

1 From Macros to DSLs: The Evolution of Racket*

2 Ryan Culpepper

3 PLT

4 ryanc@racket-lang.org

5 Matthias Felleisen

6 PLT

7 matthias@racket-lang.org

8 Matthew Flatt

9 PLT

10 mflatt@racket-lang.org

11 Shriram Krishnamurthi

12 PLT

13 sk@racket-lang.org

14 — Abstract —

15 The Racket language promotes a language-oriented style of programming. Developers create many
16 domain-specific languages, write programs in them, and compose these programs via Racket code.
17 This style of programming can work only if creating and composing little languages is simple and
18 effective. While Racket's Lisp heritage might suggest that macros suffice, its design team discovered
19 significant shortcomings and had to improve them in many ways. This paper presents the evolution
20 of Racket's macro system, including a false start, and assesses its current state.

21 **2012 ACM Subject Classification** Software and its engineering → Semantics

22 **Keywords and phrases** design principles, macros systems, domain-specific languages

23 **Digital Object Identifier** 10.4230/LIPIcs...0

* Over 20 years, this work was partially supported by our host institutions (Brown University, Northeastern University, Prague Technical University, and University of Utah) as well as several funding organizations (AFOSR, Cisco, DARPA, Microsoft, Mozilla, NSA, and NSF).



© Culpepper, Felleisen, Flatt, Krishnamurthi;
licensed under Creative Commons License CC-BY

Editors: ; Article No. 0; pp. 0:1–0:19



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

24 1 Macros and Domain-Specific Languages

25 The Racket manifesto [19, 20] argues for a *language-oriented programming* (LOP) [12, 43]
 26 approach to software development. The idea is to take Hudak’s slogan of “languages [as]
 27 the ultimate abstractions” [32] seriously and to program with domain-specific languages
 28 (DSLs) as if they were proper abstractions within the chosen language. As with all kinds
 29 of abstractions, programmers wish to create DSLs, write programs in them, embed these
 30 programs in code of the underlying language, and have them communicate with each other.

31 According to the Lisp worldview, a language with macros supports this vision particularly
 32 well. Using macros, programmers can tailor the language to a domain. Because programs in
 33 such tailored languages sit within host programs, they can easily communicate with the host
 34 and each other. In short, creating, using, and composing DSLs looks easy.

35 Macros alone do not make DSLs, however, a lesson that the Racket team has learned
 36 over 20 years of working on a realization of language-oriented programming. This paper
 37 recounts Racket’s history of linguistic mechanisms designed to support language-oriented
 38 programming; it formulates desiderata for DSL support based on, and refined by, the Racket
 39 team’s experiences; it also assesses how well the desiderata are met by Racket’s current
 40 capabilities. The history begins with Scheme’s hygienic macros, which, in turn, derive
 41 from Lisp (see sec. 2). After a false start (see sec. 4), the Racket designers switched to
 42 procedural, hygienic macros and made them work across modules; they also strictly separated
 43 expansion time from run time (see sec. 5). Eventually, they created a meta-DSL for writing
 44 macros that could properly express the grammatical constraints of an extension, check them,
 45 and synthesize proper error messages (see sec. 6). A comparison between general DSL
 46 implementation desiderata (see sec. 3) and Racket’s capabilities shows that the language’s
 47 support for a certain class of DSLs still falls short in several ways (see sec. 7).

48 *Note* This paper does *not* address the safety issues of language-oriented programming. As
 49 a similar set of authors explained in the Racket Manifesto [19], language-oriented programming
 50 means programmers use a host language to safely compose code from many small pieces
 51 written in many different DSLs and use the very same host language to implement the DSLs
 52 themselves. Hence language-oriented programming clearly needs the tools to link DSLs safely
 53 (e.g., via contracts [21] or types [40]) and to incorporate systems-level protection features
 54 (e.g., sandboxes and resource custodians [29]).

55 Some Hints For Beginners on Reading Code

56 We use Racket-y constructs (e.g., `define-syntax-rule`) to illustrate Lisp and Scheme macros.
 57 Readers familiar with the original languages should be able to reconstruct the original ideas;
 58 beginners can experiment with the examples in Racket.

Lisp’s S-expression construction	Racket’s code construction
<code>'</code> S-expression quote	<code>#'</code> code quote
<code>`</code> Quine quasiquote	<code>#`</code> code quasiquote
<code>,</code> Quine unquote	<code>#,</code> code unquote
<code>@,</code> list splicing	<code>#@,</code> code splicing

■ **Figure 1** Hints on strange symbols

59 For operations on S-expressions, i.e., the nested and heterogeneous lists that represent
 60 syntax trees, Lisp uses `car` for **first**, `cdr` for **rest**, and `cadr` for **second**. For the convenient
 61 construction of S-expressions, Lisp comes with an implementation of Quine’s quasiquotation

62 idea that uses the symbols shown on the left of fig. 1 as short-hands: quote, quasiquote,
 63 unquote, and unquote-splicing. By contrast, Racket introduces a parallel data structure
 64 (syntax objects). To distinguish between the two constructions, the short-hand names are
 65 prefixed with #.

66 2 The Lisp and Scheme Pre-history

67 Lisp has supported macros since 1963 [31], and Scheme inherited them from Lisp in 1975 [41].
 68 Roughly speaking, a Lisp or Scheme implementation uses a *reader* to turn the sequence of
 69 characters into a concrete tree representation: *S-expressions*. It then applies an *expander* to
 70 obtain the abstract syntax tree(s) (AST). The expander traverses the S-expression to find
 71 and eliminate uses of macros. A macro rewrites one S-expression into a different one, which
 72 the expander traverses afresh. Once the expander discovers a node with a syntax constructor
 73 from the core language, say `lambda`, it descends into its branches and recursively expands
 74 those. The process ends when the entire S-expression is made of core constructors.

75 A bit more technically, macros are functions of type

76 $S\text{-expression} \rightarrow S\text{-expression}$

77 The `define-macro` form defines macros, which are written using operators on S-expressions.
 78 Each such definition adds a macro function to a table of macros. The expander maps an
 79 S-expression together with this table to an intermediate abstract syntax representation:

80 $S\text{-expression} \times \text{TableOf}[\text{MacroId}, (S\text{-expression} \rightarrow S\text{-expression})] \rightarrow \text{AST}$

81 The AST is an internal representation of an S-expression of only core constructors.

82 See the left-hand side of fig. 2 for the simple example of a `let` macro. As the comments
 83 above the code say, the `let` macro extends Racket with a block-structure construct for local
 84 definitions. The macro implementation assumes that it is given an S-expression of a certain
 85 shape. Once the definition of the `let` macro is recognized, the expander adds the symbol
 86 `'let` together with the specified transformer function to the macro table. Every time the
 87 macro expander encounters an S-expression whose head symbol is `'let`, it retrieves the
 88 macro transformer and calls it on the S-expression. The function deconstructs the given
 89 S-expression into four pieces: `decl`, `lhs`, `rhs`, `body`. From these, it constructs an S-expression
 90 that represents the immediate application of a `lambda` function.

```

;; PURPOSE extend Racket with block-oriented, local bindings
;;
;; ASSUME the given S-expression has the shape
;; (let ((lhs rhs) ...) body ...)
;; FURTHERMORE ASSUME:
;; (1) lhs ... is a sequence of distinct identifiers
;; (2) rhs ..., body ... are expressions
;; PRODUCE
;; ((lambda (lhs ...) body ...) rhs ...)

(define-macro (let e)
  (define decl (cadr e))
  (define lhs (map car decl))
  (define rhs (map cadr decl))
  (define body (caddr e))
  ;; return
  `((lambda ,lhs ,@body) ,@rhs))

(define-syntax-rule
  (let ((lhs rhs) ...) body ...)
  ;; rewrites above pattern to template below
  ((lambda (lhs ...) body ...) rhs ...))

```

■ **Figure 2** Macros articulated in plain Lisp vs Kohlbecker's macro DSL

91 Macros greatly enhance the power of Lisp, but their formulation as functions on S-
 92 expressions is both error-prone and inconvenient. As fig. 2 shows, the creator of the function
 93 makes certain assumption about the shape of the given S-expression, which are not guaranteed
 94 by the macro expansion process. Yet even writing just a transformation from the assumed
 95 shape of the S-expression to the properly shaped output requires bureaucratic programming
 96 patterns, something a macro author must manage and easily causes omissions and oversights.

97 For concreteness, consider the following problems in the context of `let` macro:

- 98 1. The S-expression could be an improper list. The transformer, as written, does not notice
 99 such a problem, meaning the compilation process ignores this violation of the implied
 100 grammar of the language extension.
- 101 2. The S-expression could be too short. Its second part might not be a list. If it is a list, it
 102 may contain an S-expression without a `cadr` field. In these cases, the macro transformer
 103 raises an exception and the compilation process is aborted.
- 104 3. The S-expression has the correct length but its second part may contain lists that contain
 105 too many S-expressions. Once again, the macro transformer ignores this problem.
- 106 4. The S-expression may come with something other than an identifier as the `lhs` part of a
 107 local declaration. Or, it may repeat the same identifier as an `lhs` part of the second clause.
 108 In this case, the macro generates code anyways, relying on the rest of the compilation
 109 process to discover the problems. When these problems are discovered,
 110 a. it may have become impossible to report the error in terms of source code, meaning a
 111 programmer might not understand where to look for the syntax error.
 112 b. it has definitely become impossible to report errors in terms of the language extension,
 113 meaning a programmer might not comprehend the error message.
- 114 5. The author of the macro might forget the unquote `,` to the left of `lhs`. In many members
 115 of the Lisp family, the resulting code would be syntactically well formed but semantically
 116 rather different from the intended one. In particular, conventional Lisp would generate a
 117 function that binds all occurrences of `lhs` in `body` via this newly created `lambda`—a clear
 118 violation of the intended scoping arrangements expressed in the comments.

119 In short, if the S-expression fails to live up to the stated assumptions, the macro transformation
 120 may break, ignore code, or generate code that some later step in the compilation process
 121 recognizes as an error but describes in inappropriate terms. If the programmer makes even a
 122 small mistake, strange code may run and is likely to cause inexplicable run-time errors.

123 Kohlbecker's dissertation research on macros greatly improves this situation [34, 35, 36].
 124 His macro system for Scheme 84 adds two elements to the macro writer's toolbox. The first
 125 is a DSL for articulating macro transformations as rewriting rules consisting of a pattern
 126 and a template. The revised macro expander matches S-expressions against the specified
 127 patterns; if there is a match, the template is instantiated with the resulting substitution. This
 128 DSL removes programming patterns from macro definitions and, to some extent, eliminates
 129 problems 1 through 3 from above. For an example, see the right-hand side of fig. 2.

130 *Note* We call Lisp-style macros *procedural* and Kohlbecker's approach *declarative*.

131 The second novel element is hygienic expansion. By default, Kohlbecker's macro expander
 132 assumes that identifiers contained in the source must be distinct from macro-generated
 133 identifiers in binding positions. As such, it eliminates the need to explicitly protect against
 134 accidental interference between the macro's lexical scopes and those of its use contexts—that
 135 is, yet another programming pattern from macro code. At a minimum, this hygienic expander
 136 would not bind `lhs` in `body` as indicated in problem 5 above.

137 Further work [2, 7, 14] refined and improved the pattern-oriented approach to specifying
 138 macros as well as hygienic macro expansion. The `define-syntax-rule` construct and

139 hygienic expansion became part of the Scheme standard by the late 1990s [1]. Starting
140 in 1988 and in parallel to the Scheme standardization process, Dybvig et al. [14] designed
141 and implemented a macro definition construct, `define-syntax-cases` (in actual code it
142 requires a combination of `define-syntax` and `syntax-case`) that merged the procedural
143 and declarative elements of the Lisp world. Dybvig et al. also switched from S-expressions
144 to trees of syntax objects. These trees included source locations so that the error handling
145 code could try to point back to the surface code (problem 4a above).

146 Starting in the late 80s, researchers explored other ways to facilitate the work of macro
147 authors, including two relevant to creating DSLs from macros. Dybvig et al. [13] invented
148 expander-passing macros. Macro authors would write their own expanders and use different
149 ones in different macros. At an abstract level, expansion-passing style anticipates the need
150 for checking static attributes. Blume [3] as well as Kelsey and Reese [33] added modules that
151 could export and import macros. Such modules allow macro programmers to encapsulate
152 bundles of macros, a first step towards encapsulating a DSL's design and implementation.

153 **3 DSLs Require More than Bunches of Macros**

154 Scheme-style macros greatly improve on Lisp's as far as the *extension* of an existing language
155 is concerned. A developer can add concise and lexically correct macros to a program and may
156 immediately use them, for writing either ordinary run-time code or additional macros. This
157 immediacy is powerful and enticing because a programmer never has to leave the familiar
158 programming environment, use external tools, or mess with scaffolding setups.

159 The idea of macros is also easy to comprehend at the abstract level. Conceptually, a macro
160 definition adds a new alternative to one of Racket's grammatical productions: definitions
161 or expressions. The declarative approach makes it easy to specify simple S-expression
162 transformers in a straightforward manner; hygienic macro expansion guarantees the integrity
163 of the program's lexical scope.

164 The problem is that a language extension provides only a false sense of a purpose-tailored
165 language. On one hand, a programmer who uses a bunch of macro-based language extensions
166 as if it were a self-contained DSL must code with an extreme degree of self-discipline. On
167 the other hand, the macro system fails to support some of the traditional advantages of
168 using DSLs: catching mistakes in the parsing or type-checking stage; exploiting constraints
169 to generate optimized code; or link with/target tailor-made run-time functions.

170 Conventionally, the creation of DSLs demands a pipeline of compiler passes:

- 171 1. a *parser*, based on explicit specification of a domain-specific *vocabulary* and a *grammar*,
172 that reports errors at the DSL's source level;
- 173 2. a *static semantics*, because one goal of migrating from an application interface to a DSL
174 is to enforce certain constraints statically;
- 175 3. a *code generation* and *optimization* pass, because another goal of introducing DSLs is to
176 exploit the static or linguistic constraints for improved performance; and,
- 177 4. a *run-time system*, because (1) the host language may lack pieces of functionality or (2)
178 the target language might be distinct from the host language.

179 Scheme macros *per se* do not support the creation of such pipelines or its proper encapsulation.

180 The Racket designers noticed some of these problems when they created their first
181 teaching languages [17, 18]. In response, they launched two orthogonal efforts to support the
182 development of DSLs via language-extension mechanisms with the explicit goal of retaining
183 the ease of use of the latter:

184 ■ One concerned the encapsulation of DSLs and support for some traditional passes. This
 185 idea was to develop a module system that allows the export and import of macros and
 186 functions while also retaining a notion of separate compilation for modules.
 187 ■ The other aimed at a mechanism for easily expressing a macro’s assumptions about its
 188 input and synthesizing error messages at the appropriate level, i.e., the problems from
 189 sec. 2. The results would also help with implementing DSLs via modules.
 190 While sec. 4 reports on an ambitious, and abandoned, attempt to address these problems all
 191 at once, secs. 5 and 6 describe the tools that Racket eventually provided to DSL designers
 192 and implementors.

193 4 Ambitious Beginnings

194 When the Racket designers discovered the shortcomings of a traditional Scheme macro system,
 195 they decided to address them with three innovations. First, they decided to move beyond
 196 the traditional S-expression representation of syntax and instead use a richly structured
 197 one (see sec. 4.1). Second, they realized that macros needed to work together to implement
 198 context-sensitive checks. To this end, they supplemented declarative macros with procedural
 199 *micros* that could deal with attributes of the expansion context (see sec. 4.2). Finally they
 200 decided to use modules as the containers of macro-based DSL implementations as well as the
 201 units of DSL use (see sec. 4.3).

202 4.1 From S-expressions to Syntax Objects

203 To track source locations across macro expansion, Racket—like Dybvig’s Chez Scheme—
 204 introduced a *syntax object* representation of the surface code, abandoning the conventional
 205 S-expression representation. Roughly speaking, a syntax object resembles an S-expression
 206 with a structure wrapped around every node. At a minimum, this structure contains source
 207 locations of the various tokens in the syntax. Using this information, a macro expander can
 208 often pinpoint the source location of a syntax error, partially solving problem 4a from sec. 3.

209 4.2 The Vocabularies of Micros

210 Recall that a macro is a function on the syntax representation. Once this representation
 211 uses structures instead of just S-expressions, the signature of a macro has to be adapted:

212 `Syntax-Object` \longrightarrow `Syntax-Object`

213 Of course, this very signature says that macros cannot naturally express¹ communication
 214 channels concerning attributes of the expansion context.

215 Krishnamurthi et al.’s work [38] supplements macros with *micros* to solve this issue. Like
 216 `define-macro`, `define-micro` specifies a function that consumes the representation of a
 217 syntax. Additionally, it may absorb any number of `Attribute` values so that collections of
 218 micros can communicate contextual information to each other explicitly:

219 `Syntax-Object` \longrightarrow (`Attribute` ... \longrightarrow `Output`)

220 As this signature shows, a micro also differs from a macro in that the result is some arbitrary
 221 type called `Output`. This type must be the same for all micros that collaborate but may

¹ A macro author could implement this form of communication via a protocol that encodes attributes as syntax objects. We consider an encoding unnatural and therefore use the phrase “naturally express.”

222 differ from one collection of micros to another. For macro-like micros, Output would equal
 223 `Syntax-Object`. By contrast, for an embedded compiler Output would be AST, meaning the
 224 type of abstract syntax trees for the target language. This target language might be Racket,
 225 but it could also be something completely different, such as GPU assembly code. The Racket
 226 team did not explore this direction at the time.

227 As this explanation points out, micros for DSLs must be thought of as members of a
 228 collection. To make this notion concrete, Krishnamurthi et al. also introduce the notion of
 229 a *vocabulary*. Since collections of macros and micros determine the “words” and “sentence
 230 structure” of a DSL, a vocabulary represents the formal equivalent of a dictionary and
 231 grammar rules. The micros themselves transform “sentences” in an embedded language into
 232 meaningful—that is, executable—programs.

233 In Krishnamurthi et al.’s setting, a vocabulary is created with `(make-vocabulary)` and
 234 comes with two operations: `define-micro`, which adds a micro function to a given vocabulary,
 235 and `dispatch`, which applies a micro to an expression in the context of a specific vocabulary.

```

;; type Output = RacketAST                                ;; type Output = RacketType
                                                         (define type-check (make-vocabulary))

(define compiler (make-vocabulary))                       - - - elided - - -

- - - elided - - -                                       (define-micro type-check
                                                         (if cond then else)
                                                         ==>
                                                         (lambda (Γ)
                                                         ;; first block
                                                         (define (tc t)
                                                         ((dispatch t type-check) Γ))
                                                         (define cond-type (tc cond))
                                                         (unless (type== cond-type Boolean)
                                                         (error _ _ _ elided _ _ _))
                                                         (define then-type (tc then))
                                                         (define else-type (tc else))
                                                         (unless (type== then-type else-type)
                                                         (error _ _ _ elided _ _ _))
                                                         then-type))

- - - elided - - -                                       - - - elided - - -

(define compiler-language
  (extend-vocabulary
    base-language
    compiler))

```

■ **Figure 3** Micros and vocabularies

236 Fig. 3 collects illustrative excerpts from a vocabulary-micro code base. The left-hand
 237 column sets up a `compiler` vocabulary, which expresses transformations from the surface
 238 syntax into Racket’s core language. Among other micros, the `if` micro is added to `compiler`
 239 because it is a core construct. The final definition shows how to construct a complete
 240 language implementation by mixing in vocabularies into the common base language.

241 Like Scheme’s macro definitions, micro definitions use a pattern DSL for specifying inputs.
 242 As for the `Attribute ...` sequence, micros consume those via an explicit `lambda`. To
 243 create its output, the `if` micro allocates an AST node via `make-AST-if`. The pieces of this
 244 node are the results of expanding the three pieces that make up the original `if` expression.
 245 The expansions of these sub-expressions employ `dispatch`, a function that consumes the
 246 expression to be expanded together with the contextual vocabulary and the attributes (none
 247 here) in a staged fashion.

248 The right-hand side of fig. 3 shows how to add an `if` micro for a type-checking variant of
 249 the DSL. The code introduces a second vocabulary for the type checker. The `if` micro for
 250 this additional vocabulary implements the type checking rule for `if` in a standard manner,
 251 reporting an error as soon it is discovered.

252 Once the `type-check` vocabulary is in place, a developer can use it independently or
 253 in combination with the `compiler` vocabulary. For example, Racket’s soft typing sys-
 254 tem [23] needed a distinct interpretation for the language’s `letrec` construct, i.e., a distinct
 255 `type-check` vocabulary unrelated to the actual compiler. A variant of Typed Racket [42]
 256 could be implemented via the composition of these two vocabularies; in this context, the
 257 composition would discard the result of the pass based on the `type-check` vocabulary.

258 In general, DSL creators get two advantages from vocabularies and micros. First, they
 259 can now specify the syntax of their languages via explicit collections of micros. Each micro
 260 denotes a new production in the language’s expression language, and the input patterns
 261 describe its shape. Second, they can naturally express and implement static checking. The
 262 micro’s secondary arguments represent “inherited” attribute, and the flexible `Output` type
 263 allows the propagation of “synthesized” ones.

264 Implementing complete DSLs from vocabularies becomes similar to playing with Legos:
 265 (1) vocabularies are like mixins [30], (2) languages resemble classes, and (3) `dispatch` is
 266 basically a method invocation. Hence creating a variety of similar DSLs is often a game of
 267 linking a number of pieces from a box of language-building blocks. For the team’s rapid
 268 production and modification of teaching languages in the mid 1990s, vocabularies were a
 269 critical first step.

270 4.3 Languages for Semantic Modules

271 According to sec. 3 the implementation of any language combines a compiler with a run-time
 272 system. This dictum also applies to DSLs, whether realized with macros or micros. Both
 273 translate source code to target code, which refers to run-time values (functions, objects,
 274 constants, and so on). Such run-time values often collaborate “via conspiracy,” meaning their
 275 uses in target code satisfies logical statements—invariants that would not hold if all code
 276 had free access to these values. That is, the implementor of a DSL will almost certainly wish
 277 to hide these run-time values and even some of the auxiliary compile-time transformations.
 278 All of this suggests that macros, micros and vocabularies should go into a module, and such
 279 modules should make up a DSL implementation.

280 Conversely, the implementors of DSLs do not think of deploying individual constructs
 281 but complete languages. Indeed, conventional language implementors imagine that DSL
 282 programmers create self-contained programs. By contrast, Lispers think of their language
 283 extensions and imagine that DSL programmers may wish to escape into the underlying host
 284 language or even integrate constructs from different DSL-like extensions at the expression
 285 level. The question is whether a macro-micro based approach can move away from the “whole
 286 program” thinking of ordinary DSLs and realize a Lisp-ish approach of deploying languages
 287 for small units of code.

288 Krishnamurthi’s dissertation [37] presents answers to these two questions and thus
 289 introduces the first full framework for a macro-oriented approach to language-oriented
 290 programming. It combines macros with the first-class modules of the 1990s Racket, dubbed
 291 units [28], where the latter becomes both the container for DSL implementations as well as
 292 the one for DSL deployment. Technically, these units have the shape of fig. 4. They are
 293 parameterized over a `Language` and link-time imports, and they export values in response.


```
(unit/lang Language
  (ImportIdentifier ...)
  (ExportIdentifier ...)
  Definitions-and-Expressions ...)
```

■ **Figure 4** Language-parameterized, first-class units

294 A DSL implementation is also just a `unit/lang` that combines macros, micros, and
 295 run-time values. It is not recognized as a valid `Language` until it is registered with a *language*
 296 *administrator*. The latter compiles `unit/lang` expressions separately to plain `units`. For
 297 this compilation, the language administrator expands all uses of macros and micros and then
 298 resolves all the names in the generated code—without exposing any of them to other code.
 299 In particular, the developer does not need to take any action, such as adding the names of
 300 run-time values to export specifications of `Languages` or to `unit/langs` that use a `Language`.
 301 The result of a compilation is a collection of plain Racket `units`, and the Racket compiler
 302 turns this collection into a running program.

303 In principle, Krishnamurthi’s `unit/lang` system addresses all four DSL criteria listed
 304 in sec. 3. The micro-vocabulary combination can enforce syntax constraints beyond what
 305 macros can do. They are designed to express static processing in several passes and explicitly
 306 accommodate target languages distinct from Racket. And, the implementations of DSLs as
 307 `unit/langs` encapsulates the compiler component with a run-time component.

308 What this system fails to satisfy is the desire to synthesize DSL implementation techniques
 309 with Lisp’s incremental language-extension approach. The main problem is that a programmer
 310 has to parameterize an entire unit over a complete language. It is impossible to selectively
 311 import individual macros and micros from a `unit/lang`, which is what Racket developers
 312 truly want from a modular macro system. After a few years of using plain `units`, the Racket
 313 team also realized that first-class units provided more expressive power than they usually
 314 needed, meaning the extra complexity of programming the linking process rarely ever paid
 315 off in the code base.

316 Additionally, the `unit/lang` system was a step too far on the social side. Racket—then
 317 called PLT Scheme—was firmly in the Scheme camp and, at the time, the Scheme community
 318 had developed and open-sourced a new syntax system [14] that quickly gained in popularity.
 319 This macro system combined the declarative form of Krishnamurthi’s macros with the proce-
 320 dural form of his micros into a single `define-syntax-cases` form. Furthermore, this new
 321 macro system came with the same kind of syntax-object representation as Krishnamurthi’s,
 322 allowing source tracking, hygienic expansion, and other cross-expansion communication. In
 323 other words, the new system seemed to come with all the positive aspects of Krishnamurthi’s
 324 without its downsides. Hence, the Racket team decided to adapt this macro system and
 325 create a module system around it.

326 5 Languages From Syntactic Modules

327 The Racket designers started this rebuilding effort in 2000. The goal was to create a module
 328 system where a developer could write down each module in a DSL that fit the problem domain
 329 and where a module could export/import individual macros to/from other modules—and
 330 this second point forced them to reconsider the first-class nature of modules.

331 Flatt’s “you want it when” module-macro system [24] realizes this goal. It introduces a
 332 module form, which at first glance looks like `unit/lang`. Like the latter, `module` explicitly
 333 specifies the language of a module body, as the grammar in fig. 5 shows. Otherwise the

0:10 From Macros to DSLs: The Evolution of Racket

<pre>(module Name Language { ProvideSpeification RequireSpecification Definition Expression }*)</pre>	<pre>#lang Language { ProvideSpeification RequireSpecification Definition Expression }*</pre>	<pre>Name.rkt</pre>
--	---	---------------------

■ **Figure 5** Language-parameterized, first-order modules and their modern abbreviation

334 grammar appears to introduce a new expression form whose internals consist of a sequence
335 of exports, imports, definitions and expressions. A small difference concerns the organization
336 of the module body. The import and export specifications no longer need to show up as
337 the first element of the module; they can appear anywhere in the module. Appearances
338 are deceiving, however, and the `Name` part suggests the key difference. A module is *not* an
339 expression but a first-order form, known to the expander.

340 When the expander encounters `module`, it imports the `Language`'s provided identifiers.
341 This step establishes the base syntax and semantics of the module's expressions, definitions,
342 imports, and exports. Next the expander finds the imported and locally-defined macros in
343 the module body. The search for imported macros calls for the expansion and compilation of
344 the referenced modules. It is this step that requires the restriction to first-order modules,
345 because the expander must be able to identify the sources of imported macros and retrieve
346 their full meaning. Finally, the expander adds those imported and local macros to the
347 language syntax and then expands the module body properly, delivering an abstract-syntax
348 representation in the Racket core language.

349 One consequence of this arrangement is that the expansion of one module may demand
350 the evaluation of an entire tower of modules. The first module may import and use a macro
351 from a second module, whose definition relies on code that also uses language extensions.
352 Hence, this second module must be compiled after expanding and compiling the module that
353 supplies these auxiliary macros.

<pre>#lang racket (loop.rkt) (provide inf-loop) (define-syntax-cases [(inf-loop e) (begin (displayln "generating inf-loop") #'(do-it (lambda () e)))]]) (define (do-it th) (th) (do-it th))</pre>	<pre>#lang racket (use-loop.rkt) (provide display-infinitely-often) (require "loop.rkt") (define (display-infinitely-often x) (inf-loop (do-it x))) (define (do-it x) (displayln x))</pre>
---	--

■ **Figure 6** Exporting macros from, and importing them into, modules

354 The right-hand side of fig. 5 also shows the modern, alternative syntax for modules.
355 The first line of code specifies only the language of the module via a `#lang` specification;
356 the name of the file (boxed) determines the name of the module. Fig. 6 illustrates how
357 two modules interact at the syntax and run-time level. The module on the left defines the
358 language extension `inf-loop`, whose implementation generates code with a reference to the
359 function `do-it`. The module on the right imports this language extension via the `require`
360 specification. The Racket compiler retrieves the macro during compile time and uses it to

361 expand the body of the `display-infinitely-often` function—including a reference to the
 362 `do-it` function in module `loop.rkt`. Cross-module hygienic expansion [24, 26] ensures that
 363 this macro-introduced name does not conflict in any way with the `do-it` function definition
 364 of the `use-loop.rkt` module. Conceptually, the expansion of `display-infinitely-often`
 365 looks like the following definition:

```
366 (define (display-infinitely-often x)
367   (loop.rkt-do-it (lambda () (use-loop.rkt-do-it x))))
```

368 with the two distinct, fully-resolved names guaranteeing the proper functioning of the code
 369 according to the intuitive surface meaning.

370 Flatt's module-macro system allows the use of both declarative and procedural language
 371 extensions. To illustrate the latter kind, the `inf-loop` macro uses `define-syntax-cases`.
 372 If the expander can match a piece of syntax against one of the left-hand-side patterns
 373 of `define-syntax-cases`, it evaluates the expression on the right. This evaluation must
 374 produce code, which is often accomplished via the use of templates (introduced by `#'`) whose
 375 pattern variables are automatically replaced by matching pieces of syntax. But, as the
 376 definition of `inf-loop` suggests, the right-hand side may contain side-effecting expressions
 377 such as `displayln`. While these expressions do not become a part of the generated code as
 378 the above snippet shows, their side effects are observable during compile time.

379 To enable separate compilation, Racket discards the effects of the expansion phase before
 380 it moves on to running a module. Discarding such effects reflects the Racket designers' under-
 381 standing that language-extensions are like compilers, which do not have to be implemented
 382 in the same language as the one that they compile and which are not run in the same phase
 383 as the program that they translate. Phase separation greatly facilitates reasoning about
 384 compilation, avoiding a lot of the pitfalls of Lisp's and Scheme's module systems with their
 385 explicit `eval-when` and `compile-when` instructions [24].

<pre>#lang racket (provide Ack) ;; Number Number -> Number (define (Ack x y) (cond [(zero? x) (+ y 1)] [(and (> x 0) (zero? y)) (Ack (- x 1) 1)] [else (Ack (- x 1) (Ack x (- y 1)))]))</pre>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">math.rkt</div>	<pre>#lang racket (require (for-syntax "math.rkt")) (define-syntax-cases () [(static-Ack x y) ;; rewrites the pattern to a template ;; via some procedural processing (let* ((x-e (syntax-e #'x)) (y-e (syntax-e #'y))) (unless (and (number? x-e) (number? y-e)) (raise-syntax-error #f "not numbers")) (define ack (Ack x-e y-e)) #'(printf "the Ack # is -a" #,ack)))] (static-Ack 1 2)</pre>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">use-acker.rkt</div>
--	---	---	--

■ **Figure 7** Importing at a different phase

386 Phase separation imposes some cost on developers, however. If a module needs run-
 387 time functions for the definition of a language construct, the import specification must
 388 explicitly request a phase shift. For an example, see fig. 7. The module on the right defines
 389 `static-Ack`, which computes the Ackermann function of two numbers at compile time. Since
 390 the Ackermann function belongs into a different library module, say `math`, the `use-ack`
 391 module must import it from there. But, because this function must be used at compile time,
 392 the `require` specification uses the (underlined) `for-syntax` annotation to shift the import to
 393 this early phase. The Racket designers' experience shows that phase-shifting annotations are

394 still significantly easier to work with than Lisp’s and Scheme’s `expand-when` and `eval-when`
 395 annotations.

396 Like Krishnamurthi’s `unit/langs`, Flatt’s `modules` allow developers to write different
 397 components in different languages. In the case of `modules`, the `Language` position points
 398 to a module itself. The exports of this `Language` module determine the initial syntax and
 399 semantics of a client module.

400 In contrast to an ordinary `module`, a `Language` module must export certain macros, called
 401 interposition points; it may export others. An interposition point is a keyword that the
 402 macro expander adds to some forms during its traversal of the source tree. Here are the two
 403 most important ones for `Language` modules:

- 404 ■ `#:module-begin` is the (invisible) keyword that introduces the sequence of definitions
 405 and expressions in a module body. A `Language` module must export this form.
- 406 ■ `#:top-interaction` enables the read-eval-print loop for a `Language`, i.e., dynamic loading
 407 of files and interactive evaluation of expressions.

408 Other interposition points control different aspects of a `Language`’s meaning:

- 409 ■ `#:app` is inserted into function applications. In source code, an application has the shape
 410 `(fun arg ...)`, which expands to the intermediate form `(#:app fun arg ...)`.
- 411 ■ `#:datum` is wrapped around every literal constant.
- 412 ■ `#:top` is used to annotate module-level variable occurrences.

413 In practice, a developer creates a `Language` by adding features to a base language,
 414 subtracting others (by not exporting them), and re-interpreting some. Here “features” covers
 415 both macros and run-time values. The `#:module-begin` macro is commonly re-interpreted
 416 for a couple of reasons. Its re-definition often helps with the elimination of boilerplate
 417 code but also the communication of context-sensitive information from one source-level
 418 S-expression (including `modules`) to another during expansion.

<pre>#lang racket</pre>	<pre>lazy.rkt</pre>	<pre>#lang "lazy.rkt"</pre>	<pre>no-error.rkt</pre>
<pre>(provide (except-out (all-from-out racket) #:app) (rename-out [lazy-app #:app])) (define-syntax-rule (lazy-app f a ...) (#:app f (lambda () a) ...))</pre>	<pre>; a constant function (define (f x y) 10) ; called on two erroneous terms (f (/ 1 0) (first '())) ; evaluates to 10</pre>	<pre>; a constant function (define (f x y) 10) ; called on two erroneous terms (f (/ 1 0) (first '())) ; evaluates to 10</pre>	<pre>; a constant function (define (f x y) 10) ; called on two erroneous terms (f (/ 1 0) (first '())) ; evaluates to 10</pre>

■ **Figure 8** Building an embedded DSL from modules and macros

419 Fig. 8 indicates how a developer could quickly build a language that looks like Racket but
 420 uses call-by-name instead of call-by-value. The module on the left is the language implemen-
 421 tation. It starts from Racket and re-exports all of its features, including `#:module-begin`,
 422 except for function application. The module re-interprets function application via the sec-
 423 ond part of `provide`. Technically, a re-interpretation consists of a macro definition that is
 424 re-named in a `provide`. The `lazy` module comes with a `lazy-app` macro, which rewrites
 425 `(lazy-app fun arg ...)` to `(#:app fun (lambda () arg) ...)`. By static scope, the
 426 `#:app` in the expansion refers to the function application form of Racket. Since this macro
 427 is provided under the name `#:app`, a client module’s function applications—into which the
 428 expander inserts `#:app`—eventually expand according to `lazy-app`. In particular, the two
 429 exception-raising expressions in the `no-error` module are wrapped in `lambda`; because `f` is a
 430 constant function that does not evaluate its arguments, these errors are never reported. (For
 431 additional details on `lazy`, see the last chapter of *Realm of Racket* [15].)

```
#lang racket all-in-one.rkt

(module lazy-impl racket

  (provide
    (except-out (all-from-out racket) #%app)
    (rename-out [lazy-app #%app])))

(define-syntax-rule
  (lazy-app f a ...)
  (#%app f (lambda () a) ...))

(module lazy-client (submod "." lazy-impl)

  (define (f x y)
    10)

  (f (/ 1 0) (first '())))

(require (submod "." lazy-client))
```

(a) DSL development in one module

```
DrRacketCS  File  Edit
felleisen-algol60.alg.rkt (define ...)
1: Untitled
1 | #lang algol60
2 |
3 | begin
4 |   real procedure INVERSE(x);
5 |     real x;
6 |     begin
7 |       INVERSE := 1/x;
8 |     end;
9 |
10 |   println(INVERSE(5));
11 | end
```

(b) Algol 60 as a Racket DSL

■ **Figure 9** Developing and deploying DSLs in Racket

432 Modules and macros jointly make DSL development an interactive activity in the Racket
 433 ecosystem. A programmer can open two tabs or windows in an IDE to use one for the
 434 DSL implementation and another for a DSL program. Or, a programmer can place a DSL-
 435 implementing submodule [25] and a DSL-using submodule into a single file, which can then
 436 be edited and executed within a single editor window of the preferred IDE. Fig. 9a shows
 437 how to combine the modules of fig. 8 into a single file. This program consists of three pieces.
 438 The first one is a submodule that implements the lazy language, while the second uses the
 439 first one in the **Language** position. Hence the first submodule is the programming language
 440 of the second. The last piece of the program requires and thus evaluates the client module.
 441 Any change to the first submodule is immediately visible in the second.

442 A developer may also equip a DSL with any desired syntax, not just build on top of
 443 Racket's beautiful parentheses. To support this kind of syntax, a **Language** module may
 444 export a new reader. Recall from sec. 2 that a Lisp reader turns the stream of characters
 445 into a sequence of S-expressions (or **Syntax-Objects**, in the case of Racket). The rest of the
 446 implementation can then use the usual mix of macros and functions. Butterick's *Beautiful*
 447 *Racket* [4] is a comprehensive introduction to this strategy and comes with a powerful library
 448 package for lexing and parsing.

449 In the context of modular macros, a developer may also create a conventional compiler
 450 with the macro infrastructure. Instead of just expanding to Racket, a DSL implementation
 451 may use a combination of macros and compile-time functions to perform conventional type
 452 checking or other context-sensitive checks.

453 Fig. 9b presents a simple example of a Racket DSL program in conventional syntax.
 454 Except for the first line, the code represents a standard Algol 60 program. The first line
 455 turns this program into a Racket DSL and thus allows Racket to parse, type check, compile,
 456 and run this program. Because the DSL implementation turns the Algol 60 program into
 457 syntax objects and implements its semantics via macro expansion, DrRacket (the Racket
 458 IDE [22]) automatically adapts itself to this new language. For example, fig. 9b illustrates
 459 how DrRacket connects the binding occurrence of **INVERSE**'s parameter to its bound ones.

In sum, Racket’s modules simultaneously allow the incremental and interactive construction of language extensions as well as the construction of complete DSLs with their own vocabulary. The key design decision is to turn macros into entities that first-order modules can export, import, hide, and re-interpret. It does necessitate the introduction of strict phase separation between the expansion phase and run-time phase to obtain separate compilation.

6 Syntax Done Properly With Parsing Macros

The implementation of a DSL’s syntax consists of two essential parts: parsing syntactically legitimate sentences, and reporting violations of the syntactic rules. Both aspects are equally important, but for 40 years, the macro community mostly neglected the second one.

Sec. 2 lists five problems with parsing via Lisp-style macros. Kohlbecker’s rewriting DSL—based on patterns and templates—eliminates all of them except for problem 4. To appreciate the complexity of this particular problem, consider the actual grammatical production of `let` expressions in classical BNF notation:

```
(let ({[id expression]}*) expression+)
```

Kohlbecker’s pattern-based meta-DSL addresses this context-free shape specification with the elegant trick of using ellipses (...) for * and unrolling for +:

```
(let ([id expression] ...) expression expression ...)
```

What Kohlbecker’s notation cannot express is the side condition of fig. 2:

```
id ... is a sequence of distinct identifiers
```

Indeed, Kohlbecker’s notation cannot even specify that `id` must stand for an identifier.

So now imagine a programmer who writes

```
(let (((+ 1 2) x)) (* x 3)) ;; => ((lambda ((+ 1 2)) (* x 3)) x)
```

OR

```
(let ((x 1) (x 2)) (* x 3)) ;; => ((lambda (x x) (* x 3)) 1 2)
```

In either case, a pattern-oriented language generates the `lambda` expression to the right of the `=>` arrow. Hence, the resulting syntax errors speak of `lambda` and parameters, concepts that the grammatical description of `let` never mentions. While a reader might be tempted to dismiss this particular error message as “obvious,” it is imperative to keep in mind that this `let` expression might have been generated by the use of some other macro, which in turn might be the result of some macro-defining macro, and so on.

Dybvig’s `define-syntax-cases` slightly improves on Kohlbecker’s DSL. It allows the attachment of *fenders*—Boolean expressions—to a macro’s input patterns. With such fenders, a macro developer can manually formulate conditions that check such side conditions. Even in such simple cases as `let`, however, the error-checking code is many times the size of the rewriting specification. And this is why most macro authors fail to add this code or, if they do, fail to write comprehensive checks that also generates good error messages.

Culpepper’s DSL for defining macros solves this problem with two innovations [8,9,10,11]. The first is an augmentation of the pattern-matching DSL with “words” for articulating classification constraints such as those of the `let` macro. The second is a DSL for specifying new classifications. Together, these innovations allow programmers to easily enforce assumptions about the surface syntax, synthesize error messages in terms of the specification, and deconstruct the inputs of a macro into relevant pieces.

```

(define-syntax-class distinct-bindings
  #:description "sequence of distinct binding pairs"
  (pattern (b:binding ...))
  #:fail-when (check-duplicate-id #'(b.lhs ...))
    "duplicate variable name"
  #:with (lhs* ...) #'(b.lhs ...)
  #:with (rhs* ...) #'(b.rhs ...))

(define-syntax-class binding
  #:description "binding pair"
  (pattern (lhs:id rhs:expr)))

```

■ **Figure 10** Syntax classifications

502 Following our discussion above, the specification of `let` needs two syntax classifications:
 503 one to say that the second part of `let`'s input is a sequence and another one to say that the
 504 elements of this sequence are identifier-expression pairs. Fig. 10 shows how a programmer
 505 can define these classifications in Culpepper's meta-DSL. A classification must come with at
 506 least one `pattern` clause, which spells out the context-free shape of the form and names its
 507 pieces. For example, the `binding` class uses the pre-defined classifications `id` (for identifier)
 508 and `expr` (for expression) to say that a binding has the shape `(id expr)` and that the name
 509 of the `id` is `lhs` and the name of `expr` is `rhs`. Any use of such a syntax class, for example the
 510 one in the definition of `distinct-bindings`, may refer to these attributes of the input via a
 511 dot notation. Thus, `b.lhs ...` in `distinct-bindings` denotes the sequence of identifiers.
 512 As this example also shows, a syntax-class definition may also defer to procedural code, such
 513 as `check-duplicate-id` to process the input. A `fail-when` clause allows macro developers
 514 to specify a part of the synthesized error message (when the default is not sufficiently clear).

```

(define-syntax-parser let
  [(_ bs:distinct-bindings body:expr ...+)
   ;; rewrites the pattern to a template
   #'((lambda (bs.lhs* ...) body ...) bs.rhs* ...)])

```

■ **Figure 11** Macros via parsing macros

515 Using these two syntax classes, specifying the complete shape of `let` is straightforward;
 516 see fig. 11. The `:distinct-bindings` classification of `bs` introduces names for two pieces
 517 of the input syntax: a sequence of identifiers (`bs.lhs*`) and a sequence of right-hand-side
 518 expressions (`bs.rhs*`), one per variable. The syntax template uses these pieces to generate
 519 the same target code as the macros in fig. 2.

520 A comparison of figs. 2 and 11 illustrates the advantages as well as the disadvantages of
 521 Culpepper's DSL for writing macros. On the positive side, the size of the Culpepper-style
 522 macro definition appears to remain the same as the one for the Kohlbecker-style one. The
 523 revised definition merely adds classifications to the macro's pattern and attribute selections
 524 to the macro's template. This shallow size comparison camouflages that these small changes
 525 cause the macro to check *all* constraints on the shape of `let` and formulate syntax errors in
 526 terms of the specified surface syntax. As Culpepper [9, page 469] explains, implementing the
 527 same level of assumption checking and error reporting via procedural macros increases the
 528 code size by "several factors." Furthermore the "primary benefit [of this meta-DSL] ... is
 529 increased clarity" of a macro's input specification and its code template.

530 On the negative side, macro programmers are now expected to develop syntax classi-
 531 fications such as those in fig. 10 and use them properly in macro definitions, as in fig. 11.
 532 While the development of syntax classifications clearly poses a new obstacle, their use comes
 533 with a significant payoff *and* most end up as reusable elements in libraries. Hence the cost
 534 of developing them is recouped through reuse. As for the use of syntax classifications in

535 macro templates and patterns, experience shows that most macro programmers consider the
 536 annotations as type-like notions and the attribute selections as a natural by-product.

537 In short, Culpepper’s meta-DSL completely replaces the `define-syntax-cases` meta-
 538 DSL for macro definitions. By now, the large majority of Racket programmers develop macros
 539 in Culpepper’s DSL and contribute to the ever-expanding collection of syntax classifications.

	lang. extens. (sec. 3)	lexical scope (2)	classify syntax (1)	error messages (1)	separate compil. (4)	run-time encaps. (4)	code gen. opt. (3)
Lisp macros	✓	–	–	–	–	–	–
Scheme							
– syntax-rules	✓	✓	patterns	–	–	–	–
– syntax-case	✓	✓	patterns & fenders	–	–	–	–
Racket	✓	✓	patterns & syn. classes	✓	✓& phases	✓	module only
– syntax-parse							

– means programmers have the tools to design manual solutions

■ **Figure 12** A concise overview of Lisp-family language extension features

540 7 DSL Creators Need Still More Than Modular, Parsing Macros

541 Racket has made great progress in improving the state of the art of macros with an eye
 542 toward both language extension and DSL implementation. Fig. 12 surveys the progress
 543 in roughly the terms of sec. 3’s criteria. The `syntax-parse` DSL for defining macros can
 544 express almost every context-free and -sensitive constraint; macro developers get away with
 545 a few hints and yet get code that reports syntax errors in terms of the macro-defined variant.
 546 The `module` system supports both the fine-grained export/import of macros for language
 547 extensions and the whole-cloth implementation of DSLs.

548 At this point, implementing DSLs is well within reach for Racket beginners [4, 15] and
 549 easy for experts. While beginners may focus on `module`-based DSLs, experts use macros
 550 to create fluidly embedded DSLs. Examples are the DSL of pattern-matching for run-time
 551 values, the `syntax-parse` DSL itself, and Redex [16, 39]. In this domain, however, the macro
 552 framework falls short of satisfying the full list of desiderata for from sec. 3.

553 To explain this gap, let us concisely classify DSLs and characterize Racket’s support:

554 ■ *stand-alone DSLs*

555 These are the most recognized form in the real world. Racket supports those via `module`
 556 languages with at least the same conveniences as other DSL construction frameworks.

557 ■ *embedded DSLs with a fixed interface*

558 All programming languages come with numerous such sub-languages. For example,
 559 `printf` interprets the format DSL—usually written as an embedded string—for rendering
 560 some number of values for an output device. In Racket, such DSLs instead come as a
 561 new set of expression forms with which programmers compose domain-specific programs.
 562 Even in Racket, though, such DSLs allow only restricted interactions with the host.

563 ■ *embedded and extensible DSLs with an expression-level interface*

564 Racket’s DSLs for pattern matching and structure declarations are illuminating examples
 565 of this kind. The former allows programmers to articulate complex patterns, with
 566 embedded Racket expressions. The latter may contain patterns, which contain expressions,
 567 etc. The pattern DSL is extensible so that, for example, the DSL of structure definitions


```
(define-typed-syntax (if cond then else)
  [- cond >> cond-ir ==> cond-type]
  [- then >> then-ir ==> then-type]
  [- else >> else-ir <=< else-type]
  -----
  [(AST-if cond-ir then-ir else-ir) -> then-type])
```

■ **Figure 13** Type-checking from macros

568 can automatically generate patterns for matching structure instances. Naturally, this
 569 DSL for structure declarations can also embed Racket expressions at a fine granularity.

570 With regard to the third kind of DSL, Racket’s macro approach suffers from several
 571 problems. A comparison with the criteria in sec. 3 suggests three obvious ones.

572 The first concerns DSLs that demand new syntactic categories and, in turn, good parsing
 573 and error reporting techniques. While syntax classes allow DSL creators to enumerate
 574 the elements of a new syntactic category, this enumeration is fixed. Experienced DSL
 575 implementors can work around this restriction, just like programmers can create extensible
 576 visitor patterns in object-oriented languages to allow the blind-box extension of data types.

577 The second problem is about context-sensitive language processing. The existing macro
 578 framework makes it difficult to implement context-sensitive static checking, translations, and
 579 optimizing transformations—even for just Racket’s macros, not to mention those that define
 580 new syntactic categories. Chang and his students [5, 6] have begun to push the boundaries
 581 in the realm of type checking, a particular rich form of context-sensitivity. Specifically,
 582 the team has encoded the rich domain of type checking as a meta-DSL. In essence, this
 583 meta-DSL enables DSL creators to formulate type checking in the form of type elaboration
 584 rules from the literature (see fig. 13), instead of the procedural approach of fig. 3. However,
 585 their innovation exploits brittle protocols to make macros work together [27]. As a result,
 586 it is difficult to extend their framework or adapt it to other domains without using design
 587 patterns for macro programming.

588 Finally, the DSL framework fails to accommodate languages whose compilation target is
 589 not Racket. Consider an embedded DSL for Cuda programming that benefits from a fluid
 590 integration with Racket. Such a DSL may need two interpretations: on computers with
 591 graphical co-processors it should compile to GPU code, while on a computer without such a
 592 processor it may need to denote a plain Racket expression. Implementing a dependent-type
 593 system in the spirit of Chang et al.’s work supplies a second concrete example. The language
 594 of types does not have the semantics of Racket’s expressions and definitions. Although it is
 595 possible to expand such DSLs through Racket’s core, it forces DSL developers to employ
 596 numerous macro-design patterns.

597 The proposed work-arounds for these three problems reveal why the Racket team does
 598 not consider the problem solved. Racket is all about helping programmers avoid syntactic
 599 design patterns. Hence, the appearance of design patterns at the macro level is antithetical to
 600 the Racket way of doing things, and the Racket team will continue to look for improvements.

601 **Acknowledgements** The authors thank Michael Ballantyne, Eli Barzilay, Stephen Chang,
 602 Robby Findler, Alex Knauth, Alexis King, and Sam Tobin-Hochstadt for contributing at
 603 various stages to the evolution of Racket’s macro system and how it supports LOP. They
 604 also gratefully acknowledge the suggestions of the anonymous SNAPL ’19 reviewers, Sam
 605 Caldwell, Ben Greenman for improving the presentation.

606 — References

- 607 1 H. Abelson, R.K. Dybvig, C.T. Haynes, G.J. Rozas, N.I. Adams, D.P. Friedman, E. Kohlbecker,
608 G.L. Steele, D.H. Bartley, R. Halstead, D. Oxley, G.J. Sussman, G. Brooks, C. Hanson, K.M.
609 Pitman, and M. Wand. Revised5 report on the algorithmic language scheme. *Higher-Order
610 and Symbolic Computation*, 11(1):7–105, Aug 1998.
- 611 2 Alan Bawden and Jonathan Rees. Syntactic closures. In *Symposium on Lisp and Functional
612 Programming*, pages 86–95, 1988.
- 613 3 Matthias Blume. Refining hygienic macros for modules and separate compilation. Technical
614 report tr-h-171, ATR Human Information Processing Research Laboratories, Kyoto, Japan,
615 1995. people.cs.uchicago.edu/~blume/papers/hygmac.pdf.
- 616 4 Matthew Butterick. *Beautiful Racket*. 2013. URL: <https://beautifulracket.com/>.
- 617 5 Stephen Chang, Alex Knauth, and Ben Greenman. Type systems as macros. In *Symposium
618 on Principles of Programming Languages*, pages 694–705, 2017.
- 619 6 Stephen Chang, Alex Knauth, and Emina Torlak. Symbolic types for lenient symbolic execution.
620 In *Symposium on Principles of Programming Languages*, pages 40:1–40:29, 2018.
- 621 7 William Clinger and Jonathan Rees. Macros that work. In *Symposium on Principles of
622 Programming Languages*, pages 155–162, 1991.
- 623 8 Ryan Culpepper. *Refining Syntactic Sugar: Tools for Supporting Macro Development*. PhD
624 thesis, Northeastern University, 2010.
- 625 9 Ryan Culpepper. Fortifying macros. *Journal of Functional Programming*, 22(4–5):439–476,
626 2012.
- 627 10 Ryan Culpepper and Matthias Felleisen. Taming macros. In *Generative Programming and
628 Component Engineering*, pages 225–243, 2004.
- 629 11 Ryan Culpepper and Matthias Felleisen. Fortifying macros. In *International Conference on
630 Functional Programming*, pages 235–246, 2010.
- 631 12 Sergey Dmitriev. Language-oriented programming: the next programming paradigm, 2004.
- 632 13 R. Kent Dybvig, Daniel P. Friedman, and Christopher T. Haynes. Expansion-passing style: A
633 general macro mechanism. *Lisp and Symbolic Computation*, 1(1):53–75, January 1988.
- 634 14 R. Kent Dybvig, Robert Hieb, and Carl Bruggeman. Syntactic abstraction in Scheme. *Lisp
635 and Symbolic Computation*, 5(4):295–326, December 1993.
- 636 15 Matthias Felleisen, Forrest Bice, Rose DeMaio, Spencer Florence, Feng-Yun Mimi Lin, Scott
637 Lindeman, Nicole Nussbaum, Eric Peterson, Ryan Plessner, David Van Horn, and Conrad
638 Barski. *Realm of Racket*. No Starch Press, 2013. URL: <http://www.realmofracket.com/>.
- 639 16 Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. *Semantics Engineering with
640 PLT Redex*. MIT Press, 2009.
- 641 17 Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. *How
642 to Design Programs. Second Edition*. MIT Press, 2001–2018. URL: <http://www.htdp.org/>.
- 643 18 Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. The
644 structure and interpretation of the computer science curriculum. *Journal of Functional
645 Programming*, 14(4):365–378, 2004.
- 646 19 Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, Shriram Krishnamurthi, Eli Barzilay,
647 Jay McCarthy, and Sam Tobin-Hochstadt. The Racket Manifesto. In *First Summit on Advances
648 in Programming Languages*, pages 113–128, 2015.
- 649 20 Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, Shriram Krishnamurthi, Eli Barzi-
650 lay, Jay McCarthy, and Sam Tobin-Hochstadt. A programmable programming language.
651 *Communications of the ACM*, 61(3):62–71, February 2018.
- 652 21 R. Findler and M. Felleisen. Contracts for higher-order functions. In *International Conference
653 on Functional Programming*, pages 48–59, 2002.
- 654 22 Robert Bruce Findler, John Clements, Cormac Flanagan, Matthew Flatt, Shriram Krishna-
655 murthi, Paul Steckler, and Matthias Felleisen. DrScheme: A programming environment for
656 Scheme. *Journal of Functional Programming*, 12(2):159–182, 2002.

- 657 23 Cormac Flanagan, Matthew Flatt, Shriram Krishnamurthi, Stephanie Weirich, and Matthias
658 Felleisen. Catching bugs in the web of program invariants. In *Conference on Programming
659 Language Design and Implementation*, pages 23–32, 1996.
- 660 24 Matthew Flatt. Composable and compilable macros: You want it *when*? In *International
661 Conference on Functional Programming*, pages 72–83, 2002.
- 662 25 Matthew Flatt. Submodules in Racket: you want it when, again? In *Generative Programming
663 and Component Engineering*, pages 13–22, 2013.
- 664 26 Matthew Flatt. Binding as sets of scopes. In *Symposium on Principles of Programming
665 Languages*, pages 705–717, 2016.
- 666 27 Matthew Flatt, Ryan Culpepper, David Darais, and Robert Bruce Findler. Macros that
667 work together: Compile-time bindings, partial expansion, and definition contexts. *Journal of
668 Functional Programming*, 22(2):181–216, March 2012.
- 669 28 Matthew Flatt and Matthias Felleisen. Cool modules for HOT languages. In *Conference on
670 Programming Language Design and Implementation*, pages 236–248, 1998.
- 671 29 Matthew Flatt, Robert Bruce Findler, Shriram Krishnamurthi, and Matthias Felleisen. Pro-
672 gramming languages as operating systems (or, Revenge of the Son of the Lisp Machine). In
673 *International Conference on Functional Programming*, pages 138–147, September 1999.
- 674 30 Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. Classes and mixins. In
675 *Symposium on Principles of Programming Languages*, pages 171–183, 1998.
- 676 31 Timothy P. Hart. MACROS for LISP. Technical Report 57, MIT Artificial Intelligence
677 Laboratory, 1963.
- 678 32 Paul Hudak. Modular domain specific languages and tools. In *Fifth International Conference
679 on Software Reuse*, pages 134–142, 1998.
- 680 33 Richard Kelsey and Jonathan Rees. A tractable Scheme implementation. *Lisp and Symbolic
681 Computation*, 5(4):315–335, 1994.
- 682 34 Eugene E. Kohlbecker. *Syntactic Extensions in the Programming Language Lisp*. PhD thesis,
683 Indiana University, 1986.
- 684 35 Eugene E. Kohlbecker, Daniel P. Friedman, Matthias Felleisen, and Bruce F. Duba. Hygienic
685 macro expansion. In *Symposium on Lisp and Functional Programming*, pages 151–161, 1986.
- 686 36 Eugene E. Kohlbecker and Mitchell Wand. Macros-by-example: Deriving syntactic transfor-
687 mations from their specifications. In *Symposium on Principles of Programming Languages*,
688 pages 77–84, 1987.
- 689 37 Shriram Krishnamurthi. *Linguistic Reuse*. PhD thesis, Rice University, 2001.
- 690 38 Shriram Krishnamurthi, Matthias Felleisen, and Bruce F. Duba. From macros to reusable
691 generative programming. In *International Symposium on Generative and Component-Based
692 Software Engineering*, pages 105–120, 1999.
- 693 39 Jacob Matthews, Robert Bruce Findler, Matthew Flatt, and Matthias Felleisen. A visual
694 environment for developing context-sensitive term rewriting systems. In *Rewriting Techniques
695 and Applications*, pages 2–16, 2004.
- 696 40 Daniel Patterson and Amal Ahmed. Linking types for multi-language software: Have your
697 cake and eat it too. In *Summit on Advances in Programming Languages*, pages 12:1–12:15,
698 2017.
- 699 41 Gerald L. Sussman and Guy Lewis Steele Jr. Scheme: An interpreter for extended lambda
700 calculus. Technical Report 349, MIT Artificial Intelligence Laboratory, 1975.
- 701 42 Sam Tobin-Hochstadt and Matthias Felleisen. The design and implementation of Typed
702 Scheme. In *Symposium on Principles of Programming Languages*, pages 395–406, 2008.
- 703 43 Martin P. Ward. Language oriented programming. *Software Concepts and Tools*, 15:147–161,
704 April 1994.