

Soundness of the Simply Typed Lambda Calculus in ACL2

Sol Swords
sswords@cs.utexas.edu

William R. Cook
wcook@cs.utexas.edu

Department of Computer Sciences
University of Texas at Austin

ABSTRACT

To make it practical to mechanize proofs in programming language metatheory, several capabilities are required of the theorem proving framework. One must be able to represent and efficiently reason about complex recursively-defined expressions, define arbitrary induction schemes including mutual inductions over several objects and inductions over derivations, and reason about variable bindings with minimal overhead. We introduce a method for performing these proofs in ACL2, including a macro which automates the process of defining functions and theorems to facilitate reasoning about recursive data types. To illustrate this method, we present a proof in ACL2 of the soundness of the simply typed λ -calculus.

1. INTRODUCTION

Programming language metatheory is a tempting target for mechanized theorem proving. Desirable properties can usually be succinctly stated, but the complexity of the proofs grows with the number of syntactic constructs in the language. Often such proofs are strategically simple but can be long enough that they are tedious to write and difficult to check. Unfortunately, the details which make hand proofs tedious can also impede efforts to mechanize the proofs. Further discussion of the motivations for and challenges facing efforts to mechanize metatheoretic proofs can be found in [1].

In this paper we examine a proof in ACL2 of the soundness of the simply-typed λ -calculus, discussing problems which present themselves in going from the well-known hand proof to a mechanized proof. We present a helpful method of guiding proofs in this area. We also discuss a representation framework generated by a macro which facilitates reasoning about expressions with complicated abstract syntax.

2. BACKGROUND

To demonstrate our method for metatheory proofs, we show a proof of the soundness of the simply-typed λ -calculus with booleans. The abstract syntax of types and expressions for this language are as follows:

T	::=	Bool		$T \rightarrow T$	<i>Types</i>		
X, Y, Z	::=	True		False		v	<i>Simple Expressions</i>
				$\lambda v:T. X$	<i>Abstraction</i>		
				$X Y$	<i>Application</i>		
				if Z then X else Y	<i>Conditionals</i>		

Bool is the single base type, and the arrow type $A \rightarrow B$ is the type of a function which takes an argument of type A and produces a result of type B . Expressions consist of the Boolean constants True and False, variable references, λ abstractions, function applications, and conditional expressions.

To formalize this description, we define the semantics of the language in terms of an evaluation relation, a two-element relation between expressions. Elements of the relation are written $X \rightsquigarrow Y$, meaning X evaluates to Y in one step. We consider the Boolean constants and λ abstractions to be values, meaning they are considered to be fully evaluated.

These rules define the evaluation relation, where $[v \mapsto Y]X$ denotes a capture-avoiding substitution of Y for v in X :

$$\frac{X_1 \rightsquigarrow X_2}{X_1 Y \rightsquigarrow X_2 Y} \quad (\text{E-APP1})$$

$$\frac{X \text{ is a value} \quad Y_1 \rightsquigarrow Y_2}{X Y_1 \rightsquigarrow X Y_2} \quad (\text{E-APP2})$$

$$\frac{Y \text{ is a value}}{(\lambda v:T. X) Y \rightsquigarrow [v \mapsto Y]X} \quad (\text{E-APPABS})$$

$$\frac{Z_1 \rightsquigarrow Z_2}{\text{if } Z_1 \text{ then } X \text{ else } Y \rightsquigarrow \text{if } Z_2 \text{ then } X \text{ else } Y} \quad (\text{E-IFCOND})$$

$$\text{if True then } X \text{ else } Y \rightsquigarrow X \quad (\text{E-IFTRUE})$$

$$\text{if False then } X \text{ else } Y \rightsquigarrow Y \quad (\text{E-IFFALSE})$$

Requiring that X or Y be a value, although not absolutely necessary, has the effect of determining the order of evaluation of sub-expressions. Note that types are ignored by the evaluation relation.

The evaluation relation does not define transitions for values, because they are fully evaluated. But there are other expressions for which no transition is defined. Examples of such expressions are $\text{if } (\lambda v:\text{Bool}. X) \text{ then True else False}$, in which a function appears where a boolean value is expected, and True False , which uses the boolean True where a function is expected. These expressions are syntactically well-formed

but semantically nonsensical. They are examples of *type errors*, because of the mismatch between the type of value and the type of value that is needed. Such expressions are called *stuck*: they are not values but they cannot be evaluated further. The type system’s purpose is to ensure that stuck expressions do not occur anywhere during evaluation.

The type system is based on a three-place relation, traditionally written as $\Gamma \vdash X : T$, which means “expression X has type T under typing context Γ .” The typing context Γ (also called the type environment) is a list of assumptions of the types of variables, each written $v : T$, meaning variable v has type T . A term is considered well-typed if it has a type under the empty context, which is written as a blank, as in $\vdash X : T$.

The following rules define the typing relation:

$$\begin{array}{c}
\Gamma \vdash \text{True} : \text{Bool} \qquad\qquad\qquad (\text{T-TRUE}) \\
\Gamma \vdash \text{False} : \text{Bool} \qquad\qquad\qquad (\text{T-FALSE}) \\
\frac{v : T \in \Gamma}{\Gamma \vdash v : T} \qquad\qquad\qquad (\text{T-VAR}) \\
\frac{\Gamma, v : T_1 \vdash X : T_2}{\Gamma \vdash \lambda v : T_1 . X : T_1 \rightarrow T_2} \qquad\qquad\qquad (\text{T-ABS}) \\
\frac{\Gamma \vdash X : T_1 \rightarrow T \quad \Gamma \vdash Y : T_1}{\Gamma \vdash X Y : T} \qquad\qquad\qquad (\text{T-APP}) \\
\frac{\Gamma \vdash Z : \text{Bool} \quad \Gamma \vdash X : T \quad \Gamma \vdash Y : T}{\Gamma \vdash \text{if } Z \text{ then } X \text{ else } Y : T} \qquad\qquad\qquad (\text{T-IF})
\end{array}$$

We will prove the soundness of this type system: that if a term is well-typed, its repeated evaluation will never give rise to a stuck expression (a non-value which has no possible evaluations). The statement of soundness is in the Progress and Preservation theorems, discussed in section 4.

3. REPRESENTING TERMS AND TYPES

It is easy to design a simple representation in ACL2 of the λ expressions and types using a list representation with distinguished symbols to denote different syntactic forms. However, in doing preliminary work on this problem we discovered that it could become very cumbersome to reason about such recursive data structures. In our case, a function which takes a λ -expression as input will have to distinguish between the six different syntactic forms and typically break down the components of the expression to construct new ones. To streamline these operations, we defined constructor and accessor functions for these structures as simple list operations. Even so, when we left their definitions enabled we found that proofs were cumbersome and difficult to read. In order to reason about them with the function definitions disabled, we found that numerous trivial theorems were necessary and that with an incomplete set of them most proofs would quickly fail. We therefore settled on creating a macro

which, given a description of the desired structure, would define all the functions and the theorems necessary to reason about them, allowing us to leave the function definitions disabled. After many revisions to the group of events submitted by this macro, we find that only in rare and specific cases is it necessary to re-enable the function definitions.

The macro we defined which generates these structures is named `defsum`, because we are defining a type which is a sum (or disjoint union) of several Cartesian products of types. The macro’s syntax is a Lisp adaptation of the syntax for defining similar datatypes in languages such as ML and Haskell. `Defsum` is different from the `defstructure` book in ACL2 [2], which allows update of components of a structure, but does not support mutual recursion between structures. In contrast, `defsum` does not support update but does support mutual recursion. The abstract syntax of terms and types can be represented in a straightforward manner using `defsum`, as shown below.

```

(defsum type
  (BOOL)
  (FUN (stype-p domain) (stype-p range)))

(defsum expression
  (TRUE)
  (FALSE)
  (LAM (varname-p var)
    (stype-p type)
    (expression-p body))
  (APP (expression-p fun)
    (expression-p arg))
  (VAR (varname-p name))
  (IFELSE (expression-p cond)
    (expression-p case1)
    (expression-p case2)))

```

As an aid to writing functions involving sum types, we have written a companion macro which performs pattern matching, binding variables to corresponding parts of an input term when the term is of the correct form. This macro, called `pattern-match`, is similar in function to `case-match`, but uses user-defined recognizers and accessors rather than operating directly on the list structure. `Defsum` produces a form for each product defined which allows `pattern-match` to recognize each product.

An additional difficulty impeding the adaptation of language metatheory to mechanical theorem proving technology is the representation of variable bindings. Hand proofs in this field usually ignore the technically difficult problems involved in formalizing these notions and instead assume the nonexistence of variable name conflicts. When we defined the evaluation relation in section 2, we did not define the concept of capture-avoiding substitution. In discussing a more complicated language, there would be many more such omissions. In formalizing a language’s semantics, it isn’t possible to gloss over these. Several approaches for modeling variables and binding have been developed. One approach is *higher-order abstract syntax* [4], in which binding constructs are represented by functions in the underlying language – this approach requires support for high-order functions, so it cannot be used in ACL2. Another popular strategy is to adopt

a canonical form, such as de Bruijn notation [3], which replaces names with binding offsets. To follow the traditional presentation of syntax and typing rules as closely as possible, we have adopted the traditional approach in which variables are named and may be renamed in case of a conflict. This approach involves a manageable amount of work, namely defining the capture-avoiding substitution and proving a lemma about the effect of α -substitution on the typing relation.

4. PROGRESS AND PRESERVATION

The soundness of the simply-typed λ -calculus is stated as two theorems. In combination they suffice to show that repeated evaluations of well-typed expressions can never result in an expression which is stuck in the sense that it has no evaluations but is not a value. The theorems are stated as follows.

Progress. If X is a well-typed expression, then either it is a value or there exists some expression X' such that $X \rightsquigarrow X'$.

Preservation. If X is a well-typed expression and $X \rightsquigarrow X'$, then X' is also a well-typed expression.

A “well-typed expression” here is X such that $\vdash X : T$ for some T ; that is, the expression must have a type in the empty context. However, the preservation theorem as stated is too weak to prove by induction; we instead prove the theorem for an arbitrary context instead of just the empty context, and also require that the type of X' is the same as the type of X .

Preservation. If $\Gamma \vdash X : T$ and $X \rightsquigarrow X'$, then $\Gamma \vdash X' : T$.

In proving these properties, we found that it is extremely helpful to use an explicit representation for derivations of the typing and evaluation relations. The following defsum forms define their syntax.

```
(defsum type-deriv
  (T-TRUE)
  (T-FALSE)
  (T-VAR)
  (T-ABS (type-deriv-p body))
  (T-APP (stype-p argtype)
         (type-deriv-p fun)
         (type-deriv-p arg))
  (T-IF (type-deriv-p cond)
        (type-deriv-p case1)
        (type-deriv-p case2)))
```

```
(defsum eval-deriv
  (E-APPABS)
  (E-APP1 (eval-deriv-p fun))
  (E-APP2 (eval-deriv-p arg))
  (E-IFCOND (eval-deriv-p cond))
  (E-IFTRUE)
  (E-IFFALSE))
```

A `type-deriv` object represents a proof that a certain typing relation holds. Given a typing context, expression, type,

and type derivation, we can recursively check that the type derivation corresponds to correct instantiations of the rules defining the typing relation. We define this operation in the function `valid-typing`. Similarly, we check for a valid evaluation relation in the function `valid-evaluation`, which takes two expressions and an evaluation derivation.

The progress and preservation theorems, stated in terms of typing and evaluation derivations, illuminate a new path toward proving them:

Progress. If there exists a valid derivation of $\vdash X : T$, then either X is a value or there exists an expression X' for which there is a valid derivation of $X \rightsquigarrow X'$.

Preservation. If there exist valid derivations of $\Gamma \vdash X : T$ and $X \rightsquigarrow X'$, then there exists a valid derivation of $\Gamma \vdash X' : T$.

The statements of these theorems make the method of proof clear. Given the objects mentioned in the hypotheses, we need to exhibit the objects postulated in the conclusions. We therefore define functions that produce the necessary objects. These functions are `next-expr` and `progress-deriv` which produce the new expression and evaluation derivation for progress, and `preservation-deriv` which produces the typing derivation for preservation. These are the statements of our progress and preservation theorems:

```
(defthm progress
  (implies
    (and (valid-typing nil expr type deriv)
         (not (value-p expr)))
    (valid-evaluation
     expr
     (next-expr expr type deriv)
     (progress-deriv expr type deriv))))

(defthm preservation
  (implies
    (and (valid-typing cntxt expr type type-deriv)
         (valid-evaluation expr expr2 eval-deriv))
    (valid-typing
     cntxt expr2 type
     (preservation-deriv
      cntxt expr type type-deriv eval-deriv))))
```

The progress theorem is simple to prove using the induction schema of the derivation-producing function, which recurses on the structure of the typing derivation. The proof of preservation is a straightforward induction on the evaluation derivation except in the `E-APPABS` case, in which it must be shown that a substitution of an expression for a variable of the same type preserves the type of the outer expression. This substitution lemma is the largest part of the proof effort for the soundness theorems. It requires three sublemmas, listed below, which must be invoked in appropriate locations within its derivation building function. The textbook proof of preservation (as in [5], for example) uses all of these lemmas but the third, which is necessitated by our use of explicitly named variables.

Permutation. If $\Gamma \vdash X : T$ and Γ' assumes the same types

as Γ for all variables appearing in Γ , then $\Gamma' \vdash X : T$.

```
(defthm permutation
  (implies
    (and (valid-typing cntxt1 expr type deriv)
         (env-same-bindings cntxt1 cntxt2))
    (valid-typing
      cntxt2 expr type
      (permutation-deriv cntxt1 deriv))))
```

Weakening. If $\Gamma, \Delta \vdash X : T$ and v does not appear in X , then $\Gamma, v : T_v, \Delta \vdash X : T$.

```
(defthm weakening
  (implies
    (and (valid-typing cntxt1 expr type deriv)
         (not (is-used-in var expr))
         (is-suffix suffix cntxt1)
         (equal cntxt2
              (insert-assoc
                var t2 suffix cntxt1)))
    (valid-typing
      cntxt2 expr type
      (weakening-deriv
        cntxt1 var t2 suffix deriv))))
```

Variable name substitution. If $\Gamma \vdash X : T$ and v_2 does not appear in X , then $[v_1 \mapsto v_2]\Gamma \vdash [v_1 \mapsto v_2]X : T$.

```
(defthm alpha-subst-ok
  (implies
    (and (valid-typing cntxt expr type deriv)
         (not (is-used-in var2 expr))
         (alpha-subst-env-okp
          var1 var2 suff cntxt)
         (equal cntxt2
              (env-subst-up-to
                var1 var2 suff cntxt)))
    (valid-typing
      cntxt2
      (alpha-subst var1 var2 expr)
      type
      (alpha-subst-deriv
        suff cntxt expr var1 var2 deriv))))
```

Substitution. If $\Gamma, v : T_v \vdash X : T$ and $\Gamma \vdash Y : T_v$, then $\Gamma \vdash [v \mapsto Y]X : T$.

```
(defthm substitution
  (implies
    (and (valid-typing cntxt val vtype vderiv)
         (valid-typing
          (cons (cons var vtype) cntxt)
          expr type deriv))
    (valid-typing
      cntxt
      (subst-expression val var expr)
      type
      (substitution-deriv
        cntxt var val vtype vderiv expr
        type deriv))))
```

For the first three of these lemmas, it is not necessary to construct a new derivation: in fact, the derivation function simply returns the derivation it is given, ignoring the other variables. In these cases it is still convenient to define a specific function to do this for each lemma. The rewrite rule

resulting from each lemma operates only on terms which are calls of **valid-typing** on calls of the appropriate derivation function. Use of these derivation functions inside other derivation-producing functions causes the theorem prover to apply the corresponding lemma as a rewrite rule. This leads to a proof style that is similar to that of hand proofs: the user specifies the high-level strategy for each proof by defining the derivation function, and ACL2 grinds through the intuitive but tedious details which might be omitted in a hand proof.

Because these derivation functions are only used when we expect their corresponding lemmas to be applicable, we find it aids in both proof debugging and speed to force the hypotheses of such lemmas and set their backchain limits to zero. The hypotheses for each application of a lemma are then relieved in a separate forcing round, so that if there is a problem relieving one of the hypotheses, it is easy to see what lemma was being tried and why the proof failed.

5. CONCLUSIONS

Our method of reasoning about language metatheory consists of three major choices: first, a systematic and automated representation of terms, types, and derivations; second, explicit symbolic naming of variables; and third, the phrasing of lemmas as rewrite rules which are triggered by the presence of a particular function call, allowing us to explicitly and systematically guide the theorem prover. In future work, we hope to improve our treatment of variable naming in order to scale our method to more complex languages. One particular goal is to solve the PoplMark Challenge [1]. $F_{<}$, the language used in the challenge, extends the λ -calculus with type abstractions and a subtyping relation. We have observed that these extensions significantly complicate reasoning about the naming of bound variables.

6. REFERENCES

- [1] AYDEMIR, B. E., BOHANNON, A., FAIRBAIRN, M., FOSTER, J. N., PIERCE, B. C., SWEWLL, P., VYTINIOTIS, D., WASHBURN, G., WEIRICH, S., AND ZDANCEWIC, S. Mechanized metatheory for the masses: The PoplMark challenge. In *Proc. of the 18th International Conference on Theorem Proving in Higher Order Logics (TPHOLs)* (2005).
- [2] BROCK, B. **defstructure** for ACL2 Version 2.0, 1997.
- [3] DE BRUIJN, N. G. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation. In *Indagationes Mathematicae* (1972), vol. 34, pp. 381–392.
- [4] PFENNING, F., AND ELLIOT, C. Higher-order abstract syntax. In *Proc. of the Conf. on Programming Language Design and Implementation (PLDI)* (1988), ACM Press, pp. 199–208.
- [5] PIERCE, B. C. *Types and Programming Languages*. MIT Press, 2002.