church-Rosser theorem for LF with $\beta\eta$ -reduction. *Unpublished notes to a talk given at the First Workshop on Logical Frameworks*.
- Schürmann, Carsten. (2000). *Automating the meta-theory of deductive systems*. Ph.D. thesis, Carnegie Mellon University.
- Schürmann, Carsten, & Pfenning, Frank. (1998). Automated theorem proving in a simple meta-logic for LF. *Pages 286–300 of: Kirchner, Claude, & Kirchner, Hélène (eds), International Conference on Automated Deduction*. Lecture Notes in Computer Science, vol. 1421. Springer-Verlag.
- Schürmann, Carsten, & Pfenning, Frank. (2003). A coverage checking algorithm for LF. *International Conference on Theorem Proving in Higher-Order Logics*. Springer-Verlag.
- Schürmann, Carsten, & Stehr, Mark-Oliver. (2005). An executable formalization of the HOL/NuPRL connection in Twelf. *International conference on logic for programming artificial intelligence and reasoning*. Springer-Verlag.
- Schürmann, Carsten, Yu, Dachuan, & Ni, Zhaozhong. (2001). A representation of F_ω in LF. *Electronic notes in Theoretical Computer Science*, **58**(1).

- Schürmann, Carsten, Poswolsky, Adam, & Sarnat, Jeffrey. (2005). The ∇ -calculus: Functional programming with higher-order encodings. *International Conference on Typed Lambda Calculi and Applications*. Springer-Verlag.
- Simmons, Robert. (2005). *Twelf as a unified framework for language formalization and implementation*. Tech. rept. Princeton University. Undergraduate Senior Thesis 18679.
- van Daalen, D. T. (1980). *The language theory of AUTOMATH*. Ph.D. thesis, Technical University of Eindhoven, Eindhoven, Netherlands.
- Virga, Roberto. (1999). *Higher-order rewriting with dependent types*. Ph.D. thesis, Carnegie Mellon University.
- Watkins, Kevin, Cervesato, Iliano, Pfenning, Frank, & Walker, David. (2002). *A concurrent logical framework I: Judgments and properties*. Tech. rept. CMU-CS-02-101. Department of Computer Science, Carnegie Mellon University. Revised May 2003.
- Watkins, Kevin, Cervesato, Iliano, Pfenning, Frank, , & Walker, David. (2004a). A concurrent logical framework: the propositional fragment. *Pages 355–377 of: Berardi, S., Coppo, M., & Damiani, F. (eds), Types for proofs and programs*. Lecture Notes in Computer Science, vol. 3085. Springer-Verlag.
- Watkins, Kevin, Cervesato, Iliano, Pfenning, Frank, & Walker, David. (2004b). Specifying properties of concurrent computations in CLF. Schürmann, Carsten (ed), *International workshop on logical frameworks and meta-languages*.