

BEYOND NOTATIONS: HYGIENIC MACRO EXPANSION FOR THEOREM PROVING LANGUAGES *

SEBASTIAN ULLRICH AND LEONARDO DE MOURA

^a Karlsruhe Institute of Technology, Germany
e-mail address: sebastian.ullrich@kit.edu

^b Microsoft Research, USA
e-mail address: leonardo@microsoft.com

ABSTRACT. In interactive theorem provers (ITPs), extensible syntax is not only crucial to lower the cognitive burden of manipulating complex mathematical objects, but plays a critical role in developing reusable abstractions in libraries. Most ITPs support such extensions in the form of restrictive “syntax sugar” substitutions and other ad hoc mechanisms, which are too rudimentary to support many desirable abstractions. As a result, libraries are littered with unnecessary redundancy. Tactic languages in these systems are plagued by a seemingly unrelated issue: accidental name capture, which often produces unexpected and counterintuitive behavior. We take ideas from the Scheme family of programming languages and solve these two problems simultaneously by proposing a novel *hygienic macro system* custom-built for ITPs. We further describe how our approach can be extended to cover type-directed macro expansion resulting in a single, uniform system offering multiple abstraction levels that range from supporting simplest syntax sugars to elaboration of formerly baked-in syntax. We have implemented our new macro system and integrated it into the new version of the Lean theorem prover, Lean 4. Despite its expressivity, the macro system is simple enough that it can easily be integrated into other systems.

1. INTRODUCTION

Mixfix notation systems have become an established part of many modern ITPs for attaching terse and familiar syntax to functions and predicates of arbitrary arity.

<code>_⊢_:_ = Typing</code>	<i>Agda</i>
<code>Notation "Ctx ⊢ E : T" := (Typing Ctx E T).</code>	<i>Coq</i>
<code>notation typing ("_ ⊢ _ : _")</code>	<i>Isabelle</i>
<code>notation Γ ⊢ e : τ := Typing Γ e τ</code>	<i>Lean 3</i>

As a further extension, all shown systems also allow *binding names* inside mixfix notations.

Key words and phrases: lean macro hygiene.

* This is an extended version of [UdM20], with the most significant changes being the addition of Section 4.1 and Section 7.

<code>syntax</code> $\exists A (\lambda x \rightarrow P) = \exists [x \in A] P$	<i>Agda</i>
<code>Notation</code> " $\exists x, P$ " := (exists (fun x => P)).	<i>Coq</i>
<code>notation</code> exists (binder " \exists ")	<i>Isabelle</i>
<code>notation</code> `∃` binder ``,` r:(scoped P, Exists P) := r	<i>Lean 3</i>

While these extensions differ in the exact syntax used, what is true about all of them is that at the time of the notation declaration, the system already, statically knows what parts of the term are bound by the newly introduced variable. This is in stark contrast to *macro systems* in Lisp and related languages where the expansion of a *macro* (a syntactic substitution) can be specified not only by a *template expression* with placeholders like above, but also by arbitrary *syntax transformers*, i.e. code evaluated at compile time that takes and returns a syntax tree.¹ As we move to more and more expressive notations and ideally remove the boundary between built-in and user-defined syntax, we argue that we should no more be limited by the static nature of existing notation systems and should instead introduce syntax transformers to the world of ITPs.

However, as usual, with greater power comes greater responsibility. By using arbitrary syntax transformers, we lose the ability to statically determine what parts of the macro template can be bound by the macro input (and vice versa). Thus it is no longer straightforward to avoid *hygiene* issues (i.e. accidental capturing of identifiers; [KFFD86]) by automatically renaming identifiers. We propose to learn from and adapt the macro hygiene systems implemented in the Scheme family of languages for interactive theorem provers in order to obtain more general but still well-behaved notation systems.

After giving a practical overview of the new, macro-based notation system we implemented in Lean 4 [dMU21] in Section 2, we describe the issue of hygiene and our general hygiene algorithm, which should be just as applicable to other ITPs, in Section 3. Section 4 gives a detailed description of the implementation of this algorithm, and macros in general, in Lean 4. In Section 5, we extend the use case of macros from mere syntax substitutions to type-aware elaboration.² Finally, we have already encountered hygiene issues in the current version of Lean in a different part of the system: the tactic framework. We discuss how these issues are inevitable when implementing reusable tactic scripts and how our macro system can be applied to this hygiene problem as well in Section 6.

Contributions. We present a system for hygienic macros optimized for theorem proving languages as implemented³ in the new version of the Lean theorem prover, Lean 4.

- We describe a novel, efficient hygiene algorithm to employ macros in ITP languages at large: a combination of a white-box, effect-based approach for detecting newly introduced identifiers and an efficient encoding of scope metadata.
- We show how such a macro system can be seamlessly integrated into existing elaboration designs to support type-directed expansion even if they are not based on homogeneous source-to-source transformations.
- We show how hygiene issues also manifest in tactic languages and how they can be solved with the same macro system. To the best of our knowledge, the tactic language in Lean 4 is the first tactic language in an established theorem prover that is automatically hygienic in this regard.

¹These two macro declaration styles are commonly referred to as *pattern-based* vs. *procedural*.

²By “elaboration”, we mean the transformation of surface-level syntax into the explicit kernel term/declaration representation, including type and typeclass inference.

³<https://github.com/leanprover/lean4/blob/IJCAR20-LMCS/src/Lean/Elab>

2. THE NEW MACRO SYSTEM

Lean's previous notation system as shown in Section 1 is still supported in Lean 4, but based on a much more general macro system; in fact, the `notation` command⁴ itself has been reimplemented as a macro, more specifically as a *macro-generating macro* making use of our tower of abstraction levels. The corresponding Lean 4 command⁵ for the example from the previous section

```
notation Γ "⊢" e ":" τ => Typing Γ e τ
```

expands to the macro declaration

```
macro Γ:term "⊢" e:term ":" τ:term : term => `(Typing $Γ $e $τ)
```

where the *syntactic category* (`term`) of placeholders and of the entire macro is now specified explicitly. The right-hand side uses an explicit *syntax quasiquotation* to construct the syntax tree, with syntax placeholders (*antiquotations*) prefixed with `$`. As suggested by the explicit use of a quotation, the right-hand side may now be an arbitrary Lean term computing a syntax object; in other words, there is no distinction between pattern-based and procedural macros in our system. We can now use this abstraction level to implement simple macros in syntactic categories other than `term`, such as for commands.

```
macro "defthunk" id:ident ":@" e:term : command =>
  `(def $id := Thunk.mk (fun _ => $e))
defthunk big := mkArray 100000 true
```

`macro` itself is another command-level macro that, for our `notation` example, expands to two commands

```
syntax term "⊢" term ":" term : term
macro_rules
  | `($Γ ⊢ $e : $τ) => `(Typing $Γ $e $τ)
```

that is, a pair of parser extension (which we will not further discuss in this paper) and syntax transformer. Our reason for ultimately separating these two concerns is that we can now obtain a well-structured syntax tree pre-expansion, i.e. a *concrete* syntax tree, and use it to implement source code tooling such as auto-completion, go-to-definition, and refactorings. Implementing even just the most basic of these tools for the Lean 3 frontend that combined parsing and notation expansion meant that they had to be implemented right inside the parser, which was not an extensible or even maintainable approach in our experience.

Both `syntax` and `macro_rules` are in fact further macros for regular Lean definitions encoding procedural metaprograms, though users should rarely need to make use of this lowest abstraction level explicitly. Both commands can only be used at the top level; we are not currently planning support for local macros.

There is no more need for the complicated `scoped` syntax since the desired translation can now be specified naturally, without any need for further annotations.

```
notation "∃" b "," P => Exists (fun b => P)
```

⁴A Lean file is made up of a sequence of commands, which are processed in turn and can extend the environment with new declarations or metadata (such as custom parsers).

⁵All examples including full context can be found in the supplemental material at <https://github.com/Kha/macro-supplement>.

The lack of static restrictions on the right-hand side ensures that this works just as well with custom binding notations, even ones whose translation cannot statically be determined before substitution.

```

syntax "{ term "|" term }" : term
macro_rules
  | `({$x ∈ $s | $p}) => `(setOf (fun $x => $x ∈ $s ∧ $p))
  | `({$b      | $p}) => `(setOf (fun $b => $p))

notation "∪" b "," p => Union {b | p}

```

Here we explicitly make use of the `macro_rules` abstraction level for its convenient syntactic pattern matching syntax. `macro_rules` are “open” in the sense that multiple transformers for the same `syntax` declaration can be defined; they are tried up to the first match, starting with the newest declaration (though this can be customized using explicit priority annotations). Thus the following extension will not be shadowed by the `$b` default case above:

```

macro_rules
  | `({$x ≤ $e | $p}) => `(setOf (fun $x => $x ≤ $e ∧ $p))

```

As a final example, we present a partial reimplementaion of the arithmetic “bigop” notations found⁶ in Coq’s Mathematical Components library [MT] such as

```
\sum_ (i <- [0, 2, 4] | i != 2) i
```

for summing over a filtered sequence of elements. The specific bigop notations are defined in terms of a single `\big_` fold operator; however, because Coq’s notation system is unable to abstract over the indexing syntax, every specific bigop notation has to redundantly repeat every specific index notation before delegating to `\big_`. In total, the 12 index notations for `\big_` are duplicated for 3 different bigops in the file.

```

Notation "\sum_ ( i <- r ) F"      := (\big[addn/0]_(i <- r) F).
Notation "\sum_ ( i <- r | P ) F" := (\big[addn/0]_(i <- r | P) F).
...
Notation "\prod_ ( i <- r ) F"      := (\big[muln/1]_(i <- r) F).
Notation "\prod_ ( i <- r | P ) F" := (\big[muln/1]_(i <- r | P) F).
...

```

In contrast, using our system, we can introduce a new syntactic category for index notations, interpret it once in `\big_`, and define new bigops on top of it without any redundancy.

```

declare_syntax_cat index
syntax ident "<-" term : index
syntax ident "<-" term "|" term : index
...

macro "$Σ" "(" idx:index ")" F:term : term =>
  `(\big_ [Add.add, 0] ($idx) $F)

macro "$Π" "(" idx:index ")" F:term : term =>
  `(\big_ [Mul.mul, 1] ($idx) $F)

```

⁶<https://github.com/math-comp/math-comp/blob/master/mathcomp/ssreflect/bigop.v>

The full example is included in the supplement. We reiterate that this is not merely a showcase of a parsing extension, but that abstracting over binding syntax in this manner is fundamentally incompatible with any static approach of ensuring hygiene. The dynamic nature of Scheme-like macros allows us to always apply such factorings without being burdened by static restrictions while still preserving hygiene.

3. HYGIENE ALGORITHM

In this section, we will give a mostly self-contained description of our algorithm for automatic hygiene applied to a simple recursive macro expander; we postpone comparisons to existing hygiene algorithms to Section 8.

Hygiene issues occur when transformations such as macro expansions lead to an unexpected capture (rebinding) of identifiers. For example, we would expect the notation

```
notation "const" e => fun x => e
```

to always produce a constant function regardless of the specific e . We would not expect the term `const x` to be the identity function because intuitively there is no x in scope at the argument position of `const`; that the implementation of the macro makes use of the name internally should be of no concern to the macro user.

Thus hygiene issues can also be described as a *confusion of scopes* when syntax parts are removed from their original context and inserted into new contexts, which makes name resolution strictly after macro expansion (such as in a compiler preceded by a preprocessor) futile. Instead we need to track scopes *as metadata* before and during macro expansion so as not to lose information about the original context of identifiers. Specifically,

- (1) when an identifier captured in a syntax quotation matches one or more⁷ top-level symbols, the identifier is annotated with a list of these symbols as *top-level* scopes to preserve its *extra-macro* context (which, because of the lack of local macros, can only contain top-level bindings), and
- (2) when a macro is expanded, all identifiers freshly introduced by the expansion are annotated with a new *macro* scope to preserve the *intra-macro* context. In particular, different expansions of the same macro introduce different annotations. Macro scopes are appended to a list, i.e. ordered by expansion time. This full “history of expansions” is necessary to treat macro-producing macros correctly, as we shall see in Section 3.2.

Thus, the expansion of the above term `const x` should be (an equivalent of) `fun x.1 => x` where `1` is a fresh macro scope appended to the macro-introduced x , preventing it from capturing the x from the original input. In general, we will present hygienic identifiers in the following as `n.msc1.msc2...mscn{tsc1,...,tscn}` where `n` is the original name, `msc` are macro scopes, and `tsc` top-level scopes, eliding the braces if there are no top-level scopes as in the example above. We use the dot notation to suggest both the ordered nature of macro scopes and their eventual implementation in Section 4. We will now describe how to implement these operations in a standard macro expander.

⁷Lean allows overloaded top-level bindings whereas local bindings are shadowing.

3.1. Expansion Algorithm. A Scheme-style macro expander takes a syntax tree as input and produces a fully expanded tree, that is, where all macro uses have been reduced to *core forms* that cannot be described as macros and are instead handled by the later stages, such as an elaborator. The expander should furthermore rename bindings where necessary to avoid hygiene issues such that later stages do not have to know anything about the implementation of hygiene, or indeed that it was applied at all.

Given a *global context* (a set of symbols), our expansion algorithm does so by a conventional top-down expansion, keeping track of an initially-empty *local context* (another set of symbols). When a binding core form is encountered, the local context is extended with the bound symbol(s); existing top-level scopes on bindings are discarded at this step since they are only meaningful on references. Thus we will formally define a symbol as an identifier together with a list of macro scopes, such as `x.1` above. As we shall see in Section 4, this definition of symbol is covered by the existing one in Lean, so later stages indeed do not have to concern themselves with it.

When a reference (another core form), i.e. an identifier not in binding position, is encountered, it is resolved according to the following rules:

- (1) If the local context has an entry for the same symbol, the reference binds to the corresponding local binding; any top-level scopes are again discarded.
- (2) Otherwise, if the identifier is annotated with one or more top-level scopes or matches one or more symbols in the global context, it binds to all of these (to be disambiguated by the elaborator).
- (3) Otherwise, the identifier is unbound and an error is generated.

In the common incremental compilation mode of ITPs, every command is fully processed before subsequent commands. Thus, an expander for such a system will never extend the global context by itself, but pass the fully expanded command to the next compilation stage before being called again with the next command's unexpanded syntax tree and a possibly extended global context.

Notably, our expander does not introduce macro scopes by itself, either, much in contrast to other expansion algorithms. We instead delegate this task to the macro's implementation, though in a completely transparent way for all pattern-based and for many procedural macros. We claim that a macro should in fact be interpreted as an *effectful* computation since two expansions of the same identifier-introducing macro should not return the same syntax tree to avoid unhygienic interactions between them. Thus, as a *side effect*, it should apply a fresh macro scope to each newly introduced identifier. In particular, a syntax quotation should not merely be seen as a datum, but as an effectful value that obtains and applies this fresh scope to all the identifiers captured by it to immediately ensure hygiene for pattern-based macros. Procedural macros producing identifiers not originating from syntax quotations might need to obtain and make use of the fresh macro scope explicitly. Note that forgoing to do so is not sufficient to reliably implement *anaphoric* or other hygiene-bending macros that make an identifier (conventionally `it` in Lisp languages) available in the scope of the macro caller, as discussed in [BCF11]. Instead, we believe that the correct translation of anaphoric macros to Lean is to change such identifiers to *keywords* that do not participate in hygiene at all, analogous to the *syntax parameters* of [BCF11]. An example for this is the `this` keyword introduced by tactics such as `have` that can be used to refer to the just-proved fact.⁸

⁸<https://github.com/leanprover/lean4/blob/6d0c91c/src/Init/Notation.lean#L204-L207>

We give a specific monad-based [Mog91] implementation of effectful syntax quotations as a regular macro in Section 4. For the remainder of this section, we will simply assume that macros have been implemented in this way and observe their interaction with the expansion algorithm described above.

3.2. **Examples.** Given the following input,

```
def x := 1
def e := fun y => x
notation "const" e => fun x => e
def y := const x
```

a Lean-like system using the presented expansion algorithm should incrementally parse, expand, and elaborate each declaration before advancing to the next one. For a first, trivial example, let us focus on the expansion of the second line. At this point, the global context contains the symbol `x` (plus any default imports that we will ignore here). Descending into the right-hand side of the definition, the expander first adds `y` to the local context. The reference `x` does not match any local definitions, so it binds to the matching top-level definition.

In the next line, the built-in `notation` macro expands to the definitions

```
syntax "const" term : term
macro_rules
| `(const $e) => `(fun x => $e)
```

When a top-level macro application unfolds to multiple declarations, we expand and elaborate these incrementally as well to ensure that declarations are in the global context of subsequent declarations from the same expansion. When recursively expanding the `macro_rules` declaration (we will assume for this example that `macro_rules` itself is a core form) in the global context $\{x, e\}$, we first visit the syntax quotation on the left-hand side. The identifier `e` inside of it is in an antiquotation and thus not captured by the quotation. It is in binding position for the right-hand side, so we add `e` to the local context. Visiting the right-hand side, we find the quotation-captured identifier `x` and annotate it with the matching top-level definition of the same name; we do not yet know that it is in a binding position. When visiting the reference `e`, we see that it matches a local binding and do not add top-level scopes.

```
macro_rules
| `(const $e) => `(fun x{x} => $e)
```

Visiting the last line

```
def y := const x
```

with the global context $\{x, e\}$, we descend into the right-hand side. We expand the `const` macro given a fresh macro scope `1`, which is applied to any captured identifiers.

```
def y := fun x.1{x} => x
```

We add the symbol `x.1` (discarding the top-level scope `x`) to the local context and finally visit the reference `x`. The reference does not match the local binding `x.1` but does match the top-level binding `x`, so it binds to the latter.

```
def y := fun x.1 => x
```


Now let us briefly look at a more complex macro-macro example demonstrating use of the macro scopes stack:

```
macro "m" n:ident : command => `(
  def f := 1
  macro "mm" : command => `(
    def $n := f + 1
    def f := $n + 1))
```

If we use this macro as in `m f`, we apply a fresh macro scope 1 to all captured identifiers, then incrementally process the two new declarations.

```
def f.1 := 1
macro "mm" : command => `(
  def f := f.1{f.1} + 1
  def f.1{f.1} := f + 1)
```

If we use the new macro `mm`, we apply one more macro scope 2.

```
def f.2 := f.1.2{f.1} + 1
def f.1.2{f.1} := f.2 + 1
```

When processing these new definitions, we see that the scopes ensure the expected name resolution.

```
def f.1 := 1
...
def f.2 := f.1 + 1
def f.1.2 := f.2 + 1
```

In particular, we now have global declarations `f.1`, `f.2`, and `f.1.2` that show that storing only a single macro scope would have led to a collision.

4. IMPLEMENTATION

Syntax objects in Lean 4 are represented as an inductive type of *nodes* (or nonterminals), *atoms* (or terminals), and, as a special case of terminals, *identifiers*.

```
inductive Syntax where
| node (kind : Name) (args : Array Syntax)
| atom (info : SourceInfo) (val : String)
| ident (info : SourceInfo) (rawVal : String) (val : Name) (preresolved :
  List (Nat × List String))
| missing
```

An additional constructor represents *missing* parts from syntax error recovery. Atoms and identifiers are annotated with source location metadata unless generated by a macro. Identifiers carry macro scopes inline in their `Name` while top-level scopes are held in a separate list. The additional `Nat` is an implementation detail of Lean's hierarchical name resolution.

The type `Name` of hierarchical names precedes the implementation of the macro system and is used throughout Lean's implementation for referring to (namespaced) symbols.


```

1  partial def expand : Syntax → ExpanderM Syntax
2    | `($id:ident) => do
3      let val : Name := getIdentVal id
4      let gctx ← getGlobalContext
5      let lctx ← getLocalContext
6      if lctx.contains val then
7        pure (mkIdent val)
8      else match resolve gctx val ++ getPreresolved id with
9        | []          => throw ("unknown identifier " ++ toString val)
10       | [(id, _)] => pure (mkIdent id)
11       | ids        => pure (mkOverloadedIds ids)
12  | `(fun ($id : $ty) => $e) => do
13    let val := getIdentVal id
14    let ty ← expand ty
15    let e ← withLocal val (expand e)
16    `(fun ($(mkIdent val) : $ty) => $e)
17  | ... -- other core forms
18  | _ => do
19    let t ← getTransformerFor stx.getKind
20    let stx ← withFreshMacroScope (t stx)
21    expand stx

```

Figure 1: Abbreviated implementation of a recursive expander for our macro system

```

inductive Name where
  | anonymous
  | str (base : Name) (s : String)
  | num (base : Name) (n : Nat)

```

The syntax ``a.b` is a literal of type `Name` for use in meta-programs. The numeric part of `Name` is not accessible from the surface syntax and reserved for internal names; similar designs are found in other ITPs. By reusing `Name` for storing macro scopes, but not top-level scopes, we ensure that the new definition of *symbol* from Section 3.1 coincides with the existing Lean type and no changes to the implementation of the local or global context are necessary for adopting the macro system.

A Lean 4 implementation of the expansion algorithm described in the previous section is given in Fig. 1; the full implementation including examples is included in the supplement. As a generalization, syntax transformers in the full implementation have the type `Syntax →`

`TransformerM Syntax` where the `TransformerM` monad gives access to the global context and a fresh macro scope per macro expansion. The expander itself uses an extended `ExpanderM` monad based on `TransformerM` that also stores the local context and the set of registered macros. We use the Lean equivalent of Haskell’s `do` notation [M⁺10] to program in these monads.

As described in Section 3.1, the expander in Fig. 1 has built-in knowledge of some “core forms” (lines 2-16) with special expansion behavior, while all other forms are assumed to be macros and expanded recursively (lines 19-21). Identifiers form one base case of the recursion. As described in the previous section, they are first looked up in the local context (recall that

the `Name` of an identifier includes macro scopes), then as a fall back in the global context plus its own top-level scopes. `mkIdent : Name → Syntax` creates an identifier without source information or top-level scopes, which are not needed after expansion. `mkOverloadedIds` implements the Lean special case of overloaded symbols to be disambiguated by elaboration; systems without overloading support should throw an ambiguity error instead in this case.

As an example of a core binding form, the expansion of a single-parameter `fun` is shown in lines 12-16 of Fig. 1. It recursively expands the given parameter type, then expands the body in a new local context extended with the value of `id`. Here `getIdentVal : Syntax → Name` in particular implements the discarding of top-level scopes from binders.

Finally, in the macro case, we fetch the syntax transformer for the given node kind, run it in a new context with a fresh current macro scope, and recurse.

```

1 partial def quoteSyntax : Syntax → TransformerM Syntax
2   | Syntax.ident info rawVal val preresolved => do
3     let gctx ← getGlobalContext
4     let preresolved := resolve gctx val ++ preresolved
5     `(Syntax.ident SourceInfo.none $(quote rawVal)
6       (addMacroScope $(quote val) msc) $(quote preresolved))
7   | stx@(Syntax.node k args) =>
8     if isAntiquot stx then pure (getAntiquotTerm stx)
9     else do
10      let args ← args.mapM quoteSyntax
11      `(Syntax.node $(quote k) $(quote args))
12   | Syntax.atom info val => `(Syntax.atom SourceInfo.none $(quote val))
13   | Syntax.missing => pure Syntax.missing
14
15 def expandStxQuot (stx : Syntax) : TransformerM Syntax := do
16   let stx ← quoteSyntax (stx.getArg 1)
17   `(do msc ← getCurrMacroScope; pure $stx)

```

Figure 2: Simplified syntax transformer for syntax quotations

Syntax quotations are given as one example of a macro: they do not have built-in semantics but transform into code that constructs the appropriate syntax tree (`expandStxQuot` in Fig. 2). More specifically, a syntax quotation will, at run time (of the surrounding macro), query the current macro scope `msc` from the surrounding `TransformerM` monad (code generated by `expandStxQuot`) and apply it to all captured identifiers (code generated by `quoteSyntax`). `quoteSyntax` recurses through the quoted syntax tree, reflecting its constructors. Basic datatypes such as `String` and `Name` are turned into `Syntax` via the typeclass method `quote`. For antiquotations, we return their contents unreflected. In the case of identifiers, we resolve possible global references at compile time and reflect them, while `msc` is applied at run time. Thus a quotation ``(a + $b)` inside a global context where the symbol `a` matches declarations

`a.a` and `b.a` is transformed to the equivalent of

```
do let msc ← getCurrMacroScope
  pure (Syntax.node `plus [
    Syntax.ident SourceInfo.none "a" (addMacroScope `a msc)
    [ `a.a, `b.a ],
    Syntax.atom SourceInfo.none "+",
    b])
```

This implementation of syntax quotations itself makes use of syntax quotations for simplicity and thus is dependent on its own implementation in the previous stage of the compiler. Indeed, the helper variable `msc` must be renamed should the name already be in scope and used inside an antiquotation.⁹ Note that `quoteSyntax` is allowed to reference the same `msc` as `expandStxQuot` because they are part of the same macro call and the current macro scope is unchanged between them. While alternative approaches that use fresh macro scopes on function calls *within* a macro are thinkable, we prefer the presented behavior, which matches that of the Scheme family, because it preserves referential transparency: if `quoteSyntax` is inlined into `expandStxQuot`, the behavior is unchanged.

4.1. Extended Quasiquotations. Automatic hygiene can greatly simplify development of macros, but a convenient way for con- and destructing syntax is at least as important. Before we get to more complex macro examples below, we want to describe some syntactic extensions to quotations and antiquotations we have implemented that will come in useful.

A first obvious such extension is to allow quotations including antiquotations as patterns such as after `match` or `fun`.

```
fun
  | `()          => ...
  | `($e)       => ...
  | `($e, $f)   => ...
```

Because every use of patterns eventually unfolds to a `match` in Lean, this is in fact implemented as a macro that expands `match` terms with quotation patterns into ones without such patterns. Note also that because no new syntactic identifiers are generated while matching against a quotation, there is no issue of hygiene in this case.

For syntax with repeated parts, quotation *splices* enable us to match or introduce these parts as a whole. For example, a recursive macro for n -tuple syntax can be written as

```
macro_rules
  | `()          => `(Unit.unit)
  | `($e)       => e
  | `($e, $es,*) => `(Prod.mk $e ($es,*))
```

Here `$es,*` matches/introduces the remaining elements of the tuple, including its separators. Analogous splicing syntax exists for other separators, as well as `$x*` for separator-less iteration. At most one splice can be used per sequence, but it can be pre- and suffixed with an arbitrary (but fixed) number of other elements.

In `$x*`, `x` has type `Array Syntax`, the same type as the second argument of the `node` constructor. For `$x,*` and similar we instead use the dependent wrapper type `SepArray ", "`

⁹As long as no such case exists, a hygienic implementation of syntax quotations can be bootstrapped from an unhygienic one, which is what we did in the case of Lean.

that provides convenience access functions for the sequence both with and without separator elements. Finally, we also provide implicit coercions between these types that automatically insert/remove/replace the separators accordingly.

While exposing splices as typed values in this way ensures that we can comfortably process or synthesize them procedurally as well, it is often more convenient to inspect splice contents immediately as part of the quotation. For this we support extended splices `[$[..]*` etc. where the splice content is parsed like an element of the sequence and can contain nested antiquotations. If used as a pattern, the match succeeds if and only if the nested pattern matches every element, in which case the contained antiquotations are each bound to an `Array` of all corresponding element-wise matches.

```
match stx with
| `(match $discr with $[| $patss,* => $branches]*) =>
  -- discr : Syntax
  -- patss : Array (SepArray ",")
  -- branches : Array Syntax
  ...
```

By default, quotations are parsed as either terms or top-level commands, since these syntactic categories are both commonly used and should usually be disjoint. Other syntactic categories, e.g. the category of universe levels that heavily overlaps with `term`, can be specified explicitly at the beginning of a quotation. Similarly, antiquotations can be suffixed with a colon followed by a category or named parser where otherwise ambiguous.

```
match levelStx with
| `(level| $id:ident) => ... -- a universe variable
| `(level| _)         => ... -- a universe placeholder
| `(level| $!)       => ... -- any (other) universe term
```

For a full example of using these and other features, we can look at a macro rule unfolding syntax such as

```
fun
| some a, some b => some (a + b)
| _,           _   => none
```

into

```
fun x.1 x.2 =>
  match x.1, x.2 with
  | some a, some b => some (a + b)
  | _,           _   => none
```

The macro rule (Fig. 3) derives the number of *discriminants* (`x.1`, `x.2` in the example) from the number of patterns of the first alternative; if other alternatives have differing number of patterns, it will lead to an elaboration error in `match` later on. The macro then introduces a lambda abstraction over a sequence of fresh variable names of this number and subsequently matches on them using the given patterns.

It does so by matching on the first alternative of the match in detail, then capturing the remaining ones in `alts : Array Syntax`. The left-hand side `ps1` of the first alternative is a `SepArray` `",",` so we use `getElems` to access its elements and generate a fresh variable

```

1 macro_rules
2   | `(fun
3     | $ps1,* => $rhs1
4     $alts:matchAlt*) => do
5     let discrs ← ps1.getElems.mapM (fun _ => withFreshMacroScope `(x))
6     `(fun $discrs* =>
7       match [$discrs:ident],* with
8       | $ps1,* => $rhs1
9       $alts:matchAlt*)

```

Figure 3: A macro rule for expanding a combined fun-match syntax.

for each of them, which we do by running the single quotation ``(x)` repeatedly under `withFreshMacroScope`, annotating the variable with a unique macro scope each time¹⁰.

With `discrs : Array Syntax` generated, we insert it into the final quotation, once as a straight sequence after `fun` and once separated by commas after `match`, followed by the alternatives copied from the input without changes. Note that `$discrs:ident*` would not have worked in this case because identifiers are merely a special case of the more general `match` discriminant syntax that allows prefixing a discriminant with `h:`, where the identifier `h` will then hold the proof that the discriminant matched the corresponding pattern. Thus there is no direct identifier sequence to insert and we have to instead say that we insert a sequence of general discriminants, each one built up of an identifier without a proof variable prefix, which internally will wrap each element in an additional syntax tree node of the `matchDiscr` kind. This necessary disambiguation of overlapping syntax sadly is a price we have to pay for our preference of such syntax over more regular but verbose one such as S-expressions.

Finally, for the sake of completeness we will mention the *token antiquotation* ``${x}` that any token can be suffixed with to extract/set its `SourceInfo` metadata. This kind of antiquotation is mostly useful for displaying errors on specific tokens and preserving metadata in transformations.

```

| `(tactic| case $tag =>`${arrTk $tac) => do
...
  reportUnsolvedGoalsAt arrTk
...

case cons => skip
  --^ unsolved goals displayed here

```

5. INTEGRATING MACROS INTO ELABORATION

The macro system as described so far can handle most syntax sugars of Lean 3 except for ones requiring type information. For example, the *anonymous constructor* `<e, ...>` is sugar for `(c e ...)` if the expected type of the expression is known and it is an inductive type with a single constructor `c`. While trivial to parse, there is no way to implement this syntax as

¹⁰This is comparable to a call to the `gensym` function found in many Lisp systems.

a macro if expansion is done strictly prior to elaboration. A more complex example is the *structure instance notation* `{ field1 := e, ... }` that must analyze the definition of the given or inferred structure type in order to expand to the correct constructor call. To the best of our knowledge, none of the ITPs listed in the introduction support hygienic elaboration extensions of this kind, but we will show how to extend their common elaboration scheme in that way in this section.

Elaboration¹¹ can be thought of as a function `elabTerm : Syntax → ElabM Expr` in an appropriate monad `ElabM`¹² from a (concrete or abstract) surface-level syntax tree type `Syntax` to a fully-specified core term type `Expr` [dMAKR15]. We have presented the (concrete) definition of `Syntax` in Lean 4 in Section 4; the particular definition of `Expr` is not important here. While such an elaboration system could readily be composed with a type-insensitive macro expander such as the one presented in Section 3, we would rather like to *intertwine* the two to support type-sensitive but still hygienic-by-default macros (henceforth called *elaborators*) without having to reimplement macros of the kind discussed so far. Indeed, these can automatically be adapted to the new type given an adapter between the two monads, similarly to the adaption of macros to *expanders* in [DFH86]:

```
def transformerToElaborator (t : Syntax → TransformerM Syntax) :
  Syntax → ElabM Expr :=
fun stx => do
  let stx' ← (transformerMToElabM t) stx
  elabTerm stx'
```

Because most parts of our hygiene system are implemented by the expander for syntax quotations, the only changes to an elaboration system necessary for supporting hygiene are storing the current macro scope in the elaboration monad (to be passed to the expansion monad in the adapter) and allocating a fresh macro scope whenever a macro or elaborator is invoked. Thus elaborators immediately benefit from hygiene as well whenever they use syntax quotations to construct unelaborated helper syntax objects to pass to `elabTerm`. In order to support syntax quotations in these two and other monads, we generalize their implementation to a new monad typeclass implemented by both monads.

```
class MonadQuotation (m : Type → Type) where
  getCurrMacroScope : m MacroScope
  withFreshMacroScope : m α → m α
```

The second operation is not used by syntax quotations directly, but can be used by procedural macros and elaborators to manually enter new macro call scopes.

As an example, the following is a simplified implementation of the anonymous constructor syntax mentioned above.

```
@[termElab anonymousCtor] def elabAnonymousCtor : Syntax → ElabM Expr
| `(<$args,*>) => do
  let expectedType ← getExpectedType
  match Expr.getAppFn expectedType with
  | Expr.const constName _ _ => do
    let ctors ← getCtors constName
```

¹¹At the term level; elaboration of other syntactic categories work analogously but with different output types.

¹²Or some other encoding of effects.

```

match ctors with
| [ctor] => do
  let stx ← `($ (mkCIdent ctor) $args*)
  elabTerm stx
... -- error handling

```

The `[termElab]` attribute registers this elaborator for the given syntax node kind. The function `mkCIdent : Name → Syntax` synthesizes a hygienic reference to the given constant name by storing it as a top-level scope and applying a reserved macro scope to the constructed identifier. Note the monadic binding of the syntax quotation, and that the separators of `$args,*` are implicitly discarded when it is used as a plain sequence `$args*`.

This implementation fails if the expected type is not yet sufficiently known at this point. The actual implementation¹³ of this elaborator extends the code by *postponing* elaboration in this case. When an elaborator requests postponement, the system returns a fresh metavariable as a placeholder and associates the input syntax tree with it. Before finishing elaboration of a command, postponed elaborators associated with unsolved metavariables are retried until they all ultimately succeed, or else elaboration is stuck because of cyclic dependencies and an error is signed.

6. TACTIC HYGIENE

Lean 3 includes a tactic framework that, much like macros, allows users to write custom automation either procedurally inside a tactic monad or “by example” using tactic language quotations, or in a mix of both [EUR⁺17]. For example, Lean 3 uses a short tactic block to prove injection lemmas for data constructors.

```

def mkInjEq : TacticM Unit :=
  `[intros; apply propext; apply Iff.intro; ...]

```

Unfortunately, this code unexpectedly broke in Lean 3 when used from a library for homotopy type theory that defined its own `propext` and `Iff.intro` declarations;¹⁴ in other words, Lean 3 tactic quotations are unhygienic and required manual intervention in this case. Just like with macros, the issue with tactics is that binding structure in such embedded terms is not known at declaration time. Only at tactic run time do we know all local variables in the current context that preceding tactics may have added or removed, and therefore the scope of each captured identifier.

Arguably, the Lean 3 implementation also exhibited a lack of hygiene in the handling of tactic-introduced identifiers: it did not prevent users from referencing such an identifier outside of the scope it was declared in.

```

def myTac : TacticM Unit := `[intro h]
lemma triv (p : Prop) : p → p := begin myTac; exact h end

```

Coq’s similar Ltac tactic language [Del00] exhibits the same issue and users are advised not to introduce fixed names in tactic scripts but to generate fresh names using the `fresh` tactic first,¹⁵ which can be considered a manual hygiene solution.

¹³<https://github.com/leanprover/lean4/blob/IJCAR20-LMCS/src/Lean/Elab/BuiltinNotation.lean#L16>

¹⁴<https://github.com/leanprover/lean/pull/1913>

¹⁵<https://github.com/coq/coq/issues/9474>

Lean 4 instead extends its automatically hygienic macro implementation to tactic scripts by allowing regular macros in the place of tactic invocations.

```
macro "myTac" : tactic => `(intro h; exact h)
theorem triv (p : Prop) : p → p := by myTac
```

By the same hygiene mechanism described above, introduced identifiers such as `h` are renamed so as not to be accessible outside of their original scope, while references to global declarations are preserved as top-level scope annotations. Thus Lean 4's tactic framework resolves both hygiene issues discussed here without requiring manual intervention by the user. Expansion of tactic macros in fact does not precede but is integrated into the *tactic evaluator* `evalTactic : Syntax → TacticM Unit` such that recursive macro calls are expanded lazily, allowing for combinators like `repeat` that would otherwise lead to infinite recursion during expansion.

```
syntax "repeat" tactic : tactic
macro_rules
| `(tactic| repeat $t) => `(tactic| try ($t; repeat $t))
```

Note that `macro` cannot be used here because the parser for `repeat` would not yet be available in the right-hand side. When `$t` eventually fails, the recursion is broken without visiting and expanding the subsequent `repeat` macro call. The `try` tactical is used to ignore this eventual failure.

While we believe that macros will cover most use cases of Lean 3's tactic quotations in Lean 4, their use within larger `TacticM` metaprograms can be recovered by passing such a quotation to `evalTactic`:

```
def myTac2 : TacticM Unit := do
  let stx ← `(tactic| intro h; exact h)
  evalTactic stx
```

`TacticM` implements the `MonadQuotation` typeclass for this purpose.

7. BEST-EFFORT EAGER NAME ANALYSIS IN MACROS

The dynamic nature of binding in macros has enabled us to implement many Lean language features as macros, with hygiene guaranteeing that bindings within the macro do not interfere with ones outside of it. However, while knowing that a mistyped identifier in a macro will not be accidentally be bound by bindings at the use site is great in theory, in practice it would be even better to be told about the typo immediately while writing the macro! This is especially true when renaming a declaration, either manually where we might accidentally miss an occurrence of it inside a macro and then must track name binding errors at use sites back to the responsible macro, or automatically where refactoring tools must treat macros as black boxes. After all, a static view of a notation's binding structure is exactly what we gave up in Section 1 in exchange for the ability to arbitrarily abstract over bindings. For example, there is no general way to statically analyze whether `x` inside the following quotation is well-scoped given an arbitrary computation resulting in `stx`:

```
let stx ← ...
  `(fun $stx => x)
```

With this theoretical limitation in mind, and given that this is more of a practical issue of maintenance, perhaps a practical, best-effort solution is sufficient as long as we retain all

hygiene guarantees. We have done so with an *opt-in*, partial but extensible approach to eager name resolution in macros based on a variant of our quotation syntax.

```
``(fun x => x + $y + id z) -- error: unknown identifier 'z'
```

A *double-backtick* quotation eventually unfolds to the basic, single-backtick version and thus retains all its semantics. Before that, however, it recursively checks for identifiers that can statically be assumed to be unbound using the following heuristics:

- (1) If there is a special *precheck* hook registered for the syntax kind in question, we use it. Precheck hooks can signal binding errors as well as recursively continue the precheck on nested syntax, possibly with an extended (untyped) *quotation context*, which is initially empty. For example, we provide a precheck hook for `fun x => e` that recurses into `e` after adding `x` to the quotation context. The central *identifier* precheck hook ultimately raises an error if an identifier is reached that is neither in the global, extra-macro context, nor in the quotation context. Other examples for precheck hooks we have added are for `match` and function application.
- (2) Otherwise, if there are no identifiers in the quoted syntax, we assume that there is no risk of unbound ones, and the check succeeds. In particular, antiquotations (which contain *unquoted* identifiers only) are always skipped.
- (3) Otherwise, if the quoted syntax is a macro, we unfold it and precheck the result. Here we assume that the macro is sufficiently good-natured: while macros are pure functions by definition in Lean, their behavior and in particular binding structure could in theory change drastically enough in the presence of antiquotations that it could lead to false positives of this analysis. If that is the case, the user either has to provide a custom precheck hook for it, or fall back to the basic quotation syntax.

In contrast to macros, we cannot typically run elaborators directly during this check because they usually depend on type information that is not yet available, and are not guaranteed to be pure.

- (4) Otherwise the check fails since we do not know how to analyze the syntax at hand. Again, users would either have to provide a precheck hook, or use to the single-backtick quotation syntax.

This approach is clearly best-effort with many opportunities for unhandled cases. However, since it is guaranteed that after the check the semantics are the same as for the basic quotation syntax, including unchanged hygiene guarantees, we believe that the practical advantages of using the syntax where possible are significant. In particular, we have observed that quotations capturing global identifiers, which are most at risk of breaking during refactorings, are usually quite simple while complicated quotations that are used to translate one syntax into a more general one often do not reference any identifiers at all, and thus would not benefit from the additional check anyway.

This is especially true for notations, which are usually exceedingly simple in structure, most often consisting of nothing but an application of a global function symbol to the notation arguments, which is covered by the identifier and application precheck hooks. Thus we have modified the `notation` macro to use double-backtick quotations by default in order to make users immediately aware of any unbound identifiers inside of them, restoring its Lean 3 behavior (where the absence of macros naturally allowed for such a check).

```
notation "∃" x "," e => Exits.intro (fun x => e) -- error: unknown identifier
      'Exits.intro'
```

We provide an option to disable this change, though all notations in the standard library passed the additional check without modification. The check did on the other hand find a notation example in our documentation that contained an accidentally unbound identifier.

8. RELATED WORK

The main inspiration behind our hygiene implementation was Racket’s *Sets of Scopes* [Fla16] hygiene algorithm. Much like in our approach, Racket annotates identifiers both with scopes from their original context as well as with additional macro scopes when introduced by a macro expansion. However, there are some significant differences: Racket stores both types of scopes in a homogeneous, unordered set and does name resolution via a maximum-subset check. For both simplicity of implementation and performance, we have reduced scopes to the bare minimal representation using only strict equality checks, which we can easily encode in our existing `Name` implementation. In particular, we only apply scopes to matching identifiers and only inside syntax quotations. This optimization is of special importance because top-level declarations in Lean and other ITPs are not part of a single, mutually recursive scope as in Racket, but they each open their own scope over all subsequent declarations, which would lead to a total number of scope annotations quadratic in the number of declarations using the Sets of Scopes algorithm. Finally, Racket detects macro-introduced identifiers using a “black-box” approach without the macro’s cooperation following the marking approach of [KFFD86]: a fresh macro scope is applied to all identifiers in the macro input, then inverted on the macro output. While elegant, a naive implementation of this approach can again result in quadratic runtime compared to unhygienic expansion and requires further optimizations in the form of lazy scope propagation [DHB93], which is difficult to implement in a pure language like Lean. Our “white-box” approach based on the single primitive of an effectful syntax quotation, while slightly easier to escape from in procedural syntax transformers, is simple to implement, incurs minimal overhead, and is equivalent for pattern-based macros.

The idea of automatically handling hygiene in the macro, and not in the expander, was introduced in [CR91], though only for pattern-based macros. MetaML [TS00] refined this idea by tying hygiene more specifically to syntax quotations that could be used in larger metaprogram contexts, which Template Haskell [SJ02] interpreted as effectful (monadic) computations requiring access to a fresh-names generator, much like in our design. However, both of the latter systems should perhaps be characterized more as metaprogramming frameworks than Scheme-like macro systems: there are no “macro calls” but only explicit splices and so only built-in syntax with known binding semantics can be captured inside syntax quotations. Thus the question of which captured identifiers to rename becomes trivial again, just like in the basic notation systems discussed in Section 1.

While the vast majority of research on hygienic macro systems has focused on S-expression-based languages, there have been previous efforts on marrying that research with non-parenthetical syntax, with different solutions for combining syntax tree construction and macro expansion. The Dylan language requires macro syntax to use predefined terminators and eagerly scans for the end of a macro call using this knowledge [BPS99], while in Honu [RF12] the syntactic structure of a macro call is discovered during expansion by a process called “enforestation”. The Fortress [ACH⁺05] language strictly separates the two concerns into grammar extensions and transformer declarations, much like we do. Dylan and Fortress are restricted to pattern-based macro declarations and thus can make use of simple hygiene algorithms while Honu uses the full generality of the Racket macro expander. On the

other hand, Honu’s authors “explicitly trade expressiveness for syntactic simplicity” [RF12]. In order to express the full Lean language and desirable extensions in a macro system, we require both unrestricted syntax of macros and procedural transformers.

Many of the antiquotation extensions presented in Section 4.1 have been inspired by similar syntax in Rust’s pattern-based macro language, though we have opened their use to procedural macros as well, using appropriate type representations.

Many theorem provers such as Agda, Coq, Idris, and Isabelle not already based on a macro-powered language provide restricted syntax extension mechanisms, circumventing hygiene issues by statically determining binding as seen in Section 1. Extensions that go beyond that do not come with automatic hygiene guarantees. Agda’s macros,¹⁶ for example, operate on the De Bruijn index-based core term level and are not hygienic.¹⁷ The ACL2 prover in contrast uses a subset of Common Lisp as its input language and adapts the hygiene algorithm of [DHB93] based on renaming [EF10]. The experimental Cur [CBTB19] theorem prover is a kind of dual to our approach: it takes an established language with hygienic macros, Racket, and extends it with a dependent type system and theorem proving tools. ACL2 does not support tactic scripts, while in Cur they can be defined via regular macros. However, this approach does not currently provide tactic hygiene as defined in Section 6.¹⁸ The Eisbach [MMW16] proof method language for Isabelle is notable in that while it allows for reusable proof methods (comparable to Lean’s tactics) to be abstracted over terms, these terms are analyzed and typechecked before being passed to the proof method, and there is no interaction between names in different proof methods either when using the “structured” proof style, so the question of hygiene is moot. Using the “unstructured” proof style, the same hygiene issue as noted for other system can be triggered.

```
method my_allI =
  rule_tac allI, rename_tac escaped

lemma "∀ x. x = x"
  apply my_allI
  apply (rule_tac t = escaped in refl)
```

9. CONCLUSION

We have proposed a new macro system for interactive theorem provers that enables syntactic abstraction and reuse far beyond the usual support of mixfix notations. Our system is based on a novel hygiene algorithm designed with a focus on minimal runtime overhead as well as ease of integration into pre-existing codebases, including integration into standard elaboration designs to support type-directed macro expansion. Despite that, the algorithm is general enough to provide a complete hygiene solution for pattern-based macros and provides flexible hygiene for procedural macros. We have also demonstrated how our macro system can address unexpected name capture issues that haunt existing tactic frameworks. We have implemented our method in the new version of the Lean theorem prover, Lean 4; it should be sufficiently attractive and straightforward to implement to be adopted by other interactive theorem proving systems as well.

¹⁶<https://agda.readthedocs.io/en/v2.6.0.1/language/reflection.html#macros>

¹⁷<https://github.com/agda/agda/issues/3819>

¹⁸<https://github.com/wilbowma/cur/issues/104>

Acknowledgments. We are very grateful to the anonymous reviewers, David Thrane Christiansen, Gabriel Ebner, Matthew Flatt, Sebastian Graf, Alexis King, Daniel Selsam, and Max Wagner for extensive comments, corrections, and advice.

REFERENCES

- [ACH⁺05] Eric Allen, David Chase, Joe Hallett, Victor Luchangco, Jan-Willem Maessen, Sukyoung Ryu, Guy L Steele Jr, Sam Tobin-Hochstadt, Joao Dias, Carl Eastlund, et al. The Fortress language specification. *Sun Microsystems*, 139(140):116, 2005.
- [BCF11] Eli Barzilay, Ryan Culpepper, and Matthew Flatt. Keeping it clean with syntax parameters. *Proc. Wksp. Scheme and Functional Programming*, 2011.
- [BPS99] Jonathan Bachrach, Keith Playford, and C Street. D-expressions: Lisp power, Dylan style. *Style DeKalb IL*, 1999.
- [CBTB19] Stephen Chang, Michael Ballantyne, Milo Turner, and William J Bowman. Dependent type systems as macros. *Proceedings of the ACM on Programming Languages*, 4(POPL):1–29, 2019.
- [CR91] William Clinger and Jonathan Rees. Macros that work. In *Proceedings of the 18th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 155–162, 1991.
- [Del00] David Delahaye. A tactic language for the system Coq. In *Logic for Programming and Automated Reasoning, 7th International Conference, LPAR 2000, Proceedings*, pages 85–95, 2000.
- [DFH86] R Kent Dybvig, Daniel P Friedman, and Christopher T Haynes. Expansion-passing style: Beyond conventional macros. In *Proceedings of the 1986 ACM conference on LISP and functional programming*, pages 143–150, 1986.
- [DHB93] R Kent Dybvig, Robert Hieb, and Carl Bruggeman. Syntactic abstraction in Scheme. *Lisp and symbolic computation*, 5(4):295–326, 1993.
- [dMAKR15] Leonardo de Moura, Jeremy Avigad, Soonho Kong, and Cody Roux. Elaboration in dependent type theory, 2015. [arXiv:1505.04324](https://arxiv.org/abs/1505.04324).
- [dMU21] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *International Conference on Automated Deduction*, pages 625–635. Springer, 2021.
- [EF10] Carl Eastlund and Matthias Felleisen. Hygienic macros for ACL2. In *International Symposium on Trends in Functional Programming*, pages 84–101. Springer, 2010.
- [EUR⁺17] Gabriel Ebner, Sebastian Ullrich, Jared Roesch, Jeremy Avigad, and Leonardo de Moura. A metaprogramming framework for formal verification. *Proc. ACM Program. Lang.*, 1(ICFP), September 2017. doi:[http://dx.doi.org/10.1145/3110278](https://dx.doi.org/10.1145/3110278).
- [Fla16] Matthew Flatt. Binding as sets of scopes. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '16*, pages 705–717, New York, NY, USA, 2016. ACM. URL: [http://doi.acm.org/10.1145/2837614.2837620](https://doi.acm.org/10.1145/2837614.2837620), doi:10.1145/2837614.2837620.
- [KFFD86] Eugene Kohlbecker, Daniel P Friedman, Matthias Felleisen, and Bruce Duba. Hygienic macro expansion. In *Proceedings of the 1986 ACM conference on LISP and functional programming*, pages 151–161, 1986.
- [M⁺10] Simon Marlow et al. Haskell 2010 language report. 2010. URL: <https://www.haskell.org/onlinereport/haskell2010>.
- [MMW16] Daniel Matichuk, Toby Murray, and Makarius Wenzel. Eisbach: A proof method language for Isabelle. *Journal of Automated Reasoning*, 56(3):261–282, 2016.
- [Mog91] Eugenio Moggi. Notions of computation and monads. *Information and computation*, 93(1):55–92, 1991.
- [MT] Assia Mahboubi and Enrico Tassi. Mathematical components. URL: <https://math-comp.github.io/mcb/>.
- [RF12] Jon Ralfkind and Matthew Flatt. Honu: syntactic extension for algebraic notation through enforestation. In *ACM SIGPLAN Notices*, volume 48, pages 122–131. ACM, 2012.
- [SJ02] Tim Sheard and Simon Peyton Jones. Template meta-programming for Haskell. In *Proceedings of the 2002 ACM SIGPLAN workshop on Haskell*, pages 1–16. ACM, 2002.
- [TS00] Walid Taha and Tim Sheard. MetaML and multi-stage programming with explicit annotations. *Theoretical computer science*, 248(1-2):211–242, 2000.

- [UdM20] Sebastian Ullrich and Leonardo de Moura. Beyond notations: Hygienic macro expansion for theorem proving languages. In *International Joint Conference on Automated Reasoning*, pages 167–182. Springer, 2020.