

A positive perspective on term representation: Extended paper*

Dale Miller  

Inria Saclay & LIX, Institut Polytechnique de Paris

Jui-Hsuan Wu  

LIX, Institut Polytechnique de Paris

Abstract

We use the focused proof system *LJF* as a framework for describing term structures and substitution. Since the proof theory of *LJF* does not pick a canonical polarization for primitive types, two different approaches to term representation arise. When primitive types are given the negative polarity, *LJF* proofs encode terms as tree-like structures in a familiar fashion. In this situation, cut elimination also yields the familiar notion of substitution. On the other hand, when primitive types are given the positive polarity, *LJF* proofs yield a structure in which explicit sharing of term structures is possible. Such a representation of terms provides an explicit method for sharing term structures. In this setting, cut elimination yields a different notion of substitution. We illustrate these two approaches to term representation by applying them to the encoding of untyped λ -terms. We also exploit concurrency theory techniques—namely traces and simulation—to compare untyped λ -terms using such different structuring disciplines.

2012 ACM Subject Classification Theory of computation \rightarrow Proof theory

Keywords and phrases term representation, sharing, focused proof systems

Digital Object Identifier [10.4230/LIPIcs.CSL.2023.1](https://doi.org/10.4230/LIPIcs.CSL.2023.1)

Related Version *Full Version:* <https://hal.inria.fr/hal-03843587>

1 Introduction

The structure of terms and expressions are represented variously via labeled trees or directed acyclic graphs (DAGs). When such expressions contain bindings, additional devices are needed. We follow a familiar line of research in which the design of term representations is motivated by proof-theoretic considerations. Accordingly, we rely on proof theory in the hope that it will provide careful and formal descriptions of the term structures that serve as the foundation of theorem provers, semantic specifications, interpreters, parsers, and compilers.

Applications of structural proof theory usually start with either natural deduction or sequent calculus proof systems, both of which were introduced by Gentzen in [15]. For our purposes here, these two proof systems are inadequate: natural deduction (with or without *generalized elimination rules* [36]) seems too restrictive, and sequent calculus seems too unstructured. Instead, we use the *LJF* proof system [24], which results from applying the notions of *polarity* and *focusing* [1, 17] to Gentzen’s *LJ* proof system as a framework for studying term structures with and without bindings and with and without explicit sharing constructs. By examining the dynamics of cut-elimination in *LJF*, we can also describe substitution into such term structures.

Focused proof theories have been successfully applied within the Curry-Howard correspondence: for example, call-by-value, call-by-name, and call-by-push-value evaluation strategies have been related to different choices in polarizing type expressions [34, 39, 42]. Instead of dealing with functional programming style evaluation, we address a more primitive notion: what is a *term* (used to encode programs), and what is *substitution* (used to describe various kinds of evaluation). Even this

* This technical report adds appendixes to our paper that appears in the *Proceedings of CSL 2023*.



© Dale Miller and Jui-Hsuan Wu;

licensed under Creative Commons License CC-BY 4.0

31st EACSL Annual Conference on Computer Science Logic (CSL 2023).

Editors: Bartek Klin and Elaine Pimentel; Article No. 1; pp. 1:1–1:25

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

more narrow setting has been addressed using the focused proof systems *LJT* [12, 22] and *LJQ* [11]. As in those papers, we shall consider proofs-as-terms and not proofs-as-programs since we do not consider the evaluation of such structures in this paper.

2 The LJF proof system

Gentzen’s sequent calculus proof system *LJ* [15] employs rather tiny and slippery inference rules: they are tiny since most of them deal with at most one logical connective at a time, and they are slippery since the order in which they are applied can often be freely reorganized. As a result, it is hard to use *LJ* to identify large-scale structures in proofs. For example, capturing the two phases of proof search in logic programming—*goal-reduction* and *backchaining*—with sequent calculus rules requires rather technical arguments about permutations of inference rules [30]. Andreoli’s invention of a *focused proof system* for linear logic [1] provided means for adding more structure to sequent calculus proofs for linear logic. This notion of focusing was eventually moved to intuitionistic and classical logic as the *LJF* and *LKF* proof systems [24]. This paper considers only intuitionistic logic and only for formulas built using implication as the only logical connective.

2.1 LJF inference rules

There are two kinds of sequents in *LJF*: the \uparrow -sequents $\Gamma \uparrow \Theta \vdash \Delta \uparrow \Delta'$ and the \downarrow -sequents $\Gamma \downarrow \Theta \vdash \Delta \downarrow \Delta'$. Here, all four *zones* Γ , Θ , Δ , and Δ' are multisets of formulas. Given that we are in an intuitionistic proof system, we require that the multiset union $\Delta \cup \Delta'$ is always a singleton. The zones Γ and Δ' are called the left and right *storage zones*. The zones Θ and Δ are called the left and right *staging areas*. As these names suggest, formulas in the storage zones can persist during the construction of a proof, while formulas in the staging area are intended to have a limited impact on the long-term construction of such a proof. Occurrences of logical connectives introduced by the left and right introduction rules only appear in the staging area. As we shall see, the only sequents with empty staging areas are of the form $\Gamma \uparrow \cdot \vdash \cdot \uparrow \Delta$. Such sequents are called *border sequents*: when we define synthetic inference rules in Definition 4, these sequents form the borders (the conclusion and premises) of synthetic inference rules.

Notational conventions: We usually denote an empty zone by explicitly using the dot \cdot . Also, while every *LJF* sequent has two occurrences of either \uparrow or \downarrow , we write fewer of these arrows by adopting the convention that we drop writing $\cdot \downarrow$ and $\cdot \uparrow$ when they appear on the right, and we drop writing $\downarrow \cdot$ and $\uparrow \cdot$ when they appear on the left. Finally, since the right side of sequents have exactly one formula, we replace writing $\downarrow \cdot$ with \downarrow and $\uparrow \cdot$ with \uparrow . Thus, the sequent $\Gamma \uparrow \cdot \vdash \cdot \uparrow E$ can be written as $\Gamma \vdash E$ and the sequent $\Gamma \downarrow \cdot \vdash E \downarrow \cdot$ can be written as $\Gamma \vdash E \downarrow$. As a result of these conventions, border sequents in *LJF* will resemble sequents in *LJ*, which is a completely desirable resemblance.

In the general setting, Gentzen’s *LJ* proof system involves *unpolarized* formulas, while the *LJF* focused proof system involves *polarized formulas*. Since, in this paper, we are only interested in one logical connective, the implication \supset , and since the polarization of an implication is unambiguous (it is *negative*), the only distinction between unpolarized and polarized formulas falls on atomic formulas: in polarized formulas, we need to describe the polarity of atomic formulas explicitly. (We remind the reader that a logical connective is polarized negatively if its right introduction rule is invertible.)

► **Definition 1.** An atomic bias assignment is a function, δ , that maps atomic formulas to either $+$ or $-$. A formula B is negative if it is either an implication or atomic and $\delta(B) = -$. A formula B is positive if it is atomic and $\delta(B) = +$.

$$\begin{array}{c}
\text{DECIDE, RELEASE, AND STORE RULES} \\
\frac{N, \Gamma \Downarrow N \vdash A}{N, \Gamma \vdash A} D_l \quad \frac{\Gamma \vdash P \Downarrow}{\Gamma \vdash P} D_r \quad \frac{\Gamma \Uparrow P \vdash A}{\Gamma \Downarrow P \vdash A} R_l \quad \frac{\Gamma \vdash N \Uparrow}{\Gamma \vdash N \Downarrow} R_r \quad \frac{\Gamma, C \Uparrow \Theta \vdash \Delta' \Uparrow \Delta}{\Gamma \Uparrow \Theta, C \vdash \Delta' \Uparrow \Delta} S_l \quad \frac{\Gamma \Uparrow \Theta \vdash A}{\Gamma \Uparrow \Theta \vdash A \Uparrow} S_r \\
\\
\begin{array}{cc}
\text{INITIAL RULES} & \text{INTRODUCTION RULES FOR IMPLICATION} \\
\frac{\delta(A) = +}{A, \Gamma \vdash A \Downarrow} I_r & \frac{\delta(A) = -}{\Gamma \Downarrow A \vdash A} I_l \quad \frac{\Gamma \vdash B \Downarrow \quad \Gamma \Downarrow B' \vdash A}{\Gamma \Downarrow B \supset B' \vdash A} \supset L \quad \frac{\Gamma \Uparrow \Theta, B \vdash B' \Uparrow}{\Gamma \Uparrow \Theta \vdash B \supset B' \Uparrow} \supset R
\end{array}
\end{array}$$

■ **Figure 1** The rules of (cut-free) *LJF* for the implicational fragment of propositional intuitionistic logic.

The *LJF* proof system for intuitionistic propositional logic over just implication is given in Figure 1. This figure uses the following schema variables: P is a positive formula, N is a negative formula, A is an atomic formula, and B , B' , and C denote arbitrary formulas. Note that in our simplified setting, a positive formula is also atomic.

An \Uparrow -phase is a collection of occurrences of \Uparrow -sequents that are all connected via inference rules; similarly, a \Downarrow -phase is a collection of occurrences of \Downarrow -sequents that are all connected via inference rules. The *decide* rules D_l and D_r are the only inference rules (in a cut-free proof) that have a border sequent as a conclusion. Thus, the decide rules sit on top of an \Uparrow -phase while their premises are at the bottom of a \Downarrow -phase. The *store* rules S_l and S_r work within an \Uparrow -phase. The *release* rules R_l and R_r are dual to the decide rules in that they sit on top of a \Downarrow -phase while their premises are at the bottom of an \Uparrow -phase.

When searching for a proof in *LJF*, the only occurrence of *don't know nondeterminism* occurs with the decide rules. There is a possible choice to decide on the right or left (D_l or D_r), and if D_l is selected, then another important choice is which formula in the left storage should be selected.

The following soundness and completeness theorem for *LJF* can be found in [24], where these theorems are proved for full first-order intuitionistic logic. (That paper also proves that *LJF* captures the structure of *LJT* and *LJQ* as well as some hybrid forms of those two proof systems.)

► **Theorem 2.** *Let B be an unpolarized formula composed only of implications and atomic formulas.*

1. *If B is provable in *LJ* and if $\delta(\cdot)$ is any atomic bias assignment for the atoms in B , then the sequent $\cdot \Uparrow \cdot \vdash B \Uparrow \cdot$ is provable in *LJF*.*
2. *If $\delta(\cdot)$ is an atomic bias assignment for the atoms in B and if $\cdot \Uparrow \cdot \vdash B \Uparrow \cdot$ is provable in *LJF* then B is provable in *LJ*.*

The theorem above implies that polarization of atomic formulas does not affect *provability* in *LJF*: in particular, if $\cdot \Uparrow \cdot \vdash B \Uparrow \cdot$ is provable for some atomic bias assignment then that sequent is provable for all such polarizations. On the other hand, different choices of atomic bias assignments can make a big difference in the shape and size of *LJF* proofs. We illustrate this difference in the next section. Before doing that, we define the *order of a formula* and state two technical results about the order of formulas within *LJF* proofs.

► **Definition 3** (Order of a formula). *The order of the formula B , written $\text{ord}(B)$, is defined as follows: $\text{ord}(A) = 0$ if A is atomic and $\text{ord}(B_1 \supset B_2) = \max(\text{ord}(B_1) + 1, \text{ord}(B_2))$.*

Note that if we claim that all formulas in a multiset must have an order that is less than 0, then that multiset must necessarily be empty. The following proposition is proved by a straightforward induction on the structure of formulas.

► **Proposition 1.** *Let B be a formula such that $\text{ord}(B) \leq n$. There is an \Uparrow -phase that has $\cdot \Uparrow \cdot \vdash B \Uparrow \cdot$ as its conclusion and has premises that are border sequents. Those border sequents are of the form*

$\Gamma \uparrow \cdot \vdash \cdot \uparrow A$, where A is atomic (i.e., $\text{ord}(A) = 0$) and the formulas in Γ have order less than or equal to $n - 1$.

The following proposition is proved by a simple induction on the structure of proofs.

► **Proposition 2.** Let Ξ be a (cut-free) *LJF* proof of the border sequent $\Gamma \uparrow \cdot \vdash \cdot \uparrow A$, where A is atomic and the formulas in Γ have order less than or equal to n . Then every border sequent in Ξ is of the form $\Gamma, \Delta \uparrow \cdot \vdash \cdot \uparrow A'$ where A' is an atomic formula and all formulas in Δ are of order less than or equal to $n - 2$.

2.2 Synthetic inference rules

The following definitions are based on similar definitions in [27].

► **Definition 4** (Synthetic inference rule). A left synthetic inference rule is a rule of the form

$$\frac{\Gamma_1 \vdash \Delta_1 \quad \dots \quad \Gamma_n \vdash \Delta_n}{\Gamma \vdash \Delta} B, \quad \text{This rule is justified by a} \quad \frac{\Gamma_1 \uparrow \cdot \vdash \cdot \uparrow \Delta_1 \quad \dots \quad \Gamma_n \uparrow \cdot \vdash \cdot \uparrow \Delta_n}{\Gamma \uparrow \cdot \vdash \cdot \uparrow \Delta} D_l$$

derivation of the form

Here, $B \in \Gamma$, $n \geq 0$, and within Π , a \Downarrow -sequent never occurs above an \uparrow -sequent. The structure of *LJF* proofs also forces $B \in \Gamma_i$ for all $1 \leq i \leq n$. Given our use of only implications, it is the case that there is a unique left synthetic inference rule for a given formula. The formula B can be used to name this left synthetic inference rule, and we say that this is the left synthetic inference rule for B . We can similarly define the notion of right synthetic inference rule: however, in our context where the only positive formulas that can be used in a D_r rule are atomic, the only inference rule that can be applied to the premise of the D_r rule is I_r . As a result, right synthetic inference rules in our setting have zero premises and appear only at the leaves of proofs. We often write just synthetic inference rule to mean the left variant.

► **Definition 5** (Bipole). A bipole is a (left) synthetic inference rule in which all formulas stored using the store rules within the derivation justifying this synthetic inference rule are atomic formulas.

Bipoles are, therefore, synthetic inference rules in which the only difference between the concluding sequent and any one of its premises is the presence or absence of atomic formulas. Note that, since we are only admitting implications, the synthetic rule for B is a bipole if and only if $\text{ord}(B) \leq 2$.

The primary use of the *LJF* proof system in this paper is to build large-scale, synthetic rules that we can add to the unpolarized proof system *LJ*: focusing gives us a framework to create such rules from Gentzen's micro rules. We define such an extension of *LJ* below and then illustrate it with two examples.

► **Definition 6** (Rules from polarized theory). Let \mathcal{T} be a finite set of formulas of order two or less, and let δ be an atomic bias assignment. We define $\text{LJ}[\delta, \mathcal{T}]$ to be the two-sided proof system built as follows. The only sequents in the $\text{LJ}[\delta, \mathcal{T}]$ proof system are of the form $\Gamma \vdash A$ where A is atomic and Γ is a multiset of atomic formulas. The inference rules of $\text{LJ}[\delta, \mathcal{T}]$ correspond to synthetic inference rules in the following way. For every left synthetic inference rule

$$\frac{\mathcal{T}, \Gamma_1 \vdash A_1 \quad \dots \quad \mathcal{T}, \Gamma_n \vdash A_n}{\mathcal{T}, \Gamma \vdash A} B, \quad \text{where } B \in \mathcal{T} \text{ is a negative formula, then the rule} \quad \frac{\Gamma_1 \vdash A_1 \quad \dots \quad \Gamma_n \vdash A_n}{\Gamma \vdash A} B$$

is added to $\text{LJ}[\delta, \mathcal{T}]$. The right synthetic inference rules are of the form

$$\frac{}{\Gamma \vdash A} \text{ provided } A \in \mathcal{T} \text{ and } \delta(A) = +,$$

and these are added as well. The only other inference rule added to $LJ[\delta, \mathcal{T}]$ is the rule

$$\frac{}{\Gamma, A \vdash A} \textit{init}$$

where A is atomic. If $\delta(A) = -$, this rule is justified by the D_l rule, while if $\delta(A) = +$, this rule is justified by the D_r rule.

► **Example 7.** Let $n \geq 0$ and let a_0, a_1, \dots, a_n be a sequence of distinct atomic (propositional) formulas. Let \mathcal{T} be the multiset of formulas $\{d_0, \dots, d_n\}$, where d_0 is a_0 , d_1 is $a_0 \supset a_1$, and so on until d_n is $a_0 \supset \dots \supset a_{n-1} \supset a_n$. Let $\delta^-(\cdot)$ be the atomic bias assignment that gives all atomic formulas the negative polarity and let $\delta^+(\cdot)$ be the atomic bias assignment that gives all atomic formulas the positive polarity. The inference rules in $LJ[\delta^-, \mathcal{T}]$ contains the *init* rule along with the following $n + 1$ rules

$$\frac{}{\Gamma \vdash a_0} d_0 \quad \frac{\Gamma \vdash a_0}{\Gamma \vdash a_1} d_1 \quad \frac{\Gamma \vdash a_0 \quad \Gamma \vdash a_1}{\Gamma \vdash a_2} d_2 \quad \dots \quad \frac{\Gamma \vdash a_0 \quad \dots \quad \Gamma \vdash a_{n-1}}{\Gamma \vdash a_n} d_n.$$

Given these inference rules, there is a *unique* proof of $\vdash a_n$, and that proof has 2^n occurrences of these inference rules. The negative bias assignment to atomic formulas yields the *backchaining* interpretation of these formulas. In contrast, the inference rules in $LJ[\delta^+, \mathcal{T}]$ contains the *init* rule along with the following n rules

$$\frac{\Gamma, a_0, a_1 \vdash A}{\Gamma, a_0 \vdash A} d_1 \quad \frac{\Gamma, a_0, a_1, a_2 \vdash A}{\Gamma, a_0, a_1 \vdash A} d_2 \quad \dots \quad \frac{\Gamma, a_0, \dots, a_{n-1}, a_n \vdash A}{\Gamma, a_0, \dots, a_{n-1} \vdash A} d_n.$$

Given these inference rules, it is easy to note that there are an infinite number of proofs of $a_0 \vdash a_n$, and that the shortest of these proofs contains n occurrences of synthetic inference rules plus one occurrence of the initial rule. The positive bias assignment to atomic formulas yields the *forward-chaining* interpretation of these clauses.

There are several proofs in the literature that show how cut can be eliminated within focusing proofs. Some proofs, such as the one given by Liang & Miller in [24, 26], introduce different variants of the cut rule and follow a rather tedious argument detailing how these cut rules move through individual inference rules within \uparrow and \downarrow phases. Bruscoli & Guglielmi [4] provided a different style of proof of cut elimination in a focused proof system for linear logic in which they showed how cuts could move through entire phases at a time. Other phase-based cut-elimination proofs appear in [6, 25, 35, 41, 42]. Part II of Graham-Lengrand's HdR dissertation [20] and Simmons's article [37] provide overviews of such cut-elimination proofs. While none of the proofs mentioned above deal directly with *synthetic* rules in the sense that we have defined them here, the paper [27] does have a cut-elimination theorem that directly deals with eliminating cuts in proofs with synthetic rules. The following theorem follows directly from [27, Theorem 16]. It can also be proved directly by simple induction. Let the *atomic cut rule* be the rule

$$\frac{\Gamma \vdash A \quad \Gamma, A \vdash B}{\Gamma \vdash B} \textit{Acut},$$

where A is atomic. When we add this rule to $LJ[\delta, \mathcal{T}]$, then B is also atomic.

► **Theorem 8** (Cut admissibility for $LJ[\delta, \mathcal{T}]$). *Let \mathcal{T} be a set of formulas of order two or less and let $\delta(\cdot)$ be an atomic bias assignment. The atomic cut rule is admissible for the proof system $LJ[\delta, \mathcal{T}]$.*

3 Synthetic inference rules as term constructors

Many approaches to relating proofs and terms start with the assumption that terms have been defined, say, as λ -term, and that the notion of proof is derived from terms, e.g., proofs are often identified as dependently typed λ -terms [21]. As we motivated in the introduction, we plan to reverse this influence by starting with proof structures as given and then deriving term structures from them. In particular, we shall map synthetic inference rules directly to the constructors used to build terms.

The first step in translating proofs to terms is to annotate sequents properly. We shall annotate assumptions (i.e., formulas on the left-hand side of sequents) with variables (i.e., token of type *var*) and annotate the conclusion (i.e., the formula on the right-hand side of sequents) with a term of type *tm*. The term constructor that corresponds to the *init* rule (Definition 6) is written $[\cdot]$ and has syntactic type $var \rightarrow tm$. The corresponding annotated inference rule is simply

$$\frac{}{\Gamma \vdash [x]: A} \textit{init}, \text{ provided that } x: A \in \Gamma.$$

The following example illustrates how the proofs in Example 7 can be seen as terms.

► **Example 9.** Consider changing the naming of the formulas used in Example 7 into typing assumptions, as follows:

$$d_0: a_0, \quad d_1: a_0 \supset a_1, \quad \dots, \quad d_n: a_0 \supset \dots \supset a_{n-1} \supset a_n.$$

Here, d_0, \dots, d_n are all of syntactic type *var*. Under the δ^- atomic bias assignment, the synthetic rules can be annotations as

$$\frac{\Gamma \vdash t_0: a_0 \quad \dots \quad \Gamma \vdash t_{i-1}: a_{i-1}}{\Gamma \vdash (E_i t_0 \dots t_{i-1}): a_i} d_i,$$

where $0 \leq i \leq n$. The syntactic type of constructor E_i is $tm \rightarrow \dots \rightarrow tm \rightarrow tm$ (where *tm* occurs $i + 1$ times). On the other hand, if all atomic formulas are polarized positively, then the annotated synthetic rules are

$$\frac{\Gamma, x_0: a_0, \dots, x_{i-1}: a_{i-1}, y: a_i \vdash t: A}{\Gamma, x_0: a_0, \dots, x_{i-1}: a_{i-1} \vdash (F_i x_0 \dots x_{i-1} (\lambda y.t)): A} \quad (\text{provided } y \text{ is new}),$$

where $i \geq 1$. The syntactic type of F_i ($i \geq 1$) is $var \rightarrow \dots \rightarrow var \rightarrow (var \rightarrow tm) \rightarrow tm$ (where *var* occurs $i + 1$ times). Below, we display the unique proof of a_4 using the E_n constructors and the shortest proof of a_4 using the F_n constructors: the structure on the right allows for explicit sharing of subterms while the structure on the left must repeat these subterms.

$$\begin{array}{ll} (E_4 (E_3 (E_2 (E_1 E_0) (E_1 E_0)) (E_2 (E_1 E_0) (E_1 E_0)))) & (F_1 d_0 \quad (\lambda x_1. \\ (E_3 (E_2 (E_1 E_0) (E_1 E_0)) (E_2 (E_1 E_0) (E_1 E_0)))) & (F_2 d_0 x_1 \quad (\lambda x_2. \\ & (F_3 d_0 x_1 x_2 \quad (\lambda x_3. \\ & (F_4 d_0 x_1 x_2 x_3 \quad (\lambda x_4. [x_4])))))) \end{array}$$

More generally, we can describe the structure of synthetic inference rules as follows. In order to compute the left synthetic inference rule that corresponds to using D_l on the indexed formula $f: B$, where $\text{ord}(B) \leq 2$, we must know how the atomic formulas in B are polarized. In what follows, we assume that there are exactly two atomic bias assignments, namely, δ^- and δ^+ . We consider these two cases separately. In the following discussion, we use Δ as a schematic variable to range over

multisets of atomic formulas and use \bar{w} to denote a list of variables, say, w_1, \dots, w_n for some $n \geq 0$. The notation $\bar{w} : \Delta$ denote a type assignment, say, $w_1 : A_1, \dots, w_n : A_n$, where \bar{w} is w_1, \dots, w_n and Δ is A_1, \dots, A_n . Finally, whenever we construct typing assignments for the left-hand side of sequents, we will always assume that the names assigned types are all distinct.

Consider first the case that we are using δ^- . Any formula B such that $\text{ord}(B) \leq 2$ can be written as

$$(\Delta_1 \supset A_1) \supset \dots \supset (\Delta_m \supset A_m) \supset A_0 \quad (m \geq 0)$$

where $\Delta_1, \dots, \Delta_m$ are multisets of atomic formulas. (If Δ is empty then $\Delta \supset B$ is simply B .) If the displayed formula above has order 0 then $m = 0$. If that formula has order 1 then $m \geq 1$ and $\Delta_1, \dots, \Delta_m$ are empty. If that formula has order 2 then $m \geq 1$ and at least one of $\Delta_1, \dots, \Delta_m$ is nonempty. In any case, all of the formulas $\Delta_1 \supset A_1, \dots, \Delta_m \supset A_m$ have negative polarity. In this case, the left synthetic inference rule corresponding to B can be written as

$$\frac{\Gamma, \bar{w}_1 : \Delta_1 \vdash t_1 : A_1 \quad \dots \quad \Gamma, \bar{w}_m : \Delta_m \vdash t_m : A_m}{\Gamma \vdash (f (\lambda \bar{w}_1. t_1) \dots (\lambda \bar{w}_m. t_m)) : A_0}$$

By $\bar{w} : \{D_1, \dots, D_n\}$ we mean $w_1 : D_1, \dots, w_n : D_n$, assuming that all the tokens w_1, \dots, w_n are all distinct and do not occur in Γ . If the list of variables \bar{w} is empty, then we abbreviate $\lambda \bar{w}. t$ as simply t .

Now consider the second case where we are using δ^+ . In this case, we note that any formula B such that $\text{ord}(B) \leq 2$ can be written as

$$\Delta_0 \supset (\Delta_1 \supset A_1) \supset \dots \supset (\Delta_m \supset A_m) \supset A_0 \quad (m \geq 0)$$

(We assume that all the atomic argument types are listed before the non-atomic argument types.) Here, $\Delta_0, \dots, \Delta_m$ are all multisets of atomic formulas, and we assume that $\Delta_1, \dots, \Delta_m$ are nonempty. If the displayed formula above has order 0 then $m = 0$ and Δ_0 is empty. If that formula has order 1 then $m = 0$ and Δ_0 is nonempty. If that formula has order 2 then $m \geq 1$. In this way of presenting the type of f , the (atomic) argument types in Δ_0 all have positive polarity, while the remaining argument types all have negative polarity. Thus, the left synthetic inference rule corresponding to B can be written as

$$\frac{\Gamma, \bar{w}_1 : \Delta_1 \vdash t_1 : A_1 \quad \dots \quad \Gamma, \bar{w}_m : \Delta_m \vdash t_m : A_m \quad \Gamma, y : A_0 \vdash t : A}{\Gamma \vdash \mathbf{name} \ y = (f \bar{u} (\lambda \bar{w}_1. t_1) \dots (\lambda \bar{w}_m. t_m)) \ \mathbf{in} \ t : A} \text{ provided } \bar{u} : \Delta_0 \subseteq \Gamma$$

The proviso to this rule means that every type assumption $u_i : D_i, \dots, u_n : D_n$ is a member of Γ (here, \bar{u} is u_1, \dots, u_n and Δ_0 is $\{D_1, \dots, D_n\}$). We do not assume that the variables in u_1, \dots, u_n are distinct. The first m premises above construct abstracted terms, and the last premise continues to construct a term for type A but this time with an additional variable y that names the structure $(f \bar{u} (\lambda \bar{w}_1. t_1) \dots (\lambda \bar{w}_m. t_m))$. Note also that the first arguments of f (written here as \bar{u}) must already be present on the left-hand side of the conclusion.


► **Example 10.** Using this notation, the last term in Example 9 can be written as follows.

$$\begin{aligned} \mathbf{name} \ x_1 &= (F_1 \ d_0) \ \mathbf{in} \\ \mathbf{name} \ x_2 &= (F_2 \ d_0 \ x_1) \ \mathbf{in} \\ \mathbf{name} \ x_3 &= (F_3 \ d_0 \ x_1 \ x_2) \ \mathbf{in} \\ \mathbf{name} \ x_4 &= (F_4 \ d_0 \ x_1 \ x_2 \ x_3) \ \mathbf{in} \ [x_4] \end{aligned}$$

To provide a simple, graphical illustration of how atomic bias assignments can affect term structure, consider the typing assignment $\{z : i, f : i \supset i \supset i\}$. The following two terms denote proofs built with synthetic inference rules in which i has a negative and positive polarity, respectively.

```
(f (f z z) (f z z))
name y1 = (f z z) in name y2 = (f y1 y1) in y2.
```

These terms could also be displayed as the following labeled tree and DAG. Here, we display the tree associated with a term with its root at the top and its leaves below: this is, of course, the opposite convention to how proof trees are commonly displayed.



In Appendix C, we show how these constructors—corresponding to synthetic inference rules—can be implemented using the $\lambda\kappa$ -calculus, a term representation for *LJF* introduced by Brock-Nannestad et al. in [3].

In the following subsections, we note two applications of the term representations that we have described above.

3.1 Intermediate representation of programs

The name expressions used above resemble the more common `let` expressions, but they are different in at least two important ways. First, the `let`-expression “let $x = r$ in t ” is often considered to be a non-normal term since it sometimes abbreviates the β -redex $((\lambda x.t)r)$. In contrast, the term “name $x = r$ in t ” is normal if t is normal. (Note that the synthetic inference rules are built without the use of the cut inference rule.) Second, the term r in a `let`-expression can be an arbitrary term, whereas in the name-expression, the term r must have a particular structure: when we are considering only first-order signatures, r must be the application of a constructor to one or more variables (and not to general terms).

Given an expression E using negative bias syntax, there might be a subexpression, say E' , that has many occurrences in E . If we let $F(x)$ denote the result of replacing every occurrence of E' in E with the variable x , then the expression `let $x = E'$ in $F(x)$` might be a more appropriate presentation of E in which the subformula E' is named and explicitly shared. Such an operation is often called *common subexpression elimination*. The positive bias syntax that we have described here is orthogonal to this processing in the sense that the positive bias assignment syntax forces *all* function applications (not just those that are repeated) to be named while there is no guarantee that naming is not redundant (i.e., a given subterm might be given two names). Of course, if one takes care in building terms using positive bias assignment, it is possible to build terms where common subexpressions are explicitly shared.

A useful intermediate representation of programs in compilers of functional programming languages is the *administrative normal form (ANF)* [13] in which all arguments to functions are values, that is, either constants, variables, or λ -abstractions. Clearly, when we are using positive bias assignment syntax, the expressions that result are in ANF.

Various other term representations have been developed for focused proofs that contain more logical connectives and inference rules than we have considered here. See, for example, [5, 23, 37]. However, atomic formulas are given the negative polarity in these references.

3.2 Encoding functional expressions as relational queries

When a Prolog programmer needs to compute the value of a mathematical expression, such as $\sqrt{b^2 - 4ac}$, it is necessary to explicitly convert the calls to various functions (here, subtraction, addition, multiplication, and square root) into associated relations. For example, addition on real numbers is usually represented by the three-place predicate *plus* such that the atomic formula (*plus* x y z) holds if and only if $x + y = z$. Now assume that relations are available to encode each primitive function. One way to organize the relations needed to compute the expression above involves converting that expression into positive bias syntax. For example, the function-based expression above can be written as

name $n_1 = b \times b$ **in** **name** $n_2 = 4 \times a$ **in** **name** $n_3 = n_2 \times c$ **in**
name $n_4 = n_1 - n_3$ **in** **name** $n_5 = \sqrt{n_4}$ **in** [n_5].

As described in [16], it is straightforward to convert such an expression into a series of calls to predicates. In particular, we can rewrite this expression by replacing [**name** $n = f$ $x_1 \cdots x_i$ **in** \bullet] with $[\exists n.(R_f x_1 \cdots x_i n) \wedge \bullet]$, where R_f is a relation that computes the function f . Assuming that *times*, *minus*, and *sqr*t are all relations that compute multiplication, subtraction, and the (positive) square root, then the relational presentation can be given as

$\exists n_1. \text{times } b \ b \ n_1 \wedge \exists n_2. \text{times } 4 \ a \ n_2 \wedge \exists n_3. \text{times } n_2 \ c \ n_3 \wedge \exists n_4. \text{minus } n_1 \ n_3 \ n_4 \wedge \exists n_5. \text{sqr}t \ n_4 \ n_5,$

which is an expression that is easily written as a Prolog goal formula.

4 The untyped lambda-calculus

Let D be an atomic formula and let Γ_0 be the theory $\{(D \supset D) \supset D, D \supset (D \supset D)\}$. This theory is *inconsistent* in that every formula built from implications and D is provable from Γ_0 . We choose to consider Γ_0 because cut-free proofs in *LJF* of $\Gamma_0 \vdash D$ correspond to closed untyped λ -terms. The following derivations result from applying D_l to these two formulas (assuming that $\Gamma_0 \subseteq \Gamma$).

$$\frac{\frac{\frac{\Xi_1}{\Gamma \vdash D \Downarrow} \quad \frac{\Xi_2}{\Gamma \vdash D \Downarrow} \quad \frac{\Xi_3}{\Gamma \Downarrow D \vdash D}}{\Gamma \Downarrow D \supset (D \supset D) \vdash D} \supset L \times 2 \quad \frac{\frac{\frac{\Xi_4}{\Gamma, D \vdash D} \quad \frac{\Gamma \uparrow D \vdash D \uparrow}{\Gamma \vdash D \supset D \uparrow} \quad S_l, S_r \quad \frac{\Xi_5}{\Gamma \Downarrow D \vdash D}}{\Gamma \Downarrow (D \supset D) \supset D \vdash D} R_r, \supset R}{\Gamma \vdash D} \supset L}{\Gamma \vdash D} D_l$$

These derivations can only be extended to have border sequents as premises if we reveal the polarity assigned to D . If D is polarized negatively, then Ξ_1 and Ξ_2 are both R_r while Ξ_3 and Ξ_5 are I_l . In this case, the resulting synthetic inference rules are

$$\frac{\Gamma \vdash D}{\Gamma \vdash D} \quad \text{and} \quad \frac{\Gamma, D \vdash D}{\Gamma \vdash D} .$$

On the other hand, if D is polarized positively, then Ξ_1 and Ξ_2 are both I_r while Ξ_3 and Ξ_5 is R_l . In this case, the resulting synthetic inference rules are

$$\frac{\Gamma, D, D, D \vdash D}{\Gamma, D, D \vdash D} \quad \text{and} \quad \frac{\Gamma, D \vdash D \quad \Gamma, D \vdash D}{\Gamma \vdash D} .$$

Without annotations, these inference rules are not illuminating. We provide such annotations next.

4.1 Using negative bias syntax for the untyped lambda-calculus

If we polarize D negatively, the (annotated) synthetic inference rules based on Γ_0 yield the standard tree-like representation for the untyped λ -calculus. In particular, consider an annotated *LJF* proof of a sequent of the form $\Gamma_0, x_1 : D, \dots, x_n : D \vdash t : D$. The synthetic inference rules in such a proof are either the result of deciding on $x_i : D$ (for $1 \leq i \leq n$) followed by I_l or on one of the two formulas in Γ_0 . The three inference rules annotate these three choices for using the D_l rule.

$$\frac{\Gamma \vdash t : D \quad \Gamma \vdash s : D}{\Gamma \vdash \text{napp } t s : D} \quad \frac{\Gamma, x : D \vdash t : D}{\Gamma \vdash \text{nabs } (\lambda x.t) : D} \quad \frac{}{\Gamma \vdash \text{nvar } x_i : D} \text{ provided } x_i : D \in \Gamma$$

The constructors represented by these synthetic rules are listed below, along with their syntactic types (using typing declarations from λ Prolog [29]).

```
type napp  tm -> tm -> tm.
type nabs  (var -> tm) -> tm.
type nvar  var -> tm.
```

As we have assumed before, `var` and `tm` are primitive types: annotations on the left side of sequents have type `var` while annotations on the right side have type `tm`. The prefix `n` on these names is meant to remind us that we assigned D the negative polarity. As an example of using these constructors, the untyped λ -term $((\lambda x.xx)(\lambda x.xx))$ can be written as the following term of type tm :¹

```
(napp (nabs x \ napp (nvar x) (nvar x))
      (nabs x \ napp (nvar x) (nvar x))).
```

Many computational logic systems, such as λ Prolog, Abella, LF, and Isabelle, encode the untyped λ -calculus in this fashion.

4.2 Using positive bias syntax for the untyped lambda-calculus

If we polarize D positively, we get different synthetic rules based on using Γ_0 and, hence, we get a different format for encoding the untyped λ -calculus. Again, there are exactly three synthetic inference rules for *LJF* proofs of sequents of the form $\Gamma_0, x_1 : D, \dots, x_n : D \uparrow \cdot \vdash \cdot \uparrow t : D$, but this time there are two left synthetic inference rules and one right synthetic inference rules.

$$\frac{\Gamma, y : D \vdash t : D}{\Gamma \vdash \text{papp } x_i x_j (\lambda y.t) : D} \text{ provided } \Gamma \text{ contains } x_i : D \text{ and } x_j : D$$

$$\frac{\Gamma, x : D \vdash t : D \quad \Gamma, y : D \vdash s : D}{\Gamma \vdash \text{pabs } (\lambda x.t) (\lambda y.s) : D} \quad \frac{}{\Gamma \vdash \text{pvar } x_i : D} \text{ provided } x_i : D \in \Gamma$$

The constructors represented by these synthetic rules can be given the following syntactic typing.

```
type papp  var -> var -> (var -> tm) -> tm.
type pabs  (var -> tm) -> (var -> tm) -> tm.
type pvar  var -> tm.
```

The prefix `p` on these names is meant to remind us that, in this case, we assigned D the positive polarity. Using these constructors, the untyped λ -term $((\lambda x.xx)(\lambda x.xx))$ can be written as the term

```
(pabs (x \ papp x x (y \ pvar y)) (u \ papp u u (z \ pvar z)))
```

¹ The λ -abstraction of λ Prolog is written using an infix backslash.

Note that this untyped λ -term contains two occurrences of `papp` while in the previous representation, this term contains three occurrences of `napp`.

In order to make terms in this syntax easier to read, we will often use the name-expressions presented before. In particular, the expression `(papp u v (w\ Body))`, which denotes the application of the variable `u` to the variable `v`, and then *naming* that application as `w` in the scope of `Body`. Thus, this expression can also be written using the expression `name w = (app u v) in Body`. Similarly, the expression `(pabs R (w\ Body))` can be written as `name w = (abs R) in Body`. If we used this syntax, then the expression above denoting $((\lambda x.xx)(\lambda x.xx))$ is written as

```
name u = (abs x\ name y = (app x x) in y) in
name z = (app u u) in z
```

That is, `u` is used to name the encoding for the term $(\lambda x.xx)$, and then that name is used twice in building the final application that is named `z`. Extending this example slightly, we see that the expression $((\lambda x.xx)(\lambda x.xx)(\lambda x.xx))$ can be written as

```
(pabs (x\ papp x x (y\ pvar y)) (u\ papp u u (z\ papp z u v\ pvar v)))
```

or, using the name syntax, as

```
name u = (abs x\ name y = (app x x) in y) in
name z = (app u u) in
name v = (app z u) in v.
```

As this alternative syntax suggests, the syntax that results from making the primitive type D positive makes sharing explicit by its requirement that all applications are built from *named* structures and that that application is named itself. It is also clear that the following two expressions denote the same untyped λ -term.

```
name u = (abs x\ name y = (app x x) in y) in
name z1 = (app u u) in
name z2 = (app u u) in
name v = (app z1 u) in v
```

```
name u1 = (abs x\ name y = (app x x) in y) in
name z = (app u1 u1) in
name u2 = (abs x\ name y = (app x x) in y) in
name v = (app z u2) in v
```

The first of these terms illustrates that a named structure might not be used in its scope: we call this *vacuous naming*. The second of these terms illustrates that the same structure can be named twice: we call this *redundant naming*. The proof theory behind *LJF* allows for both vacuous and redundant naming: we currently see no proof-theoretic device that can cleanly eliminate these kinds of naming expressions.

This style of syntactic representation seems relatively low-level since it uses names to designate all constructors in a term. Such a representation of terms resembles, in fact, the use of pointers to encode terms in memory: a pointer (name) indicates a unit of memory that contains the name of a constructor followed by a vector of pointers to that constructor's arguments.

4.3 Tracing untyped lambda-terms

Given that we have two different formats for untyped λ -terms, it is a natural question whether or not two such expressions denote the same untyped λ -term. For example, it seems sensible to consider the

```

type ntrace, ptrace   tm -> trace -> o.

ntrace (napp M _) (left P) :- ntrace M P.
ntrace (napp _ N) (right P) :- ntrace N P.
ntrace (nabs R) (bnd P) :- pi x\pi p\ ntrace (nvar x) p =>
                           ntrace (R x) (P p).

ptrace (papp U V K) P :-
  pi x\ (pi P\ ptrace (pvar x) (left P) :- ptrace (pvar U) P) =>
        (pi P\ ptrace (pvar x) (right P) :- ptrace (pvar V) P) =>
  ptrace (K x) P.
ptrace (pabs R K) P :-
  pi x\ (pi Q\ ptrace (pvar x) (bnd Q) :-
        pi p\ pi u\ ptrace (pvar u) p => ptrace (R u) (Q p))
  => ptrace (K x) P.

```

■ **Figure 2** Traces through negative and positive bias syntax.

last three expressions in Section 4.2 (based on positive bias assignment) as equivalent in some sense to each other and to the following expression (based on negative bias assignment).

```

(napp (napp (nabs x\ napp (nvar x) (nvar x))
         (nabs x\ napp (nvar x) (nvar x)))
 (nabs x\ napp (nvar x) (nvar x)))

```

Broadly speaking, there are two approaches to answering this question. The “white box” approach examines the actual syntax of proof expressions to see if they should be considered equal. For example, in the setting of natural deduction, two proofs are often considered equal if they reduce to the same normal form. Given that we are considering proofs built with difference sets of (synthetic) inference rules, a different approach needs to be taken, such as the approach described in [35], where proofs based on positive bias assignment to atomic formulas are systematically converted to proofs based on negative bias assignment.

Instead, we propose to use a “black box” approach in which we probe a term to describe *traces* within expressions. For example, we can ask whether or not the term denotes a top-level application or not. This check is easy to make for negative bias syntax by simply examining the top-level symbol of the expression. It is also easy to check for the expressions using the positive bias assignment: simply examine the top-level naming structure, say,

$$[\text{name } x_1 = E_1 \text{ in name } x_2 = E_2 \text{ in } \dots \text{ name } x_n = E_n \text{ in } x_j]$$

and check if E_j is an application or not. If two expressions denote an application, we can continue to develop a trace by examining either the first or second argument of that application. Similarly, we can examine two expressions to see if they denote a λ -abstraction. If they are both λ -abstractions, then we can probe the body of those abstractions, taking appropriate care when descending under a binding.

A formal specification of such trace predicates is easy in a language such as λ Prolog. In particular, the following declarations define the datatype of traces through untyped λ -terms.

```

kind trace           type.
type left, right    trace -> trace.
type bnd            (trace -> trace) -> trace.

```

The specification of traces within both variants of expressions for (closed) untyped λ -terms is given in Figure 2. Note that the orders of the clauses for `ntrace` are 0, 1, 2 while for `ptrace` the orders are 0, 3, and 4.² If we wish to treat open expressions, we can add a constant, say `w`, to denote a free variable, along with the declarations

```
type w      var.
type wtrace trace.
```

```
ntrace (nvar w) wtrace.
ptrace (pvar w) wtrace.
```

We say that two expressions denoting untyped λ -terms (using either positive or negative bias assignment) are *trace equivalent* if they both have the same traces.³ It is easy to prove that two expressions using the negative bias syntax are trace equivalent if and only if they are α -equivalent.⁴ This statement is not true for positive bias syntax: in particular, the examples at the end of Section 4.2 that illustrate vacuous and duplicate naming all have the same traces but are not α -convertible expressions.

We note that in λ Prolog, it is possible to synthesize an expression from a list of traces using, for example, a query such as

```
?- forall (ntrace T) [(bnd (u\ left (bnd (v\ left u))),
                    (bnd (u\ left (bnd (v\ right v))),
                    (bnd (u\ right u))].
T = nabs (u\ napp (nabs (v\ napp (nvar u) (nvar v))) (nvar u))
```

Here, `forall` is a higher-order predicate that applies its predicate argument (here, `(ntrace T)`) to all members in its second argument. A black box method of converting an expression using the positive bias assignment into an expression using negative bias assignment proceeds as follows: first, list all possible traces in the positive bias assignment expression, and second, synthesize the negative bias assignment expression using the technique illustrated above (see also [29, Section 7.4.2]).

4.4 Sharing bisimulation

Determining that two untyped λ -term expressions are trace equivalent by enumerating every trace in them has an exponential cost since all sharing structures are removed when listing traces. Condoluci et al. [9] developed a graphic representation of sharing in the untyped λ -calculus using *λ graphs*. When an appropriate bisimulation is defined on nodes in such *λ graphs*, it is possible to check the bisimilarity in such graphs in linear time. As is known from concurrency theory, bisimilarity implies (maximal) trace equivalence. In our setting, if two terms—represented as two nodes in a *λ graph*—are bisimilar, then those two terms are also trace equivalent.

To illustrate how one can manipulate positive bias syntax effectively, we used Abella [2] to specify a simulation relation (closely related to the bisimulation relation defined in [9]) that compares two expressions in such a way that unfolding of the sharing does not happen.

The top-level of an untyped λ -term expression based on the positive bias assignment for D can be described as a pair containing an association list of naming variables and operations (such as `app` and

² Standard techniques can be used to rewrite the last two of these clauses to clauses of order 2 at the expense of adding new predicate constants. See Appendix A for such a specification.

³ In concurrency theory, this notion is more often called *maximal trace equivalence*.

⁴ See <http://abella-prover.org/examples/lambda-calculus/term-structure/path.html> for a short, formal proof of this claim in Abella.

```

Kind op      type.
Type app     val -> val -> op.
Type abs     (val -> tm) -> op.

Kind pair    type.
Type pr      val -> op -> pair.

Kind node    type.
Type nd      list pair -> tm -> node.

```

■ **Figure 3** An Abella specification of sharing simulation

abs used above) and the name of a particular naming variable. For example, the last term displayed in Section 4.2 can be seen as a list of four pairs, namely,

$$[\langle u1, (\text{abs } x \text{ name } y = (\text{app } x \ x) \text{ in } y) \rangle, \langle z, (\text{app } u1 \ u1) \rangle, \\ \langle u2, (\text{abs } x \text{ name } y = (\text{app } x \ x) \text{ in } y) \rangle, \langle v, (\text{app } z \ u2) \rangle].$$

along with the designation of one particular variable, here v . This presentation suggests encoding such terms using the Abella declarations (which are similar to λ Prolog declarations) found in Figure 3. For example, the example term displayed above can be written as the Abella term of type name

```

(nd ((pr n4 (app n2 n3)) ::
     (pr n3 (abs x\ papp x x y\ pvar y)) ::
     (pr n2 (app n1 n1)) ::
     (pr n1 (abs x\ papp x x y\ pvar y)) :: nil)
 (pvar n4))

```

(Here, the symbols $n1, \dots, n4$ are examples of nominal constants in Abella.) In this setting, expressions using positive bias syntax are encoded as terms of type `node`.

Structures of type `node` can be used to generate a labeled transition system in which some arcs are labeled. These labels, called *actions* here, are of the following three kinds (see the full specification in Figure 4).

1. Primitive actions are described using the predicate `paction`. Such actions name a variable.
2. Bound actions, specified using `baction`, carry an abstraction node to the body of that abstraction.
3. Application actions, specified using `faction`, carry an application node to either its left or right argument: this action names that direction.

The ∇ -quantifier [14, 31] (written in Abella as `nable`) is used to manage nominal constants and binding mobility [28]. The distinction between *free action* and *bound action* here is essentially the same as has been used to specify various simulations in the π -calculus [33]. Our specification of simulation in the presence of both free and bound actions follows the specification technique of Tiu & Miller [38] that also relied on the ∇ -quantifier. Given our simulation specification, the bisimulation specification is also easy to write.

5 Cut elimination at the level of synthetic rules

Theorem 8 states that cut elimination holds for proofs built using synthetic inference rules. Our goal in this section is to use cut elimination to determine what substitution into terms should be. In particular, if we have the term t and we have the abstraction of x over term s , how do we compute the

```

Define paction : list val -> node -> val -> prop by
  paction Vs (nd C (pvar w)) w ;
  paction Vs (nd C (pvar V)) V := member V Vs.

Define baction : node -> (val -> node) -> prop by
  nabla n, baction (nd ((pr n (abs R)):: C) (pvar n))
    (u\ nd C (R u)) ;
  nabla n, baction (nd ((pr M (Op n)):: (C n)) (pvar n)) Nd :=
    nabla n, baction (nd (C n) (pvar n)) Nd.

Kind direction      type.
Type right, left   direction.

Define faction : node -> direction -> node -> prop by
  nabla n, faction (nd ((pr n (app U V)):: C) (pvar n)) right
    (nd C (pvar V)) ;
  nabla n, faction (nd ((pr n (app U V)):: C) (pvar n)) left
    (nd C (pvar U)) ;
  nabla n, faction (nd ((pr M (Op n)):: (C n)) (pvar n)) A T :=
    nabla n, faction (nd (C n) (pvar n)) A T.

Define sim : list val -> node -> node -> prop,
  simm : list val -> node -> node -> prop by
  sim Vs (nd C (papp U V K)) Nd :=
    nabla n, sim Vs (nd ((pr n (app U V)):: C) (K n)) Nd ;
  sim Vs (nd C (pabs R K)) Nd :=
    nabla n, sim Vs (nd ((pr n (abs R)):: C) (K n)) Nd ;
  sim Vs (nd D (pvar T)) (nd C (papp U V K)) :=
    nabla n, sim Vs (nd D (pvar T)) (nd ((pr n (app U V)):: C) (K n)) ;
  sim Vs (nd D (pvar T)) (nd C (pabs R K)) :=
    nabla n, sim Vs (nd D (pvar T)) (nd ((pr n (abs R)):: C) (K n)) ;
  sim Vs (nd C (pvar U)) (nd D (pvar V)) :=
    simm Vs (nd C (pvar U)) (nd D (pvar V)) ;

  simm Vs P Q :=
    (forall N,   paction Vs P N -> paction Vs Q N) /\
    (forall A R, faction P A R -> exists S, faction Q A S /\
                                     sim Vs R S) /\
    (forall R,   baction P R -> exists S,
                                     baction Q S /\
                                     nabla u, sim (u::Vs) (R u) (S u)).

```

■ **Figure 4** An Abella specification of sharing simulation

result of substituting t for x in s ? Clearly, the answer will depend on which polarity assumption we are using for primitive types.

In order to see how cut-elimination can yield substitution, consider the following instance of the cut rule: here, E and E' are atomic formulas, and the *LJF* proofs Ξ_L and Ξ_R are cut-free.

$$\frac{\Xi_L \quad \Xi_R}{\Gamma \vdash t : E' \quad \Gamma, x : E' \vdash s : E} \text{Cut}_0$$

$$\Gamma \vdash \text{Cut}_0(x.s, t) : E$$

While the term $\text{Cut}_0(x.s, t)$ is not a term, it denotes the result of substituting u for x in t . By performing cut-elimination on this proof, we will arrive at a cut-free proof and the term annotating that proof should denote the result of such a substitution.

We illustrate here how the cut-elimination procedure works on our two encodings of untyped λ -terms.⁵ In particular, we will provide specifications (using *λProlog* code) to define the predicates

```
type nsubst, psubst   tm -> (var -> tm) -> tm -> o.
```

that have the following specification: given S and T of type tm and R of type $\text{var} \rightarrow \text{tm}$ then $(\text{nsubst } T \ R \ S)$ is provable if and only if T , R , and S use the constructors napp , nabs , and nvar and S is the result of substituting T into the bound variable of R . Similarly, we wish to have the same kind of specification for $(\text{psubst } T \ R \ S)$ but with the arguments S , T , and R built using the constructors papp , pabs , and pvar .

When terms are built using negative bias syntax, the cut always moves to the right branch, which means that substitution can be defined recursively on R . Moreover, substitution is applied to all the arguments of constructors recursively. Thus, we have the following *λProlog* specification of nsubst .

```
nsubst T (x \ nvar x) T.
nsubst T (x \ nvar Y) (nvar Y).
nsubst T (x \ napp (R x) (S x)) (napp R' S') :-
  nsubst T R R', nsubst T S S'.
nsubst T (x \ nabs y \ R x y) (nabs y \ R' y) :-
  pi y \ nsubst T (x \ R x y) (R' y).
```

Note that this substitution predicate moves recursively through its second (abstracted) argument. This style of substitution is, of course, the familiar one. Using the following functional equations, we can write this substitution operation as a postfix operator.

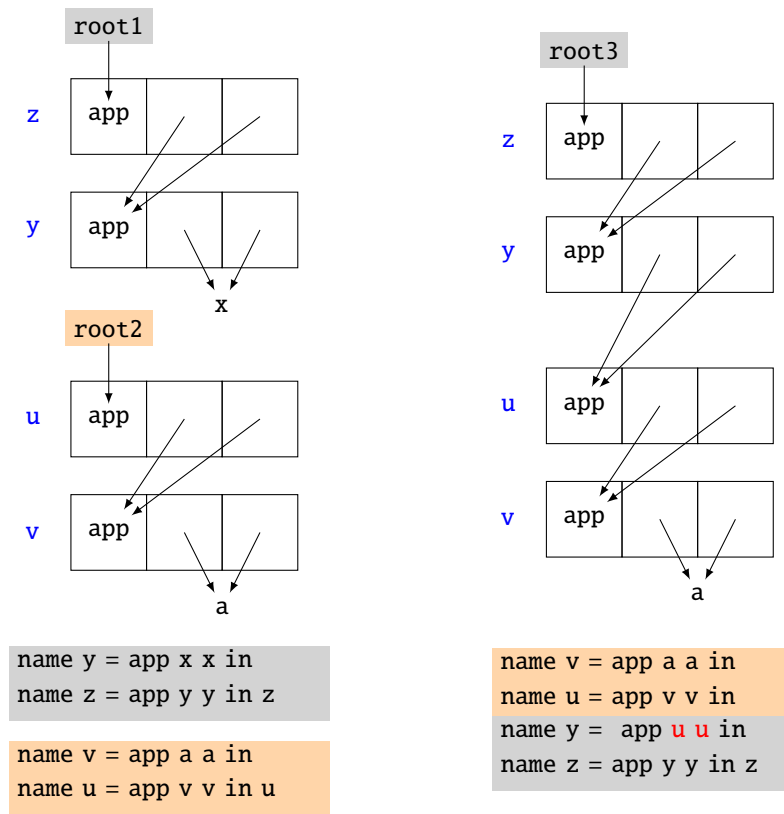
- $(\text{nvar } x)[x/t] = t$;
- $(\text{nvar } y)[x/t] = (\text{nvar } y)$, provided x and y are different;
- $(\text{napp } R \ S)[x/t] = (\text{napp } R[x/t] \ S[x/t])$, and
- $(\text{nabs } (\lambda y.R))[x/t] = (\text{nabs } \lambda y. (R[x/t]))$, provided that x and y are different and y is not free in t .

When terms are built using the positive polarity, the cut moves to the left branch, which means that the substitution can be defined recursively on the first argument.

```
psubst (papp U V K) R (papp U V H) :- pi x \ psubst (K x) R (H x).
psubst (pabs S K)   R (pabs S H)   :- pi x \ psubst (K x) R (H x).
psubst (pvar U)     R (R U).
```

Note that the last line of this specification uses a meta-level β -reduction but only to effect a variable renaming substitution. An example query using this last predicate is the following.

⁵ The cut-elimination procedure is given in Appendix B.



■ **Figure 5** The term at root3 is the result of substituting the term at root2 for x into the term at root1. Operationally speaking, the pointers to x are redirected to point to u instead.

```
?- psubst (papp w w y\ papp y y z\ papp z z v\ pvar v)
      (x\ papp x x pvar) R.
R = papp w w (y\ papp y y (z\ papp z z (v\ papp v v u\ pvar u)))
```

Another example of substitution using this style of syntax is given graphically (and using the name notation) in Figure 5.

We can instead write this substitution operation as a prefix operator using the following functional equations. Below, the operation $t[x := u]$ denotes the replacement of every free occurrence of x in t by the variable u , provided that no free occurrence of x is in the scope of a binding on u .

- $[(\text{papp } u \ v \ (\lambda y.K))/x]R = (\text{papp } u \ v \ (\lambda y.[K/x]R))$, provided x and y are different;
- $[(\text{pabs } S \ (\lambda y.K))/x]R = (\text{pabs } S \ (\lambda y.[K/x]R))$, provided x and y are different; and
- $[(\text{pvar } u)/x]R = R[x := u]$, provided no free occurrence of x is in the scope of a binding on u .

Given our discussion of checking the simulation of two untyped λ -terms in Section 4.4, we know that such terms can be represented as a pair, say, $\langle \Gamma, u \rangle$ where Γ is an association list between a variable (of type var) and an operation (of type op) that indicates the kind of node that variable names (either an application or an abstraction). An equivalent description for substitution can then be given as follow: The result of substituting the term $\langle \Gamma, u \rangle$ for x in the term $\langle \Gamma', u' \rangle$ is the term $\langle (\Gamma'[x := u]) \sqcup \Gamma, u' \rangle$, where \sqcup denotes appending of two lists.

6 Related and future work

One goal of developing a logical framework, such as *LJF*, is to account for many other calculi within that framework. The literature contains other approaches to term representation that have also been motivated by focused proof systems [11, 12, 23]. Since *LJF* can easily account for several other focusing proof systems for intuitionistic logic [24], this framework should similarly account for such term calculi. Other frameworks have been used to justify term structures: for example, it would be interesting to see if there are any overlaps with the terms-as-graphs work of Grabmayer [19].

While many constructors for building term structures are only second-order, it is natural and occasionally important to be able to treat constructors of order greater than 2. Of course, the proof theory of *LJF* can treat formulas of all orders. However, the notion of synthetic inference rule (which is here limited to second order) would need to be generalized. In a setting with such higher-order constructors, cut elimination should be able to derive and generalize *hereditary substitutions* [40].

As mentioned in Section 4.4, the λ graphs in [9] represent sharing as DAG structures in the untyped λ -calculus. We conjecture that by using a multifocused version of the *LJF* proof system, we should be able to prove that maximal multifocused *LJF* proofs correspond to λ graphs. Maximal multifocused proofs have been shown elsewhere to correspond to graphical proof systems, such as proof nets [8], expansion proofs [7], and natural deduction proofs [35].

The black box methods of probing the structure of terms with sharing (see Sections 4.3 and 4.4) is closely related to well-established results in concurrency theory. However, these methods are not based on proof-theoretic principles, at least not that we have established here. We hope to find a way to describe trace equivalence and bisimilarity via proof-theoretic concepts. These notions seem related to Girard's Ludics project [18]. One promising approach might start with recognizing that simulations can be seen as winning strategies and that winning strategies can be related to focused proofs [10].

Acknowledgments: We thank Beniamino Accattoli and Kaustuv Chaudhuri for their valuable discussions and suggestions. We also thank anonymous reviewers for their comments on an earlier draft of this paper.

References

- 1 Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *J. of Logic and Computation*, 2(3):297–347, 1992. doi:10.1093/logcom/2.3.297.
- 2 David Baelde, Kaustuv Chaudhuri, Andrew Gacek, Dale Miller, Gopalan Nadathur, Alwen Tiu, and Yuting Wang. Abella: A system for reasoning about relational specifications. *Journal of Formalized Reasoning*, 7(2):1–89, 2014. doi:10.6092/issn.1972-5787/4650.
- 3 Taus Brock-Nannestad, Nicolas Guenot, and Daniel Gustafsson. Computation in focused intuitionistic logic. In Moreno Falaschi and Elvira Albert, editors, *Proceedings of the 17th International Symposium on Principles and Practice of Declarative Programming, Siena, Italy, July 14–16, 2015*, pages 43–54. ACM, 2015. doi:10.1145/2790449.2790528.
- 4 Paola Bruscoli and Alessio Guglielmi. On structuring proof search for first order linear logic. *Theoretical Computer Science*, 360(1-3):42–76, 2006. doi:10.1016/j.tcs.2005.11.047.
- 5 Iliano Cervesato and Frank Pfenning. A linear spine calculus. *Journal of Logic and Computation*, 13(5):639–688, 2003. doi:10.1093/logcom/13.5.639.
- 6 Kaustuv Chaudhuri. Focusing strategies in the sequent calculus of synthetic connectives. In Iliano Cervesato, Helmut Veith, and Andrei Voronkov, editors, *LPAR: International Conference on Logic, Programming, Artificial Intelligence and Reasoning*, volume 5330 of *LNCS*, pages 467–481. Springer, November 2008. doi:10.1007/978-3-540-89439-1_33.
- 7 Kaustuv Chaudhuri, Stefan Hetzl, and Dale Miller. A multi-focused proof system isomorphic to expansion proofs. *J. of Logic and Computation*, 26(2):577–603, 2016. doi:10.1093/logcom/exu030.

- 8 Kaustuv Chaudhuri, Dale Miller, and Alexis Saurin. Canonical sequent proofs via multi-focusing. In G. Ausiello, J. Karhumäki, G. Mauri, and L. Ong, editors, *Fifth International Conference on Theoretical Computer Science*, volume 273 of *IFIP*, pages 383–396. Springer, September 2008. doi:10.1007/978-0-387-09680-3_26.
- 9 Andrea Condoluci, Beniamino Accattoli, and Claudio Sacerdoti Coen. Sharing equality is linear. In *Proceedings of the 21st International Symposium on Principles and Practice of Declarative Programming*, pages 1–14, 2019. doi:10.1145/3354166.3354174.
- 10 Olivier Delandé, Dale Miller, and Alexis Saurin. Proof and refutation in MALL as a game. *Annals of Pure and Applied Logic*, 161(5):654–672, February 2010. doi:10.1016/j.apal.2009.07.017.
- 11 Roy Dyckhoff and Stéphane Lengrand. Call-by-value λ -calculus and LJQ. *J. of Logic and Computation*, 17(6):1109–1134, 2007. doi:10.1093/logcom/exm037.
- 12 José Espírito Santo. Completing herbelin’s programme. In Simona Ronchi Della Rocca, editor, *Typed Lambda Calculi and Applications, 8th International Conference, TLCA 2007, Paris, France, June 26-28, 2007, Proceedings*, volume 4583 of *LNCS*, pages 118–132. Springer, 2007. doi:10.1007/978-3-540-73228-0_10.
- 13 Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. The essence of compiling with continuations. *ACM SIGPLAN Notices*, 28(6):237–247, 1993. doi:10.1145/155090.155113.
- 14 Andrew Gacek, Dale Miller, and Gopalan Nadathur. Nominal abstraction. *Information and Computation*, 209(1):48–73, 2011. doi:10.1016/j.ic.2010.09.004.
- 15 Gerhard Gentzen. Investigations into logical deduction. In M. E. Szabo, editor, *The Collected Papers of Gerhard Gentzen*, pages 68–131. North-Holland, Amsterdam, 1935. Translation of articles that appeared in 1934–35. Collected papers appeared in 1969. doi:10.1007/BF01201353.
- 16 Ulysse Gérard and Dale Miller. Separating functional computation from relations. In Valentin Goranko and Mads Dam, editors, *26th EACSL Annual Conference on Computer Science Logic (CSL 2017)*, volume 82 of *LIPICs*, pages 23:1–23:17, 2017. doi:10.4230/LIPICs.CSL.2017.23.
- 17 Jean-Yves Girard. A new constructive logic: classical logic. *Math. Structures in Comp. Science*, 1:255–296, 1991. doi:10.1017/S0960129500001328.
- 18 Jean-Yves Girard. Locus solum: From the rules of logic to the logic of rules. *Mathematical Structures in Computer Science*, 11(3):301–506, June 2001. doi:10.1017/S096012950100336X.
- 19 Clemens Grabmayer. Modeling terms by graphs with structure constraints (two illustrations). In Maribel Fernández and Ian Mackie, editors, *TERMGRAPH@FSCD*, volume 288 of *EPTCS*, pages 1–13, 2018. doi:10.48550/arXiv.1902.02010.
- 20 Stéphane Graham-Lengrand. Polarities and focussing: a journey from realisability to automated reasoning, December 2014. Habilitation à diriger des recherches. URL: <https://tel.archives-ouvertes.fr/tel-01094980>.
- 21 Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993. doi:10.1145/138027.138060.
- 22 Hugo Herbelin. A lambda-calculus structure isomorphic to Gentzen-style sequent calculus structure. In *Computer Science Logic, 8th International Workshop, CSL ’94*, volume 933 of *Lecture Notes in Computer Science*, pages 61–75. Springer, 1995. doi:10.1007/BFb0022247.
- 23 Hugo Herbelin. *Séquents qu’on calcule: de l’interprétation du calcul des séquents comme calcul de lambda-termes et comme calcul de stratégies gagnantes*. PhD thesis, Université Paris 7, 1995. URL: <https://tel.archives-ouvertes.fr/tel-00382528>.
- 24 Chuck Liang and Dale Miller. Focusing and polarization in linear, intuitionistic, and classical logics. *Theoretical Computer Science*, 410(46):4747–4768, 2009. Abstract Interpretation and Logic Programming: In honor of professor Giorgio Levi. doi:10.1016/j.tcs.2009.07.041.
- 25 Chuck Liang and Dale Miller. A focused approach to combining logics. *Annals of Pure and Applied Logic*, 162(9):679–697, 2011. doi:10.1016/j.apal.2011.01.012.
- 26 Chuck Liang and Dale Miller. Focusing Gentzen’s LK proof system. In Thomas Piecha and Kai Wehmeier, editors, *Peter Schroeder-Heister on Proof-Theoretic Semantics*, Outstanding Contributions to Logic. Springer, 2022. To appear. URL: <https://hal.archives-ouvertes.fr/hal-03457379>.

- 27 Sonia Marin, Dale Miller, Elaine Pimentel, and Marco Volpe. From axioms to synthetic inference rules via focusing. *Annals of Pure and Applied Logic*, 173(5):1–32, 2022. doi:10.1016/j.apal.2022.103091.
- 28 Dale Miller. Mechanized metatheory revisited. *Journal of Automated Reasoning*, 63(3):625–665, October 2019. doi:10.1007/s10817-018-9483-3.
- 29 Dale Miller and Gopalan Nadathur. *Programming with Higher-Order Logic*. Cambridge University Press, June 2012. doi:10.1017/CB09781139021326.
- 30 Dale Miller, Gopalan Nadathur, Frank Pfenning, and Andre Scedrov. Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic*, 51(1–2):125–157, 1991. doi:10.1016/0168-0072(91)90068-W.
- 31 Dale Miller and Alwen Tiu. A proof theory for generic judgments. *ACM Trans. on Computational Logic*, 6(4):749–783, October 2005. doi:10.1145/1094622.1094628.
- 32 Dale Miller and Jui-Hsuan Wu. A positive perspective on term representations: Extended paper. Technical report, Inria, 2022. URL: <https://hal.inria.fr/hal-03843587>.
- 33 Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, Part II. *Information and Computation*, 100(1):41–77, 1992. doi:10.1016/0890-5401(92)90009-5.
- 34 Guillaume Munch-Maccagnoni and Gabriel Scherer. Polarised intermediate representation of lambda calculus with sums. In *30th Symp. on Logic in Computer Science*, pages 127–140. IEEE Computer Society, 2015. doi:10.1109/LICS.2015.22.
- 35 Elaine Pimentel, Vivek Nigam, and João Neto. Multi-focused proofs with different polarity assignments. In Mario Benevides and Rene Thiemann, editors, *Proc. of the Tenth Workshop on Logical and Semantic Frameworks, with Applications (LSFA 2015)*, volume 323 of *ENTCS*, pages 163–179, July 2016. doi:10.1016/j.entcs.2016.06.011.
- 36 Jan von Plato. Natural deduction with general elimination rules. *Archive for Mathematical Logic*, 40(7):541–567, 2001. doi:10.1007/s001530100091.
- 37 Robert J. Simmons. Structural focalization. *ACM Trans. on Computational Logic*, 15(3):21, 2014. doi:10.1145/2629678.
- 38 Alwen Tiu and Dale Miller. Proof search specifications of bisimulation and modal logics for the π -calculus. *ACM Trans. on Computational Logic*, 11(2):13:1–13:35, 2010. doi:10.1145/1656242.1656248.
- 39 Philip Wadler. Call-by-value is dual to call-by-name. In *8th Int. Conf. on Functional Programming*, pages 189–201, New York, NY, 2003. ACM. doi:10.1145/944705.944723.
- 40 Keven Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. A concurrent logical framework I: The propositional fragment. In *Post-proceedings of TYPES 2003 Workshop*, number 3085 in *LNCS*. Springer, 2003. doi:10.1007/978-3-540-24849-1_23.
- 41 Noam Zeilberger. Focusing and higher-order abstract syntax. In George C. Necula and Philip Wadler, editors, *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, pages 359–369. ACM, 2008. doi:10.1145/1328897.1328482.
- 42 Noam Zeilberger. On the unity of duality. *Annals of Pure and Applied Logic*, 153(1), 2008. Special issue on classical logic and computation. doi:10.1016/j.apal.2008.01.001.

A A second-order specification of pptrace

The following λ Prolog specification of the predicate `pptrace` describes traces in expressions with positive bias assignment to D : that is, it is a reimplementaion of the `ptrace` predicate specified in Section 4.3. Note that this specification involves clauses of only order two and contains three additional auxiliary predicate constants.

```

type pptrace          tm -> trace -> o.
type pppleft, pppright var -> var -> o.
type pppabs          var -> (var -> tm) -> o.

pptrace (papp U V K) P :- pi v \ pppleft v U =>
                          pppright v V => pptrace (K v) P.
pptrace (pabs R K)    P :- pi v \ pppabs v R => pptrace (K v) P.

pptrace (pvar V) (left P) :- pppleft V U, pptrace (pvar U) P.
pptrace (pvar V) (right P) :- pppright V U, pptrace (pvar U) P.
pptrace (pvar V) (bnd Q) :- pppabs V R,
                          pi p \ pi u \ pptrace (pvar u) p => pptrace (R u) (Q p).

```

B The cut-elimination procedure for atomic formulas

While several cut rules can be written for *LJF* [3, 24], we only need to consider cuts for atomic formulas among border sequents in order to formally derive the substitution operators `nsbst` and `psbst` in Section 5.

We shall show the elimination of atomic cut formulas at the level of synthetic connectives. In order to help describe that elimination procedure, we note the following observations about the structure of a cut-free *LJF* proof, say Ξ , of a border sequent of the form $\Gamma \vdash A$, where A is an atomic formula. The proof Ξ ends in either a right or left synthetic inference rule. If that last rule is a right synthetic rule, then A must be polarized positively and $A \in \Gamma$. If that last rule is a left synthetic rule, then there is $B_1 \supset \cdots \supset B_n \supset B_0$ in Γ that is chosen for the focus. (If $n = 0$, we write the expression $B_1 \supset \cdots \supset B_n \supset B_0$ as simply B_0 : thus, this formula is an atomic formula if and only if $n = 0$.) The structure of this synthetic inference rule depends, in part, on the polarity of B_0 . If B_0 has negative polarity, then Ξ ends in a synthetic inference rule of the form

$$\frac{\Xi_1 \quad \cdots \quad \Xi_m}{\Gamma, \Gamma_1 \vdash P_1 \quad \cdots \quad \Gamma, \Gamma_m \vdash P_m} \Gamma \vdash B_0,$$

where $0 \leq m \leq n$. On the other hand, if B_0 has positive polarity, then $n \geq 1$ and Ξ ends in a synthetic inference rule of the form

$$\frac{\Xi_1 \quad \cdots \quad \Xi_m \quad \Xi_0}{\Gamma, \Gamma_1 \vdash P_1 \quad \cdots \quad \Gamma, \Gamma_m \vdash P_m \quad \Gamma, B_0 \vdash A} \Gamma \vdash A$$

where $0 \leq m \leq n$.

We now describe how the cut rule can be permuted upwards in a proof. Let Π be the proof

$$\frac{\Pi_L \quad \Pi_R}{\Gamma \vdash E' \quad \Gamma, E' \vdash E} \text{Cut}$$

where both Π_L and Π_R are cut-free. Let $|\Pi|$ be the height of the cut-free proof Π : this is the maximum number of occurrences of synthetic inference rules in a path from an initial rule to the endsequent in Π . The *measure* of the above occurrence of *Cut* is $|\Pi_L| + |\Pi_R|$.

We distinguish the following two cases, depending on whether or not the last rule in Π_L is D_r or D_l . The first case is immediate. If Π_L ends with D_r/I_r , then Π_R is the result of this step (once the redundant occurrences of E' in the context are removed). Otherwise, assume that Π_L ends with D_l on $B_1 \supset \dots \supset B_n \supset B_0$. This proof ends in one of two different synthetic rules, depending on the polarity assigned to B_0 .

1. If B_0 is polarized positively, then Π_L has the following structure.

$$\frac{\frac{\Pi_1}{\Gamma, \Gamma_1 \vdash P_1} \quad \dots \quad \frac{\Pi_m}{\Gamma, \Gamma_m \vdash P_m} \quad \frac{\Pi_0}{\Gamma, B_0 \vdash E'}}{\Gamma \vdash E'}$$

The *Cut* rule can be pushed to the following position.

$$\frac{\frac{\Pi_1}{\Gamma, \Gamma_1 \vdash P_1} \quad \dots \quad \frac{\Pi_m}{\Gamma, \Gamma_m \vdash P_m} \quad \frac{\frac{\Pi_0}{\Gamma, B_0 \vdash E'} \quad \frac{\Pi_R}{\Gamma, E' \vdash E}}{\Gamma, B_0 \vdash E'}}{\Gamma \vdash E'} \text{Cut}$$

Note that the measure of this occurrence of *Cut* is strictly less than the original measure since $|\Pi_0| + |\Pi_R| < |\Pi_L| + |\Pi_R|$.

2. B_0 is polarized negatively. Thus, $B_0 = E'$ and Π_L have the following structure.

$$\frac{\frac{\Pi_1}{\Gamma, \Gamma_1 \vdash P_1} \quad \dots \quad \frac{\Pi_m}{\Gamma, \Gamma_m \vdash P_m}}{\Gamma \vdash B_0}$$

In this case, we must now consider the three possible last inference rules of Π_R .

- a. If Π_R ends with a right synthetic inference rule, then, since E must be positive and since $B_0 = E'$ is negative, it is the case that $E \in \Gamma$. We can then eliminate the cut by discarding Π_L and using the proof that consists of just Π_R .
- b. If Π_R ends with D_l on the cut formula E' , then $E = E'$, and we eliminate the cut by replacing it with Π_L .
- c. If Π_R ends with D_l on $D_1 \supset \dots \supset D_n \supset D_0 \in \Gamma$, for some formulas D_1, \dots, D_n ($n \geq 0$), and atomic formula D_0 . Then the structure of Π_R is

$$\frac{\frac{\Pi_1}{\Gamma, E', \Gamma_1 \vdash P_1} \quad \dots \quad \frac{\Pi_p}{\Gamma, E', \Gamma_p \vdash P_p} *}{\Gamma, E' \vdash E}$$

Here, the number and structure of the premises' border sequents depend on the polarity of the formulas D_0, \dots, D_n (we do not single out the particular polarity of D_0). In every case, however, the cut rule can be moved to all premises, yielding the following proof

$$\frac{\frac{\frac{\Pi_L}{\Gamma \vdash E'} \quad \frac{\Pi_1}{\Gamma, E', \Gamma_1 \vdash P_1}}{\Gamma, \Gamma_1 \vdash P_1} \text{Cut} \quad \dots \quad \frac{\frac{\Pi_L}{\Gamma \vdash E'} \quad \frac{\Pi_p}{\Gamma, E', \Gamma_p \vdash P_p}}{\Gamma, \Gamma_p \vdash P_p} \text{Cut}}{\Gamma \vdash E} *$$

Here, the synthetic inference rules marked by $*$ in the two preceding proof figures are the same. Note also that for all $1 \leq i \leq p$, $|\Pi_i| < |\Pi_R|$ so the measure of each of the p occurrences of *Cut* in this proof is strictly less than the measure of the original occurrence of *Cut*.

Let LJF^+ be the proof system that results from adding the *Cut* rule to LJF . We can now complete the argument that we can remove atomic cuts from proofs in LJF^+ to get a (cut-free) proof in LJF . Let Ξ be an LJF^+ proof. If there is an occurrence of a *Cut* in Ξ , then pick such an occurrence as high as possible. Thus the premises of that occurrence of *Cut* are cut-free proofs. Applying the procedure above, this occurrence of *Cut* can be removed with possibly several occurrences of *Cut*, all with strictly smaller measures. By recursively applying this procedure on all these occurrences of *Cut*, we can successfully replace this one occurrence of *Cut* with a cut-free proof. In this way, we have eliminated one occurrence of *Cut*. Repeating this process will leave us with a cut-free proof.

C A term calculus for LJF

A term calculus for LJF , called the $\lambda\kappa$ -calculus, was introduced by Brock-Nannestad et al. in [3] in order to study how cut-elimination and polarization can be used to describe different evaluation strategies. The portion of their calculus that corresponds to cut-free LJF proofs contains three syntactic categories and seven constructors: these are displayed below, along with their declarations using λ Prolog syntax.

Terms: $t, u ::= \lambda x.t \mid x \hat{\ } k \mid [p]$
 Values: $p, q ::= x \mid [t]$
 Continuations: $k ::= \epsilon \mid p :: k \mid \kappa x.t$

```
kind tm, var, cont                                type.
type abs      (var -> tm) -> tm.
type app      var -> cont -> tm.
type cl       var -> tm.
type fl       tm -> var.
type epsilon                                     cont.
type cns      var -> cont -> cont.
type kappa    (var -> tm) -> cont.
```

As is clear from that λ Prolog specification, we are using meta-level binding to capture $\lambda\kappa$ -calculus bindings. The expression $\kappa(\lambda x.t)$ is abbreviated as $\kappa x.t$. The term constructor written using the boldface λ is of type $(var \rightarrow tm) \rightarrow tm$: the expression $\lambda x.t$ is encoded as `(abs x\ t)` in λ Prolog. At the same time, we are using meta-level application to capture the application of constructors. For example, the term constructor written as $x \hat{\ } k$ is encoded using two meta-level applications `((app x) k)`. Note that the syntactic category for *values* contains *variables* as well as expressions of the form $[t]$: given that we are restricting ourselves to second-order types only, this second kind of value occurs only within the construction of synthetic inference rules. For this reason, we use the primitive type `var` (for variables) to denote values since this seems to be a more accurate typing for the constructors associated to synthetic inference rules.

Figure 6 contains an annotated version of the LJF proof system. Since the order in which \uparrow rules are applied is not important, we have merged the R_l and S_l rules into one rule, and we have restricted the S_r rule so that it is applied only when the left staging area is empty.

Since synthetic inference rules are collections of LJF inference rules, synthetic inference rules can be annotated by collections of $\lambda\kappa$ -calculus constructors, which we shall refer to here as *combinators*. The syntactic types of combinators will contain `var` and `tm` but not `cont`: this latter type is used only *internally* to construct synthetic inference rules.

N.B.: We encode syntactic categories as type expressions (built using primitive types and the arrow constructor). From the point of view of LJF , however, type expressions are formally encoded as propositional formulas (built using atomic formulas and implications). Thus, we find it convenient to conflate the notions of syntactic categories and types, as well as primitive types and atomic formulas.

To illustrate how the $\lambda\kappa$ -calculus can be used to build combinators, we provide the $\lambda\kappa$ -calculus specifications of some of the constructors used to encode synthetic inference rules we have seen in

DECIDE, RELEASE, AND STORE RULES

$$\frac{x : N, \Gamma \Downarrow N \vdash k : A}{x : N, \Gamma \vdash x \hat{\ } k : A} D_l \quad \frac{\Gamma \vdash p : P \Downarrow}{\Gamma \vdash [p] : P} D_r \quad \frac{\Gamma \vdash t : N \Uparrow}{\Gamma \vdash [t] : N \Downarrow} R_r$$

$$\frac{\Gamma, x : P \vdash t : A}{\Gamma \Downarrow P \vdash \kappa x.t : A} R_l/S_l \quad \frac{\Gamma \vdash t : A}{\Gamma \vdash t : A \Uparrow} S_r$$

INITIAL RULES

$$\frac{\delta(A) = +}{\Gamma, x : A \vdash x : A \Downarrow} I_r \quad \frac{\delta(A) = -}{\Gamma \Downarrow A \vdash \epsilon : A} I_l$$

INTRODUCTION RULES FOR IMPLICATION

$$\frac{\Gamma \vdash p : B \Downarrow \quad \Gamma \Downarrow B' \vdash k : A}{\Gamma \Downarrow B \supset B' \vdash p :: k : A} \supset L \quad \frac{\Gamma, x : B \vdash t : B' \Uparrow}{\Gamma \vdash \lambda x.t : B \supset B' \Uparrow} \supset R/S_l$$

■ **Figure 6** $\lambda\kappa$ -terms annotating some *LJF* inference rules

Sections 3 and 4. Recall the annotated theory from Example 9,

$$d_1 : a_0 \supset a_1, \quad \dots, \quad d_n : a_0 \supset \dots \supset a_{n-1} \supset a_n,$$

where $n \geq 0$. If all atomic formulas are polarized negatively, then the annotated synthetic rules are of the form

$$\frac{\Gamma \vdash t_0 : a_0 \quad \dots \quad \Gamma \vdash t_{n-1} : a_{n-1}}{\Gamma \vdash (E_n t_0 \dots t_{n-1}) : a_n},$$

where the combinators corresponding to these rules are given by

$$E_0 = d_0 \hat{\ } \epsilon, \quad E_1 = \lambda x_0. d_1 \hat{\ } ([x_0] :: \epsilon), \quad \dots, \quad E_n = \lambda x_0 \dots \lambda x_{n-1}. d_n \hat{\ } ([x_0] :: \dots :: [x_{n-1}] :: \epsilon).$$

The syntactic type of E_n ($n \geq 1$) is $tm \rightarrow \dots \rightarrow tm \rightarrow tm$ (where tm occurs $n + 1$ times). On the other hand, if all atomic formulas are polarized positively, then the annotated synthetic rules are

$$\frac{\Gamma, x_0 : a_0, \dots, x_{n-1} : a_{n-1}, y : a_n \vdash t : A}{\Gamma, x_0 : a_0, \dots, x_{n-1} : a_{n-1} \vdash (F_n x_0 \dots x_{n-1} (\lambda y.t)) : A} \quad (\text{provided } y \text{ is new}),$$

where $n \geq 1$ and the combinators corresponding to these rules are given by

$$F_1 = \lambda x_0 \lambda k. d_1 \hat{\ } (x_0 :: \kappa k), \quad \dots, \quad F_n = \lambda x_0 \dots \lambda x_{n-1} \lambda k. d_n \hat{\ } (x_0 :: \dots :: x_{n-1} :: \kappa k).$$

Since the expression $\kappa x.t$ is an abbreviation for $\kappa(\lambda x.t)$, when we write above $\dots \lambda k. \dots (\kappa k) \dots$, the bound variable k has syntactic type $var \rightarrow tm$. The syntactic type of F_n ($n \geq 1$) is $var \rightarrow \dots \rightarrow var \rightarrow (var \rightarrow tm) \rightarrow tm$ (where var occurs $n + 1$ times).

In order to use the $\lambda\kappa$ -calculus to encode the untyped λ -calculus, we take the *LJF* proofs in Section 4 and annotate their border sequents as

$$\Psi : (D \rightarrow D) \rightarrow D, \quad \Phi : D \rightarrow (D \rightarrow D), \quad x_1 : D, \quad \dots, \quad x_n : D \vdash t : D,$$

where Ψ and Φ are two $\lambda\kappa$ -calculus variables of type var . In this case, the term t represents an untyped λ -term with free variables among x_1, \dots, x_n .

N.B.: There are two notions of typing occurring here. The symbol, say, Ψ has a type within the $\lambda\kappa$ -calculus setting of var , whereas in our annotated sequent setting, Ψ is also *associated* with the formula $D \supset (D \supset D)$: for convenience, that association is written using the familiar syntax of typing.

$$\begin{array}{c}
\frac{\frac{\Gamma \vdash t : D}{\Gamma \vdash [t] : D \Downarrow} R_r \quad \frac{\Gamma \vdash s : D}{\Gamma \vdash [s] : D \Downarrow} R_r \quad \frac{}{\Gamma \Downarrow D \vdash \epsilon : D} I_l}{\frac{\Gamma \Downarrow D \supset D \supset D \vdash [t] :: [s] :: \epsilon : D}{\Gamma \vdash \Phi^\wedge([t] :: [s] :: \epsilon) : D} D_l} \supset L \times 2 \\
\frac{\frac{\Gamma, x : D \vdash t : D}{\Gamma, x : D \vdash t : D \Uparrow} S_r \quad \frac{}{\Gamma \vdash \lambda x.t : D \supset D \Uparrow} \supset R/S_l \quad \frac{\Gamma \vdash [\lambda x.t] : D \supset D \Downarrow}{\Gamma \Downarrow (D \supset D) \supset D \vdash [\lambda x.t] :: \epsilon : D} R_r \quad \frac{}{\Gamma \Downarrow D \vdash \epsilon : D} I_l}{\frac{\Gamma \Downarrow (D \supset D) \supset D \vdash [\lambda x.t] :: \epsilon : D}{\Gamma \vdash \Psi^\wedge([\lambda x.t] :: \epsilon) : D} D_l} \supset L \quad \frac{}{\frac{\Gamma \Downarrow D \vdash \epsilon : D}{\Gamma \vdash x_i \frown \epsilon : D} D_l} I_l
\end{array}$$

■ **Figure 7** The three derivations that justify the three synthetic rules that arise when D has negative polarity.

$$\begin{array}{c}
\frac{\frac{}{\Gamma \vdash x_i : D \Downarrow} I_l \quad \frac{}{\Gamma \vdash x_j : D \Downarrow} I_l \quad \frac{\Gamma, y : D \vdash t : D}{\Gamma \Downarrow D \vdash \kappa y.t : D} R_l/S_l}{\frac{\Gamma \Downarrow D \supset D \supset D \vdash x_i :: x_j :: \kappa y.t : D}{\Gamma \vdash \Phi^\wedge(x_i :: x_j :: \kappa y.t) : D} D_l} \supset L \times 2 \\
\frac{\frac{\Gamma, x : D \vdash t : D}{\Gamma, x : D \vdash t : D \Uparrow} S_r \quad \frac{}{\Gamma \vdash \lambda x.t : D \supset D \Uparrow} \supset R/S_l \quad \frac{\Gamma, y : D \vdash s : D}{\Gamma \Downarrow D \vdash \kappa y.s : D} R_l/S_l}{\frac{\Gamma \Downarrow (D \supset D) \supset D \vdash [\lambda x.t] :: \kappa y.s : D}{\Gamma \vdash \Psi^\wedge([\lambda x.t] :: \kappa y.s) : D} D_l} \supset L \quad \frac{}{\frac{\Gamma \vdash x_i : D \Downarrow}{\Gamma \vdash [x_i] : D} D_r} I_r
\end{array}$$

■ **Figure 8** The three derivations that justify the three synthetic rules that arise when D has positive polarity.

If the primitive type (atomic formula) D is polarized negatively, then the three synthetic inference rules of such border sequents are displayed in Figure 7. Thus, the following three combinators are implemented by the following $\lambda\kappa$ -calculus terms.

type napp $\text{tm} \rightarrow \text{tm} \rightarrow \text{tm}.$ $\lambda t \lambda s. \Phi^\wedge([t] :: [s] :: \epsilon)$
type nabs $(\text{var} \rightarrow \text{tm}) \rightarrow \text{tm}.$ $\lambda r. \Psi^\wedge([\lambda r] :: \epsilon)$
type nvar $\text{var} \rightarrow \text{tm}.$ $\lambda x. x \frown \epsilon$

On the other hand, if D is polarized positively, then the three synthetic inference rules of such border sequents are displayed in Figure 8. Thus, the following three combinators are implemented by the following $\lambda\kappa$ -calculus terms.

type papp $\text{var} \rightarrow \text{var} \rightarrow (\text{var} \rightarrow \text{tm}) \rightarrow \text{tm}.$ $\lambda u \lambda v \lambda k. \Phi^\wedge(u :: v :: \kappa k)$
type pabs $(\text{var} \rightarrow \text{tm}) \rightarrow (\text{var} \rightarrow \text{tm}) \rightarrow \text{tm}.$ $\lambda r \lambda k. \Psi^\wedge([\lambda r] :: \kappa k)$
type pvar $\text{var} \rightarrow \text{tm}.$ $\lambda x. [x]$