

# Monoids model extensibility

or, Moxy: a language with extensibly extensible syntax

Michael Arntzenius

Carnegie Mellon University  
daekharel@gmail.com

## Abstract

Many languages, libraries, or systems advertise themselves as extensible, but these claims are usually informal. We observe that *monoids* formally characterize much of the essence of extensibility. We present extensibility monoids for several pre-existing systems, including grammars, macros, and monad transformers. To show the utility of this approach, we present Moxy, a syntactically extensible programming language. Moxy’s implementation is surprisingly simple for the power it offers. We hope the use of monoids as an organizing principle of extensibility will offer similar utility in other domains.

**Keywords** Extensibility, extensible languages, macros, monoids, parser-combinators, parsing, syntax.

## 1. Introduction

The concept of extensibility recurs frequently in programming languages research, but lacks a concrete definition. The most one can say is “you know it when you see it”. This paper observes that *monoids* capture some of the essential properties of extensibility. In Section 2, several common examples of extensibility are fruitfully analysed in terms of their monoids.

To show this approach’s utility on a larger scale, in Section 3 we present Moxy, a language designed to be syntactically extensible — to permit the programmer to add syntax for new data types, operations, embedded DSLs, and so on — by using monoids to abstract extensibility. Moxy’s design has several interesting qualities:

1. Moxy is not a Lisp or Scheme; extensibility is not achieved by the sacrifice of syntactic conveniences such as infix operators.<sup>1</sup>
2. Moxy is *self-extensible*: extensions are written in Moxy itself.
3. Moxy extensions are *modular*, that is to say, they are imported just like ordinary library definitions.
4. Moxy extensions are *scoped*; an extension can be imported within a whole file, within a module, or just within a single **let-expression**.

<sup>1</sup>Opinions on the desirability of this property may differ.

5. Moxy extensions are not limited to adding new forms of *expression*. They can also add new forms of *pattern* for use in pattern-matching, for example.
6. Moxy is *meta-extensible*: extensions can *themselves* be extended. For example, a Moxy programmer could create an embedded DSL as an extension library, and leave that DSL open to further extension by future programmers.
7. Moxy is *homogeneously extensible*: All forms of extension — new expressions, new patterns, extensions to some library-defined extensible extension, and so forth — are accomplished by the same mechanism; no special privilege is given to built-in forms of extension, or to expressions over patterns.
8. Finally, Moxy itself is simple: about 2,000 lines of Racket code.

## 2. Monoids as a framework for extensibility

A monoid  $\langle A, \cdot, \varepsilon \rangle$  consists of a set  $A$ , an operation  $(\cdot) : A \times A \rightarrow A$ , and an element  $\varepsilon : A$  satisfying:

$$\begin{aligned} a \cdot (b \cdot c) &= (a \cdot b) \cdot c && \text{associativity} \\ a \cdot \varepsilon &= a && \text{right-identity} \\ \varepsilon \cdot a &= a && \text{left-identity} \end{aligned}$$

A notion of extensibility may be defined by giving a monoid where  $A$  is the set of all possible extensions,  $(\cdot)$  is an operator that *combines* or *composes* two extensions, and  $\varepsilon$  is the “null extension”, representing the lack of any additional behavior.

We show that monoids are a good model for extensibility by analysing five distinct forms of extensibility into their respective monoids: composition of context-free grammars, lisp-style macros, open recursion, middleware in a web framework, and monad transformers.

### 2.1 Composition of context-free grammars

Consider the following grammar for a  $\lambda$ -calculus:

**expressions**  $e ::= x \mid \lambda x. e \mid e e$   
**variables**  $x ::= \dots$  (omitted) ...

It is easy enough to add infix arithmetic to this language:

**expressions**  $e ::= x \mid \lambda x. e \mid e e$   
 $\mid n \mid e \otimes e$   
**numerals**  $n ::= d \mid dn$   
**digits**  $d ::= 0 \mid 1 \mid \dots \mid 9$   
**operators**  $\otimes ::= + \mid - \mid \times \mid /$

It is similarly easy to add lists:

**expressions**  $e ::= x \mid \lambda x. e \mid e e$   
 $\mid [] \mid [es]$   
**expression lists**  $es ::= e \mid e, es$

[Copyright notice will appear here once ‘preprint’ option is removed.]

Formally, each of these three grammars is a separate, complete object. But there is a clear sense in which the two latter languages are composed out of parts: “ $\lambda$ -calculus *plus* infix arithmetic”, “ $\lambda$ -calculus *plus* lists”. And, without even writing its grammar, it is obvious there exists a fourth language: “ $\lambda$ -calculus *plus* infix arithmetic *plus* lists”. Can we formalize this colloquial “*plus*” operator? We can; and it is even a monoid!

**Context-free grammars.** A context-free grammar (CFG) is a tuple  $\langle \Sigma, N, P, e \rangle$ , where  $\Sigma$  is a finite set of terminal symbols,  $N$  is a set (disjoint from  $\Sigma$ ) of nonterminal symbols,  $P$  is a finite set of *production rules*, and  $e \in N$  is an initial non-terminal. Production rules  $p \in P$  have the form  $N \hookrightarrow (\Sigma \cup N)^*$ ; that is,  $P \subseteq N \times (\Sigma \cup N)^*$ .

Fix a terminal set  $\Sigma$ , a (possibly infinite) nonterminal set  $N$ , and an initial nonterminal  $e$ . That is, consider grammars over a known alphabet (e.g. Unicode characters). The non-terminal set is also fixed, but this is no great limitation, since it is permitted to be infinite; we will consider nonterminals to come from an inexhaustible set of abstract names. And we fix an initial nonterminal  $e$ , intended to represent well-formed expressions in some language.

Having fixed all this, a grammar is fully specified by a finite set of production rules  $P$ . Our original  $\lambda$ -calculus is represented by:

$$P_\lambda = \{e \hookrightarrow x, \\ e \hookrightarrow \lambda x.e, \\ e \hookrightarrow e e, \\ \dots \text{(omitted rules for } x) \dots\}$$

**Composing CFGs.** To compose two grammars represented as production-rule-sets, simply take their union. For example, consider the sets of rules *added* when extending our language with infix arithmetic and lists respectively:

$$P_\otimes = \{e \hookrightarrow n, e \hookrightarrow e \otimes e, \\ n \hookrightarrow d, n \hookrightarrow dn, \\ d \hookrightarrow 0, \dots, d \hookrightarrow 9, \\ \otimes \hookrightarrow +, \otimes \hookrightarrow -, \otimes \hookrightarrow \times, \otimes \hookrightarrow /\}$$

$$P_{[]} = \{e \hookrightarrow [], e \hookrightarrow [es], \\ es \hookrightarrow e, es \hookrightarrow e, es\}$$

If we take  $P_\lambda \cup P_\otimes$ , it gives us precisely the grammar of our second language, “ $\lambda$ -calculus *plus* infix arithmetic”.  $P_\lambda \cup P_{[]} is our third language, “ $\lambda$ -calculus *plus* lists”. And  $P_\lambda \cup P_\otimes \cup P_{[]} is the hypothesized fourth language, “ $\lambda$ -calculus *plus* infix arithmetic *plus* lists”.$$

We have thus successfully separated our languages into parts — one which represents  $\lambda$ -calculus, one which represents infix arithmetic, one which represents lists — and found an operator that can combine them again. Finally, note that this operator (union of production-rule-sets) is a (commutative, idempotent) monoid, with  $\emptyset$  as identity.

**Limitations.** While a CFG specifies a syntax, it does not supply an algorithm for parsing. Moreover, it specifies *only* syntax; a programming language also needs *semantics*. Section 3.3 covers one approach to parsing a monoidally-extensible syntax. Section 3.4 describes how Moxy permits extensible semantics.

## 2.2 Macro-expansion

Consider a simple lisp-like syntax:

<b>s-expressions</b>	$e$	::=	$a \mid (es)$
<b>s-expression lists</b>	$es$	::=	$\mid e es$
<b>atoms</b>	$a$	::=	$s \mid n$
<b>numerals</b>	$n$	::=	$0 \mid 1 \mid \dots$
<b>symbols</b>	$s$	::=	$\dots$

We formalize a simplistic version of macro-expansion by defining  $\text{expand}(env, e)$ , which takes an environment  $env$  mapping symbols to their definitions, an s-expression  $e$  to expand, and returns  $ep$  with all macro invocations recursively replaced by their expansions. In pseudocode:

```
expand(env, (s es)) if s in env = expand(env[s](es))
expand(env, (es)) = map(expand(env, .), (es))
expand(env, a) = a
```

(We write  $env[s]$  for the macro-definition of  $s$  in  $env$ , which we take to be a function from an argument-list  $es$  to its immediate expansion under that macro.)

**Where’s the monoid?** If macros yield a form of extensibility, how do they fit into our monoidal framework? First, we must find the set of possible extensions. Intuitively, *macros* are the extensions we are concerned with. Perhaps the set of all macro-definitions? However, there is no obvious operator to “merge” two macro-definitions.

This suggests that our notion of extension is not powerful enough. So instead of single macros, we take our extensions to be *environments* mapping macro-names to their definitions, like the  $env$  argument to  $\text{expand}$ . Our binary operator  $(\cdot)$  merges two environments, so that:

$$s \text{ in } (env_1 \cdot env_2) \quad \text{iff} \quad (s \text{ in } env_1) \vee (s \text{ in } env_2)$$

$$(env_1 \cdot env_2)[s] = \begin{cases} env_1[s] & \text{if } s \text{ in } env_1 \\ env_2[s] & \text{otherwise} \end{cases}$$

By convention,  $(\cdot)$  is left-biased: if the same symbol is defined in both arguments, it uses the left one.

Finally, note that  $(\cdot)$  is associative, and forms an identity with the empty environment; that is,  $(\cdot)$  is a monoid.

**Monoid-parameterized functions.** Observe that  $\text{expand}$  is a function parameterized by a value of the monoid representing extensions. It demonstrates one way of giving *semantics* to a notion of extensibility: interpret extensions into functions. This is a pattern we will see again in Moxy’s parse and compile phases.

## 2.3 Open recursion and mixins

*Open recursion* is the problem of, first, defining a set of mutually-recursive functions *by parts*; that is, permitting the programmer to define only some of the functions, and complete the set by defining the rest later; and, second, of allowing the behavior of these functions to be *overridden* or *extended*.

For example, consider the functions  $\text{even}(n)$  and  $\text{odd}(n)$ , of type  $\mathbb{N} \rightarrow \text{Bool}$ , defined by:

```
even(0) = true
even(n) = odd(n - 1)
odd(0) = false
odd(n) = even(n - 1)
```

Definition *by parts* means that we could define  $\text{even}$  and  $\text{odd}$  separately, as *mixins*, and later combine them into an *instance* that implements both:

```
mixin EvenMixin where
  inherit odd
  even(0) = true
  even(n) = self.odd(n - 1)
end

mixin OddMixin where
  inherit even
  odd(0) = false
  odd(n) = self.even(n - 1)
end
```

**instance** EvenOdd **where**

**use** EvenMixin  
**use** OddMixin

**end**

EvenOdd.even(10) — *returns true*

*Overriding* means that a mixin can extend the behavior of a function in a fashion similar to traditional object-oriented subclassing:

**mixin** EvenSpyMixin **where**

**inherit** odd  
even( $n$ ) =  
  print( $n$ );  
  **super**.even( $n$ )

**end**

**instance** EvenOddDebug **where**

**use** EvenMixin  
**use** OddMixin  
**use** EvenSpyMixin

**end**

— *prints 2, then prints 0, then returns true*

EvenOddDebug.even(2)

### 2.3.1 Semantics of mutual recursion.

Setting aside for a moment the problem of open recursion, consider ordinary, “closed”, mutual recursion. Suppose we wish to implement some signature  $\Sigma = \langle N, \tau \rangle$  of mutually recursive functions, where  $N$  is a set of names and the function  $\tau : N \rightarrow \text{type}$  assigns each name a type. An *implementation* of  $\Sigma$  is a function giving to each name  $n$  a value of type  $\tau(n)$ . We write  $\text{Impl}_\Sigma$  for the type of implementations of the signature  $\Sigma$ :

$$\text{Impl}_{\langle N, \tau \rangle} = \Pi(n : N). \tau(n)$$

We represent a group of mutually-recursive function *definitions* as a function from implementations to implementations: we take a self object, and to recursively call the function  $n$ , we call  $\text{self}(n)$ . In this manner we make self-reference explicit. We write  $\text{Defn}_\Sigma$  for the type of mutually-recursive function definitions represented this way:

$$\text{Defn}_\Sigma = \text{Impl}_\Sigma \rightarrow \text{Impl}_\Sigma$$

For example, our original mutually-recursive definition of even and odd is represented by:

$$\begin{aligned} \Sigma &= \langle N, \tau \rangle \\ N &= \{\text{even}, \text{odd}\} \\ \tau(\text{even}) &= \mathbb{N} \rightarrow \text{Bool} \\ \tau(\text{odd}) &= \mathbb{N} \rightarrow \text{Bool} \\ \text{EvenOdd} &: \text{Defn}_\Sigma \\ \text{EvenOdd}(\text{self})(\text{even})(0) &= \text{true} \\ \text{EvenOdd}(\text{self})(\text{even})(n) &= \text{self}(\text{odd})(n-1) \\ \text{EvenOdd}(\text{self})(\text{odd})(0) &= \text{false} \\ \text{EvenOdd}(\text{self})(\text{odd})(n) &= \text{self}(\text{even})(n-1) \end{aligned}$$

To obtain the desired implementations of even and odd, we take the least-fixed-point of the EvenOdd function:

$$\begin{aligned} \text{even} &= \text{fix}(\text{EvenOdd})(\text{even}) \\ \text{odd} &= \text{fix}(\text{EvenOdd})(\text{odd}) \end{aligned}$$

For some function  $\text{fix} : (\alpha \rightarrow \alpha) \rightarrow \alpha$  satisfying  $\text{fix}(f) = f(\text{fix}(f))$  for appropriate  $\alpha$  (here,  $\alpha = \text{Impl}_\Sigma$ ).

### 2.3.2 Mixins as a monoid

Let’s apply our monoid methodology to the problem of semantics for open recursion. We’ll tackle definition by parts first, and later expand our semantics to cover overriding as well.

First, we must determine our set of extensions. The extensions we are concerned with are “mixins”: partial definitions of a set of mutually recursive functions. The definitions in a mixin have access to a self object, allowing mutual- or self-recursion.

At first glance, this seems very similar to the way we formalized closed sets of mutually-recursive definitions. The only difference is that mixins are permitted to be *partial*, and omit a definition for a particular function in the signature.

However, we can use the same type  $\text{Defn}_\Sigma$  to represent mixins if for names  $n$  which the mixin  $M$  does not define, we simply let  $M(\text{self})(n) = \text{self}(n)$ , passing the undefined method through unchanged. Thus a mixin is represented by a function from implementations to implementations: it takes a self object, and returns that object updated with each of the functions the mixin implements.

For example, EvenMixin and OddMixin are represented by the following functions:

$$\begin{aligned} \text{EvenMixin}(\text{self})(\text{even})(0) &= \text{true} \\ \text{EvenMixin}(\text{self})(\text{even})(n) &= \text{self}(\text{odd})(n-1) \\ \text{EvenMixin}(\text{self})(\text{odd}) &= \text{self}(\text{odd}) \\ \text{OddMixin}(\text{self})(\text{odd})(0) &= \text{false} \\ \text{OddMixin}(\text{self})(\text{odd})(n) &= \text{self}(\text{even})(n-1) \\ \text{OddMixin}(\text{self})(\text{even}) &= \text{self}(\text{even}) \end{aligned}$$

At first, this representation seems pointless, like trying to fit a square peg (mixins) into a round hole ( $\text{Defn}_\Sigma$ ). After all, if we take  $\text{fix}(\text{OddMixin})(\text{odd})$ , the resulting function diverges for all non-zero inputs!

However, observe that

$$(\text{EvenMixin} \circ \text{OddMixin}) = \text{EvenOdd}$$

(as can be verified by some tedious calculations)!

That is, that we can “combine” mixins by composing them as functions; our monoidal operator is  $\circ$ . Our identity is just the identity function at type  $\text{Impl}_\Sigma$ . Function composition is associative, so we are done.

### 2.3.3 Mixins with overriding

$$\begin{aligned} \text{Mixin}_\Sigma &= \text{Impl}_\Sigma \rightarrow \text{Impl}_\Sigma \rightarrow \text{Impl}_\Sigma \\ \varepsilon(\text{self})(\text{super}) &= \text{super} \\ (f \cdot g)(\text{self})(\text{super}) &= f(\text{self})(g(\text{self})(\text{super})) \end{aligned}$$

$$\begin{aligned} \text{EvenMixin}(\text{self})(\text{super})(\text{even})(0) &= \text{true} \\ \text{EvenMixin}(\text{self})(\text{super})(\text{even})(n) &= \text{self}(\text{odd})(n-1) \\ \text{EvenMixin}(\text{self})(\text{super})(\text{odd}) &= \text{super}(\text{odd}) \\ \text{OddMixin}(\text{self})(\text{super})(\text{odd})(0) &= \text{false} \\ \text{OddMixin}(\text{self})(\text{super})(\text{odd})(n) &= \text{self}(\text{even})(n-1) \\ \text{OddMixin}(\text{self})(\text{super})(\text{even}) &= \text{super}(\text{even}) \\ \text{EvenSpyMixin}(\text{self})(\text{super})(\text{odd}) &= \text{super}(\text{odd}) \\ \text{EvenSpyMixin}(\text{self})(\text{super})(\text{even})(n) &= \mathbf{print}(n); \\ &\quad \text{super}(\text{even}(n)) \end{aligned}$$

To actually produce an implementation:

$$\begin{aligned} \text{impl} &: \text{Mixin}_\Sigma \rightarrow \text{Impl}_\Sigma \\ \text{impl}(f) &= \text{fix}(\lambda s. f(s)(\perp)) \end{aligned}$$

## 2.4 Web framework middleware

We observe that middleware in many web frameworks (for example, Django) is essentially a stack of pairs of functions,  $(\text{Request} \rightarrow \text{Request}) \times (\text{Response} \rightarrow \text{Response})$ , which are composed together via the monoid:

$$\begin{aligned} \varepsilon &= \langle \text{id}_{\text{Request}}, \text{id}_{\text{Response}} \rangle \\ \langle f_{\text{req}}, f_{\text{resp}} \rangle \cdot \langle g_{\text{req}}, g_{\text{resp}} \rangle &= \langle f_{\text{req}} \circ g_{\text{req}}, g_{\text{resp}} \circ f_{\text{resp}} \rangle \end{aligned}$$

<b>top-level expressions</b>	$t$	$::=$	$d \mid \mathbf{module} N \{t^*\}$
	$e$	$::=$	$Pn \mid PN \mid l \mid (e) \mid e \oplus e$ $\mid (p, \dots) e \mid e(e, \dots)$ $\mid \mathbf{let} d^* \mathbf{in} e$ $\mid \mathbf{case} e [ \mid p \rightarrow e ]^*$
<b>declarations</b>	$d$	$::=$	$\mathbf{val} p = e$ $\mid \mathbf{fun} n(p, \dots) = e [ \mid n(p, \dots) = e ]^*$ $\mid \mathbf{tag} N[(n, \dots)]$ $\mid \mathbf{rec} d [\mathbf{and} d]^*$ $\mid \mathbf{open} PX$
<b>patterns</b>	$p$	$::=$	$x \mid l \mid PX[(p, \dots)]$
<b>operators</b>	$\oplus$	$::=$	$+ \mid - \mid * \mid / \mid == \mid <= \mid >= \mid < \mid > \mid ;$ $\dots$
<b>module paths</b>	$P$	$::=$	$[N..]^*$
<b>capital names</b>	$N$		
<b>names</b>	$n$		
<b>literals</b>	$l$		

**Figure 1.** Grammar of Moxy, sans extensions

And then composed onto a base handler function  $h : \text{Request} \rightarrow \text{Response}$ :

$$\text{withMiddleware}(\langle m_{\text{in}}, m_{\text{out}} \rangle, h) = m_{\text{out}} \circ h \circ m_{\text{in}}$$

Middleware “in the wild” is more complicated, having to deal with details such as error-handling, the possibility of early exit, and so forth, but still forms a monoid. The fact that it forms a monoid can even be observed by the way one specifies what middleware to use: via a list, i.e. an element of the free monoid.

## 2.5 Monad transformers

Broadly speaking, monads as they appear in Haskell represent *effects* which a computation has access to: for example, `Maybe` represents the possibility of failure; `State s` represents access to a single mutable location of type  $s$ ; `List` represents backtracking nondeterminism.

It is often useful to have multiple kinds of effect; for example, failure *and* state. For this Haskell uses *monad transformers*, type-level functions from monads to monads. Letting `Hask` be the category of Haskell types, a monad such as `State s` has kind `Hask → Hask`; a monad transformer such as `StateT s` has kind `(Hask → Hask) → (Hask → Hask)`.

To combine multiple effects, one creates a stack of monad transformers, terminating it with the identity monad `Identity`. To combine failure (`Maybe`) with binary state (`State Bool`) and logging (`Writer [String]`), we write:

```
WriterT [String] (StateT Bool (MaybeT Identity))
```

Clearly monad transformers represent a notion of extensibility, namely “adding effects” to a pure base language. Unsurprisingly, they form a monoid; the operator is simply composition, and the identity is the monad transformer `IdentityT`. The above can be rewritten (assuming a type-level composition operator  $\circ$ ) as:

```
(WriterT [String] ◦ StateT Bool ◦ MaybeT) Identity
```

## 3. Moxy

### 3.1 A brief introduction to Moxy, sans extensions

Moxy exists to test the hypothesis that monoids are a good way of structuring extensibility. In other respects it is a banal, humdrum language.

The basic grammar of Moxy is given in figure 1. “ $e, \dots$ ” indicates a comma-separated list of expressions  $e$ ; similarly for patterns, “ $p, \dots$ ”.

Syntactically, Moxy is moderately MLish; semantically, Moxy is somewhat Schemelike. Evaluation is eager; functions take multiple arguments and return one value; recursion is the only looping construct; pattern-matching is the only conditional construct (`if` is implemented as a syntax extension). Moxy has a simple module system that only handles namespacing (no ML-style functors). Strings and numbers are built-in; booleans are defined in the implicitly-imported prelude.

As an example, here is a naïve recursive implementation of the fibonacci function in Moxy:

```
fun fib(0) = 1
  | fib(1) = 1
  | fib(n) = fib(n-1) + fib(n-2)
```

Case is syntactically significant in Moxy: capitalized names are used for modules, tags, and extension points, uncapitalized for everything else.

### 3.2 Moxy’s approach to extensibility: Extension points

Moxy aims to be *meta-extensible*: to permit extensions (for example, a DSL for parsers) that are themselves extensible (for example, defining `a+` as an abbreviation for `aa*`). It therefore seems difficult to specify in advance what the monoid representing Moxy extensions should be.

Moxy’s solution is to let the programmer define their own extension monoids, which Moxy calls *extension points*. To define an extension point, the programmer gives it a name (a Moxy identifier), a value representing the identity element, and a function representing the monoid operator. For example:

```
extension ExtCounts(0, \ (x,y) x+y)
```

This defines an extension point named `ExtCounts` with the monoid  $(\mathbb{N}, +, 0)$ . As Moxy is dynamically typed, the fact that the intended domain of `ExtCounts` is  $\mathbb{N}$  exists only in the programmer’s mind.<sup>2</sup>

Having defined an extension point, we can extend it with an **extend** declaration:

```
extend ExtCounts with 2 + 2
```

Within the scope of this declaration, `ExtCounts` will have a value 4 higher than it otherwise would. Of course, this is quite useless without some way to observe the value of an extension point and use it in parsing.

For this purpose, Moxy comes with built-in extension points which allow extending Moxy’s built-in syntax classes: expressions, declarations, and patterns. For example, the `InfixExprs` extension point permits adding new infix operators to the expression language:

```
extend InfixExprs with
{ TSYM("."): { precedence = 10,
              parse = [...omitted...] } }
```

### 3.3 Parsing Moxy

Moxy uses monadic parser-combinators after the style of `Parsec` to implement a Pratt-style recursive descent parser. In practice what the latter means is that infix operators are handled by means of a table mapping the token for a given operator (say, `+` for addition) to its precedence, paired with a parser for its right-hand-side. When parsing an expression we keep track of the precedence we are currently parsing at, and if we encounter an operator of looser

<sup>2</sup>Indeed, there is nothing to say that the intended domain is not  $\mathbb{Z}$  instead.

precedence we stop parsing and return. Otherwise we recursively invoke the parser for the operator we found, with the expression we've parsed so far as an additional argument.

This permits arbitrary infix, suffix, and mixfix operators. However, it requires that each infix operator have a unique token identifying it; no overloading is possible (at least, not in the parser).

Moxy currently has a separate tokenization phase, which unfortunately limits the expressiveness of parse extensions. We believe this can be eliminated in future by means of a scannerless-parsing technique, such as that used by Parsec's `Text.Parsec.Token` module.

The monad which Moxy uses for parsing, in addition to behaving like Parsec, also acts as a Reader `ParseEnv`, where `ParseEnv` is a datastructure representing all extensions to all extension points currently in scope. A `ParseEnv` is effectively a mapping from extension points to their values. It is the parameterization of the parser by this value that lets extensions affect parsing (just as the parameterization of `expand` by a macro-environment lets macro-definitions affect macro-expansion).

For example, when the parser tries to parse the start of an expression, it first examines the current value of the `Exprs` extension point. `Exprs`' domain is dictionaries mapping tokens to parsers; its monoid operation is right-biased merge and its identity is the empty dictionary. If the next token is present in the value of `Exprs`, then we invoke the parser it is bound to. Otherwise we try built-in parse productions such as literals, variables, and function application.

Similar techniques are used for extending declarations and patterns.

### 3.4 Compiling Moxy

Nodes in Moxy's AST are instances of abstract interfaces. For example, an expression is merely something which *knows how to compile itself*. That is, expressions are "records"<sup>3</sup> with a functional compile field. This compile function, when supplied with a *resolve environment*, returns the intermediate-representation<sup>4</sup> compilation of the expression in question.

The *resolve environment* is effectively merely the lexical environment of the expression, a dictionary which tells the compiler which variable names are in scope and how references to them are to be compiled.

## 4. Prior work

Lisp and Scheme are obvious predecessors in the search for extensible forms of programming. Moxy is in large part an attempt to replicate the ease and power of Lisp-family macros in a non-Lisp setting.

SugarJ does almost everything Moxy does and more [1]. However, it weighs in at 26k LOC, while Moxy is a mere 2k LOC. Moxy can use extensions per-scope, not just per-file. Moxy has a REPL, but SugarJ's approach is probably compatible with a REPL as well. Moxy allows for non-syntactic notions of scoped extensibility (but I have no motivating examples). The Moxy approach currently has no story for integrating with existing languages (but it probably could be done).

Moxy's use of a Pratt-style parser to aid extensibility is similar to that of Magpie. [4]

Other previous syntactically extensible systems include OMeta [3], Sweet.js, Seed7, and Omar et al's Type-Specific Languages [2].

## References

- [1] **SugarJ: Library-based Syntactic Language Extensibility**. Sebastian Erdweg, Tillman Rendel, Christian Kästner and Klaus Ostermann. In *Proceedings of Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 391-406. ACM, 2011.
- [2] **Safely Composable Type-Specific Languages**. C. Omar, D. Kurilova, L. Nistor, B. Chung, A. Potanin and J. Aldrich. *European Conference on Object-Oriented Programming (ECOOP 2014) Uppsala, Sweden, July 28 – August 1, 2014*.
- [3] **Experimenting with Programming Languages**. Alessandro Warth. Technical Report TR-2008-003, Viewpoints Research Institute, 2009.
- [4] **Extending Syntax from Within a Language**. Bob Nystrom. <http://journal.stuffwithstuff.com/2011/02/13/extending-syntax-from-within-a-language/>, 2011-02-13, accessed 2014-09-25.
- [5] **Parsec: Direct Style Monadic Parser Combinators For The Real World**. Daan Leijen and Erik Meijer. Technical Report UU-CS-2001-27, Department of Computer Science, Universiteit Utrecht, 2001.

<sup>3</sup> In this case, represented as hash-tables.

<sup>4</sup> In this case, s-expressions representing Racket code.