

A Linear Logical Framework¹

Iliano Cervesato

*Advanced Engineering and Sciences Division
ITT Industries, Inc.
Alexandria, VA 22303-1410*

E-mail: iliano@itd.nrl.navy.mil

and

Frank Pfenning

*Department of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891*

E-mail: fp@cs.cmu.edu

Received ; revised ; accepted

We present the linear type theory $\lambda^{\Pi-\circ\&\top}$ as the formal basis for *LLF*, a conservative extension of the logical framework *LF*. *LLF* combines the expressive power of dependent types with linear logic to permit the natural and concise representation of a whole new class of deductive systems, namely those dealing with state. As an example we encode a version of *Mini-ML* with mutable references including its type system and its operational semantics, and describe how to take practical advantage of the representation of its computations.

© Academic Press

1. INTRODUCTION

A *logical framework* [27, 37] is a formal meta-language specifically designed to represent and reason about programming languages, logics and other formalisms that can be described as deductive systems. These frameworks consist of a *meta-representation language* with desirable computational and representational properties, normally a logic or a type theory, and of a *meta-representation methodology* that suggests how to best take advantage of the underlying meta-language to encode

¹ This work was supported by NSF Grant CCR-9303383. This research was completed while the first author was at the Department of Computer Science of Carnegie Mellon University. The second author was supported by the *Alexander-von-Humboldt-Stiftung* when working on this paper during a visit to the Department of Mathematics of the Technical University Darmstadt.

a given formal system. The logical framework LF [27] is among the most successful such proposals: it is based on the dependent type theory λ^Π , relies on the *judgments-as-types* representation methodology, and has been implemented as the higher-order constraint logic programming language Elf [38, 41]. LF and Elf have been widely used to study logical formalisms [43] and programming languages [33, 39] (see [44] for a survey).

Unfortunately, many constructs and concepts needed in common programming practice cannot be represented in a satisfactory way in meta-languages based on intuitionistic logic and intuitionistic type theory, such as LF . In particular, constructs based on the notion of state as found in imperative languages often escape an elegant formalization by means of these tools. Similarly, logical systems that, by definition (e.g. substructural logics) or by presentation (e.g. Dyckhoff’s contraction-free intuitionistic sequent calculus [17]), rely on destructive context operations require awkward encodings in an intuitionistic framework. Consequently the adequacy of the representation is difficult to prove and the formal meta-theory quickly becomes intractable.

Linear logic [21] provides a view of context formulas as resources, which can be exploited to model the notion of state, as described for example in [12, 28, 34, 53]. The current proposals put the emphasis on the issue of *representing* imperative constructs and resource-based logics, but appear inadequate for *reasoning* effectively about these representations. For example, the linear specification formalism *Forum* [34] has been used to give an immediate representation of the semantics of imperative programming languages [12]; however imperative computations are not effectively representable in this formalism and therefore meta-theoretic properties have not been encoded. On the other hand, intuitionistic type-theoretic frameworks such as LF make the representation of meta-reasoning easy, but do not have any notion of linearity built in. For example, the computations of the functional programming language *Mini-ML* can easily be expressed in LF , which permits automating the meta-theory of that language [33]. However, LF is not equipped to handle imperative computations as effectively, causing the meta-reasoning task to become a major challenge [42].

In this paper, we propose a conservative extension of the logical framework LF that permits representing linear objects and reasoning about them. We call this formalism *Linear LF*, or LLF for short. The language underlying LLF is the dependent type theory $\lambda^{\Pi \rightarrow \& \top}$, which extends LF ’s λ^Π with the linear connectives \rightarrow (linear implication), $\&$ (additive conjunction), and \top (additive truth), seen in this setting as type constructors. The language of objects of λ^Π is consequently extended with linear functional abstraction, additive pairs and unit, the corresponding destructors, and their equational theory. In order to keep the system simple we restrict the indices of type families to be linearly closed so that a type can depend only on intuitionistic assumptions, but not on linear variables. While at first this may appear to be a strong restriction, the expressive power of the resulting system does not seem to be hindered by this limitation.

The meta-representation methodology of LLF extends the judgments-as-types technique adopted in LF with a direct way to map state-related constructs and behaviors onto the linear operators of $\lambda^{\Pi \rightarrow \& \top}$. The resulting representations retain the elegance and immediacy that characterize LF encodings, and the ease of

proving their adequacy. *LLF* has so far been used to encode the syntax of linear logic, sequent calculus and natural deduction presentations of its semantics, imperative programming languages and their operational behavior, and a number of state-based games. We have also applied *LLF* to formalize aspects of the meta-theory of these systems such as the proof of cut elimination for classical linear logic, translations between linear natural deduction and sequent calculus, and properties of imperative languages such as type preservation [6].

The principal contributions of this paper are: (1) the definition of a uniform type theory admitting linear entities in conjunction with dependent types; (2) a thorough meta-theoretical investigation of this framework; and (3) the use of this system as a logical framework to represent and reason about problems that are not handled well by previous formalisms, either linear or intuitionistic. To our knowledge, $\lambda^{\Pi-\circ\&\top}$ is the first example in the literature of a linear type theory with dependencies. The case of simple types has been analyzed for example in [1, 2, 5, 32]. Subsequent work along the same lines of thought has been proposed by Ishtiaq and Pym in [30]. Both type theories were inspired by ideas in [35].

The paper is organized as follows. Section 2 describes the linear type theory $\lambda^{\Pi-\circ\&\top}$ on which *LLF* is founded. It also presents major results in its meta-theory, such as the strong normalization theorem and the decidability of verifying whether a term has a given type (type-checking) and of computing a type for a given term (type synthesis). Finally, it describes a canonical formulation of this language, which forms the basis for the meta-representation methodology adopted in *LLF*. Section 3 demonstrates the expressive power of *LLF* as a logical framework by providing an encoding of the syntax and the semantics of an imperative programming language, and by showing how to take practical advantage of the resulting representation of computations. Finally, Section 4 assesses the results and outlines future work. Further details about the work presented in this paper can be found in [6].

2. THE LINEAR TYPE THEORY $\lambda^{\Pi-\circ\&\top}$

In this section, we define the linear type theory $\lambda^{\Pi-\circ\&\top}$ on which *LLF* is founded. More importantly, we present the major results in its meta-theory that justify adopting it as a meta-representation language. In order to facilitate the description of *LLF* in the available space, we must assume that the reader is familiar with both the logical framework *LF* [27] and various presentations of linear logic [21, 22] and linear λ -calculi [1, 2]. We will also take advantage of the natural extension of the Curry-Howard isomorphism to linear logic by viewing types as formulas. Due to space constraints, we limit our discussion to the main results in the meta-theory of *LLF* and, moreover, present only sketches of their proofs. The interested reader is invited to consult [6] for further details about the presentation and the proofs in this section.

The discussion proceeds as follows: we first describe the syntax of $\lambda^{\Pi-\circ\&\top}$ in Section 2.1. Then, in Section 2.2, we introduce its semantics as a *pre-canonical* typing system, where typable terms are expected to be in η -long form, although they may contain β -redices. In Section 2.3, we focus our attention on the equational theory of this language. We present some basic properties in Section 2.4 and prove

strong normalization for this system in Section 2.5. In Section 2.6, we exploit this result to simplify the pre-canonical presentation of the semantics of $\lambda^{\Pi-\circ\&\top}$ as an equivalent *algorithmic* system, which allows easy proofs of the decidability of type-checking and type synthesis. These properties are presented Section 2.7. On the basis of these results, we devise in Section 2.8 a *canonical* system for $\lambda^{\Pi-\circ\&\top}$ whose only typable terms are both η -long and β -normal. The way $\lambda^{\Pi-\circ\&\top}$ is used as the language of the logical framework *LLF* relies on this formulation. Finally, in order to simplify the treatment of the case study in the next section, we extend the concrete syntax of *Elf*, the major implementation of *LF*, to the linear operators of $\lambda^{\Pi-\circ\&\top}$ in Section 2.9.

2.1. Language and Basic Operations

The linear type theory $\lambda^{\Pi-\circ\&\top}$ underlying *LLF* extends the language λ^{Π} of the logical framework *LF* with three connectives from linear logic, seen in this context as type constructors, namely *multiplicative implication* (\multimap), *additive conjunction* ($\&$), and *additive truth* (\top). The language of objects is augmented accordingly with the respective constructors and destructors. Linear types manipulate *linear assumptions* which we represent as distinguished declarations of the form $x \hat{?} A$ in the context; we write $x:A$ for context elements à la λ^{Π} and call them *intuitionistic assumptions*. The syntax of $\lambda^{\Pi-\circ\&\top}$ is given by the following annotated grammar, where we have separated the constructs not present in λ^{Π} with a double bar $\|$:

<i>Kinds:</i>	$K ::= \text{TYPE}$	(Class of types)
	$ \Pi x:A. K$	(Class of dependent type families)
<i>Types:</i>	$A ::= P$	(Base types)
	$ \Pi x:A_1. A_2$	(Intuitionistic function types)
	$\ A_1 \multimap A_2$	(Linear function types)
	$ A_1 \& A_2$	(Additive product types)
	$ \top$	(Additive unit type)
<i>Type families:</i>	$P ::= a$	(Type family constants)
	$ P M$	(Type family instantiation)
<i>Objects:</i>	$M ::= c \mid x$	(Object constants and variables)
	$ \lambda x:A. M \mid M_1 M_2$	(Intuitionistic functions)
	$\ \hat{\lambda} x:A. M \mid M_1 \hat{\sim} M_2$	(Linear functions)
	$ \langle M_1, M_2 \rangle \mid$	(Additive pairing)
	$\text{FST } M \mid \text{SND } M$	(Additive unit element)
	$ \langle \rangle$	(Additive unit element)
<i>Contexts:</i>	$\Psi ::= \cdot$	(Empty context)
	$ \Psi, x:A$	(Intuitionistic assumption)
	$\ \Psi, x \hat{?} A$	(Linear assumption)
<i>Signatures:</i>	$\Sigma ::= \cdot$	(Empty signature)
	$ \Sigma, a:K$	(Type family constant declaration)
	$ \Sigma, c:A$	(Object constant declaration)

Here x , c and a range over object-level variables, object constants, and type family constants, respectively. We adopt the convention of denoting linear variables with

$$\boxed{
\begin{array}{c}
\frac{}{\cdot = \cdot \bowtie \cdot} \text{s_dot} \qquad \frac{\Psi = \Psi' \bowtie \Psi''}{(\Psi, u \hat{\cdot} A) = (\Psi', u \hat{\cdot} A) \bowtie \Psi''} \text{s_lin1} \\
\frac{\Psi = \Psi' \bowtie \Psi''}{(\Psi, x : A) = (\Psi', x : A) \bowtie (\Psi'', x : A)} \text{s_int} \qquad \frac{\Psi = \Psi' \bowtie \Psi''}{(\Psi, u \hat{\cdot} A) = \Psi' \bowtie (\Psi'', u \hat{\cdot} A)} \text{s_lin2}
\end{array}
}$$

FIG. 1. Context Splitting

the letter u , possibly subscripted; we will however continue to write x for generic variables. In particular, we write $x \hat{\cdot} A$ for a context assumption whose exact nature (linear $— x \hat{\cdot} A$ — or intuitionistic $— x : A$) is unimportant. In addition to the names displayed above, we will often use N and B to range over objects and types respectively. Moreover, we denote generic terms, i.e., objects, types or kinds, with the letters U and V , possibly subscripted. As usual, we write $A \rightarrow U$ for $\Pi x : A. U$ whenever x does not occur in the type or kind U . Finally, an *index* is an argument M_i to a type family in a base type $P = a M_1 \dots M_n$.

The notions of free and bound variables are adapted from LF (notice the presence of a new binding construct: linear $\hat{\lambda}$ -abstraction). We denote with $\text{FV}(U)$ the free (linear or intuitionistic) variables of a term U . We extend this notation to contexts and write $\text{FV}(\Psi)$ to denote the union of $\text{FV}(A)$ for all $x \hat{\cdot} A$ in Ψ . As usual, we identify terms that differ only by the names of their bound variables and write $[M/x]U$ for the capture-avoiding substitution of M for x in the term U ; note that x can be either linear or intuitionistic. We extend this notation to contexts and write $[M/x]\Psi$ for the result of substituting M for x in the type of every assumption in Ψ . Finally we require variables and constants to be declared at most once in a context and in a signature, respectively.

In the following discussion, we will as usual drop the leading *empty sequence* (\cdot) from the representation of a context. Similarly, we overload the context constructor (\cdot) and use it to denote *sequence concatenation* as well. We do not state or prove the usual properties of this operation. Whenever we concatenate two contexts Ψ_1 and Ψ_2 we assume they do not declare common variables so that the resulting context (Ψ_1, Ψ_2) contains just one assumption for each declared variable. We denote with $\text{dom } \Psi$ the domain of context Ψ , defined as the set of variables declared in it, and write $\Psi|_\chi$ for the restriction of Ψ to the variables appearing in χ .

Below we will often need to refer to the *intuitionistic part* of a context Ψ . Therefore, we introduce the function $\overline{\Psi}$, defined as follows:

$$\left\{ \begin{array}{l} \overline{\cdot} = \cdot \\ \overline{\Psi, x : A} = \overline{\Psi}, x : A \\ \overline{\Psi, x \hat{\cdot} A} = \overline{\Psi} \end{array} \right.$$

We overload this notation and use $\overline{\Psi}$ to express the fact that the linear portion of the denoted context is constrained to be empty (e.g., in the all rules for type families in Figure 3).

Multiplicative connectives in linear logic require the context to be split among the premises of a binary rule (or the contexts in the premises to be merged in

the conclusion, depending on the point of view). We rely on the *context splitting judgment* to specify that the linear assumptions in a context Ψ must be distributed in the contexts Ψ_1 and Ψ_2 , while the intuitionistic assumptions should be shared. Whenever this is the case, the judgment $\Psi = \Psi_1 \bowtie \Psi_2$ is derivable. The rules in Figure 1 define this judgment.

Notice that, whenever the judgment $\Psi = \Psi_1 \bowtie \Psi_2$ is derivable, Ψ_1 and Ψ_2 differ from Ψ only by missing linear assumptions. In particular, the relative order of the declarations still mentioned in these contexts corresponds to the order in which they occur in Ψ . We anticipate that assumptions, either intuitionistic or linear, cannot depend on linear variables in $\lambda^{\Pi \multimap \circ \& \top}$. Therefore, splitting a context that is valid according to the specifications in the next subsection yields two valid contexts. Similarly, merging valid contexts having distinct names for their linear variables produces a valid context.

2.2. Pre-Canonical Forms

The meaning of the syntactic entities of a language can be presented in various forms, the choice being dictated by the aspects we want to emphasize. In this section, we define the semantics of $\lambda^{\Pi \multimap \circ \& \top}$ by means of a deductive system that we call *pre-canonical*, which enforces derivable terms to be in η -long form, although they might contain β -redices. The aim of the present section is to study the main properties of the type theory underlying *LLF*, ultimately the decidability of type checking. Relying on a pre-canonical system is particularly convenient for this purpose since it cleanly separates the practical desideratum of having extensionality as part of our language, commonly expressed by means of η -conversion rules, from the role of β -redices as the foundation of the equational theory of $\lambda^{\Pi \multimap \circ \& \top}$.

The main properties of *LF* were originally stated and proved for a type theory that did not enforce extensionality, but whose notion of definitional equality was restricted to β -equivalence [27]. However, the adequacy theorems that relate an object system to its *LF* encoding and the efficient implementation of this formalism as a logic programming language require considering λ^{Π} terms in canonical form. Therefore, the type theory that is used as a meta-representation language is based on $\beta\eta$ -equivalence. This discrepancy was known to Harper, Honsell and Plotkin when they first presented *LF* in 1987 [27]. A full treatment of the meta-theory of *LF* with $\beta\eta$ -equivalence was successively devised by various authors [14, 20, 49] and resulted in non-trivial complications.

The formulation of the semantics of *LLF* as a pre-canonical system has the advantage of forcing all derivable judgments to mention only terms in η -long form, as formally expressed below in Lemma 2.11. Indeed, all the issues concerning η -conversion are hardwired into the system and do not require explicit treatment. Consequently, the terms that we will ultimately produce are exactly the $\beta\eta$ -long terms needed in the adequacy theorems. That property allows us to focus on β -conversion and in particular to retain the simple techniques used in [27], without the above mentioned anomaly. Our approach was inspired by Felty's *canonical LF* [18].

The pre-canonical system for *LLF* is specified by means of a number of judgments. We first have seven judgments defining *pre-canonical* terms and the auxiliary notion of *pre-atomic* expressions, as well as the natural extension of these concepts to

contexts and signatures. The inference rules describing how to derive them are distributed over Figures 2 and 3. These rules will be discussed in detail shortly. The shape of these expressions is reported here:

- | | | |
|-------|---|---|
| (i) | $\vdash_{\Sigma}^p \Sigma \uparrow \text{Sig}$ | $(\Sigma \text{ is a pre-canonical signature})$ |
| (ii) | $\vdash_{\Sigma}^p \Psi \uparrow \text{Ctx}$ | $(\Psi \text{ is a pre-canonical context in } \Sigma)$ |
| (iii) | $\bar{\Psi} \vdash_{\Sigma}^p K \uparrow \text{Kind}$ | $(K \text{ is a pre-canonical kind in } \bar{\Psi} \text{ and } \Sigma)$ |
| (iv) | $\bar{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE}$ | $(A \text{ is a pre-canonical type in } \bar{\Psi} \text{ and } \Sigma)$ |
| | | |
| (v) | $\bar{\Psi} \vdash_{\Sigma}^p P \downarrow K$ | $(P \text{ is a pre-atomic type family of kind } K \text{ in } \bar{\Psi} \text{ and } \Sigma)$ |
| (vi) | $\Psi \vdash_{\Sigma}^p M \uparrow A$ | $(M \text{ is a pre-canonical object of type } A \text{ in } \Psi \text{ and } \Sigma)$ |
| (vii) | $\Psi \vdash_{\Sigma}^p M \downarrow A$ | $(M \text{ is a pre-atomic object of type } A \text{ in } \Psi \text{ and } \Sigma)$ |

Note that the judgments referring to types and kinds operate on purely intuitionistic assumptions, expressed by using the notation $\bar{\Psi}$ for their context. We will gain in clarity and conciseness in the following by relying on some abbreviations. Whenever the same property holds for each of the judgments (iii–vii) when applied to terms of the appropriate syntactic category, we write $\Psi \vdash_{\Sigma}^p U \uparrow \downarrow V$ and then refer to the generic terms U and V if needed. Moreover, if two or more such expressions occur in a statement, we assume that the arrows of the actual judgments match, unless explicitly stated otherwise. We take the liberty of adopting this notation also in the case of kinds, even though *Kind* is not a term and there are no pre-atomic kinds. Whenever the judgment $\Psi \vdash_{\Sigma}^p U \uparrow \downarrow V$ has a derivation \mathcal{P} , a fact that we will sometimes write as $\mathcal{P} :: (\Psi \vdash_{\Sigma}^p U \uparrow \downarrow V)$, we will often refer to U as the term being *validated* in this judgment, and call \mathcal{P} a *validation* of U .

The notion of *definitional equality* we consider is β -equivalence and it operates at the level of objects, type families, and kinds. Among the various possible presentations, we adopt *parallel nested reduction* (\longrightarrow), defined in Figure 4, and discussed shortly. We write \longrightarrow^* and \equiv for its transitive closure and the corresponding equivalence relation. We omit the obvious rules defining them. Cumulatively, we have the following nine judgments:

- | | | |
|-----------|-------------------------|---|
| (viii–x) | $U \longrightarrow V$ | $(U \text{ reduces to } V)$ |
| (xi–xiii) | $U \longrightarrow^* V$ | $(U \text{ transitively reduces to } V)$ |
| (xiv–xvi) | $U \equiv V$ | $(U \text{ is definitionally equal to } V)$ |

We need one further ingredient to cope with the multiplicative type constructor \multimap , namely the context splitting judgment presented in the previous section and defined in Figure 1.

- | | | |
|--------|-------------------------------|---|
| (xvii) | $\Psi = \Psi' \bowtie \Psi''$ | $(\Psi \text{ splits into } \Psi' \text{ and } \Psi'')$ |
|--------|-------------------------------|---|

We will now go through the rules defining $\lambda^{\Pi-\multimap\&\top}$ and describe the main ideas behind this formulation of the semantics of our language. We first concentrate on the typing judgments in Figures 2–3, and then discuss the notion of definitional

equality, founded on the reduction relation in Figure 4, to be discussed in the next section.

A term M of type A is in η -long form, or *pre-canonical*, if it is structured as a sequence consisting solely of constructors (abstractions, pairing, and unit) that matches the structure of the type A , applied to *pre-atomic* terms in those positions where objects of base type are required. A pre-atomic term consists of a sequence of destructors (applications and projections) that ends with a constant, a variable or another pre-canonical term, where the argument part of each application is also required to be pre-canonical. Note that this allows β -redices. This definition extends the usual notion of η -long forms of λ^Π to the linear type operators \multimap , $\&$ and \top of $\lambda^{\Pi\multimap\&\top}$ without insisting on β -normal forms.

It is characteristic for η -long forms that the type alone determines the structure of a term until we reach a base type, so, for example, any η -long term N of type $A = a \& (a \multimap a)$ will have the form

$$\langle N_1, \hat{\lambda}y : a. N_2 \rangle$$

where N_1 and N_2 are pre-atomic terms of type a . A variable x of the above type A is pre-atomic, but not pre-canonical. However it can be rewritten as

$$\langle \text{FST } x, \hat{\lambda}y : a. (\text{SND } x) \hat{\ } y \rangle$$

which is pre-canonical.

Consider first the rules displayed in the upper part of Figure 2. They validate pre-canonical objects M by deriving judgments of the form $\Psi \vdash_{\Sigma}^P M \uparrow A$. Rules **opc_unit**, **opc_pair**, **opc_llam**, and **opc_ilam** allow the construction of terms of the form $\langle \rangle$, $\langle M_1, M_2 \rangle$, $\hat{\lambda}u : A. M$, and $\lambda x : A. M$, respectively. These four operators are the *object constructors* of our calculus. We call these inference patterns *introduction rules* since, if we focus our attention on their type component, they introduce each of the type constructors of $\lambda^{\Pi\multimap\&\top}$ in their conclusion. The manner they handle their context is familiar from linear logic. Notice in particular that **opc_unit** (for \top) is applicable with any valid context and that the premises of rule **opc_pair** (for $\&$) share the same context, which also appears in its conclusion. These two are therefore additive constructors, in the sense of linear logic.

Rules **opc_llam** (for \multimap) and **opc_ilam** (for \rightarrow) differ only by the nature of the assumption they add to the context in their premise: linear in the case of the former, intuitionistic for the latter. The two remaining rules defining the object-level pre-canonical judgment leave the term in their central part unchanged. The type conversion rule **opc_eq** simply allows replacing the type component of a pre-canonical judgment with another type, under the condition that it is valid and definitionally equivalent to the original type.

Rule **opc_a** is the coercion from pre-atomic to pre-canonical terms. It is restricted to base types P . As a result, there is exactly one rule for each type constructor and one rule for base type, if we ignore type conversion for the moment. This guarantees the property stated above, namely, that the structure of a pre-canonical term is determined by the structure of its type. Type conversion (rule **opc_eq**) does not destroy this property since it affects only objects embedded as indices in base types, as it will become clear shortly.

Objects	
$\frac{\Psi \vdash_{\Sigma}^p M \downarrow P}{\text{opc_a}}$	$\frac{\Psi \vdash_{\Sigma}^p M \uparrow B \quad A \equiv B \quad \bar{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE}}{\text{opc_eq}}$
$\Psi \vdash_{\Sigma}^p M \uparrow P$	$\Psi \vdash_{\Sigma}^p M \uparrow A$
$\frac{\vdash_{\Sigma}^p \Psi \uparrow \text{Ctx}}{\text{opc_unit}}$	$\frac{\Psi \vdash_{\Sigma}^p M \uparrow A \quad \Psi \vdash_{\Sigma}^p N \uparrow B}{\text{opc_pair}}$
$\Psi \vdash_{\Sigma}^p \langle \rangle \uparrow \top$	$\Psi \vdash_{\Sigma}^p \langle M, N \rangle \uparrow A \& B$
$\frac{\Psi, u \hat{:} A \vdash_{\Sigma}^p M \uparrow B}{\text{opc_llam}}$	$\frac{\Psi, x : A \vdash_{\Sigma}^p M \uparrow B}{\text{opc_ilam}}$
$\Psi \vdash_{\Sigma}^p \hat{\lambda} u : A. M \uparrow A \multimap B$	$\Psi \vdash_{\Sigma}^p \lambda x : A. M \uparrow \Pi x : A. B$
.....	
$\frac{\Psi \vdash_{\Sigma}^p M \uparrow A}{\text{opa_c}}$	$\frac{\Psi \vdash_{\Sigma}^p M \downarrow B \quad A \equiv B \quad \bar{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE}}{\text{opa_eq}}$
$\Psi \vdash_{\Sigma}^p M \downarrow A$	$\Psi \vdash_{\Sigma}^p M \downarrow A$
	$\frac{\vdash_{\Sigma, c: A, \Sigma'}^p \bar{\Psi} \uparrow \text{Ctx}}{\text{opa_con}}$
	$\bar{\Psi} \vdash_{\Sigma, c: A, \Sigma'}^p c \downarrow A$
$\frac{\vdash_{\Sigma}^p \bar{\Psi}, u \hat{:} A, \bar{\Psi}' \uparrow \text{Ctx}}{\text{opa_lvar}}$	$\frac{\vdash_{\Sigma}^p \bar{\Psi}, x : A, \bar{\Psi}' \uparrow \text{Ctx}}{\text{opa_ivar}}$
$\bar{\Psi}, u \hat{:} A, \bar{\Psi}' \vdash_{\Sigma}^p u \downarrow A$	$\bar{\Psi}, x : A, \bar{\Psi}' \vdash_{\Sigma}^p x \downarrow A$
(No rule for \top)	$\frac{\Psi \vdash_{\Sigma}^p M \downarrow A \& B}{\text{opa_fst}}$
	$\Psi \vdash_{\Sigma}^p \text{FST } M \downarrow A$
	$\frac{\Psi \vdash_{\Sigma}^p M \downarrow A \& B}{\text{opa_snd}}$
	$\Psi \vdash_{\Sigma}^p \text{SND } M \downarrow B$
$\frac{\Psi' \vdash_{\Sigma}^p M \downarrow A \multimap B \quad \Psi'' \vdash_{\Sigma}^p N \uparrow A \quad \Psi = \Psi' \bowtie \Psi''}{\text{opa_lapp}}$	
	$\Psi \vdash_{\Sigma}^p M \wedge N \downarrow B$
$\frac{\Psi \vdash_{\Sigma}^p M \downarrow \Pi x : A. B \quad \bar{\Psi} \vdash_{\Sigma}^p N \uparrow A}{\text{opa_iapp}}$	
	$\Psi \vdash_{\Sigma}^p M N \downarrow [N/x]B$

FIG. 2. Pre-Canonical Deduction System for $\lambda^{\Pi \multimap \& \top}$, Objects

The rules defining the pre-atomic judgment at the level of objects, $\Psi \vdash_{\Sigma}^p M \downarrow A$, are displayed in the lower part of Figure 2. They validate constants (rule **opa_con**) and linear and intuitionistic variables (rules **opa_lvar** and **opa_ivar**, respectively). They also allow the formation of the terms $M N$, $M \wedge N$, $\text{FST } M$, and $\text{SND } M$ (rules **opa_iapp**, **opa_lapp**, **opa_fst**, and **opa_snd**, respectively), whose main operators we call *destructors*. The latter four inference figures are called *elimination rules* since they permit taking apart each of the type constructors of $\lambda^{\Pi \multimap \& \top}$ from one of their premise, with the exception of \top . The role played by linear assumptions in $\lambda^{\Pi \multimap \& \top}$ is particularly evident in these rules. Indeed, an axiom rule (**opa_con**, **opa_lvar**, and **opa_ivar**) can be applied only if the linear part of its context is empty, or contains just the variable to be validated, with the proper type; this is expressed by using the $\hat{\cdot}$ notation. Linearity appears also in the elimination rule for \multimap , where we rely on the splitting judgment defined in Figure 1 to manage the context for this connective in rule **opa_lapp**. Observe also that the context of the argument part of an intuitionistic application, in rule **opa_iapp**, is constrained not

to contain any linear assumption. Two remaining rules define pre-atomic derivability for the level of objects. The semantics of the equivalence rule **opa_eq** is similar to its pre-canonical counterpart.

The coercion from pre-canonical to pre-atomic objects, **opa_c**, is unrestricted in its type. This means that destructors can be directly applied to constructors, that is, objects may contain redices. If we omit this rule (or restrict it to base type, which is equivalent), we obtain precisely the canonical forms, that is, those η -long forms which contain no β -redices.

The rules concerning linear objects in Figure 2 define the behavior of linear types. If we ignore the objects and the distinction between pre-canonical and pre-atomic judgments, they correspond to the specification of the familiar rules for the linear connectives \top , $\&$, and \multimap , presented in a *natural deduction* style. It is easy to prove the equivalence to the usual sequent formulation. The objects that appear on the left of these types record the structure of a natural deduction proof for the corresponding linear formulas. The dependent function type $\Pi x : A. B$ that $\lambda^{\Pi \multimap \& \top}$ inherits from λ^{Π} generalizes both intuitionistic implication $A \rightarrow B$ (customarily defined as $!A \multimap B$ in linear logic) and the universal quantifier $\forall x. B$, where A plays the role of the type of the (intuitionistic) variable x . With this interpretation, $\lambda^{\Pi \multimap \& \top}$ encompasses all the connectives and quantifiers of the freely generated fragment of the language of linear hereditary Harrop formulas, on which the programming language *Lolli* is based [28]. Additionally, $\lambda^{\Pi \multimap \& \top}$ offers the characteristic features of a type theory: higher-order functions, proof terms, and type families indexed by arbitrary objects, possibly higher-order and linear.

Admitting other linear connectives in this language is problematic since the remaining operators of linear logic would introduce in the equational theory a form of reductions known as *commuting conversions*, which would destroy the possibility of achieving unique normal forms. On the other hand, the semantics of the other linear connectives can be easily emulated in $\lambda^{\Pi \multimap \& \top}$, as shown in [42].

We now turn to the judgments validating types, kinds, contexts, and signatures, treated in Figure 3. The rules defining the pre-canonical judgment $\overline{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE}$ simply specify that every valid type in the system should result from the combination of base types (rule **fpc_a**) by means of the type constructors Π , \multimap , $\&$, and \top (rules **fpc_dep**, **fpc_limp**, **fpc_with**, and **fpc_top**, respectively). Notice the differences in the rules concerning the two function type constructors: the validity of A in $\Pi x : A. B$ is implicitly ascertained when checking the validity of the context in the premise; instead, the type A in $A \multimap B$ is to be validated explicitly since no assumption is inserted in the context. The rules for the pre-atomic type family judgment $\overline{\Psi} \vdash_{\Sigma}^p A \downarrow K$ simply verify that base types are syntactically well-formed and respect the kind declaration of their leading type family constant (rules **fpa_iapp** and **fpa_con**). Notice the presence of an equivalence rule: **fpa_eq**. Finally, the rules defining the pre-canonical kind judgment, $\overline{\Psi} \vdash_{\Sigma}^p K \uparrow \text{Kind}$, check that every type appearing in K is valid. Note that this judgment is invoked only when validating a signature. The remaining rules in Figure 3 consider signatures and context. They specify that a signature is valid if the type or kind of every item declared in it is itself valid. Similarly, a context is valid if the type of each of its assumptions is valid.

Signatures	$\frac{}{\vdash^p \cdot \uparrow \text{Sig}} \text{sp_dot}$
	$\frac{\vdash^p \Sigma \uparrow \text{Sig} \quad \cdot \vdash_{\Sigma}^p A \uparrow \text{TYPE}}{\vdash^p \Sigma, c: A \uparrow \text{Sig}} \text{sp_obj} \qquad \frac{\vdash^p \Sigma \uparrow \text{Sig} \quad \cdot \vdash_{\Sigma}^p K \uparrow \text{Kind}}{\vdash^p \Sigma, a: K \uparrow \text{Sig}} \text{sp_fam}$
Contexts	$\frac{\vdash^p \Sigma \uparrow \text{Sig}}{\vdash_{\Sigma}^p \cdot \uparrow \text{Ctx}} \text{cp_dot}$
	$\frac{\vdash_{\Sigma}^p \Psi \uparrow \text{Ctx} \quad \bar{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE}}{\vdash_{\Sigma}^p \Psi, x: A \uparrow \text{Ctx}} \text{cp_int} \qquad \frac{\vdash_{\Sigma}^p \Psi \uparrow \text{Ctx} \quad \bar{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE}}{\vdash_{\Sigma}^p \Psi, u \hat{?} A \uparrow \text{Ctx}} \text{cp_lin}$
Kinds	$\frac{\vdash_{\Sigma}^p \bar{\Psi} \uparrow \text{Ctx}}{\bar{\Psi} \vdash_{\Sigma}^p \text{TYPE} \uparrow \text{Kind}} \text{kpc_type} \qquad \frac{\bar{\Psi}, x: A \vdash_{\Sigma}^p K \uparrow \text{Kind}}{\bar{\Psi} \vdash_{\Sigma}^p \Pi x: A. K \uparrow \text{Kind}} \text{kpc_dep}$
Types/type families	$\frac{\bar{\Psi} \vdash_{\Sigma}^p P \downarrow \text{TYPE}}{\bar{\Psi} \vdash_{\Sigma}^p P \uparrow \text{TYPE}} \text{fpc_a}$
	$\frac{\vdash_{\Sigma}^p \bar{\Psi} \uparrow \text{Ctx}}{\bar{\Psi} \vdash_{\Sigma}^p \top \uparrow \text{TYPE}} \text{fpc_top} \qquad \frac{\bar{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE} \quad \bar{\Psi} \vdash_{\Sigma}^p B \uparrow \text{TYPE}}{\bar{\Psi} \vdash_{\Sigma}^p A \& B \uparrow \text{TYPE}} \text{fpc_with}$
	$\frac{\bar{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE} \quad \bar{\Psi} \vdash_{\Sigma}^p B \uparrow \text{TYPE}}{\bar{\Psi} \vdash_{\Sigma}^p A \multimap B \uparrow \text{TYPE}} \text{fpc_limp} \qquad \frac{\bar{\Psi}, x: A \vdash_{\Sigma}^p B \uparrow \text{TYPE}}{\bar{\Psi} \vdash_{\Sigma}^p \Pi x: A. B \uparrow \text{TYPE}} \text{fpc_dep}$
	<hr style="border-top: 1px dotted black;"/>
(No fpc_eq , no fpa_c)	$\frac{\bar{\Psi} \vdash_{\Sigma}^p P \downarrow K \quad K \equiv K' \quad \bar{\Psi} \vdash_{\Sigma}^p K' \uparrow \text{Kind}}{\bar{\Psi} \vdash_{\Sigma}^p P \downarrow K'} \text{fpa_eq}$
	$\frac{\vdash_{\Sigma, a: K, \Sigma'}^p \bar{\Psi} \uparrow \text{Ctx}}{\bar{\Psi} \vdash_{\Sigma, a: K, \Sigma'}^p a \downarrow K} \text{fpa_con} \qquad \frac{\bar{\Psi} \vdash_{\Sigma}^p P \downarrow \Pi x: A. K \quad \bar{\Psi} \vdash_{\Sigma}^p N \uparrow A}{\bar{\Psi} \vdash_{\Sigma}^p P N \downarrow [N/x]K} \text{fpa_iapp}$

FIG. 3. Pre-Canonical Deduction System for $\lambda^{\Pi \multimap \& \top}$, Kinds and Types

In the rules in Figures 2–3, types and kinds are always checked using a purely intuitionistic context. This has the effect of preventing valid types from containing free linear variables (although bound linear variables are admitted). Therefore, although the indices of type families are in general linear objects, these terms can refer only to context variables that are intuitionistic. We say that indices are *linearly closed*. Loosening this restriction would require admitting linear dependent function types in our language, corresponding to linear quantifiers. Preliminary investigations indicate that this would lead to tremendous complications in the typing rules of the language, not to speak of its meta-theory. For example, we could not expect a purely intuitionistic context anymore when looking up a variable, context

splitting would rely on the typing judgments since a blind split might violate linear dependencies, and linear dependent types have been observed to interact in a complex manner with the other type constructors, in particular \top . On the other hand, very few of our examples could have taken advantage of a linear version of Π . In every case, using its intuitionistic counterpart in its place did not substantially alter the resulting representation, nor its adequacy. In conclusion, although a dependent version of \multimap appears beneficial for certain applications, we are lead to believe that the consequent complications in the meta-theory of the language might outweigh the potential advantages.

We classify the rules in Figures 2–3 into *essential* and *non-essential* rules. We count among the latter class the type conversion rules (**opc.eq**, **opa.eq**, and **fpa.eq**) and coercion **opa.c**. All other rules are considered essential. A major task in this section of the paper will be to either hide or eliminate the non-essential rules of $\lambda^{\Pi \multimap \& \top}$. Hiding the equivalence rules will permit an easy proof of the decidability of type-checking for this language. Showing that the rule **opa.c** can be eliminated (in the sense that a type is inhabited with this rule if and only if it is inhabited without it) amounts to showing that canonical forms exist for all objects and types. This is a necessary condition for adopting $\lambda^{\Pi \multimap \& \top}$ as the underlying type theory of the logical framework *LLF*.

We now turn to the reduction semantics of $\lambda^{\Pi \multimap \& \top}$, partially defined in Figure 4. The notion of definitional equality that we consider is the equivalence relation \equiv constructed on the congruence relation \longrightarrow . The basis of this congruence consists of the following β -reduction rules:

$$\begin{array}{ll} \beta_{fst} : \text{FST } \langle M, N \rangle \longrightarrow M & \beta_{lapp} : (\hat{\lambda}u : A. M) \wedge N \longrightarrow [N/u]M \\ \beta_{snd} : \text{SND } \langle M, N \rangle \longrightarrow N & \beta_{rapp} : (\lambda x : A. M) N \longrightarrow [N/x]M \end{array}$$

As usual, we call the expressions appearing on the left-hand side of the arrow *redices*. The only possible redex in λ^{Π} is β_{lapp} . We adopt the standard terminology and call a term U that does not contain β -redices *normal*, or β -*normal*. This definition extends immediately to contexts and signatures. Another task of this section will be to show that every valid entity in $\lambda^{\Pi \multimap \& \top}$ can be reduced to a normal form, and that this normal form is itself valid in *LLF*. Finally, a term U is *reducible* if there is a derivation of the judgment $U \longrightarrow V$ for some V .

2.3. Equational Theory

We now attack the formal study of $\lambda^{\Pi \multimap \& \top}$. We aim at proving that this type theory has all the desirable properties that a formalism should possess in order to be a suitable meta-language for a logical framework. In particular, we will show that $\lambda^{\Pi \multimap \& \top}$ is strongly normalizing, admits unique normal forms, and that type-checking is decidable for this language. These properties rely on a large number of lemmas and it is a challenge of its own to organize the overall meta-theory in a linear sequence of results with simple dependencies. This organization relies crucially on the details of the formulation of the semantics of our type-theory. The adoption of a pre-canonical system that imposes extensionality, rather than just permitting it via additional rules, simplifies the development of the meta-theory of $\lambda^{\Pi \multimap \& \top}$ considerably. The principal consequence of this choice is that

Objects	<i>Congruences</i>
$\frac{}{c \longrightarrow c} \text{ or_con} \qquad \frac{}{x \longrightarrow x} \text{ or_var} \qquad \frac{}{\langle \rangle \longrightarrow \langle \rangle} \text{ or_unit}$	
$\frac{M \longrightarrow M'}{\text{FST } M \longrightarrow \text{FST } M'} \text{ or_fst} \qquad \frac{M \longrightarrow M'}{\text{SND } M \longrightarrow \text{SND } M'} \text{ or_snd} \qquad \frac{M \longrightarrow M' \quad N \longrightarrow N'}{\langle M, N \rangle \longrightarrow \langle M', N' \rangle} \text{ or_pair}$	
$\frac{M \longrightarrow M' \quad N \longrightarrow N'}{M \wedge N \longrightarrow M' \wedge N'} \text{ or_lapp} \qquad \frac{A \longrightarrow A' \quad M \longrightarrow M'}{\hat{\lambda}u : A. M \longrightarrow \hat{\lambda}u : A'. M'} \text{ or_llam}$	
$\frac{M \longrightarrow M' \quad N \longrightarrow N'}{MN \longrightarrow M' N'} \text{ or_iapp} \qquad \frac{A \longrightarrow A' \quad M \longrightarrow M'}{\lambda x : A. M \longrightarrow \lambda x : A'. M'} \text{ or_ilam}$	
<i>β-reductions</i>	
$\frac{M \longrightarrow M'}{\text{FST } \langle M, N \rangle \longrightarrow M'} \text{ or_beta_fst} \qquad \frac{N \longrightarrow N'}{\text{SND } \langle M, N \rangle \longrightarrow N'} \text{ or_beta_snd}$	
$\frac{M \longrightarrow M' \quad N \longrightarrow N'}{(\hat{\lambda}u : A. M) \wedge N \longrightarrow [N'/u]M'} \text{ or_beta_lin} \qquad \frac{M \longrightarrow M' \quad N \longrightarrow N'}{(\lambda x : A. M) N \longrightarrow [N'/x]M'} \text{ or_beta_int}$	
Types/types families	<i>Congruences</i>
$\frac{}{a \longrightarrow a} \text{ fr_con} \qquad \frac{P \longrightarrow P' \quad N \longrightarrow N'}{PN \longrightarrow P' N'} \text{ fr_iapp}$	
$\frac{}{\top \longrightarrow \top} \text{ fr_top} \qquad \frac{A \longrightarrow A' \quad B \longrightarrow B'}{A \& B \longrightarrow A' \& B'} \text{ fr_with}$	
$\frac{A \longrightarrow A' \quad B \longrightarrow B'}{A \multimap B \longrightarrow A' \multimap B'} \text{ fr_limp} \qquad \frac{A \longrightarrow A' \quad B \longrightarrow B'}{\Pi x : A. B \longrightarrow \Pi x : A'. B'} \text{ fr_dep}$	
(No β -reductions for types and type families)	
Kinds	<i>Congruences</i>
$\frac{}{\text{TYPE} \longrightarrow \text{TYPE}} \text{ kr_type} \qquad \frac{A \longrightarrow A' \quad K \longrightarrow K'}{\Pi x : A. K \longrightarrow \Pi x : A'. K'} \text{ kr_dep}$	
(No β -reductions for kinds)	

FIG. 4. Parallel Nested Reduction for $\lambda^{\Pi \multimap \& \top}$

definitional equality can be based entirely on β -reduction and, more importantly, can be defined independently from typing, as shown in Figure 4. Therefore, the analysis of the equational theory of $\lambda^{\Pi \multimap \& \top}$ can be conducted in a totally self-contained manner. Indeed, the results that we will present in this subsection, in particular the Church-Rosser theorem, apply to arbitrary $\lambda^{\Pi \multimap \& \top}$ terms and not just to those that are valid according to the typing rules of our language. This apparently unnecessary generality is the first step towards disentangling the meta-theory of $\lambda^{\Pi \multimap \& \top}$. Had we relaxed extensionality by considering an equational theory containing η -conversion rules, as normally done in the literature, we would have been forced to provide mutually recursive definitions for the typing and the

definitional equality judgments. In particular, the equational theory of the resulting formalism could not be studied in isolation and most of its meta-theory would collapse in one dense theorem consisting of a discouraging number of mutually dependent properties. We will come back to this point at the end of this subsection.

As we already anticipated, the principal result of this section will be the Church-Rosser theorem for parallel nested reduction. We will rely on this property in order to prove the uniqueness of normal forms, and therefore the decidability of the equational theory of *LLF*.

The proofs of the results in this section adapt the technique originally devised by Tait and Martin-Löf for the traditional untyped λ -calculus [3]. A very detailed presentation of that method, as well as its formalization in *Elf*, can be found in [40]. We deviate from this presentation in order to take into account all the entities of *LLF* that participate in the definition of parallel nested reduction. Specifically, we treat the linear constructs of our language and the new forms of β -reduction they introduce; we also need to consider types and kinds.

The parallel nested reduction strategy defined in Figure 4 is based on the four β -reduction rules **or_beta_fst**, **or_beta_snd**, **or_beta_lin**, and **or_beta_int**. All the other rules are congruences that allow applying reductions to subterms. Notice that the β -reduction rules are directional: the expression on the left-hand side of the arrow is a β -redex, and we like to think of the expression on the right-hand side as “simpler”, even if it may be larger, or contain more β -redices, than the term on the other side of the arrow.

A key result in the study of the definitional equality of $\lambda^{\Pi \rightarrow \& \top}$ is that substituting a variable in a reducible term maintains its reducibility. This property is formalized and generalized in the following lemma, where $\mathcal{R} :: J$ is used as an abbreviation for “there is a derivation \mathcal{R} of the judgment J ”.

LEMMA 2.1 (Substitution).

Assume that there exist a derivation for $N \longrightarrow N'$. Then,

- i. if $\mathcal{R} :: U \longrightarrow U'$, then $[N/x]U \longrightarrow [N'/x]U'$;
- ii. if $\mathcal{R} :: U \longrightarrow^* U'$, then $[N/x]U \longrightarrow^* [N'/x]U'$;
- iii. if $\mathcal{R} :: U \equiv U'$, then $[N/x]U \equiv [N'/x]U'$.

Proof. We proceed by induction on \mathcal{R} in each case. ■

A term can contain several β -redices and the parallel reduction strategy can reduce any of them, possibly zero or more than one. Therefore, a term U can in general reduce to a number of distinct terms U_1, \dots, U_n . However, a fundamental property of this strategy is that there always exist a common term V to which all these terms are reducible. This is known in the literature as the *Diamond property* and it is stated below.

LEMMA 2.2 (Diamond property).

If $\mathcal{R}' :: U \longrightarrow U'$ and $\mathcal{R}'' :: U \longrightarrow U''$, then there is a term V such that $U' \longrightarrow V$ and $U'' \longrightarrow V$.

Proof. By induction on the structure of U and inversion on \mathcal{R}' and \mathcal{R}'' . Functional β -reduction is handled through the Substitution Lemma. ■

Although one run of parallel nested reduction has the possibility of reducing several β -redices, it is not sufficient in general to produce the normal form for a term, even when it exists. For example, the following judgment shows on the right-hand side the simplest term the expression on the left-hand side can be reduced to in one step:

$$(\lambda x:(a \rightarrow a) \rightarrow (a \rightarrow a).x c) (\lambda y:a \rightarrow a.y) \longrightarrow (\lambda y:a \rightarrow a.y) c$$

Notice that one further step would suffice to obtain the normal form of that expression, which is c .

In order to achieve normal forms when they exist, we need to chain parallel nested reductions by taking their transitive closure. *Confluence* extends the Diamond Property to \longrightarrow^* , while the *Church-Rosser Theorem* states that it is always possible to reduce equivalent entities back to a common term. These two properties are stated below. They follow from the Diamond Property by virtue of general techniques [16].

THEOREM 2.1 (Church-Rosser).

Confluence: If $U \longrightarrow^* U'$ and $U \longrightarrow^* U''$, then there is a term V such that $U' \longrightarrow^* V$ and $U'' \longrightarrow^* V$.

Church-Rosser: If $U' \equiv U''$, then there is a term V such that $U' \longrightarrow^* V$ and $U'' \longrightarrow^* V$. □

The properties above apply to arbitrary terms, possibly ill-typed or in general invalid according to the pre-canonical system above (indeed, our example above contained the subterm $\lambda y:a \rightarrow a.y$ which is not η -expanded). Although definitional equality is always invoked with valid terms in the rules in Figures 2–3, intermediate terms participating in the equivalence derivation might not be pre-canonical. We will show that it is possible to limit the intermediate terms produced during a definitional equality test to entities that are valid.

Arbitrary terms do not have in general a normal form. A classical example is the term $(\lambda x:a.x x) (\lambda x:a.x x)$. This term reduces to itself and therefore it is not possible to eliminate the β -redex it contains by reduction. However, every valid term, i.e., every term that appears in a derivable typing judgment in our pre-canonical system admits a normal form. Furthermore, the strong normalization theorem proved below will show that the order in which β -redices are reduced is not important.

We conclude this subsection with a short discussion about notions of definitional equality that includes rules for η -conversion. If we were simply to add η -rules based on the following reductions:

$$\begin{aligned} \eta_{pair} &: \langle \text{fst } M, \text{snd } M \rangle \longrightarrow M \\ \eta_{lam} &: \hat{\lambda}u:A. M \hat{u} \longrightarrow M \\ \eta_{ilam} &: \lambda x:A. M x \longrightarrow M \end{aligned}$$

with u and x not occurring free in M in the last two rules, then the Church-Rosser property would cease to hold: it is observed in [27] that the term $\lambda x:A. (\lambda y:B. M) x$ reduces via η_{ilam} to $\lambda y:B. M$ and by β to $\lambda x:A. [x/y]M$, i.e. $\lambda y:A. M$; however, the two resulting terms can be equated only if A and B have a common reduct, which is not the case for certain ill-typed terms.

Therefore, the simple η -conversion rules above are insufficient to capture the notion of definitional equality we are interested in. The situation is actually more serious than that: in the presence of a unit type, \top in our case, the equational theory we need cannot be axiomatized by means of schematic reduction rules such as the above. Indeed, $\langle \rangle$ is the only inhabitant of type \top and therefore every object of type \top is η -equivalent to $\langle \rangle$. However, adding the rule $M \longrightarrow \langle \rangle$ is clearly unsound since it does not take typing information into account. The correct approach to this problem requires typed definitional equality judgments. The specific inference rules that handle η -expansion are called *extensionality* rules and only vaguely resemble the η -conversions presented above. It seems that (untyped) η -rules can serve as a foundation of the definitional equality of type theories only in strongly restricted circumstances.

2.4. Fundamental Properties

The purpose of this subsection is to illustrate the basic properties of the pre-canonical deduction system presented earlier. Many of these results are interesting by themselves since they provide insight about the type theory of *LLF* beyond what is apparent from the inference rules. Moreover, most of these properties will play a role in the development of the proof of the decidability of type checking for our language.

First, we summarize the principal properties regarding the occurrences of free variables. The long proof of this lemma can be found in [6]. Only a limited number of further results will mention free variables or domains explicitly. In reading this statement, recall we can tacitly rename bound variables and that the names of all variables declared in a context are distinct. Moreover, we remind the reader that $\Psi \vdash_{\Sigma}^p U \Downarrow V$ is an abbreviation for the pre-canonical or pre-atomic judgments at the level of objects, types, and kinds.

LEMMA 2.3 (Free variables).

- i.* If $\Psi \vdash_{\Sigma}^p U \Downarrow V$, then $\text{FV}(U) \subseteq \text{dom } \Psi$ and $\text{FV}(V) \subseteq \text{dom } \bar{\Psi}$.
- ii.* If $\Psi \vdash_{\Sigma}^p U \Downarrow V$ or $\vdash_{\Sigma}^p \Psi \uparrow \text{Ctx}$, then $\text{FV}(\Psi) \subseteq \text{dom } \bar{\Psi}$.
- iii.* If $\Psi \vdash_{\Sigma}^p U \Downarrow V$ or $\vdash_{\Sigma}^p \Psi \uparrow \text{Ctx}$ or $\vdash^p \Sigma \uparrow \text{Sig}$, then $\text{FV}(\Sigma) = \emptyset$.
- iv.* If $\Psi, x \dot{\vdash} A, \Psi' \vdash_{\Sigma}^p U \Downarrow V$ or $\vdash_{\Sigma}^p \Psi, x \dot{\vdash} A, \Psi' \uparrow \text{Ctx}$, then $\text{FV}(A) \cup \text{FV}(\Psi) \subseteq \text{dom } \bar{\Psi}$. □

This lemma provides some insight about how and where free variables can appear in a derivable judgment in $\lambda^{\Pi-\circ\&\top}$. By (i) and (ii), we have that all free variables occurring in a valid judgment must be declared in its context. Property (iv) specifies that the free variables in an assumption must have been declared in the part of the context that is to its left, i.e., assumptions can depend only on declarations made before them. Item (iii) in the lemma states instead that a signature cannot contain free variables. All these properties already hold in λ^{Π} . The peculiarities of our

language appear when analyzing the role of free variables that are assumed linearly. Since the context of the judgments for types and kinds are strictly intuitionistic, (i) entails that free linear variables are permitted only in valid terms from the level of objects. Moreover, (ii) implies that no assumption in the context can depend on a linear variable. These strict constraints are a consequence of the property of our language of permitting only linearly closed expressions as indices to type families.

We now present a number of properties that can be seen as admissible rules of inference. First we have that assumptions in the context can be exchanged freely as long as they do not violate dependencies among them. More precisely, if $x :: A$ immediately precedes $y :: B$ and x does not occur free in B , then the relative order of these two assumptions can be exchanged. This idea is generalized in the lemma below. Permutation depends on weakening, which itself requires permutation in its proof. Therefore, we need to state and prove both properties at the same time. Notice that weakening forbids adding linear assumptions into the context.

LEMMA 2.4 (Structural properties of contexts).

Permutation:

- i. If $\mathcal{P}' :: (\Psi, \Psi', x :: A, \Psi'' \vdash_{\Sigma}^p U \Downarrow V)$ and $\mathcal{P}'' :: (\overline{\Psi} \vdash_{\Sigma}^p A \Uparrow \text{TYPE})$, then $\Psi, x :: A, \Psi', \Psi'' \vdash_{\Sigma}^p U \Downarrow V$.
- ii. If $\mathcal{P}' :: (\vdash_{\Sigma}^p \Psi, \Psi', x :: A, \Psi'' \Uparrow Ctx)$ and $\mathcal{P}'' :: (\overline{\Psi} \vdash_{\Sigma}^p A \Uparrow \text{TYPE})$, then $\vdash_{\Sigma}^p \Psi, x :: A, \Psi', \Psi'' \Uparrow Ctx$.

Weakening:

If $\mathcal{P}' :: (\Psi \vdash_{\Sigma}^p U \Downarrow V)$ and $\mathcal{P}'' :: (\vdash_{\Sigma}^p \Psi, \overline{\Psi}' \Uparrow Ctx)$, then $\Psi, \overline{\Psi}' \vdash_{\Sigma}^p U \Downarrow V$.

Proof. By simultaneous induction on \mathcal{P}' and, in the case of permutation, on the length of Ψ'' . ■

The permutation property has important consequences on the linear assumptions in the context. As we described earlier, no assumption in a context Ψ can depend on a linear variable. Therefore, the permutation lemma allows us to shift all these variables to end of Ψ . Let us write $\widehat{\Psi}$ for the linear assumptions of a context Ψ . Notice that, because of possible dependencies on intuitionistic variables, $\widehat{\Psi}$ is not necessarily a valid context. We can then write Ψ as $(\overline{\Psi}, \widehat{\Psi})$, or even as “ $\overline{\Psi}; \widehat{\Psi}$ ” as in recent presentations of linear logic that maintain intuitionistic and linear assumptions in different contexts (separated by “;”) [22, 28, 45]. Furthermore, since the variables in the domain of $\widehat{\Psi}$ cannot occur in $\overline{\Psi}$, we are free to permute the contents of this part of the context. Therefore, $\widehat{\Psi}$ can be treated as a multiset.

We could expect a further property, *strengthening*, to be part of the statement of the lemma above; the presence of this property would actually make that proof easier, but unfortunately we do not have yet the tools to prove it.

Strengthening states that, whenever a variable is declared in a context but does not occur free anywhere in a judgment, then it can safely be dropped and the judgment will still be provable. However, strengthening requires the strong normalization theorem, proved in Section 2.5. Even though a variable does not occur in a derivable judgment, it is possible that the application of one of the definitional equivalence rules produces a term containing it, so that not the judgment itself,

but its derivation mentions it. These uses of equivalence are non-essential and can be removed from the derivation. However, we will be able to prove this only as a by-product of strong normalization.

Next, we present a technical result that, although of minor importance in itself, plays an important role in the statement of the adequacy theorems for the representation of an object language. Specifically, when the meta-theory of the object language expects certain objects to behave as if they were linear hypotheses, an adequate encoding would require free linear variables in the indices of base types. This is not achievable in *LLF* since our language does not admit linear dependent types. We bypass the problem by encoding these linear entities as intuitionistic assumptions. Linearity conditions can be checked at an earlier stage of the computation, or be kept as intrinsic invariants of the object deductive system. This technique permits us to give an effective representation to complex linear judgments without dealing with the complications of linear dependent types. Examples showing how this issue is handled in *LLF* can be found in [6].

The lemma below states that whenever a derivable term mentions a linear variable, we can safely make it intuitionistic. Intuitively, linear variables must be used once while intuitionistic variables can be used as many times as desired.

LEMMA 2.5 (Promotion).

If $\mathcal{P} :: (\Psi, u \dot{?} A, \Psi' \vdash_{\Sigma}^p M \Downarrow B)$, then $\Psi, u : A, \Psi' \vdash_{\Sigma} M \Downarrow B$.

Proof. By induction on the structure of \mathcal{P} . ■

An important ingredient of the proofs of the theorems below are the lemmas that we call transitivity, following the terminology in [27]. These results permit interpreting assumptions as place-holders for unspecified derivations. Whenever a provable judgment depends on the assumption $x \dot{?} A$, any derivation of a term N of type A satisfying certain context conditions can be substituted into the original derivation and maintain its validity. Therefore, judgments containing assumptions can be thought of as parametric expressions. The transitivity lemmas specify how to instantiate these parameters.

These results contribute to the suitability of $\lambda^{\Pi \rightarrow \circ \& \top}$ as the meta-language of the logical framework *LLF*. They are the formal justification of the representation of the hypothetical and parametric judgments, so common in formal systems, as simple and dependent function types, respectively. The transitivity lemmas, together with the inversion lemma below, postulate that these operators have a semantics that mimics the behavior of those forms of judgment. *LLF* extends this correspondence, already present in *LF*, to capture hypothetical judgments where the hypotheses are linear.

The transitivity lemmas are tightly connected to several results in logic and type theory. The interpretation depends on which part of the judgments we focus our attention on. From the point of view of the λ -calculus embedded in our language, these lemmas can be seen as substitution principles since they describe how variables can be substituted into valid terms while preserving validity. In this, they are closely related to the notion of subject reduction for functional objects. From the logical perspective, under the interpretation of types as formulas, the transitivity lemmas state the admissibility of the cut rule for intuitionistic and linear formulas.

Whenever a formula B relies on an assumption A , any evidence of the validity of A , possibly on the basis of further assumptions, can be included directly in an equivalent proof of B that does not mention A among its hypotheses.

Linear and intuitionistic assumptions need to be treated separately since they require a different structuring of the context. Therefore, we distinguish two transitivity lemmas. This corresponds to differentiating two substitution principles or to having a linear and an intuitionistic cut rule (see [29, 42]).

LEMMA 2.6 (Intuitionistic transitivity).

- i.* If $\overline{\Psi} \vdash_{\Sigma}^p N \uparrow A$ and $\mathcal{P} :: (\Psi, x:A, \Psi' \vdash_{\Sigma}^p U \Downarrow V)$,
then $\Psi, [N/x]\Psi' \vdash_{\Sigma}^p [N/x]U \Downarrow [N/x]V$.
- ii.* If $\overline{\Psi} \vdash_{\Sigma}^p N \uparrow A$ and $\mathcal{P} :: (\vdash_{\Sigma}^p \Psi, x:A, \Psi' \uparrow Ctx)$, then $\vdash_{\Sigma}^p \Psi, [N/x]\Psi' \uparrow Ctx$.

Proof. By induction on the structure of \mathcal{P} . ■

We now state the linear transitivity lemma. Notice that, in contrast to the intuitionistic case, the context in the resulting judgment can be larger than the contexts mentioned in either premise.

LEMMA 2.7 (Linear transitivity).

- If $\Psi \vdash_{\Sigma}^p N \uparrow A$ and $\mathcal{P} :: (\overline{\Psi}, u:A, \Psi' \vdash_{\Sigma}^p M \Downarrow B)$, then $\Psi, \Psi' \vdash_{\Sigma}^p [N/u]M \Downarrow B$.

Proof. We proceed by induction on the structure of \mathcal{P} . Weakening is required in this proof. ■

The cumulative validity lemma below states that whenever a judgment is derivable, all entities mentioned in it are themselves valid, i.e., have derivations validating them.

LEMMA 2.8 (Consistency).

- i.* If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p M \Downarrow A$, then $\overline{\Psi} \vdash_{\Sigma}^p A \uparrow \text{TYPE}$.
- ii.* If $\mathcal{P} :: \overline{\Psi} \vdash_{\Sigma}^p A \Downarrow K$, then $\overline{\Psi} \vdash_{\Sigma}^p K \uparrow \text{Kind}$.
- iii.* If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p U \Downarrow V$, then $\vdash_{\Sigma}^p \Psi \uparrow Ctx$.
- iv.* If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p U \Downarrow V$ or $\mathcal{P} :: \vdash_{\Sigma}^p \Psi \uparrow Ctx$, then $\vdash^p \Sigma \uparrow \text{Sig}$.

Proof. By induction on the structure of \mathcal{P} . We need intuitionistic transitivity in order to handle the dependent function type constructor. ■

The traditional Curry-Howard interpretation associates types with formulas and terms with proofs of these formulas. Clearly, a single formula can have more than one proof, expressed in type theory by admitting several objects of the same type, possibly infinitely many. This is also consistent with the view of types as sets and terms as their elements. In the logical interpretation, we expect every proof to be the proof of a single formula. This property might not be desirable in all type theories, but it holds in the case of languages such as LF and LLF so that objects have meaning independent of the type ascribed to them. Uniqueness, in these frameworks, is considered modulo definitional equality. Indeed, every valid $\lambda^{\Pi \rightarrow \circ \& \top}$ object-level term has a unique type. This property, which extends naturally to

kinds, is essential in our proof of the decidability of type checking. It is formally stated in the following lemma.

LEMMA 2.9 (Uniqueness of types and kinds).

- i. If $\mathcal{P}' :: \Psi' \vdash_{\Sigma}^p M \Downarrow A'$ and $\mathcal{P}'' :: \Psi'' \vdash_{\Sigma}^p M \Downarrow A''$ with $\Psi'_{|\text{FV}(M)} = \Psi''_{|\text{FV}(M)}$ and where the arrows do not need to match, then $A' \equiv A''$.
- ii. If $\mathcal{P}' :: \bar{\Psi}' \vdash_{\Sigma}^p A \Downarrow K'$ and $\mathcal{P}'' :: \bar{\Psi}'' \vdash_{\Sigma}^p A \Downarrow K''$ where the arrows do not need to match, then $K' \equiv K''$.

Proof. By induction on the structure of \mathcal{P}' and \mathcal{P}'' . The idea is to examine these derivations from the bottom up until an introduction or an elimination rule is exposed; then we apply the induction hypothesis on the subderivations. ■

Given a particular instance of a judgment, the proof technique known as *inversion* allows identifying a limited number of inference rules whose conclusion matches this judgment. Each matching rule constitutes an alternative case and the judgments obtained by instantiation of its premises can be used in order to draw further inferences. In order to prove that the original judgment is derivable, it is sufficient to exhibit derivations for the premises of all the matching rules. This technique is general and can be applied in our system. Deductive systems having the characteristic that every rule of inference is fully determined by the shape of a particular term in its conclusion are called *syntax-directed*. They are particularly useful since matching this term yields a single rule. Therefore, further inferences can be drawn on the basis of its premises, without having to consider alternatives. The essential rules in the pre-canonical deductive system for $\lambda^{\Pi \rightarrow \circ \& \top}$ are syntax-directed with respect to the term they validate. However, the equivalences and the rules that bridge the pre-atomic and pre-canonical judgments do not change the derived term and can therefore be seen as filters or pipelines that connect these essential rules, from the standpoint of this term. A detailed analysis of these rules shows that we can indirectly recover the stronger form of inversion. This desirable property is expressed by the following lemma.

LEMMA 2.10 (Inversion).

- i. If $\mathcal{P} :: (\Psi \vdash_{\Sigma}^p \langle M, N \rangle \Downarrow A \& B)$, then $\Psi \vdash_{\Sigma}^p M \Downarrow A$ and $\Psi \vdash_{\Sigma}^p N \Downarrow B$.
- ii. If $\mathcal{P} :: (\Psi \vdash_{\Sigma}^p \hat{\lambda}u:A. M \Downarrow A' \multimap B)$, then $\Psi, u:A \vdash_{\Sigma}^p M \Downarrow B$ and $A \equiv A'$.
- iii. If $\mathcal{P} :: (\Psi \vdash_{\Sigma}^p \lambda x:A. M \Downarrow \Pi x:A'. B)$, then $\Psi, x:A \vdash_{\Sigma}^p M \Downarrow B$ and $A \equiv A'$.

Proof. By induction on the structure of \mathcal{P} . All these results apply the same technique: the derivation is unfolded until an introduction or an elimination appears as its last inference rule. ■

The last property we present in this section is extensionality. It confirms that all terms that are valid according to the pre-canonical judgment of the object level are indeed pre-canonical, i.e., in η -long form. It forbids, for example, a constant c of compound type A to be derivable by means of a judgment of the form $\Psi \vdash_{\Sigma}^p c \Downarrow A$, whatever the signature Σ and the context Ψ are.

LEMMA 2.11 (Extensionality).

- i. If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p M \uparrow \top$, then $M = \langle \rangle$.
- ii. If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p M \uparrow A \& B$, then $M = \langle M_1, M_2 \rangle$.
- iii. If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p M \uparrow A \multimap B$, then $M = \hat{\lambda}u:A'. M_1$ with $A \equiv A'$.
- iv. If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p M \uparrow \Pi x:A. B$, then $M = \lambda x:A'. M_1$ with $A \equiv A'$.

Proof. By induction on the structure of \mathcal{P} . ■

2.5. Strong Normalization

The aim of this subsection is to prove strong normalization for $\lambda^{\Pi \multimap \& \top}$. This property implies that every valid $\lambda^{\Pi \multimap \& \top}$ term U has a normal form $\text{NF}(U)$, that this normal form is unique, and that it can be obtained by performing β -reductions in arbitrary order in U . We adapt the technique originally proposed for LF in [27]. We only sketch it here. The interested reader is referred to [6] for details.

The proof proceeds via a translation of $\lambda^{\Pi \multimap \& \top}$ into the simply typed λ -calculus with pairs $\lambda^{\times \rightarrow}$. The effect of this encoding will be to eliminate dependencies and linearity, considerably simplifying the treatment of the calculus. This translation has two fundamental properties: first it maintains well-typedness, so that valid terms in $\lambda^{\Pi \multimap \& \top}$ are mapped to terms that are valid in $\lambda^{\times \rightarrow}$. Second, it preserves reductions so that every reduction sequence in $\lambda^{\Pi \multimap \& \top}$ corresponds to a reduction sequence in $\lambda^{\times \rightarrow}$. The strong normalization theorem for $\lambda^{\Pi \multimap \& \top}$ is then a consequence of the same property of $\lambda^{\times \rightarrow}$.

The first step towards proving the strong normalization theorem is given by the following lemma. It states that derivability is closed under reduction, i.e., if a term U is valid in $\lambda^{\Pi \multimap \& \top}$, then every term U' that differs from U only by the application of a β -reduction step is also valid. This property is known as *subject reduction*. We write $U \longrightarrow^+ V$ if $U \longrightarrow^* V$ and $U \neq V$. Notice that the symmetric property considering β -expansion does not hold in general.

LEMMA 2.12 (Subject reduction).

If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p U \uparrow \downarrow V$ and $\mathcal{R} :: U \longrightarrow^+ U'$, then $\Psi \vdash_{\Sigma}^p U' \uparrow \downarrow V$.

Proof. By induction on the structure of \mathcal{P} and inversion on \mathcal{R} . ■

We do not present in detail the simply typed λ -calculus with pairs, $\lambda^{\times \rightarrow}$ [51]. We overload some of the operators of $\lambda^{\Pi \multimap \& \top}$ to indicate analogous symbols of $\lambda^{\times \rightarrow}$. For our purposes, we will need a single base type, that we denote as ω . We base the equational theory of $\lambda^{\times \rightarrow}$ on the one-step reduction relation \longrightarrow_1 , which is more appropriate for our purposes than parallel nested reduction. We write \longrightarrow_1^+ for its transitive closure. We will rely on some basic properties for these judgments. We do not state or prove them formally since they resemble similar properties of LLF and are well-known from the literature. A further property that $\lambda^{\times \rightarrow}$ enjoys is strong normalization: if $\Gamma \vdash_{\Sigma} M : \sigma$, then every reduction sequence on M terminates and, by confluence, yields a unique normal form for this term. Proofs of this and stronger properties for extensions of this language can be found in the literature [19, 51].

The encoding we propose transforms LLF judgments $\Psi \vdash_{\Sigma}^p U \uparrow \downarrow V$ into $\lambda^{\times \rightarrow}$ judgments of the form $\Gamma \vdash_{\Sigma'} M : \sigma$. It maps the generic term U to an object M in

$\lambda^{\times\rightarrow}$, V to a simple type σ , Σ to a signature Σ' , and Ψ to a context Γ . We now present the four parts that constitute this translation.

We use the function $\tau(_)$ to denote the translation of a term that appears on the right-hand side of the arrow of an *LLF* judgment. These terms can be either types, kinds or the symbol *Kind*, which we map to ω . Given a type or kind U , $\tau(U)$ is a simple type of $\lambda^{\times\rightarrow}$ that maintains the structure of U , but forgets dependencies and linearity. Type families are mapped to the base type ω , the additive product type constructor $\&$ of *LLF* is encoded into the (intuitionistic) product types \times of $\lambda^{\times\rightarrow}$, and both function type constructors \multimap and Π of our language are represented by the unique arrow of that calculus. Kinds are treated similarly. Specifically, we have the following definition for $\tau(_)$:

<i>Types:</i>	<i>Kinds:</i>
$\tau(P) = \omega$	$\tau(\text{TYPE}) = \omega$
$\tau(\top) = \omega$	$\tau(\Pi x:A. K) = \tau(A) \rightarrow \tau(K)$
$\tau(A \& B) = \tau(A) \times \tau(B)$	
$\tau(A \multimap B) = \tau(A) \rightarrow \tau(B)$	
$\tau(\Pi x:A. B) = \tau(A) \rightarrow \tau(B)$	

As an example, the $\lambda^{\Pi\multimap\&\top}$ type

$$A = a \multimap ((\Pi x:a \multimap a. a) \& (a \multimap \Pi y:a. a))$$

has the following encoding in $\lambda^{\times\rightarrow}$:

$$\tau(A) = \omega \rightarrow (((\omega \rightarrow \omega) \rightarrow \omega) \times (\omega \rightarrow \omega \rightarrow \omega))$$

A term U appearing immediately to the left of the arrow of an *LLF* judgment $\Psi \vdash_{\Sigma}^p U \Downarrow V$ is mapped to a $\lambda^{\times\rightarrow}$ object by means of the function $|_$. U can be an object, a type, a type family or a kind.

The encoding of objects maps variables to variables, constants to constants, and constructors and destructors of $\lambda^{\Pi\multimap\&\top}$ to the corresponding operator of $\lambda^{\times\rightarrow}$. The two forms of λ -abstraction of *LLF* must be mapped to $\lambda^{\times\rightarrow}$ in a way which preserve the redices in the type label. We cope with this issue by encoding $\hat{\lambda}x:A. M$, for example, as $(\lambda y:\omega. \lambda x:\tau(A). |M|) |A|$. The expected translation of the former term is $\lambda x:\tau(A). |M|$. We embed it in the β -redex $(\lambda y:\omega. _) |A|$ in order to account for possible reductions performed in the $\lambda^{\Pi\multimap\&\top}$ type A . This redex is vacuous since y is a fresh variable not appearing in A or M .

The encoding of types and kinds translates each $\lambda^{\Pi\multimap\&\top}$ operator as a constant in $\lambda^{\times\rightarrow}$, which is applied to the encoding of the arguments. The dependent type and kind constructor Π requires a functional second argument since its semantics introduces assumptions in the context.

We have the following definition for this encoding function, where π , possibly subscripted, denotes constants in $\lambda^{\times\rightarrow}$:

Objects:

$$\begin{aligned}
|x| &= x \\
|c| &= \pi_c \\
|\langle \rangle| &= \pi_{\langle \rangle} \\
|\langle M, N \rangle| &= \langle |M|, |N| \rangle \\
|\text{FST } M| &= \text{FST } |M| \\
|\text{SND } M| &= \text{SND } |M| \\
|\hat{\lambda}u:A. M| &= (\lambda y:\omega. \lambda u:\tau(A). |M|) |A| \\
|M \hat{\wedge} N| &= |M| |N| \\
|\lambda x:A. M| &= (\lambda y:\omega. \lambda x:\tau(A). |M|) |A| \\
|M N| &= |M| |N|
\end{aligned}$$

Types and kinds:

$$\begin{aligned}
|a| &= \pi_a \\
|P N| &= |P| |N| \\
|\top| &= \pi_{\top} \\
|A \& B| &= \pi_{\&} |A| |B| \\
|A \multimap B| &= \pi_{\multimap} |A| |B| \\
|\Pi x:A. B| &= \pi_{\tau(A)} |A| (\lambda x:\tau(A). |B|) \\
|\text{TYPE}| &= \pi_{\text{TYPE}} \\
|\Pi x:A. K| &= \pi_{\tau(A)} |A| (\lambda x:\tau(A). |K|)
\end{aligned}$$

with y not appearing in either A or M . For example, the translation of the functional identity

$$M = \lambda x: (\Pi z: a. a). x$$

is

$$|M| = (\lambda y:\omega. \lambda x:\omega \rightarrow \omega. x)(\pi_{\omega} \pi_a (\lambda z:\omega. \pi_a)).$$

The encoding of contexts inductively extends $\tau(\cdot)$, eliminating the distinction between intuitionistic and linear assumptions. The encoding $\tau(\Sigma)$ of an *LLF* signature Σ consists of two parts: a variable part defining the type of every declaration $v:U$ in Σ as $\pi_v:\tau(U)$, and a fixed part declaring the type of the constants needed to represent the type and kind operators of our language. The treatment of Π yields an infinite family of declarations, one for each simple type in $\lambda^{\times \rightarrow}$. We have the following definitions:

Context:

$$\begin{aligned}
\tau(\cdot) &= \cdot \\
\tau(\Psi, u \hat{\vdash} A) &= \tau(\Psi), u:\tau(A) \\
\tau(\Psi, x:A) &= \tau(\Psi), x:\tau(A)
\end{aligned}$$

Signature:

$$\begin{aligned}
\tau(\Sigma) = \pi_c &: \tau(A), && \text{for } c:A \text{ in } \Sigma \\
\pi_a &: \tau(K), && \text{for } a:K \text{ in } \Sigma \\
\pi_{\sigma} &: \omega \rightarrow (\sigma \rightarrow \omega) \rightarrow \omega, \\
\pi_{\langle \rangle} &: \omega, \\
\pi_{\top} &: \omega, \\
\pi_{\&} &: \omega \rightarrow \omega \rightarrow \omega, \\
\pi_{\multimap} &: \omega \rightarrow \omega \rightarrow \omega, \\
\pi_{\text{TYPE}} &: \omega
\end{aligned}$$

The encoding just presented preserves well-typedness: Whenever a term U is valid in $\lambda^{\Pi \multimap \& \top}$, the object $|U|$ is valid in $\lambda^{\times \rightarrow}$. This property is formally stated in the following lemma.

LEMMA 2.13 (Adequacy of the translation).

- i. If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p M \uparrow\downarrow A$, then $\tau(\Psi) \vdash_{\tau(\Sigma)} |M| : \tau(A)$.
- ii. If $\mathcal{P} :: \bar{\Psi} \vdash_{\Sigma}^p A \uparrow\downarrow K$, then $\tau(\bar{\Psi}) \vdash_{\tau(\Sigma)} |A| : \tau(K)$.
- iii. If $\mathcal{P} :: \bar{\Psi} \vdash_{\Sigma}^p K \uparrow \text{Kind}$, then $\tau(\bar{\Psi}) \vdash_{\tau(\Sigma)} |K| : \omega$.

Proof. We proceed by induction on the structure of \mathcal{P} . ■

The proposed encoding has the further property of preserving reductions. Therefore, whenever a term U reduces to U' in $\lambda^{\Pi\text{-}\circ\&\top}$, the term $|U|$ reduces to $|U'|$ in $\lambda^{\times\rightarrow}$ in at least as many steps. The extra β -redex in the representation of λ -abstraction causes individual reductions in our language to be mapped to multi-step reductions in the target language, in general.

LEMMA 2.14 (Preservation of reduction sequences).

If $\mathcal{R} :: U \rightarrow^+ V$, then $|U| \rightarrow_1^+ |V|$.

Proof. By induction on the structure of \mathcal{R} . ■

We now have all the ingredients to prove the strong normalization theorem for $\lambda^{\Pi\text{-}\circ\&\top}$. A term U is *normalizing* if there exist a term U' in normal form such that $U \rightarrow^* U'$. U is *strongly normalizing* if every reduction sequence yields a normal term.

The strong normalization theorem states that every derivable term is strongly normalizing. This property holds in $\lambda^{\times\rightarrow}$, as proved for example in [19, 51], and we use this fact to prove that it is valid also for $\lambda^{\Pi\text{-}\circ\&\top}$.

THEOREM 2.2 (Strong normalization).

If $\Psi \vdash_{\Sigma}^p U \uparrow\downarrow V$, then U is strongly normalizing.

Proof. By the adequacy of the translation (Lemma 2.13), we have that $\tau(\Psi) \vdash_{\tau(\Sigma)}$ $|U| : \tau(V)$ is derivable in $\lambda^{\times\rightarrow}$.

Assume we have a (possibly infinite) reduction sequence in $\lambda^{\Pi\text{-}\circ\&\top}$ starting from U :

$$U = U_0 \rightarrow^+ U_1 \rightarrow^+ \dots$$

By reduction preservation (Lemma 2.14) there is a corresponding reduction sequence

$$|U_0| \rightarrow_1^+ |U_1| \rightarrow_1^+ \dots$$

in $\lambda^{\times\rightarrow}$. Since the latter must be finite, the former will also be finite. ■

The validity of strong normalization permits the derivation of a number of further properties for our language. A first result is that the normal form of a derivable term is unique, stated below.

COROLLARY 2.1 (Uniqueness of normal forms).

If $\Psi \vdash_{\Sigma}^p U \uparrow\downarrow V$, $U \rightarrow^* U'$ and $U \rightarrow^* U''$ with both U' and U'' in normal form, then $U' = U''$.

Proof. By confluence, there exist a term V such that $U' \rightarrow^* V$ and $U'' \rightarrow^* V$. However, since U' and U'' are normal, they do not contain β -redices, and therefore $U' = V = U''$. ■

This property allows us to define a function $\text{NF}(\cdot)$ in order to denote *the* normal form of a valid term U . $\text{NF}(U)$ is computed from U by applying β -reductions in this

term until a normal form is eventually reached. Strong normalization guarantees that this normal form arises after a finite number of steps, and the lemma above ensures that the resulting term is unique.

A further consequence of the strong normalization theorem is that the equational theory of LLF is decidable, i.e. it can be effectively decided whether there exists a derivation for the judgment $U \equiv U'$, for U and U' valid terms. The idea is to check that $\text{NF}(U)$ and $\text{NF}(U')$ are identical.

COROLLARY 2.2 (Decidability of the equational theory).

If $\Psi \vdash_{\Sigma}^p U' \Downarrow V$ and $\Psi \vdash_{\Sigma}^p U'' \Downarrow V$, then it can be recursively decided whether $U' \equiv U''$ is derivable.

Proof. By the Church-Rosser property, U' and U'' have a common reduct U . By subject reduction, U is valid. Therefore, by uniqueness (Corollary 2.1), U' , U and U'' share the same normal form $\text{NF}(U)$.

By the strong normalization theorem, every sequence of reduction on U' and U'' eventually produces $\text{NF}(U)$ after a finite number of steps. Therefore, a possible decision procedure for definitional equality is as follows: compute the normal forms of U' and U'' and then check whether they are syntactically equal. If they are, then U' is definitionally equal to U'' . Otherwise, they are not equivalent. ■

Yet another consequence of the strong normalization theorem is that every derivable $\lambda^{\Pi-\circ\&\top}$ judgment can be constrained to mention only objects in normal form. Although not strictly needed, it is a common practice to write LLF signatures in normal form.

COROLLARY 2.3 (Normal forms).

- i. If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p U \Downarrow V$, then $\text{NF}(\Psi) \vdash_{\text{NF}(\Sigma)}^p \text{NF}(U) \Downarrow \text{NF}(V)$.*
- ii. If $\mathcal{P} :: \vdash_{\Sigma}^p \Psi \Uparrow \text{Ctx}$, then $\vdash_{\text{NF}(\Sigma)}^p \text{NF}(\Psi) \Uparrow \text{Ctx}$.*
- iii. If $\mathcal{P} :: \vdash^p \Sigma \Uparrow \text{Sig}$, then $\vdash^p \text{NF}(\Sigma) \Uparrow \text{Sig}$.*

Proof. We first reduce U to normal form in (i) by means of the previous corollary, and then proceed by induction on the structure of \mathcal{P} . ■

In an implementation of the language, converting terms to normal form as soon as β -redices appear as the result of substitutions is not always necessary. It is usually more efficient to work with *weak head-normal forms*, which differ from normal forms by permitting redices in the argument of applications.

A final consequence of the strong normalization theorem is that rule **opa.c** can be dropped as soon as we are only interested in valid normal terms. As we briefly motivated earlier, only the presence of this rule permits the formation of β -redices in valid terms (i.e. in the terms immediately to the left of the arrow in a pre-canonical or pre-atomic judgment). Eliminating this rule is beneficial in order to use $\lambda^{\Pi-\circ\&\top}$ as the language of the logical framework LLF since it prevents non-normal terms from being validated without losing any valid normal term.

2.6. Algorithmic System

A proof of the decidability of type checking for LLF is difficult to achieve directly in the pre-canonical system in Figures 2–3. Indeed, it is not possible to predict the

Signatures	$\frac{}{\vdash^a \cdot \uparrow \text{Sig}} \text{sa_dot}$
	$\frac{\vdash^a \Sigma \uparrow \text{Sig} \quad \cdot \vdash^a A \uparrow \text{TYPE}}{\vdash^a \Sigma, c: A \uparrow \text{Sig}} \text{sa_obj} \qquad \frac{\vdash^a \Sigma \uparrow \text{Sig} \quad \cdot \vdash^a K \uparrow \text{Kind}}{\vdash^a \Sigma, a: K \uparrow \text{Sig}} \text{sa_fam}$
Contexts	$\frac{\vdash^a \Sigma \uparrow \text{Sig}}{\vdash^a \cdot \uparrow \text{Ctx}} \text{ca_dot}$
	$\frac{\vdash^a \Psi \uparrow \text{Ctx} \quad \bar{\Psi} \vdash^a A \uparrow \text{TYPE}}{\vdash^a \Psi, x: A \uparrow \text{Ctx}} \text{ca_int} \qquad \frac{\vdash^a \Psi \uparrow \text{Ctx} \quad \bar{\Psi} \vdash^a A \uparrow \text{TYPE}}{\vdash^a \Psi, u \hat{?} A \uparrow \text{Ctx}} \text{ca_lin}$
Kinds	$\frac{\vdash^a \bar{\Psi} \uparrow \text{Ctx}}{\bar{\Psi} \vdash^a \text{TYPE} \uparrow \text{Kind}} \text{kca_type} \qquad \frac{\bar{\Psi}, x: A \vdash^a K \uparrow \text{Kind}}{\bar{\Psi} \vdash^a \Pi x: A. K \uparrow \text{Kind}} \text{kca_dep}$
Types/type families	$\frac{\bar{\Psi} \vdash^a P \downarrow \text{TYPE}}{\bar{\Psi} \vdash^a P \uparrow \text{TYPE}} \text{fca_a}$
	$\frac{\vdash^a \bar{\Psi} \uparrow \text{Ctx}}{\bar{\Psi} \vdash^a \top \uparrow \text{TYPE}} \text{fca_top} \qquad \frac{\bar{\Psi} \vdash^a A \uparrow \text{TYPE} \quad \bar{\Psi} \vdash^a B \uparrow \text{TYPE}}{\bar{\Psi} \vdash^a A \& B \uparrow \text{TYPE}} \text{fca_with}$
	$\frac{\bar{\Psi} \vdash^a A \uparrow \text{TYPE} \quad \bar{\Psi} \vdash^a B \uparrow \text{TYPE}}{\bar{\Psi} \vdash^a A \multimap B \uparrow \text{TYPE}} \text{fca_limp} \qquad \frac{\bar{\Psi}, x: A \vdash^a B \uparrow \text{TYPE}}{\bar{\Psi} \vdash^a \Pi x: A. B \uparrow \text{TYPE}} \text{fca_dep}$
	<hr style="border-top: 1px dotted black;"/> $\frac{\vdash^a_{\Sigma, a: K, \Sigma'} \bar{\Psi} \uparrow \text{Ctx}}{\bar{\Psi} \vdash^a_{\Sigma, a: K, \Sigma'} a \downarrow \text{NF}(K)} \text{faa_con} \qquad \frac{\bar{\Psi} \vdash^a P \downarrow \Pi x: A. K \quad \bar{\Psi} \vdash^a N \uparrow A}{\bar{\Psi} \vdash^a P N \downarrow \text{NF}([N/x]K)} \text{faa_iapp}$

FIG. 5. Algorithmic Deduction System for $\lambda^{\Pi \multimap \& \top}$, Kinds and Types

size of a derivation for a judgment since the rules that we called “non-essential” in Section 2.1 (the equivalence rules and **opa.c**) can be chained arbitrarily. The strong normalization theorem and the admissibility of **opa.c** limit the need for these rules drastically. In this section, we will embed the first of these results in a deductive system for $\lambda^{\Pi \multimap \& \top}$ having the characteristic that every derivable judgment has a derivation whose size is bounded by a function of the terms constituting this judgment; we will come back to this aspect in Section 2.7. Following the terminology of [27], we call this system *algorithmic*. The properties of this system will also permit us to eventually prove the validity of strengthening for our language.

The algorithmic system for $\lambda^{\Pi \multimap \& \top}$ consists of judgments similar to the pre-canonical presentation; indeed, we use the same expressions, only annotating the turnstile symbol with the letter a instead of p . The notion of definitional equality

Objects	
	$\frac{\Psi \vdash_{\Sigma}^a M \downarrow P}{\Psi \vdash_{\Sigma}^a M \uparrow P} \text{ oca_a}$
$\frac{\vdash_{\Sigma}^a \Psi \uparrow Ctx}{\Psi \vdash_{\Sigma}^a \langle \rangle \uparrow \top} \text{ oca_unit}$	$\frac{\Psi \vdash_{\Sigma}^a M \uparrow A \quad \Psi \vdash_{\Sigma}^a N \uparrow B}{\Psi \vdash_{\Sigma}^a \langle M, N \rangle \uparrow A \& B} \text{ oca_pair}$
$\frac{\Psi, u \hat{\vdash} A \vdash_{\Sigma}^a M \uparrow B}{\Psi \vdash_{\Sigma}^a \hat{\lambda} u : A. M \uparrow \text{NF}(A) \multimap B} \text{ oca_llam}$	$\frac{\Psi, x : A \vdash_{\Sigma}^a M \uparrow B}{\Psi \vdash_{\Sigma}^a \lambda x : A. M \uparrow \Pi x : \text{NF}(A). B} \text{ oca_ilam}$

$\frac{\Psi \vdash_{\Sigma}^a M \uparrow A}{\Psi \vdash_{\Sigma}^a M \downarrow A} \text{ oaa_c}$	$\frac{\vdash_{\Sigma, c : A, \Sigma'}^a \bar{\Psi} \uparrow Ctx}{\bar{\Psi} \vdash_{\Sigma, c : A, \Sigma'}^a c \downarrow \text{NF}(A)} \text{ oaa_con}$
$\frac{\vdash_{\Sigma}^a \bar{\Psi}, u \hat{\vdash} A, \bar{\Psi}' \uparrow Ctx}{\bar{\Psi}, u \hat{\vdash} A, \bar{\Psi}' \vdash_{\Sigma}^a u \downarrow \text{NF}(A)} \text{ oaa_lvar}$	$\frac{\vdash_{\Sigma}^a \bar{\Psi}, x : A, \bar{\Psi}' \uparrow Ctx}{\bar{\Psi}, x : A, \bar{\Psi}' \vdash_{\Sigma}^a x \downarrow \text{NF}(A)} \text{ oaa_ivar}$
(No rule for \top)	$\frac{\Psi \vdash_{\Sigma}^a M \downarrow A \& B}{\Psi \vdash_{\Sigma}^a \text{FST } M \downarrow A} \text{ oaa_fst}$
	$\frac{\Psi \vdash_{\Sigma}^a M \downarrow A \& B}{\Psi \vdash_{\Sigma}^a \text{SND } M \downarrow B} \text{ oaa_snd}$
$\frac{\Psi' \vdash_{\Sigma}^a M \downarrow A \multimap B \quad \Psi'' \vdash_{\Sigma}^a N \uparrow A \quad \Psi = \Psi' \boxtimes \Psi''}{\Psi \vdash_{\Sigma}^a M \wedge N \downarrow B} \text{ oaa_lapp}$	
$\frac{\Psi \vdash_{\Sigma}^a M \downarrow \Pi x : A. B \quad \bar{\Psi} \vdash_{\Sigma}^a N \uparrow A}{\Psi \vdash_{\Sigma}^a M N \downarrow \text{NF}([N/x]B)} \text{ oaa_iapp}$	

FIG. 6. Algorithmic Deduction System for $\lambda^{\Pi \multimap \& \top}$, Objects

is the same, but we do not access the judgments defining the equational theory directly. We rely instead on the normalization function $\text{NF}(_)$, which is known to exist from the previous section. We remind the reader that this function is defined only for valid terms, and it will be easy to check that whenever it is used in the algorithmic system, its argument is valid.

The inference rules defining the behavior of the algorithmic system are displayed in Figures 5–6. This deductive system shares with the pre-canonical system of Section 2.2 the property that every derivable term is in η -long form. This aspect will be a consequence of the soundness theorem below. However, the algorithmic system has the further characteristic that all terms mentioned in any well-formed derivation are themselves valid. As we said, ill-formed terms could appear in equivalence subderivations by using β -expansions. In the algorithmic system, the equivalence relation \equiv has been eliminated in favor of applications of the normalization function in the rules introducing the dependent type or kind constructor Π (**faa.iapp** and **oaa.iapp**), which involve the application of a substitution. The algorithmic system has also the property that the terms appearing on the right of the arrow are always in canonical form. We achieve this effect by normalizing types and kinds when

fetching them from the signature or the context (rules **faa_con**, **oaa_con**, **oaa_lvar**, **oaa_ivar**, **oca_llam**, and **oca_ilam**).

The correspondence between the algorithmic and the pre-canonical systems is formalized by means of the following soundness and completeness theorems. First, every valid term in the algorithmic system is also valid in the pre-canonical formulation.

THEOREM 2.3 (Soundness of the algorithmic system).

- i.* If $\mathcal{A} :: \Psi \vdash_{\Sigma}^{\alpha} U \Downarrow V$, then $\Psi \vdash_{\Sigma}^p U \Downarrow V$ and V is in normal form.
- ii.* If $\mathcal{A} :: \vdash_{\Sigma}^{\alpha} \Psi \Uparrow Ctx$, then $\vdash_{\Sigma}^p \Psi \Uparrow Ctx$.
- iii.* If $\mathcal{A} :: \vdash^{\alpha} \Sigma \Uparrow Sig$, then $\vdash^p \Sigma \Uparrow Sig$.

Proof. We proceed by induction on the structure of \mathcal{A} . ■

The completeness theorem states that every judgment having a derivation in the pre-canonical system is also derivable in the system presented in this section. Therefore, no valid term is lost by moving to the algorithmic system. Notice however that the type or kind appearing on the right-hand side of the main judgments of the pre-canonical system must be normalized in the algorithmic system.

THEOREM 2.4 (Completeness of the algorithmic system).

- i.* If $\mathcal{P} :: \Psi \vdash_{\Sigma}^p U \Downarrow V$, then $\Psi \vdash_{\Sigma}^{\alpha} U \Downarrow \text{NF}(V)$.
- ii.* If $\mathcal{P} :: \vdash_{\Sigma}^p \Psi \Uparrow Ctx$, then $\vdash_{\Sigma}^{\alpha} \Psi \Uparrow Ctx$.
- iii.* If $\mathcal{P} :: \vdash^p \Sigma \Uparrow Sig$, then $\vdash^{\alpha} \Sigma \Uparrow Sig$.

Proof. By induction on the structure of \mathcal{P} . ■

The proofs of these results are constructive and therefore specify an effective transformation procedure.

The strict access to definitional equality, and in particular the impossibility of using it for β -expansions permits a direct proof of the strengthening lemma in the algorithmic system.

LEMMA 2.15 (Strengthening).

- i.* If $\Psi, x \hat{::} A, \Psi' \vdash_{\Sigma}^p U \Downarrow V$ and $x \notin \text{FV}(\Psi') \cup \text{FV}(U) \cup \text{FV}(V)$, then $\Psi, \Psi' \vdash_{\Sigma}^p U \Downarrow V$.
- ii.* If $\vdash_{\Sigma}^p \Psi, x \hat{::} A, \Psi' \Uparrow Ctx$ and $x \notin \text{FV}(\Psi')$, then $\vdash_{\Sigma}^p \Psi, \Psi' \Uparrow Ctx$.

Proof. We first prove the analogous lemma for the algorithmic system by induction on a derivation of the given derivations, and then use the above soundness and completeness results to transfer it to the pre-canonical setting. ■

2.7. Decidability

The absence of explicit equivalences in the algorithmic system limits the choices of the inference rules that can be used at every step of a derivation considerably. Every well-formed judgment matches the conclusion of at most one rule, with the

only exception of judgments of the form $\Psi \vdash_{\Sigma} M \downarrow A$, for which the coercion **oaa_c** from canonical terms is always available. Moreover, the terms appearing in the central part of a validity judgment become smaller when going from the conclusion to the premises in all rules in Figure 5–6 except **fca_a**, **oca_a** and **oaa_c**, for which they remain the same. Notice that possible cycles generated by the last two can be easily detected and removed. In practice, we restrict the **oaa_c** coercion to types A which are not base types P , thereby also avoiding cycles while retaining completeness.

In this section, we take advantage of these characteristics in order to prove the decidability in *LLF* of verifying whether a fully specified $\lambda^{\Pi \rightarrow \circ \& \top}$ judgment is derivable or not — *type checking* — and of computing a type or a kind for a judgment whose rightmost term is left unspecified, or declaring that no such term exists — *type synthesis*. Both issues need to be faced simultaneously in our language.

In order to achieve this goal, a preliminary step consists of defining a complexity measure for an algorithmic judgment. This number yields an upper bound for the size of at least one of its derivations. For this purpose, we rely on a family of *size functions* that we denote uniformly as $\|\cdot\|$. We designed these functions so that the size of the conclusion of every essential rule in the algorithmic system is strictly larger than the size of each of its premises.

The size of terms, types and kinds is defined in the upper part of Figure 7. The numerical constants in this definition ensure that a term has larger size than its subterms. Notice that the size expressions of constructs that bind variables rely on the constant 2 rather than 1. This measure ensures that the size of the conclusion of their introduction rule is larger than the size of the premise, which mentions an extended context.

This definition to contexts and signatures in the central part of Figure 7. We then combine the size of terms, contexts, and signature in order to define the size of all the judgments participating in the algorithmic system in the lower part of Figure 7. Notice that the size of a judgment does not refer to the term appearing to the right of the arrow. This is necessary for our purposes since the size of this term in the premises of the elimination rules can in general be larger than in the conclusion.

We designed the functions above so that the size of a judgment is an upper bound on the height of at least one of its derivations in the algorithmic system (which does not directly access the definitional equality judgments). This property is expressed by the following lemma. It is proved in two steps: first we eliminate all sequences of rules consisting only of the alternation of **oca_a** and **oaa_c**; second, we show that the size of the premises of an introduction or elimination rule is always smaller than the size of the conclusion.

LEMMA 2.16 (Upper bound on the size of a derivation).

- i. Let $h = \|\Psi \vdash_{\Sigma}^a U \Downarrow V\|$. If $\mathcal{A} :: \Psi \vdash_{\Sigma}^a U \Downarrow V$, then $\mathcal{A}' :: \Psi \vdash_{\Sigma}^a U \Downarrow V$.
- ii. Let $h = \|\vdash_{\Sigma}^a \Psi \uparrow Ctx\|$. If $\mathcal{A} :: \vdash_{\Sigma}^a \Psi \uparrow Ctx$, then $\mathcal{A}' :: \vdash_{\Sigma}^a \Psi \uparrow Ctx$.
- iii. Let $h = \|\vdash^a \Sigma \uparrow Sig\|$. If $\mathcal{A} :: \vdash^a \Sigma \uparrow Sig$, then $\mathcal{A}' :: \vdash^a \Sigma \uparrow Sig$.

In each case \mathcal{A}' has height less than $2h$ and contains at most 3^{2h} nodes. \square

<p style="text-align: center;"><i>Objects:</i></p> $\ x\ = 1$ $\ c\ = 1$ $\ \langle \rangle\ = 1$ $\ \langle M, N \rangle\ = 1 + \ M\ + \ N\ $ $\ \text{FST } M\ = 1 + \ M\ $ $\ \text{SND } M\ = 1 + \ M\ $ $\ \hat{\lambda}u : A. M\ = 2 + \ A\ + \ M\ $ $\ M \hat{\wedge} N\ = 1 + \ M\ + \ N\ $ $\ \lambda x : A. M\ = 2 + \ A\ + \ M\ $ $\ M N\ = 1 + \ M\ + \ N\ $	<p style="text-align: center;"><i>Types:</i></p> $\ a\ = 1$ $\ P N\ = 1 + \ P\ + \ N\ $ $\ \top\ = 1$ $\ A \& B\ = 1 + \ A\ + \ B\ $ $\ A \multimap B\ = 1 + \ A\ + \ B\ $ $\ \Pi x : A. B\ = 2 + \ A\ + \ B\ $ <hr style="border: 0.5px solid black;"/> <p style="text-align: center;"><i>Kinds:</i></p> $\ \text{TYPE}\ = 1$ $\ \Pi x : A. K\ = 2 + \ A\ + \ K\ $
<p style="text-align: center;"><i>Contexts:</i></p> $\ \cdot\ = 1$ $\ \Psi, u \hat{\wedge} A\ = 1 + \ \Psi\ + \ A\ $ $\ \Psi, x : A\ = 1 + \ \Psi\ + \ A\ $	<p style="text-align: center;"><i>Signatures:</i></p> $\ \cdot\ = 1$ $\ \Sigma, c : A\ = 1 + \ \Sigma\ + \ A\ $ $\ \Sigma, a : K\ = 1 + \ \Sigma\ + \ K\ $
<p><i>Judgments:</i></p> $\ \vdash^a \Sigma \uparrow \text{Sig}\ = \ \Sigma\ $ $\ \vdash_\Sigma^a \Psi \uparrow \text{Ctx}\ = \ \Sigma\ + \ \Psi\ $ $\ \bar{\Psi} \vdash_\Sigma^a K \uparrow \text{Kind}\ = \ \Sigma\ + \ \Psi\ + \ K\ $ $\ \bar{\Psi} \vdash_\Sigma^a A \uparrow \text{TYPE}\ = \ \Sigma\ + \ \Psi\ + \ A\ $ $\ \bar{\Psi} \vdash_\Sigma^a A \downarrow K\ = \ \Sigma\ + \ \Psi\ + \ A\ $ $\ \Psi \vdash_\Sigma^a M \uparrow A\ = \ \Sigma\ + \ \Psi\ + \ M\ $ $\ \Psi \vdash_\Sigma^a M \downarrow A\ = \ \Sigma\ + \ \Psi\ + \ M\ $ $\ \Psi = \Psi' \bowtie \Psi''\ = \ \Psi\ $	

FIG. 7. Size Computation for the Algorithmic System

We have now all the necessary ingredients to prove that the type checking problem is decidable in *LLF*. Given a judgment whose validity we want to decide, a first naive idea is to match it against the conclusion of the inference rules defining the algorithmic system. If none of these rules apply, then the judgment is not derivable, otherwise, we check recursively that the instantiated premises of the viable rules are derivable. The lemma above provides an upper bound on the number of rules that need to be considered.

Unfortunately, such a bound is not enough since the types in the premises of the elimination rules are larger than in the conclusion and would have to be guessed in a pure bottom-up strategy. However, they are determined by the signature and context using *type synthesis* [4]. Proving decidability of type synthesis requires type checking in order to validate contexts, so we need to prove these two properties simultaneously.

THEOREM 2.5 (Decidability of type checking).

- i. (Type checking)*
It can be recursively decided whether $\Psi \vdash_{\Sigma}^p U \Downarrow V$, $\vdash_{\Sigma}^p \Psi \Uparrow \text{Ctx}$ and $\vdash_{\Sigma}^p \Sigma \Uparrow \text{Sig}$ are derivable.
- ii. (Type synthesis)*
Given a signature Σ , a context Ψ and a term U , there is a recursive procedure that computes a term V such that the judgment $\Psi \vdash_{\Sigma}^p U \Downarrow V$ is derivable, or determines that no such V exists.

Proof. We prove the analogous property for the algorithmic system and rely the constructive aspects of the soundness and completeness theorems above to transfer it to the pre-canonical setting. The idea, in order to prove the algorithmic formulation of this result, is to apply inference rules that match U (and V for type checking) until either a derivation is produced, or no rule is applicable, or the upper bound on the size of the derivation at hand has been reached. ■

The mutually recursive parts of this theorem yield effective procedures for type checking and type synthesis. Once this result has been proved, type checking can be conveniently reduced to type synthesis: In order to check whether $\Psi \vdash_{\Sigma}^p M \Downarrow A$ is derivable, it suffices to check that A is valid, infer a type A' such that $\Psi \vdash_{\Sigma}^p M \Downarrow A'$ is derivable, and check whether $A \equiv A'$ holds, which, by Corollary 2.2, is decidable. An application of the equivalence rules yields $\Psi \vdash_{\Sigma}^p M \Downarrow A$. The subproblem of checking whether A is valid is also reduced to finding a kind for it, without indirections this time.

The decidability of type checking is a necessary property for using a formalism as a meta-representation language. Like *LF*, *LLF* encodes judgments from an object language as types and their derivations as object level terms. The decidability of type checking for $\lambda^{\Pi \rightarrow \circ \& \top}$ permits determining effectively whether a given object is the representation of a valid derivation from a (potentially linear) object formalism.

2.8. Canonical Forms

As in the pre-canonical case, the algorithmic system presented in Section 2.6 prevents terms which are not in η -long form from being validated. The only exception concerns the process of constructing a term U by means of the judgments $\Psi \vdash_{\Sigma}^{\circ} U \Downarrow V$, where U will not necessarily be η -expanded, in general. In this section, we will present a proof system that forces all entities appearing in a derivable judgment to be also in normal form. Therefore, all valid terms will be canonical, i.e. in η -long form and without β -redices. We will achieve this property by removing rule **oaa_c**, that, as we saw, permits the formation of β -redices in derivable terms. In this system, we will also customize the rules of the algorithmic system to make

Signatures	$\frac{}{\vdash \cdot \uparrow \text{Sig}} \text{sc_dot}$ $\frac{\vdash \Sigma \uparrow \text{Sig} \quad \cdot \vdash_{\Sigma} A \uparrow \text{TYPE}}{\vdash \Sigma, c : A \uparrow \text{Sig}} \text{sc_obj} \qquad \frac{\vdash \Sigma \uparrow \text{Sig} \quad \cdot \vdash_{\Sigma} K \uparrow \text{Kind}}{\vdash \Sigma, a : K \uparrow \text{Sig}} \text{sc_fam}$
Contexts	$\frac{}{\vdash_{\Sigma} \cdot \uparrow \text{Ctx}} \text{cc_dot}$ $\frac{\vdash_{\Sigma} \Psi \uparrow \text{Ctx} \quad \bar{\Psi} \vdash_{\Sigma} A \uparrow \text{TYPE}}{\vdash_{\Sigma} \Psi, x : A \uparrow \text{Ctx}} \text{cc_int} \qquad \frac{\vdash_{\Sigma} \Psi \uparrow \text{Ctx} \quad \bar{\Psi} \vdash_{\Sigma} A \uparrow \text{TYPE}}{\vdash_{\Sigma} \Psi, u \hat{?} A \uparrow \text{Ctx}} \text{cc_lin}$
Kinds	$\frac{}{\bar{\Psi} \vdash_{\Sigma} \text{TYPE} \uparrow \text{Kind}} \text{kc_type} \qquad \frac{\bar{\Psi} \vdash_{\Sigma} A \uparrow \text{TYPE} \quad \bar{\Psi}, x : A \vdash_{\Sigma} K \uparrow \text{Kind}}{\bar{\Psi} \vdash_{\Sigma} \Pi x : A. K \uparrow \text{Kind}} \text{kc_dep}$
Types/type families	$\frac{\bar{\Psi} \vdash_{\Sigma} P \downarrow \text{TYPE}}{\bar{\Psi} \vdash_{\Sigma} P \uparrow \text{TYPE}} \text{fc_a}$ $\frac{}{\bar{\Psi} \vdash_{\Sigma} \top \uparrow \text{TYPE}} \text{fc_top} \qquad \frac{\bar{\Psi} \vdash_{\Sigma} A \uparrow \text{TYPE} \quad \bar{\Psi} \vdash_{\Sigma} B \uparrow \text{TYPE}}{\bar{\Psi} \vdash_{\Sigma} A \& B \uparrow \text{TYPE}} \text{fc_with}$ $\frac{\bar{\Psi} \vdash_{\Sigma} A \uparrow \text{TYPE} \quad \bar{\Psi} \vdash_{\Sigma} B \uparrow \text{TYPE}}{\bar{\Psi} \vdash_{\Sigma} A \multimap B \uparrow \text{TYPE}} \text{fc_limp} \qquad \frac{\bar{\Psi} \vdash_{\Sigma} A \uparrow \text{TYPE} \quad \bar{\Psi}, x : A \vdash_{\Sigma} B \uparrow \text{TYPE}}{\bar{\Psi} \vdash_{\Sigma} \Pi x : A. B \uparrow \text{TYPE}} \text{fc_dep}$ <hr style="border-top: 1px dotted black;"/> $\frac{}{\bar{\Psi} \vdash_{\Sigma, a : K, \Sigma'} a \downarrow K} \text{fa_con} \qquad \frac{\bar{\Psi} \vdash_{\Sigma} P \downarrow \Pi x : A. K \quad \bar{\Psi} \vdash_{\Sigma} N \uparrow A}{\bar{\Psi} \vdash_{\Sigma} P N \downarrow \text{NF}([N/x]K)} \text{fa_iapp}$

FIG. 8. Canonical Deduction System for $\lambda^{\Pi \multimap \& \top}$, Kinds and Types

them more suited for automation. We will rely on this deductive system in the remainder of the paper.

When deriving a judgment of the form $\Psi \vdash_{\Sigma} U \uparrow \downarrow V$, the *canonical system* in Figures 8–9 presupposes the validity of both the signature Σ and the context Ψ . Consequently, the rules dealing with constants, variables, and other atomic terms do not need their premise anymore; they are the leaves of the derivations. When applying the rules for λ -abstraction or when checking that a dependent type is valid (rules **fc_dep** and **kc_dep**), we need to check that the type of the assumption added to the context is valid. This ensures the validity of the context in the premise of the rules. In rule **oa_lapp**, which relies on context splitting, the contexts occurring in the premises differ from the context in the conclusion only by the removal of some linear assumptions. These contexts are however valid since in our language, linear variables cannot occur free in types, and therefore no assumption can depend on them (see Lemma 2.3).

Objects		
	$\frac{\Psi \vdash_{\Sigma} M \downarrow P}{\Psi \vdash_{\Sigma} M \uparrow P} \text{oc_a}$	
$\frac{}{\Psi \vdash_{\Sigma} \langle \rangle \uparrow \top} \text{oc_unit}$	$\frac{\Psi \vdash_{\Sigma} M \uparrow A \quad \Psi \vdash_{\Sigma} N \uparrow B}{\Psi \vdash_{\Sigma} \langle M, N \rangle \uparrow A \& B} \text{oc_pair}$	
	$\frac{\bar{\Psi} \vdash_{\Sigma} A \uparrow \text{TYPE} \quad \Psi, u \hat{:} A \vdash_{\Sigma} M \uparrow B}{\Psi \vdash_{\Sigma} \hat{\lambda} u : A. M \uparrow A \multimap B} \text{oc_llam}$	
	$\frac{\bar{\Psi} \vdash_{\Sigma} A \uparrow \text{TYPE} \quad \Psi, x : A \vdash_{\Sigma} M \uparrow B}{\Psi \vdash_{\Sigma} \lambda x : A. M \uparrow \Pi x : A. B} \text{oc_ilam}$	
	$\frac{}{\bar{\Psi} \vdash_{\Sigma, c : A, \Sigma'} c \downarrow A} \text{oa_con}$	
$\frac{}{\bar{\Psi}, u \hat{:} A, \bar{\Psi}' \vdash_{\Sigma} u \downarrow A} \text{oa_lvar}$	$\frac{}{\bar{\Psi}, x : A, \bar{\Psi}' \vdash_{\Sigma} x \downarrow A} \text{oa_ivar}$	
(No rule for \top)	$\frac{\Psi \vdash_{\Sigma} M \downarrow A \& B}{\Psi \vdash_{\Sigma} \text{FST } M \downarrow A} \text{oa_fst}$	$\frac{\Psi \vdash_{\Sigma} M \downarrow A \& B}{\Psi \vdash_{\Sigma} \text{SND } M \downarrow B} \text{oa_snd}$
	$\frac{\Psi' \vdash_{\Sigma} M \downarrow A \multimap B \quad \Psi'' \vdash_{\Sigma} N \uparrow A \quad \Psi = \Psi' \boxtimes \Psi''}{\Psi \vdash_{\Sigma} M \wedge N \downarrow B} \text{oa_lapp}$	
	$\frac{\Psi \vdash_{\Sigma} M \downarrow \Pi x : A. B \quad \bar{\Psi} \vdash_{\Sigma} N \uparrow A}{\Psi \vdash_{\Sigma} M N \downarrow \text{NF}([N/x]B)} \text{oa_iapp}$	

FIG. 9. Canonical Deduction System for $\lambda^{\Pi \multimap \& \top}$, Objects

Another novelty of the canonical system is the absence of a rule corresponding **oa_c**, the coercion from pre-canonical to pre-atomic terms. We already saw that only through this rule can we show the validity of terms containing β -redices. Without it, the canonical system can only derive η -long terms that are β -normal, i.e., canonical, as the name of this system implies.

An important consequence of the elimination of **oa_c** is that the resulting system is syntax-directed: any judgment matches the conclusion of at most one inference rule.

The equivalence between the algorithmic and the canonical system in Figures 8–9 is expressed by means of the following soundness and completeness theorems.

THEOREM 2.6 (Soundness of the canonical system).

- i. If $\vdash \Sigma \uparrow \text{Sig}$, $\vdash_{\Sigma} \Psi \uparrow \text{Ctx}$ and $C :: \Psi \vdash_{\Sigma} U \Downarrow V$, then $\Psi \vdash_{\Sigma}^a U \Downarrow V$.
- ii. If $\vdash \Sigma \uparrow \text{Sig}$ and $C :: \vdash_{\Sigma} \Psi \uparrow \text{Ctx}$, then $\vdash_{\Sigma}^a \Psi \uparrow \text{Ctx}$.
- iii. If $C :: \vdash \Sigma \uparrow \text{Sig}$, then $\vdash^a \Sigma \uparrow \text{Sig}$.

Moreover, Σ , Ψ , U and V are in normal form.

Proof. By induction on the structure of \mathcal{C} . ■

The converse of this statement is not true in general. It holds, however, whenever all the entities appearing in a derivable algorithmic judgment are in normal form. We know by the normal form corollary 2.3 that every derivable judgment in that system is equivalent to a judgment that mentions only normal terms. Therefore, the correspondence between the algorithmic and the canonical system is not perfect, since it preserves derivability but not derivations, in general. This is however acceptable for our purposes since, when performing search and when encoding a deductive system, we are only interested in terms that are η -long and β -normal, i.e., canonical.

THEOREM 2.7 (Completeness of the canonical system).

- i. If $\Psi \vdash_{\Sigma}^a M \Downarrow A$, then
 - if A is a base type, then $\text{NF}(\Psi) \vdash_{\text{NF}(\Sigma)} \text{NF}(M) \Downarrow A$;
 - if $\text{NF}(M) = \langle \rangle$ or $\langle M', M'' \rangle$ or $\hat{\lambda}u:A'. M'$ or $\lambda x:A'. M'$, then $\text{NF}(\Psi) \vdash_{\text{NF}(\Sigma)} \text{NF}(M) \Uparrow A$;
 - if $\text{NF}(M) = x$ or c or $\text{FST } M'$ or $\text{SND } M'$ or $M' \wedge M''$ or $M' M''$, then $\text{NF}(\Psi) \vdash_{\text{NF}(\Sigma)} \text{NF}(M) \Downarrow A$.
- ii. If $\bar{\Psi} \vdash_{\Sigma}^a A \Downarrow K$, then $\text{NF}(\bar{\Psi}) \vdash_{\text{NF}(\Sigma)} \text{NF}(A) \Downarrow K$.
- iii. If $\bar{\Psi} \vdash_{\Sigma}^a K \Uparrow \text{Kind}$, then $\text{NF}(\bar{\Psi}) \vdash_{\text{NF}(\Sigma)} \text{NF}(K) \Uparrow \text{Kind}$.
- iv. If $\vdash_{\Sigma}^a \Psi \Uparrow \text{Ctx}$, then $\vdash_{\text{NF}(\Sigma)} \text{NF}(\Psi) \Uparrow \text{Ctx}$.
- v. If $\vdash_{\Sigma}^a \Sigma \Uparrow \text{Sig}$, then $\vdash \text{NF}(\Sigma) \Uparrow \text{Sig}$.

Proof. We proceed by induction on the given judgments after applying the Normal Forms Corollary 2.3 and taking into consideration the admissibility of rule **oaa_c**, which derives from the analogous property of rule **opa_c** in the pre-canonical system. ■

We will adopt the system presented in this section to compare the features of LLF to LF . We will rely on a canonical system adapted from [39] for our comparison. Details can be found in [6]. We will distinguish the λ^{Π} equivalents of the judgments presented earlier by annotating them with the superscript LF .

All syntactic entities of λ^{Π} are available in our language. This embedding is maintained at the level of judgments. The canonical system for $\lambda^{\Pi \rightarrow \circ \& \top}$ in Figures 8–9 differs from the corresponding λ^{Π} system only by the addition of rules that deal with the linear entities of our language. Therefore, every judgment derivable in λ^{Π} has an isomorphic derivation in $\lambda^{\Pi \rightarrow \circ \& \top}$.

THEOREM 2.8 (Extension over LF).

- i. If $\mathcal{LF} :: \Gamma \vdash_{\Sigma}^{LF} U \Downarrow V$, then $\Gamma \vdash_{\Sigma} U \Downarrow V$.
- ii. If $\mathcal{LF} :: \vdash_{\Sigma}^{LF} \Gamma \Uparrow \text{Ctx}$, then $\vdash_{\Sigma} \Gamma \Uparrow \text{Ctx}$.
- iii. If $\mathcal{LF} :: \vdash_{\Sigma}^{LF} \Sigma \Uparrow \text{Sig}$, then $\vdash \Sigma \Uparrow \text{Sig}$.

Proof. We proceed by induction on the structure of \mathcal{LF} . All cases are immediate as soon as we notice that, for a λ^{Π} context Γ , we have that $\bar{\Gamma} = \Gamma$. ■

	<i>Abstract syntax</i>	<i>Concrete syntax</i>	
Kinds	TYPE $\Pi x:A. K$	type $\{x:A\}K$	$A \multimap K$ $K \multimap A$
Types	$P M$ \top $A \& B$ $A \multimap B$ $\Pi x:A. B$	$P M$ $\langle \top \rangle$ $A \& B$ $A \multimap B$ $\{x:A\}B$	$B \multimap A$ $A \multimap B$ $B \multimap A$
Objects	$\langle \rangle$ $\langle M, N \rangle$ $\text{fst } M$ $\text{snd } M$ $\hat{\lambda}x:A. M$ $M \wedge N$ $\lambda x:A. M$ $M N$	$()$ M, N $\langle \text{fst} \rangle M$ $\langle \text{snd} \rangle M$ $[x \hat{\sim} A]M$ $M \wedge N$ $[x:A]M$ $M N$	

FIG. 10. Concrete Syntax for *LLF*

LLF has also the converse property of being *conservative* over *LF*, i.e., every derivable $\lambda^{\Pi \multimap \& \top}$ judgment that mentions only entities in the λ^{Π} fragment of the syntax has a corresponding derivation in λ^{Π} . This also entails that every judgment that is not derivable in *LF* remains such in *LLF*.

THEOREM 2.9 (Conservativity over *LF*).

Let Σ, Γ, U and V be an *LF* signature, an *LF* context and two *LF* terms, respectively, then

- i. If $\mathcal{C} :: \Gamma \vdash_{\Sigma} U \Downarrow V$, then $\Gamma \vdash_{\Sigma}^{\text{LF}} U \Downarrow V$.
- ii. If $\mathcal{C} :: \vdash_{\Sigma} \Gamma \uparrow \text{Ctx}$, then $\vdash_{\Sigma}^{\text{LF}} \Gamma \uparrow \text{Ctx}$.
- iii. If $\mathcal{C} :: \vdash \Sigma \uparrow \text{Sig}$, then $\vdash^{\text{LF}} \Sigma \uparrow \text{Sig}$.

Proof. We proceed by induction on the structure of \mathcal{C} . We need to remember that for an *LF* context Γ , we have that $\bar{\Gamma} = \Gamma$. ■

These properties have important consequences. Not only every judgment derivable in *LF* is derivable also in our language, but, more importantly, all the representation techniques, adequacy theorems, and examples developed for *LF* remain valid for *LLF*.

2.9. A Concrete Syntax for *LLF*

In this section, we extend the concrete syntax of *Elf* [41] to express the linear operators of *LLF*. In doing so, we want to fulfill two constraints: first of all, existing

<i>Precedence</i>	<i>Operator</i>	<i>Position</i>
<i>highest</i>	<fst> _ <snd> _	prefix
	_ - - ^ _	left associative
	_ & _	right associative
	_ -o _ - -> _	right associative
	_ o- _ - <- _	left associative
<i>lowest</i>	{_ : _}_ [_ : _]_ [_ ^ _]_	prefix
	_, -	right associative

FIG. 11. Concrete Syntax for *LLF*

Elf programs should not undergo any syntactic alteration (unless they declare some of the reserved identifiers that we will introduce) if we were to execute them in an implementation of *LLF* relying on the new syntax. In other words, the extension we propose should be conservative with respect to the syntax of *Elf*. Second, we want to avoid a proliferation of operators: keeping their number as small as possible will make future extensions easier to accommodate if their inclusion appears beneficial.

The set of special characters of *Elf* consists of % : .) (] [} { . We extend these with two symbols: , and ^ . $\lambda^{\Pi \rightarrow \circ \& \top}$ object and type family constants are consequently represented as identifiers consisting of any non-empty string that does not contain spaces or the characters % : .) (] [} { , ^ . As in *Elf*, identifiers must be separated from each other by whitespace (i.e., blanks, tabs, and new lines) or special characters. We augment the set of reserved identifiers of *Elf* (`type`, `->` and `<-`) with `<T>`, `&`, `-o`, `o-`, `<fst>` and `<snd>`. Although not properly an identifier, the symbol `()` is also reserved; this string is forbidden in *Elf*.

Figure 10 associates every $\lambda^{\Pi \rightarrow \circ \& \top}$ operator to its concrete representation. Terms in the λ^{Π} sublanguage of *LLF* are mapped to the syntax of *Elf*. This language offers the convenience of writing `->` as `<-` with the arguments reversed in order to give a more operational reading to a program, when desired: under this perspective, we read the expression `A <- B` as “*A if B*”. We extend this possibility to linear implication, `-o`. Clearly, when we use `o-`, the arguments should be swapped: `A o- B` is syntactic sugar for `B -o A`.

The table in Figure 11 gives the relative precedence and associativity of these operators. As in *Elf*, parentheses are available to override these behaviors.

As in *Elf*, a signature declaration `c : A` is represented by the program clause:

$$c : A.$$

Type family constants are declared similarly. For practical purposes, it is convenient to provide a means of declaring linear assumptions. Indeed, whenever the object formalism we want to represent requires numerous linear hypotheses, it is simpler to write them as program clauses than to rely on some initialization routine that assumes them in the context during its execution. To this end, we permit

declarations of the form

$$c \sim A.$$

with the intent that this declaration should be inserted in the context as a linear assumption.

We retain from *Elf* the use of % for comments and interpreter directives. We adopt the conventions available in that language in order to enhance the readability of *LLF* programs [38]. In particular, we permit keeping the type of bound variables implicit whenever they can be effectively reconstructed by means techniques akin to those currently implemented in *Elf* [38]. We write $\{x\}B$, $[x]B$ and $[x^\sim]B$ when maintaining implicit the type A of the variable x in $\{x:A\}B$, $[x:A]B$ and $[x^\sim A]B$, respectively. Similar conventions apply to dependent kinds. As in *Elf*, the binders for variables quantified at the head of a clause can be omitted altogether if we write these variables with identifiers starting with a capital letter. Moreover, the arguments instantiating them can be kept implicit when using these declarations.

Finally, we relax the requirement of writing *LLF* declarations only in η -long form. With sufficient typing information it is always possible to transform a signature to that format.

3. THE METHODOLOGY OF LINEAR META-REPRESENTATION

LF and *Elf* constitute a useful tool for studying existing logics and programming languages, and an ideal playground for experimenting with alternative constructs in the design phase of new languages. The range of practical applicability of these formalisms is limited by their foundation on intuitionistic type theory. All the formal systems that have been successfully encoded in *LF* (functional and logic programming languages [33, 39], λ -calculi [40], and a number of logics [27, 43]) share a fundamental characteristic: whenever a judgment mentions a context, a bottom-up reading of the inference rules for it may add items, but it never removes assumptions. We call contexts with this property *permanent*, in contrast with *volatile* contexts, free from this restriction. Object formalisms admitting arbitrary operations on their context cannot be effectively encoded in *LF*: the standard technique, representing object context items as *LF* assumptions, is not sound in this case since *LF* assumptions are permanent. The alternative is to represent the object context as a term in *LF* and implement explicitly the operations required to access and manipulate it. This is undesirable since it makes the adequacy results difficult to prove and often complicates the encoding of meta-theoretic properties to the point of making it hardly manageable in practice.

This situation is quite unfortunate since most formalisms of practical significance rely on a volatile context in an essential manner. The languages used for programming commercial applications are imperative: they have a store and assignment instructions to change the value of variables. Most real-world problems carry a state that changes with time. Many new logics and type theories are inherently bound to destructive context manipulations. Permanent contexts are insufficient even for more traditional formalisms, for example when studying efficient proof-search procedures for intuitionistic logic [17].

The linear type theory $\lambda^{\top \& -\circ \Pi}$ presented in the previous section retains all the desirable properties of LF and also augments this formalism with linear assumptions, admitting volatile manipulations, and with a suitable set of operators to manage them. These new features overcome the above deficiency of LF : if we represent the volatile context of an object language as linear assumptions in $\lambda^{\top \& -\circ \Pi}$, destructive context operations in the object formalism can be modeled by an appropriate combination of linear operators.

The linear logical framework LLF is founded on the type theory $\lambda^{\top \& -\circ \Pi}$ and combines as its meta-representation methodology the *judgments-as-types* technique of LF with the above observation. The present section illustrates the added expressiveness of LLF as a logical framework by describing the meta-representation methodology it adopts, first abstractly and then on a concrete case study. The formalism we want to represent is an imperative extension of *Mini-ML* [25, 33, 39], a purely functional restriction of the programming language ML [23, 36]. More precisely, we augment that language with a store and imperative instructions to access and modify the values it contains, we formalize the typing and evaluation semantics of these constructs and we show that this extended language enjoys the type preservation property. We call this language MLR , for *Mini-ML with References*. The linear assumptions of LLF can be used to encode individual memory cells and the linear operators of our type theory offer effective tools to model manipulations on them.

We review the judgments-as-types representation methodology and extend it to handle volatile assumptions in Section 3.1. Then, we give a detailed but informal presentation of the syntax, semantics and of the type preservation property for MLR in Section 3.2. Finally, we show how to encode these different aspects in LLF in Section 3.7. Appendix A contains the complete LLF signature for this example. In the following, we will concentrate mainly on the novel constructions available in MLR , referring the reader to the literature [13, 25, 33, 39] for aspects already present in *Mini-ML*.

3.1. Judgments-as-Types Revisited

We will review the technique of *judgments-as-types* of LF [27] by analyzing the following simplified rule of inference from the case study in this section:

$$\frac{\Gamma, x:\tau \vdash^e e : \tau}{\Gamma \vdash^e \mathbf{fix} \ x.e : \tau} \mathbf{tpe_fix}$$

Ignoring for the moment the context Γ , it specifies that the fix-point expression $\mathbf{fix} \ x.e$ has type τ if e has type τ assuming that the variable x has also type τ . We will emphasize the fact that x can occur in e by writing $e(x)$. Given a closed expression $\mathbf{fix} \ x.e(x)$, the judgment in the conclusion of **tpe_fix** postulates that $\mathbf{fix} \ x.e(x)$ has type τ (we need to provide a derivation to ascertain that this is indeed the case). We call such a judgment *simple*. The judgments-as-types representation methodology encodes simple judgments as λ^{Π} base types. In Section 3.7, we will use the type family constants EXP and TP , both of kind TYPE to classify the expressions and the types of the object language, respectively. The general form of the typing judgment above relates an expression and a type, and therefore we encode it as a type family TPE , of kind $\text{EXP} \rightarrow \text{TP} \rightarrow \text{TYPE}$. Given

representations $(\text{FIX } \lambda x:\text{EXP}. \ulcorner e \urcorner x)$ and $\ulcorner \tau \urcorner$ (to be explained below) for the closed expression $\mathbf{fix } x.e(x)$ and for the object language type τ , the simple judgment $\Gamma \vdash^e \mathbf{fix } x.e(x) : \tau$ is represented as

$$\text{TPE } (\text{FIX } (\lambda x:\text{EXP}. \ulcorner e \urcorner x)) \ulcorner \tau \urcorner.$$

The judgment in the premise of rule **tpe_fix** is different in nature. Indeed, it specifies that the expression $e(x)$ has type τ *if* we assume that the variable x has also type τ . A judgment of this form is called *hypothetical*. Notice also that x is a bound variable in $\mathbf{fix } x.e(x)$, but it is free in $e(x)$. Therefore, that premise expresses the fact that $e(x)$ has type τ for a generic expression x of type τ . The judgment $\Gamma, x:\tau \vdash^e e(x) : \tau$ is therefore said to be also *parametric* in x . The judgments-as-types representation methodology encodes hypothetical and parametric judgments by means of simple and dependent function types respectively. The premise of the rule above, which is parametric in x and hypothetical in $x:\tau$, is represented as follows:

$$\Pi x:\text{EXP}. \text{TPE } x \ulcorner \tau \urcorner \rightarrow \text{TPE } (\ulcorner e \urcorner x) \ulcorner \tau \urcorner$$

Notice that instantiating the parameter x with some term e' yields a hypothetical judgment postulating that $e(e')$ has type τ assuming that e' has type τ . This reduces to a simple judgment as soon as we provide a derivation for this hypothesis.

An attempt at finding a canonical *LF* derivation with the above type reduces to searching for a derivation for the base type $\text{TPE } (\ulcorner e \urcorner x) \ulcorner \tau \urcorner$ after having added the assumptions $x:\text{EXP}$ and $t_x:\text{TPE } x \ulcorner \tau \urcorner$ to the context of *LF*. Viewing this as an alternate encoding for the premise of rule **tpe_fix** illustrates the manner an object context is encoded according to the judgments-as-types methodology: each item in the context of the object formalism is represented as one or more assumptions in the context of *LF*. This technique offers the further advantage that we can rely on the primitive operations of *LF* to simulate the lookup of object level assumptions. Less sophisticated representations, for example those that encode the object context as a term, must provide explicit access operations.

Observe that rule **tpe_fix** can be read as a judgment that is parametric in the (functional) expression e and the type τ , and hypothetical in the derivability of its premise. Indeed, it is encoded as the following declaration:

$$\begin{aligned} \text{TPE_FIX} : \Pi e:\text{EXP} \rightarrow \text{EXP}. \Pi \tau:\text{TP}. \\ (\Pi x:\text{EXP}. \text{TPE } x \tau \rightarrow \text{TPE } (e x) \tau) \\ \rightarrow \text{TPE } (\text{FIX } (\lambda x:\text{EXP}. e x)) \tau \end{aligned}$$

or, taking advantage of the concrete syntax of *Elf* (see Section 2.9),

```
tpe_fix :  ({x:exp} tpe x T -> tpe (E x) T)
          -> tpe (fix ([x:exp] E x)) T.
```

In summary, the judgments-as-types representation methodology for *LF* encodes simple judgments as base types, hypothetical and parametric judgments as simple and dependent function types respectively, and element of the object context as items in the context of *LF*. Moreover, derivations for a simple judgments are naturally represented as terms of the corresponding base type.

The judgments-as-types methodology interacts particularly well with *higher-order abstract syntax*, a technique for the representation of the syntactic level of an object formalism that encodes object variables as meta-variables and relies on the λ -abstraction of λ^Π to emulate generic object-level binding constructs. Above, we encoded the fix-point expression $\mathbf{fix} \ x. e(x)$, that binds the variable x in $e(x)$ as $(\mathbf{FIX} (\lambda x : \mathbf{EXP}. \ulcorner e \urcorner x))$. We used the λ -abstraction of LF to express binding, and consequently encoded the operator \mathbf{fix} by means of the LF constant \mathbf{fix} that accepts a functional operator ($\mathbf{fix} : (\mathbf{exp} \rightarrow \mathbf{exp}) \rightarrow \mathbf{exp}$).

The faithfulness of the representation of an object formalism is captured by means of *adequacy theorems* that relate the entities being represented to their encoding. An important advantage of the judgments-as-types technique with respect to less sophisticated approaches is that it produces encodings very close to the notations being formalized. This makes the adequacy theorems easy to prove.

Here, and in the remainder of this paper, we view and describe operations on the context as they arise when we construct derivations “bottom-up”, that is, from the judgment in question towards the axioms. This view is the most natural one to elucidate the examples and anticipates the logic programming interpretation of LLF . For example, instead of saying that we *discharge* a hypothesis in rule **opc_ila**m in Figure 2 we say that we *introduce* a hypothesis. From this point of view, λ^Π offers two operations on its context: insertion and lookup. In particular, the context can only grow during the bottom-up construction of a derivation. Therefore, the judgments-as-types methodology in λ^Π cannot capture object languages that perform deletion on their context. Consider as an example the following inference rule, taken from the case study in the next section:

$$\frac{(S, c = v) \triangleright K \vdash \mathbf{return} \langle \rangle \hookrightarrow a}{(S, c = v') \triangleright K \vdash c :=_2^* v \hookrightarrow a} \mathbf{ev_assign}_2^*$$

This rule describes the semantics of assignment in an imperative programming language (further details will be given in the next section). It specifies that, in order to assign the value v to the cell c , we must update the binding $c = v'$ in the store with $c = v$; some uninteresting value is returned. An elegant encoding of this system in LF would represent each cell-value pair in the store as a meta-level assumption. However, λ^Π does not provide means to simulate the deletion of the old binding, $c = v'$.

In contrast, we can easily achieve this effect in LLF . Indeed, looking up a *linear* assumption in $\lambda^{\Pi \multimap \& \top}$ removes it from the context. This suggests encoding each cell-value pair $c = v$ present at any instant in the store of the object language as an LLF linear assumption $\mathbf{Cn} \ ? \ \mathbf{contains} \ c \ \ulcorner v \urcorner$.

The linear type constructors of $\lambda^{\Pi \multimap \& \top}$ provide the necessary means to manipulate such assumptions. We rely on \multimap to enter them in the context of LLF and take advantage of the context splitting semantics of this operator to isolate them in order to access them. The additive product type constructor, $\&$, offers means to duplicate or share linear assumptions among its two conjuncts. This operator can also be used to express selection between exclusive alternatives, although we will not take advantage of this feature here. Finally, the unit type, \top , permits discarding unused linear hypotheses.

These different features will be illustrated in detail in Section 3.7. We just show the encoding of the rule above:

```
ev_assign*2 : (contains C V -o ev K (return unit) A)
              -o (contains C V' -o ev K (assign*2 (rf C) V) A).
```

The linearity of our logical framework can be integrated into higher-order abstract syntax as a convenient manner of encoding languages relying on linear binders [6]. When they are not needed we can just use the *LF* fragment of *LLF* exactly as before.

3.2. *Mini-ML* with References

Critical choices in the implementation of programming languages depend on the validity of meta-theoretic properties. Type preservation in *Standard ML* [23, 36], for example, guarantees that no typing error can arise during evaluation; therefore execution can be sped up significantly by disregarding type information at run-time. Meta-theoretic properties in the presence of non-functional features, included in most concrete languages, are difficult to prove and therefore the formal analysis of imperative extensions of purely functional programming languages has received great attention in the literature. The addition of references and their interaction with polymorphism has been analyzed with different tools, ranging from the complex domain-theoretic approach of Damas [15] to the syntactic formulation of Harper [26]. The latter idea was adapted from Wright and Felleisen, who additionally consider continuations and exceptions [54].

The proofs of these properties are long and error-prone. Therefore, recent work has investigated the possibility of partially automating their generation or at least their verification. Chirimar gives *Forum* specifications for a language with references, exceptions, and continuations and uses the meta-theory of *Forum* [34] to study program equivalence [12]. VanInwegen [52] formally proves properties such as value soundness (the fact that evaluating an expression yields a value, if it terminates) for most of *Standard ML* with the help of the *HOL* theorem prover [24].

In this section, we define *MLR* as an extension of *Mini-ML* with references and imperative instructions, and study aspects of its meta-theory. Although our principal objective is to demonstrate the expressive power of *LLF*, our presentation differs in some aspects from the formulations and proofs in the literature and therefore might be interesting in itself. We will point out differences and similarities with other approaches as they arise.

Expressions and Store

Since its introduction in [13], the language *Mini-ML* and variants of it have been used for case studies in the presentation of logical frameworks [25, 33, 39]. *Mini-ML* is a purely functional restriction of the programming language *ML* [23, 36]. More specifically, it is a small statically typed functional programming language including numerals, conditional expressions, pairs, polymorphic definitions, recursion, and functional expressions.

We consider an extension of *Mini-ML* with a store and imperative instructions in the style of *ML* to access and modify the values it contains. We call this language *Mini-ML with References*, or *MLR* for short. The *store* of an *MLR* program is

defined as a collection of *cells* each containing a value. We will sometimes use *location* or *address* as synonyms of cell. *MLR* makes available all the constructs of *Mini-ML* but enriches the syntax of its expressions with the necessary operations to manipulate individual cells. The resulting language is specified by the following grammar, where we have separated out the constructs not present in standard presentations of *Mini-ML* with a double bar (\parallel). Cells c and stores S are not directly accessible to the programmer, but it is customary and convenient to enrich the syntax in order to represent intermediate stages during computation.

<i>Expressions:</i>	$e ::= x$	(<i>Variables</i>)
	$ \mathbf{z} \mid \mathbf{s} e$	(<i>Natural numbers</i>)
	$ \mathbf{case} e \mathbf{of} \mathbf{z} \Rightarrow e_1 \mid \mathbf{s} x \Rightarrow e_2$	(<i>Conditionals</i>)
	$ \langle \rangle$	(<i>Unit element</i>)
	$ \langle e_1, e_2 \rangle \mid \mathbf{fst} e \mid \mathbf{snd} e$	(<i>Pairs</i>)
	$ \mathbf{lam} x. e \mid e_1 e_2$	(<i>Functions</i>)
	$ \mathbf{letval} x = e_1 \mathbf{in} e_2 \mid$	(<i>Definitions</i>)
	$\mathbf{letname} x = e_1 \mathbf{in} e_2$	
	$ \mathbf{fix} x. e$	(<i>Recursion</i>)
	$\parallel c \mid \mathbf{ref} e \mid !e$	(<i>References</i>)
	$ e_1 := e_2 \mid e_1; e_2$	(<i>Commands</i>)
<i>Stores:</i>	$S ::= \cdot \mid S, c = v$	

In these productions, c ranges over the lexical category of memory locations, while we use the letter x for variables. The meta-variable v denotes values, that we will define shortly. We will treat stores as multisets, omit the leading “,” from a non-empty store, and overload “,” to denote the union of two stores. Finally, we require the cells appearing on the left-hand side of a store item to be distinct.

The polymorphism in *MLR* is restricted to values, which is generally accepted as superior to the imperative type variables present in previous versions of *SML* [31]. We achieve this by distinguishing two forms of **let**. The expression **ref** e dynamically allocates a cell and initializes it with the value of e . The contents of a cell can be inspected by dereferencing it with **!** and modified with an assignment ($:=$). Differently from [54], but consistently with the main stream in the literature (including the definition of *Standard ML* [36]), we choose this operation not to return the assigned object, but the unit element $\langle \rangle$. The sequencing operator ($;$) is typically used as a means of chaining a series of assignments with some interesting final value; it is syntactic sugar for the expression (**letval** $x = e_1$ **in** e_2) when x does not occur in e_2 . As is normally the case in functional languages, *MLR* does not offer explicit means to deallocate memory cells.

All these constructs are available in *Standard ML* [36] with the exception of addresses themselves (c), which cannot be manipulated directly in that language. We require *MLR* programs not to mention locations directly so that cells are always guaranteed to be initialized. Thus cells are created dynamically with **ref** and can be named by binding them to variables with one of the two **let** constructs of *MLR*.

As in *ML*, the reference cells of *MLR* encompass two distinct features of imperative programming languages such as *C* or *Pascal*. First of all, they play the role of the imperative variables of these languages and can be used as such (except for the necessity of dereferencing them explicitly in order to access their value). Second,

we can use them as pointers in data structures, although their usefulness is rather limited in this respect due to the absence of recursive data types in *MLR*. Such data structures could be easily added to the language.

Typing

The language of types of *MLR* augments the typing constructs typically present in *Mini-ML*, namely natural numbers, unit, pairing, and functional types, with one new constructor: for each type τ , the type τ **ref** for references to objects of type τ . The syntax of types is summarized in the following grammar:

$$\begin{aligned} \text{Types: } \tau ::= & \alpha \mid \mathbf{nat} \mid \mathbf{1} \mid \tau_1 \times \tau_2 \mid \tau_1 \rightarrow \tau_2 \\ & \parallel \tau \mathbf{ref} \end{aligned}$$

We use type variables to express schematic polymorphism. We eliminate an explicit quantifier in favor of substitution in the typing rule for the **letname** construct (see Figure 12). On the basis of this definition, the static semantics of *MLR* naturally extends the traditional typing rules of *Mini-ML*. The possibility of expressions to mention cells requires introducing a *store context* as a means to declare the type of free locations. More precisely, the item $c:\tau$ in a store context declares τ as the type of the values that c can contain; c itself has consequently type τ **ref**. Contexts, as usual, assign types to free variables. They are constructed according to the following grammar:

$$\begin{aligned} \text{Contexts: } \Gamma ::= & \cdot \mid \Gamma, x:\tau \\ \text{Store contexts: } \Omega ::= & \cdot \mid \Omega, c:\tau \end{aligned}$$

We rely on the usual convention that the names of the variables and the cells declared in stores and context stores, respectively, are distinct. Moreover, we treat both forms of contexts as multisets.

We express the fact that the *MLR* expression e has type τ with respect to a store context Ω and a context Γ with the judgment

$$\Omega; \Gamma \vdash^e e : \tau.$$

The presence of a store context in the typing rules for *MLR* is necessary even if we forbid the users to write addresses directly in their programs. It accounts for cells dynamically allocated during evaluation, which may appear in intermediate results and in the final answer.

The inference rules for the typing judgment are displayed in Figure 12. The upper part of this figure shows the rules for the functional core of *MLR*. The changes with respect to the usual rules for *Mini-ML* are limited to the systematic inclusion of a store context in the judgments.

The central part of Figure 12 shows the rules for the novel features of *MLR*. As for the functional case, they express the conditions under which an expression can be statically accepted as meaningful. For example, rule **tpe_deref** enforces that only references be dereferenced.

In the lower part of Figure 12, we present the rules for typing a store. The judgments we consider have the form

$$\Omega' \vdash^S S : \Omega$$

Expressions	
$\frac{}{\Omega; \Gamma, x: \tau \vdash^e x : \tau} \text{tpe_x}$	$\frac{}{\Omega; \Gamma \vdash^e \langle \rangle : \mathbf{1}} \text{tpe_unit}$
$\frac{}{\Omega; \Gamma \vdash^e \mathbf{z} : \mathbf{nat}} \text{tpe_z}$	$\frac{\Omega; \Gamma \vdash^e e : \mathbf{nat}}{\Omega; \Gamma \vdash^e \mathbf{s} e : \mathbf{nat}} \text{tpe_s}$
$\frac{\Omega; \Gamma \vdash^e e : \mathbf{nat} \quad \Omega; \Gamma \vdash^e e_1 : \tau \quad \Omega; \Gamma, x: \mathbf{nat} \vdash^e e_2 : \tau}{\Omega; \Gamma \vdash^e \mathbf{case} e \text{ of } \mathbf{z} \Rightarrow e_1 \mid \mathbf{s} x \Rightarrow e_2 : \tau} \text{tpe_case}$	
$\frac{\Omega; \Gamma \vdash^e e_1 : \tau_1 \quad \Omega; \Gamma \vdash^e e_2 : \tau_2}{\Omega; \Gamma \vdash^e \langle e_1, e_2 \rangle : \tau_1 \times \tau_2} \text{tpe_pair}$	
$\frac{\Omega; \Gamma \vdash^e e : \tau_1 \times \tau_2}{\Omega; \Gamma \vdash^e \mathbf{fst} e : \tau_1} \text{tpe_fst}$	$\frac{\Omega; \Gamma \vdash^e e : \tau_1 \times \tau_2}{\Omega; \Gamma \vdash^e \mathbf{snd} e : \tau_2} \text{tpe_snd}$
$\frac{\Omega; \Gamma, x: \tau_1 \vdash^e e : \tau_2}{\Omega; \Gamma \vdash^e \mathbf{lam} x. e : \tau_1 \rightarrow \tau_2} \text{tpe_lam}$	$\frac{\Omega; \Gamma \vdash^e e_1 : \tau_2 \rightarrow \tau_1 \quad \Omega; \Gamma \vdash^e e_2 : \tau_2}{\Omega; \Gamma \vdash^e e_1 e_2 : \tau_1} \text{tpe_app}$
$\frac{\Omega; \Gamma \vdash^e e_1 : \tau_1 \quad \Omega; \Gamma, x: \tau_1 \vdash^e e_2 : \tau_2}{\Omega; \Gamma \vdash^e \mathbf{letval} x = e_1 \text{ in } e_2 : \tau_2} \text{tpe_letval}$	
$\frac{\Omega; \Gamma \vdash^e [e_1/x]e_2 : \tau}{\Omega; \Gamma \vdash^e \mathbf{letname} x = e_1 \text{ in } e_2 : \tau} \text{tpe_letname}$	
$\frac{\Omega; \Gamma, x: \tau \vdash^e e : \tau}{\Omega; \Gamma \vdash^e \mathbf{fix} x. e : \tau} \text{tpe_fix}$	
$\frac{}{\Omega, c: \tau; \Gamma \vdash^e c : \tau \mathbf{ref}} \text{tpe_cell}$	
$\frac{\Omega; \Gamma \vdash^e e : \tau}{\Omega; \Gamma \vdash^e \mathbf{ref} e : \tau \mathbf{ref}} \text{tpe_ref}$	$\frac{\Omega; \Gamma \vdash^e e : \tau \mathbf{ref}}{\Omega; \Gamma \vdash^e !e : \tau} \text{tpe_deref}$
$\frac{\Omega; \Gamma \vdash^e e_1 : \tau_1 \quad \Omega; \Gamma \vdash^e e_2 : \tau_2}{\Omega; \Gamma \vdash^e e_1; e_2 : \tau_2} \text{tpe_seq}$	
$\frac{\Omega; \Gamma \vdash^e e_1 : \tau \mathbf{ref} \quad \Omega; \Gamma \vdash^e e_2 : \tau}{\Omega; \Gamma \vdash^e e_1 := e_2 : \mathbf{1}} \text{tpe_assign}$	
Store	
$\frac{}{\Omega \vdash^S \cdot : \cdot} \text{tpS_empty}$	$\frac{\Omega \vdash^S S : \Omega' \quad \Omega; \cdot \vdash^e v : \tau}{\Omega \vdash^S (S, c = v) : (\Omega', c: \tau)} \text{tpS_cell}$

FIG. 12. Typing Rules in *MLR*, Expressions and Store

that we interpret as requiring that the type of each value v stored in S coincides with the type of the corresponding cell as specified in Ω . The store context Ω' gives the type of the cells v may mention. We will always be interested in top-level judgments of the form $\Omega \vdash^S S : \Omega$ since a store will in general refer circularly to its

own cells. Rule **tpS.cell** prevents expressions containing free variables from being inserted in the store.

Evaluation

An *MLR* expression e will in general mention reference cells whose values are contained in the store. The evaluation of e will typically not only retrieve these values, but also change them or create new cells. Therefore, as e is evaluated, the store will undergo transformations, and by the time a value for e is eventually produced, it might appear very different from the store we started with. This observation suggests an evaluation judgment of the form

$$S; e \hookrightarrow S'; v$$

where S is the store prior to evaluating e , and S' results from the evaluation of e to v : cells in e refer to S while cells in v refer to S' . This formulation extends the traditional evaluation judgment for *Mini-ML* [25, 33, 39].

The dynamic semantics of functional languages enriched with imperative features, such as *MLR*'s references, is normally expressed in the literature in this manner. We will instead adopt a different strategy and present the reductions occurring during the execution of an *MLR* program as continuation-based evaluation rules. This choice has been dictated by our intention to encode the semantics of *MLR* in *LLF*. A direct representation of the judgment above, although possible, would have resulted in a less elegant encoding. For similar reasons, Chirimar [12] also chose a continuation-based formulation.

Differently from more declarative formulations, a continuation-based execution strategy imposes a strict order of evaluation on the different subexpressions of any given construct in the language. This order respects the expected flow of data and is therefore natural. For example, when computing the value of an expression of the form (**letval** $x = e_1$ **in** e_2) we will first evaluate e_1 , obtain a value v' , substitute it for x in e_2 and only then evaluate the resulting expression.

An effective implementation of this strategy requires sequentializing the evaluation of the subexpressions of constructs with more than one argument. One of them is evaluated immediately while the evaluation of the others is postponed until a value has been produced for it. Clearly, if a subexpression depends on the value of another, we process it last. We realize this idea by maintaining a stack of expressions to be evaluated, called a *continuation*.

Postponing the evaluation of an expression e_2 in favor of another expression e_1 is achieved by pushing the former into the continuation. Since, as when evaluating (**letval** $x = e_1$ **in** e_2) for example, the value of e_1 might need to be substituted for some free variable x in e_2 , we wrap a binder for x around e_2 and thus insert an object of the form $\lambda x. e_2$ into the continuation (or compose it with the current continuation, depending on whether the continuation is viewed as a stack of functions, or as a single function corresponding to their composition). For uniformity, it is convenient to take this measure every time we insert an item into the stack. As soon as e_1 has been fully evaluated to a value v , $\lambda x. e_2$ is extracted from the continuation, v is substituted for the variable x in e_2 , and $[v/x]e_2$ is evaluated in turn.

The necessity of distinguishing expressions still to be evaluated from values being returned requires the introduction of the new syntactic layer of *instructions*. Specifically, we write **eval** e for the request to evaluate an expression e and denote the intention to return a value v as **return** v . Instructions are needed also for the purpose of handling partially evaluated expressions.

While evaluating a *Mini-ML* expression simply yielded a value, *MLR* expressions will in general produce objects mentioning cells. Therefore the result of the evaluation of an instruction i must include not only a final value v but also a reification $[S']$ of the final store S' it draws its references from; moreover, as a measure of hygiene, we mark the cells c that have been introduced during the evaluation process by binding them in front of the pair $([S'], v)$ by means of the **new** $c.$ operator. The resulting object is called an *answer* and is indicated with the letter a . For our purposes, $[S']$ will be a sequence obtained by ordering the elements of S' according to some arbitrary order. It is however conceivable that only the cells that contribute to the final value be kept, realizing in this way a form of garbage collection.

The structure of instructions, continuations and answers is given by the following grammar, where we have indicated with the double bar the instructions introduced in correspondence to the imperative constructs of *MLR*.

$$\begin{aligned}
\text{Instructions: } \quad i &::= \mathbf{eval} \ e \mid \mathbf{return} \ v \\
&\quad \mid \mathbf{case}^* \ v \ \mathbf{of} \ \mathbf{z} \Rightarrow e_1 \mid \mathbf{s} \ x \Rightarrow e_2 \\
&\quad \mid \langle v, e \rangle^* \mid \mathbf{fst}^* \ v \mid \mathbf{snd}^* \ v \\
&\quad \mid \mathbf{app}^* \ v \ e \\
&\quad \parallel \mathbf{ref}^* \ v \mid \mathbf{deref}^* \ v \mid v :=_1^* \ e \mid v_1 :=_2^* \ v_2 \\
\text{Continuations: } \quad K &::= \mathbf{init} \mid K, \lambda x. i \\
\text{Answers: } \quad a &::= ([S], v) \mid \mathbf{new} \ c. a
\end{aligned}$$

The typing rules for objects in these three categories are displayed in Figure 13. Notice that the type of an answer coincides with the type of the embedded value. Rule **tpa_val** requires that the store it is paired with be well-typed, while rule **tpa_new** constrains every occurrence of the cells bound in an answer to be consistently typed.

Values constitute the subclass of expressions that evaluate to themselves. They are specified by the following grammar.

$$\begin{aligned}
\text{Values: } \quad v &::= x \mid \mathbf{z} \mid \mathbf{s} \ v \mid \langle \rangle \mid \langle v_1, v_2 \rangle \mid \mathbf{lam} \ x. e \\
&\quad \parallel c
\end{aligned}$$

On the basis of this definition, we can justify the uses of the term “value” in the above presentation. Not only does **return** operate only on values, but computation places a value at the heart of answers and the contents of every cell in the store is a value. See [6] for a formal statement of these properties.

We model the continuation-based semantics of the imperative constructs of *MLR* by means of a judgment of the form

$$S \triangleright K \vdash i \leftrightarrow a$$

where i is the instruction to be executed, K is the current continuation, S is the store with respect to which i is to be evaluated and a is the final answer produced as the result of the evaluation.

Instructions	
$\frac{\Omega; \Gamma \vdash^e e : \tau}{\Omega; \Gamma \vdash^i \mathbf{eval} e : \tau} \mathbf{tpi_eval}$	$\frac{\Omega; \Gamma \vdash^e v : \tau}{\Omega; \Gamma \vdash^i \mathbf{return} v : \tau} \mathbf{tpi_return}$
$\frac{\Omega; \Gamma \vdash^e v : \mathbf{nat} \quad \Omega; \Gamma \vdash^e e_1 : \tau \quad \Omega; \Gamma, x : \mathbf{nat} \vdash^e e_2 : \tau}{\Omega; \Gamma \vdash^i \mathbf{case}^* v \mathbf{of} \mathbf{z} \Rightarrow e_1 \mid \mathbf{s} x \Rightarrow e_2 : \tau} \mathbf{tpi_case}^*$	
$\frac{\Omega; \Gamma \vdash^e v : \tau_1 \quad \Omega; \Gamma \vdash^e e : \tau_2}{\Omega; \Gamma \vdash^i \langle v, e \rangle^* : \tau_1 \times \tau_2} \mathbf{tpi_pair}^*$	
$\frac{\Omega; \Gamma \vdash^e v : \tau_1 \times \tau_2}{\Omega; \Gamma \vdash^i \mathbf{fst}^* v : \tau_1} \mathbf{tpi_fst}^*$	$\frac{\Omega; \Gamma \vdash^e v : \tau_1 \times \tau_2}{\Omega; \Gamma \vdash^i \mathbf{snd}^* v : \tau_2} \mathbf{tpi_snd}^*$
$\frac{\Omega; \Gamma \vdash^e v : \tau_2 \rightarrow \tau_1 \quad \Omega; \Gamma \vdash^e e : \tau_2}{\Omega; \Gamma \vdash^i \mathbf{app}^* v e : \tau_1} \mathbf{tpi_app}^*$	
.....	
$\frac{\Omega; \Gamma \vdash^e v : \tau}{\Omega; \Gamma \vdash^i \mathbf{ref}^* v : \tau \mathbf{ref}} \mathbf{tpi_ref}^*$	$\frac{\Omega; \Gamma \vdash^e v : \tau \mathbf{ref}}{\Omega; \Gamma \vdash^i \mathbf{deref}^* v : \tau} \mathbf{tpi_deref}^*$
$\frac{\Omega; \Gamma \vdash^e v : \tau \mathbf{ref} \quad \Omega; \Gamma \vdash^e e : \tau}{\Omega; \Gamma \vdash^i v :=_1^* e : \mathbf{1}} \mathbf{tpi_assign1}^*$	
$\frac{\Omega; \Gamma \vdash^e v_1 : \tau \mathbf{ref} \quad \Omega; \Gamma \vdash^e v_2 : \tau}{\Omega; \Gamma \vdash^i v_1 :=_2^* v_2 : \mathbf{1}} \mathbf{tpi_assign2}^*$	
Continuations	
$\frac{}{\Omega \vdash^K \mathbf{init} : \tau \Rightarrow \tau} \mathbf{tpK_init}$	$\frac{\Omega; x : \tau_1 \vdash^i i : \tau \quad \Omega \vdash^K K : \tau \Rightarrow \tau_2}{\Omega \vdash^K K, \lambda x. i : \tau_1 \Rightarrow \tau_2} \mathbf{tpK_lam}$
Answers	
$\frac{\Omega \vdash^S S : \Omega \quad \Omega; \cdot \vdash^e v : \tau}{\Omega \vdash^a ([S], v) : \tau} \mathbf{tpa_val}$	$\frac{\Omega, c : \tau' \vdash^a a : \tau}{\Omega \vdash^a \mathbf{new} c. a : \tau} \mathbf{tpa_new}$

FIG. 13. Typing Rules in *MLR*, Instructions, Continuations and Answers

The inference rules for evaluation are given in Figures 14–15. The evaluation of most instructions in the functional core of *MLR* does not access the store. The only exception is rule **ev_init** which must package the current store together with the produced value in order to construct the final answer. More generally, this operation could also include garbage collection, but we do not pursue this possibility here.

The inference rules concerned with non-functional expressions of *MLR* and the corresponding instructions are separated out by a dotted line in Figures 14 and 15, respectively.

Cells (rule **ev_cell**) simply evaluate to themselves, like any value. The sequencing instruction $e_1; e_2$ has a simple semantics too: it evaluates e_1 , disregards the returned value, and then proceed with the evaluation of e_2 (rule **ev_seq**).

Expressions	
(No ev_x)	$\frac{S \triangleright K \vdash \mathbf{return\ } z \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } z \hookrightarrow a} \mathbf{ev_z} \quad \frac{S \triangleright K, \lambda x. \mathbf{return\ } s\ x \vdash \mathbf{eval\ } e \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } s\ e \hookrightarrow a} \mathbf{ev_s}$
	$\frac{S \triangleright K, \lambda y. \mathbf{case}^* y \mathbf{ of\ } z \Rightarrow e_1 \mid s\ x \Rightarrow e_2 \vdash \mathbf{eval\ } e \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ case\ } e \mathbf{ of\ } z \Rightarrow e_1 \mid s\ x \Rightarrow e_2 \hookrightarrow a} \mathbf{ev_case}$
	$\frac{S \triangleright K \vdash \mathbf{return\ } \langle \rangle \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } \langle \rangle \hookrightarrow a} \mathbf{ev_unit} \quad \frac{S \triangleright K, \lambda x. \langle x, e_2 \rangle^* \vdash \mathbf{eval\ } e_1 \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } \langle e_1, e_2 \rangle \hookrightarrow a} \mathbf{ev_pair}$
	$\frac{S \triangleright K, \lambda x. \mathbf{fst}^* x \vdash \mathbf{eval\ } e \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ fst\ } e \hookrightarrow a} \mathbf{ev_fst} \quad \frac{S \triangleright K, \lambda x. \mathbf{snd}^* x \vdash \mathbf{eval\ } e \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ snd\ } e \hookrightarrow a} \mathbf{ev_snd}$
	$\frac{S \triangleright K \vdash \mathbf{return\ lam\ } x. e \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ lam\ } x. e \hookrightarrow a} \mathbf{ev_lam} \quad \frac{S \triangleright K, \lambda x. \mathbf{app}^* x\ e_2 \vdash \mathbf{eval\ } e_1 \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } e_1\ e_2 \hookrightarrow a} \mathbf{ev_app}$
	$\frac{S \triangleright K, \lambda x. \mathbf{eval\ } e_2 \vdash \mathbf{eval\ } e_1 \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ letval\ } x = e_1 \mathbf{ in\ } e_2 \hookrightarrow a} \mathbf{ev_letval}$
	$\frac{S \triangleright K \vdash \mathbf{eval\ } [e_1/x]e_2 \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ letname\ } x = e_1 \mathbf{ in\ } e_2 \hookrightarrow a} \mathbf{ev_letname}$
	$\frac{S \triangleright K \vdash \mathbf{eval\ } [\mathbf{fix\ } x. e/x]e \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ fix\ } x. e \hookrightarrow a} \mathbf{ev_fix}$
.....	
	$\frac{S \triangleright K \vdash \mathbf{return\ } c \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } c \hookrightarrow a} \mathbf{ev_cell}$
	$\frac{S \triangleright K, \lambda x. \mathbf{ref}^* x \vdash \mathbf{eval\ } e \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ ref\ } e \hookrightarrow a} \mathbf{ev_ref} \quad \frac{S \triangleright K, \lambda x. \mathbf{deref}^* x \vdash \mathbf{eval\ } e \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } !e \hookrightarrow a} \mathbf{ev_deref}$
	$\frac{S \triangleright K, \lambda x. \mathbf{eval\ } e_2 \vdash \mathbf{eval\ } e_1 \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } e_1; e_2 \hookrightarrow a} \mathbf{ev_seq} \quad \frac{S \triangleright K, \lambda x. x :=_1^* e_2 \vdash \mathbf{eval\ } e_1 \hookrightarrow a}{S \triangleright K \vdash \mathbf{eval\ } e_1 := e_2 \hookrightarrow a} \mathbf{ev_assign}$

FIG. 14. Evaluation in *MLR*, Expressions

The evaluation of $\mathbf{ref}\ e$ computes the value v of e (rule $\mathbf{ev_ref}$), allocates a new cell c in the store, initializes it with v and finally returns c itself (rule $\mathbf{ev_ref}^*$). Notice that rule $\mathbf{ev_ref}^*$ has also the effect of binding c in the final answer by means of $\mathbf{new\ } c. \dots$. The argument part of $!e$ is evaluated to a reference cell (rule $\mathbf{ev_deref}$) and the value associated to it is returned (rule $\mathbf{ev_deref}^*$). We rely on the auxiliary *read judgment* $S \vdash c = v$ in order to retrieve the value of a cell (rule $\mathbf{read_val}$). The evaluation of $e_1 := e_2$ first evaluates e_1 to a store location c (rule $\mathbf{ev_assign}$), computes the value v of e_2 (rule $\mathbf{ev_assign}_1^*$), and replaces the former contents of c with v (rule $\mathbf{ev_assign}_2^*$). The returned value is $\langle \rangle$.

We conclude our discussion about evaluation with a few words about the interaction of references and polymorphism. The question is subtle and has received great attention in the literature [50, 26, 31]. Consider for example the following *MLR*

Values
$\frac{}{S \triangleright \mathbf{init} \vdash \mathbf{return} v \hookrightarrow ([S], v)} \text{ev_init} \qquad \frac{S \triangleright K \vdash [v/x]i \hookrightarrow a}{S \triangleright K, \lambda x. i \vdash \mathbf{return} v \hookrightarrow a} \text{ev_cont}$
Auxiliary instructions
$\frac{S \triangleright K \vdash \mathbf{eval} e_1 \hookrightarrow a}{S \triangleright K \vdash \mathbf{case}^* \mathbf{z} \mathbf{of} \mathbf{z} \Rightarrow e_1 \mid \mathbf{s} x \Rightarrow e_2 \hookrightarrow a} \text{ev_case}_1^*$
$\frac{S \triangleright K \vdash \mathbf{eval} [v/x]e_2 \hookrightarrow a}{S \triangleright K \vdash \mathbf{case}^* \mathbf{s} v \mathbf{of} \mathbf{z} \Rightarrow e_1 \mid \mathbf{s} x \Rightarrow e_2 \hookrightarrow a} \text{ev_case}_2^*$
$\frac{S \triangleright K, \lambda x. \mathbf{return} \langle v, x \rangle \vdash \mathbf{eval} e \hookrightarrow a}{S \triangleright K \vdash \langle v, e \rangle^* \hookrightarrow a} \text{ev_pair}^*$
$\frac{S \triangleright K \vdash \mathbf{return} v_1 \hookrightarrow a}{S \triangleright K \vdash \mathbf{fst}^* \langle v_1, v_2 \rangle \hookrightarrow a} \text{ev_fst}^* \qquad \frac{S \triangleright K \vdash \mathbf{return} v_2 \hookrightarrow a}{S \triangleright K \vdash \mathbf{snd}^* \langle v_1, v_2 \rangle \hookrightarrow a} \text{ev_snd}^*$
$\frac{S \triangleright K, \lambda x. \mathbf{eval} e_1 \vdash \mathbf{eval} e_2 \hookrightarrow a}{S \triangleright K \vdash \mathbf{app}^* (\mathbf{lam} x. e_1) e_2 \hookrightarrow a} \text{ev_app}^*$
<hr style="border-top: 1px dotted black;"/>
$\frac{(S, c = v) \triangleright K \vdash \mathbf{return} c \hookrightarrow a}{S \triangleright K \vdash \mathbf{ref}^* v \hookrightarrow \mathbf{new} c. a} \text{ev_ref}^* \qquad \frac{S \vdash c = v \quad S \triangleright K \vdash \mathbf{return} v \hookrightarrow a}{S \triangleright K \vdash \mathbf{deref}^* c \hookrightarrow a} \text{ev_deref}^*$
$\frac{S \triangleright K, \lambda x. c :=_2^* x \vdash \mathbf{eval} e \hookrightarrow a}{S \triangleright K \vdash c :=_1^* e \hookrightarrow a} \text{ev_assign}_1^*$
$\frac{(S, c = v) \triangleright K \vdash \mathbf{return} \langle \rangle \hookrightarrow a}{(S, c = v') \triangleright K \vdash c :=_2^* v \hookrightarrow a} \text{ev_assign}_2^*$
Read
$\frac{}{(S, c = v) \vdash c = v} \text{read_val}$

FIG. 15. Evaluation in *MLR*, Values and Auxiliary Instructions

expression:

$$\begin{aligned} & \mathbf{letname} f = \mathbf{ref} (\mathbf{lam} x. x) \\ & \mathbf{in} f := \mathbf{lam} x. \mathbf{s} x; \\ & !f \langle \rangle \end{aligned}$$

At first sight, this expression allocates a cell and initializes it with the identity function, which has polymorphic type $\alpha \rightarrow \alpha$. In the body of **letname**, we first update it to the successor function, of type **nat** \rightarrow **nat**, and then apply it to $\langle \rangle$, of type **1**. Clearly, something is wrong, but the typing rules of *MLR* accept the program above as a correct expression of type **1**. Is there a flaw in the definition of the static semantics of our language? Fortunately, no. A closer analysis reveals that, since the evaluation of **letname** substitutes **ref** (**lam** $x. x$) for every occurrences of

f in its body, the expression above reduces to:

$$\begin{aligned} & \mathbf{ref} (\mathbf{lam} \ x. x) := \mathbf{lam} \ x. \mathbf{s} \ x; \\ & (!\mathbf{ref} (\mathbf{lam} \ x. x)) \langle \rangle \end{aligned}$$

Each occurrence of $\mathbf{ref} (\mathbf{lam} \ x. x)$ evaluates to a different cell that is typed according to its use. The expression above would not be typable if we had used \mathbf{letval} in place of $\mathbf{letname}$.

Languages with explicit type variables solve the same problem by distinguishing between *applicative* and *imperative* type variables in order to avoid problems such as the above [23, 36, 26]. Restricting polymorphism to values has also been proposed as a solution to this problem [50] and has been adopted in the current definition of *Standard ML* [36]. This language offers only one form of \mathbf{let} , but it takes different courses of action depending on whether it defines a value or an arbitrary expression. Our treatment is slightly more general since it makes the call-by-name semantics of $\mathbf{letname}$ directly available: for example, the above expression does not type-check in *SML*.

Type Preservation

We conclude this section with the statement of the type preservation theorem for *MLR* and of the lemmas it depends on. For reasons of space, we will not formalize the proof of these results in *LLF*. The interested reader can find an encoding of this proof in our linear logical framework in [6].

The type preservation theorem states that the type of an expression does not change as the result of evaluation. The proof of the type preservation theorem relies on a number of auxiliary lemmas. The first is *weakening*: whenever an expression is well-typed in a given context and store, it remains well-typed under further assumptions and additional cells. This is easily proved by induction on typing derivations.

LEMMA 3.1 (Weakening).

- i. If $\mathcal{T} :: \Omega; \Gamma \vdash^e e : \tau$, then $\Omega, \Omega'; \Gamma, \Gamma' \vdash^e e : \tau$.
- ii. If $\mathcal{T} :: \Omega; \Gamma \vdash^i i : \tau$, then $\Omega, \Omega'; \Gamma, \Gamma' \vdash^i i : \tau$.
- iii. If $\mathcal{T} :: \Omega \vdash^K K : \tau_1 \Rightarrow \tau_2$, then $\Omega, \Omega' \vdash^K K : \tau_1 \Rightarrow \tau_2$.
- iv. If $\mathcal{T} :: \Omega \vdash^S S : \bar{\Omega}$, then $\Omega, \Omega' \vdash^S S : \bar{\Omega}$.
- v. If $\mathcal{T} :: \Omega \vdash^a a : \tau$, then $\Omega, \Omega' \vdash^a a : \tau$.

Proof. We proceed by induction on the structure of \mathcal{T} . The parts of this lemma should be proved in the order they are presented. ■

The second auxiliary property we need is the *substitution lemma*: it states that free variables in a well-typed expression can be substituted for expressions of the same type and the result will be well-typed.

LEMMA 3.2 (Substitution).

- i. If $\mathcal{T} :: \Omega; \Gamma, x : \tau' \vdash^e e : \tau$ and $\Omega; \Gamma \vdash^e e' : \tau'$, then $\Omega; \Gamma \vdash^e [e'/x]e : \tau$.
- ii. If $\mathcal{T} :: \Omega; \Gamma, x : \tau' \vdash^i i : \tau$ and $\Omega; \Gamma \vdash^e e' : \tau'$, then $\Omega; \Gamma \vdash^i [e'/x]i : \tau$.

Proof. We proceed by induction on the structure of \mathcal{T} . ■

As in the functional case, type preservation ensures that the type of an expression is identical to the type of its value. Intermediate evaluation steps require us to take into account arbitrary continuations and stores. We have the following generalization.

THEOREM 3.1 (Type preservation).

If $S \triangleright K \vdash i \hookrightarrow a$ with $\Omega; \cdot \vdash^i i : \tau$, $\Omega \vdash^K K : \tau \Rightarrow \tau'$ and $\Omega \vdash^S S : \Omega$, then $\Omega \vdash^a a : \tau'$.

Proof. We proceed by induction on the structure of a derivation of the evaluation judgment and inversion on the derivations of the typing judgments. ■

The type preservation result is formalized as follows at the top level of evaluation.

COROLLARY 3.1 (Type preservation).

If $\cdot \triangleright \mathbf{init} \vdash \mathbf{eval} e \hookrightarrow a$ with $\cdot; \cdot \vdash^e e : \tau$, then $\cdot \vdash^a a : \tau$. □

3.7. Representation in *LLF*

In this subsection, we give an *LLF* representation of the syntax of *MLR*, of its static and dynamic semantics and show how to exploit the resulting encoding of computations. The representation we propose is a natural extension of the *LF* code for *Mini-ML* found in the literature [33]. In particular, it retains its structure, its elegance, and the ease of proving its adequacy with respect to the informal presentation we just concluded. We describe the main issues in the representation by displaying fragments of the code and a limited number of adequacy statements. A complete treatment can be found in Appendix A. It is interesting to compare the result of our encoding with similar endeavors in the literature.

VanInwegen used the *HOL* theorem prover [24] to verify properties about a substantial portion of *Standard ML* [52]. She adopted a brute-force approach to the meta-representation problem, encoding, for example, contexts as terms. This choice resulted in a complex representation, and only partial achievement of the main goal of this endeavor: a formal proof of type preservation for that language. Although on a much simpler fragment, our use of higher-order abstract syntax, of parametric and hypothetical judgments, and of the linear features of *LLF* avoids these difficulties completely.

Chirimar used *Forum* [34] to represent a language similar to *MLR* with the addition of exceptions and continuations [12], but without any emphasis on typing. He took advantage of the higher-order nature of *Forum* and of its linear constructs. The resulting program is as elegant as our code and is proved adequate with respect to the informal specification of the object language. The absence of proof-terms in *Forum* prevents the direct manipulation of object-level derivations and no attempt is made to use that meta-language to investigate meta-theoretic properties such as type preservation.

Syntax

The representation of the syntactic level of *MLR* is based on higher-order abstract syntax and does not require the expressive power of the linear constructs of *LLF*. It lies therefore in the *LF* fragment of this language.

As is normally done in *LF*, every syntactic category of the object language is mapped to a distinguished base type. The type families necessary to encode the syntactic categories of *MLR* are given by the following declarations:

exp : type.	cell : type.
tp : type.	store : type.
instr : type.	cv : type.
cont : type.	answer : type.

The four declarations on the left encode expressions, instructions, types, and continuations. The four on the right are needed to represent the imperative features of *MLR* programs. **cell** corresponds to the lexical category of memory cells. **cv** and **store** will be used to represent the store. Finally, **answer** encodes final answers.

We encode the abstract syntax of *MLR* expressions, as described in the grammar of Section 3.7, by means of the representation function $\ulcorner _ \urcorner$. This function maps every production to an *LLF* object constant that, when applied to the representation of the subexpressions that it relates, yields an object that has type **exp**. The function $\ulcorner _ \urcorner$ is inductively defined on the left-hand side of the table below (we have separated out the treatment of the imperative constructs); its right-hand side gives the type of the constants used in the encoding.

$\ulcorner x \urcorner = x$	z : exp.
$\ulcorner z \urcorner = z$	s : exp \rightarrow exp.
$\ulcorner s e \urcorner = s \ulcorner e \urcorner$	case : exp
$\ulcorner \text{case } e \quad \urcorner = \text{case } \ulcorner e \urcorner$	\rightarrow exp
of z $\Rightarrow e_1$	\rightarrow (exp \rightarrow exp) \rightarrow exp.
s $x \Rightarrow e_2$	$([x:\text{exp}]\ulcorner e_2 \urcorner)$
$\ulcorner \langle \rangle \urcorner = \text{unit}$	unit : exp.
$\ulcorner \langle e_1, e_2 \rangle \urcorner = \text{pair } \ulcorner e_1 \urcorner \ulcorner e_2 \urcorner$	pair : exp \rightarrow exp \rightarrow exp.
$\ulcorner \text{fst } e \urcorner = \text{fst } \ulcorner e \urcorner$	fst : exp \rightarrow exp.
$\ulcorner \text{snd } e \urcorner = \text{snd } \ulcorner e \urcorner$	snd : exp \rightarrow exp.
$\ulcorner \text{lam } x. e \urcorner = \text{lam } ([x:\text{exp}]\ulcorner e \urcorner)$	lam : (exp \rightarrow exp) \rightarrow exp.
$\ulcorner e_1 e_2 \urcorner = \text{app } \ulcorner e_1 \urcorner \ulcorner e_2 \urcorner$	app : exp \rightarrow exp \rightarrow exp.
$\ulcorner \text{letval } x = e_1 \urcorner = \text{letval } \ulcorner e_1 \urcorner$	letval : exp
in e_2	\rightarrow (exp \rightarrow exp) \rightarrow exp.
$\ulcorner \text{letname } x = e_1 \urcorner = \text{letname } \ulcorner e_1 \urcorner$	letname : exp
in e_2	\rightarrow (exp \rightarrow exp) \rightarrow exp.
$\ulcorner \text{fix } x. e \urcorner = \text{fix } ([x:\text{exp}]\ulcorner e \urcorner)$	fix : (exp \rightarrow exp) \rightarrow exp.
$\ulcorner c \urcorner = \text{rf } c$	rf : cell \rightarrow exp.
$\ulcorner \text{ref } e \urcorner = \text{ref } \ulcorner e \urcorner$	ref : exp \rightarrow exp.
$\ulcorner !e \urcorner = !\ulcorner e \urcorner$! : exp \rightarrow exp.
$\ulcorner e_1 := e_2 \urcorner = \text{assign } \ulcorner e_1 \urcorner \ulcorner e_2 \urcorner$	assign : exp \rightarrow exp \rightarrow exp.
$\ulcorner e_1; e_2 \urcorner = \text{seq } \ulcorner e_1 \urcorner \ulcorner e_2 \urcorner$	seq : exp \rightarrow exp \rightarrow exp.

The representation of most expressions reflects directly the abstract syntax of *MLR*. We take advantage of higher-order abstract syntax in the representation of cells, variables, and binding constructs of *MLR*. Variables are encoded as *LLF* variables (of type `exp`). The fact that an object-level construct binds a variables x in a sub-expression e is then modeled by using the λ -abstraction of *LLF* in order to bind x in $\ulcorner e \urcorner$. Cells appear as hypotheses $c:\text{cell}$ in the context of *LLF*, similarly to free variables. Their representation as expressions is mediated by the constant `rf`, which maps entities of type `cell` to objects of type `exp`.

As an example, consider again the following *MLR* expression from the previous section:

```
letname f = ref (lam x.x)
in f := lam x.s x;
!f  $\langle$ 
```

It is represented by the following *LLF* term of type `exp`:

```
letname
  (ref (lam ([x] x)))
  ([f] (seq
    (assign f (lam ([x] (s x))))
    (app (! f) unit))))))
```

The faithfulness of this representation with respect to the object level syntax of expressions consists of a number of properties that we summarize in the following adequacy theorem, where Σ corresponds to the signature in Appendix A:

THEOREM 3.2 (Adequacy of the representation of *MLR* expressions).

*The function $\ulcorner _ \urcorner$ above is a compositional bijection between *MLR* expression with free variables among x_1, \dots, x_n and cells c_1, \dots, c_m , and canonical *LLF* objects M such that the judgment*

$$x_1:\text{exp}, \dots, x_n:\text{exp}, c_1:\text{cell}, \dots, c_m:\text{cell} \vdash_{\Sigma} M \uparrow \text{exp}$$

is derivable.

□

Compositionality in this statement means that the representation function commutes with substitution, i.e., that for every *MLR* expression e and e' , $\ulcorner [e'/x]e \urcorner = \ulcorner [e'/x]e \urcorner$. It confirms the correct application of higher-order abstract syntax in our encoding. Note that compositionality is not needed for cells since they are never subject to substitution.

Due to the complexity of our object language, we do not display the simple but long and somewhat tedious inductive proof of this statement. The interested reader is referred to [6] for a full treatment; the proof of a different adequacy statement is sketched at the end of this section. The techniques used in order to prove adequacy theorems for *LLF* encodings naturally extends the methods successfully applied for years in the more restricted setting of *LF*. In particular, they retain their simplicity in our richer application area. This contrasts with other proposals, e.g., the treatment of linearity in *LF* itself [42], where adequacy theorems have complex proofs even for simple object languages.

Types, instructions and continuations are represented in a similar way. The *LLF* declarations for the constants needed in their encoding can be found in Appendix A. We omit displaying the statements of the respective adequacy theorems since they do not introduce new concepts. They can be found in [6].

MLR makes a dual usage of the collection of cell-value pairs that we informally referred to as its store: as a repository from which to retrieve the value associated with a cell during evaluation (the proper store we indicated as S), and as a term to be eventually returned with the final answer (the reified store we denoted $[S]$). We will correspondingly have two distinct *LLF* representations of the store. We will discuss the *internal* encoding $\ulcorner S \urcorner$ of a proper store S when considering evaluation. A reified store $[S]$ is given the following *external* representation $\ulcorner [S] \urcorner$:

$$\begin{array}{ll} \ulcorner [\cdot] \urcorner = \text{estore} & \text{estore} : \text{store} \\ \ulcorner [S, c = v] \urcorner = \text{with } \ulcorner [S] \urcorner (\text{holds } c \ulcorner v \urcorner) & \text{with} : \text{store} \rightarrow \text{cv} \rightarrow \text{store}. \\ & \text{holds} : \text{cell} \rightarrow \text{exp} \rightarrow \text{cv}. \end{array}$$

Here and in the following, we systematically overload the notation $\ulcorner _ \urcorner$ used for expressions to denote the various representation function that are required in this example. The nature of its argument should always permit disambiguating which specific function we are referring to in each case.

The representation of answers directly expresses the grammatical rules:

$$\begin{array}{ll} \ulcorner ([S], v) \urcorner = \text{close } \ulcorner [S] \urcorner \ulcorner v \urcorner & \text{close} : \text{store} \rightarrow \text{exp} \rightarrow \text{answer}. \\ \ulcorner \text{new } c. a \urcorner = \text{new } ([c:\text{cell}] \ulcorner a \urcorner) & \text{new} : (\text{cell} \rightarrow \text{answer}) \rightarrow \text{answer}. \end{array}$$

The declarations for these constants are repeated in Appendix A. The adequacy theorems that link them to the syntax of *MLR* are reported in [6].

Static Semantics

On the basis of the above encoding of the syntax of *MLR*, we will now describe the meta-representation of the static semantics of this language.

As for syntax, the representation of the static semantics of *MLR* does not rely on the linear features of *LLF*. The resulting code lies therefore in the *Elf* fragment of our logical framework. We have the following declarations for the type families that model the various typing judgments presented in Section 3.2:

$$\begin{array}{l} \text{tpe} : \text{exp} \rightarrow \text{tp} \rightarrow \text{type}. \\ \text{tpi} : \text{instr} \rightarrow \text{tp} \rightarrow \text{type}. \\ \text{tpK} : \text{cont} \rightarrow \text{tp} \rightarrow \text{tp} \rightarrow \text{type}. \\ \text{tpc} : \text{cell} \rightarrow \text{tp} \rightarrow \text{type}. \\ \text{tpS} : \text{store} \rightarrow \text{type}. \\ \text{tpa} : \text{answer} \rightarrow \text{tp} \rightarrow \text{type}. \end{array}$$

Again we have separated out the declarations that suffice for the functional core of *MLR* from the type families that are required to handle the imperative aspects of this language. The first three represent the typing judgments for expressions, instructions, and continuations, while tpS and tpa encode respectively the store

and answer typing judgments, and **tpc** records the type of individual cells in the store.

We illustrate the representation of the static semantics of *MLR* by displaying how to encode a typing derivation for expressions. The remaining typing judgments are treated similarly and the resulting *LLF* declarations are presented in Appendix A.

In Section 3.2, we denoted the fact that an expression e has type τ assuming the types given in Γ for its free variables and the types given in Ω for its reference cells as the hypothetical judgment $\Omega; \Gamma \vdash^e e : \tau$. We represent the schematic form of this judgment by means of the *LLF* type family **tpe**. This family accepts two parameters: the representation of an expression and the representation of a type. Therefore, the instance above will be encoded as the *LLF* base type **tpe** $\ulcorner e \urcorner \ulcorner \tau \urcorner$. The context Γ is taken into consideration only when checking that this term is indeed derivable in *LLF*. Then, we will encode each pair $x_i : \tau_i$ in Γ by means of an *LLF* hypothesis

$$t_i : \mathbf{tpe} \ x_i \ \ulcorner \tau_i \urcorner$$

where the free variable x_i is declared as an expression ($x_i : \mathbf{exp}$). Similarly, we encode every item $c_j : \tau'_j$ in Ω as the (intuitionistic) assumption

$$t'_j : \mathbf{tpc} \ c_j \ \ulcorner \tau'_j \urcorner$$

in the context of *LLF*, where c_j is declared as a cell ($c_j : \mathbf{cell}$). Note that **tpc** only serves the purpose of making typing assumptions for cells. We write $\ulcorner \Gamma \urcorner$ and $\ulcorner \Omega \urcorner$ for the encoding we just outlined for the context Γ and the store context Ω , respectively.

The inference rules defining the derivability of the typing judgment for *MLR* are encoded according to the technique presented in Section 3.1. We consider two rules as additional examples, the remaining declarations can be found in Appendix A. Rule **tpe_z** associates the type **nat** to the numeral **z**. We represent it by means of the *LF* constant **tpe_z** that relate **z** to **nat**:

$$\ulcorner \frac{}{\Omega; \Gamma \vdash^e \mathbf{z} : \mathbf{nat}} \mathbf{tpe_z} \urcorner = \mathbf{tpe_z} : \mathbf{tpe} \ \mathbf{z} \ \mathbf{nat}.$$

Rule **tpe_case** specifies how to type-check a conditional expression. We repeat it from Figure 12:

$$\frac{\Omega; \Gamma \vdash^e e : \mathbf{nat} \quad \Omega; \Gamma \vdash^e e_1 : \tau \quad \Omega; \Gamma, x : \mathbf{nat} \vdash^e e_2 : \tau}{\Omega; \Gamma \vdash^e \mathbf{case} \ e \ \mathbf{of} \ \mathbf{z} \ \Rightarrow \ e_1 \ \mid \ \mathbf{s} \ x \ \Rightarrow \ e_2 : \tau} \mathbf{tpe_case}$$

This rule has multiple premises. It is hypothetical because the rightmost premise inserts the assumption $x : \mathbf{nat}$ into the context. It is also parametric since the variable x is bound in the case construct, but it appears as a new symbol both in the added hypothesis and in the expression e_2 type-checked by the rightmost premise. We represent this rule by means of the *LLF* constant **tpe_case** and encode its structure in the associated type. We have the declaration:

```
tpe_case :   tpe E nat
            -> tpe E1 T
```

```

-> ({x:exp} tpe x nat -> tpe (E2 x) T)
-> tpe (case E E1 ([x:exp] E2 x)) T.

```

Notice the quantification over x and the embedded implication with antecedent $\mathbf{tpe\ x\ nat}$ in the encoding of the third premise. In this declaration, the *LLF* variables E , $E1$, $E2$ and T correspond to the schematic variables e , e_1 , e_2 and τ respectively. They are implicitly quantified at the front of the declaration.

It is worth noticing that there is no declaration in correspondence of rule $\mathbf{tpe_x}$:

$$\frac{}{\Omega; \Gamma, x:\tau \vdash^e x : \tau} \mathbf{tpe_x}$$

Since assumptions are represented directly in the context of *LLF*, a judgment of the form $\mathbf{tpe\ x\ T}$, where T is the representation of some concrete type τ , will be validated by accessing the context of *LLF* rather than the signature, and succeed precisely when $t_x : \mathbf{tpe\ x\ T}$ appears in it as an assumption. Similar considerations hold for reference cells.

We have now the means for representing derivations of expression typing judgments. The adequacy theorem below ensures that whenever \mathcal{T} is a (valid) derivation for the *MLR* judgment $\Omega; \Gamma \vdash^e e : \tau$, it is a canonical inhabitant of the *LLF* type $\mathbf{tpe\ } \ulcorner e \urcorner \ulcorner \tau \urcorner$ with respect to the proper encoding of Γ and Ω , and vice versa.

THEOREM 3.3 (Adequacy of the representation of *MLR* expression typing).

Given an MLR expression e and a type τ , there is a compositional bijection $\ulcorner _ \urcorner$ between derivations of

$$\Omega; \Gamma \vdash^e e : \tau$$

and LLF objects M such that

$$\ulcorner \Omega \urcorner, \ulcorner \Gamma \urcorner \vdash_{\Sigma} M \uparrow \mathbf{tpe\ } \ulcorner e \urcorner \ulcorner \tau \urcorner$$

is derivable. □

Dynamic Semantics

Unlike syntax and static semantics, the representation of evaluation relies heavily on the linear features of *LLF*. It is based on the following four type families:

```

ev          : cont -> instr -> answer -> type.
contains    : cell -> exp -> type.
collect     : store -> type.
read        : cell -> exp -> type.

```

which we will describe in turn.

Assuming the appropriate representation functions $\ulcorner _ \urcorner$ for continuations, instructions, and answers, we model the continuation-based judgment $S \triangleright K \vdash i \hookrightarrow a$ as the *LLF* base type $\mathbf{ev\ } \ulcorner K \urcorner \ulcorner i \urcorner \ulcorner a \urcorner$. The store S mentioned by this judgment is

represented in a distributed fashion in the context of *LLF*. Each item $c = v$ in S is modeled by two assumptions: first of all, we need to declare c as a cell and we do so by means of the assumption $c:\mathbf{cell}$, second, we represent the fact that the current contents of c is v by a *linear* hypothesis of the form $h \hat{\mathbf{contains}} c \ulcorner v \urcorner$. The first assumption should clearly be intuitionistic since c may be mentioned many times in K , i , a and S . In contrast, the second must be linear since assignment updates the value associated with a cell destructively. If h were an intuitionistic hypothesis, we would have no means of prohibiting the old value from being accessed. In summary, we associate to every proper store $S = (c_1 = v_1, \dots, c_n = v_n)$ the following *internal representation*:

$$\ulcorner S \urcorner = \left[\begin{array}{l} c_1:\mathbf{cell}, \quad \dots, c_n:\mathbf{cell}, \\ h_1 \hat{\mathbf{contains}} c_1 \ulcorner v_1 \urcorner, \dots, h_n \hat{\mathbf{contains}} c_n \ulcorner v_n \urcorner \end{array} \right]$$

Four rules in the deductive system for continuation-based evaluation presented in Figures 14–15 access the store directly: $\mathbf{ev_ref}^*$, $\mathbf{ev_assign}_2^*$, $\mathbf{ev_deref}^*$, and $\mathbf{ev_init}$. We will illustrate the use of the linear features of *LLF* on their encoding. However, in order to gain familiarity with our representation technique, we will first analyze rule $\mathbf{ev_z}$. All other inference figures are treated similarly to this rule. The complete code is displayed in Appendix A.

A bottom up reading of rule $\mathbf{ev_z}$, shown below on the left, specifies that evaluating \mathbf{z} simply amounts to returning it as a value. We represent this rule by means of the declaration for the constant $\mathbf{ev_z}$ shown on the right:

$$\frac{\ulcorner S \triangleright K \vdash \mathbf{return} \mathbf{z} \hookrightarrow a \urcorner}{S \triangleright K \vdash \mathbf{eval} \mathbf{z} \hookrightarrow a} \mathbf{ev_z} = \left[\begin{array}{l} \mathbf{ev_z} : \quad \mathbf{ev} \ K \ (\mathbf{return} \ \mathbf{z}) \ A \\ \quad \quad \quad \mathbf{-o} \ \mathbf{ev} \ K \ (\mathbf{eval} \ \mathbf{z}) \ A. \end{array} \right]$$

The linear arrow in the representation of rule $\mathbf{ev_z}$ enables its antecedent and its consequent to access the same linear assumptions in the context. This accounts for the fact that the premise and the conclusion of this rule mention the same store. Had we used an intuitionistic implication, the antecedent (and therefore the whole expression) would have been applicable only with contexts deprived of any linear assumptions, corresponding to empty stores.

Rule $\mathbf{ev_ref}^*$, repeated below on the left, creates a new location c in the store and initializes it with the argument v of \mathbf{ref}^* . Its representation on the right models these actions on the context of *LLF*: the new cell is intuitionistically assumed when processing the dependent type $\{\mathbf{c}:\mathbf{cell}\}$, while the resolution of the embedded linear implication has the effect of asserting $\mathbf{contains} \ c \ V$ in the linear part of the context. Since this assumption is made linearly, it will be possible to remove it from the context, for example in order to update the value contained in c in response to an assignment. Notice how the newly created cell c is bound in the final answer.

$$\frac{\ulcorner (S, c = v) \triangleright K \vdash \mathbf{return} \ c \hookrightarrow a \urcorner}{S \triangleright K \vdash \mathbf{ref}^* \ v \hookrightarrow \mathbf{new} \ c. \ a} \mathbf{ev_ref}^* = \left[\begin{array}{l} \mathbf{ev_ref}^* : \ (\{\mathbf{c}:\mathbf{cell}\} \ \mathbf{contains} \ c \ V \\ \quad \quad \quad \mathbf{-o} \ \mathbf{ev} \ K \ (\mathbf{return} \ (\mathbf{rf} \ c)) \ (A \ c)) \\ \quad \quad \quad \mathbf{-o} \ \mathbf{ev} \ K \ (\mathbf{ref}^* \ V) \ (\mathbf{new} \ ([\mathbf{c}:\mathbf{cell}] \ A \ c)). \end{array} \right]$$

Of the three rules that realize assignment in *MLL*, only **ev_assign*₂** accesses the store. The declaration **ev_assign*₂** below mimics the destructive update of the contents of the cell c (written C in the clause) in the store in two steps. First the old value is retrieved by **contains C V'**. Since it appears as a linear assumption, accessing it causes its removal from the linear context of *LLF*. Since the other antecedent of this clause is reached through the multiplicative connective \multimap , the remaining linear hypotheses will be passed to it. This term inserts the new value v (i.e. V) of c in the representation of the store by means of the antecedent **contains C V** of the embedded linear implication.

$$\frac{\ulcorner (S, c = v) \triangleright K \vdash \mathbf{return} \langle \rangle \hookrightarrow a \urcorner}{(S, c = v') \triangleright K \vdash c :=_2^* v \hookrightarrow a} \mathbf{ev_assign}_2^* =$$

$$\left[\begin{array}{l} \mathbf{ev_assign*2} : (\mathbf{contains} \ C \ V \ \multimap \ \mathbf{ev} \ K \ (\mathbf{return} \ \mathbf{unit}) \ A) \\ \multimap (\mathbf{contains} \ C \ V' \ \multimap \ \mathbf{ev} \ K \ (\mathbf{assign*2} \ (\mathbf{rf} \ C) \ V) \ A). \end{array} \right.$$

Dereferencing a cell c is naturally modeled in *LLF* through the use of the additive operators of our language. In order to encode rule **ev_deref***, we need two copies of the store representation: one to retrieve the contents of c , and one to proceed with the evaluation. This is immediately achieved by means of the additive conjunction of *LLF*. We have the following declaration:

$$\frac{\ulcorner S \vdash c = v \quad S \triangleright K \vdash \mathbf{return} \ v \hookrightarrow a \urcorner}{S \triangleright K \vdash \mathbf{deref}^* \ c \hookrightarrow a} \mathbf{ev_deref}^* =$$

$$\left[\begin{array}{l} \mathbf{ev_deref*} : \quad (\mathbf{read} \ C \ V \\ \quad \& \ \mathbf{ev} \ K \ (\mathbf{return} \ V) \ A) \\ \multimap \ \mathbf{ev} \ K \ (\mathbf{deref*} \ (\mathbf{rf} \ C)) \ A. \end{array} \right.$$

The conjunct **read C V**, which implements the read judgment $S \vdash c = v$, looks up its copy of the linear context in search of the assumption **contains C V** and relies on the additive unit of *LLF*, written $\langle T \rangle$, to discard the rest. This technique is generally applicable to every situation that involves looking up the encoding of volatile information. The definition of **read** consists of a single clause encoding rule **read_val**:

$$\frac{\ulcorner \quad \urcorner}{(S, c = v) \vdash c = v} \mathbf{read_val} = \left[\begin{array}{l} \mathbf{read_val} : \quad \mathbf{contains} \ C \ V \\ \multimap \ \langle T \rangle \\ \multimap \ \mathbf{read} \ C \ V. \end{array} \right.$$

We could have alternatively modeled dereferencing similarly to assignment by first accessing the linear assumption **contains C V** directly. In order to balance its consequent removal from the linear context of *LLF*, this same assumption should be re-entered in the context before returning the value V . We would have the following declaration:

$$\mathbf{ev_deref*}' : \quad (\mathbf{contains} \ C \ V \ \multimap \ \mathbf{ev} \ K \ (\mathbf{return} \ V) \ A) \\ \multimap (\mathbf{contains} \ C \ V \ \multimap \ \mathbf{ev} \ K \ (\mathbf{deref*} \ (\mathbf{rf} \ C)) \ A).$$

Although it achieves a similar effect, this declaration does not encode rule **ev_deref***, or **read_val**, or any combination of these rules. Instead, it is a transliteration of

the following inference rule, which we could have used to formalize dereferencing:

$$\frac{(S, c = v) \triangleright K \vdash \mathbf{return} \ v \hookrightarrow a}{(S, c = v) \triangleright K \vdash \mathbf{deref}^* \ c \hookrightarrow a} \mathbf{ev_deref}^{**'}$$

Finally, rule **ev_init** pairs up the store and the final value in order to produce the answer. We model this behavior by means of the auxiliary procedure **collect** which translates the internal representation of the store S , as linear LLF assumptions, to its external representation $[S]$, as an object of type **store**.

$$\frac{\ulcorner \quad \urcorner}{S \triangleright \mathbf{init} \vdash \mathbf{return} \ v \hookrightarrow ([S], v)} \mathbf{ev_init} \urcorner = \left[\begin{array}{l} \mathbf{ev_init} : \quad \mathbf{collect} \ S \\ \quad \quad \quad \mathbf{-o} \ \mathbf{ev} \ \mathbf{init} \ (\mathbf{return} \ V) \ (\mathbf{close} \ S \ V). \end{array} \right.$$

The code for **collect** is displayed below.

```
col_empty : collect estore.
col_cv    :   contains C V
           -o collect S
           -o collect (with S (holds C V)).
```

Since the use of multiplicatives removes the assumptions **contains C V** as they are retrieved, each recursive access to **collect** adds a different item to the external representation of the store. Clause **col_empty** is provable only when the linear part of the context of LLF is empty, and therefore only when the complete store of MLR has been externalized.

The effectiveness of the representation we just illustrated relies on the ability to remove objects from the context of LLF . Using LF on this problem would have produced awkward encodings with prohibitive consequence for the development of the meta-theory of MLR [6]: a first alternative would have relied entirely on the external representation of the store, implementing all the operations required to access and modify it explicitly. A second alternative would be to proceed as we did, with the tedious addition of declarations aimed at checking the linearity of the resulting derivations a posteriori.

We will now make the above motivating discussion more precise. The faithfulness of our representation of evaluation is expressed by the following adequacy theorem.

THEOREM 3.4 (Adequacy of the representation of MLR evaluation).

Given an MLR continuation K , an instruction i , a store S and an answer a , where K , i , S and a are closed except for the possible presence of free cells, there is a bijection $\ulcorner _ \urcorner$ between derivations of

$$S \triangleright K \vdash i \hookrightarrow a$$

and LLF objects M such that

$$\ulcorner S \urcorner \vdash_{\Sigma} M \uparrow \mathbf{ev} \ulcorner K \urcorner \ulcorner i \urcorner \ulcorner a \urcorner$$

is derivable.

□

In order to prove the above theorem, we will decompose it into four parts. Again, Σ is the signature contained in Appendix A. We need to prove the following properties:

Functionality: $\ulcorner _ \urcorner$ is a total function from *MLR* evaluation derivations to *LLF* objects over Σ .

Soundness: The representation of a derivation of a given *MLR* evaluation judgment is an *LLF* object whose type is the representation of this judgment.

Completeness: Whenever a canonical *LLF* object over Σ inhabits the type corresponding to the encoding of an *MLR* evaluation judgment, this object is the representation of a derivation of that judgment.

Bijectivity: $\ulcorner _ \urcorner$ is a bijection between evaluation derivations in *MLR* and canonical *LLF* objects whose type encodes the corresponding evaluation judgment.

Differently from expressions and typing derivations, the representation function $\ulcorner _ \urcorner$ is trivially compositional (it involves closed expressions only), otherwise we should prove it as an additional property.

Detailed proofs of these properties are long and rather tedious, although conceptually simple. We will sketch them by using the declaration for `ev_assign*2` as a representative case. In order to do so, we repeat it complete with the Π -quantifiers we omitted in the above presentation:

```
ev_assign*2 : {C:cell}{V':exp}{V:exp}{K:cont}{A:answer}
              (contains C V -o ev K (return unit) A)
              -o (contains C V' -o ev K (assign*2 (rf C) V) A).
```

In the specific case of this example, it is convenient to state and prove the functionality and soundness properties together. We have the following result:

LEMMA 3.3 (Functionality and soundness of *MLR* evaluation's representation).

*Given a store S , a continuation K , an instruction i and an answer a , where K , i , S and a are closed except for the possible presence of free cells, for every derivation \mathcal{E} of the judgment $S \triangleright K \vdash i \hookrightarrow a$, $\ulcorner \mathcal{E} \urcorner$ is defined and unique, and the *LLF* judgment*

$$\ulcorner S \urcorner \vdash_{\Sigma} \ulcorner \mathcal{E} \urcorner \uparrow \text{ev} \ulcorner K \urcorner \ulcorner i \urcorner \ulcorner a \urcorner$$

is derivable.

Proof. This proof proceeds by induction on the structure of the derivation \mathcal{E} . We illustrate only the case in which it ends with an application of rule `ev_assign*2`. Therefore,

$$\mathcal{E}' = \frac{(S^*, c = v) \triangleright K \vdash \mathbf{return} \langle \rangle \hookrightarrow a}{(S^*, c = v') \triangleright K \vdash c :=_2^* v \hookrightarrow a} \text{ev_assign}_2^*$$

where $S = (S^*, c = v')$ and $i = (c :=_2^* v)$. Let us also denote as S' the store $(S^*, c = v)$. Notice that $\ulcorner S \urcorner = \ulcorner S' \urcorner$.

By induction hypothesis on \mathcal{E}' , we deduce that there is a unique *LLF* object M' such that $M' = \ulcorner \mathcal{E}' \urcorner$ and there is a derivation of the judgment $\ulcorner S' \urcorner \vdash_{\Sigma} M' \uparrow \text{ev} \ulcorner K \urcorner (\text{return unit}) \ulcorner a \urcorner$.

Iterated applications of the *LLF* rule **oa_iapp** are used to instantiate the arguments of the declaration for **ev_assign*2**. Indeed, there is an atomic derivation \mathcal{A}'' of the following judgment:

$$\overline{\ulcorner S' \urcorner} \vdash_{\Sigma} \text{ev_assign*2 } c \ulcorner v' \urcorner \ulcorner v \urcorner \ulcorner K \urcorner \ulcorner a \urcorner$$

$$\downarrow \left[\begin{array}{l} (\text{contains } c \ulcorner v \urcorner \\ \multimap \text{ev} \ulcorner K \urcorner (\text{return unit}) \ulcorner a \urcorner) \\ \multimap (\text{contains } c \ulcorner v' \urcorner \\ \multimap \text{ev} \ulcorner K \urcorner (\text{assign*2 } (\text{rf } c) \ulcorner v \urcorner) \ulcorner a \urcorner) \end{array} \right]$$

Let $t \hat{=} \text{contains } c \ulcorner v \urcorner$ be the assumption in $\ulcorner S' \urcorner$ corresponding to the pair $(c = v)$ in S' . We can abstract it over M' in the *LLF* derivation for this object, obtaining a derivation \mathcal{C}'' of the judgment $\ulcorner S^{\urcorner*} \vdash_{\Sigma} \hat{\lambda}t : \text{contains } c \ulcorner v \urcorner. M' \uparrow (\text{contains } c \ulcorner v \urcorner \multimap \text{ev} \ulcorner K \urcorner (\text{return unit}) \ulcorner a \urcorner)$, where $\ulcorner S^{\urcorner*}$ differs from $\ulcorner S' \urcorner$ only by the removal of assumption t . Since $\overline{\ulcorner S' \urcorner} = \ulcorner S^{\urcorner*}$, we can then apply rule **oa_lapp** to \mathcal{A}'' and \mathcal{C}'' , obtaining a derivation \mathcal{A}' of:

$$\ulcorner S^{\urcorner*} \vdash_{\Sigma} \text{ev_assign*2 } c \ulcorner v' \urcorner \ulcorner v \urcorner \ulcorner K \urcorner \ulcorner a \urcorner \wedge (\hat{\lambda}t : \text{contains } c \ulcorner v \urcorner. M')$$

$$\downarrow (\text{contains } c \ulcorner v' \urcorner \multimap \text{ev} \ulcorner K \urcorner (\text{assign*2 } (\text{rf } c) \ulcorner v \urcorner) \ulcorner a \urcorner)$$

Let $t' \hat{=} \text{contains } c \ulcorner v' \urcorner$ be the assumption in $\ulcorner S^{\urcorner}$ corresponding to the pair $(c = v')$ in the store S . Then, there is a derivation \mathcal{C}' of the *LLF* judgment $(\ulcorner S^{\urcorner*}, t' \hat{=} \text{contains } c \ulcorner v' \urcorner) \vdash_{\Sigma} t' \uparrow \text{contains } c \ulcorner v' \urcorner$. We can then apply rule **oa_lapp** again to \mathcal{A}' and \mathcal{C}' to obtain a derivation \mathcal{A} of:

$$\ulcorner S^{\urcorner*}, t' \hat{=} \text{contains } c \ulcorner v' \urcorner \vdash_{\Sigma}$$

$$\text{ev_assign*2 } c \ulcorner v' \urcorner \ulcorner v \urcorner \ulcorner K \urcorner \ulcorner a \urcorner \wedge (\hat{\lambda}t : \text{contains } c \ulcorner v \urcorner. M') \wedge t'$$

$$\downarrow \text{ev} \ulcorner K \urcorner (\text{assign*2 } (\text{rf } c) \ulcorner v \urcorner) \ulcorner a \urcorner$$

In order to understand this formula, observe that $\ulcorner S^{\urcorner} = (\ulcorner S^{\urcorner*}, t' \hat{=} \text{contains } c \ulcorner v' \urcorner)$. We now apply rule **oc_a** to \mathcal{A} to get the desired canonical derivation.

At this point, it is enough to notice that the *LLF* object M appearing on the left of the arrow in this canonical judgment is the representation of the *MLR* derivation \mathcal{E} above and that the type on the right of the arrow is the representation of its type. It is also easy to ascertain that M is unique, given the uniqueness of M' . ■

We now consider the completeness of the encoding of *MLR* evaluation derivations. We have the following lemma.

LEMMA 3.4 (Completeness of the representation of *MLR* evaluation).

Given a store S , a continuation K , an instruction i and an answer a , where K , i , S and a are closed except for the possible presence of free cells, for every *LLF* object M such that the judgment

$$\ulcorner S \urcorner \vdash_{\Sigma} M \uparrow \text{ev} \ulcorner K \urcorner \ulcorner i \urcorner \ulcorner a \urcorner$$

has a derivation \mathcal{C} , there is a derivation \mathcal{E} of the *MLR* judgment $S \triangleright K \vdash i \hookrightarrow a$ such that $M = \ulcorner \mathcal{E} \urcorner$.

Proof. We proceed by induction on the structure of M . Since the type in \mathcal{C} is a base type, M can either be a constant, a variable or start with a destructor. Then M has the following structure:

$$M = c_M * M_1 * \dots * M_n,$$

where c_M is a constant in Σ of some appropriate type, $*$ represents either linear or intuitionistic application, and M_1, \dots, M_n are objects of some type. The proof now distinguishes cases on the basis of possible constants c_M . We consider only the case in which this constant is `ev_assign*2`.

If c_M is `ev_assign*2`, then it must be the case that $i = (c :=_2^* v)$ for some cell c and expression v , and moreover

$$M = \text{ev_assign*2 } c M_{v'} \ulcorner v \urcorner \ulcorner K \urcorner \ulcorner a \urcorner \wedge M^* \wedge M_{t'}$$

By analyzing the types of the objects $M_{v'}$, M^* and $M_{t'}$, we deduce that there is an expression v' such that $M_{v'} = \ulcorner v' \urcorner$, that $M^* = \hat{\lambda}t : \text{contains } c \ulcorner v \urcorner . M'$ for some term M' of type `ev` $\ulcorner K \urcorner$ (`return unit`) $\ulcorner a \urcorner$, and that $M_{t'} = t'$ for some linear assumption $t' \text{? contains } c \ulcorner v' \urcorner$. Moreover, we have that $S = (S^*, c = v')$.

We can apply the induction hypothesis on M' relative to a store representation that differs from $\ulcorner S \urcorner$ by the replacement of assumption t' with t . The corresponding *MLR* store S' is $(S^*, c = v)$. We deduce in this way that there exist a derivation \mathcal{E}' of the judgment $(S^*, c = v) \triangleright K \vdash \text{return } \langle \rangle \hookrightarrow a$. An application of rule `ev_assign*2` suffices to obtain the desired derivation. ■

We conclude the treatment of the adequacy of the representation of *MLR* evaluation derivations by showing that the function $\ulcorner _ \urcorner$ is indeed bijective.

LEMMA 3.5 (Bijectivity of the representation of *MLR* evaluation).

Given a store S , a continuation K , an instruction i and an answer a , where K , i , S and a are closed except for the possible presence of free cells, the representation function $\ulcorner _ \urcorner$ is a bijection between derivations \mathcal{E} of *MLR* the judgment $S \triangleright K \vdash i \hookrightarrow a$, and *LLF* objects of type `ev` $\ulcorner K \urcorner \ulcorner i \urcorner \ulcorner a \urcorner$ in the context $\ulcorner S \urcorner$.

Proof. Lemma 3.3 establishes that $\ulcorner _ \urcorner$ is a total function from the set of *MLR* derivations mentioned in the statement to the specified set of *LLF* objects. By the completeness lemma, we deduce that this function is surjective. It therefore remains solely to prove that it is also injective. Given two derivations \mathcal{E}_1 and \mathcal{E}_2 such that $\ulcorner \mathcal{E}_1 \urcorner = \ulcorner \mathcal{E}_2 \urcorner$, the proof that $\mathcal{E}_1 = \mathcal{E}_2$ proceeds by induction on these derivations. ■

A derivation \mathcal{E} for an evaluation judgment $S \triangleright K \vdash i \hookrightarrow a$ is a trace of the computation that an continuation-based *MLR* interpreter performs when evaluating the instruction i and the continuation K to the final answer a with respect to the store contents S . According to the above adequacy theorem, such derivations are faithfully represented by the terms M inhabiting the *LLF* type encoding this judgment. We conclude this section by illustrating how to take advantage of

this internal representation of *MLR* computations. We only give a small example here—more interesting examples such as the proof of type preservation, or a cut elimination procedure for classical linear logic can be found in [6]. Specifically, we will give *LLF* declarations that permit counting the number of reference cells dynamically allocated during the evaluation.

In order to achieve this purpose, we first give the following declarations for natural numbers:

```
num : type.
zero : num.
succ : num -> num.
```

The counting judgment relates an *MLR* computation to the number of cells it allocates. It is represented by the following type family

```
count : ev K I A -> num -> type.
```

We implement the counting procedure in *LLF* by unfolding the representation of an *MLR* computation. We ignore the steps that do not allocate memory cells, but increment by one the counter every time rule **ev_ref*** is applied. We show three declarations, corresponding to the initialization step performed by rule **ev_init**, the allocation of a new cell by rule **ev_ref***, and one of the numerous cases where nothing happens (rule **ev_z**):

```
cnt_init : count (ev_init ^ C) zero.
cnt_ref* : ({c:cell}{d:contains c V} count (C c ^ d) N)
          -> count (ev_ref* ^ ([c:cell] [d^contains c V] C c ^ d)) (succ N).
cnt_z    : count C N
          -> count (ev_z ^ C) N.
```

We conclude this section by displaying the *LLF* base type that ascertains that the *MLR* expression

$$\begin{aligned} &\mathbf{letname} \ f = \mathbf{ref} \ (\mathbf{lam} \ x. x) \\ &\mathbf{in} \ f := \mathbf{lam} \ x. s \ x; \\ &\quad !f \langle \rangle \end{aligned}$$

allocates two cells. The concrete syntax of this term is as follows, where the first argument is the representation of the evaluation of the above expression:

```
count
(ev_letname ^
 (ev_seq ^
 (ev_assign ^
 (ev_ref ^
 (ev_lam ^
 (ev_cont ^
 (ev_ref* ^
 ([c1:cell] [Cn1^contains c1 (lam [x:exp] x)] ev_cont ^
 (ev_assign*1 ^
 (ev_lam ^
 (ev_cont ^
 (ev_assign*2 ^
 ([Cn1'^contains c1 (lam [x:exp] s x)] ev_cont ^
 (ev_app ^
 (ev_deref ^
```

```

(ev_ref ^
 (ev_lam ^
  (ev_cont ^
   (ev_ref* ^
    ([c2:cell] [Cn2^contains c2 (lam [x:exp] x)] ev_cont ^
     (ev_deref* ^
      (read_val ^ () ^ Cn2 ,
       ev_cont ^
        (ev_app* ^
         (ev_unit ^
          (ev_cont ^
           (ev_unit ^
            (ev_init ^ (col_cv ^ Cn2 ^ (col_cv ^ Cn1' ^ col_empty))))
            ))))))))))) ^ Cn1))))))))) (succ (succ zero))

```

4. CONCLUSION AND FUTURE WORK

In this paper, we have presented the linear logical framework *LLF* as an extension of *LF* with internal support for the representation of state-based problems. We have demonstrated its expressive power by providing a usable representation of the syntax and the semantics of an imperative variant of the functional programming language *Mini-ML*; space reasons prevented us from extending this encoding to aspects of the meta-theory of this language, such as a proof of its type preservation property [6]. Additional substantial case studies we have completed include the formalization of a proof of cut elimination for classical linear logic, translations between minimal linear natural deduction and sequent calculus, as well as a number of puzzles and solitaires. The interested reader may access them on the World-Wide Web at [9] or in [6].

The representation language of *LLF*, $\lambda^{\Pi \dashv \multimap \& \top}$, conservatively extends *LF*'s λ^{Π} with constructs from linear logic. We can think of it as the type theory obtained from the type constructors \top , $\&$, \dashv , and Π . This choice of constructors is complete in the sense that they suffice to represent full intuitionistic or classical linear logic. Further, adding any other linear connective as a free type constructor destroys the property that usable canonical forms exist by introducing commuting conversions. This property is necessary in the proofs of adequacy theorems for encodings and also for the interpretation of *LLF* as an abstract logic programming language.

The meta-representation methodology of *LLF* extends the judgments-as-types technique adopted in *LF* with a direct way to map state-related constructs and behaviors onto the linear operators of $\lambda^{\Pi \dashv \multimap \& \top}$. The resulting representations retain the elegance and immediacy that characterize *LF* encodings, and the ease of proving their adequacy.

LLF maintains the computational nature of *LF* as an abstract logic programming language. The implementation of *LLF* combines the experience with higher-order logic programming languages gained with *Elf* [38, 41], an older realization of *LF*, on previous research work on linearity as in the language *Lolli* [28, 8], and on new experimental term representation [11] and compilation [7] techniques. Among the new problems is the necessity of performing higher-order unification between linear terms [10].

LLF generalizes other formalisms based on linear logic such as *Forum* [34] by making linear objects available for representations, by permitting proof terms and by providing linear types. It is closely related to the system *RLF* of Ishtiaq and

Pym [30], which allows dependencies on linear variables, but does not have \top as an operator. Linear dependent types are potentially useful but not essential in our experience, while \top is a necessary tool in many representation problems. The meta-theory of *LLF* appears significantly simpler than that of *RLF*, a fact that might imply that proving the adequacy of an encoding may be substantially more complex in this formalism. Finally, our approach is orthogonal to general logics in the style of *LU* [22].

In the near future, we intend to gain experience with the use of *LLF* as a representation language by encoding state-based deductive systems such as imperative programming languages constructs, hardware systems, security protocols, and real-time systems. The availability of an implementation will be of great help in doing so since it will enable us to concentrate on high-level representation issues. We would also like to extend the tools available in *Twelf* [48, 46], notably the theorem proving component of this system [47], to handle the possibilities offered by the linear operators of *LLF*. Finally, we are interested in investigating a generalization of the type constructors $\&$ and \multimap of $\lambda^{\Pi \multimap \& \top}$ to linear Σ and Π types, respectively, although it currently appears that this would greatly complicate the type theory while it is not clear how much would be gained.

APPENDIX A

Formalization of *MLR*

A.1. SYNTAX

```

exp      : type.
tp       : type.
instr    : type.
cont     : type.
cell     : type.
cv       : type.
store    : type.
answer   : type.

% Expressions

z        : exp.
s        : exp -> exp.
case     : exp -> exp -> (exp -> exp) -> exp.
unit     : exp.
pair     : exp -> exp -> exp.
fst      : exp -> exp.
snd      : exp -> exp.
lam      : (exp -> exp) -> exp.
app      : exp -> exp -> exp.
letval   : exp -> (exp -> exp) -> exp.
letname  : exp -> (exp -> exp) -> exp.
fix      : (exp -> exp) -> exp.

rf       : cell -> exp.
ref      : exp -> exp.
!        : exp -> exp.
seq      : exp -> exp -> exp.
assign   : exp -> exp -> exp.

```

```

% Types

nat   : tp.
one   : tp.
cross : tp -> tp -> tp.
arrow : tp -> tp -> tp.

tref  : tp -> tp.

% Instructions

eval   : exp -> instr.
return : exp -> instr.
case*  : exp -> exp -> (exp -> exp) -> instr.
pair*  : exp -> exp -> instr.
fst*   : exp -> instr.
snd*   : exp -> instr.
app*   : exp -> exp -> instr.

ref*   : exp -> instr.
deref* : exp -> instr.
assign*1 : exp -> exp -> instr.
assign*2 : exp -> exp -> instr.

% Continuations

init : cont.
klam : cont -> (exp -> instr) -> cont.

% Store

estore : store.
with   : store -> cv -> store.
holds  : cell -> exp -> cv.

% Answers

close : store -> exp -> answer.
new   : (cell -> answer) -> answer.

```

A.2. TYPING

```

tpc : cell -> tp -> type.
tpe : exp -> tp -> type.
tpi : instr -> tp -> type.
tpK : cont -> tp -> tp -> type.
tpS : store -> type.
tpa : answer -> tp -> type.

% Expressions

tpe_z   : tpe z nat.
tpe_s   : tpe E nat
        -> tpe (s E) nat.
tpe_case : tpe E nat
        -> tpe E1 T

```

```

-> ({x:exp} tpe x nat -> tpe (E2 x) T)
-> tpe (case E E1 E2) T.
tpe_unit : tpe unit one.
tpe_pair : tpe E1 T1
-> tpe E2 T2
-> tpe (pair E1 E2) (cross T1 T2).
tpe_fst  : tpe E (cross T1 T2)
-> tpe (fst E) T1.
tpe_snd  : tpe E (cross T1 T2)
-> tpe (snd E) T2.
tpe_lam  : ({x:exp} tpe x T1 -> tpe (E x) T2)
-> tpe (lam E) (arrow T1 T2).
tpe_app  : tpe E1 (arrow T2 T1)
-> tpe E2 T2
-> tpe (app E1 E2) T1.
tpe_letval : tpe E1 T1
-> ({x:exp} tpe x T1 -> tpe (E2 x) T2)
-> tpe (letval E1 E2) T2.
tpe_letname : tpe (E2 E1) T
-> tpe (letname E1 E2) T.
tpe_fix   : ({x:exp} tpe x T -> tpe (E x) T)
-> tpe (fix E) T.

tpe_cell  : tpc C T
-> tpe (rf C) (tref T).
tpe_ref   : tpe E T
-> tpe (ref E) (tref T).
tpe_deref : tpe E (tref T)
-> tpe (! E) T.
tpe_seq   : tpe E1 T1
-> tpe E2 T2
-> tpe (seq E1 E2) T2.
tpe_assign : tpe E1 (tref T)
-> tpe E2 T
-> tpe (assign E1 E2) one.

% Instructions

tpe_eval  : tpe E T
-> tpe (eval E) T.
tpe_return : tpe V T
-> tpe (return V) T.
tpe_case* : tpe V nat
-> tpe E1 T
-> ({x:exp} tpe x nat -> tpe (E2 x) T)
-> tpe (case* V E1 E2) T.
tpe_pair* : tpe V T1
-> tpe E T2
-> tpe (pair* V E) (cross T1 T2).
tpe_fst*  : tpe V (cross T1 T2)
-> tpe (fst* V) T1.
tpe_snd*  : tpe V (cross T1 T2)
-> tpe (snd* V) T2.
tpe_app*  : tpe V (arrow T2 T1)
-> tpe E T2
-> tpe (app* V E) T1.

tpe_ref*  : tpe V T
-> tpe (ref* V) (tref T).
tpe_deref* : tpe V (tref T)
-> tpe (deref* V) T.
tpe_assign*1: tpe V (tref T)
-> tpe E T

```

```

        -> tpi (assign*1 V E) one.
tpi_assign*2:  tpe V1 (tref T)
              -> tpe V2 T
              -> tpi (assign*2 V1 V2) one.

% Continuations

tpK_init : tpK init T T.
tpK_lam  :  ({x:exp} tpe x T1 -> tpi (I x) T)
          -> tpK K T T2
          -> tpK (klam K I) T1 T2.

% Store

tpS_empty :  tpS estore.
tpS_with  :  tpS S
            -> tpc C T
            -> tpe V T
            -> tpS (with S (holds C V)).

% Answers

tpa_close :  tpS S
            -> tpe V T
            -> tpa (close S V) T.
tpa_new   :  ({c: cell} tpc c T' -> tpa (A c) T)
            -> tpa (new A) T.

```

A.3. EVALUATION

```

%%% Contents of a cell

contains: cell -> exp -> type.

% Only run-time assumptions

%%% Collection of all the cells in the store

collect: store -> type.

col_empty :  collect estore.
col_cv    :  contains C V
            -o collect S
            -o collect (with S (holds C V)).

%%% Reading the value of a cell from the store

read: cell -> exp -> type.

read_val  :  contains C V
            -o <T>
            -o read C V.

%%% Evaluation

ev : cont -> instr -> answer -> type.

```

% Expressions

```

ev_z      :   ev K (return z) A
           -o ev K (eval z) A.
ev_s      :   ev (klam K ([x:exp] return (s x))) (eval E) A
           -o ev K (eval (s E)) A.
ev_case   :   ev (klam K ([x:exp] case* x E1 E2)) (eval E) A
           -o ev K (eval (case E E1 E2)) A.
ev_unit   :   ev K (return unit) A
           -o ev K (eval unit) A.
ev_pair   :   ev (klam K ([x:exp] pair* x E2)) (eval E1) A
           -o ev K (eval (pair E1 E2)) A.
ev_fst    :   ev (klam K ([x:exp] fst* x)) (eval E) A
           -o ev K (eval (fst E)) A.
ev_snd    :   ev (klam K ([x:exp] snd* x)) (eval E) A
           -o ev K (eval (snd E)) A.
ev_lam    :   ev K (return (lam E)) A
           -o ev K (eval (lam E)) A.
ev_app    :   ev (klam K ([x:exp] app* x E2)) (eval E1) A
           -o ev K (eval (app E1 E2)) A.
ev_letval :   ev (klam K ([x:exp] eval (E2 x))) (eval E1) A
           -o ev K (eval (letval E1 E2)) A.
ev_letname : ev K (eval (E2 E1)) A
           -o ev K (eval (letname E1 E2)) A.
ev_fix    :   ev K (eval (E (fix E))) A
           -o ev K (eval (fix E)) A.

ev_cell   :   ev K (return (rf C)) A
           -o ev K (eval (rf C)) A.
ev_ref    :   ev (klam K ([x:exp] ref* x)) (eval E) A
           -o ev K (eval (ref E)) A.
ev_deref  :   ev (klam K ([x:exp] deref* x)) (eval E) A
           -o ev K (eval (! E)) A.
ev_seq    :   ev (klam K ([x:exp] eval E2)) (eval E1) A
           -o ev K (eval (seq E1 E2)) A.
ev_assign :   ev (klam K ([x:exp] assign*1 x E2)) (eval E1) A
           -o ev K (eval (assign E1 E2)) A.

```

% Values

```

ev_init   :   collect S
           -o ev init (return V) (close S V).
ev_cont   :   ev K (I V) A
           -o ev (klam K I) (return V) A.

```

% Auxiliary instructions

```

ev_case*1 :   ev K (eval E1) A
           -o ev K (case* z E1 E2) A.
ev_case*2 :   ev K (eval (E2 V)) A
           -o ev K (case* (s V) E1 E2) A.
ev_pair*  :   ev (klam K ([x:exp] return (pair V x))) (eval E) A
           -o ev K (pair* V E) A.
ev_fst*   :   ev K (return V1) A
           -o ev K (fst* (pair V1 V2)) A.
ev_snd*   :   ev K (return V2) A
           -o ev K (snd* (pair V1 V2)) A.
ev_app*   :   ev (klam K ([x:exp] eval (E1 x))) (eval E2) A
           -o ev K (app* (lam E1) E2) A.

```

```

ev_ref*      :  ({c:cell} contains c V -o ev K (return (rf c)) (A c))
               -o ev K (ref* V) (new A).
ev_deref*    :  (read C V
               & ev K (return V) A)
               -o ev K (deref* (rf C)) A.
ev_assign*1  :  ev (klam K ([x:exp] assign*2 (rf C) x)) (eval E) A
               -o ev K (assign*1 (rf C) E) A.
ev_assign*2  :  (contains C V -o ev K (return unit) A)
               -o (contains C V' -o ev K (assign*2 (rf C) V) A).

```

ACKNOWLEDGMENT

We would like to thank the anonymous referees for their valuable comments, which helped improve this paper as a whole.

REFERENCES

1. Samson Abramsky. Computational interpretations of linear logic. *Theoretical Computer Science*, 111:3–57, 1993.
2. Andrew Barber. Dual intuitionistic linear logic. Technical Report ECS-LFCS-96-347, Laboratory for Foundations of Computer Sciences, University of Edinburgh, 1996.
3. H. P. Barendregt. *The Lambda-Calculus: Its Syntax and Semantics*. North-Holland, 1980.
4. Henk Barendregt. Lambda calculi with types. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume II, pages 118–309. Oxford University Press, 1992.
5. Nick Benton, G. M. Bierman, J. Martin E. Hyland, and Valeria de Paiva. A term calculus for intuitionistic linear logic. In M. Bezem and J. F. Groote, editors, *Proceedings of the International Conference on Typed Lambda Calculi and Applications*, pages 75–90. Springer-Verlag LNCS 664, 1993.
6. Iliano Cervesato. *A Linear Logical Framework*. PhD thesis, Dipartimento di Informatica, Università di Torino, February 1996.
7. Iliano Cervesato. Proof-theoretic foundation of compilation in logic programming languages. In J. Jaffar, editor, *Proceedings of the 1998 Joint International Conference and Symposium on Logic Programming — JICSLP'98*, pages 115–129, Manchester, UK, 16–19 June 1998. MIT Press.
8. Iliano Cervesato, Joshua S. Hodas, and Frank Pfenning. Efficient resource management for linear logic proof search. *Theoretical Computer Science*, 232(1–2):133–163, February 2000.
9. Iliano Cervesato and Frank Pfenning. The linear logical framework *LLF*. Accessible on the World-Wide Web as <http://www.cs.cmu.edu/~fp/llf/>.
10. Iliano Cervesato and Frank Pfenning. Linear higher-order pre-unification. In Glynn Winskel, editor, *Proceedings of the Twelfth Annual Symposium on Logic in Computer Science (LICS'97)*, pages 422–433, Warsaw, Poland, June 1997. IEEE Computer Society Press.
11. Iliano Cervesato and Frank Pfenning. A linear spine calculus. Technical Report CMU-CS-97-125, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, April 1997.
12. Jawahar Lal Chirimar. *Proof Theoretic Approach to Specification Languages*. PhD thesis, University of Pennsylvania, May 1995.
13. Dominique Clément, Joëlle Despeyroux, Thierry Despeyroux, and Gilles Kahn. A simple applicative language: Mini-ML. In *Proceedings of the 1986 Conference on LISP and Functional Programming*, pages 13–27. ACM Press, 1986.
14. Thierry Coquand. An algorithm for testing conversion in type theory. In Gérard Huet and Gordon Plotkin, editors, *Logical Frameworks*, pages 255–279. Cambridge University Press, 1991.
15. Luis M.M. Damas. *Type Assignment in Programming Languages*. PhD thesis, University of Edinburgh, 1985.
16. Nachum Dershowitz and Jean-Pierre Jouannaud. *Handbook of Theoretical Computer Science*, volume B, chapter Rewrite Systems, pages 243–320. MIT Press, 1990.

17. Roy Dyckhoff. Contraction-free sequent calculi for intuitionistic logic. *Journal of Symbolic Logic*, 57(3):795–807, September 1992.
18. Amy Felty. Encoding dependent types in an intuitionistic logic. In Gérard Huet and Gordon D. Plotkin, editors, *Logical Frameworks*, pages 214–251. Cambridge University Press, 1991.
19. R.O. Gandy. Proofs of strong normalization. In J.P. Seldin and J.R. Hindley, editors, *To H.B. Curry: Essays in Combinatory Logic, Lambda Calculus and Formalism*, pages 457–477. Academic Press, 1980.
20. Herman Geuvers. *Logics and Type Systems*. PhD thesis, Katholieke Universiteit Nijmegen, 1993.
21. Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
22. Jean-Yves Girard. On the unity of logic. *Annals of Pure and Applied Logic*, 59:201–217, 1993.
23. Michael J. Gordon, Robin Milner, and Christopher P. Wadsworth. *Edinburgh LCF*. Springer-Verlag LNCS 78, 1979.
24. M.J.C. Gordon and T.F. Melham. *Introduction to HOL: a Theorem Proving Environment for Higher-order Logic*. Cambridge University Press, 1993.
25. John Hannan and Dale Miller. From operational semantics to abstract machines: Preliminary results. In M. Wand, editor, *Proceedings of the 1990 ACM Conference on Lisp and Functional Programming*, pages 323–332, Nice, France, 1990.
26. Robert Harper. A simplified account of polymorphic references. *Information Processing Letter*, 51:201–206, 1994.
27. Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the Association for Computing Machinery*, 40(1):143–184, January 1993.
28. Joshua Hodas and Dale Miller. Logic programming in a fragment of intuitionistic linear logic. *Information and Computation*, 110(2):327–365, 1994. A preliminary version appeared in the Proceedings of the Sixth Annual IEEE Symposium on Logic in Computer Science, pages 32–42, Amsterdam, The Netherlands, July 1991.
29. Joshua S. Hodas. *Logic Programming in Intuitionistic Linear Logic: Theory, Design, and Implementation*. PhD thesis, University of Pennsylvania, Department of Computer and Information Science, 1994.
30. Samin Ishtiaq and David Pym. A relevant analysis of natural deduction. *Journal of Logic and Computation*, 8(6), 1998.
31. Xavier Leroy and Pierre Weis. Polymorphic type inference and assignment. In *Conference Record of the Eighteenth Annual ACM Symposium on Principles of Programming Languages*, pages 291–302, Orlando, Florida, January 21–23 1991. ACM press.
32. Patrick Lincoln and John Mitchell. Operational aspects of linear lambda calculus. In *Seventh Annual Symposium on Logic in Computer Science*, pages 235–246, Santa Cruz, California, June 1992. IEEE Computer Society Press.
33. Spiro Michaylov and Frank Pfenning. Natural semantics and some of its meta-theory in Elf. In L.-H. Eriksson, L. Hallnäs, and P. Schroeder-Heister, editors, *Proceedings of the Second International Workshop on Extensions of Logic Programming*, pages 299–344, Stockholm, Sweden, January 1991. Springer-Verlag LNAI 596.
34. Dale Miller. A multiple-conclusion meta-logic. In S. Abramsky, editor, *Ninth Annual IEEE Symposium on Logic in Computer Science*, pages 272–281, Paris, France, July 1994.
35. Dale Miller, Gordon Plotkin, and David Pym. A relevant analysis of natural deduction. Talk given at the Workshop on Logical Frameworks, Båstad, Sweden, May 1992.
36. Robin Milner, Mads Tofte, Robert Harper, and Dave McQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.
37. Frank Pfenning. Logical frameworks. Accessible on the World-Wide Web as <http://www.cs.cmu.edu/~fp/lfs.html>.
38. Frank Pfenning. Logic programming in the LF logical framework. In Gérard Huet and Gordon Plotkin, editors, *Logical Frameworks*, pages 149–181. Cambridge University Press, 1991.
39. Frank Pfenning. Computation and deduction. Unpublished lecture notes, 277 pp. Revised May 1994, April 1996, May 1992.
40. Frank Pfenning. A proof of the Church-Rosser theorem and its representation in a logical framework. Technical Report CMU-CS-92-186, Department of Computer Science, Carnegie Mellon University, September 1992.

41. Frank Pfenning. Elf: A meta-language for deductive systems. In A. Bundy, editor, *Proceedings of the 12th International Conference on Automated Deduction*, pages 811–815, Nancy, France, June 1994. Springer-Verlag LNAI 814. System abstract.
42. Frank Pfenning. Structural cut elimination in linear logic. Technical Report CMU-CS-94-222, Department of Computer Science, Carnegie Mellon University, December 1994.
43. Frank Pfenning. Structural cut elimination. In D. Kozen, editor, *Proceedings of the Tenth Annual Symposium on Logic in Computer Science*, pages 156–166, San Diego, California, June 1995. IEEE Computer Society Press.
44. Frank Pfenning. The practice of logical frameworks. In H el ene Kirchner, editor, *Proceedings of the Colloquium on Trees in Algebra and Programming*, pages 119–134, Link oping, Sweden, April 1996. Springer-Verlag LNCS 1059. Invited talk.
45. Frank Pfenning. Structural cut elimination I. intuitionistic and classical logic. *Information and Computation*, 157(1/2):84–141, March 2000.
46. Frank Pfenning and Carsten Sch urmann. The *Twelf* project. Accessible on the World-Wide Web as <http://www.cs.cmu.edu/~twelf/>.
47. Frank Pfenning and Carsten Sch urmann. Automated theorem proving in a simple meta-logic for LF. In C. Kirchner and H. Kirchner, editors, *Proceedings of the 15th Conference on Automated Deduction — CADE-15*, pages 286–300, Lindau, Germany, July 1998. Springer-Verlag LNAI 1421.
48. Frank Pfenning and Carsten Sch urmann. System description: Twelf — a meta-logical framework for deductive systems. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, pages 202–206, Trento, Italy, July 1999. Springer-Verlag LNAI 1632.
49. Anne Salvesen. The Church-Rosser theorem for LF with $\beta\eta$ -reduction. Unpublished notes to a talk given at the First Workshop on Logical Frameworks in Antibes, France, May 1990.
50. Mads Tofte. Type inference for polymorphic references. *Information & Computation*, 89:1–34, November 1990.
51. Anne S. Troelstra. Strong normalization for typed terms with surjective pairing. *Notre Dame Journal of Formal Logic*, 27(4):547–550, October 1986.
52. Myra VanInwegen. *The Machine-Assisted Proof of Programming Language Properties*. PhD thesis, University of Pennsylvania, Department of Computer and Information Science, 1996.
53. Philip Wadler. Linear types can change the world. In M. Broy and C. B. Jones, editors, *IFIP TC 2 Working Conference on Programming Concepts and Methods*, pages 561–581, Sea of Gallilee, Israel, April 1990. North-Holland.
54. Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information & Computation*, 115(1):38–94, November 1994.