Section 1: Introduction

Most theorem proving paradigms for classical logic are centered around *ad hoc* proof structures which are designed to support a particular search procedure. Proof structures, such as resolution refutations or connection graphs, are not intended to be first-class *values*: they are very large, implementation dependent structures which are generally discarded after their discovery. This is very unfortunate since there is much information which could automatically be extract from such proofs. Such theorem provers are, for example, incapable of rendering natural justifications of their proofs to a human reader.

One obvious solution to this situation is to represent proofs by natural deduction or sequential proof trees. Such proof structures have been extensively studied and many structural manipulations are known. There are, however, at least two drawbacks to using such proof trees in a classical logic setting. First, such trees are also very large and awkward objects because they contain far more explicit information than is generally of interest. For example, natural deduction proofs record the order in which every logical connective and quantifier is introduced and eliminated. Secondly, Herbrand's Theorem states that it is substitution which is the key element in classical proofs; logical connectives play a secondary and simplier role. Hence, it should be possible to greatly simplify representation of proofs in classical logic by simply recording the role substitutions play in building proofs.

In this paper, we present just such a representation for classical proofs, called *expansion tree proofs*. These proof structures record in a very compact form the essential information, namely substitutions, of classical proofs. We feel that these structures make suitable *values* within computational settings, and we demonstrate this by presenting several computations which can be performed directly on them. In particular, we show how to convert expansion tree proofs to H-proofs [8] (derivations from tautologies using universal and existential generalization), cut-free sequential proofs [7, 13, 16], and linear reasoning [5]. In the latter system, when interpolants exist, a very simple computation on expansion tree proofs will produce them. Finally, since many classical logic proof systems are designed to use Skolem terms to simplify the role of quantifiers in proofs, we present a version of expansion trees which use Skolem terms. We show that these two version are equivalent by presenting the transformations between them.

For traditional theorem proving systems, the conversion of expansion tree proofs to sequential proofs is very valuable. In particular, if a given resolution-style theorem prover built an expansion tree from its resolution refutation, the transformation to the "natural" proof structures described in Section 4 would provide a means by which a human readable presentation of a resolution refutation could be generated. Just such a practical use of expansion tree proofs has been demonstrated in [6, 11].

Since substitutions are central to understanding classical proofs and since the λ-calculus is an elegant formalism for representing substitutions, we have choose to use

a version of classical logic which is based on Church's Simple Theory of Types [3]. This logic can represent quantification at all finite types, and, hence, all the results of this paper are valid for higher-order logic as well as for first-order logic. Furthermore, we have been able to provide two results for this higher-order logic which have not appeared before in the literature: a strengthened form of the (first-order) interpolation theorem, and a correct description of Skolem functions and the Herbrand Universe.

Section 2: Logical Preliminaries

The higher-order logic, $\mathcal{T}$, which we shall consider here is essentially the Simple Theory of Types described by Church in [3], except that we do not use the axioms of extensionality, choice, descriptions, or infinity. $\mathcal{T}$ contains two base types, o for boolean and ι for individuals, although adding any number of other base types can easily be done. All other types are functional types, *i.e.* the type $(\beta\alpha)$ is the type of a function with domain type α and codomain type β. Such types are often written elsewhere as $\alpha \to \beta$. The type $(o\alpha)$, being the type of a function from type α to booleans, *i.e.* a characteristic function, is used in $\mathcal{T}$ to represent the type for sets and predicates of elements of type α. Formulas are built up from logical constants, variables, and parameters (nonlogical constants) by λ-abstraction and function application. Hence, the type of $[\lambda x_\alpha A_\beta]$ (where x_α is a variable) is $(\beta\alpha)$ while the type for $[A_{(\beta\alpha)} B_\alpha]$ is β. (Here, type subscripts provide for type assignments.) We shall seldom adorn formulas with type symbols, but rather, when the type of a formula, say A, cannot be determined from context, we will add the phrase "where A is a formula$_\alpha$ (variable$_\alpha$) (constant$_\alpha$)" to indicate that A has type α. When we do provides types as subscripts within larger formulas, we shall only provide an explicit type for the first occurrence of a variable or constant — we shall assume that all other occurrences will be implicitly typed the same. The logical constants of $\mathcal{T}$ are $\sim_{oo}$ (negation), $\vee_{(oo)o}$ (disjunction), and, for each type α, $\forall_{o(o\alpha)}$ (the "universal α-type set recognizer"). We also use the following abbreviations: $A \wedge B$ stands for $\sim[\sim A \vee \sim B]$, $A \supset B$ stands for $\sim A \vee B$, $\forall x\, P$ stands for $\forall[\lambda x P]$, and $\exists x\, P$ stands for $\sim\forall[\lambda x \sim P]$. Since the type of a set is of the form $(o\alpha)$, we write $L_{o\alpha} x_\alpha$ to denote the set-theoretic expression $x \in L$.

We shall present a few simple facts about λ-conversion. The reader is referred to [3] and [4] for more details. If x is a variable$_\alpha$ and t is a formula$_\alpha$, $\underset{\cdot}{\mathbf{S}}^x_t A$ denotes the formula which is the result of replacing all free occurrences of x in A with t. We shall assume that bound variable names are systematically changed to avoid variable capture. The operation of replacing a subformula of A of the form $[\lambda x C]E$ with $\underset{\cdot}{\mathbf{S}}^x_E C$ is called λ-*contraction*. We write A red B if B is the result of zero or more alphabetic changes in bound variables and λ-contractions of A. The converse of λ-contraction is λ-*expansion*. We write A conv B and say that A is λ-convertible to B if B is the result of zero or more alphabetic changes in bound variables, λ-contractions, and λ-expansions. A formula is in λ-*normal form* if it contains no *contractible part*, *i.e.* a subformula of the form $[\lambda x C]E$. For every typed

λ-calculus formula A there is a formula B in λ-normal form such that A red B. By *principle normal formula* we shall mean a formula B such that for each subformula $\lambda x C$ of B, x is the first variable in alphabetical order which is not free in C. Clearly, for every formula A there is a unique principle normal formula B such that A red B. The left-most, non-bracket symbol of the principle normal form of a formula A is called the *head* of A.

Definition 2.1. Let A be a λ-normal formula$_o$. An occurrence of a subformula B in A is a *boolean subformula occurrence* if it is in the scope of only $\sim$ and $\vee$, or if A is B. A boolean subformula occurrence is either positive or negative, depending on whether it is in the scope of an even or odd number of occurrences of $\sim$. The head of A is either $\sim$, $\vee$, $\forall$, or a variable or parameter. A is an *atom* if its head is a variable or parameter, and a *boolean atom* (*b-atom*, for short) if its head is a variable, parameter, or $\forall$. ∎

Definition 2.2. Let B be a boolean atom occurrence in the formula$_o$ A. If the head of B is not a $\forall$, then we say that B is *neutral*. Otherwise, we say that B is *existential* if it is in the scope of an odd number of negations and *universal* if it is in the scope of an even number of negations. ∎

Below we list the axioms and rules of inference for the logical calculus $\mathcal{T}$. First the axioms:

(1) All propositional tautologies.

(2) $\forall_{o(o\alpha)} f_{o\alpha} \supset f x_\alpha$

(3) $\forall x_\alpha\, [p \vee f_{o\alpha} x] \supset p \vee \forall_{o(o\alpha)} f$

The rules of inference are the following:

(1) *λ-conversion*: From A to infer B provided that A conv B.

(2) *Substitution:* From $F_{o\alpha} x_\alpha$ to infer $F_{o\alpha} A_\alpha$ provided that x_α is not a free variable of $F_{o\alpha}$.

(3) *Modus Ponens:* From $A \supset B$ and A to infer B.

(4) *Generalization:* From $F_{o\alpha} x_\alpha$ to infer $\forall_{o(o\alpha)} F_{o\alpha}$, provided that x_α is not a free variable of $F_{o\alpha}$.

Those axioms and rules of inference which contain the type variable α are considered schemata. We say that a formula$_o$, A, is a theorem of $\mathcal{T}$, written $\vdash_{\mathcal{T}} A$, if there is a list of formulas$_o$, $A_1, \ldots, A_n = A$ $(n \geq 1)$ such that for each i, $1 \leq i \leq n$, A_i is either an axiom or is derived from one or two previous formulas by a rule of inference. The deduction theorem holds for $\mathcal{T}$.

The axioms and inferences of $\mathcal{T}$ are different from those for a first-order system only in the inclusion of λ-conversion and the richer structure of formulas. These two extensions, however, substantially change the character of inferences. Consider the following example of the interaction between λ-contractions and the logical connectives in higher-order logic.

Let Y be a variable$_{o\iota}$, and let D and T be variables$_{o(o\iota)}$. If we are given the formula

$$\forall D\,[DY \supset TY]$$

and we wished to do a universal instantiation of this formula with the term $\lambda Z_{o\iota}[TZ \wedge \forall x_\iota\ [Zx \supset Ax]]$, *i.e.* the term representing the set of all sets of individuals which are members of T and are subsets of A, we would then have

$$[\lambda Z[TZ \wedge \forall x\ [Zx \supset Ax]]]Y \supset TY.$$

A λ-normal form of this formula is then

$$[TY \wedge \forall x\ [Yx \supset Ax]] \supset TY.$$

The structure of this last formula is much more complex than that of the formula from which it was deduced, since it contains occurrences of logical connectives and quantifiers which are not present in the original formula. Also, Y now has the role of a predicate where this was not the case in the first formula. None of these structural changes can occur in first-order logic.

Also consider the problems of extending Herbrand's Fundamental Theorem [8] to higher-order logic. In first-order logic, Skolem functions and prenex normal forms can be used to reduce theoremhood for any formula to theoremhood of a formula of the form $\exists x_1\ \ldots \exists x_n\ A(x_1, \ldots, x_n)$ where A contains no quantifiers. Herbrand's Theorem then states that this formula is a theorem if and only if there is a collection of n-tuples $(t_1^1, \ldots, t_n^1), \ldots, (t_1^m, \ldots, t_n^m)$ of substitution terms from the Herbrand Universe, such that the *compound Herbrand instance* $A(t_1^1, \ldots, t_n^1) \vee \ldots \vee A(t_1^m, \ldots, t_n^m)$ is tautologous. In the higher-order setting, there are two important problems with this formulation of Herbrand's Theorem. The first is that the Herbrand Universe of terms can not be constructed as simply a free term algebra. The interaction between λ-abstraction and Skolem functions must be restricted. We avoid using Skolem functions entirely until this issue is clarified in Section 6.

The second, more difficult problem in extending the first-order version of Herbrand's Theorem is that in higher-order logic, the result of substituting terms for quantified variables in a given formula may yield new quantifiers which were not in the original formula, and hence have not been provided with a substitution term. For example, $A(t_1^i, \ldots, t_n^i)$ may contain quantifiers even if $A(x_1, \ldots, x_n)$ did not. Some of these quantifiers could also contribute new Skolem functions to the Herbrand Universe which might be required in some other substitution term, say t_1^j for $j \neq i$. None of these possibilities are anticipated by the first-order version of this theorem.

This form of Herbrand's Theorem fails to generalize to higher-order logic because it relies on two normal forms — prenex normal form and Skolem normal form — neither of which are preserved under substitution and λ-contraction. Our generalization of compound Herbrand instances, called expansion trees, will abandon these normal forms.

Section 3: Expansion Tree Proofs

In this section, structures which generalize formulas are defined. These structures, called expansion trees, may contain logical connectives as well as the new connective $+^t$, where t is a formula. Informally, an expression of the kind $\forall B +^{t_1} Q_1 +^{t_2} \ldots +^{t_n} Q_n$, for $n \geq 0$, represents the result of instantiating the quantified expression $\forall B$ with the terms t_i to get the structure Q_i, $i = 1, \ldots, n$. Such an expression represents a labeled, ordered tree in which the root, labeled with the formula $\forall B$, has n out-arcs, each labeled in left-to-right fashion with the formulas $t_1, \ldots, t_n$, which connect the root to the structures $Q_1, \ldots, Q_n$. Such structures which contain no labeled arcs are interpreted as formulas$_o$.

Definition 3.1. *Expansion trees*, *dual expansion trees*, *selected variables*, *expansion terms*, and two functions Sh and Dp (for *shallow* and *deep*, resp.) which both map expansion trees and dual expansion trees to formulas are defined by the following mutual recursion.

(1) Let A be a λ-normal b-atom. If A is of the form $\forall B$ then A is a dual expansion tree, otherwise it is both an expansion tree and a dual expansion tree. In either case, $Sh(A) := A$ and $Dp(A) := A$.

(2) If Q is an expansion tree, then $\sim Q$ is a dual expansion tree. If Q is a dual expansion tree, then $\sim Q$ is an expansion tree. In either case, $Sh(\sim Q) := \sim Sh(Q)$ and $Dp(\sim Q) := \sim Dp(Q)$.

(3) Assume that Q_1 and Q_2 do not share selected variables. If Q_1 and Q_2 are expansion trees then so is $Q_1 \vee Q_2$. If Q_1 and Q_2 are dual expansion trees then so is $Q_1 \vee Q_2$. In either case $Sh(Q_1 \vee Q_2) := Sh(Q_1) \vee Sh(Q_2)$ and $Dp(Q_1 \vee Q_2) := Dp(Q_1) \vee Dp(Q_2)$.

(4) If Q is an expansion tree and $Sh(Q)$ is a λ-normal form of By for some λ-normal formula$_{o\alpha}$ B and some variable$_\alpha$ y which is not selected in Q, then $Q' := \forall B +^y Q$ is an expansion tree. The variable y is a *selected variable* of Q'. Also, $Sh(Q') := \forall B$ and $Dp(Q') := Dp(Q)$.

(5) Let B be a λ-normal formula$_{o\alpha}$ and let $t_1, \ldots, t_n$ be a list of λ-normal formulas$_\alpha$ $(n \geq 1)$. If $Q_1, \ldots, Q_n$ is a list of dual expansion tree where no variable occurs selected more than once in all these trees, and for each $i = 1, \ldots n$, $Sh(Q_i)$ is a λ-normal form of Bt_i, then $Q' := \forall B +^{t_1} Q_1 +^{t_2} \ldots +^{t_n} Q_n$ is a dual expansion tree. The formulas $t_1, \ldots, t_n$ are *expansion terms* of Q'. Also, $Sh(Q') := \forall B$ and $Dp(Q') := Dp(Q_1) \wedge \ldots \wedge Dp(Q_n)$. ∎

It is very natural to think of both expansion trees and dual expansion trees as finite, ordered trees in which non-terminal nodes are labeled with either $\sim$, $\vee$, or $\forall B$, for some formula$_{o\alpha}$, B, and where terminal nodes are labeled with b-atoms. A non-terminal node labeled with a formula of the form $\forall B$ will also have out-arcs labeled with either occurrences of expansion terms or selected variables. These trees are viewed with their roots at the top

and their leaves (terminal nodes) at the bottom. A node *dominates* another node if both are on a common branch and the first node is *higher* in the tree than the second. Arcs can dominate other arcs in the same fashion. This dominance relation will be considered reflexive. Also, an arc dominates a node if the node which terminates the arc dominates the given node. In particular, an arc dominates the node in which it terminates. Just as it is possible to classify subformula$_o$ occurrences of formulas$_o$ as either positive or negative, nodes in an expansion tree are classified as either positive if it is dominated by an even number of occurrence of $\sim$, or negative otherwise.

Definition 3.2. Let Q be either an expansion tree or a dual expansion tree, and let N be a node in Q which is labeled with $\forall B$ for some formula$_{o\alpha}$ B. If Q is an expansion tree (dual expansion tree) then N is *universal* if it occurs positively (negatively) in Q and *existential* if it occurs negatively (positively) in Q.

In either an expansion tree or a dual expansion tree, a non-terminal node N labeled with a formula of the form $\forall B$ is an *instantiated node* which is instantiated by the formulas labeling its out-arcs. It is easy to see that universal nodes are instantiated by selected variable and that existential nodes can only be instantiated by expansion terms. In the first case, the node N is *selected by* the selected variable, and in the second case N is *expanded by* the expansion terms. A universal (existential) node which is not dominated by any other universal or existential node is called a *top-level universal (existential) node.* A labeled arc is a *top-level labeled arc* if it is not dominated by any other labeled arc. ∎

Definition 3.3. Let Q be an expansion tree or a dual expansion tree. $\mathbf{S}_Q$ is the set of all selected variables in Q and Θ_Q is the set of *occurrences of expansion terms* in Q. Let x be a variable$_\alpha$ which is not selected in Q, and t a formula$_\alpha$. $\mathbf{S}^x_t Q$ is defined as the result of applying $\mathbf{S}^x_t$ to (and then normalizing) all formulas labeling either arcs or nodes in Q. $\mathbf{S}^x_t Q$ is an expansion tree (dual expansion) if Q is an expansion tree (dual expansion tree), and $Sh(\mathbf{S}^x_t Q)$ conv $\mathbf{S}^x_t Sh(Q)$ and $Dp(\mathbf{S}^x_t Q)$ conv $\mathbf{S}^x_t Dp(Q)$. Finally, a variable is *new to* Q if it has no occurrence in any formula which is a label in Q. ∎

The following relation is needed for our definition of Herbrand instances since Skolem functions are not being used.

Definition 3.4. Let A be an expansion tree and let $<^0_Q$ be the binary relation on Θ_Q such that $t <^0_Q s$ if there exists a variable which is selected for a node dominated by t and which is free in s. $<_Q$, the transitive closure of $<^0_Q$, is called the *dependency relation* for Q. ∎

ET-proofs, which are our generalization of tautologous, compound, Herbrand instances, is defined below.

Definition 3.5. An expansion tree or dual expansion tree is *sound* if the free variables of $Sh(Q)$ are not selected in Q. An expansion tree Q is an expansion tree *for* A if $Sh(Q)$ is a

λ-normal form of A and Q is sound. An expansion tree is *grounded* if none of its terminal nodes are labeled with formulas of the form $\forall B$. An *ET-proof* is an expansion tree Q such that $Dp(Q)$ is tautologous and $<_Q$ is acyclic. An ET-proof is a *grounded ET-proof* if it is also a grounded expansion tree. ∎

Sound expansion trees are those trees which are expansion trees *for* some formula. In particular, if Q is a sound expansion tree, then Q is an expansion tree for $Sh(Q)$.

Example 3.6. Let A be the theorem $\exists y\, \forall x\, [Px \supset Py]$. A grounded ET-proof for A would be the tree Q given as:

$$[\exists y\, \forall x\, [Px \supset Py] +^u [\forall x\, [Px \supset Pu] +^v [Pv \supset Pu]]$$
$$+^v [\forall x\, [Px \supset Pv] +^w [Pw \supset Pv]]].$$

Here, $Dp(Q) = [Pv \supset Pu] \vee [Pw \supset Pv]$, $\Theta_Q = \{u, v\}$, and $\mathbf{S}_Q = \{v, w\}$. The dependency relation is given by the pair $u <_Q v$. If u was used in place of w, $<_Q$ would have been cyclic. ∎

The following proposition is straightforward.

Proposition 3.7. *Let $\mathcal{B}$ be a finite set of variables. If A has an ET-proof, then it has an ET-proof in which no selected variable is a member of $\mathcal{B}$.*

A common way to present Herbrand's Theorem is to show that any theorem can be proved in a system which contains tautologies as axioms and universal and existential generalizations as the sole inference rules [5, 8]. In particular, let H be the following proof system. The axioms of H are instances of tautologies. The inference rules are the following: anti-prenexing, existential and universal generalization, λ-expansion, and existential contraction (*i.e.* replace a negatively occurring subformula $\forall B \wedge \forall B$ with $\forall B$). The proof of the following theorem shows how to translate expansion tree proofs into H-proofs.

Theorem 3.8. *If the formula$_o$ A has an ET-proof then it has an H-proof.*

ET-proofs can be seen as directly encoding an H-proof. To prove this requires the following definitions and lemmas, many of which will also be used in the following section.

Definition 3.9. A term t is *admissible* in Q if no variable free in t is contained in $\mathbf{S}_Q$. Let Q be either an expansion tree or a dual expansion tree, and let N be a top-level, instantiated node in Q. A labeled out-arc of N can be *eliminated* in one of the following two ways.

(1) If N is a universal node, then it is the root of a subtree of Q of the form $\forall B +^y Q_1$. The tree which results by replacing this subtree by Q_1 is the *result of eliminating y from Q.*

(2) If N is an existential node, then it is the root of a subtree $Q_0 := \forall B +^{t_1} Q_1 + \ldots +^{t_n} Q_n$. If $n = 1$ and t_1 is admissible in Q, then let Q' be the result of

replacing Q_0 with Q_1. If $n > 1$ and for some i, $1 \leq i \leq n$, t_i is admissible in Q, then let Q' be the result of replacing Q_0 with the tree

$$[\forall B +^{t_1} Q_1 + \ldots +^{t_{i-1}} Q_{i-1} +^{t_{i+1}} Q_{i+1} + \ldots +^{t_n} Q_n] \wedge Q_i.$$

If in the first case $i := 1$, then in either case, Q' is the *result of eliminating* t_i *from* Q. ∎

Lemma 3.10. *If Q' is the result of eliminating a top-level labeled arc from the expansion tree Q then*

(1) *if $<_Q$ is acyclic, then so is $<_{Q'}$,*
(2) *$Dp(Q')$ and $Dp(Q)$ are truth-functionally equivalent,*
(3) *if Q is sound then Q' is sound, and*
(4) *if Q is sound then $Sh(Q)$ can be derived from $Sh(Q')$ by some combination of rule from the proof system H.*

Proof. Proofs of (1) and (2) are straightforward. To prove (3), assume that Q is sound. If Q' arises by eliminating a selected variable $y \in \mathbf{S}_Q$, then Q' must also be sound, since the selected variable y, which may now be free in $Sh(Q')$, is not selected in Q'. Otherwise, assume Q' arises by eliminating an admissible expansion term $t \in \Theta_Q$ from Q. $Sh(Q')$ can be formed by replacing an existential b-atom $\forall B$ with a λ-normal form of either $\forall B \wedge Bt$ or Bt. Assume that Q' is not sound. Then there must be some $z \in \mathbf{S}_{Q'} = \mathbf{S}_Q$ which is free in $Sh(Q')$. But then z is free in t, which contradicts the fact that t was admissible. Hence, Q' is sound. Thus, if Q is an ET-proof then so is Q'.

Let Q be sound, so by (3), Q' is sound. If Q' is the result of eliminating a top-level selected variable, then $Sh(Q)$ follows from $Sh(Q')$ by universal generalization, λ-expansion and anti-prenexing. If Q' is the result of eliminating a top-level admissible, expansion term, then $Sh(Q)$ can be derived from $Sh(Q')$ by using λ-expansion, existential generalization, anti-prenexing, and, possibly, existential contraction. Hence, (4) is proved. □

Lemma 3.11. *If the expansion tree Q has a labeled arc and $<_Q$ is acyclic, then some top-level labeled arc can be eliminated.*

Proof. If Q has a top-level selected variable, then this arc can be eliminated. Assume that Q has no top-level instantiated universal nodes. Let $t_1, \ldots, t_m$ be the list of all the occurrences of top-level expansion terms of Q. Since $<_Q$ is acyclic, this list must contains a $<_Q$-minimal element. It is easy to verify that such minimal elements are admissible and can, therefore, be eliminated. □

The proof of Theorem 3.8 can now be completed. Let Q be an ET-proof for A. The preceding lemmas guarantee the existence of a list of ET-proofs, $Q_1, \ldots, Q_m$, such that $Q_1 = Q$, Q_m contains no labeled arcs, and for $1 \leq i < m$, Q_{i+1} is the result of eliminating a top-level, labeled arc from Q_i. Clearly, Q_m $(= Sh(Q_m))$ is actually a tautologous formula

and therefore has a (one line) H-proof. By induction and Lemma 3.10 (4), $Sh(Q_0)$ has an H-proof. Since A is derivable from $Sh(Q_0)$ by λ-expansion, A has an H-derivation.

This last theorem demonstrates how information in an ET-proof can be eliminated to yield a simpler ET-proof. It is also possible to take two expansion trees and combine their information to make a single expansion tree which collects their information. Such combining is needed in the next section.

Definition 3.12. Let Q_1 and Q_2 be both either grounded expansion trees or grounded dual expansion trees such that $Sh(Q_1)$ and $Sh(Q_2)$ are equal up to alphabetic changes of bound variables. Assume that selected variables which label arcs not dominated by expansion terms are the same in both trees. All other selected variables of Q_1 and Q_2 will be assumed to be different. In each step of the following recursive definition of *merging* Q_1 and Q_2 to get Q_3, the following facts are easily verified.

(a) If Q_1 and Q_2 are expansion trees then Q_3 is an expansion tree, $Sh(Q_1) = Sh(Q_2) = Sh(Q_3)$, and $[Dp(Q_1) \vee Dp(Q_2)] \supset Dp(Q_3)$ is tautologous.

(b) If Q_1 and Q_2 are dual expansion trees then Q_3 is a dual expansion tree, $Sh(Q_1) = Sh(Q_2) = Sh(Q_3)$, and $Dp(Q_3) \supset [Dp(Q_1) \wedge Dp(Q_2)]$ is tautologous.

The following recursively describes how to construct the merge of trees Q_1 and Q_2.

(1) If Q_1 is a one-node tree, then so is Q_2, and $Q_3 := Q_1$ is the merge of Q_1 and Q_2.

(2) If $Q_1 = {\sim}Q_1'$ then $Q_2 = {\sim}Q_2'$. If Q_3' is the merge of Q_1' and Q_2', then $Q_3 := {\sim}Q_3'$ is the merge of Q_1 and Q_2.

(3) If $Q_1 = Q_1' \vee Q_1''$ then $Q_2 = Q_2' \vee Q_2''$. If Q_3' and Q_3'' are the result of merging Q_1' with Q_2' and Q_1'' with Q_2'', then merging Q_1 with Q_2 yields $Q_3 := Q_3' \vee Q_3''$.

(4) If the root of Q_1 is a existential node then $Q_1 = \forall B_1 +^{t_1} Q_1^1 + \ldots +^{t_n} Q_1^n$ and $Q_2 = \forall B_2 +^{s_1} Q_2^1 + \ldots +^{s_m} Q_2^m$, where B_1 and B_2 are equal up to alphabetic change of bound variables, $t_1, \ldots, t_n, s_1, \ldots, s_m$ are formulas, and $n, m \geq 1$. The result of merging Q_1 and Q_2 is then simply

$$Q_3 := \forall B_1 +^{t_1} Q_1^1 + \ldots +^{t_n} Q_1^n +^{s_1} Q_2^1 + \ldots +^{s_m} Q_2^m.$$

(5) If the root of Q_1 is a universal node then $Q_1 = \forall B_1 +^y Q_1'$ and $Q_2 = \forall B_2 +^y Q_2'$ where B_1 and B_2 are equal up to alphabetic changes of bound variables. If Q_3' is the result of merging Q_1' and Q_2' then $Q_3 := \forall B_1 +^y Q_3'$ is the result of merging Q_1 and Q_2. ∎

Lemma 3.13. *Let A, B, and C be formulas$_o$. We then have the following:*

(1) If $A \vee B$ has a grounded ET-proof and $C \vee B$ has a grounded ET-proof, then $[A \wedge C] \vee B$ has a grounded ET-proof.

(2) If $A \vee B \vee B$ has a grounded ET-proof, then $A \vee B$ has a grounded ET-proof.

Proof. Only the proof of (1) is given since the proof of (2) is similar and easier. Thus, assume that $A \vee B$ and $C \vee B$ have grounded ET-proofs $Q_1 \vee Q_2$ and $Q_3 \vee Q_4$. These expansion trees can be picked so that Q_1 and Q_3 can be merged to obtain Q_5 and that $Q := Q_5 \vee [Q_2 \wedge Q_4]$ contains no variable selected twice. Clearly, Q is a grounded expansion tree for $S \vee [A \wedge B]$. Since $Q_1 \vee Q_2$ and $Q_3 \vee Q_4$ are ET-proofs, $Dp(Q_1) \vee Dp(Q_2)$ and $Dp(Q_3) \vee Dp(Q_4)$ are tautologous. By the facts noted in the definition of merging, $[Dp(Q_1) \vee Dp(Q_3)] \supset Dp(Q_5)$ is tautologous. Hence, so too is $Dp(Q) = Dp(Q_5) \vee [Dp(Q_2) \wedge Dp(Q_4)]$. Finally, if the dependency relation for Q contains a cycle, it is easy to show that that cycle must appear in either $Q_1 \vee Q_2$ or $Q_3 \vee Q_4$. Hence, the dependency relation for Q is acyclic, and Q is indeed a grounded ET-proof of $[A \wedge C] \vee B$. □

Section 4: Sequential Proofs

Our higher-order version of the sequential calculus, called $\mathcal{L}$, is very similar to the L-systems given in [7, 13, 16]. For convenience the succedent and antecedent of sequents will be multisets. The structural inference figures for $\mathcal{L}$ are listed in Figure 1 and the introduction inference figures for $\sim$ and $\vee$ are listed in Figure 2. Here, Γ and Θ denote (possibly empty) multisets of formulas$_o$.

$$\frac{\Gamma \longrightarrow \Theta}{A, \Gamma \longrightarrow \Theta}\ \text{Thinning} \qquad \frac{\Gamma \longrightarrow \Theta}{\Gamma \longrightarrow \Theta, A}\ \text{Thinning}$$

$$\frac{A, A, \Gamma \longrightarrow \Theta}{A, \Gamma \longrightarrow \Theta}\ \text{Contraction} \qquad \frac{\Gamma \longrightarrow \Theta, A, A}{\Gamma \longrightarrow \Theta, A}\ \text{Contraction}$$

Figure 1: Structural inference figures.

$$\frac{A, \Gamma \longrightarrow \Theta \qquad C, \Gamma \longrightarrow \Theta}{A \vee C, \Gamma \longrightarrow \Theta}\ \vee\text{–IA} \qquad \frac{\Gamma \longrightarrow \Theta, A, C}{\Gamma \longrightarrow \Theta, A \vee C}\ \vee\text{–IS}$$

$$\frac{\Gamma \longrightarrow \Theta, A}{\sim A, \Gamma \longrightarrow \Theta}\ \sim\text{–IA} \qquad \frac{A, \Gamma \longrightarrow \Theta}{\Gamma \longrightarrow \Theta, \sim A}\ \sim\text{–IS}$$

Figure 2: Inference figures for introducing $\vee$ and $\sim$.

Figure 3 lists the inference figures for the introduction of $\forall$ and the use of λ-convertibility. Here, A conv A'. The inference figure $\forall$–IS must be restricted so that y is not free in any of the formulas in its lower sequent. An *$\mathcal{L}$-derivation* is defined, in the usual fashion, to

be a tree structure of instances of these inference figures. An $\mathcal{L}$-derivation is said to be an *$\mathcal{L}$-proof* of its root sequent (endsequent) if its leaves are all axioms, *i.e.* they are of the form $A \to A$, where A is any formula$_o$. A formula B has an $\mathcal{L}$-proof if $\to B$ is the root sequent of an $\mathcal{L}$-proof.

$$\frac{A, \Gamma \longrightarrow \Theta}{A', \Gamma \longrightarrow \Theta}\ \lambda \qquad \frac{\Gamma \longrightarrow \Theta, A}{\Gamma \longrightarrow \Theta, A'}\ \lambda$$

$$\frac{Bt, \Gamma \longrightarrow \Theta}{\forall B, \Gamma \longrightarrow \Theta}\ \forall\text{–IA} \qquad \frac{\Gamma \longrightarrow \Theta, By}{\Gamma \longrightarrow \Theta, \forall B}\ \forall\text{–IS}$$

Figure 3: Inference figures for λ and $\forall$.

The transformation from ET-proofs to sequent proofs is based on *q-sequents*, which are structures

$$P_1, \ldots, P_r \to Q_1, \ldots, Q_s$$

where $\{Q_1, \ldots, Q_r\}$ is a possibly empty multiset of expansion trees and $\{P_1, \ldots, P_r\}$ is a possibly empty multiset of dual expansion trees, and $\sim P_1 \vee \ldots \vee \sim P_r \vee Q_1 \vee \ldots \vee Q_s$ is an ET-proof. (The choice of how these multisets are enumerated is not important.) This ET-proof is the *ET-proof associated with* this q-sequent, and the sequent

$$Sh(P_1), \ldots, Sh(P_r) \to Sh(Q_1), \ldots, Sh(Q_s)$$

is the *sequent associated with* this q-sequent. Notice that if Q is an ET-proof for A then $\to Q$ is a q-sequent, and the associated sequent would be $\to A'$, where A' is a λ-normal form of A. A q-sequent is *grounded* if its associated ET-proof is grounded.

Now let Σ be a (possibly empty) multiset of dual expansion trees, and let Ω be a (possibly empty) multiset of expansion trees. The following statements are true for q-sequents.

(1) If $P_1 \vee P_2, \Sigma \to \Omega$ is a q-sequent then both $P_1, \Sigma \to \Omega$ and $P_2, \Sigma \to \Omega$ are q-sequents.

(2) If $\Sigma \to \Omega, Q_1 \vee Q_2$ is a q-sequent then $\Sigma \to \Omega, Q_1, Q_2$ is a q-sequent.

(3) If $\Sigma \to \Omega, \sim P$ is a q-sequent then $P, \Sigma \to \Omega$ is a q-sequent.

(4) If $\sim Q, \Sigma \to \Omega$ is a q-sequent then $\Sigma \to \Omega, Q$ is a q-sequent.

These facts can be considered *q-inference figures* and are represented as such in Figure 4.

The following facts about q-sequents follow easily from the property of elimination of top-level selected variables and expansion terms (see Proposition 3.10).

$$\frac{P_1, \Sigma \longrightarrow \Omega \qquad\qquad P_2, \Sigma \longrightarrow \Omega}{P_1 \vee P_2, \Sigma \longrightarrow \Omega}\ \vee\text{-IA}_q \qquad \frac{\Sigma \longrightarrow \Omega, Q_1, Q_2}{\Sigma \longrightarrow \Omega, Q_1 \vee Q_2}\ \vee\text{-IS}_q$$

$$\frac{\Sigma \longrightarrow \Omega, Q}{\sim Q, \Sigma \longrightarrow \Omega}\ \sim\text{-IA}_q \qquad \frac{P, \Sigma \longrightarrow \Omega}{\Sigma \longrightarrow \Omega, \sim P}\ \sim\text{-IS}_q$$

Figure 4: The q-analogs for the inference figures in Figure 2.

(1) If $\Sigma \to \Omega, \forall B +^y Q$ is a q-sequent then $\Sigma \to \Omega, Q$ is a q-sequent.

(2) If $\forall B +^{t_1} P_1, \Sigma \to \Omega$ is a q-sequent and if t_1 is admissible in the associated ET-proof, then $P_1, \Sigma \to \Omega$ is a q-sequent. Similarly, if $[\forall B +^{t_1} P_1 + \ldots +^{t_n} P_n], \Sigma \to \Omega$ $(n > 1)$ is a q-sequent and t_i is admissible for some i, $1 \leq i \leq n$, then

$$P_i, [\forall B +^{t_1} P_1 + \ldots +^{t_{i-1}} P_{i-1} +^{t_{i+1}} P_{i+1} + \ldots +^{t_n} P_n], \Sigma \to \Omega$$

is a q-sequent.

Figure 5 shows the q-inference figures based on these relationships between q-sequents. A *q-derivation* is a tree structure of q-sequents, each of whose nonterminal nodes are instances of one of these seven q-inference figures.

$$\frac{\Sigma \longrightarrow \Omega, Q}{\Sigma \longrightarrow \Omega, \forall B +^y Q}\ \forall\text{-IS}_q \qquad \frac{P_1, \Sigma \longrightarrow \Omega}{\forall B +^{t_1} P_1, \Sigma \longrightarrow \Omega}\ \forall\text{-IA}_q$$

$$\frac{P_i, [\forall B +^{t_1} P_1 + \ldots +^{t_{i-1}} P_{i-1} +^{t_{i+1}} P_{i+1} + \ldots +^{t_n} P_n], \Sigma \longrightarrow \Omega}{[\forall B +^{t_1} P_1 + \ldots +^{t_n} P_n], \Sigma \longrightarrow \Omega}\ \forall\text{-IA}^*_q$$

Figure 5: The q-analogs for the $\forall$ inferences figures.

Let the *degree* of a q-sequent be the number of nonterminal nodes labeled with $\vee$ or $\sim$ plus the number of selected variables and occurrences of expansion terms in its expansion trees and dual expansion trees. In all the q-inference figures in Figures 4 and 5, the upper q-sequents have strictly smaller degree than the lower sequents. Thus, any q-sequent $\Sigma \to \Omega$ can be placed at the root of a q-derivation where all its leaves are q-sequents with zero degree. Theorem 3.11 must be used to ensure that if any expansion terms are present in a q-sequent then at least one of them is admissible. The following theorem now follows easily.

Theorem 4.1. *If A is a formula$_o$ with a grounded ET-proof, then there is an $\mathcal{L}$-derivation of the sequent $\rightarrow A$.*

Proof. Let Q be a grounded ET-proof for A. From the above discussion, $\rightarrow Q$ is the end q-sequent of a q-derivation, Ξ_q, in which the leaves are q-sequents of zero degree and in which the only q-inference figures used are those in Figures 4 and 5. Below, Ξ_q is transformed into an $\mathcal{L}$-derivation tree Ξ whose endsequent is $\rightarrow A'$, where A' is some λ-normal form of A, and whose leaves contain only atomic formulas.

For any q-inference figure in Ξ_q in Figure 4, replace it with the $\mathcal{L}$ inference figure obtained by replacing the upper and lower q-sequents with their associated sequent and dropping the "q" subscript from the name of the inference figure. Hence, if a q-inference figure in Ξ_q is $\vee$$-IA_q$ then the corresponding figure is Ξ is

$$\frac{Sh(P_1), Sh(\Sigma) \longrightarrow Sh(\Omega) \qquad\qquad Sh(P_2), Sh(\Sigma) \longrightarrow Sh(\Omega)}{Sh(P_1) \vee Sh(P_2), Sh(\Sigma) \longrightarrow Sh(\Omega)}\ \vee\text{-IA}$$

Here, $Sh(\Sigma)$ and $Sh(\Omega)$ are the multisets whose elements are the shallow formulas of the trees in Σ and Ω, respectively.

Now consider the inference figures in Figure 5. For any occurrence of a $\forall$$-IS_q$ inference figure in Ξ_q, place in Ξ the inference figures

$$\cfrac{\cfrac{Sh(\Sigma) \longrightarrow Sh(\Omega), A'}{Sh(\Sigma) \longrightarrow Sh(\Omega), By}\ \lambda}{Sh(\Sigma) \longrightarrow Sh(\Omega), \forall B}\ \forall\text{-IS}$$

where A' is a λ-normal form of By. Similarly, for any occurrence of a $\forall$$-IA_q$ inference figure in Ξ_q, place in Ξ the inference figures

$$\cfrac{\cfrac{A', Sh(\Sigma) \longrightarrow Sh(\Omega)}{Bt_1, Sh(\Sigma) \longrightarrow Sh(\Omega)}\ \lambda}{\forall B, Sh(\Sigma) \longrightarrow Sh(\Omega)}\ \forall\text{-IA}$$

where A' is a λ-normal form of Bt_1. Finally, for any occurrence of a $\forall$$-IA^*_q$ inference figure in Ξ_q, place in Ξ the inference figures

$$\cfrac{\cfrac{\cfrac{A', \forall B, Sh(\Sigma) \longrightarrow Sh(\Omega)}{Bt_i, \forall B, Sh(\Sigma) \longrightarrow Sh(\Omega)}\ \lambda}{\forall B, \forall B, Sh(\Sigma) \longrightarrow Sh(\Omega)}\ \forall\text{-IA}}{\forall B, Sh(\Sigma) \longrightarrow Sh(\Omega)}\ \text{Contraction}$$

where A' is a λ-normal form of Bt_i.

Ξ is an $\mathcal{L}$-derivation of $\rightarrow A'$ (where A' is a λ-normal form of A) in which the leaves

are sequents containing only atoms. Let Ξ' be the $\mathcal{L}$-derivation which results from placing at the endsequent of Ξ the inference figure

$$\frac{\longrightarrow\ A'}{\longrightarrow\ A}\ \lambda.$$

Each leaf of Ξ' is a sequent of the form $\Gamma \rightarrow \Theta$ in which Γ and Θ are multisets of atomic formulas$_o$ and $[\vee{\sim}\Gamma] \vee [\vee\Theta]$ is tautologous. Thus Γ and Θ contain a common atomic formula, *i.e.* $\Gamma \rightarrow \Theta$ is the result of thinning an axiom. Let Ξ'' be the $\mathcal{L}$-proof of $\rightarrow A$ which is constructed from Ξ' by placing above each leaf of Ξ' an axiom sequent and enough thinning inferences to infer that leaf. Ξ'' is then an $\mathcal{L}$-proof of $\rightarrow A$. □

Now consider the converse of this theorem.

Theorem 4.2. *If the sequent $\Gamma \rightarrow \Theta$ has an $\mathcal{L}$-proof, there is a grounded q-sequent $\Sigma \rightarrow \Omega$ such that $\Gamma = Sh(\Sigma)$ and $\Theta = Sh(\Omega)$. Thus, if A has an $\mathcal{L}$-proof, it has a grounded ET-proof.*

Proof. Assume that the sequent $\Gamma \rightarrow \Theta$ has an $\mathcal{L}$-proof Ξ. This $\mathcal{L}$-proof can be assumed to have axioms which involve only λ-normal atomic formulas, since otherwise it could be reduced to such a proof by applying various inference figures to the non-atomic axioms. We prove by induction on the height of Ξ that there is a grounded q-sequent, $\Sigma \rightarrow \Omega$, such that $\Gamma = Sh(\Sigma)$ and $\Theta = Sh(\Omega)$.

If the height of Ξ is 0, then Ξ is an axiom instance, say $C \rightarrow C$. If C' is a λ-normal form of C, then $C' \rightarrow C'$ is the desired grounded q-sequent.

Now assume the height of Ξ is greater than 0. We need to show how a q-sequent for the final sequent of Ξ can be constructed from the q-sequent(s) associated with the premise(s) of that sequent. To do this, we need to consider 12 cases; one for each inference rule which may terminate the proof tree Ξ. The construction for most of these cases is very simple. For example, the q-sequent for the premise of either λ inference rule is also a q-sequent for its conclusion. For the $\forall$ introduction rules, the q-sequent for the conclusion is simply the result of adding a labeled arc and new root node to the q-sequent for the premise. It is only the contraction and $\vee-$IA rule where the construction is not immediate. Both these cases, however, were already anticipated and solved by Lemma 3.13.

Thus there is a quite simple inductive algorithm which can construct a ground q-sequent for the root sequent of an $\mathcal{L}$-proof. Hence, if Ξ is an $\mathcal{L}$-proof of the sequent $\rightarrow A$, this algorithm constructs a q-sequent, say $\rightarrow Q$, such that Q is grounded and A red $Sh(Q)$. Hence, Q is a grounded ET-proof for A. □

Another way to state the cut-elimination result for $\mathcal{T}$ is that for any formula$_o$ A, $\vdash_{\mathcal{T}} A$ if and only if there is a (cut-free) proof of A in $\mathcal{L}$. Using the above two theorems, the soundness and completeness for ET-proofs follows immediately, *i.e.* $\vdash_{\mathcal{T}} A$ if and only if A has a grounded ET-proof.

The proof transformation algorithms presented in this section have very practical value for the implementation of theorem proving systems. If a resolution-style theorem prover recorded its substitution information in an expansion tree, the transformation from ET-proofs to $\mathcal{L}$-proof could be used to constitute a "natural" rendering of a resolution refutation. Also, consider the following kind of theorem proving system: let a user interactively edit a sequent-style proof. Once a proof is completed, the transformation of $\mathcal{L}$-proofs to ET-proofs could be applied. The resulting ET-proof would record just the "essential" information of the proof. If this latter ET-proof was to be transformed back to a sequential proof, there are many such proofs which could be produced: the orginal sequential proof is only one of many. If the sequential proof builder encompassed notions of good proof style, the newly created sequential proof could be a stylistically improved version of the orginally entered sequential proof. Such proof revision has been experimented with in [6].

Section 5: Linear Reasoning

If the formula $C \supset D$ has an ET-proof, it is of the form $P \supset Q$, where P is a dual expansion tree for C and Q is an expansion tree for D. (For convenience, expansion trees shall be denoted by a (possibly) ornamented "Q" while a dual expansion tree shall be denoted by a (possibly) ornamented "P.") As was shown in the proof of Theorem 3.8, there is a list of ET-proofs

$$P \supset Q = P_1 \supset Q_1, \quad \ldots \quad , P_n \supset Q_n$$

such that $Dp(P_n) \supset Dp(Q_n)$ is tautologous, and for $1 \leq i < n$, either P_{i+1} is equal to P_i and Q_{i+1} is the result of eliminating a top-level labeled arc from Q_i, or Q_{i+1} is equal to Q_i and P_{i+1} is the result of eliminating a top-level labeled arc from P_i. In the first case, either $\vdash_{\mathcal{T}} Sh(Q_{i+1}) \supset Sh(Q_i)$ or $\vdash_{\mathcal{T}} \forall y\, Sh(Q_{i+1}) \supset Sh(Q_i)$, depending on whether Q_{i+1} is the result of eliminating an expansion term or a universal variable y from Q_i. In the latter case, either $\vdash_{\mathcal{T}} Sh(P_i) \supset Sh(P_{i+1})$ or $\vdash_{\mathcal{T}} Sh(P_i) \supset \exists y\, Sh(P_{i+1})$, depending on whether P_{i+1} is the result of eliminating an expansion term or a universal variable y from P_i. It is this relationship between $Sh(P_i)$ and $Sh(P_{i+1})$ and between $Sh(Q_{i+1})$ and $Sh(Q_i)$ which is the basis of Craig's system of linear reasoning.

Let A be a formula$_o$ and ψ be a (possibly empty) list of quantifier occurrences, $\forall x$ or $\exists x$, for any variable x. We define a *prefixed formula* to be a pair $\langle \psi, A \rangle$, also written as ψA. A prefixed formula $\langle \psi, A \rangle$ *represents the formula* B if B is the result of attaching to A the quantifiers listed in ψ. While $\langle \psi, A \rangle$ represents a unique formula, the converse is not true. For example, the formula $\exists x \forall y\, Pxy$ is represented by the three prefixed formulas $\langle \exists x \forall y, Pxy \rangle$, $\langle \exists x, \forall y\, Pxy \rangle$, and $\langle \emptyset, \exists x \forall y\, Pxy \rangle$.

Let A be a formula$_o$ and let B be a boolean subformula occurrence in A. We will use the symbol $A[B]^{\pm}$ to denote this fact. The symbol $\pm$ will be $+$ if the occurrence of B is

positive and will be $-$ otherwise. Once the position and sign of an occurrence of B in A is established, that occurrence of B can be changed. This is done by using the symbol $A[C]^{\pm}$, *i.e.* the instance of B which was marked out in $A[B]^{\pm}$ is now changed to C. The symbols $A[B]^{\pm}$ and $A[C]^{\pm}$ only make sense when they occur in pairs.

Below is the list of L-deductions used in linear reasoning. Each L-deduction takes one prefixed formula as a premise and yields one prefixed formula as a conclusion.

Duplication: From $\psi\ A[\forall B]^+$ infer $\psi\ A[\forall B \wedge \forall B]^+$.

Simplication: From $\psi\ A[\forall B \wedge \forall B]^-$ infer $\psi\ A[\forall B]^-$.

$\exists$-*exportation:* From $\psi\ A[\forall B]^-$ infer $\psi \exists y\ A[By]^-$.

$\forall$-*importation:* From $\psi \forall y\ A[By]^+$ infer $\psi\ A[\forall B]^+$, when y is not free in $A[\forall B]$.

$\forall$-*vacuous-introduction:* From $\psi_1 \psi_2\ A$ infer $\psi_1 \forall y \psi_2\ A$.

$\exists$-*vacuous-removal:* From $\psi_1 \exists y \psi_2\ A$ infer $\psi_1 \psi_2\ A$, when y is not free in $\psi_2 A$.

$\exists$-*generalization:* From $\psi\ A[Bt]^-$ infer $\psi\ A[\forall B]^-$.

$\forall$-*instantiation:* From $\psi\ A[\forall B]^+$ infer $\psi\ A[Bt]^+$.

Matrix-change: From $\psi\ A$ infer $\psi\ A'$, provided that $A \supset A'$ is tautologous.

λ-*conversion:* From $\psi\ A$ infer $\psi\ A'$, provided that A conv A'.

Notice that if $\psi\ A, \psi'\ A'$ is a L-deduction then for any variable y, both $\forall y \psi\ A$, $\forall y \psi'\ A'$ and $\exists y \psi\ A, \exists y \psi'\ A'$ are justified by the same L-deduction. Furthermore, in the deductions $\forall$-instantiation, $\exists$-generalization, and matrix-change, the premise implies the conclusion. In all the remaining deductions, the premise is equivalent to conclusion. A list of prefixed formulas $\psi_1 A_1 \ldots , \psi_n A_n$ is called an *L-derivation of* $\psi_n A_n$ *from* $\psi_1 A_1$ if for each i, $1 \leq i < n$, $\psi_i\ A_i, \psi_{i+1}\ A_{i+1}$ is one of the above L-deductions. If B_1 represents $\psi_1 A_1$ and B_n represents $\psi_n A_n$, then $\vdash_{\mathcal{T}}\ B_1 \supset B_n$.

Definition 5.1. An L-derivation is *balanced* if there exists exactly one matrix-change deduction and all $\forall$-instantiations and Duplications occur before and all $\exists$-generalization and Simplications after this matrix-change, and all λ-conversions prior to the matrix-change are λ-contractions while all those after the matrix-change are λ-expansions. Given a balanced L-derivation there is a prefix ψ and two formulas$_o$ M_1 and M_2 such that the L-deduction $\psi M_1, \psi M_2$ is the matrix-change deduction for this L-derivation. We shall call ψ the *matrix prefix*, M_1 the *left matrix formula*, and M_2 the *right matrix formula* of this L-derivation. ∎

Theorem 5.2. *If* $\vdash_{\mathcal{T}}\ C \supset D$ *then there is a balanced L-derivation of* D *from* C*. In fact, let* $P \supset Q$ *be an ET-proof of* $C \supset D$ *and let* $\mathbf{S}$ *be the set of selected variables of* $P \supset Q$ *which are free in* $Dp(P \supset Q)$ *and let* $\prec$ *be the embedding relation for* $P \supset Q$*. There is a balanced L-derivation of* D *from* C *such that its left and right matrix formulas are equal (modulo associativity of conjunction) to* $Dp(P)$ *and* $Dp(Q)$*, respectively, and its matrix prefix is*

$$\mathcal{Q}_m y_m, \ldots , \mathcal{Q}_1 y_1,$$

where $y_1, \ldots, y_m$ is a topological sort of $\mathbf{S}$ *with respect to* $\prec$ *(i.e. if $y_i \prec y_j$ then $i < j$) and $\mathcal{Q}_i$ is $\exists$ if y_i is selected in the tree P and is $\forall$ if y_i is selected in the tree Q, for $i = 1, \ldots, m$.*

The proof of this theorem closely parallels the one given by Craig in [5]. There are two main differences, however. The first is the obvious addition of λ-formulas and λ-conversion. This complicates the nature of $\forall$-instantiation and $\exists$-generalization deductions and forces the second difference. Since the substitution of higher-order terms can result in the appearance of embedded quantifiers, formulas in an L-derivation cannot be required to be in prenex normal form. Dropping prenex normal forms, although necessary, does not make doing linear reasoning any harder. L-derivations cannot, however, be require to satisfy the property that any Duplication, $\exists$-exportation, and $\forall$-vacuous-introduction precedes any $\forall$-instantiation deduction and that any Simplication, $\forall$-importation, and $\exists$-vacuous-removal deduction follows any $\exists$-generalization deduction. Although this is possible in the first-order setting, here these various deductions may well need to get mixed on the left side (for the first four deductions) and the right side (for the second four deductions) of the matrix-change.

Let M_1 and M_2 be two formulas such that $M_1 \supset M_2$ is tautologous. From [5] it is easy to see that there is a third formula M such that both $M_1 \supset M$ and $M \supset M_2$ are tautologous, and that every parameter or free variable occurring in M occurs in both M_1 and M_2. Such a formula M is a *propositional interpolant for* $M_1 \supset M_2$. The following weak interpolation theorem holds for $\mathcal{T}$.

Theorem 5.3. *If the closed formula $C \supset D$ has an ET-proof $P \supset Q$ then there exists a formulas X such that $\vdash_{\mathcal{T}} C \supset X$, $\vdash_{\mathcal{T}} X \supset D$, and if a parameter c occurs in X and not in both C and D, then c occurred either in D and some expansion term of P or in C and some expansion term of Q.*

Proof. Let $P \supset Q$ be an ET-proof for $C \supset D$, M be a propositional interpolant for $Dp(P) \supset Dp(Q)$, and ψ be the prefix described in Theorem 5.2. Then ψM is a formula such that $\vdash_{\mathcal{T}} C \supset \psi M$ and $\vdash_{\mathcal{T}} \psi M \supset D$. The only parameters occurring in M must be in both $Dp(P)$ and $Dp(Q)$. Since the only parameters occurring in $Dp(P)$ occur in either C or expansion terms of P and the only parameters occurring in $Dp(Q)$ occur in either D or an expansion term of Q, ψM satisfies the conditions of the theorem. □

This theorem can be viewed as an extension of the interpolation theorem for first-order logic since it immediately yields that theorem: Since expansion terms in first-order logic can not contain predicates, the predicates in the formula ψM must be common to both C and D, no matter which ET-proof was used to construct ψM.

Section 6: Skolemization

The nature of Skolem functions in higher-order logic is more complex than it is for first-order logic. In first-order logic, the Herbrand Universe of terms is freely generated

from Skolem functions using application. Terms containing Skolem functions can be used identical to terms not containing Skolem functions. In higher-order logic, however, the use of Skolem functions must be modified. For example, the formula

$$\forall x_\iota\, \exists y_\iota\, P_{(o\iota)\iota} xy \supset \exists f_{\iota\iota}\, \forall z_\iota\, Pz[fz]$$

is not provable in $\mathcal{T}$, although the following Skolemized form of it is provable:

$$\forall x_\iota\, P_{(o\iota)\iota} x[gx] \supset \exists f_{\iota\iota}\, P[hf][f[hf]].$$

Here, g is a Skolem function$_{\iota\iota}$ and h is a Skolem function$_{\iota(\iota\iota)}$. If such functions are not restricted, this latter formula would have an ET-proof in which f would be substituted with either g or $\lambda x\, gx$. The lack of such a restriction in the resolution refutation system in [1] caused that system to be unsound. Assuming the Axiom of Choice would have made that refutation system sound but then not complete.

Skolem functions should simply play the syntactic role of providing for a parametric collection of new "objects." They should not be used as an actual function such as g is when used in the substitution terms g and $\lambda x\, gx$ in the above example. The necessary restriction on Skolem functions is given in the following definition.

Definition 6.1. The list $\sigma := \langle \alpha, \beta_1, \ldots, \beta_p \rangle$, where $\alpha, \beta_1, \ldots, \beta_p$ are type symbols ($p \geq 0$), is called a *signature* (for a Skolem function). For each signature, σ, let $\mathcal{K}_\sigma$ be a denumerably infinite set of function symbols all of type $(\ldots(\alpha\beta_1)\ldots\beta_p)$ which are not in the formulation of $\mathcal{T}$ and such that if σ_1 and σ_2 are two different signatures then $\mathcal{K}_{\sigma_1}$ and $\mathcal{K}_{\sigma_2}$ are disjoint. $f \in \mathcal{K}_\sigma$ is called a *Skolem function of signature σ with arity p*. Let $\mathcal{T}^*$ be the formulation of $\mathcal{T}$ in which these Skolem functions are added. The *Herbrand Universe* of terms for $\mathcal{T}^*$ is the set, $\mathcal{U}$, of all formulas A of $\mathcal{T}^*$ such that whenever a Skolem function of arity p has an occurrence in A, it is applied to at least p arguments. These arguments are called its *necessary* arguments, and a formula with a Skolem function of arity p with p argument attached is called a *Skolem term*. Furthermore, if a variable has a free occurrence in any of these necessary arguments, that occurrence is also free in A. $\mathcal{U}_\alpha$ will denote the set of all formulas in $\mathcal{U}$ of type α. ∎

Two Skolem functions may have the same type while they have different arities. For example, if α is of the form $\alpha'\beta_0$, then a Skolem term with signature $\langle \alpha, \beta_1, \ldots, \beta_p \rangle$ and one with the signature $\langle \alpha', \beta_0, \beta_1, \ldots, \beta_p \rangle$ have different arities but have the same type. Since types can generally be determined from context while arity often cannot be, Skolem functions are frequently written with a superscripted nonnegative integer to denote its arity, *i.e.* f^p.

Example 6.2. If f, g are Skolem terms with signature $\langle \iota, \iota \rangle$, x, w are variables$_\iota$, and A is a variable$_{o(o\iota)}$ then $f[gx] \in \mathcal{U}_\iota$, $\lambda x x \in \mathcal{U}_{\iota\iota}$, and $\lambda w[Aw[gx]] \in \mathcal{U}_{o\iota}$, while $f \notin \mathcal{U}$, $\lambda x[fx] \notin \mathcal{U}$, and $\lambda w[A[gx][fw]] \notin \mathcal{U}$. ∎

The following proposition can easily be established using standard methods of the λ-calculus.

Proposition 6.3. *If $A \in \mathcal{U}_\alpha$ and A red B then $B \in \mathcal{U}_\alpha$.*

Expansion trees can now be modified to use Skolem functions in place of selected variables.

Definition 6.4. Skolem expansion trees and dual Skolem expansion trees are defined similarly to expansion trees and dual expansion trees, except as follows. When building trees, Skolem terms are used to instantiate universal nodes in place of selected variables. Second, expansion terms must be members of $\mathcal{U}$. Finally, there are two global requirements. Let $f^p t_1 \ldots t_p$ be the Skolem term which labels the sole out-arc of a universal node, N. Then f^p appears as the head of no other Skolem term labeling the out arc of universal node. Also, on the path from N to the root of the tree, there are exactly p expansion terms $t_1, \ldots, t_p$, in that order of dominance, which dominate N. The term $f^p t_1 \ldots t_p$ is said to be used to do a *Skolem instantiation* of the node N.

The definitions for the functions Dp and Sh are essentially the same as they are presented in Definition 3.1. A Skolem expansion tree, Q, is a tree *for* a formula A of $\mathcal{T}$ if $Sh(Q)$ conv A and Q is an *ST-proof* if $Dp(Q)$ is tautologous. ∎

There are two fundamental differences between expansion trees and Skolem expansion trees: First, subtrees of expansion trees are either expansion trees or dual expansion trees, while subtrees of Skolem expansion trees are not necessarily Skolem expansion trees or their duals. Second, a substitution instance of a Skolem expansion tree is another Skolem expansion tree, while the substitution instance of an expansion tree with an acyclic dependency relation may no longer be an expansion tree with an acyclic dependency relation. This fact is the main reason why proof structures using Skolemization are frequently used in automated proof systems. It is important in many of these automated systems that proof systems remain proof systems under substitution.

Example 6.5. Let A be the theorem $\exists y\, \forall x\, [Px \supset Py]$, and let f and g be Skolem functions with signature $\langle \iota, \iota \rangle$. A Skolem expansion tree for A would then be the tree Q_1 given as (compare with Example 3.6):

$$\begin{aligned}[\exists y\, \forall x\, [Px \supset Py] &+^{u} [\forall x\, [Px \supset Pu] +^{fu} [P[fu] \supset Pu]] \\ &+^{v} [\forall x\, [Px \supset Pv] +^{gv} [P[gv] \supset Pv]]].\end{aligned}$$

An ST-proof for A would then be the tree Q_2 given as:

$$\begin{aligned}[\exists y\, \forall x\, [Px \supset Py] &+^{u} [\forall x\, [Px \supset Pu] +^{fu} [P[fu] \supset Pu]] \\ &+^{v} [\forall x\, [Px \supset Pv] +^{g[fu]} [P[g[fu]] \supset P[fu]]]].\end{aligned}$$

Q_2 is the result of substituting fu for v in Q_1. ∎

The definition of ST-proofs has no condition similar the one for ET-proofs which required the dependency relation be acyclic. The restriction it specified for ET-proofs is explicitly encoded into Skolem terms used in ST-proofs. To make this relationship more explicit, the following binary relation for the selected variables in expansion trees is provided.

Definition 6.6. Let Q be an expansion tree. Let $\prec_Q^0$ be the binary relation on $\mathbf{S}_Q$ such that $z \prec_Q^0 y$ if there exists a $t \in \Theta_Q$ such that z is free in t and a node dominated by (the arc labeled with) t is selected by y. $\prec_Q$, the transitive closure of $\prec_Q^0$, is called the *embedding relation.* ∎

This relation is closely related to the dependency relation. The following proposition and its proof reveals this connection.

Proposition 6.7. *$<_Q$ is acyclic if and only if $\prec_Q$ is acyclic.*

Proof. Let $<_Q$ be cyclic. That is, assume that there are expansion term occurrences $t_1, \ldots, t_m \in \Theta_Q$ such that $t_1 <_Q^0 \ldots <_Q^0 t_m <_Q^0 t_{m+1} = t_1$ for $m \geq 1$. Let y_i, for $i = 1, \ldots, m$, be chosen from $\mathbf{S}_Q$ so that y_i is selected for a node dominated by t_i and y_i is free in t_{i+1}. If y_{m+1} is identified with y_1, then $y_i \prec_Q^0 y_{i+1}$, for $i = 1, \ldots, m$, since y_{i+1} is selected for a node dominated by t_{i+1} and y_i is free in the formula t_{i+1}. Hence, $y_1 \prec_Q^0 \ldots \prec_Q^0 y_m \prec_Q^0 y_1$, and $\prec_Q$ is cyclic. The proof in the other direction is very similar. □

As will be shown, an ET-proof Q can be converted to an ST-proof Q', and conversely. In such a case, two selected variables z and y of Q satisfy $z \prec_Q y$ if and only if the Skolem term corresponding to z is a subformula of the Skolem term corresponding to y. Since the subformula relation is guaranteed to be acyclic, it is not necessary to assume any additional acyclic conditions for ST-proofs.

Theorem 6.8. *If A has an ET-proof then A has an ST-proof.*

Proof. Let A have an ET-proof Q. For any selected variable$_\alpha$ $y \in \mathbf{S}_Q$ define an *associated Skolem term* as a term $ft_1 \ldots t_p$, such that (a) $t_1, \ldots, t_p$ are the expansion terms in Q which dominate the arc labeled with y, (b) these expansion terms are ordered so that $1 \leq l < k \leq p$ implies t_l dominates t_k, and (c) f is some Skolem function with signature $\langle \alpha, \beta_1, \ldots, \beta_p \rangle$, where β_j is the type of t_j, $j = 0, \ldots, p$. Since none of the formulas, $t_1, \ldots, t_p$ contain Skolem functions, associated Skolem terms are all members of $\mathcal{U}$.

Let $\langle y_1, \ldots, y_r \rangle$ be a list of the variables in $\mathbf{S}_Q$ such that $y_i \prec_Q y_j$ implies $i < j$, and let $\langle s_1, \ldots, s_r \rangle$ be a corresponding list of associated Skolem terms all of which have different Skolem functions for their heads. Notice that $1 \leq i \leq j \leq r$ implies that y_j is not free in s_i.

Let $\varphi := \mathbf{S}^{y_1}_{s_1} \circ \cdots \circ \mathbf{S}^{y_r}_{s_r}$. A simple induction argument shows that if C is a formula of $\mathcal{T}$, then φC will be a formula of $\mathcal{T}^*$ in which none of the variables $y_1, \ldots, y_r$ are free.

Also, φy_i is a Skolem term with top-level Skolem function f_i. It is easy to verify that φ commutes with Dp, Sh, and $+^t$. From this it follows immediately that φQ is an ST-proof for A. $\square$

To convert ST-proofs to ET-proofs, Skolem terms must be replaced with selected variables. The following operator is used to do this.

Definition 6.9. Let $A \in \mathcal{U}_\alpha$, s be a Skolem term$_\beta$, and y be a variable$_\beta$ which does not appear in A or in s. Let $D^s_y A$, the *deskolemizing operator*, be the result of replacing in A every subformula, t, such that t conv s, by y. If Q is a Skolem expansion tree, then $D^s_y Q$ is the result of applying D^s_y to all formulas labeling arcs and nodes in Q. ▮

Example 6.10. If f be a Skolem function with signature $\langle \iota\iota, \iota \rangle$, then

$$D^{fv}_y[\lambda z_\iota [f v_\iota z]] = \lambda z[y_{\iota\iota} z] \text{ and } D^{fv}_y Pz[f[[\lambda w_\iota\ w]v]] = Pzy.$$

▮

The following proposition can be proved using standard methods of λ-calculus.

Lemma 6.11. *Let $A, B \in \mathcal{U}_\alpha$, s be a Skolem term$_\beta$, and y be a variable$_\beta$ which does not appear in A, B, or s. The all the following are true.*

(1) $D^s_y A \in \mathcal{U}_\alpha$ *and* $\mathbf{S}^s_y D^s_y A$ conv A.

(2) A red B *implies* $D^s_y A$ red $D^s_y B$.

(3) *If B is a λ-normal form of A, then $D^s_y B$ is a λ-normal form of $D^s_y A$.*

The soundness of Skolem functions is proved by the following theorem.

Theorem 6.12. *If A has an ST-proof then A has an ET-proof.*

Proof. Let Q be an ST-proof for A and let $\mathcal{V}$ be the set of principle normal forms of Skolem terms which are subformulas of formulas used to do expansions or Skolem instantiations in Q. Let $\langle s_1, \ldots, s_r \rangle$ be an ordering of $\mathcal{V}$ such that whenever s_j is an alphabetic variant of a subformula of s_i then $i < j$. Let $y_1, \ldots, y_r$ be r distinct variables new to Q and A such that y_i has the same type as s_i, $i = 1 \ldots, r$. Let ρ be the compound deskolemizing operator

$$\rho := D^{s_n}_{y_n} \circ \cdots \circ D^{s_1}_{y_1}.$$

Now $\rho s_i = y_i$, for all $i = 1, \ldots, r$, since $j < i$ implies that $D^{s_j}_{y_j} s_i = s_i$. By Lemma 6.11, it is easy to see that ρ commutes with Dp, Sh, and $+^t$. Since $Sh(Q)$ contains no Skolem functions, $Sh(\rho Q) = Sh(Q)$. Hence, ρQ is an expansion tree for A. We need only show that $\prec_{\rho Q}$ is acyclic.

Assume that $y_i \prec^0_{\rho Q} y_j$ for some $y_i, y_j \in \mathbf{S}_{\rho Q}$. Then the selection arc labeled with y_j in ρQ is dominated by some expansion term t which contains a free occurrence of y_i. Let t', s'_j, s'_i be the formula labeling the arcs in Q which correspond to the arc labeled with t, y_j, y_i, respectively, in ρQ. Here, the principle normal form of s'_j and s'_i are s_j and s_i,

respectively. Clearly, t' would contain s_i' as a subformula, while t' is a subformula of s_j'. Thus s_i is an alphabetic variant of a subformula of s_j and, therefore, $j < i$. Therefore, if $y_i \prec_{\rho Q} y_j$ then $j < i$, and $\prec_{\rho Q}$ is acyclic. □

The soundness and completeness for ST-proofs follow immediately from the above two theorems.

Acknowledgements

Much of the material presented here was taken from my dissertation done in the Mathematics Department at Carnegie-Mellon University. This thesis, titled "Proofs in Higher-Order Logic," is available as technical report MS-CIS-83-37 from the Computer and Information Science Department, University of Pennsylvania. I am greatly indebted to Peter Andrews, my dissertation advisor, for my formal training in logic and for a good deal of encouragement. I was also helped by numerous conversations with Frank Pfenning, Rick Statman, Daniel Leivant, and Dana Scott. I would also like to thank Gopalan Nadathur, Dafa Li, and Greg Hager for their comments on an early draft of this paper. This research was supported by NSF grant MCS81-02870.

Bibliography

[1] Peter B. Andrews, *Resolution in Type Theory*, Journal of Symbolic Logic **36** (1971), 414 – 432.

[2] Peter B. Andrews, Dale A. Miller, Eve Longini Cohen, Frank Pfenning, "Automating Higher-Order Logic" in *Automated Theorem Proving: After 25 Years*, AMS Contemporary Mathematics Series **29** (1984).

[3] Alonzo Church, *A Formulation of the Simple Theory of Types* Journal of Symbolic Logic **5** (1940), 56 – 68.

[4] Alonzo Church, *The Calculi of Lambda-Conversion*, Princeton University Press, 1941.

[5] William Craig, *Linear reasoning. A new form of the Herbrand-Gentzen theorem*, Journal of Symbolic Logic **22** (1957), 250 – 268.

[6] Amy P. Felty, "Using Extended Tacticals to Do Proof Transformations," MSE Thesis, University of Pennsylvania, December 1986.

[7] Gerhard Gentzen, *Investigations into Logical Deductions* in *The Collected Papers of Gerhard Gentzen* edited by M. E. Szabo, North-Holland Publishing Co., Amsterdam, 1969, 68 – 131.

[8] J. Herbrand, *Recherches sur la théorie de la démonstration*, Travaux de la Société des Sciences et des Lettres de Varsovie, Classe III Sciences Mathematiques et Phyiques 33 (1930). Translated in [9].

[9] J. Herbrand, *Logical Writings*, Harvard University Press, Cambridge, Mass, 1972.

[10] Gérard P. Huet, "A Mechanization of Type Theory," *Proceedings of the Third International Joint Conference on Artificial Intelligence* 1973, 139 – 146.

[11] Dale A. Miller, "Expansion Tree Proofs and Their Conversion to Natural Deduction Proofs," 7th International Conference on Automated Deduction (ed. by R. E. Shostak), Springer-Verlag, LNCS 170, 375 –393.

[12] Tomasz Pietrzykowski, "A Complete Mechanization of Second-Order Type Theory," Journal of the Association for Computing Machinery **20** (1973), 333 – 364.

[13] Dag Prawitz, *Natural Deduction*, Almqvist & Wiksell, 1965.

[14] Dag Prawitz, *Hauptsatz for higher order logic*, Journal of Symbolic Logic **33** (1968), 452 – 457.

[15] J. A. Robinson, "Mechanizing Higher-Order Logic," *Machine Intelligence 4*, Edinburgh University Press, 1969, 151 – 170.

[16] Raymond M. Smullyan, *First-Order Logic*, Springer-Verlag, New York, 1968.

[17] Moto-o-Takahashi, *A proof of cut-elimination theorem in simple type-theory*, Journal of the Mathematical Society of Japan **19** (1967), 399 – 410.

A Compact Representation of Proofs

Dale A. Miller
Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389, USA

23 February 1993

Abstract: A structure which generalizes formulas by including substitution terms is used to represent proofs in classical logic. These structures, called *expansion trees*, can be most easily understood as describing a tautologous substitution instance of a theorem. They also provide a computationally useful representation of classical proofs as first-class *values*. As values they are compact and can easily be manipulated and transformed. For example, we present an explicit transformations between expansion tree proofs and cut-free sequential proofs. A theorem prover which represents proofs using expansion trees can use this transformation to present its proofs in more human-readable form. Also a very simple computation on expansion trees can transform them into Craig-style linear reasoning and into interpolants when they exist. We have chosen a sublogic of the Simple Theory of Types for our classical logic because it elegantly represents substitutions at all finite types through the use of the typed λ-calculus. Since all the proof-theoretic results we shall study depend heavily on properties of substitutions, using this logic has allowed us to strengthen and extend prior results: we are able to prove a strengthen form of the first-order interpolation theorem as well as provide a correct description of Skolem functions and the Herbrand Universe. The latter are not straightforward generalization of their first-order definitions.

Table of Contents

This paper appears in *Studia Logica*, Vol. 46, No. 4, 1987.