Resolution and Type Theory

Leen Helmink

Philips Research Laboratories
P.O. Box 80.000, 5600 JA Eindhoven, the Netherlands

Abstract.
In this paper, an inference mechanism is proposed for proof construction in Constructive Type Theory. An interactive system that implements this method has been developed.

Key words. Type Theory, Calculus of Constructions, Typed Lambda Calculus, Natural Deduction.

1 Introduction

A method is presented to perform unification based top down proof construction for Constructive Type Theory, thus offering a well-founded, elegant and powerful underlying formalism for a proof development system. It combines the advantages of Horn clause resolution and higher order natural deduction style theorem proving. No theoretical contribution to Constructive Type Theory is claimed. The method is demonstrated for the Calculus of Constructions [Co85], but is applicable to other variants of type theory as well, e.g. systems in the families of AUTOMATH [Br73][Da80], Martin-Löf [Ma84], LF [Ha87], Elf [Pf89]. A full derivation example is included.

The problem addressed in this paper is to construct an object in a given context, given its type. This amounts to higher order theorem proving. This paper demonstrates that this construction problem can be handled by Horn clause resolution, provided that the set of available Horn clauses is continuously adapted to the context in which the proof is conducted. This rests on a mechanism that provides a simple clausal interpretation for the assumptions in a context. The method is not complete, due to the expressive power of type theory. Although the provided inference steps suggest certain search strategies (*tactics*), these issues are outside the scope of this paper. A proof environment based on the method, named *Constructor*, has been developed within Esprit project 1222: 'Genesis' [He88][He89]. Experiments with this system demonstrate the power and efficiency of the method.

2 Constructive Type Theory (CTT)

We assume familiarity with Constructive Type Theory [Br73][Co85][Cq85][Fo83][Hu87], a variant of the *propositions as types* paradigm. This section describes the particular system of interest. We shall use a system developed by Coquand, a version of the Calculus of Constructions. It must be emphasized that this is just one of the possible variants for which the presented proof construction method is applicable.

2.1 Terms

The syntactic formation rules for the terms in the system are defined as:

- constant, viz. one of {*prop, type, kind*}.
- variable, denoted by an identifier.
- $\lambda[x{:}A].B$, typed abstraction, where A and B are terms, and x a variable.
- $\Pi[x{:}A].B$, generalized Cartesian product of types B indexed over x of type A, where A and B are terms, and x a variable. [1]
- $(A\ B)$, application, where A and B are terms (function and argument). [2]

A *typing* is a construction of the form $[t{:}T]$, where t and T are terms. The intuition behind this is that T is the type of t. Types may have types themselves, and we will refer to such expressions as domains. Four levels of expressions are distinguished in the hierarchy of types: 0-, 1-, 2-, and 3-expressions, where n-expressions serve as types of $(n+1)$-expressions $(n = 0, 1, 2)$. There is only one 0-expression: the constant 'supertype' *kind*. We introduce two predefined constants as primitive 1-expressions (kinds) of the system: the kind *type*, the set containing all 'plain' types (2-expressions), and the kind *prop*, that plays an identical role and is treated similarly, but which is inhabited by propositions (2-expressions) [3]. The level of 3-expressions contains objects and proofs.

2.2 Correctness

A *context* is a list of assumptions, introducing variables with their type. A context is of the form:

$$[x_1{:}A_1],...,[x_n{:}A_n] \quad (n \geq 0)$$

No variable may be declared more than once in a context. Moreover, the free variables occurring in A_i must have been declared earlier in the context, i.e. $\mathcal{FV}(A_i) \subseteq \{x_1, ..., x_{i-1}\}$, where $\mathcal{FV}(A)$ stands for the set of variables occurring free in A. Apart from assumptions, contexts may contain *definitions*. A definition introduces a variable as an abbreviation for an arbitrary correct expression. A definition is of the form $[c \equiv a{:}A]$, where c is a fresh identifier, and establishes that c abbreviates the term a of type A.

We will use Γ as a metavariable over contexts. Well-formedness of a context Γ will be denoted *well-formed*(Γ). We will write $\Gamma, [x{:}A]$ to denote the context Γ extended with the assumption $[x{:}A]$, and Γ_1, Γ_2 for the concatenation of contexts Γ_1 and Γ_2. The symbol '$\emptyset$' represents the empty context. We will write $E \in \Gamma$ to denote that E is a member of context Γ. A *sequent* is an expression of the form $\Gamma \vdash [a{:}A]$, denoting that $[a{:}A]$ is a correct typing in the well-formed context Γ. The predicate *constant*(x) denotes that x is one of the constants *type*, *prop* or *kind*. We will write $B[a/x]$ to denote substitution of the term a for the free occurrences of the variable x in the expression B. We will write '$=_{\beta\delta}$' to denote the transitive reflexive closure of β- and δ-reduction. β-reduction corresponds to the usual notion in typed lambda-calculus, and δ-reduction denotes local expansion of definitions that have been introduced in the context. Note that both β- and δ-reduction are context-dependent. We assume α-conversion whenever necessary and all equality is modulo α-conversion. This can be achieved by using De Bruijn indices [Br72] or Barendregt's variable convention [Ba81].

[1] In type theory, typed abstraction is often denoted $[x{:}A]B$, while typed product is denoted $(x{:}A)B$ or $\{x{:}A\}B$.

[2] Application associates to the left, so we will write (a b c) for ((a b) c).

[3] In many versions of this system, *type* and *prop* are identified (usually denoted '$\star$' or '*Type*'). Here we will explicitly distinguish between them to avoid possible confusion. This is however not essential to the formalism.

Well-formedness of contexts and correct typing of terms is inductively defined as:

[0] $\mathit{well\text{-}formed}(\emptyset)$

[1a] If $\mathit{well\text{-}formed}(\Gamma)$ then $\Gamma \vdash [\mathit{type}{:}\mathit{kind}]$

[1b] If $\mathit{well\text{-}formed}(\Gamma)$ then $\Gamma \vdash [\mathit{prop}{:}\mathit{kind}]$

[2a] If $\Gamma \vdash [A{:}K]$ and $\mathit{constant}(K)$ then $\mathit{well\text{-}formed}(\Gamma, [x{:}A])$

[2b] If $\Gamma \vdash [a{:}A]$ then $\mathit{well\text{-}formed}(\Gamma, [c \equiv a{:}A])$

[3a] If $\mathit{well\text{-}formed}(\Gamma)$ and $[x{:}A] \in \Gamma$, then $\Gamma \vdash [x{:}A]$

[3b] If $\mathit{well\text{-}formed}(\Gamma)$ and $[c \equiv a{:}A] \in \Gamma$, then $\Gamma \vdash [c{:}A]$

[4] If $\Gamma, [x{:}A] \vdash [b{:}B]$ then $\Gamma \vdash [\lambda[x{:}A].b : \Pi[x{:}A].B]$

[5] If $\Gamma, [x{:}A] \vdash [B{:}K]$ and $\mathit{constant}(K)$ then $\Gamma \vdash [\Pi[x{:}A].B{:}K]$

[6] If $\Gamma \vdash [a{:}A]$ and $\Gamma \vdash [b : \Pi[x{:}A].B]$ then $\Gamma \vdash [(b\ a){:}B[a/x]]$

[7] If $\Gamma \vdash [a{:}A]$ and $\Gamma \vdash [B{:}K]$ and $\mathit{constant}(K)$ and $A =_{\beta\delta} B$ then $\Gamma \vdash [a{:}B]$

Details and properties of the described system can be found in [Co85][Cq85][Ha89][Hu87][Ju86].

2.3 Interpretation and Use

If, for a dependent product type $\Pi[x{:}A].B$, x does not occur free in B ($x \notin \mathcal{FV}(B)$), the type simplifies, as usual, to the ordinary *function type* $A \rightarrow B$. In case $A{:}\mathit{prop}$, A is considered a proposition and a typing $[a{:}A]$ is interpreted as: a is a proof for A, i.e. a proposition plays the role of the type of its proofs. This means that a proposition is considered valid if and only if it is inhabited. The system described contains intuitionistic higher order predicate logic. For example, if $B{:}\mathit{prop}$, then $\Pi[x{:}A].B$ can be interpreted as the universally quantified proposition $\forall x{:}A.B$. If x does not occur free in B and $A{:}\mathit{prop}$ and $B{:}\mathit{prop}$, then $\Pi[x{:}A].B$ can be interpreted as the intuitionistic implication $A \Rightarrow B$. The Calculus of Constructions formalism thus provides a definition language that can be and has been used to formalize and mechanically verify many parts of mathematics. Texts in this language are written in the form of *theories* (*books* in AUTOMATH terminology). Theories are CTT contexts. For a field of interest, assumptions allow the axiomization of the primitive notions, whereas the definitions allow abbreviation of derived notions like lemmas.

3 CTT Proof Construction Method

It is a well-known fact, that correctness of sequents $\Gamma \vdash [t{:}T]$ (t has type T in context Γ) in Constructive Type Theory and related systems is decidable, even feasibly decidable, and several proof checkers exist that mechanically determine correctness for given CTT theories [Ju76], [Da80], [Co85]. For a proof construction system, the objective is not to verify whether a given object has a certain given type, but, for a given type, to attempt construction of an inhabitant of this type. More precisely: given a context (theory) Γ and a type A, the objective is to construct an object p such that $\Gamma \vdash [p{:}A]$. For propositions, this corresponds to finding a proof object.

The central problem with goal directed proof construction in Constructive Type Theory is that direct backward chaining with the correctness rules of section 2.2 is hardly possible. Therefore, the approach is to extract from the given formalization a sound set of derived rules, that do allow easy backward inference. These derived rules then serve as the primitive proof steps of the system.

In the method, CTT sequents will be derived using Horn clause resolution. In goal directed proving, the idea is to start off with a goal to be proven, and to replace goals by appropriate subgoals by resolution with inference rules [Ro65]. Horn clause inference rules consist of a (possibly empty) set of antecedents $S_1...S_k$ and one conclusion S. Horn clauses will be denoted: $S \Leftarrow S_1...S_k$. In the method, antecedents and conclusions will all be CTT sequents. Sequents may contain logical variables over CTT terms. To avoid confusion with CTT variables, we will denote logical variables by identifiers prefixed will a '$\sharp$' symbol. Logical variables are considered to be universally quantified over clauses. A term is grounded if it does not contain logical variables. A context will be called grounded if it contains grounded terms only. The meaning of a Horn clause is that *if* instantiations for the logical variables can be found, such that the antecedent sequents are correct, *then* the associated consequent is a correct sequent. A Horn clause will be considered *valid* if its meaning is correct with respect to the correctness rules of CTT. Unification determines how two rules can be combined into derivations. A derivation for a sequent is either a Horn clause concluding that sequent, or it is a derivation where an antecedent sequent S_i is replaced by the antecedents of a Horn clause concluding S_i', after unifying S_i and S_i'. If no antecedents are left in a derivation, the derivation is complete and serves as a justification for the correctness of its conclusion sequent. Derivations correspond to derived Horn clauses and have the same interpretation.

The queries considered consist of one goal of the form $\Gamma \vdash [\sharp P{:}A]$, where Γ must be a grounded, well-formed context, A must be a grounded, correct domain, and $\sharp P$ is a logical variable. For a given query S, the derivation process starts with the trivial derivation '$S \Leftarrow S$', which is obviously valid. The objective is to transform this derivation by resolution with valid Horn clauses, until the derivation '$S' \Leftarrow$' is reached. S' is then a correct instance of S.

Derivations and Horn clauses will be of the form: $\Gamma \vdash [p{:}A] \Leftarrow \Gamma_1 \vdash [p_1{:}A_1] \ ... \ \Gamma_n \vdash [p_n{:}A_n]$. The following invariant properties will hold for all derivations (but not necessarily for Horn clauses):

1. all contexts Γ and Γ_1 ... Γ_n will be grounded and well-formed.
2. $\Gamma \subseteq \Gamma_1$, ... , $\Gamma \subseteq \Gamma_n$.
3. for any logical variable $\sharp P$ occurring in the type field A_i of a subgoal $\Gamma_i \vdash [p_i{:}A_i]$, $\sharp P \in \{p_1...p_{i-1}\}$. If an object field p_i of a subgoal $\Gamma_i \vdash [p_i{:}A_i]$ is not a logical variable, then for any logical variable $\sharp P$ occurring in p_i, $\sharp P \in \{p_1...p_{i-1}\}$.
4. for any logical variable $\sharp P$ occurring in the object field p of the conclusion $\Gamma \vdash [p{:}A]$, $\sharp P \in \{p_1...p_n\}$.

The first property reflects the fact that construction always takes place within a known context. The second property states that contexts can be extended during backward proving. The third property ensures that logical variables are 'introduced before use'. The fourth property guarantees that the conclusion of a derivation will be grounded when all subgoals have been solved (the type field A of a derivation conclusion is grounded from the start). For our queries $\Gamma \vdash [\sharp P{:}A]$ this implies that an object $\sharp P$ of the requested type A has been constructed. Note that the trivial derivations that correspond to our queries of interest have all the required properties.

It turns out that we can avoid a problem that usually arises with inference rules over sequents, viz. unification over contexts. The reason for this is that in contrast to general purpose higher order theorem provers as presented in e.g. [Pa86], [Pa89] and [Fe88], the method inferences at the object level, not at the meta level. This is possible because the method is specialized for type theory

only; it is not a generic inference method over arbitrary sequents. Contexts will be treated in a special way, and it is sufficient to unify over typings. Our goals, that denote CTT sequents, will be regarded as tuples with a typing and a context.

For the method to work, it is necessary that contexts occurring in derivations are grounded, so that (1) we do not unify over contexts at all, and (2) we can extract the necessary object level horn clauses directly from contexts. The exact consequences of this restriction will be alluded later.

In the subsequent sections, valid Horn clauses will be derived. These Horn clauses will not violate the given invariant property for derivations. Some of the Horn clauses correspond directly to correctness rules for type theory and are of interest to all subgoals in a derivation. The problem is that no suitable Horn clause can be given for the application rule (rule 6), on account of the substitution. A solution for this problem is presented, that rests on a mechanism to provide a direct clausal interpretation for the assumptions in a context. The Horn clauses thus obtained cover derivation steps that can construct the necessary application terms. This is the crux of the method. For any given subgoal $\Gamma_i \vdash [p_i{:}A_i]$, this mechanism, when applied to the context Γ_i, allows derivation of a set of valid Horn clauses that are candidates for resolution on this particular subgoal. A proposal for this mechanism was first suggested by [Ah85].

3.1 Kinds rule

The *kinds* rule (rule 1) is directly equivalent to the valid Horn clauses:

$\Gamma \vdash [\mathit{type} : \mathit{kind}] \Leftarrow .$
$\Gamma \vdash [\mathit{prop} : \mathit{kind}] \Leftarrow .$

For any subgoal $\Gamma_i \vdash [p_i{:}A_i]$, such a rule applies if $[p_i{:}A_i]$ unifies with $[\mathit{type}{:}\mathit{kind}]$ or $[\mathit{prop}{:}\mathit{kind}]$. Γ is not treated as a logical variable. Because contexts in our derivations are always grounded and well-formed, the well-formedness check on Γ_i (required in rule 1) is needless.

3.2 Lambda abstraction rule

The lambda abstraction rule (rule 4) corresponds to the valid Horn clause:

$$\Gamma \vdash [\lambda[x{:}\sharp A].\sharp B : \Pi[x{:}\sharp A].\sharp T] \Leftarrow \Gamma, [x{:}\sharp A] \vdash [\sharp B{:}\sharp T].$$

The typing of a goal of the form $\Gamma_i \vdash [P : \Pi[x{:}A].T]$ may be unified with the typing in the conclusion of this rule, unifying P with $\lambda[x{:}\sharp A].\sharp B$ and resulting in the new stripped subgoal $[\sharp B{:}\sharp T]$, to be solved in the context Γ_i extended with the typing $[x{:}\sharp A]$. Γ does not play the role of a logical variable. To ensure that this new context is grounded and well-formed, the restriction is imposed that A must be grounded and correct domain, thus preventing logical variables over domains to be introduced in the context. For example, this rule can not be used to find a proof for an implication $\sharp A \Rightarrow B$, because this would introduce an unknown assumption $[p{:}\sharp A]$ in the context.

3.3 Pi abstraction rule

The pi abstraction rule (rule 5) corresponds to the valid Horn clause:

$$\Gamma \vdash [\Pi[x{:}\sharp A].\sharp B : \sharp K] \Leftarrow \Gamma, [x{:}\sharp A] \vdash [\sharp B{:}\sharp K].$$

For goals of the form $\Gamma_i \vdash [\Pi[x{:}A].B : K]$, application of this rule results, after unification of typings, in a stripped subgoal $[B{:}K]$, to be solved in the context Γ_i extended with the typing $[x{:}A]$. K must be a constant. Though this check has to be postponed if K is not grounded, it can be demonstrated that this need never occur, and K must be either *type* or *prop*. To ensure that the new context is grounded and well-formed, again the restriction is imposed that A must be a grounded and correct domain. For example, this prevents using this rule on a goal $\Gamma_i \vdash [\sharp P{:}prop]$.

3.4 Derived Clauses

The application rule (rule 6) cannot be translated directly to a valid Horn clause, on account of the substitution. A solution is offered, for which the following theorem is essential:

Correctness of a sequent of the form

$$\Gamma \vdash [C : \Pi[x_1{:}A_1] \ldots \Pi[x_n{:}A_n].B]$$

is equivalent to the validity of the Horn clause:

$$\begin{aligned}\Gamma, \Gamma' \vdash [(C\ \sharp x_1 \ldots \sharp x_n) : B[\sharp x_1/x_1, \ldots, \sharp x_n/x_n]] \Leftarrow\ & \Gamma, \Gamma' \vdash [\sharp x_1{:}A_1] \\ & \Gamma, \Gamma' \vdash [\sharp x_2{:}A_2[\sharp x_1/x_1]] \\ & \ldots \\ & \Gamma, \Gamma' \vdash [\sharp x_n{:}A_n[\sharp x_1/x_1, \ldots, \sharp x_{n-1}/x_{n-1}]].\end{aligned}$$

where $\sharp x_1 \ldots \sharp x_n$ are the logical variables of the Horn clause. The context Γ, Γ' denotes any well-formed context extension of Γ. Note that all possible occurrences of the CTT variables x_i have been replaced by corresponding logical variables $\sharp x_i$. For a complete proof of this theorem the reader is referred to [He88]. For n=0, the set of premises is empty and the application in the consequent simplifies to the object C. The selection rule (rule 3a and 3b) justifies that the theorem is in particular applicable to all introductions and definitions occurring in any well-formed context Γ, i.e.

For all CTT variables c, if

$$[c : \Pi[x_1{:}A_1] \ldots \Pi[x_n{:}A_n].B] \in \Gamma \quad \text{or} \quad [c \equiv C : \Pi[x_1{:}A_1] \ldots \Pi[x_n{:}A_n].B] \in \Gamma$$

this theorem guarantees that:

$$\begin{aligned}\Gamma \vdash [(c\ \sharp x_1 \ldots \sharp x_n) : B[\sharp x_1/x_1, \ldots, \sharp x_n/x_n]] \Leftarrow\ & \Gamma \vdash [\sharp x_1{:}A_1] \\ & \Gamma \vdash [\sharp x_2{:}A_2[\sharp x_1/x_1]] \\ & \ldots \\ & \Gamma \vdash [\sharp x_n{:}A_n[\sharp x_1/x_1, \ldots, \sharp x_{n-1}/x_{n-1}]].\end{aligned}$$

This result is now used to interpret the introductions and definitions in a context in this clausal form. The idea is to 'unfold' the top level Π-abstractions for context elements to clauses. For a goal $\Gamma_i \vdash [p_i{:}A_i]$, all clauses thus obtained from the grounded context Γ_i are available as valid Horn clauses for resolution on this goal. Note that resolving a goal with such a Horn clause does not affect the associated context, i.e. all subgoals that arise are to be solved in the original context Γ_i. Although for a given context element of type $\Pi[x_1{:}A_1] \ldots \Pi[x_n{:}A_n].B$ (where B is not itself a Π-abstraction) the theorem gives $n+1$ different valid Horn clauses, it is sufficient to provide the completely unfolded clause with n antecedents. The small price to pay is that the resulting proofs may contain η-redexes. If desired, these can be reduced immediately. See also section 5.

3.5 Unification and Type Conversion

Unification determines whether a clause is applicable in a given situation. Unification will also handle the type conversion rule (rule 7), dealing with equality of types. It is important to observe that it is sufficient to provide unification over typings, not over contexts (although context information is of course relevant for β- and δ-reductions during unification). Unifying typings $[P{:}A]$ and $[P'{:}A']$ can be achieved by unifying the objects P and P' and subsequently unifying the types A and A'. For types, the unification is with respect to β- and δ-equality. Although β-equality for objects is not explicit in the correctness rules, it is also desirable to identify β- (and δ-) equivalent terms. This is justified by the closure under reduction property, and corresponds to proof normalization [Co85][Da80][Ha87]. Because we unfold derived clauses completely, it is desirable to augment object unification with outermost η-equality, to ensure reachability of objects in η-normal form. Unification for expressions in typed λ-calculus with respect to α, β and possibly η-conversion requires complete higher order unification. This problem has a possibly infinite set of solutions and is known to be semidecidable, in the sense that if two terms do not unify, search algorithms for unifiers may diverge [Hu75]. [Hu75] also gives an algorithm that can be used for implementations of the method. Sound approximations for higher order unification may also be used, although this affects completeness, of course.

4 Completeness

An interesting question is whether the method is complete in the sense that a top down derivation can be constructed for all correct inhabitants (modulo object conversion) of a given type (checking is complete). Note that completeness is of course determined by the completeness of the higher-order unification procedure. But what exactly are the consequences of the restriction that we impose on the context, viz. that it is always grounded?

We already saw that our queries of interest are not affected by the restriction. Now consider the effect of the restriction during the inferencing process. It should be clear that only the lambda rule and the pi rule are affected by the restriction, as they may extend a context during resolution.

For issues related to completeness (at least in a non-deterministic sense), the following observation is important: due to the third invariant property on derivations, we only need to consider goals where the type field of the conclusion is grounded, because any logical variable $\sharp P$ occurring there can be instantiated by first solving the associated goal where $\sharp P$ is introduced. This implies that the partiality of the lambda-abstraction rule poses no fundamental restrictions, because it can be circumvented by postponing resolution on the goal in question. It is clear that the restriction on the applicability of the pi-abstraction rule poses real limitations: it explicitly restricts querying for arbitrary 2-expressions, i.e. it refuses to enumerate all Π-abstracted propositions or types. The expressive power of CTT is such, that it does allow the construction of proofs that involve e.g. induction loading, where a stronger proposition is sought to construct a proof for a weaker one. Top down construction of such proofs is unattainable in general. The limitations imposed on the pi abstraction rule are related to this fundamental problem.

Thus, the method as presented is not complete. The restriction affects construction of propositions and types. Since it is possible to enumerate all propositions (or types), the method can be made complete by replacing the pi rule by an enumerator algorithm, in those places where construction is desired.

5 Derivation Example

As an example, consider the following correct theory Γ_0:

$$
\begin{array}{l}
[nat : type],\\
[0 : nat],\\
[s : \Pi[x{:}nat].\ nat],\\
[< : \Pi[x{:}nat].\ \Pi[y{:}nat].\ prop],\\
[axiom1 : \Pi[x{:}nat].\ (<\ x\ (s\ x))],\\
[trans : \Pi[x{:}nat].\ \Pi[y{:}nat].\ \Pi[z{:}nat].\\
\qquad \Pi[p{:}(<\ x\ y)].\ \Pi[q{:}(<\ y\ z)].\ (<\ x\ z)],\\
[ind : \Pi\ [p : \Pi[x{:}nat].\ prop].\\
\qquad \Pi\ [g : (p\ 0)].\\
\qquad \Pi\ [h : \Pi[n{:}nat].\ \Pi[hyp{:}(p\ n)].\ (p\ (s\ n))].\\
\qquad \Pi\ [z : nat].\ (p\ z)],\\
[pred1 \equiv \lambda[x{:}nat].\ (<\ 0\ (s\ x)) : \Pi[x{:}nat].\ prop]
\end{array}
$$

To elucidate the method, a top down derivation is now presented for the theorem $\forall y{:}nat.\ (pred1\ y)$, i.e. a proof object $\sharp P$ is sought, such that $\Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].\ (pred1\ y)]$. The associated trivial derivation for the query is:

$$\Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].\ (pred1\ y)] \Leftarrow \Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].\ (pred1\ y)].$$

The Horn clauses available for resolution are:

$$
\begin{array}{l}
\Gamma \vdash [\lambda[x{:}\sharp A].\sharp B : \Pi[x{:}\sharp A].\sharp T] \Leftarrow \Gamma, [x{:}\sharp A] \vdash [\sharp B{:}\sharp T].\\
\Gamma \vdash [\Pi[x{:}\sharp A].\sharp B : \sharp K] \Leftarrow \Gamma, [x{:}\sharp A] \vdash [\sharp B{:}\sharp K].
\end{array}
$$

viz. the lambda abstraction rule and the pi abstraction rule (the kinds rule is of no relevance to this example). These Horn clauses are always available for goals. For a given goal, they are extended with Horn clauses that can be obtained from the unfolded context elements of the goal. For the antecedent in the given trivial derivation this means:

$$
\begin{array}{l}
[nat : type] \Leftarrow .\\
[0 : nat] \Leftarrow .\\
[(s\ \sharp X) : nat] \Leftarrow [\sharp X : nat].\\
[(<\ \sharp X\ \sharp Y) : prop] \Leftarrow [\sharp X : nat]\ [\sharp Y : nat].\\
[(axiom1\ \sharp X) : (<\ \sharp X\ (s\ \sharp X))] \Leftarrow [\sharp X : nat].\\
[(trans\ \sharp X\ \sharp Y\ \sharp Z\ \sharp P\ \sharp Q) : (<\ \sharp X\ \sharp Z)] \Leftarrow [\sharp X{:}nat]\ [\sharp Y{:}nat]\ [\sharp Z{:}nat]\\
\qquad\qquad [\sharp P{:}(<\ \sharp X\ \sharp Y)]\ [\sharp Q{:}(<\ \sharp Y\ \sharp Z)].\\
[(ind\ \sharp P\ \sharp G\ \sharp H\ \sharp Z) : (\sharp P\ \sharp Z)] \Leftarrow [\sharp P : \Pi[x{:}nat].\ prop]\ [\sharp G : (\sharp P\ 0)]\\
\qquad\qquad [\sharp H : \Pi[n{:}nat].\ \Pi[hyp{:}(\sharp P\ n)].\ (\sharp P\ (s\ n))]\ [\sharp Z : nat].\\
[(pred1\ \sharp X) : prop] \Leftarrow [\sharp X : nat].
\end{array}
$$

where Γ_0 has been omitted from the sequents. Similar clauses can be constructed for extensions of Γ_0. The only rule applicable to our derivation is the lambda abstraction rule. Resolution gives:

$$\Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].\ (pred1\ y)] \Leftarrow \Gamma_0, [y{:}nat] \vdash [\sharp P' : (pred1\ y)].$$

instantiating $\sharp P$ to $\lambda[y{:}nat].\sharp P'$. In the extended context of the subgoal, a derived clause for y ($[y{:}nat] \Leftarrow .$) is now available. Resolution with the induction clause (*ind*) from the context gives:

$$\Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].(pred1\ y)] \Leftarrow \Gamma_0, [y{:}nat] \vdash [pred1 : \Pi[x{:}nat].\ prop]$$
$$\Gamma_0, [y{:}nat] \vdash [\sharp G : (pred1\ 0)]$$
$$\Gamma_0, [y{:}nat] \vdash [\sharp H : \Pi[n{:}nat].\ \Pi[hyp{:}(pred1\ n)].\ (pred1\ (s\ n))]$$
$$\Gamma_0, [y{:}nat] \vdash [y : nat].$$

instantiating $\sharp P'$ to $(ind\ \sharp P''\ \sharp G\ \sharp H\ \sharp Z)$, $\sharp P''$ to *pred1*, and $\sharp Z$ to y. Note that this requires higher order unification. The first subgoal is grounded and can be checked, but this goal can also be solved after resolution with the lambda abstraction rule, provided that the unification knows that *pred1* is equivalent to $[x{:}nat](pred1\ x)$ (outermost η-equivalence). The last subgoal is solved directly with the Horn clause for y from the context. The derivation thus becomes:

$$\Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].(pred1\ y)] \Leftarrow \Gamma_0, [y{:}nat] \vdash [\sharp G : (pred1\ 0)]$$
$$\Gamma_0, [y{:}nat] \vdash [\sharp H : \Pi[n{:}nat].\ \Pi[hyp{:}(pred1\ n)].\ (pred1\ (s\ n))].$$

The first subgoal resolves with the Horn clause for *axiom1* from the context, instantiating $\sharp G$ to $(axiom1\ 0)$ and leaving $[0{:}nat]$ as trivial subgoal that can be resolved immediately. The remaining subgoal is stripped twice with the lambda abstraction rule. The derivation is now:

$$\Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].\ (pred1\ y)] \Leftarrow \Gamma_1 \vdash [\sharp H' : (pred1\ (s\ n))].$$

instantiating $\sharp H$ to $\lambda[n{:}nat].\lambda[hyp{:}(pred1\ n)].\ \sharp H'$. Γ_1 stands for $\Gamma_0, [y{:}nat], [n{:}nat], [hyp{:}(pred1\ n)]$. The remaining proof obligation is now resolved with the context clause for transitivity (*trans*):

$$\Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].\ (pred1\ y)] \Leftarrow \Gamma_1 \vdash [0 : nat]$$
$$\Gamma_1 \vdash [\sharp Y : nat]$$
$$\Gamma_1 \vdash [(s\ (s\ n)) : nat]$$
$$\Gamma_1 \vdash [\sharp P1 : (<\ 0\ \sharp Y)]$$
$$\Gamma_1 \vdash [\sharp Q : (<\ \sharp Y\ (s\ (s\ n)))].$$

instantiating $\sharp H'$ to $(trans\ 0\ \sharp Y\ (s\ (s\ n))\ \sharp P1\ \sharp Q)$. The first and third subgoal are eliminated with context clauses from Γ_1 for 0, s and n. Because these subgoals are grounded, this amounts to checking. Resolving the last subgoal with *axiom1* instantiates $\sharp Q$ to $(axiom1\ n)$ and $\sharp Y$ to $(s\ n)$. The derivation has become:

$$\Gamma_0 \vdash [\sharp P : \Pi[y{:}nat].\ (pred1\ y)] \Leftarrow \Gamma_1 \vdash [(s\ n) : nat]$$
$$\Gamma_1 \vdash [\sharp P1 : (<\ 0\ (s\ n))]$$
$$\Gamma_1 \vdash [(s\ n) : nat].$$

The proof is completed with the context clauses for s, n and *hyp*. The complete proof $\sharp P$ is now:

$$\lambda[y{:}nat].\ (ind\ pred1$$
$$(axiom1\ 0)$$
$$\lambda[n{:}nat].\ \lambda[hyp{:}(pred1\ n)].\ (trans\ 0\ (s\ n)\ (s\ (s\ n))\ hyp\ (axiom1\ (s\ n)))$$
$$y)$$

Note that this proof is an η-redex. This is due to the fact that derived clauses are unfolded as far as possible here, thus constructing applications that are provided with the full number of arguments. If outermost η-conversion of objects is provided, the η-normal proof can also be derived.

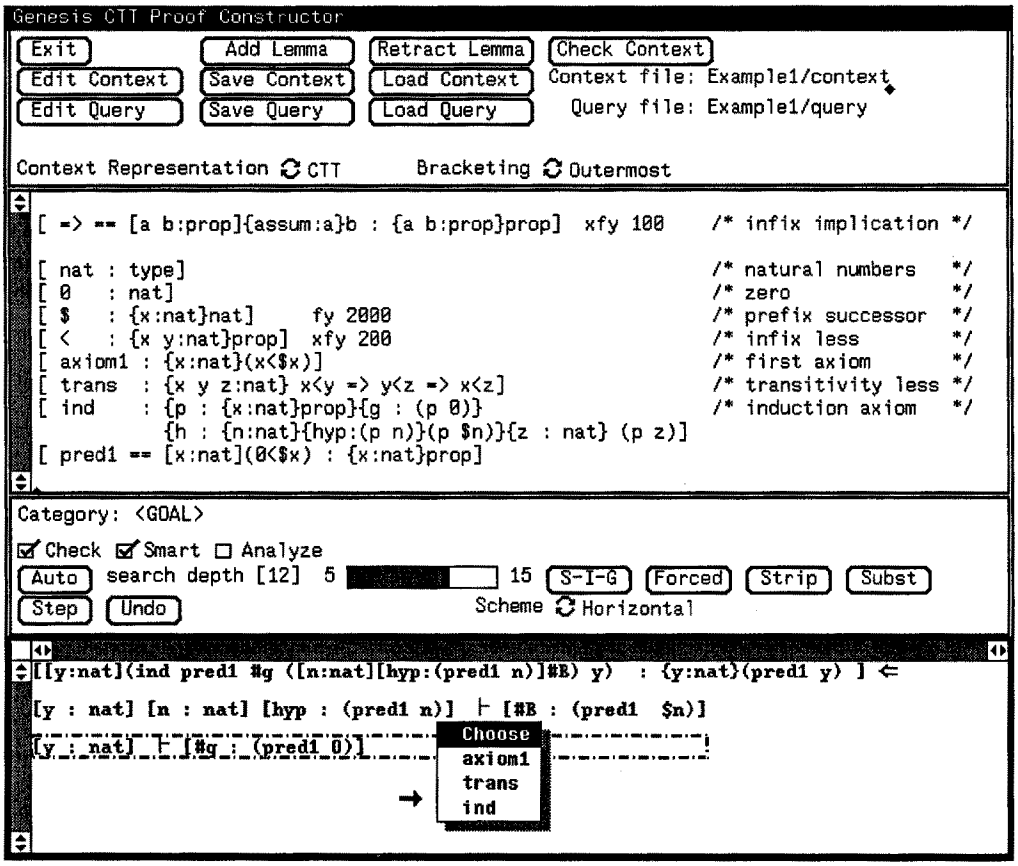

Figure 1: *Here, we are in the middle of an interactive session, solving the query from the example in the previous section. The second window contains the global context. For each goal, candidate clauses for resolution can be selected from a menu. The proof under construction is collected in the head of the derivation. Note the different local context extensions for the different subgoals.*

6 The *Constructor* Proof Environment

This section gives a short description of an interactive proof environment, named *Constructor*, that implements an inference machine based on the described method. The machine enforces correctness of proof construction in CTT. The mouse-based interface has been built using the *Genesis* system, the tool generator that resulted from Esprit project 1222. Details of the *Constructor* system can be found in [He88][He89]. Here, we will explain its most important features.

When using *Constructor*, there is always a global context present, which is the theory that formalizes a domain of interest. A proof editor is provided in which conjunctions of queries can be posed, and that admits application of correct proof steps. Queries are interpreted in the global context. Queries are typings of the form $[A{:}B]$. Figures 1 and 2 present some screen images of the system (for the syntax used in the figures see section 6.1).

Both interactive user-guided inference and automatic search are possible and may intermingle. In interactive mode the user may, for a selected goal, choose a clause from a menu with resolution candidates. Optionally, the system checks instantiated subgoals that may arise after resolution steps. Currently only one default search strategy or *tactic* (*tactical* in LCF terminology) is present for automatic search. It uses a consecutively bounded depth-first search strategy. The maximum search depth must be specified interactively. Alternative solutions are generated upon request. Facilities to 'undo' user or tactic choices are provided. The resolution method itself is used (by way of bootstrapping) for correctness checking of contexts and local context introductions. Completed derivations may be added to the global context as lemmas. To this end, they must be given a name and will be available for use in subsequent queries. It is possible to 'freeze' definitions, i.e. to hide their contents and treat them as axioms. In case of clash of variable names (α-clash), unique variable names are generated by numbering. Textual editing of theories and queries is provided in the environment itself.

The special handling of contexts can be implemented efficiently. For example, translation of context elements to clauses only needs to be done once, because the main theorem guarantees that clauses remain valid in extended contexts (this is due to the fact that CTT is monotonic). Verifying well-formedness of contexts can be done incrementally. Contexts can be shared amongst goals.

Apart from the application rule (rule 6) and the selection rule (rule 3), which are handled by the method proposed, the *Constructor* system is parameterized with respect to the correctness rules for type theory, i.e. the system can handle different versions and extensions of type theory. For example, it poses no problems to interpret various dialects of AUTOMATH [Br73][Da80].

6.1 Technical Details

In *Constructor*, typed abstraction is denoted [x:A]B, whereas typed product is denoted as $\{x{:}A\}B$. Multiple variable introductions are permitted, e.g. $[x\ y{:}A]B$ denotes $[x{:}A][y{:}A]B$. Variables and definitions that are introduced can be declared as fix operators, much like in a Prolog fashion. For example, '[=> == [*a b:prop*]{*p:a*}*b* : {*a b:prop*}*prop*] *xfy 100*' declares '=>' (implication) as a right associative infix operator with priority level *100*. Application is treated as a left-associative infix operator. Bracketing is used in the usual way to over-rule priorities. Logical variables are prefixed with a '♯' symbol.

The built-in unification procedure implements a simple approximation of higher order unification with the following characteristics :

- *Higher Order Structural Matching*
  First order unification where logical variables for functors match structurally, e.g. '(♯F 0)' unifies with '(suc 0)' yielding unifier ♯F=suc.
- *Alpha conversion*
  The unification is modulo the name of bound variables, e.g. '[x:nat]x' equals '[y:nat]y'.
- *Beta conversion*
  The built in unification procedure will reduce β-redexes if necessary. To ensure that the reduction is always sound, a goal is added to demonstrate that the argument will have the required domain type, in the context in question.
- *Delta conversion*
  The built in unification procedure will do δ-reduction on definitions if necessary, i.e. it may expand abbreviating names.

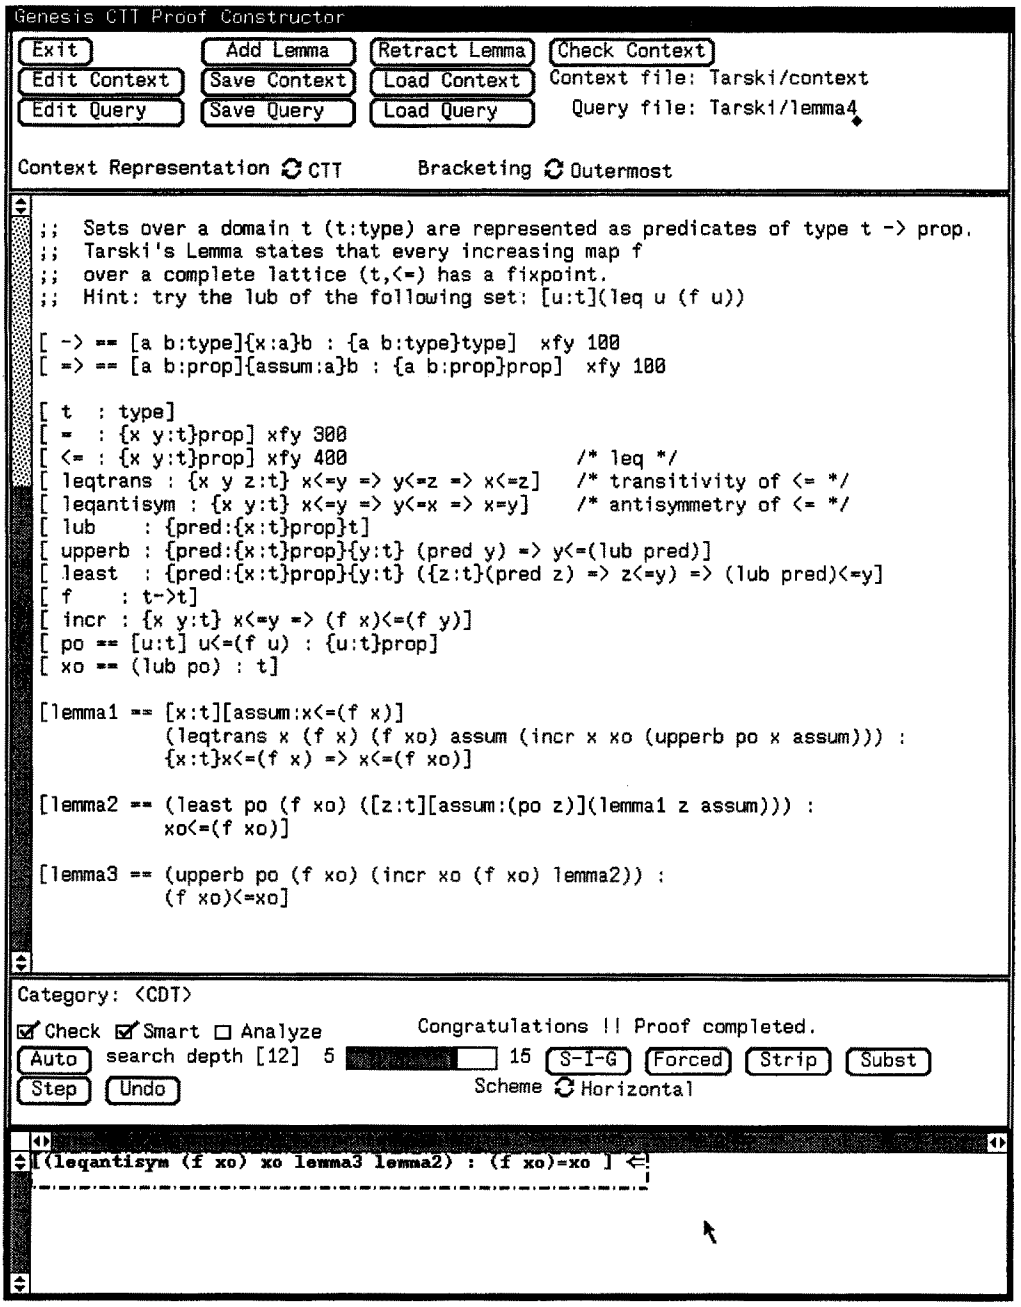

Figure 2: *Proving Tarski's Lemma [Hu87]. After automatically proving* lemma1, lemma2 *and* lemma3 *first, the fixed point property for the witness is proven correct (bottom window). Note that the proof of* lemma2 *is not in η-normal form. The search strategy confirms that the proofs given by Huet [Hu87] are indeed the shortest proofs.*

The unification also recognizes outermost η-equality for objects, so that it can use the lambda abstraction rule to verify given application objects where the functor has not been provided with the full number of arguments. This can be regarded as the inverse operation of 'unfolding' Π-abstractions to clauses. The implemented unification procedure will always yield at most one unifier. If complicated higher order unification is required, two options are available: (1) provide appropriate auxiliary definitions to obtain the desired result (cf. *pred1* in the derivation example) (2) Interactively substitute a template of the desired proof object by hand. Checking objects (also those that can not be constructed by the unification) is always possible, because all relevant terms are known then.

The *Constructor* system automatically proves the example from the previous section within seconds. As another example, the system can construct the proof for Tarski's Lemma [Hu87] [4] , a famous example from constructive mathematics, by first proving the lemmas as proposed in [Hu87], and adding those to the context. In fact, on SUN 4 workstations, the system only needs the second of the three lemmas used in [Hu87] to find the complete proof. See also figure 2.

The performance of the system is good. On SUN 4 workstations, the system performs on an average 4000 unifications every second (including possible β and δ-reductions). Currently, no clause compilation takes place. Automatic construction of the proofs shown in the figures runs in seconds.

The *Constructor* system is implemented in Common Lisp. If we disregard the interface, it consists of 1800 lines of code. User licences for the *Constructor* system can be obtained at no cost. The system runs on SUN 3/4 workstations with 8+ Megabytes of memory. A full Sun Common Lisp Licence is required. A version without the window interface is also available and runs on any Common Lisp.

7 Related Systems

The *Nuprl* system ([Co86]) offers an interesting and impressive interactive proof development environment that is based on Martin-Löf's Type Theory [Ma84]. It is a significant improvement over the LCF proof system ([Go79]) that strongly influenced it. There is an inconvenience in the Nuprl system that may have been overcome by the method proposed in this document. The problem referred to is that proof construction in Nuprl is not based on unification based clausal resolution. The inference rules are the correctness rules for the underlying type theory, and implications in a context can not be used directly as derived rules to resolve goals. Defining a new rule requires a detailed knowledge of the system and the programming language ML. As unification is not directly available, derivation of new hypotheses by instantiating others is often demanded, i.e. variables need to be given that could have been calculated. Tactics in Nuprl are written in the meta language ML. Because unification based resolution with hypotheses is not provided, writing tactics is difficult. Automated theorem proving has not yet been accomplished with Nuprl ([Co86], p.13).

One of the most powerful existing proof systems is *Isabelle* [Pa86][Pa89]. Comparison to this system is difficult, because Isabelle is a generic theorem prover, whereas the proof method proposed here is dedicated to only one single proof formalism, viz. CTT. The same remark can be made for a comparision to the work of Miller on theorem provers [Fe88].

[4] Actually, Huet uses a different version of the Calculus of Constructions where [*type*:*kind*] and [*prop*:*type*], but this is not essential to the example.

8 Discussion

The method presented combines the advantages of resolution inference with the power of type theory.

Because proof construction is not decidable, strategic information has to be provided by users, either in the form of interactive choices, or in the form of algorithms (tactics). Resolution inference allows easy writing of tactics. The method presented may have potential to be used as a logic programming language, that includes all the essential features of e.g. Prolog, but that also provides typing, higher order facilities and the use of local assumptions, thereby creating the possibility to handle queries containing universal quantification or implication (much like λProlog [Na88]). Note that the resulting derivations are in a natural deduction style.

As a meta language, the CTT formalism is suitable to specify logical systems. The method presented makes the object level inference rules of a logic directly available for resolution instead of just the underlying correctness rules of CTT, thus offering the appropriate inference level. The abbreviation mechanism provides the possibility for hiding and the use of derived lemmas.

The requirement that contexts are always grounded in derivations is essential to the method, because it avoids the problem of unification over contexts and prevents the undesired generation of new axioms, while permitting extraction of necessary derived Horn clauses. It has been demonstrated that the consequences of this restriction are directly related to a fundamental problem in higher order theorem proving, and a solution is offered if completeness is desired.

It should be noted that proofs constructed by the method are in β-normal form, aside from definitions. In other words, the proofs constructed are cut-free. The proofs are not guaranteed in η-normal form, unless outermost η-reduction on objects is provided.

Actual implementations of proof systems can efficiently handle many issues. A version of such a proof system, *Constructor*, automatically constructed many non-trivial proofs.

Acknowledgements

The author owes much gratitude to Jan Bergstra, Loe Feijs, Bert Jutting, Ton Kalker, Frank van der Linden and Rob Wieringa for numerous suggestions and corrections. Special thanks are due to René Ahn, for essential contributions to this work, to Marcel van Tien, who implemented most of *Constructor*, and to Henk Barendregt, for many stimulating and clarifying discussions.

References

[Ah85] Ahn, R.M.C. *Some extensions to Prolog based on* AUTOMATH. Internal Philips technical note nr. 173/85, 1985.

[Ba81] Barendregt, H. *The Lambda-Calculus: Its Syntax and Semantics.* North-Holland, 1981.

[Br72] De Bruijn, N.G. *Lambda-Calculus Notation with Nameless Dummies, a Tool for Automatic Formula Manipulation, with Application to the Church-Rosser Theorem.* Indag. Math. 34, 5, pp 381-392, 1972.

[Br73] De Bruijn, N.G. *A survey of the project Automath.* In: *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalisms.* (Seldin & Hindley, Eds.), Academic Press, 1980.

[Co86] Constable, R.L. et al. *Implementing Mathematics with the Nuprl Proof Development System.* Prentice Hall, 1986.

[Co85] Coquand, T. *Une Theory des Constructions.* Thèse de troisième cycle, Université de Paris VII, 1985.

[Cq85] Coquand, T. & Huet, G.P. *Constructions: A Higher Order Proof System for Mechanizing Mathematics.* EUROCAL85, Linz, Springer-Verlag LNCS 203, 1985.

[Da80] van Daalen, D.T. *The Language Theory of Automath.* Ph. D. Dissertation, Eindhoven University of Technology, Dept of Mathematics, 1980.

[Fe88] Felty, A. & Miller, D.A. *Specifying Theorem Provers in a Higher-Order Logic Programming Language.* CADE-9, Argonne, pp 61-80, Springer-Verlag LNCS 310, 1988.

[Fo83] Fortune, S. et al. *The Expressiveness of Simple and Second-Order Type Structures.* J. ACM 30, 1 (Jan.), pp 151-185, 1983.

[Go79] Gordon, M.J., Milner, R. & Wadsworth, C.P. *Edingburgh LCF.* Springer-Verlag LNCS 78, 1979.

[Ha87] Harper, R., Honsell, F. & Plotkin, G. *A Framework for defining logics.* Second Annual Symposium on Logic in Computer Science, Ithaca, IEEE, pp. 194-204, 1987.

[Ha89] Harper, R. & Pollack, R. *Type Checking, Universe Polymorphism, and Typical Ambiguity in the Calculus of Constructions.* Tapsoft '89, Barcelona, Springer-Verlag LNCS 352, Volume 2, 1989.

[He88] Helmink, L. & Ahn, R. *Goal Directed Proof Construction in Type Theory.* Internal Philips technical note nr. 229/88, 1988. Also available as document 28.3 of Esprit project 1222: 'Genesis'.

[He89] Helmink, L. & Tien, M. van. *Genesis Constructive Logic Machine: User's Guide.* Available as document 28.5 of Esprit project 1222: 'Genesis'.

[Hu75] Huet, G.P. *A Unification Algorithm for Typed λ-calculus.* Theoret. Comput. Sci. Vol.1, pp. 27-57, 1975.

[Hu87] Huet, G.P. *Induction Principles Formalized in the Calculus of Constructions.* Tapsoft '87, Pisa, Springer-Verlag LNCS 250, Volume 1, 1987.

[Ju76] Jutting, L.S. *A Translation of Landau's "Grundlagen" in* AUTOMATH. Ph.D. Dissertation, Eindhoven University of Technology, Dept of Mathematics, 1976.

[Ju86] Jutting, L.S. *Normalization in Coquand's system.* Internal Philips technical note nr. 156/88, 1988.

[Ma84] Martin-Löf, P. *Intuitionistic Type Theory.* Bibliopolis, Napoli, 1984.

[Na88] Nadathur G. & Miller, D.A. *An overview of λProlog.* In: *Logic Programming: Proceedings of the Fifth International Conference and Symposium.* (Kowalski & Bowen, Eds.), MIT Press, Cambridge, Massachusetts, Volume 1, pp 820-827, August 1988.

[Pa86] Paulson, L.C. *Natural Deduction as Higher-Order Resolution.* J. Logic Programming 3, pp 237-258, 1986.

[Pa89] Paulson, L.C. *The Foundation of a Generic Theorem Prover.* J. Automated Reasoning 5, pp 363-379, 1989.

[Pf89] Pfenning, F. *Elf: a language for logic definition and verified meta-programming.* Fourth Annual Symposium on Logic in Computer Science, IEEE, pp 313-322, 1989.

[Ro65] Robinson, J.A. *A Machine Oriented Logic Based on the Resolution Principle.* J. ACM 12, 1 (Jan.), 23-49, 1965.