

Adelfa: A System for Reasoning about LF Specifications

Mary Southern

Gopalan Nadathur

University of Minnesota
Minneapolis Minnesota, USA

south163@umn.edu

ngopalan@umn.edu

We present a system called Adelfa that provides mechanized support for reasoning about specifications developed in the Edinburgh Logical Framework or LF. Underlying Adelfa is a new logic named \mathcal{L}_{LF} . Typing judgements in LF are represented by atomic formulas in \mathcal{L}_{LF} and quantification is permitted over contexts and terms that appear in such formulas. Contexts, which constitute type assignments to uniquely named variables that are modelled using the technical device of *nominal constants*, are characterized in \mathcal{L}_{LF} by *context schemas* that describe their inductive structure. We present these formulas and an associated semantics before sketching a proof system for constructing arguments that are sound with respect to the semantics. We then outline the realization of this proof system in Adelfa and illustrate its use through a few example proof developments. We conclude the paper by relating Adelfa to existing systems for reasoning about LF specifications.

1 Introduction

This paper describes a proof assistant called Adelfa that supports reasoning about specifications written in the Edinburgh Logical Framework, or LF [8]. Adelfa is based on a logic called \mathcal{L}_{LF} whose atomic formulas represent typing judgements in LF, and quantification is permitted over both contexts and terms which appear in these judgements. Term quantification is qualified by simple types that identify the functional structure of terms while forgetting dependencies. Context quantification is similarly qualified by schemas that capture the way contexts evolve during typing derivations in LF. The logic is complemented by a sequent-calculus based proof system that supports the construction of arguments of validity for relevant formulas. In addition to interpreting the logical symbols, the rules of the proof system embody an understanding of LF derivability; particular rules encode a case analysis style reasoning on LF judgements, the ability to reason inductively on the heights of such derivations, and an understanding of LF metatheorems. Adelfa is a tactics style theorem prover that implements this proof system: the development of a proof proceeds by invoking one of a set of sound tactic commands towards partially solving one of the proof obligations comprising the current proof state.

The rest of the paper is structured as follows. In Section 2 we briefly describe the underlying logic. Section 3 then outlines the associated proof system. In Section 4 we describe Adelfa and illustrate its use through a few examples. Section 5 concludes the paper by contrasting our work with other approaches that have been explored for reasoning about LF specifications.

2 A Logic for Articulating Properties of LF Specifications

We limit ourselves here to a brief overview of \mathcal{L}_{LF} , leaving its detailed presentation to other work [13, 19]. As already noted, the atomic formulas in \mathcal{L}_{LF} represent typing judgements in LF, a calculus with which we assume the reader to be already familiar. We begin by summarizing aspects of LF that are pertinent to discussions in this paper. We then present the formulas of \mathcal{L}_{LF} and illustrate their use in stating properties about LF derivability.

2.1 LF Basics and Hereditary Substitution

The variant of LF that is used in \mathcal{L}_{LF} is that presented in [9]. The only expressions permitted in this version, which is referred to as *canonical LF*, are those that are in β -normal form. Moreover, the typing rules ensure that all well-formed expressions are also in η -long form.

$$\begin{array}{ll}
 ty : Type & tm : Type \\
 unit : ty & app : tm \rightarrow tm \rightarrow tm \\
 arr : ty \rightarrow ty \rightarrow ty & abs : ty \rightarrow (tm \rightarrow tm) \rightarrow tm \\
 \\
 of : tm \rightarrow ty \rightarrow Type \\
 of_app : \Pi M_1:tm. \Pi M_2:tm. \Pi T_1:ty. \Pi T_2:ty. \\
 \quad of M_1 (arr T_1 T_2) \rightarrow of M_2 T_1 \rightarrow of (app M_1 M_2) T_2 \\
 of_abs : \Pi T_1:ty. \Pi T_2:ty. \Pi R:tm \rightarrow tm. \\
 \quad (\Pi x:tm. of x T_1 \rightarrow of (R x) T_2) \rightarrow of (abs T_1 (\lambda x. R x)) (arr T_1 T_2)
 \end{array}$$

Figure 1: An LF signature for typing in the simply typed λ -calculus

Object systems are encoded in LF by describing a signature which provides a means for representing the relevant constructs in the system as well as the relations between them. For example, if our focus is on typing judgements concerning the simply typed λ -calculus (STLC), then we might use the signature shown in Figure 1. This specification begins by identifying the LF types ty and tm and constructors for these types that serve to build LF representations of STLC types and terms; note that object language abstractions are represented using LF abstractions, following the idea of higher-order abstract syntax. Dependencies in types are then exploited to encode the typing relation between STLC terms and types. Specifically, the signature defines the type-level constant of towards this end: if T and Ty are LF terms of type tm and ty , respectively, then $(of T Ty)$ is intended to denote that the STLC expressions represented by T and Ty stand in this relation. The term-level constants of_app and of_abs that are identified as constructors for this type serve to encode the usual typing rules for applications and abstractions. In the use of such a specification, a central focus is on LF typing judgements of the form $\Gamma \vdash_{\Sigma} M \Leftarrow A$ where Σ represents a well-formed signature, A represents a well-formed type and Γ represents a well-formed context that assigns types to variables that may appear in M and A . Such a judgement is an assertion that M is typeable at the type A relative to the signature Σ and a context Γ that assigns particular LF types to the variables that appear in M and A . Such judgements are meant to translate into meaningful statements about the object systems encoded by the signature. For example, if Σ is the signature presented in Figure 1 and A is a type of the form $(of T Ty)$, then the judgement corresponds to the assertion that the STLC term represented by T is typeable at the type represented by Ty and that M encodes an STLC derivation for this judgement.

The typing rules in LF require the consideration of substitutions into LF expressions. In canonical LF, substitution must build in normalization, a requirement that is realized in [9] via the operation of *hereditary substitution*. We have generalized this operation to permit multiple simultaneous substitutions. Formally, a substitution is given by a finite set of triples of the form $\{\langle x_1, M_1, \alpha_1 \rangle, \dots, \langle x_n, M_n, \alpha_n \rangle\}$ in which, for $1 \leq i \leq n$, x_i is a variable, M_i is a term and α_i is a simple type constructed from the sole base type o . We refer to the simple types that index substitutions as *arity types*. The attempt to apply such a substitution to an LF expression always terminates. Moreover, it must yield a unique result whenever the expression being substituted into and the substituents satisfy the functional structure determined by

the arity types. We do not discuss this constraint in more detail here, noting only that the constraint is satisfied in all the uses we make of hereditary substitution in this paper. We write $E[[\theta]]$ to denote the result of applying the substitution θ to the expression E in this situation.

2.2 Formulas over LF Typing Judgements and their Meaning

The logic \mathcal{L}_{LF} is parameterized by an LF signature Σ . The basic building blocks for formulas in this context are typing judgements of the form $\Gamma \vdash_{\Sigma} M \Leftarrow A$ that are written as $\{G \vdash M : A\}$. However, the syntax of the expressions in the logic differs somewhat from that of the LF expressions. To begin with, we permit term variables that are bound by quantifiers to appear in types and terms. These variables have a different logical character from the variables that are bound in a context. More specifically, these variables may be instantiated by terms in the domains of the quantifiers, whereas the variables bound by declarations in an LF context represent fixed entities that are also distinct from all other similar entities within the typing judgement. To accurately capture the role of the variables bound in an LF context, we represent them using *nominal constants* [7, 21]; these are entities that are represented in this paper by the symbol n possibly with subscripts and that behave like constants except that they can be permuted with other such entities in an atomic formula without changing logical content. To support this treatment, we also allow nominal constants to appear in expressions corresponding to types and terms in the logic. Finally, we allow for contexts to be variables so as to permit quantification over them. More specifically the syntax for G , which constitutes a *context expression* in $\{G \vdash M : A\}$, is given by the following rule:

$$G ::= \Gamma \mid \cdot \mid G, n : A.$$

The symbol Γ here denotes the category of variables that range over contexts.

In typical reasoning scenarios, instantiations for context variables must be constrained in order to be able to articulate contentful properties. Such constraints are described via a special kind of typing for context variables. The formal realization of this idea, which is inspired by the notion of *regular worlds* in Twelf [16, 18], is based on declarations given by the following syntax rules.

$$\begin{aligned} \textbf{Block Declarations} \quad \Delta & ::= \cdot \mid \Delta, y : A \\ \textbf{Block Schema} \quad \mathcal{B} & ::= \{x_1 : \alpha_1, \dots, x_n : \alpha_n\} \Delta \\ \textbf{Context Schema} \quad \mathcal{C} & ::= \cdot \mid \mathcal{C}, \mathcal{B} \end{aligned}$$

According to this definition, a *context schema* is a collection of block schemas. The instances of a block schema $\{x_1 : \alpha_1, \dots, x_n : \alpha_n\}(y_1 : B_1, \dots, y_m : B_m)$ are sequences $n_1 : C_1, \dots, n_m : C_m$ obtained by choosing particular terms for the schematic variables x_1, \dots, x_n and particular nominal constants for the y_1, \dots, y_m . A context satisfies a context schema if the context comprises a sequence of instances for the block schemas defining it. Context schemas are the types that are associated with context variables and they play the obvious role in limiting their domains.

The formulas of \mathcal{L}_{LF} are given by the following syntax rule:

$$F ::= \{G \vdash M : A\} \mid \forall x : \alpha. F \mid \exists x : \alpha. F \mid \Pi \Gamma : \mathcal{C}. F \mid F_1 \wedge F_2 \mid F_1 \vee F_2 \mid F_1 \supset F_2 \mid \top \mid \perp$$

The symbol Π represents universal quantification pertaining to contexts; note that such quantification is qualified by a context schema. The symbol x represents a term variable, i.e. the logic permits universal and existential quantification over LF terms. The symbol α that annotates such variables represents an arity type. For a formula to be well-formed, the terms appearing in it must be well-typed and in canonical form relative to an arity typing that is determined as follows: for the quantified variables this is given by

their annotations; and for a constant in the signature or a nominal constant in the context expression that is assigned the LF type A it is the *erased form* of A , written as $(A)^-$, that is obtained by replacing the atomic types in A by o and otherwise retaining the functional structure. This well-typing requirement is intended to eliminate structural issues in the consideration of validity of a formula, thereby allowing the focus to be on the more contentful relational aspects that are manifest in dependencies in typing.

A closed formula, i.e. a formula in which no unbound context or term variable appears, that is of the form $\{G \vdash M : A\}$ is deemed to be true exactly when G is a well-formed LF context, A is a well formed LF type relative to G and the typing judgement $G \vdash_{\Sigma} M \Leftarrow A$ is derivable. This semantics is extended to closed formulas involving the logical connectives and constants by using their usual understanding. Finally, the quantifiers are accorded a substitution semantics. The formula $\Pi \Gamma : \mathcal{C}. F$ holds just in the case that $F[\{G/\Gamma\}]$ holds for every context G that satisfies the context schema \mathcal{C} , where $E[\sigma]$ denotes the result of a standard replacement substitution applied to the expression E , being careful, of course, to carry out any renaming of bound variables that is necessary to avoid inadvertent capture. The formula $\forall x : \alpha. F$ holds exactly when $F[\{\langle x, t, \alpha \rangle\}]$ holds for every closed term t that has the arity type α . Finally, the formula $\exists x : \alpha. F$ is true if there is some closed term t with arity type α such that $F[\{\langle x, t, \alpha \rangle\}]$ is true.

2.3 An Illustration of the Logic

The property of uniqueness of type assignment for the STLC can be expressed via the formula

$$\begin{aligned} \Pi \Gamma : c. \forall E : o. \forall T_1 : o. \forall T_2 : o. \forall D_1 : o. \forall D_2 : o. \{ \Gamma \vdash T_1 : ty \} \supset \{ \Gamma \vdash T_2 : ty \} \supset \\ \{ \Gamma \vdash D_1 : of E T_1 \} \supset \{ \Gamma \vdash D_2 : of E T_2 \} \supset \exists D_3 : o. \{ \cdot \vdash D_3 : eq T_1 T_2 \} \end{aligned}$$

where c represents the context schema $\{T : o\}(x : tm, y : of x T)$ and the signature parameterizing the logic is that in Figure 1 augmented with the declarations $eq : ty \rightarrow ty \rightarrow Type$ and $refl : \Pi T : ty. eq T T$. A special case of this formula is that when the context variable is instantiated with the empty context. However, the typing rule for abstractions will require us to consider non-empty contexts in the analysis. Note, though, that the signature ensures that the extended contexts that need to be considered will always satisfy the constraint expressed by c .

We can argue for the validity of the formula above in two steps. We show first the validity of

$$\Pi \Gamma : c. \forall T_1 : o. \forall T_2 : o. \forall D : o. \{ \Gamma \vdash D : eq T_1 T_2 \} \supset \{ \cdot \vdash D : eq T_1 T_2 \}.$$

This formula, which is essentially a *strengthening* property for eq , can be argued to be valid by observing that bindings in a well-formed LF context satisfying the context schema c have no role to play in constructing a term of the (well-formed) LF type $(eq T_1 T_2)$ for any terms T_1 and T_2 . We then combine this observation with the validity of the formula

$$\begin{aligned} \Pi \Gamma : c. \forall E : o. \forall T_1 : o. \forall T_2 : o. \forall D_1 : o. \forall D_2 : o. \{ \Gamma \vdash T_1 : ty \} \supset \{ \Gamma \vdash T_2 : ty \} \supset \\ \{ \Gamma \vdash D_1 : of E T_1 \} \supset \{ \Gamma \vdash D_2 : of E T_2 \} \supset \exists D_3 : o. \{ \Gamma \vdash D_3 : eq T_1 T_2 \}. \end{aligned}$$

To establish the validity of the last formula, we must show that, for a closed context expression G that instantiates the schema c and for closed expressions E, T_1, T_2, D_1 , and D_2 , if $\{G \vdash D_1 : of E T_1\}$ and $\{G \vdash D_2 : of E T_2\}$ are both valid, then T_1 and T_2 must be identical. The argument proceeds by induction on the height of the derivation of $G \vdash_{\Sigma} D_1 \Leftarrow of E T_1$ that the assumption implies must exist. An analysis of this derivation shows that there are three cases to consider, corresponding to whether the head symbol of D_1 is of_app , of_abs , or a nominal constant from G . For the last of these, the argument is that the validity of $\{G \vdash D_1 : of E T_1\}$ implies that G is a well-formed context and thus that the typing is unique. The other two cases invoke the induction hypothesis; in the case of of_abs we must consider a shorter derivation in which the context has been extended, but in a way that conforms to the definition of the context schema c .

3 A Proof System for the Logic

We now describe a sequent calculus that supports arguments of validity of the kind outlined in Section 2.3. We aim only to present the spirit of this calculus, leaving its detailed consideration again to other work [13, 19]. We note also that the main goal for the calculus is to provide a means for sound and effective reasoning rather than to be complete with respect to the semantics described for \mathcal{L}_{LF} . We begin by explaining the structure of sequents before proceeding to a discussion of the inference rules.

3.1 The Structure of Sequents

A sequent, written as $\mathbb{N}; \Psi; \Xi; \Omega \longrightarrow F$, is a judgement that relates a finite subset \mathbb{N} of nominal constants with associated arity types, a finite set Ψ of term variables also with associated arity types, a finite set Ξ of context variables with types of the kind described below, a finite set Ω of *assumption formulas* and a *conclusion* or *goal* formula F ; here, \mathbb{N} is the *support set* of the sequent, Ψ is its *eigenvariables context* and Ξ is its *context variables context*. The formulas in $\Omega \cup \{F\}$ must be formed out of the symbols in \mathbb{N} , Ψ , Ξ and the (implicit) signature Σ , and they must be well-formed with respect to these collections in the sense explained in Section 2.2. The members of Ξ have the form $\Gamma \uparrow \mathbb{N}_\Gamma : \mathcal{C}[G_1, \dots, G_n]$, where \mathbb{N}_Γ is a collection of nominal constants, \mathcal{C} is a context schema, and G_1, \dots, G_n is a listing of instances of block schemas from \mathcal{C} in which the types assigned to nominal constants are well-formed with respect to Σ , $\mathbb{N} \setminus \mathbb{N}_\Gamma$, and Ξ . This “typing” of the variable Γ is intended to limit its range to closed contexts obtained by interspersing instances of block schemas from \mathcal{C} in which nominal constants from \mathbb{N}_Γ do not appear between instances of G_1, \dots, G_n obtained by substituting terms formed from Σ and nominal constants not appearing in the support set of the sequent for the variables in Ψ ; the substitutions for the variables in Ψ must respect arity typing and the LF types in the resulting context must be well-typed in an arity sense.

The basic notion of meaning for sequents is one that pertains to closed sequents, i.e. ones of the form $\mathbb{N}; \emptyset; \emptyset; \Omega \longrightarrow F$. Such a sequent is *valid* if F is valid or one of the assumption formulas in Ω is not valid. A sequent of the general form $\mathbb{N}; \Psi; \Xi; \Omega \longrightarrow F$ is then considered valid if all of its instances obtained by substituting closed terms not containing the nominal constants in \mathbb{N} and respecting arity typing constraints for the variables in Ψ and replacing the variables in Ξ with closed contexts respecting their types in the manner described above are valid. The goal of showing that a formula F whose nominal constants are contained in the set \mathbb{N} is valid now reduces to showing the validity of the sequent $\mathbb{N}; \emptyset; \emptyset; \emptyset \longrightarrow F$.

3.2 The Rules for Deriving Sequents

The sequent calculus comprises two kinds of rules: those that pertain to the logical symbols and structural aspects of sequents and those that encode the interpretation of atomic formulas as assertions of derivability in LF. We discuss the rules under these categories below, focusing mainly on the latter kind of rules. For paucity of space, we do not present the rules explicitly but rather discuss their intuitive content. We also note that all the rules that we describe have been shown to be sound [13].

3.2.1 Structural and Logical Rules

The calculus includes the usual contraction and weakening rules pertaining to assumption formulas. Also included are rules for adding and removing entries from the support set and the eigenvariables and context variables contexts when these additions do not impact the overall well-formedness of sequents.

Finally, the cut rule, which facilitates the use of well-formed formulas as lemmas, is also present in the collection.

The most basic logical rule is that of an axiom. The main deviation from the usual form for this is the incorporation of the invariance of LF derivability under permutations of the names of context variables; this is realized in our proof system via a notion of equivalence of formulas under permutations of nominal constants. The rules for the connectives and quantifiers take the expected form. The only significant point to note is that eigenvariables that are introduced for (essential) universal quantifiers must be raised over the support set of the sequent to correctly reflect dependencies given our interpretation of sequents [11].

3.2.2 The Treatment of Atomic Formulas

The calculus builds in the understanding of formulas of the form $\{G \vdash M : A\}$ via LF derivability. If A is a type of the form $\Pi x:A_1.A_2$, then M must have the form $\lambda x.M'$ and the atomic formula can be replaced by one of the form $\{G, n : A_1 \vdash M' : A_2\}$ in the sequent; here, n must be a nominal constant not already in the support set and if G contains a context variable then its type annotation must be changed to prohibit the occurrence of n in its instantiations. If A is an atomic type and $\{G \vdash M : A\}$ is the goal formula, then the corresponding rule allows a step to be taken in the validation of the typing judgement; specifically, if M is the term $(h M_1 \dots M_n)$ where h is a constant or a nominal constant to which Σ or G assigns the LF type $\Pi x_1:A_1. \dots \Pi x_n:A_n.A'$ and A is identical to $A' \llbracket \{ \langle x_1, M_1(A_1)^- \rangle, \dots, \langle x_n, M_n, (A_n)^- \rangle \} \rrbracket$, then the rule leads to the consideration of the derivation of sequents in which the goal formula is changed to $\{G \vdash M_i \llbracket \{ \langle x_1, M_1(A_1)^- \rangle, \dots, \langle x_{i-1}, M_{i-1}, (A_{i-1})^- \rangle \} \rrbracket\}$ for $1 \leq i \leq n$. Note that if G begins with a context variable Γ , then the assignments in the blocks in the “type” of Γ are considered to be assignments in G .

The most contentful part of the treatment of atomic formulas is when the formula $\{G \vdash M : A\}$ in which A is an atomic type appears as an assumption formula in the sequent. The example in Section 2.3 demonstrates the *case analysis* style of reasoning that we would want to capture in the proof rule; we must identify all the possibilities for the valid closed instances of this formula and analyze the validity of the sequent based on these instances. The difficulty, however, is that there may be far too many closed instances to consider explicitly. This issue can be refined into two specific problems that must be addressed. First, the context G might begin with a context variable and we must then identify a realistic way to consider all the instantiations of that variable that yield an actual, closed context. Second, we must describe a manageable approach to considering all possible instantiations for the term variables that may appear in $\{G \vdash M : A\}$.

The first problem is solved in the enunciation of the rule through an *incremental elaboration* of a context variable that is driven by the atomic formula under scrutiny. Suppose that G begins with the context variable Γ corresponding to which there is the declaration $\Gamma \uparrow \mathbb{N}_\Gamma : \mathcal{C}[G_1, \dots, G_k]$ in the context variables context. We would at the outset need to consider all the constants in Σ and all the nominal constants identified explicitly in G , which includes the ones declared in G_1, \dots, G_k , as potential heads for M in the formula $\{G \vdash M : A\}$. Additionally, this head may come from a part of Γ that has not yet been made explicit. To account for this, the rule considers all the possible instances for the block declarations constituting \mathcal{C} and all possible locations for such blocks in the sequence G_1, \dots, G_k . We note that the number of such instances that have to be examined is finite because it suffices to consider exactly one representative for any nominal constant that does not appear in the support set of the sequent; this observation follows from the invariance of LF typing judgements under permutations of names for the variables bound in the context.

The second problem is addressed by first describing a notion of unification that will ensure that all closed instances will be considered and then identifying the idea of a *covering set of unifiers* that enables

us to avoid an exhaustive consideration. To elaborate a little on this approach, suppose that the (nominal) constant h with LF type $\prod x_1:A_1. \dots \prod x_n:A_n. A'$ has been identified as the candidate head for M . Further, for $1 \leq i \leq n$, let t_i be terms representing fresh variables raised over the support set of the sequent.¹ Then, based on the notion of unification described, for each substitution θ that unifies $(h \ t_1 \ \dots \ t_n)$ and M on the one hand and A and $A' \llbracket \{ \langle x_1, t_1(A_1)^- \rangle, \dots, \langle x_n, t_n, (A_n)^- \rangle \} \rrbracket$ on the other, it suffices to consider the derivability of the sequent that results from replacing $\{G \vdash M : A\}$ in the original sequent with the set of formulas

$$\{ \{ G \vdash t_i : A_i \llbracket \{ \langle x_1, t_1(A_1)^- \rangle, \dots, \langle x_{i-1}, t_{i-1}, (A_{i-1})^- \rangle \} \rrbracket \} \mid 1 \leq i \leq n \}$$

and then applying the substitution θ . However, this will still result in a large number of cases since the collection of unifiers must include all relevant closed instances for the analysis to be sound. The notion of a covering set provides a means for limiting attention to a small subset of unifiers while still preserving soundness.

3.2.3 Induction over the Heights of LF Derivations

The example in Subsection 2.3 also illustrates the role of induction over the heights of LF typing derivations in informal reasoning. This kind of induction is realized in our sequent calculus by using an annotation based scheme inspired by Abella [4, 6]. Specifically, we add to the syntax two additional forms of atomic formulas: $\{G \vdash M : A\}^{\textcircled{i}}$ and $\{G \vdash M : A\}^{*i}$. These represent, respectively, a formula that has an LF derivation of some given height and another formula of strictly smaller height; the latter formula is obtained typically by an unfolding step embodied in the use of a case analysis rule. The index i on the annotation symbol is used to identify distinct pairs of @ and * annotations. The induction proof rule then has the form

$$\frac{\mathbb{N}; \Psi; \Xi; \Omega, \mathcal{Q}_1.(F_1 \supset \dots \supset \mathcal{Q}_{k-1}.(F_{k-1} \supset \mathcal{Q}_k.(\{G \vdash M : A\}^{*i} \supset \dots \supset F_n))) \longrightarrow \mathcal{Q}_1.(F_1 \supset \dots \supset \mathcal{Q}_{k-1}.(F_{k-1} \supset \mathcal{Q}_k.(\{G \vdash M : A\}^{\textcircled{i}} \supset \dots \supset F_n)))}{\mathbb{N}; \Psi; \Xi; \Omega \longrightarrow \mathcal{Q}_1.(F_1 \supset \dots \supset \mathcal{Q}_{k-1}.(F_{k-1} \supset \mathcal{Q}_k.(\{G \vdash M : A\} \supset \dots \supset F_n)))} \textit{induction}$$

where \mathcal{Q}_i represent a sequence of context quantifiers or universal term quantifiers and the annotations \textcircled{i} and $*^i$ must not already appear in the conclusion sequent. The premise of this proof rule can be viewed as providing a proof schema for constructing an argument of validity for any particular height m , and so by an inductive argument we can conclude that the formula will be valid regardless of the derivation height. This idea is made precise in a proof of soundness for the rule [13, 19].

For this proof rule to be useful in reasoning, we will need a form of case analysis which permits us to move from a formula annotated with @ to one annotated by * when reduced. Such a mechanism is built into the proof system and is used in the examples in the next section.

3.2.4 Rules Encoding Metatheorems Concerning LF Derivability

Typing judgements in LF admit several metatheorems that are useful in reasoning about specifications: if $\Gamma \vdash_{\Sigma} M \Leftarrow A$ has a derivation then so does $\Gamma, x : A' \vdash_{\Sigma} M \Leftarrow A$ for a fresh variable x and any well-formed type A' (weakening); if $\Gamma, x : A' \vdash_{\Sigma} M \Leftarrow A$ has a derivation and x does not appear in M or A then so also does $\Gamma \vdash_{\Sigma} M \Leftarrow A$ (strengthening); if $\Gamma_1, x_1 : A_1, x_2 : A_2, \Gamma_2 \vdash_{\Sigma} M \Leftarrow A$ has a derivation and

¹If n_1, \dots, n_{ℓ} is a listing of \mathbb{N} then, for $1 \leq i \leq n$, t_i is the term $(z_i \ n_1 \ \dots \ n_{\ell})$ where z_i is a fresh variable of suitable arity type.

x_1 does not appear in A_2 then $\Gamma_1, x_2 : A_1, x_1 : A_1, \Gamma_2 \vdash_{\Sigma} M \Leftarrow A$ must have a derivation (permutation); and if $\Gamma_1 \vdash_{\Sigma} M' \Leftarrow A'$ and $\Gamma_1, x : A', \Gamma_2 \vdash_{\Sigma} M \Leftarrow A$ have derivations then there must be a derivation for $\Gamma_1, \Gamma_2 \llbracket \{ \langle x, M', (A')^- \rangle \} \rrbracket \vdash_{\Sigma} M \llbracket \{ \langle x, M', (A')^- \rangle \} \rrbracket \Leftarrow A \llbracket \{ \langle x, M', (A')^- \rangle \} \rrbracket$ (substitution). Moreover, in the first three cases, the derivations are structurally similar, e.g. they have the same heights. These metatheorems are built into the sequent calculus via (sound) axioms. For example, (one version of) the strengthening metatheorem is encoded in the axiom

$$\frac{n \text{ does not appear in } M, A, \text{ or the explicit bindings in } G}{\mathbb{N}; \Psi; \Xi; \Omega \longrightarrow \{G, n : B \vdash M : A\} \supset \{G \vdash M : A\}} \text{LF-str}$$

These axioms can be combined with the cut rule to encode the informal reasoning process.

4 The Adelfa System and its Use in Reasoning

In this section we expose the Adelfa system that provides support for mechanizing arguments of validity for formulas in \mathcal{L}_{LF} using the proof system outlined in Section 3. The first subsection below provides an overview of Adelfa. We then consider a few example reasoning tasks to indicate how the system is intended to be used and also to give a sense for its capabilities.

4.1 Overview of Adelfa

The structure of Adelfa is inspired by the proof assistant Abella [4, 6]. Each development in Adelfa is parameterized by an LF signature which is identified by an initial declaration. The interaction then proceeds to a mode in which context schemas can be identified and theorems can be posited. When a theorem has been presented, the interaction enters a *proof mode* in which the user directs the construction of a proof for a suitable sequent using a repertoire of tactics. These tactics encode the (sound) application of a combination of rules from the sequent calculus to produce a new proof state represented by a collection of sequents for which proofs need to be constructed. We note in this context the presence of a *search* tactic that attempts to construct a proof for a sequent with an atomic goal formula through the repeated use of LF typing rules. If a tactic yields more than one proof obligation then these are ordered in a predetermined way and all existing obligations are maintained in a stack, to be eventually treated in order for the proof to be completed.

The syntax of LF terms and types in Adelfa follows that of fully explicit Twelf. The type $\Pi x:A. B$ is represented by $\{x : A\} B$ and the abstraction term $\lambda x. M$ is represented by $[x] M$. A context schema definition is introduced by the keyword *Schema* and it associates an identifier with the schema that is to be used in its place in the subsequent development. The block schemas defining a context schema are separated by $;$, the schematic variable declarations of each block schema are presented within braces, and the block itself is surrounded by parenthesis. A theorem is introduced by the keyword *Theorem*, and the declaration associates an identifier with the formula to be proved. Once a theorem has been successfully proved this identifier is available for use as a lemma (via the cut rule) in later proofs.

The following table identifies the formula syntax of Adelfa.

Formula	Adelfa Syntax	Formula	Adelfa Syntax
$\forall x : \alpha. F$	forall $x : \alpha. F$	$F_1 \vee F_2$	$F_1 \vee F_2$
$\exists x : \alpha. F$	exists $x : \alpha. F$	$F_1 \supset F_2$	$F_1 \Rightarrow F_2$
$\Pi G : \mathcal{C}. F$	ctx $G : \mathcal{C}. F$	\top	true
$F_1 \wedge F_2$	$F_1 \wedge F_2$	\perp	false

A key issue in the Adelfa implementation is the realization of case analysis. In Section 3, we have discussed the basis for this rule in a special form of unification and also the role of the idea of covering sets of unifiers in its practical realization. We have shown that, in the situations where it is applicable, the notion of higher-order pattern unification [10] can be adapted to provide such covering sets of solutions [19]. Adelfa employs this approach, using the higher-order pattern unification algorithm described by Nadathur and Linnell [12] in its implementation. The benefit of following this course is that the covering set of solutions comprises a single substitution.

4.2 Example Developments in Adelfa

We begin by discussing the formalization of the proof of type uniqueness for the STLC in Adelfa, thereby demonstrating the use of both induction and case analysis in reasoning. We then show the usefulness of LF metatheorems by considering the proof of cut admissibility for a simple sequent calculus. Finally, we demonstrate the flexibility of the logic through a simple example which is of a form that cannot be directly represented as a function type; we will see in the next section that the ability to reason directly about such formulas distinguishes Adelfa from the Twelf family of systems.

4.2.1 Type Uniqueness for the STLC

The informal argument in Section 2.3 had identified a context schema that is relevant to this example. This schema would be presented to Adelfa via the following declaration:

```
Schema c := {T:o}(x:tm,y:of x T).
```

The proof of the actual type uniqueness theorem had relied on a strengthening lemma concerning the equality relation between encodings of types in the STLC. We elide the proof of this lemma and its use in the formalization, focusing instead on the proof of the formula

$$\Pi \Gamma : c. \forall E : o. \forall T_1 : o. \forall T_2 : o. \forall D_1 : o. \forall D_2 : o. \{ \Gamma \vdash T_1 : ty \} \supset \{ \Gamma \vdash T_2 : ty \} \supset \\ \{ \Gamma \vdash D_1 : of E T_1 \} \supset \{ \Gamma \vdash D_2 : of E T_2 \} \supset \exists D_3 : o. \{ \Gamma \vdash D_3 : eq T_1 T_2 \};$$

the full development of this example and the others in this paper are available with the Adelfa source from the Adelfa web page [1].

The proof development process starts with the presentation of the above formula as a theorem. This will result in Adelfa displaying the following proof state to the user:

Vars:

Nominals:

Contexts:

=====

```
ctx G:c, forall E:o T1:o T2:o D1:o D2:o,
  {G |- T1 : ty} => {G |- T2 : ty} => {G |- D1 : of E T1} =>
  {G |- D2 : of E T2} => exists d3:o, {G |- D3 : eq T1 T2}
```

This proof state corresponds transparently to the sequent whose proof will establish the validity of the formula; the components of the sequent to the left of the arrow appear above the double line and the goal formula appears below the double line. To make it easy to reference particular formulas during reasoning, a name is associated with each assumption formula in the sequent. In the general case, where there can be more than one proof obligation, the remaining obligations will be indicated by a listing of just their goal formulas below the obligation currently in focus.

The informal argument used an induction on the height of the LF derivation represented by the third atomic formula that appears in the goal formula shown. To initiate this process in Adelfa, we would invoke an induction tactic command, indicating the atomic formula that is to be the focus of the induction. This causes the goal formula with the third atomic formula annotated with a `*` to be added to the assumption formulas and the goal formula to be changed to one in which the third formula has the annotation `@`. At this stage, we may invoke a tactic command to apply a sequence of right introduction rules, followed by another command to invoke case analysis on the assumption formula $\{G \mid - D1 : \text{of } E T1\}@$. Doing so results in the following proof state:

```

Vars: a1:o -> o -> o, R:o -> o, T3:o, T4:o, D2:o, T2:o
Nominals: n2:o, n1:o, n:o
Contexts: G{n2, n1, n}:c[]
IH: ctx G:c, forall E:o T1:o T2:o D1:o D2:o,
    {G |- T1 : ty} => {G |- T2 : ty} => {G |- D1 : of E T1}* =>
    {G |- D2 : of E T2} => exists D3:o, {G |- D3 : eq T1 T2}
H1:{G |- arr T3 T4 : ty}
H2:{G |- T2 : ty}
H4:{G |- D2 : of (abs T3 ([x]R x)) T2}
H5:{G |- T3 : ty}*
H6:{G |- T4 : ty}*
H7:{G, n:tm |- R n : tm}*
H8:{G, n1:tm, n2:of n1 T3 |- a1 n1 n2 : of (R n1) T4}*

=====
exists D3, {G |- D3 : eq (arr T3 T4) T2}

```

Subgoal 2: exists D3:o, {G |- D3 : eq T1 T2}

Subgoal 3: exists D3:o, {G |- D3 : eq (T1 n n1) (T2 n n1)}

As is to be anticipated, case analysis results in the consideration of three different possibilities: that where the head of $D1$ is, respectively, the constant *of_{abs}*, the constant *of_{app}*, or a nominal constant from the context G . The first of these cases is shown in elaborated form, the other two become additional proof obligations. The analysis based on the first case replaces the original assumption formula with new ones according to the type associated with *of_{abs}*. The annotation on these formulas is changed from `@` to `*`, to represent the fact that the derivations of the LF judgements associated with them must be of smaller height. Finally, those of these formulas that represent typing judgements for (LF) abstractions are analyzed further, resulting in the introduction of the new nominal constants n , $n1$, and $n2$ into the corresponding assumption formulas, and an annotation of the context variable G that indicates that only those instantiations in which these nominal constants do not appear must be considered for it.

At this point, a case analysis based on the formula identified by $H4$ will identify only a single case, as the term structure $(\text{abs } T3 ([x]R x))$ uniquely identifies a single structure for $D2$, that where the head is *of_{abs}*. We may now use the weakening metatheorem for LF to change the context for the assumption formula identified by $H6$ and the ones resulting from the case analysis of the formula identified by $H4$ to $(G, n1:tm, n2:of n1 T3)$. Noting that this context must satisfy the context schema ctx if G satisfies it, we have the ingredients in place to utilize the induction hypothesis, i.e. the assumption formula identified by IH , to be able to add the formula

```
{G, n1:tm, n2:of n1 T3 |- D1 n5 n4 n3 n2 n1 n : eq T4 T5}
```

to the collection of assumption formulas; note that the conclusion formula for the sequent would have also become

```
exists D3, {G |- D3 : eq (arr T3 T4) (arr T3 T5)}
```

as a result of the case analysis. Analyzing this assumption formula will produce a single case in which $D1$ is bound to $(\text{refl } T5)$ and $T4$ is replaced by $T5$ throughout the sequent. This branch of the proof can then be completed by instantiating the existential quantifier in the goal formula with the term $(\text{refl } (\text{arr } T3 T5))$ and using the assumption formulas that indicate that $T3$ and $T5$ must have the type ty to conclude that $(\text{arr } T3 T5)$ must also be similarly typed.

The second subgoal, the one corresponding to the application case, will now be presented in elaborated form. We will elide a discussion of this case because it does not illustrate any capabilities beyond those already exhibited in the consideration of the abstraction case. Instead we will jump to the final subgoal which is manifest in the following proof state:

```
Vars: D2:o -> o -> o, T2:o -> o -> o, T1:o -> o -> o
Nominals: n1:o, n:o
Contexts: G{:c}:[(n:tm, n1:of n (T1 n n1))]
IH: ctx G:c, forall E:o T1:o T2:o D1:o D2:o,
    {G |- T1 : ty} => {G |- T2 : ty} => {G |- D1 : of E T1}* =>
    {G |- D2 : of E T2} => exists D3:o, {G |- D3 : eq T1 T2}
H1:{G |- T1 n n1 : ty}
H2:{G |- T2 n n1 : ty}
H4:{G |- D2 n n1 : of n (T2 n n1)}
```

```
=====
exists D3, {G |- D3 : eq (T1 n n1) (T2 n n1)}
```

As we can see in this state, a new block has been elaborated in the context variable type of G . The nominal constants that are introduced by this elaboration may be used in the terms that instantiate the eigenvariables $D2$, $T2$, and $T1$ in the original sequent, a fact that is realized here by raising these variables over the nominal constants. The key to the proof for this case is that case analysis based on the assumption formula $H4$ will only identify a single relevant case: that when $D2$ is $n1$ and $T2$ is identified with $T1$. This is because n represents a unique nominal constant that cannot be matched with any other term or nominal constant, and a name can have at most one binding in G .

4.2.2 Cut Admissibility for the Intuitionistic Sequent Calculus

The example we consider now is that of proving the admissibility of cut for an intuitionistic sequent calculus. In this proof, there is a need to consider the weakening of the antecedents of sequents. We propose an encoding of sequents below under which the weakening rule applied to sequents can be modelled by a weakening applied to LF typing judgements. Thus, this example illustrates the use of LF metatheorems encoded in Adelfa in directly realizing reasoning steps in informal proofs.

We will actually consider only a fragment of the intuitionistic sequent calculus for propositional logic in this illustration. In particular, this fragment includes only the logical constant \top and the connectives for conjunction and implication. An LF signature that encodes this fragment appears below.

<i>proptm</i> : Type	<i>top</i> : <i>proptm</i>
<i>hyp</i> : <i>proptm</i> \rightarrow Type	<i>imp</i> : <i>proptm</i> \rightarrow <i>proptm</i> \rightarrow <i>proptm</i>
<i>conc</i> : <i>proptm</i> \rightarrow Type	<i>and</i> : <i>proptm</i> \rightarrow <i>proptm</i> \rightarrow <i>proptm</i>

$$\begin{aligned}
& \text{init} : \Pi A : \text{proptm}. \Pi D : \text{hyp } A. \text{conc } A \\
& \text{topR} : \text{conc top} \\
& \text{andL} : \Pi A : \text{proptm}. \Pi B : \text{proptm}. \Pi C : \text{proptm}. \Pi D_1 : (\Pi x : \text{hyp } A. \Pi y : \text{hyp } B. \text{conc } C). \\
& \quad \Pi D_2 : \text{hyp } (\text{and } A \ B). \text{conc } C \\
& \text{andR} : \Pi A : \text{proptm}. \Pi B : \text{proptm}. \Pi D_1 : \text{conc } A. \Pi D_2 : \text{conc } B. \text{conc } (\text{and } A \ B) \\
& \text{impL} : \Pi A : \text{proptm}. \Pi B : \text{proptm}. \Pi C : \text{proptm}. \Pi D_1 : \text{conc } A. \Pi D_2 : (\Pi x : \text{hyp } B. \text{conc } C). \\
& \quad \Pi D_3 : \text{hyp } (\text{imp } A \ B). \text{conc } C \\
& \text{impR} : \Pi A : \text{proptm}. \Pi B : \text{proptm}. \Pi D : (\Pi x : \text{hyp } A. \text{conc } B). \text{conc } (\text{imp } A \ B)
\end{aligned}$$

The propositions of this calculus are encoded as terms of type *proptm*. The two LF type constructors *hyp* and *conc* that take terms of type *proptm* as arguments are used to identify propositions that are hypotheses and conclusions of a sequent respectively. More specifically, a sequent of the form $A_1, \dots, A_n \longrightarrow B$ is represented under this encoding by the LF typing judgement

$$x_1 : \text{hyp } \overline{A_1}, \dots, x_n : \text{hyp } \overline{A_n} \vdash_{\Sigma} c \Leftarrow \text{conc } \overline{B},$$

where \overline{A} denotes the encoding of the proposition A .

The contexts relevant to this example can be characterized by the following schema declaration:

Schema $\text{cutctx} := \{A : o\}(x : \text{hyp } A)$

Relative to the given signature and schema declaration, the admissibility of cut is captured by the following formula:

$$\begin{aligned}
& \Pi G : \text{cutctx}. \forall a : o. \forall b : o. \forall d_1 : o. \forall d_2 : o \rightarrow o. \\
& \quad \{\cdot \vdash a : \text{proptm}\} \supset \{G \vdash d_1 : \text{conc } a\} \supset \{G \vdash \lambda x. d_2 \ x : \Pi x : \text{hyp } a. \text{conc } b\} \supset \exists d : o. \{G \vdash d : \text{conc } b\}
\end{aligned}$$

Note that in the third antecedent the term variable d_2 is permitted to depend on the cut formula a through the explicit dependency on the bound variable x , reflecting the fact that this proof can use the “cut formula.”

The informal proof that directs our development uses a nested induction. The first induction is on the structure of the cut formula, expressed by the formula $\{\cdot \vdash a : \text{proptm}\}$. The second induction is on the height of the derivation that uses the cut formula, expressed by $\{G \vdash \lambda x. d_2 \ x : \Pi x : \text{hyp } a. \text{conc } b\}$. The full proof is based on considering the different possibilities for the last rule in the proof with the cut formula included in the assumption set. In this illustration, we consider only the case where this is the rule that introduces an implication on the right side of the sequent. In the Adelfa proof, the indicated analysis is reflected in the case analysis of the formula $\{G, n : \text{hyp } a \vdash d_2 \ n : \text{conc } b\}$, which is derived from $\{G \vdash \lambda x. d_2 \ x : \Pi x : \text{hyp } a. \text{conc } b\}$, and the case to be considered is that where the head of the term $(d_2 \ n)$ is chosen to be *impR*.

One of the characteristics of LF typing judgements is that the assignments in typing contexts are ordered so as to reflect possible dependencies. However, in the object system under consideration, the ordering of propositions on the left of the sequent arrow is unimportant. This is evident in our encoding by the fact that there are no means to represent dependencies in the kind of contexts in consideration. Thus, we may use the permutation metatheorem to realize reordering in the context in a proof development that requires this. We will need to employ this idea in order to ensure that the context structure has the form $(G, n : \text{hyp } a)$ as is needed to invoke the induction hypothesis.

To get to the details of the case under consideration, we note that the cut formula is preserved through the concluding rule in the derivation. The informal argument then uses the induction hypothesis to obtain proofs of the premises of the rule but where the cut formula is left out of the left side of the sequents.

These proofs are then combined to get a proof for the concluding sequent, again with the cut formula left out from the premises, using the same rule to introduce an implication on the right. It is this argument that we want to mimic in the Adelfa development.

The proof state at the start of the case that is the focus of our discussion is depicted below:

```

Vars: d:o -> o -> o, b1:o, b2:o, d1:o, a:o
Nominals: n1:o, n:o
Contexts: G:cutctx[]
IH:ctx G:cutctx. forall a:o, forall b:o, forall d1:o, forall d2:o -> o,
  {a : proptm}* => {G |- d1 : conc a} =>
    {G, n:hyp a |- d2 n : conc b} => exists d:o, {G |- d : conc b}
IH1:ctx G:cutctx. forall a:o, forall b:o, forall d:o1, forall d2:o -> o,
  {a : proptm}@ => {G |- d1 : conc a} =>
    {G, n:hyp a |- d2 n : conc b}** => exists d:o, {G |- d : conc b}
H1:{a : proptm}@
H2:{G |- d1 : conc a}
H3:{G, n:hyp a |- impR b1 b2 ([x] d n x) : conc (imp b1 b2)}@@
H4:{G, n:hyp a |- b1 : proptm}**
H5:{G, n:hyp a |- b2 : proptm}**
H6:{G, n:hyp a, n1:hyp b1 |- d n n1 : conc b2}**
=====
exists d, {G |- d : conc (imp b1 b2)}

```

For simplicity, we have elided the other subgoals that remain in this display. Our objective in this case is to add to the assumptions the formula $\{G, n1:hyp b1 \mid D : conc b2\}$ for some term D ; from this, we can easily construct a term that we can use to instantiate the existential quantifier in the goal formula to conclude the proof. We would like to use the inductive hypothesis identified by IH1 towards this end. To be able to do this, we would need to extend the context expression in the assumption formula identified by H2 with a new binding of type $(hyp b1)$. We can do this using a tactic command that encodes the weakening metatheorem for LF. In actually carrying out this step, we would additionally use the strengthening metatheorem for LF to conclude that the type $(hyp b)$ is well-formed in the relevant context because of the assumption identified by H4 and the fact that $b1$ cannot depend on n . We must also rearrange the type assignments in the context in the assumption formula H6 so that the assignment corresponding to the cut formula appears at the end. This can be realized using the tactic command that encodes the permutation metatheorem in LF, a use of which will be valid for the reasons identified in the earlier discussion.

The end result of applying these reasoning steps is a state of the following form:

```

Vars: d:o -> o -> o, b1:o, b2:o, d1:o, a:o
Nominals: n1:o, n:o
Contexts: G:cutctx[]
IH:ctx G:cutctx. forall a:o, forall b:o, forall d1:o, forall d2:o -> o,
  {a : proptm}* => {G |- d1 : conc a} =>
    {G, n:hyp a |- d2 n : conc b} => exists d:o, {G |- d : conc b}
IH1:ctx G:cutctx. forall a:o, forall b:o, forall d1:o, forall d2:o -> o,
  {a : proptm}@ => {G |- d1 : conc a} =>
    {G, n:hyp a |- d2 n : conc b}** => exists d:o, {G |- d : conc b}

```

```

H1:{a : proptm}@
H2:{G |- d1 : conc a}
H3:{G, n:hyp a |- impR b1 b2 ([x] d n x) : conc (imp b1 b2)}@
H4:{G, n:hyp a |- b1 : proptm}**
H5:{G, n:hyp a |- b2 : proptm}**
H6:{G, n:hyp a, n1:hyp b1 |- d n n1 : conc b2}**
H7:{G |- b1 : proptm}**
H8:{G, n1:hyp b1 |- d1 : conc a}
H9:{G, n1:hyp b1, n:hyp a |- d n n1 : conc b2}**

```

```

=====
exists d, {G |- d : conc (imp b1 b2)}

```

We can now use the induction hypothesis IH1 to the assumption formulas H1, H8, and H9 to introduce a new term variable d' and the assumption formula $\{G, n1:hyp b1 \mid - d' n1 : conc b2\}$, from which we can complete the proof as previously indicated.

4.2.3 A Disjunctive Property for Natural Numbers

The formulas that we have considered up to this point have had a “function” structure: for every term satisfying certain typing constraints they have posited the existence of a term satisfying another typing constraint. Towards exhibiting the flexibility of the logic underlying Adelfa, we now consider an example of a theorem that *does not* adhere to this structure. This example is based on the following signature that encodes natural numbers and the even and odd properties pertaining to these numbers.

$nat : Type$	$even : nat \rightarrow Type$	$odd : nat \rightarrow Type$
$z : nat$	$e-z : even z$	$e-o : \prod N : nat. even N \rightarrow odd (s N)$
$s : nat \rightarrow nat$	$o-e : \prod N : nat. odd N \rightarrow even (s N)$	

Given this signature, we can express the property that every natural number is either even or odd:

$$\forall N : o. \{ \vdash N : nat \} \supset (\exists D : o. \{ \vdash D : even N \}) \vee (\exists D : o. \{ \vdash D : odd N \}).$$

Following the obvious informal proof, this formula can be proved by induction on the height of the derivation of the LF typing judgement $\cdot \vdash_{\Sigma} N \Leftarrow nat$. The application of the induction hypothesis yields two possible cases and in each case we pick a disjunct in the conclusion that we can show to be true.

5 Conclusion

This paper has described the Adelfa system for reasoning about specifications written in LF. Adelfa is based on the logic \mathcal{L}_{LF} in which the atomic formulas represent typing judgements in LF and where quantification is permitted over terms and over contexts that are characterized by context schemas. We have sketched a proof system for this logic and have demonstrated the effectiveness of Adelfa, which realizes this proof system, through a few examples. We have developed formalizations beyond the ones discussed here, such as type preservation for the STLC and subtyping for $F_{<}$; the latter is a problem from the PoPLMark collection [3].

While this work is not unique in its objectives, it differs from other developments in how it attempts to achieve them. One prominent approach in this context is that adopted by what we refer to as the “Twelf family.” The first exemplar of this approach is the Twelf system [15] that allows properties of interest to be characterized by types and the validity of such properties to be demonstrated by exhibiting

the totality as functions of inhabitants of these types. This approach has achieved much success but it is also limited by the fact that properties that can be reasoned about must be encodable by a function type, i.e. they must be formulas of the $\forall\dots\exists\dots$ form. By contrast, the logic underlying Adelfa has a much more flexible structure, allowing us, for example, to represent properties with a disjunctive structure as we have seen in Section 4.2.3; the reader might want to consult the encoding of this property provided in the wiki page of the Twelf project that covers output factoring to understand more fully the content of this observation [2]. Another virtue of our approach is that we are able to exhibit an explicit proof for properties at the end of a development. This drawback is mitigated to an extent in the Twelf context by the presence of the logic M_2^+ [18] that provides a formal counterpart to the approach to reasoning embodied in the Twelf system [22].

The Twelf approach differs in another way from the approach described here: it does not provide an explicit means for quantifying over contexts. It is possible to parameterize a development by a context description, but it is one fixed context that then permeates the development. As an example, it is not possible to express the strengthening lemma pertaining to equality of types in the STLC that we discussed in Section 2.3. The Beluga system [17], which otherwise shares the characteristics of the Twelf system, alleviates this problem by using a richer version of type theory that allows for an explicit treatment of contexts as its basis [14].

An alternative approach, that is more in alignment with what we have described here, is one that uses a translation of LF specifications into predicate logic to then be reasoned about using the Abella system [20]. This approach has several auxiliary benefits deriving from the expressiveness of the logic underlying Abella [7]; for example, it is possible to define relations between contexts, to treat binding notions explicitly in the reasoning process through the ∇ -quantifier, and to use inductive (and co-inductive) definitions in the reasoning logic. In future work, we would like to examine how to derive some of these benefits in the context of Adelfa as well. There are, however, also some drawbacks to the translation-based approach. One of the problems is that it is based on a “theorem” about the translation [5] that we know to be false for the version of LF that it pertains to.² Another issue is that proof steps that can be taken in Abella are governed by the logic underlying that system and these allow for many more possibilities than are sensible in the LF context. Moreover, it is not immediately clear how the reasoning steps that are natural with LF specifications translate into macro Abella steps relative to the translation. In this respect, one interesting outcome of the work on the logic underlying the Adelfa system might be an understanding of how one might build within Abella (or other related proof assistants) a targeted capability for reasoning about LF specifications. This also raises another interesting possibility that is worthy of attention, that of providing an alternative justification for the proof system for \mathcal{L}_{LF} based on the translation.

Acknowledgements

This paper is based upon work supported by the National Science Foundation under Grant No. CCF-1617771. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

²The theorem in question states that a (closed) LF typing judgement is derivable if and only if the translated form of the judgement is derivable in the relevant predicate logic. This is true in the forward direction. However, in the version of LF that the theorem is about, the translation loses typing information from terms and, hence, the inverse translation is ambiguous. Thus, it is only a weaker conclusion that can be drawn, that there is *some* LF judgement that is derivable if the formula in predicate logic is derivable; there is no guarantee that this is the same LF judgement that we started out with.

References

- [1] *The Adelfa System*. <http://adelfa.cs.umn.edu/>.
- [2] *The Twelf Project*. <http://twelf.org/>.
- [3] Brian E. Aydemir, Aaron Bohannon, Matthew Fairbairn, J. Nathan Foster, Benjamin C. Pierce, Peter Sewell, Dimitrios Vytiniotis, Geoffrey Washburn, Stephanie Weirich & Steve Zdancewic (2005): *Mechanized Metatheory for the Masses: The POPLmark Challenge*. In: *Theorem Proving in Higher Order Logics: 18th International Conference, LNCS 3603*, Springer, pp. 50–65, doi:10.1007/11541868_4.
- [4] David Baelde, Kaustuv Chaudhuri, Andrew Gacek, Dale Miller, Gopalan Nadathur, Alwen Tiu & Yuting Wang (2014): *Abella: A System for Reasoning about Relational Specifications*. *Journal of Formalized Reasoning* 7(2), doi:10.6092/issn.1972-5787/4650. Available at <http://jfr.unibo.it/article/download/4650/4137>.
- [5] Amy Felty & Dale Miller (1990): *Encoding a Dependent-Type λ -Calculus in a Logic Programming Language*. In Mark Stickel, editor: *Proceedings of the 1990 Conference on Automated Deduction, LNAI 449*, Springer, pp. 221–235, doi:10.1007/3-540-52885-7_90.
- [6] Andrew Gacek (2009): *A Framework for Specifying, Prototyping, and Reasoning about Computational Systems*. Ph.D. thesis, University of Minnesota.
- [7] Andrew Gacek, Dale Miller & Gopalan Nadathur (2011): *Nominal Abstraction*. *Information and Computation* 209(1), pp. 48–73, doi:10.1016/j.ic.2010.09.004.
- [8] Robert Harper, Furio Honsell & Gordon Plotkin (1993): *A Framework for Defining Logics*. *Journal of the ACM* 40(1), pp. 143–184, doi:10.1145/138027.138060.
- [9] Robert Harper & Daniel R. Licata (2007): *Mechanizing Metatheory in a Logical Framework*. *Journal of Functional Programming* 17(4–5), pp. 613–673, doi:10.1017/S0956796807006430.
- [10] Dale Miller (1991): *A Logic Programming Language with Lambda-Abstraction, Function Variables, and Simple Unification*. *J. of Logic and Computation* 1(4), pp. 497–536, doi:10.1093/logcom/1.4.497.
- [11] Dale Miller (1992): *Unification Under a Mixed Prefix*. *Journal of Symbolic Computation* 14(4), pp. 321–358, doi:10.1016/0747-7171(92)90011-R.
- [12] Gopalan Nadathur & Natalie Linnell (2005): *Practical Higher-Order Pattern Unification With On-the-Fly Raising*. In: *ICLP 2005: 21st International Logic Programming Conference, LNCS 3668*, Springer, Sitges, Spain, pp. 371–386, doi:10.1007/11562931_28.
- [13] Gopalan Nadathur & Mary Southern (2021): *A Logic for Reasoning About LF Specifications*. Available from <http://arxiv.org/abs/2107.00111>.
- [14] Aleksandar Nanevski, Frank Pfenning & Brigitte Pientka (2008): *Contextual Model Type Theory*. *ACM Trans. on Computational Logic* 9(3), pp. 1–49, doi:10.1145/1352582.1352591.
- [15] Frank Pfenning & Carsten Schürmann (1999): *System Description: Twelf — A Meta-Logical Framework for Deductive Systems*. In H. Ganzinger, editor: *16th Conf. on Automated Deduction (CADE), LNAI 1632*, Springer, Trento, pp. 202–206, doi:10.1007/3-540-48660-7_14.
- [16] Frank Pfenning & Carsten Schürmann (2002): *Twelf User’s Guide*. Available from <http://www.cs.cmu.edu/~twelf/guide-1-4>.
- [17] Brigitte Pientka & Joshua Dunfield (2010): *Beluga: A Framework for Programming and Reasoning with Deductive Systems (System Description)*. In J. Giesl & R. Hähnle, editors: *Fifth International Joint Conference on Automated Reasoning, LNCS 6173*, pp. 15–21, doi:10.1007/978-3-642-14203-1_2.
- [18] Carsten Schürmann (2000): *Automating the Meta Theory of Deductive Systems*. Ph.D. thesis, Carnegie Mellon University. Available at <http://www.cs.yale.edu/homes/carsten/papers/S00b.ps.gz>.
- [19] Mary Southern (2021): *A Framework for Reasoning About LF Specifications*. Ph.D. thesis, University of Minnesota.

- [20] Mary Southern & Kaustuv Chaudhuri (2014): *A Two-Level Logic Approach to Reasoning About Typed Specification Languages*. In Venkatesh Raman & S. P. Suresh, editors: *34th International Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS 2014)*, *Leibniz International Proceedings in Informatics (LIPIcs)* 29, Dagstuhl, Germany, pp. 557–569, doi:10.4230/LIPIcs.FSTTCS.2014.557. Available at <http://drops.dagstuhl.de/opus/volltexte/2014/4871>.
- [21] Alwen Tiu (2006): *A Logic for Reasoning about Generic Judgments*. In A. Momigliano & B. Pientka, editors: *Int. Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP'06)*, *ENTCS* 173, pp. 3–18, doi:10.1016/j.entcs.2007.01.016.
- [22] Yuting Wang & Gopalan Nadathur (2013): *Towards Extracting Explicit Proofs from Totality Checking in Twelf*. In: *LFMTP 2013 - Proceedings of the 2013 ACM SIGPLAN Workshop on Logical Frameworks and Meta-Languages*, pp. 55–66, doi:10.1145/2503887.2503893.