

A Coverage Checking Algorithm for LF

Carsten Schürmann¹ and Frank Pfenning² *

¹ Yale University

New Haven, CT, USA, email: `carsten@cs.yale.edu`

² Carnegie Mellon University

Pittsburgh, PA, USA, email: `fp@cs.cmu.edu`

Abstract. Coverage checking is the problem of deciding whether any closed term of a given type is an instance of at least one of a given set of patterns. It can be used to verify if a function defined by pattern matching covers all possible cases. This problem has a straightforward solution for the first-order, simply-typed case, but is in general undecidable in the presence of dependent types. In this paper we present a terminating algorithm for verifying coverage of higher-order, dependently typed patterns. It either succeeds or presents a set of counterexamples with free variables, some of which may not have closed instances (a question which is undecidable). Our algorithm, together with strictness and termination checking, can be used to certify the correctness of numerous proofs of properties of deductive systems encoded in a system for reasoning about LF signatures.

1 Introduction

Coverage checking is the problem of deciding whether any closed term of a given type is an instance of at least one of a given set of patterns. This has a number of applications: in functional programming, it is used to decide if a given set of cases defining a function is exhaustive or not. In proof assistants it is used to verify if a purported proof covers all possible cases. Depending on the application, the underlying term algebra, meta-theoretic requirements, and efficiency considerations, a variety of algorithms that implement or decide properties about pattern matching emerge. In this paper we discuss one algorithm for coverage checking in the logical framework LF [8].

The choice of the underlying term algebra is essential. In traditional functional programming languages, for example, we have only simple types and possibly prefix polymorphism, and the structure of functions is not observable by pattern matching. This makes coverage checking straightforward, both in theory and practice. In LF, on the other hand, we have dependent types and functions are intensional: their structure can be observed by pattern matching. This makes coverage checking undecidable since, for example, *any* set of patterns will cover all terms of an empty type and emptiness is undecidable.

* This work was supported in part by the National Science Foundation under grants CCR-0133502 and CCR-9988281. To appear at TPHOLs, Rome, Italy, Sep.03

Nonetheless, coverage checking has many applications in logical frameworks. Coquand’s seminal paper [2] is concerned with the natural expression of proofs in a proof checker. Moreover, in certain dependent type theories, functions defined by pattern matching are strictly more expressive than schemata of primitive recursion [9]. In addition to coverage checking, proof verification requires termination analysis which we do not consider here, but has previously been treated [19, 17]. Related is the work on Delphin [22] where functions may be defined by pattern matching over LF terms in their full generality. Verifying that such functions represent correct proofs requires coverage checking. In the context of the Twelf system [16], coverage checking is used to verify that a given logic program covers all cases for its input arguments. Besides checking the exhaustiveness of definitions, coverage checking is used to verify the correctness of meta-theoretic proofs represented as relations. Numerous case studies of this style of verification have been carried out, including cut-elimination for classical and intuitionistic logics, the Church-Rosser theorem for the untyped λ -calculus, and various translations between logical systems (see [14] for an introduction and survey). The running example in this paper is the correctness of bracket abstraction, which is a critical step in the translation from natural deduction to derivations from Hilbert’s axioms.

As mentioned above, the coverage checking problem is in general undecidable. Our approach has been to design a sound approximation that always terminates and either certifies coverage or produces a set of potential counterexamples. Sometimes, these counterexamples are in effect impossible (due to dependent types). This algorithm has been fully implemented in the Twelf system and has proved enormously useful in practice to verify the correctness of meta-theoretic proofs expressed as relations. The largest project undertaken so far is Crary’s implementation of foundational typed assembly language with about 30,000 lines of Twelf code [3] with more than 1000 theorems; other examples have been mentioned above. One valuable experience gained through these experiments is that in the case of failure with spurious counterexamples it is generally possible to make a proof more explicit in such a way that it then passes the coverage checker.

The remainder of the paper is organized as follows: In Section 2 we briefly introduce the logical framework LF, and sketch the most important concepts necessary for this paper. In Section 3 we describe the coverage problem in detail and present our coverage algorithm. We discuss some related work in Section 4 and assess results and conclude in Section 5.

2 LF

The type-theoretic foundation for this paper is the logical framework LF [8]. In addition to the standard syntactic categories of objects, types, and kinds, we will also use substitutions in a critical way throughout this paper so we briefly introduce them here (see also, for example, [1]).

```

imp  : o → o → o.
forall : (i → o) → o.
k    : hil (imp A (imp B A)).
s    : hil (imp (imp A (imp B C)) (imp (imp A B) (imp A C))).
mp   : hil (imp A B) → hil A → hil B.

```

Fig. 1. The Propositional Hilbert Calculus

```

Kinds      K ::= type | Πx:A. K
Atomic Types B ::= a | B M
Types      A ::= B | Πx:A1. A2
Objects    M ::= x | c | M1 M2 | λx:A. M
Signatures Σ ::= · | Σ, a : K | Σ, c : A
Contexts   Γ ::= · | Γ, x : A
Substitutions σ ::= · | σ, M/x

```

We write a for constant type families, x or u for object-level variables, and c for constructors. A *term* may come from any of the syntactic levels. As usual, we identify α -equivalent terms. In order to state certain definitions and propositions more concisely, we write U to stand for either an object or a type and V for either a type or a kind and h for a family-level or object-level constant. We take $\beta\eta$ -conversion as the notion of definitional equality [8, 1], for which we write $U \equiv U'$ and $V \equiv V'$. Substitutions are capture-avoiding and written as $U[\sigma]$ or $V[\sigma]$ with the special form $U[M/x]$ and $V[M/x]$. Often, we write Δ for contexts that are interpreted existentially, and Γ for universal ones. When we write $\Gamma[\sigma]$, it is a shorthand for applying σ in left to right order to each variable type in Γ . Signatures, contexts, and substitutions may not declare a variable or constant more than once, and renaming of bound variables may be applied tacitly to ensure that condition. Besides equality, the main judgment is typing $\Gamma \vdash U : V$, suppressing the fixed signature Σ . We always assume our signatures, contexts and types to be valid.

Example 1 (Propositional Hilbert Calculus). An LF signature that encodes the Hilbert calculus in LF is given in Figure 1. Terms are captured by type `i`, formulas by `o`, and proofs in the Hilbert calculus by the type family `hil : o → type`. As usual we omit the leading Π -binders of variables whose types are inferable, denoted above by uppercase variable names. \square

Example 2 (Bracket abstraction). Figure 2 displays a relational representation of bracket abstraction. The type family `ded` relates hypothetical Hilbert derivations (`hil A → hil B`) with their categorical counterpart `hil (A imp B)`. Each of the four defining cases corresponds to one of the constant declarations. Again, in the interest of space we have omitted the implicit Π -abstractions. \square

```

ded : (hil A → hil B) → hil (A imp B) → type.
ded_id : ded (λx:hil A. x) (mp (mp s k) k).
ded_k : ded (λx:hil A. k) (mp k k).
ded_s : ded (λx:hil A. s) (mp k s).
ded_mp :
  ded (λx:hil A. H2 x) H'2 → ded (λx:hil A. H1 x) H'1
  → ded (λx:hil A. mp (H1 x) (H2 x)) (mp (mp s H'1) H'2).

```

Fig. 2. Bracket abstraction

Type-checking and definitional equality on well-typed terms for LF are decidable. Every term is equal to a unique β -normal η -long form which we call *canonical form*. In the remainder of the paper we assume that all terms are in canonical form, because it simplifies the presentation significantly. In the implementation this is achieved incrementally, first by an initial conversion of input terms to η -long form and later by successive weak-head normalization as terms are traversed.

Since it is perhaps not so well-known, we will give only the typing rules for substitutions, which are used pervasively in this paper.

$$\frac{}{\Gamma' \vdash \cdot : \cdot} \quad \frac{\Gamma' \vdash \sigma : \Gamma \quad \Gamma' \vdash M : A[\sigma]}{\Gamma' \vdash (\sigma, M/x) : (\Gamma, x:A)}$$

For a context $\Gamma = (x_1:A_1, \dots, x_n:A_n)$, we define $\text{id}_\Gamma = (x_1/x_1, \dots, x_n/x_n)$ so that $\Gamma \vdash \text{id}_\Gamma : \Gamma$.

Composition of substitutions is defined by $(\cdot) \circ \sigma = (\cdot)$ and $(M/x, \theta) \circ \sigma = (M[\sigma]/x, \theta \circ \sigma)$. We will only apply a substitution $\Gamma' \vdash \sigma : \Gamma$ to a term $\Gamma' \vdash U : V$ or a substitution $\Gamma' \vdash \theta : \Gamma''$ resulting in $\Gamma' \vdash U[\sigma] : V[\sigma]$ and $\Gamma' \vdash \theta \circ \sigma : \Gamma''$, respectively.

3 Coverage

In this section we first formally define the problem of coverage in the LF type theory in Section 3.1. This relies on *higher-order matching*, a problem whose decidability is an open question. We therefore identify an important subclass, the *strict coverage problems* (Section 3.2) which guarantee not only decidability but also uniqueness of matching substitutions. All examples we have ever encountered in practice belong to this class and we explain the reasons for this after the necessary definitions. Then we define *splitting* in Section 3.3, which is the second critical operation to be performed during coverage checking. Next we describe our basic coverage algorithm and prove it sound and terminating in Sections 3.4 and 3.5. The last component of our coverage checker is *finitary splitting*, discussed and proved correct in Section 3.6.

```

ded0 :  $\Pi A:o. \Pi B:o. (\text{hil } A \rightarrow \text{hil } B) \rightarrow \text{type}.$ 
ded0_id :  $\Pi A:o. \text{ded}_0 A A (\lambda x:\text{hil } A. x).$ 
ded0_k :  $\Pi A:o. \Pi A_1:o. \Pi A_2:o.$ 
          $\text{ded}_0 A (\text{imp } A_1 (\text{imp } A_2 A_1)) (\lambda x:\text{hil } A. k A_1 A_2).$ 
ded0_s :  $\Pi A:o. \Pi A_1:o. \Pi A_2:o. \Pi A_3:o.$ 
          $\text{ded}_0 A (\text{imp } (\text{imp } A_1 (\text{imp } A_2 A_3)) (\text{imp } (\text{imp } A_1 A_2) (\text{imp } A_1 A_3)))$ 
          $(\lambda x:\text{hil } A. s A_1 A_2 A_3).$ 
ded0_mp :  $\Pi A:o. \Pi A_1:o. \Pi A_2:o. \Pi H_1:\text{hil } A \rightarrow \text{hil } (\text{imp } A_2 A_1).$ 
          $\Pi H_2:\text{hil } A \rightarrow \text{hil } A_2. \text{ded}_0 A A_1 (\lambda x:\text{hil } A. \text{mp } A_2 A_1 (H_1 x) (H_2 x)).$ 

```

Fig. 3. Patterns

3.1 Definition of Coverage

A coverage goal is simply a term (object or type) with some free variables. Intuitively, a coverage goal stands for all of its closed instances. In order to emphasize the interpretation of the variables as standing for closed terms, we write Δ for such contexts and denote variables in Δ by u and v rather than x and y . The distinction between Δ and Γ can be formalized (see [18]), but this is not necessary for the present purposes.

A coverage problem is given by a goal and a set of patterns. One can think of these as the patterns of a case expression in a functional program, or the input terms in the clause heads of a logic program. In the general case, a set of patterns is just a set of terms with free variables.

Example 3. Consider the type family `ded` from Example 2 that encodes the algorithm of bracket abstraction. To satisfy coverage, `ded` must provide a case covering every possible input argument upon invocation, independent of the output. Thus, in Figure 3, the type family `ded0` is a reduction of `ded` to set of patterns for the coverage goal

$$u : o, v : o, w : (\text{hil } u \rightarrow \text{hil } v) \vdash \text{ded}_0 u v w : \text{type}. \quad (1)$$

Note that for the sake of completeness all arguments to `ded0` are explicit. \square

Definition 1 (Immediate Coverage). *We say a coverage goal $\Delta \vdash U : V$ is immediately covered by a collection of patterns $\Delta_i \vdash U_i : V_i$ if there is an i and a substitution $\Delta \vdash \sigma_i : \Delta_i$ such that $\Delta \vdash U \equiv U_i[\sigma_i] : V$.*

Example 4. Let p be a propositional symbol, $A = B = \text{imp } p p$ and $M = \lambda u : \text{hil } (\text{imp } p p). u$. The goal

$$p : o \vdash \text{ded}_0 (\text{imp } p p) (\text{imp } p p) (\lambda u : \text{hil } (\text{imp } p p). u) : \text{type}$$

is immediately covered by `ded0.id`.

Coverage has an infinitary definition, requiring immediate coverage of every ground instance of a goal.

Definition 2 (Coverage). We say $\Delta \vdash U : V$ is covered by a collection of patterns $\Delta_i \vdash U_i : V_i$ if every ground instance $\cdot \vdash U[\tau] : V[\tau]$ for $\cdot \vdash \tau : \Delta$ is immediately covered by the collection of terms $\Delta_i \vdash U_i : V_i$.

In this formulation the problem of coverage is very general, because the type of U and the type of the U_i 's may not be the same. It turns out that the algorithm is significantly easier to describe and prove correct if we restrict U and U_i to be types, and $V = V_i = \text{type}$.

The implementation in Twelf transforms any coverage problems that arise into this type-level form. This translation is straightforward and only sketched here. Given a coverage goal $\Delta \vdash M : A'$. Assume first that $A' = a' N_1 \dots N_k$ for $a' : \Pi x_1:A_1 \dots \Pi x_n:A_n.\text{type}$. In this case we declare a new type family $a : \Pi x_1:A_1 \dots \Pi x_n:A_n.a' x_1 \dots x_n \rightarrow \text{type}$. The new coverage goal is now simply $\Delta \vdash a N_1 \dots N_k M : \text{type}$. All patterns are transformed in the analogous way, using the same a to replace a' . If A' starts with some leading Π -quantifiers we carry them over from the general to the restricted form.

To summarize, without loss of generality, in the remainder of this paper we consider only coverage goals of the form $\Delta \vdash A : \text{type}$ and patterns of the form $\Delta_i \vdash A_i : \text{type}$.

3.2 Strict Patterns

To determine if a goal is immediately covered we have to solve a higher-order matching problem, instantiating the patterns A_i to match the goal A . Not incidentally, this is also the operation that is performed when matching a case subject against the patterns in each arm of a case branch, or when unifying the input arguments to a predicate with the clause head on a well-moded predicate.

In order for this pattern matching to be decidable (for the coverage algorithm) and also so that the operational semantics is well-defined (for the execution of a functional or logic program), we require the patterns to be *strict*. Strictness for a pattern $\Delta_i \vdash A_i : \text{type}$ requires that each variable in Δ_i must occur in A at least once in a rigid position [10, 15].

Definition 3 (Strictness). We say that u has a strict occurrence in U if $\Delta; \Gamma \vdash_u U$ as defined by the rules depicted in Figure 4. A pattern $\Delta_i \vdash A_i : \text{type}$ is strict if $\Delta_i; \cdot \vdash_u A_i$ for each variable u in Δ_i .

Informally, an occurrence of u is strict if it is not below another variable in Δ and if that occurrence forms a higher-order pattern in the sense of Miller [13], that is, u is applied to distinct parameters as expressed by the judgment $\Gamma \vdash u x_1 \dots x_n \text{ pat}$. Unlike higher-order patterns in the sense of Miller, however, other forms of occurrences of u are allowed, which is a practically highly significant generalization. All of the examples in Twelf are strict, but many higher-order examples are not patterns in the sense of Miller. Strictness is sufficient here because we are only interested in matching and not full unification.

$$\begin{array}{c}
\frac{\Delta; \Gamma \vdash_u A}{\Delta; \Gamma \vdash_u \lambda y : A. M} \text{ls_ld} \quad \frac{\Delta; \Gamma, y : A \vdash_u M}{\Delta; \Gamma \vdash_u \lambda y : A. M} \text{ls_lb} \\
\frac{\Delta; \Gamma \vdash_u A_1}{\Delta; \Gamma \vdash_u \Pi y : A_1. A_2} \text{ls_pd} \quad \frac{\Delta; \Gamma, y : A_1 \vdash_u A_2}{\Delta; \Gamma \vdash_u \Pi y : A_1. A_2} \text{ls_pb} \\
\frac{\Delta; \Gamma \vdash_u M_i}{\Delta; \Gamma \vdash_u c M_1 \dots M_n} \text{ls_c} \quad (1 \leq i \leq n) \quad \frac{\Delta; \Gamma \vdash_u M_i}{\Delta; \Gamma \vdash_u a M_1 \dots M_n} \text{ls_a} \quad (1 \leq i \leq n) \\
\frac{y : A \in \Gamma \quad \Delta; \Gamma \vdash_u M_i}{\Delta; \Gamma \vdash_u y M_1 \dots M_n} \text{ls_var} \quad (1 \leq i \leq n) \\
\frac{\Gamma \vdash_u x_1 \dots x_n \text{ pat}}{\Delta; \Gamma \vdash_u u x_1 \dots x_n} \text{ls_pat} \quad \begin{array}{l} \text{no rule for } \Delta; \Gamma \vdash_u v M_1 \dots M_n \\ \text{for } u \neq v, v : A \in \Delta \end{array}
\end{array}$$

Fig. 4. A formal system for strictness

Theorem 1. *Given a coverage goal $\Delta \vdash A : \text{type}$ and a strict pattern $\Delta_i \vdash A_i : \text{type}$. Then it is decidable if there is a substitution $\Delta \vdash \sigma : \Delta_i$ such that $A \equiv A_i[\sigma]$. Moreover, if such a substitution exists it is uniquely determined.*

Proof. See [21].

If such a σ exists then the coverage goal in question is covered by one of the patterns. On the other hand, if no such σ exist, either there exists a pattern that covers the coverage goal but does not happen to cover it immediately, or none of the patterns will ever cover it, in which case coverage should fail.

3.3 Splitting

In the case that the coverage goal is not immediately covered by a pattern, there is only one way to make progress on the coverage problem. Assuming we can recognize coverage failure easily (which we can, as discussed in Section 3.4), we can bridge the gap between coverage goal and patterns by splitting the coverage goal into a set of new coverage goals, by partially instantiating free variables in Δ . Since there may be infinitely many ground instances, we instantiate the coverage goal only partially, one layer at a time.

Definition 4 (Non-redundant complete set of substitutions).

Let $\Delta \vdash A : \text{type}$ be a coverage goal. We say a finite collection $\Delta_i \vdash \tau_i : \Delta$ is a non-redundant complete set of substitutions if for every $\cdot \vdash \tau : \Delta$ there exists a unique i and a unique $\cdot \vdash \sigma_i : \Delta_i$ such that $\cdot \vdash \tau = \tau_i \circ \sigma_i$.

Splitting a coverage goal using a non-redundant complete set of substitutions is conservative with respect to the coverage property, because the original coverage goal is covered if and only if the resulting set of coverage goals is covered.

Theorem 2 (Conservativity). *Let $\Delta \vdash A : \text{type}$ be a coverage goal and $\Delta_i \vdash \tau_i : \Delta$ a non-redundant complete collection of substitutions. All $\Delta_i \vdash A[\tau_i] : \text{type}$ are covered by a given set of patterns if and only if $\Delta \vdash A : \text{type}$ is covered.*

Proof. Coverage depends only on the set of ground instances of a coverage goal. But the collection of all ground instances of $\Delta_i \vdash A[\tau_i]$ is exactly the same as the set of ground instances of $\Delta \vdash A : \text{type}$ since the τ_i form a complete set. \square

The splitting operation discussed in the remainder of this section generates a non-redundant complete set of substitutions. Its definition is inspired by a similar operation in ALF [2] which in turn has its root in the basic steps of Huet’s algorithm for higher-order unification [10].

First, among all coverage goals that are not immediately covered, splitting selects coverage goal $\Delta \vdash A : \text{type}$ and a declaration $u:A_j$ from its context Δ that is below referred to as *splitting variable*. A_j may be a function type, therefore, without loss of generality, it is of the following form

$$\Delta \vdash u : \Pi \Gamma. a M_1 \dots M_m.$$

For the sake of conciseness, we consolidate all successive Π -binders into one context Γ . This is only an abbreviation and does not properly extend LF. We also use the following abbreviations $h \Gamma$ which stands for $h x_1 \dots x_m$ if $\Gamma = x_1:A_1, \dots, x_m:A_m$ where h is a constant or a parameter. Furthermore $h(\Delta \Gamma)$ is a shorthand for $h(u_1 \Gamma) \dots (u_n \Gamma)$ for $\Delta = u_1:A'_1, \dots, u_n:A'_n$.

Next, splitting considers all possible top-level structures of a term $M : \Pi \Gamma. a M_1 \dots M_m$. Because of the existence of canonical forms, each structure is determined by a constant declared in the signature or a parameter declared in the local context that occur in head position in M . Without loss of generality, the respective type of the constant or parameter has the form $\Pi u_1:A_1. \dots \Pi u_n:A_n. B$ with an atomic type B . Each u_i is intuitively interpreted as an existential variable that can be instantiated to terms valid in Γ . To account for those local dependencies, splitting raises them therefore by Γ and turns all u_i into variables of functional type.

Definition 5 (Raising). *Let Γ be a context of local parameters, A the type of a constant c . Raising A by Γ yields a $\langle \Delta \vdash A' \rangle$, a context Δ of raised existential variables and a raised type A' (that always has the form $\Pi \Gamma. B$).*

$$\text{raise}\langle \Gamma \vdash A \rangle = \begin{cases} \langle \cdot \vdash \Pi \Gamma. A \rangle & \text{if } A \text{ is atomic} \\ \langle u : \Pi \Gamma. A_1, \Delta \vdash A' \rangle & \text{if } A = \Pi u:A_1. A_2 \\ & \text{and } \langle \Delta \vdash A' \rangle = \text{raise}\langle \Gamma \vdash A_2[u \Gamma/u] \rangle \end{cases}$$

If a constant or a parameter are valid candidates for a head position of M is decided by higher-order pattern unification that postpones unresolved unification equation as constraints [4]. Although we allow constraints to arise during the process of unification, we require that after completion all constraints have been resolved. Otherwise the specified variable cannot be split, and splitting must select another variable $u : A \in \Delta$ instead. Unfortunately, successive selections of

splitting variables are not independent and it is possible that some sequences of splitting operations fail (with spurious counterexamples) while other sequences could succeed. In principle splitting could backtrack here, but this is currently not implemented.

Definition 6 (Splitting). *Let $\Delta \vdash A$: type a coverage goal, and u in $\Delta = \Delta_1, u : \Pi\Gamma. B_u, \Delta_2$ a splitting variable. The splitting operation considers each constant c declared in the signature Σ and each local parameter y declared in Γ in turn, and determines a set of substitutions σ_c, σ_y as follows.*

1. Constants: *Let $c : \Pi\Delta_c. B_c \in \Sigma$, and $\langle \Delta'_c \vdash \Pi\Gamma. B'_c \rangle = \text{raise}\langle \Gamma \vdash \Pi\Delta_c. B_c \rangle$. Let $\Delta' \vdash \sigma_c : \Delta, \Delta'_c$ be the most general unifier of the higher-order unification problem*

$$\exists \Delta. \exists \Delta'_c. (\Pi\Gamma. B_u \approx \Pi\Gamma. B'_c) \wedge (u \approx \lambda\Gamma. c (\Delta'_c \Gamma)) \quad (2)$$

if it exists.

2. Bound Variables: *Let $y : \Pi\Delta_y. B_y \in \Gamma$, and $\langle \Delta'_y \vdash \Pi\Gamma. B'_y \rangle = \text{raise}\langle \Gamma \vdash \Pi\Delta_y. B_y \rangle$. Let $\Delta' \vdash \sigma_y : \Delta, \Delta'_y$ be the most general unifier of the higher-order unification problem*

$$\exists \Delta. \exists \Delta'_y. (\Pi\Gamma. B_u \approx \Pi\Gamma. B'_y) \wedge (u \approx \lambda\Gamma. y (\Delta'_y \Gamma)) \quad (3)$$

if it exists.

Since we collect all such most general unifiers, cases for which the unification problem fails¹ simply do not contribute a substitution to the result of the splitting operation.

Example 5. Recall the original coverage goal (1): $u : \circ, v : \circ, w : (\text{hil } u \rightarrow \text{hil } v) \vdash \text{ded}_0 u v w$: type. Splitting w results in four new coverage goals, because x , k , s , or mp are all possible head constants for functions of type $\text{hil } u \rightarrow \text{hil } v$.

$$\begin{aligned} & u : \circ \vdash \text{ded}_0 u u (\lambda x : \text{hil } u. x) : \text{type} \\ & u : \circ, u_1 : \circ, u_2 : \circ \\ & \quad \vdash \text{ded}_0 u (\text{imp } u_1 (\text{imp } u_2 u_1)) (\lambda x : \text{hil } u. k u_1 u_2) : \text{type} \\ & u : \circ, u_1 : \circ, u_2 : \circ, u_3 : \circ \\ & \quad \vdash \text{ded}_0 u (\text{imp } (\text{imp } u_1 (\text{imp } u_2 u_3)) \\ & \quad \quad (\text{imp } (\text{imp } u_1 u_2) (\text{imp } u_1 u_3))) (\lambda x : \text{hil } u. s u_1 u_2 u_3) : \text{type} \\ & u : \circ, u_1 : \circ, u_2 : \circ, v : \text{hil } u \rightarrow \text{hil } (\text{imp } u_2 u_1), w : \text{hil } u \rightarrow \text{hil } u_2 \\ & \quad \vdash \text{ded}_0 u u_1 (\lambda x : \text{hil } u. \text{mp } (u_2 x) u_1 (v x) (w x)) : \text{type} \end{aligned}$$

The main result of this section is that splitting generates always a set of substitutions that is non-redundant and complete.

Theorem 3 (Splitting is non-redundant and complete). *The set of substitutions generated by splitting is non-redundant and complete.*

Proof. From properties of complete sets of most general unifiers and lemmas about raising (see [21]). \square

¹ but not those whose results are indeterminate because of residual equations, which are not permitted

3.4 The Coverage Algorithm

Recall that a coverage goal $\Delta \vdash A : \text{type}$ is immediately covered by a collection of terms $\Delta_i \vdash A_i : \text{type}$ if there is an i and $\Delta \vdash \sigma_i : \Delta_i$ such that $\Delta \vdash A \equiv A_i[\sigma_i] : \text{type}$.

Immediate coverage is central to the naive, non-deterministic coverage algorithm which we discuss next. We assume we have a set of coverage goals, all of which must be covered for the algorithm to succeed. In the first step, this is initialized with the goal $\Delta \vdash A : \text{type}$. We pick one of the coverage goals and determine, via strict higher-order matching, if it is immediately covered by any covering type A_i . If so we remove it from the set and continue. If not, we non-deterministically select a variable in the coverage goal and split it into multiple goals, which replace it in the collection of coverage goals. This coverage algorithm is naive because it may not terminate: even if types are non-empty and coverage holds, splitting the wrong variable can lead to non-termination.

The procedure we propose in this section always terminates and either indicates that coverage holds, or outputs a set of potential counterexamples. Some of these may fail to be actual counterexamples, because we may not be able to instantiate the remaining variables to a ground term that is not covered. If the counterexample is ground, however, it is guaranteed to be an actual counterexample. We analyze the possible forms of counterexamples in more detail at the beginning of Section 3.6.

The basic idea is to record *why* immediate coverage fails and not just *if* it does. Assume we are given a coverage goal $\Delta \vdash A : \text{type}$ and a pattern $\Delta' \vdash A' : \text{type}$. Instead of just applying our matching algorithm, we then construct a conjunction of equations E and the symbols \top (success) or \perp (failure) such that $\langle \Delta, \Delta'; \cdot \vdash A < A' \rangle \Longrightarrow E$. This is accomplished by using the rules for the judgment

$$\langle \Delta; \Gamma \vdash U < U' \rangle \Longrightarrow E$$

defined in Figure 5, which reads as: *Match U against pattern U' in the parameter context Γ to obtain the residual equations E .* Δ is the disjoint union of the (existential) variables in U and U' , of which only those in U may be instantiated during matching. Initially, the context Γ is always empty, and both U and U' are types. However, internally we require the context Γ of shared local parameters.

Example 6 (Rigid Matching Algorithm). The rigid matching algorithm compares the goal (1) from Example 3 to each of the four cases defining ded_0 . For each case, it produces a set of residual equations (with contexts elided).

1. Case $\text{ded}_0\text{-id}$. $u \equiv A \wedge v \equiv A \wedge w \equiv (\lambda x : \text{hil } A. x)$
2. Case $\text{ded}_0\text{-k}$. $u \equiv A \wedge v \equiv (\text{imp } A_1 (\text{imp } A_2 A_1))$
 $\wedge w \equiv (\lambda x : \text{hil } A. k A_1 A_2)$
3. Case $\text{ded}_0\text{-s}$. $u \equiv A \wedge v \equiv (\text{imp } (\text{imp } A_1 (\text{imp } A_2 A_3)))$
 $(\text{imp } (\text{imp } A_1 A_2) (\text{imp } A_1 A_3)))$
 $\wedge w \equiv (\lambda x : \text{hil } A. s A_1 A_2 A_3)$
4. Case $\text{ded}_0\text{-mp}$. $u \equiv A \wedge v \equiv A_1$
 $\wedge w \equiv (\lambda x : \text{hil } A. \text{mp } (A_2 x) A_1 (H_1 x) (H_2 x))$

$$\begin{array}{c}
\frac{\langle \Delta; \Gamma, x:A \vdash U < U' \rangle \Longrightarrow E}{\langle \Delta; \Gamma \vdash \lambda x:A. U < \lambda x:A. U' \rangle \Longrightarrow E} \\
\frac{\langle \Delta; \Gamma \vdash A < A' \rangle \Longrightarrow E_1 \quad \langle \Delta; \Gamma, x:A' \vdash B < B' \rangle \Longrightarrow E_2}{\langle \Delta; \Gamma \vdash \Pi x:A. B < \Pi x:A'. B' \rangle \Longrightarrow E_1 \wedge E_2} \\
\frac{}{\langle \Delta; \Gamma \vdash \Pi x:A. B < a \dots \rangle \Longrightarrow \perp} \quad \frac{}{\langle \Delta; \Gamma \vdash a \dots < \Pi x:A. B \rangle \Longrightarrow \perp} \\
\frac{c \neq c'}{\langle \Delta; \Gamma \vdash c \dots < c' \dots \rangle \Longrightarrow \perp} \quad \frac{h:A \in \Sigma, \Gamma \quad \langle \Delta; \Gamma \vdash U_i < U'_i \rangle \Longrightarrow E_i \text{ for } 1 \leq i \leq n}{\langle \Delta; \Gamma \vdash h U_1 \dots U_n < h U'_1 \dots U'_n \rangle \Longrightarrow E_1 \wedge \dots \wedge E_n} \\
\frac{x:A \in \Gamma \quad h:A' \in \Sigma, \Gamma \quad x \neq h}{\langle \Delta; \Gamma \vdash x \dots < h \dots \rangle \Longrightarrow \perp} \\
\frac{u:A \in \Delta \quad h:A' \in \Sigma, \Gamma}{\langle \Delta; \Gamma \vdash u U_1 \dots U_n < h U'_1 \dots U'_m \rangle \Longrightarrow \langle \Gamma \vdash u U_1 \dots U_n \approx h U'_1 \dots U'_m \rangle} \\
\frac{}{\langle \Delta; \Gamma \vdash U < u' U'_1 \dots U'_m \rangle \Longrightarrow \langle \Gamma \vdash U \approx u' U'_1 \dots U'_m \rangle}
\end{array}$$

Fig. 5. Rigid Matching Algorithm

We can think of the algorithm as a rigid decomposition, which corresponds to the `simplify` function in Huet's algorithm for higher-order unification. If all residual equations can be solved (and there is no \perp), then matching is successful. Otherwise, we have to interpret the equations to determine candidates for splitting that will make progress (as defined below).

Note that during rigid matching, no variable assignment takes place: where the two terms disagree, we record an equation. But if matching is not possible, we might either record an equation or return \perp .

In order to state the lemmas in the generality required for an inductive proof, we say that for $\Delta; \Gamma \vdash U : V$ and $\Delta'; \Gamma \vdash U' : V'$ that U' covers U if there is a substitution $\Delta, \Gamma \vdash \sigma, \text{id}_\Gamma : \Delta', \Gamma$ such that $\Delta, \Gamma \vdash U \equiv U'[\sigma, \text{id}_\Gamma] : V$.

Lemma 1. *If $\Delta, \Delta'; \Gamma \vdash U < U' \Longrightarrow E$ where E contains \perp , then U' does not immediately cover U or any instance of U .*

Proof. By induction on the given derivation.

Because U' cannot immediately cover any instance of U , we do not generate any candidate variables for splitting in Δ' in this case.

Lemma 2. *If $\Delta, \Delta'; \Gamma \vdash U < U' \Longrightarrow E$ where E does not contain \perp , but contains equations of the form $u \dots \approx c \dots$ or $u \dots \approx x \dots$. Then U' does not immediately cover U (but U' could possibly cover some instance of U).*

Proof. By induction on the given derivation. In the base cases, x and c are rigid and therefore cannot be instantiated to u .

In this case, any variable u occurring in an equation of the given form is added to the set of candidate variables for splitting, since it is possible that splitting might make progress.

Lemma 3. *If $\Delta, \Delta'; \Gamma \vdash U < U' \implies E$, where E does not contain \perp or equations for the form $u \dots \approx c \dots$ or $u \dots \approx x \dots$. Then any substitution $\Delta \vdash \sigma : \Delta'$ such that for each residual equation $\langle \Gamma_i \vdash U_i \approx U'_i \rangle$ in E we have $\Gamma_i \vdash U_i \equiv U'_i[\sigma, \text{id}_{\Gamma_i}]$ is a valid match and shows that U' covers U .*

Proof. Again, by induction on the given derivation. The base cases are evident. The tricky part in the inductive argument is that the two matched terms do not necessarily have the same type or kind (even though they do initially) because we postpone non-rigid equations. However, as in the case of higher-order dependently typed unification [5], it is enough to maintain well-typedness modulo postponed equations if we eventually solve them from left-to-right.

This means that if we have no candidates from the first two kinds of equations, we call a strict higher-order matching algorithm [21] on the residual equations. If this succeeds then A' covers A . Otherwise, A' does not cover A and we suggest no candidate variables for splitting because it would be difficult to guarantee termination.

When considering a particular coverage goal $\Delta \vdash A : \text{type}$, we apply the above algorithm with each pattern. If one of them immediately covers, we are done. If not, we take the union of all the suggested candidates and pick one non-deterministically. The current implementation picks the rightmost candidate in Δ , because internal dependencies might further constrain variables to its left during the splitting step. If splitting fails because higher-order unification with the algorithm in [4] can not determine a complete and non-redundant set of substitutions, then we try another candidate, and so on. If there are no remaining splitting candidate, we add the coverage goal to the set of potential counterexamples and pick another goal.

3.5 Termination

The overall structure of the algorithm is such that the splitting step replaces a coverage goal by several others. In order to show termination with respect to a simple multi-set ordering, we must show that each of the subgoals that replace a given goal are smaller according some well-founded measure.

We calculate this measure as follows. Given a coverage goal $\Delta \vdash A : \text{type}$ apply rigid matching against each pattern. Eliminate those equations that contain \perp . Among the remaining ones, consider only equations $u U_1 \dots U_n \approx h' U'_1 \dots U'_m$ where $h' = x$ or $h' = c$. Note that all candidates for splitting appear on the left-hand side of such an equation. Take the sum of the sizes of the right-hand sides as measured by the number of bound variable and constant occurrences.

When we apply splitting to any candidate variables in Δ , that is, one of the variables u that appears on the left-hand side of an equation as given above, then this measure decreases. Given a coverage goal $\Delta \vdash A : \text{type}$ and a fixed set of patterns proposed to cover it. If we split the coverage goal along a variable u suggested by rigid matching, each of the resulting subgoals has a smaller measure than the original goal.

Theorem 4. *Coverage checking terminates after a finite number of steps, yielding either an indication of coverage or a finite set of potential counterexamples.*

3.6 Finitary splitting

The failure-directed algorithm described above works well in most practical cases, within or outside the pattern fragment. There are two remaining difficulties: one are remaining constraints during splitting as discussed in Section 3.3, the other is that occasionally the generated counterexamples fail to be actual counterexamples. The latter is a common occurrence. In large part this is because meta-theoretic proofs represented as dependently typed functions or relations often have a number of cases that are impossible. Instead of explicitly proving that the cases are impossible, one usually just lists the cases that can arise if it is syntactically obvious that the others can not arise.

What are the types of spurious counterexamples that may be produced by the algorithm? The most obvious one is a coverage goal that is incompatible with all patterns, but has no ground instances. We explain below how to handle some of these case. A less obvious problem is that matching the residual equations fail because of a spurious dependency that cannot be an actual dependency because of subordination considerations. We treat this case by applying strengthening [21] to eliminate these spurious dependencies throughout the algorithm. Finally, it is possible that two distinct variables of the coverage goal fail to match, yet they must be identical because the type has only one element. Finitary splitting will often catch these cases and correctly report coverage.

In order to handle as many spurious counterexamples as possible, we extend the algorithm described above as follows. Once the algorithm terminates with a set of proposed counterexamples to coverage, we examine each such counterexample to see if we can determine if it is impossible, that is, if it quantifies over an empty type. More concretely, let $\Delta \vdash A : \text{type}$ be a counterexample, that is, coverage goal that is not covered and does not produce any splitting candidates. We now attempt to split each variable $u:A$ in Δ in turn, leading to a new set of coverage goals $\Delta_i \vdash U_i : V_i$ for $0 \leq i < n$. If $n = 0$ we know that the case is impossible.

If $n > 0$ we could, in principle, continue the algorithm recursively to see if each of the subgoals $\Delta_i \vdash U_i : V_i$ are impossible. However, in general this would not terminate (and cannot, because inhabitation is undecidable). Instead, we only continue to split further if all of the new variables $u_k : A_k$ in Δ_i have a type that is strictly subordinate to the type A [23, 21]. Otherwise, we fail and report the immediate supergoal as a potential counterexample.

Theorem 5. *Finitary splitting terminates, either with an indication that the given coverage goal has no ground instances, or failure.*

Proof. There are only a finite number of variables in a given coverage goal. During each step of splitting we either stop or obtain subgoals where a variable $u : A$ has been replaced by several variables $u_i : A_i$ each of which has a type strictly lower in the subordination hierarchy. Since this hierarchy is well-founded, finitary splitting will terminate.

This process can be very expensive. Fortunately, we have not found it to be a bottleneck in practice, because finitary splitting is applied only to remaining counterexamples. Usually, there are not many, and usually it is immediate to see that they are indeed possible because most types are actually inhabited. We do not presently try to verify if the types are actually inhabited (that is, start a theorem prover), although it may be useful for debugging purposes to distinguish between definite and potential counterexamples. However, in a future extension this could be done at the user’s direction if he or she cannot easily detect the source of the failure of coverage.

4 Related Work

Coquand has considered the problem of coverage for a type theory in the style of Martin-Löf [2]. He defines coverage and splitting in much the same way we do here, except that no matching against the structure of λ -expressions is allowed. He also suggests a non-deterministic semi-decision procedure for coverage by guessing the correct sequence of variable splits. In an implementation this split can be achieved interactively.

Most closely related to ours is the work by McBride [12]. He refines Coquand’s idea by suggesting an algorithm for successive splitting that is quite similar to ours in the first-order case. He also identifies the problem of empty types and suggest to recognize “obviously” empty types, which is a simpler variant of finitary splitting. Our main contribution with respect to McBride’s work is that we allow matching against the structure of higher-order terms which poses significant additional challenges.

Another related development is the theory of partial inductive definitions [7], especially in its finitary form [6] and the related notion of definitional reflection [20]. This calculus contains a rule schema that, re-interpreted in our context, would allow any (finite) complete set of unifiers between a coverage goal $\Delta \vdash A : \text{type}$ and the heads of the clauses defining A . Because of the additional condition of so-called a-sufficiency for the substitutions, this was never fully automated. Also, it appears that a simple, finite complete set of unifiers was computed as in the splitting step, but that the system could not check whether an arbitrary given set of premises could be obtained as a finite complete set of unifiers.

In the Coq system [11] functions defined by patterns can be compiled to functions defined by standard primitive recursive elimination forms. Because of

the requirement to compile such functions back into pure Coq and the lack of matching against functional expressions, the algorithm is rather straightforward compared to our coverage checker and does not handle variable dependencies, non-linearity, or empty types. It does, however, treat polymorphism which we have not considered.

5 Conclusion

We have presented a solution to the coverage checking problem for LF, generalizing and extending previous approaches. The central technical developments are strict patterns (which significantly generalize higher-order patterns in the sense of Miller), strict higher-order matching, splitting in the presence of full higher-order unification, and a two-phase control structure to guarantee termination of the algorithm.

Our coverage algorithm is sound and terminating, but it is necessarily incomplete. Applied to a given set of patterns, it either reports “yes”, or it generates a set of potential counterexamples, which often contain the vital information about why coverage has failed. Because coverage is undecidable in the case of LF, the algorithm sometimes generates spurious counterexamples, that can sometimes be removed with a highly specialized albeit incomplete algorithm called finitary splitting and has proven tremendously useful in practice.

All algorithms and techniques described in this paper are implemented in the Twelf system, Version 1.4 (December 2002). Many examples of coverage are available in the example directories of the Twelf distribution. The current implementation is somewhat more general than what we describe here since it also accounts for regular worlds [21]. We plan to extend the rigorous treatment given here to this larger class of coverage problems in a future paper.

References

1. T. Coquand. An algorithm for testing conversion in type theory. In G. Huet and G. Plotkin, editors, *Logical Frameworks*, pages 255–279. Cambridge University Press, 1991.
2. T. Coquand. Pattern matching with dependent types. In *Proceedings of the Workshop on Types for Proofs and Programs*, pages 71–83, Båstad, Sweden, 1992.
3. K. Crary. Toward a foundational typed assembly language. In G. Morrisett, editor, *Proceedings of the 30th Annual Symposium on Principles of Programming Languages*, pages 198–212, New Orleans, Louisiana, Jan. 2003. ACM Press.
4. G. Dowek, T. Hardin, C. Kirchner, and F. Pfenning. Unification via explicit substitutions: The case of higher-order patterns. In M. Maher, editor, *Proceedings of the Joint International Conference and Symposium on Logic Programming*, pages 259–273, Bonn, Germany, Sept. 1996. MIT Press.
5. C. M. Elliott. *Extensions and Applications of Higher-Order Unification*. PhD thesis, School of Computer Science, Carnegie Mellon University, May 1990. Available as Technical Report CMU-CS-90-134.

6. L.-H. Eriksson. *Finitary Partial Inductive Definitions and General Logic*. PhD thesis, Department of Computer and System Sciences, Royal Institute of Technology, Stockholm, 1993.
7. L. Hallnäs. Partial inductive definitions. *Theoretical Computer Science*, 87(1):115–142, Sept. 1991.
8. R. Harper, F. Honsell, and G. Plotkin. A framework for defining logics. *Journal of the Association for Computing Machinery*, 40(1):143–184, Jan. 1993.
9. M. Hofmann and T. Streicher. The groupoid model refutes uniqueness of identity proofs. In *Proceedings of the 9th Annual Symposium on Logic in Computer Science (LICS'94)*, pages 208–212, Paris, France, 1994. IEEE Computer Society Press.
10. G. Huet. A unification algorithm for typed λ -calculus. *Theoretical Computer Science*, 1:27–57, 1975.
11. INRIA. *The Coq Proof Assistant*, version 7.4 edition, Feb. 2003. Reference Manual.
12. C. McBride. *Dependently Typed Functional Programs and their Proofs*. PhD thesis, University of Edinburgh, 1999. Available as Technical Report ECS-LFCS-00-419.
13. D. Miller. A logic programming language with lambda-abstraction, function variables, and simple unification. *Journal of Logic and Computation*, 1(4):497–536, 1991.
14. F. Pfenning. Logical frameworks. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, chapter 17, pages 1063–1147. Elsevier Science and MIT Press, 2001.
15. F. Pfenning and C. Schürmann. Algorithms for equality and unification in the presence of notational definitions. In T. Altenkirch, W. Naraschewski, and B. Reus, editors, *Types for Proofs and Programs*, pages 179–193, Kloster Irsee, Germany, Mar. 1998. Springer-Verlag LNCS 1657.
16. F. Pfenning and C. Schürmann. System description: Twelf — a meta-logical framework for deductive systems. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, pages 202–206, Trento, Italy, July 1999. Springer-Verlag LNAI 1632.
17. B. Pientka. Termination and reduction checking for higher-order logic programs. In *First International Joint Conference on Automated Reasoning (IJCAR)*, pages 401–415, Siena, Italy, 2001. Springer Verlag, LNCS 2083.
18. B. Pientka and F. Pfenning. Optimizing higher-order pattern unification. In F. Baader, editor, *Proceedings of the 19th Conference on Automated Deduction (CADE-19)*, Miami Beach, Florida, July 2003. Springer-Verlag LNCS. To appear.
19. E. Rohwedder and F. Pfenning. Mode and termination checking for higher-order logic programs. In H. R. Nielson, editor, *Proceedings of the European Symposium on Programming*, pages 296–310, Linköping, Sweden, Apr. 1996. Springer-Verlag LNCS 1058.
20. P. Schroeder-Heister. Rules of definitional reflection. In M. Vardi, editor, *Proceedings of the Eighth Annual IEEE Symposium on Logic in Computer Science*, pages 222–232, Montreal, Canada, June 1993.
21. C. Schürmann. *Automating the Meta Theory of Deductive Systems*. PhD thesis, Department of Computer Science, Carnegie Mellon University, Aug. 2000. Available as Technical Report CMU-CS-00-146.
22. C. Schürmann, R. Fontana, and Y. Liao. Delphin: Functional programming with deductive systems. Draft.
23. R. Virga. *Higher-Order Rewriting with Dependent Types*. PhD thesis, Department of Mathematical Sciences, Carnegie Mellon University, Sept. 1999. Available as Technical Report CMU-CS-99-167.