

STITCH: The Sound Type-Indexed Type Checker

RICHARD A. EISENBERG, Bryn Mawr College, USA

A classic example of the power of generalized algebraic datatypes (GADTs) to verify a delicate implementation is the type-indexed expression AST. This tutorial paper refreshes this example, casting it in modern Haskell using many of GHC's bells and whistles. The Stitch interpreter is a full executable interpreter, with a parser, type checker, common-subexpression elimination, and a REPL. Making heavy use of GADTs and type indices, the Stitch implementation is clean, idiomatic Haskell and serves as an existence proof that Haskell's type system is advanced enough for the use of fancy types in a practical setting. The paper focuses on guiding the reader through these advanced topics, enabling them to adopt the techniques demonstrated here.

1 A SIREN FROM THE FOLKLORE

A major focus of modern functional programming research is to push the boundaries of type systems. The fancy types born of this effort allow programmers not only to specify the shape of their data—types have done *that* for decades—but also the meaning and correctness conditions of their data. In other words, while well typed programs do not go wrong, fancy typed programs always go right. By leveraging a type system to finely specify the format of their data, programmers can hook into the declarative specifications inherent in type systems to be able to reason about their programs in a compositional and familiar manner.

Though fancy types come in a great many varieties, this work focuses on an entry-level fancy type, the generalized algebraic data type, or GADT. GADTs, originally called first-class phantom types [Cheney and Hinze 2003] or guarded recursive datatypes [Xi et al. 2003], exhibit one of the most basic ways to use fancy types. When you pattern-match on a GADT value, you learn information about the type of that value. Accordingly, different branches of a GADT pattern match have access to different typing information and can make effective use of that information. In this way, a term-level, runtime operation (the pattern-match) informs the type-level, compile-time type-checking—one of the hallmarks of dependently typed programming. Indeed, GADTs, in concert with other features, can be used to effectively mimic dependent types, even without full-spectrum support [Eisenberg and Weirich 2012; Monnier and Haguenaier 2010].

It is high time for an example of what we are talking about:¹

```
data G :: Type → Type where
```

```
  BoolCon :: G Bool
```

```
  IntCon  :: G Int
```

```
match :: ∀a. G a → a
```

```
match BoolCon = True
```

```
match IntCon  = 42
```

The GADT G has two constructors. One constrains G 's index (of kind $Type$, a recent notation change from the older \star [Zavialov 2018]) to be $Bool$, the other Int . The `match` function does a GADT pattern-match on a value of type $G a$. If the value is `BoolCon`, then we learn that a is in fact

¹All the examples in this paper are type-checked in GHC during the typesetting process, with gratitude to lhs2TeX [Löh 2012].

Author's address: Richard A. Eisenberg, Computer Science Department, Bryn Mawr College, 101 N. Merion Ave, Bryn Mawr, PA, 19010, USA, rae@cs.brynmawr.edu.

2018. XXXX-XXXX/2018/10-ART \$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

50 *Bool*; our function can thus return *True :: a*. In the other branch, the value of type *G a* is *IntCon*,
 51 and thus *a* must be *Int*; we can return *42 :: Int*. The runtime pattern-match tells us the compile-time
 52 type, allowing the branches to have *different* types. In contrast, a simple pattern-match always
 53 requires every branch to have the *same* type.

54 1.1 Stitch

55
 56 This paper presents the design and implementation of Stitch, a simple extension of the simply
 57 typed λ -calculus (STLC), including integers, Booleans, basic arithmetic, conditionals, a fixpoint
 58 operator, and let-bindings. (I use “Stitch” to refer both to the language and its implementation.)
 59 The expression abstract syntax tree (AST) type in Stitch is a GADT such that only well typed Stitch
 60 expressions can be formed. That is, there is simply no representation for the expression `true 5`, as
 61 that expression is ill typed. The AST type, *Exp*, is *indexed* by the type of the expression represented,
 62 so that if *exp :: Exp ctx ty*, then the Stitch expression encoded in *exp* has the type *ty*. (Here, *ctx* is
 63 the typing context for any free variables in the expression.)

64 The example of a λ -calculus implementation using a GADT in this way is common in the folklore,
 65 and it has been explored in previous published work (see Section 10.5). However, the goal of this
 66 current work is not to present a type-indexed AST as a novel invention, but instead to methodically
 67 explore the usage of one. It is my hope that, through this example, readers can gain an appreciation
 68 for the power and versatility of fancy types and learn some techniques for how they can apply this
 69 technology in their own projects.

70 It can be easy to dismiss the example of well typed λ -calculus terms as too introspective: Can’t PL
 71 researchers come up with a better example to tout their wares than a PL implementation? However,
 72 I wish to turn this argument on its head. A PL implementation is a fantastic example, as most
 73 programmers in a functional language will quickly grasp the goal of the example, allowing them
 74 to focus on the implementation aspects instead of trying to understand the program’s behavior.
 75 Furthermore, implementing a language is a practical example. Many significant systems require PL
 76 implementations, including web browsers, database servers, editors, spreadsheets, shells, and even
 77 many games.

78 This paper will focus on the version of Haskell implemented in GHC 8.6 (the Glasgow Haskell
 79 Compiler), making critical use of GHC’s recent support for using GADT constructors at the
 80 type level [Weirich et al. 2013; Yorgey et al. 2012], type reflection (i.e. *Typeable*) [Peyton Jones
 81 et al. 2016], higher-rank type inference [Peyton Jones et al. 2007], and, of course, GADT type
 82 inference [Peyton Jones et al. 2006; Vytiniotis et al. 2011]. Accordingly, this paper can serve as an
 83 extended example of how recent innovations in GHC can power a more richly typed programming
 84 style.

85 1.2 Contributions

86
 87 While this tutorial paper does not offer new *technical* contributions, it illuminates recent innovations
 88 in Haskell—a language of importance within the PL community and gaining traction in industry—
 89 and invites intermediate programmers to use advanced PL techniques in their programs. It makes
 90 the following contributions:

- 91 • Stitch is a full executable interpreter of the STLC, available online² and suitable for classroom
 92 use.
- 93 • Section 3 is an accessible primer on Haskell’s advanced features, as used in the examples in
 94 this paper.

95
 96
 97 ²<http://cs.brynmawr.edu/~rae/papers/2018/stitch/stitch.tar.gz>

- This work offers many settings for the use of fancy types. For example, parser output is guaranteed to be well-scoped.
- Section 9 describes aspects of the common-subexpression elimination pass implemented in Stitch offered, as proof that the use of an indexed AST scales to the more complex analyses inherent in real compilers.
- The development described here serves as an existence proof that Haskell—even without full dependent types—is a suitable language in which to use practical fancy types.

2 INTRODUCING STITCH

2.1 The Simply Typed λ -Calculus

Stitch is an implementation of the simply typed λ -calculus, so we will start off with a review of this little language, including the Stitch extensions. See Figure 1.³

We see that Stitch is quite a standard implementation of the STLC with modest extensions. It has a call-by-value semantics, and the value of a let-bound variable is computed before entering the body of the let. Stitch supports general recursion by way of its (standard) fix operator, which evaluates to a fixpoint. All λ -abstractions are annotated with the type of the argument.

Stitch comes with both a small-step and big-step operational semantics, though the small-step semantics is elided here. Users of Stitch may find it interesting to compare its behavior with respect to the two presentations of semantics; commands at the Stitch REPL allow the user to choose how they wish to reduce an expression to a value, allowing users to witness that big-step semantics tell you nothing about a term during evaluation, while the small-step semantics can show you the steps the expression takes on the way to becoming a value.

2.2 The Stitch REPL

Before we jump into the implementation, it is helpful to look at the user’s experience of Stitch. The Stitch REPL allows the user to enter in expressions for evaluation, to bind new global variables, and to query aspects of an expression. An example is worth at least several hundred words here:

```

127 Welcome to the Stitch interpreter, version 1.0.
128  $\lambda$ > 1 + 1
129 2 : Int
130  $\lambda$ >  $\backslash x:\text{Int} \rightarrow \text{Int}. \backslash y:\text{Int}. x y$ 
131  $\lambda\#:\text{Int} \rightarrow \text{Int}. \lambda\#:\text{Int}. \#1 \#0 : (\text{Int} \rightarrow \text{Int}) \rightarrow \text{Int} \rightarrow \text{Int}$ 
132  $\lambda$ > expr = ( $\backslash x:\text{Int} \rightarrow \text{Int}. \backslash y:\text{Int}. x y$ ) ( $\backslash z:\text{Int}. z + 3$ ) 5
133 expr = ( $\lambda\#:\text{Int} \rightarrow \text{Int}. \lambda\#:\text{Int}. \#1 \#0$ ) ( $\lambda\#:\text{Int}. \#0 + 3$ ) 5 : Int
134  $\lambda$ > expr
135 8 : Int
136  $\lambda$ > :step expr
137 ( $\lambda\#:\text{Int} \rightarrow \text{Int}. \lambda\#:\text{Int}. \#1 \#0$ ) ( $\lambda\#:\text{Int}. \#0 + 3$ ) 5 : Int
138 --> ( $\lambda\#:\text{Int}. (\lambda\#:\text{Int}. \#0 + 3) \#0$ ) 5 : Int
139 --> ( $\lambda\#:\text{Int}. \#0 + 3$ ) 5 : Int
140 --> 5 + 3 : Int
141 --> 8 : Int

```

We see here that the syntax is straightforward and familiar, though Stitch requires a type annotation at every λ -abstraction. The REPL allows the user to create new global variables, like

³The formalization is type-checked and typeset with the help of ott [Sewell et al. 2010].

148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196

Metavariables:

x

term variables

Grammar:

$\tau ::= \tau_1 \rightarrow \tau_2 \mid \mathbf{Int} \mid \mathbf{Bool}$

types

$op ::= + \mid - \mid * \mid / \mid \% \mid < \mid \leq \mid > \mid \geq \mid \equiv$

operators

$\mathbb{Z} ::= \dots$

integers

$\mathbb{B} ::= \mathbf{true} \mid \mathbf{false}$

Booleans

$e ::= x \mid \lambda x:\tau.e \mid e_1 e_2 \mid \mathbf{let} x = e_1 \mathbf{in} e_2 \mid e_1 op e_2$
 $\quad \mid \mathbf{if} e_1 \mathbf{then} e_2 \mathbf{else} e_3 \mid \mathbf{fix} e \mid \mathbb{Z} \mid \mathbb{B}$

expressions

$v ::= \lambda x:\tau.e \mid \mathbb{Z} \mid \mathbb{B}$

values

$\Gamma ::= \emptyset \mid \Gamma, x:\tau$

typing contexts

$s ::= e \mid x = e$

statements

Other notation:

$\mathbf{result}(op)$ is the result type of an operator: **Int** for $\{+, -, *, /, \%\}$ and **Bool** for $\{<, \leq, >, \geq, \equiv\}$

$\mathbf{apply}(op, v_1, v_2)$ computes the result of using op with operands v_1 and v_2

$e_1[e_2/x]$ denotes capture-avoiding substitution of e_2 for x in e_1

$\boxed{\Gamma \vdash e : \tau}$ Typing rules

$$\frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau} \text{T_VAR} \quad \frac{\Gamma, x:\tau_1 \vdash e : \tau_2}{\Gamma \vdash \lambda x:\tau_1.e : \tau_1 \rightarrow \tau_2} \text{T_LAM} \quad \frac{\Gamma \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash e_2 : \tau_1}{\Gamma \vdash e_1 e_2 : \tau_2} \text{T_APP}$$

$$\frac{\Gamma \vdash e_1 : \tau_1 \quad \Gamma, x:\tau_1 \vdash e_2 : \tau_2}{\Gamma \vdash \mathbf{let} x = e_1 \mathbf{in} e_2 : \tau_2} \text{T_LET} \quad \frac{\Gamma \vdash e_1 : \mathbf{Int} \quad \Gamma \vdash e_2 : \mathbf{Int}}{\Gamma \vdash e_1 op e_2 : \mathbf{result}(op)} \text{T_ARITH}$$

$$\frac{\Gamma \vdash e_1 : \mathbf{Bool} \quad \Gamma \vdash e_2 : \tau \quad \Gamma \vdash e_3 : \tau}{\Gamma \vdash \mathbf{if} e_1 \mathbf{then} e_2 \mathbf{else} e_3 : \tau} \text{T_COND} \quad \frac{\Gamma \vdash e : \tau \rightarrow \tau}{\Gamma \vdash \mathbf{fix} e : \tau} \text{T_FIX}$$

$$\frac{}{\Gamma \vdash \mathbb{Z} : \mathbf{Int}} \text{T_INT} \quad \frac{}{\Gamma \vdash \mathbb{B} : \mathbf{Bool}} \text{T_BOOL}$$

$\boxed{e \Downarrow v}$ Big-step operational semantics

$$\frac{}{v \Downarrow v} \text{E_VALUE} \quad \frac{e_1 \Downarrow \lambda x:\tau.e \quad e_2 \Downarrow v_2}{e[v_2/x] \Downarrow v} \text{E_APP} \quad \frac{e_1 \Downarrow v_1 \quad e_2[v_1/x] \Downarrow v}{\mathbf{let} x = e_1 \mathbf{in} e_2 \Downarrow v} \text{E_LET}$$

$$\frac{e_1 \Downarrow v_1 \quad e_2 \Downarrow v_2}{e_1 op e_2 \Downarrow \mathbf{apply}(op, v_1, v_2)} \text{E_ARITH} \quad \frac{e \Downarrow \lambda x:\tau.e' \quad e'[\mathbf{fix}(\lambda x:\tau.e')/x] \Downarrow v}{\mathbf{fix} e \Downarrow v} \text{E_FIX}$$

$$\frac{e_1 \Downarrow \mathbf{true} \quad e_2 \Downarrow v}{\mathbf{if} e_1 \mathbf{then} e_2 \mathbf{else} e_3 \Downarrow v} \text{E_IFTRUE} \quad \frac{e_1 \Downarrow \mathbf{false} \quad e_3 \Downarrow v}{\mathbf{if} e_1 \mathbf{then} e_2 \mathbf{else} e_3 \Downarrow v} \text{E_IFFALSE}$$

Fig. 1. The simply typed λ -calculus, as embodied in Stitch.

expr. These are unevaluated. Storing *unevaluated* global variables may seem an unusual feature in a call-by-value language. This design choice is motivated by the desire, as shown above, to facilitate demonstrations of single-step reduction. There are no technical concerns around this choice; it would be just as easy to evaluate first and store globals as fully evaluated. The syntax $\mathbf{expr} = \dots$

is called a *statement*, as included in Figure 1. We can then input the global by itself or as part of a larger expression to evaluate it. However, the most distinctive aspect of this session is Stitch’s approach to variable binding, which we explore next.

2.3 De Bruijn indices

Every implementor of a programming language must make a choice of representation of variable binding. The key challenge is that, no matter which representation we choose, we must be sure that $\lambda x:\tau.x$ and $\lambda y:\tau.y$ are treated identically in all contexts. There are *many* possible choices out there: named binders [Pitts 2003], locally nameless binders [Gordon 1994], using higher-order abstract syntax [Pfenning and Elliott 1988], *parametric* higher-order abstract syntax [Chlipala 2008], UNBOUND [Weirich et al. 2011], bound [Kmett 2012], among others. The interested reader is referred to Weirich et al. [2011], where even more possibilities lie in wait. In this work, however, I choose trusty, old de Bruijn indices [de Bruijn 1972], as these serve two design goals of Stitch well: de Bruijn indices work easily with an indexed AST, and they can easily arise when teaching implementations of the λ -calculus [e.g., Pierce 2002, Chapter 6].

A de Bruijn index is a number used in the place of a variable name; it counts the number of binders that intervene between a variable occurrence and its binding site. We see above that the expression $\lambda x:\text{Int} \rightarrow \text{Int} . \lambda y:\text{Int} x y$ desugars to $\lambda \#:\text{Int} \rightarrow \text{Int} . \lambda \#:\text{Int} . \#1 \#0$, where the $\#1$ refers to the outer binder (1 intervening binding site) and the $\#0$ refers to the inner binder (0 intervening binding sites). De Bruijn indices have the enviable property of making α -equivalence utterly trivial: because variables no longer have names, we do not need to worry about renaming. However, they make other aspects of implementation harder. Specifically, two challenges come to the fore:

- (1) De Bruijn indices are hard for programmers to understand and work with.
- (2) As an expression moves into a new context, the indices may have to be shifted (increased or decreased) in order to preserve their identity, as the number of intervening binding sites might have changed. It is very easy for an implementor to make a mistake when doing these shifts.

As a partial remedy to the first problem, Stitch color-codes its output (as can be seen in this typeset document). A variable occurrence and its binding site are assigned the same color, so that a reader no longer has to count binding sites. Though only a modest innovation, this color-coding has proved to be wildly successful in practice; not only has it been helpful in my own debugging, but working functional programmers who see it have gasped, “I finally understand de Bruijn indices now!” more than once. Note that programmers never have to *write* using de Bruijn indices (the parser converts their names to indices quite handily) and so this simple reading aid goes a long way toward fixing the first drawback.

The second drawback can be more troublesome. The reason we have such a plethora of approaches to variable binding must be, in part, that implementors have been unhappy with the approaches available—they thus invent a new one. One reason for this unhappiness is that capture-avoiding substitution is a real challenge. Pierce [2002, Section 5.3] gives an instructive account of the pitfalls an implementor encounters. And it is not just substitution. As a language grows in complexity, dealing with name clashes and renaming crops up in a variety of places. Indeed, the venerable GHC implementation only recently (January, 2016) added checks to make sure its handling of variable naming is bug-free; I count 29 call sites within the GHC source code (as of October, 2018) that still use the “unchecked” variant of substitution because using the checked version fails on certain test cases. Each of these call sites is perhaps a lurking bug, waiting for a pathological program to induce an unexpected name clash that could cause GHC to go wrong.

<pre> 246 Stitch source, prime.stitch: 247 noDivisorsAbove = 248 fix \nda: Int -> Int -> Bool. 249 \tester: Int. \scrutinee: Int. 250 if tester * tester > scrutinee 251 then true 252 else if scrutinee % tester == 0 253 then false 254 else nda (tester+1) scrutinee ; 255 256 isPrime = noDivisorsAbove 2 257 258 259 260 261 262 </pre>	<pre> After parsing and type checking: noDivisorsAbove = fix λ#:Int -> Int -> Bool. λ#:Int. λ#:Int. if #1 * #1 > #0 then true else if #0 % #1 == 0 then false else #2 (#1 + 1) #0 : Int -> Int -> Bool isPrime = fix ... 2 : Int -> Bool </pre>
---	--

Fig. 2. A primality checker in Stitch.

However, a solution to this conundrum is at hand: because Stitch's expression AST type is indexed by the type of the expression represented, an erroneous or forgotten shifting of a de Bruijn index leads to a straightforward error, caught as Stitch itself is being compiled. Indeed, I shudder to think about the challenge in getting all the shifts correct without the aid of an indexed AST. Thus, using an indexed AST fully remedies the second drawback.

One twist on the second drawback remains, however: all this shifting can slow the interpreter down. A variable shift requires a full traversal and rebuild of the AST, costing precious time and allocations. Though I have not done it in my implementation, it would be possible to add a *Shift* constructor to the AST type to allow these shifts to be lazily evaluated; the design and implementation of other opportunities for optimization are left as future work.

2.4 A slightly longer example: primality checking

As a final example of a user's interaction with Stitch, I present the program in Figure 2. It implements a primality checker in Stitch. The file `prime.stitch`, included in the Stitch tarball, can be loaded into the Stitch REPL with `:load prime.stitch`.

```

278
279 λ> :load prime.stitch
280 ...
281 λ> isPrime 7
282 true : Bool
283 λ> isPrime 9
284 false : Bool
285

```

In the right half of the figure, we see Stitch's parsed and type-checked representation of the original program. This AST cannot store global variables (all variables are de Bruijn indices), so Stitch inlines `noDivisorsAbove` in the definition of `isPrime`, above.

2.5 An overview of Stitch

Before we get mired in the details, let us review the overall architecture of the Stitch interpreter. Throughout the rest of this paper, I will refer to individual modules in the package; these references are intended to help a reader who wishes to follow along in the actual codebase. However, the text

<p>295 Modules that principally define datatypes:</p> <p>296</p> <p>297</p> <p>298</p> <p>299</p> <p>300</p> <p>301</p> <p>302</p> <p>303</p> <p>304</p> <p>305</p>	<ul style="list-style-type: none"> • Type: Stitch types, (§4) • Op: Binary operators (§6.3) • Token: Lexer tokens • Unchecked: The AST for parsed, but not type checked, expressions (§5) • Exp: Expressions AST (§6) • Globals: Global variables (§7.2) • Statement: Statements (§2.2) 	<p>Modules that principally define algorithms:</p> <ul style="list-style-type: none"> • Lex: Lexer • Parse: Parser (§5) • Check: Type checker (§7) • Shift: de Bruijn index shifting (§8.2) • Eval: Operational semantics (§8) • CSE: Common-subexpression elimination (§9) • Pretty: Pretty-printing • Repl: The user-facing REPL (§2.2)
---	--	---

Fig. 3. Principal modules in Stitch. All module names are prefixed with `Language.Stitch`.

of this paper is self-contained and does not require looking at the code. The map of modules is in Figure 3.

A Stitch program travels through the interpreter in the usual fashion. The REPL module defines an interactive prompt which reads a string from the user. This string is then lexed into a series of tokens and then parsed into an expression AST that is *not* checked for type safety (defined in the Unchecked module). This expression type is then run through the type checker to be transformed into the checked AST (defined in Exp). The checked AST is optionally optimized (by performing CSE) and then evaluated according to the semantics the user chooses. A pretty-printer [Wadler 2003] renders the result back to the user.

We are now almost ready to start seeing the fancy types, but first, we need to install some necessary infrastructure.

3 FANCY-TYPED UTILITIES

Every great edifice necessarily requires some plumbing. What is fun in this case is that even the plumbing needs some fancy types in order to support what comes ahead. The definitions in this section are standard, and readers familiar with dependently typed programming may wish to skim this section quickly or skip to the next section. The utilities described here are useful beyond just Stitch, and some have implementations released separately. However, I have included them within the Stitch package in order to keep it self-contained. These modules, too, are prefixed with `Language.Stitch`. so as not to pollute the module namespace. This section introduces Peano natural numbers (useful for tracking the number of bound variables), length-indexed vectors (useful for tracking the types of in-scope variables), existentials (useful for storing the values of global variables, perhaps of different types), and singletons (useful during type checking, when we must connect a type-level context with term-level type representations).

3.1 Natural numbers

The `Data.Nat` module defines routine Peano unary natural numbers:

```
data Nat = Zero | Succ Nat
```

This datatype is used in Stitch solely in types, via Haskell’s datatype promotion mechanism [Yorkey et al. 2012]. For the last several years, GHC has allowed programmers to use data constructors (`Zero` and `Succ` in this case) in types; correspondingly, `Nat` is not only a type classifying terms, but also a kind classifying types. Indeed, recent improvements in GHC have eliminated the distinction between types and kinds [Weirich et al. 2013], and I have come to view the usage of `Zero` and

344 *Succ* in types more as a namespace issue (Haskell maintains separate “type-level” and “term-level”
 345 namespaces) than as promotion, per se. We will soon see an example of these type-level constructors
 346 in action (§3.2). Because *Nat* is used solely in types, the inefficiency of storing a unary number does
 347 not bite at runtime, slowing down only the compilation process of the Stitch interpreter, not the
 348 compiled executable.

349 One might ask: Why use unary *Nats* instead of GHC’s built-in support for type-level natural
 350 numbers [Diatchki 2014]? Unary naturals have an inherent inductive structure, making for easy
 351 definitions and proofs. While GHC cannot know, say, that $n + m$ is the same as $m + n$, the type-level
 352 arithmetic used in Stitch is quite simple and no arithmetic reasoning is necessary. In my experience,
 353 these hand-written unary naturals work better than the built-in naturals for defining vectors.
 354

355 3.2 Length-indexed vectors

356 No exploration of fancy types would be complete without the staple of length-indexed vectors, a
 357 ubiquitous example because of their perspicuity and usefulness. A length-indexed vector is simply
 358 a linked list, where the list type includes the length of the list; thus, a list of length 2 is a distinct
 359 type from a list of length 3. Here is the type definition:
 360

```
361 data Vec :: Type → Nat → Type where
362   VNil :: Vec a Zero
363   (:>) :: a → Vec a n → Vec a (Succ n)
```

365 We will take this line-by-line. We see here that *Vec* is parameterized by an element type of kind
 366 *Type* and a length index of kind *Nat*. The declaration for *VNil* states that *VNil* is always a *Vec*
 367 of length *Zero*, but it can have any element type *a*. The cons operator *:>* takes an element (of type *a*),
 368 the tail of the vector (of type *Vec a n*) and produces a vector that is one longer than the tail (of type
 369 *Vec a (Succ n)*).

370 Note the use of *Nat* as a kind and *Zero* and *Succ* as types. When GHC is resolving names used in
 371 a type, it first looks in the type-level namespace, where definitions like *Vec* and *Nat* live. Failing
 372 that lookup (for capitalized identifiers), it looks in the term-level namespace; this is what happens
 373 in the case of *Zero* and *Succ*.⁴ Finding these constructors, GHC has no trouble using them in types,
 374 where they keep their usual meaning.
 375

376 3.2.1 *Appending*. We will need to append vectors, and the two vectors may be of different lengths.
 377 Clearly, the append function should take arguments of type *Vec a n* and *Vec a m*, where the element
 378 type *a* is the same but the length indices *n* and *m* are different. However, what should the result
 379 type of appending be? Of course, the length of the concatenation of two vectors is the sum of the
 380 lengths of the vectors: the result should be *Vec a (n + m)*. We thus need to define *+* on *Nats*. What
 381 is unusual here is that we need to use *+* in types, not in terms. GHC’s approach here is to use a *type*
 382 *family* [Chakravarty et al. 2005; Eisenberg et al. 2014], which is essentially a function that works
 383 on types and type-level data. Here is the definition:
 384

```
384 type family n + m where
385   Zero + m = m
386   Succ n + m = Succ (n + m)
```

388 We are now ready to define appending two vectors:
 389

391 ⁴If the identifier exists in both namespaces, it can be prefixed with `'` to tell GHC to look only in the term-level namespace.
 392


```

393 (++) :: Vec a n → Vec a m → Vec a (n + m)
394 VNil    ++ ys = ys
395 (x :> xs) ++ ys = x :> (xs ++ ys)
396

```

397 Already, the fancy types are working for us, making sure our code is correct. In the first clause
398 of `++`, we pattern-match on `VNil`. This match tells us both that the first vector is empty, and also
399 that the type variable `n` equals `Zero`. This second fact comes from the declared type of `VNil` in the
400 definition of `Vec`. All `VNils` have a type index of `Zero`, and thus we know that if `VNil :: Vec a n`, then
401 `n` must be `Zero`. The type checker uses this fact to accept the right-hand side of that equation: it
402 must be convinced that `ys :: Vec a (n + m)`, the declared return type of `++`. Because the type checker
403 knows that `n` is `Zero`, however, it can use the definition of the type family `+` to reduce `Zero + m`
404 to `m`, and then it simply uses the fact that `ys :: Vec a m`, as `ys` is the second argument to `++`. The
405 second equation is similar, except that it uses the second equation of `+` to check the equation's
406 right-hand side. If we forgot to cons `x` onto `xs ++ ys` in this right-hand side, the definition of `++`
407 would be rejected as ill typed.

408
409 **3.2.2 Indexing.** How should we look up a value in a vector? We could use an operator like Haskell's
410 standard `!!` operator that looks up a value in a list. However, this is unsatisfactory, because the
411 `!!` throws an exception when its index is out of range. Given that we know a vector's length at
412 compile-time, we can do better.

413 The key step is to have a type that represents natural numbers less than some known bound.
414 The type `Fin` (short for "finite set"), common in dependently typed programming and declared in
415 `Data.Fin`, does the job:

```

416 data Fin :: Nat → Type where
417   FZ :: Fin (Succ n)
418   FS :: Fin n → Fin (Succ n)
419

```

420 The `Fin` type is indexed by a natural number `n`. The type `Fin n` contains exactly `n` values, corre-
421 sponding to the numbers 0 through `n - 1`. This GADT tends to be a bit harder to understand than
422 `Vec` because (unlike `Vec`), you cannot tell the type of a `Fin` just from the value. For example, the
423 value `FS FZ` can have both type `Fin 2` and `Fin 10` (where I take liberty to use decimal notation
424 instead of unary notation for `Nats`), but not `Fin 1`. Let us understand this type better by tracing
425 how we can assign a type to `FS FZ`:

- 426 • Suppose we are checking to see whether `FS FZ :: Fin 1`. We see that `FS :: Fin n → Fin (Succ n)`.
427 Thus, for `FS FZ :: Fin 1`, we must instantiate `FS` to have type `Fin Zero → Fin (Succ Zero)`. We
428 must now check `FZ :: Fin Zero`. However, this fails, because `FZ :: Fin (Succ n)`—that is, `FZ`'s
429 type index must not be `Zero`. We accordingly reject `FS FZ :: Fin 1`.
- 430 • Now say we are checking `FS FZ :: Fin 5`. This proceeds as above, but in the end, we must
431 check `FZ :: Fin 4`. The number 4 is indeed the successor of another natural, and so `FZ :: Fin 4`
432 is accepted, and thus so is `FS FZ :: Fin 5`.

433
434 Following this logic, we can see how `Fin n` really has precisely `n` values.

435 As a type whose values range from 0 to `n - 1`, `Fin n` is the perfect index into a vector of length `n`:

```

436 (!!!) :: Vec a n → Fin n → a
437 vec !!! fin = case (fin, vec) of -- reverse order due to laziness
438   (FZ, x :> _) → x
439   (FS n, _ :> xs) → xs !!! n
440

```

441

442 GHC comes with a pattern-match completeness checker [Karachalias et al. 2015] that marks
 443 this `case` as complete, even without an error case. To understand why, we follow the types. After
 444 matching `fin` against either `FZ` or `FS n`, the type checker learns that `n` must not be zero—the types
 445 of both `FZ` and `FS` end with a `Succ` index. Since `n` is not zero, then it cannot be the case that `vec` is
 446 `VNil`. Even though the pattern match includes only `:>`, that is enough to be complete.

447 Now, we can explore this match reversal. Haskell is a lazy language [Peyton Jones 2003], which
 448 means that variables can be bound to diverging computations (denoted with \perp). When matching a
 449 compound pattern, Haskell matches the patterns left-to-right, meaning that the left-most scrutinee
 450 (`fin`, in our case) is evaluated to a value and then inspected before evaluating later scrutinees, such
 451 as `vec`. Imagine matching against `vec` first. In this case, it is conceivable that `vec` would be `VNil`
 452 while `fin` would be \perp . This is not just theoretical; witness the following function:

```
453 lazinessBites :: Vec a n → Fin n → String
454 lazinessBites VNil _ = "empty vector"
455 lazinessBites _ _ = "non-empty vector"
```

457 If we try to evaluate `lazinessBites VNil undefined`, that expression is accepted by the type checker
 458 and evaluates handily to "empty vector". If we scrutinize `vec` first, then, the completeness checker
 459 correctly tells us that we must handle the `VNil` case. On the other hand, in the implementation of
 460 `!!!` with the pattern match reversed, we ensure that `fin` is not \perp before ever looking at `vec` and can
 461 thus be sure that `vec` cannot be `VNil`.

462
463

3.3 Existentials

464 Suppose we want a ragged two-dimensional vector. We might be tempted to use `Vec (Vec a n) m`,
 465 but this type requires that all inner vectors have length `n`, going against our desire for a ragged
 466 collection. Of course, we could use lists, but we stick with `Vec` for the sake of example—we will not
 467 have the easy escape of lists when we encounter this problem later.

468 What we want is a way to hide the `n` index from the type of a vector; we want a collection of
 469 vectors where every vector has *some* length, but not necessarily the same one. This is what an
 470 *existential type* does: it essentially hides a type index, allowing us to recover it only through pattern
 471 matching. Here is the quintessential existential type, defined in `Data.Exists`:

```
472 data Ex :: (k → Type) → Type where
473   Ex :: a i → Ex a
474
```

475 The `Ex` type is parameterized over the indexed type constructor `a` of the data it holds; the index
 476 itself can be of any kind `k`. Thus, `a` has kind `k → Type`. The `Ex` data constructor takes one argument
 477 of type `a i` for any `i`—note that `i` is *not* mentioned in the return type `Ex a`. This makes `i` *existentially*
 478 *bound*.

479 We can understand this better through an example:

```
480 exVecSum :: Ex (Vec Int) → Int
481 exVecSum (Ex v) = go v
482   where go :: Vec Int n → Int
483         go VNil      = 0
484         go (x :> xs) = x + go xs
```

486 The pattern match in `exVecSum` unpacks the existential to reveal a vector `v`. Naturally, `v` has type
 487 `Vec` and stores `Ints`; but, what is `v`'s length index? It is impossible to know: there exists a length,
 488 but we do not know it. Essentially, the length index is stored by the `Ex` constructor along with `v`.
 489 When we pattern-match against the `Ex` constructor, we get both the index and the term. When we

490

491 call the *go* helper method, the type of that method is instantiated to the unknown (and unnamed)
 492 index and executes as expected.

493 Now that we have *Ex*, we can make our ragged two-dimensional vector type: *Vec (Ex (Vec a)) m*.
 494 We know a value of this type has *m* rows, but each row has a different (and unknown) length.

495

496

3.4 Singletons

497

498

499

500

501

The technique of *singletons* is a well worn and well studied [Monnier and Haguenaueur 2010] way to simulate dependent types in a non-dependent language. Though at least two libraries exist for automatically generating singletons in Haskell [Eisenberg and Weirich 2012; McBride 2011], Stitch does not depend on these libraries, in order to maintain some simplicity and be self-contained. However, the design of these libraries is the direct inspiration for the definitions in Stitch.

502

503

504

505

506

507

508

509

To motivate singletons, consider writing a version of *replicate* for vectors. The *replicate* function takes a natural number *n* and an element *elt* and creates a vector of length *n* consisting of *n* copies of *elt*. Despite this simple specification, there is no easy way to write a type signature for *replicate*; you might try *replicate :: Nat → a → Vec a ?*, but you'd be stuck at the *?*. The problem is that the choice of the *type* index for the return type must be the *value* of the first parameter. This is the hallmark of dependent types. However, because Haskell does not yet support dependent types, singletons will have to do. Here is the definition of a singleton *Nat* (or, more precisely the family of singleton *Nats*):

510

511

512

513

```
data SNat :: Nat → Type where
```

```
  SZero :: SNat Zero
```

```
  SSucc :: SNat n → SNat (Succ n)
```

514

515

516

517

518

519

The type *SNat* is indexed by a *Nat* that corresponds precisely to the value of the *SNat*. That is, the type of *SSucc (SSucc SZero)* is *SNat (Succ (Succ Zero))*. Conversely, the *only* value of the type *SNat (Succ (Succ Zero))* is *SSucc (SSucc SZero)*. This last fact is why singleton types are so named: a singleton type has precisely one value. Because of the correspondence between types and terms with singleton types, matching on the values of a singleton inform the type index—exactly what we need here.

520

521

522

523

524

Here is the definition for *replicate*:

```
replicate :: SNat n → a → Vec a n
```

```
replicate SZero _ = VNil
```

```
replicate (SSucc n') elt = elt :> replicate n' elt
```

525

526

527

528

529

The GADT pattern match against *SZero* tells the type checker that *n* is *Zero* in the first equation, making *VNil* an appropriate result. Similarly, the match tells the type checker that *n* is *Succ n'* (for some *n'*) in the second equation, and thus a vector one longer than *n'* is an appropriate result. Essentially, the *n* in the type signature for *replicate* is the value of the first parameter, exactly as desired.

530

531

532

533

534

535

536

Because a singleton value is uniquely determined by its type, it is convenient to be able to pass singletons implicitly. We can take advantage of Haskell's type class mechanism to do this, via the following type class and instances:

```
class          SNat1 (n :: Nat) where snat :: SNat n
```

```
instance      SNat1 Zero   where snat = SZero
```

```
instance SNat1 n ⇒ SNat1 (Succ n) where snat = SSucc snat
```

537

538

539

Any function with a *SNat1 n* constraint can gain access to the singleton for *n* simply by calling the *snat* method.

```

540 type Ty = Ex (TypeRep :: Type → Type)
541 pattern Ty t = Ex t
542 {-# complete Ty #-}
543
544 -- Decompose a function type
545 pattern (:=) :: () ⇒ (fun ~ (arg → res)) ⇒ TypeRep arg → TypeRep res → TypeRep fun
546 isTypeRep :: ∀ a b. Typeable a ⇒ TypeRep b → Maybe (a ≈: b)
547 isTypeRep = eqTypeRep (typeRep @a)
548
549
550
551
552

```

Fig. 4. Some definitions supporting Stitch types.

The `Data.Singletons` module contains several more definitions in order to support polymorphic singletons. A full treatment of these definitions would take us too far afield, and the approach roughly mimics that taken by Eisenberg and Weirich [2012]. In this text, I avoid using these definitions; readers following along in the actual implementation may notice a few insignificant differences in the use of singletons, but these are inessential for our topics of interest.

Singletons are not the final word for dependent types in Haskell. They can be unwieldy [Lindley and McBride 2013] and conversions between singleton types and unrefined base types (such as converting from `SNat n` to `Nat`) are potentially costly at runtime. Work is under way [Eisenberg 2016; Gundry 2013; Weirich et al. 2017] to add full dependent types to Haskell. However, for our present purposes, the singletons work quite nicely, and their drawbacks do not get in our way.

4 A STITCH TYPE IS A HASKELL TYPE

An early choice in designing an interpreter for a typed language is how one will represent types. The Stitch language’s type system is very simple, as portrayed in Figure 1: it contains `Ints`, `Bools`, and functions among these. Conveniently, the Haskell type system also contains these types, and GHC’s type reflection mechanism [Peyton Jones et al. 2016] allows a programmer access to type representations.

A key aspect of GHC’s reflection mechanism is that it provides a *type-indexed* type representation, `TypeRep`. The type `TypeRep` has kind $\forall k. k \rightarrow \text{Type}$, allowing for a representation of a type of any kind. The representation for `Int` has type `TypeRep Int`; the representation for `Bool` has type `TypeRep Bool`. As such, `TypeRep` is actually the singleton type for the kind `Type`.⁵ GHC also provides a number of facilities for inspecting and building type representations, exported through its `Type.Reflection` module. By using `TypeRep` to represent Stitch types, we hook into the existing mechanism for efficient comparison of types, generation of hashes (used in Section 9), and singleton support. An excerpt of Stitch’s `Type` module appears in Figure 4.

Along with re-exporting `Type` itself, the module defines `Ty`, a type synonym for an existential package (Section 3.3) containing a `TypeRep`. The `Ty` type is used when we wish to refer to a type without doing any compile-time reasoning—for example, in the unchecked, parsed expression AST (Section 5). In order to make usage of `Ty` easier throughout Stitch, a pattern synonym [Pickering et al. 2016] is introduced. This pattern synonym, also named `Ty` (but in the term-level namespace), comes with a `{-# complete Ty #-}` pragma; this compiler directive instructs GHC that the `Ty`

⁵`TypeRep` can be viewed as a universal singleton type, because it works at all kinds. However, working with `TypeReps` for non-`Type` singletons is even more unwieldy than singletons usually are, and so I use `TypeRep` only at kind `Type → Type` and write custom singleton types for other singletons.

589 pattern, all by itself, is a complete pattern match against the `Ty` type. This pragma silences pattern-
 590 match completeness warnings, which do not yet work with pattern synonyms without the user’s
 591 help.

592 4.1 Decomposing functions

594 Next, we see the definition of the `:->` pattern synonym, which allows for decomposition of function
 595 types. For example, if we want to check whether `fun :: TypeRep ty` is a function type, we could say

```
596 case fun of arg :-> res -> ...
597           _other    -> ...
```

599 A careful reader will note the unusual type assigned to the pattern `:->`, with *two* constraints
 600 offset by \Rightarrow . (The first is empty, `()`.) While a full explanation of pattern synonym types would be
 601 a digression—and Pickering et al. [2016, Section 6] gives an accessible introduction with many
 602 examples—suffice it to say that this type indicates that a successful pattern match tells you that the
 603 scrutinee’s type index (denoted with `fun` in the type signature) will be refined to `arg -> res` in the
 604 body of the match. This is exactly what we will need in the type checker.

605 4.2 Comparing `TypeReps` using propositional equality

607 Following `:->` is `isTypeRep`, a convenient way to check whether a `TypeRep` matches a desired type.
 608 For example, this is used in the type checker when checking to see that the condition in an `if`
 609 is indeed of type `Bool`. If we are checking `rep :: TypeRep b`, then we would query `isTypeRep @Bool rep`.
 610 The `@Bool` argument is a *visible type application* [Eisenberg et al. 2016], which allows a caller of
 611 `isTypeRep` to choose the instantiation for the type variable `a`. Note that the signature for `isTypeRep`
 612 lists `a` first, meaning that the first usage of a visible type application would instantiate `a`. The body
 613 of `isTypeRep` also uses visible type application to extract an explicit `TypeRep` from the implicit
 614 `Typeable`, where we have `typeRep :: Typeable a => TypeRep a`.

615 Curiouser still is the return type of `isTypeRep`, `Maybe (a :: b)`. The type `≈:` is exported from
 616 GHC’s `Data.Type.Equality` and has this definition:

```
617 data (a :: k1) ≈: (b :: k2) where HRefl :: a ≈: a
```

618 The type `≈:` is *heterogeneous propositional equality*. It is heterogeneous because the two types
 619 related might not have the same kind.⁶ It is propositional because we must match against a value
 620 in `a ≈: b` (that is, `HRefl`) to convince the type checker that `a` is, in fact, the same as `b`. If `a :: k1`
 621 and `b :: k2`, then matching something of type `a ≈: b` against `HRefl` convinces the type checker that `a`
 622 equals `b` and `k1` equals `k2` through the usual behavior of GADT pattern-matching.

624 This is the appropriate return type provided by GHC’s `eqTypeRep :: TypeRep a -> TypeRep b ->`
 625 `Maybe (a ≈: b)`, and therefore Stitch’s `isTypeRep`. The `eqTypeRep` function is used to compare two
 626 type representations. If they are in fact equal, then it is often necessary to reflect this equality back
 627 to the type checker. Here is an example:

```
628 castTo :: ∀a b. Typeable a => a -> TypeRep b -> Maybe b
629 castTo x repB = case isTypeRep @a repB of
630   Just HRefl -> Just x
631   Nothing    -> Nothing
```

633 The idea here is that we have a value `x` of type `a`, but we wish for it to have some other type `b`. We
 634 also have the type representations of both; `a` is implicit (`Typeable`) while `b` is explicit (`TypeRep`). If

635 ⁶In the use of `TypeReps` in this paper, we have no need for heterogeneity; a homogeneous equality would do. However, as a
 636 general facility for dynamic type-checking, the `TypeRep` library exports `isTypeRep` with a heterogeneous return value.

```

638 -- Unchecked expression, indexed by the number of variables in scope
639 data UExp (n :: Nat) = UVar (Fin n) -- de Bruijn index for a variable
640 | UGlobal String
641 | ULam Ty (UExp (Succ n))
642 | UApp (UExp n) (UExp n)
643 | UArith (UExp n) UArithOp (UExp n)
644 | UIntE Int
645
646 ...
647 -- An encoding of (\x: Int. x + 1) 5, as an example
648 uexample :: UExp Zero -- Zero because the expression is closed
649 uexample = UApp (ULam (Ty (typeRep @Int)) (UArith (UVar FZ) (UArithOp Plus) (UIntE 1)))
650 (UIntE 5)
651
652
653
654
655
656

```

Fig. 5. The AST for parsed expressions, from the Unchecked module.

the type representations are equal—that is, if we can discover at runtime that both a and b are, in fact, the same—then we can return x at type b . In the *Just* case, we match against *HRefl*, a proof that a equals b . This then allows GHC to accept *Just x* as having the return type of *Maybe b*. Without the match against *HRefl*, *Just x :: Maybe b* would be rejected.

The *eqTypeRep* function must use *heterogeneous* equality (instead of the homogeneous version \sim ; which is otherwise similar) because *TypeRep* is polykinded: we might be comparing types of different kinds. Not only do we need to know the types equal, but we need to know the kinds equal as well. This heterogeneous equality is available in GHC only since version 8.0, powered by recent advances in the theory [Weirich et al. 2013].

5 SCOPE-CHECKED PARSING

Though *Stitch*'s hallmark is its indexed AST for expressions, we cannot parse into that AST directly. Type-checking can produce better error messages and is more easily engineered independent from the left-to-right nature of parsing. We thus must define an unchecked (un-indexed) AST for the result of parsing the user's program.

However, even here there is a role for fancy types. While type-checking during parsing is a challenge, name resolution during parsing works nicely. We can thus parse into an AST that can express only well-scoped terms. The AST type definition appears in Figure 5.

The type *UExp* (“unchecked expression”) is indexed by a *Nat* that denotes the number of local variables in scope in the expression. So, a *UExp 0* is a closed expression, while a *UExp 2* denotes an expression with up to two free variables. Note that *ULam* increments this index for the body of the λ -abstraction.

Variables are naturally stored in a *Fin n*—precisely the right type to store de Bruijn indices. If an expression has only 2 variables in scope, then we must make sure that a variable has an index of either 0 or 1, never more. Using *Fin* gives us this guarantee nicely.

You will see in the definition of *UExp* a few other small details:

- Occurrences of global variables are stored as strings. These will then be interpreted during type-checking to inline the stored value of the global.

- Lambda-abstractions store a $T\gamma$ —the existential wrapper around $TypeRep$ —to denote the argument type of the function. Note that there is no explicit place in the AST for the bound variable, as the bound variable always has a de Bruijn index of 0.
- The $UArith$ constructor stores a $UArithOp$, which is an existential wrapper around the indexed $ArithOp$ type, explored in more depth in Section 6.3.

The main novelty in working with $UExp$ is, of course, the $Fin\ n$ type for de Bruijn indices. Supporting this design requires accommodations in the parser. Stitch’s parser is a monadic parser built on the Parsec library [Leijen 2001]. Its input is the series of tokens, each annotated with location information, produced by the entirely unremarkable lexer (also built using Parsec). It can parse either statements or expressions.

The most interesting aspect of the parser is that the parser type must be indexed by number of in-scope variables—this is what will set the index of any parsed Fin de Bruijn indices. We thus have this definition for the parser monad:

```
type Parser n a = ParsecT [LToken] () (Reader (Vec String n)) a
```

The $ParsecT$ monad transformer [Jones 1995] is indexed by (1) the type of the input stream, which in our case is $[LToken]$; (2) the state carried by the monad, which in our case is trivial; (3) an underlying monad, which in our case is $Reader\ (Vec\ String\ n)$, where the environment is a vector of the names of the in-scope variables; and (4) the return type of computations, a . Thus, a computation of type $Parser\ n\ a$ parses a list of located tokens into something of type a in an environment with access to the names of n in-scope local variables.

5.1 A heterogeneous reader monad

The only small difficulty in working with $Parser$, as defined above, is around variable binding (naturally). Here is the relevant combinator:

```
bind :: String → Parser (Succ n) a → Parser n a
bind bound_var thing_inside
  = hlocal (bound_var :>) thing_inside
```

Given a bound variable name, $bind$ parses some type a in an extended environment (with $Succ\ n$ bound variables) and then returns the result in an environment with only n bound variables. Note that $bind$ does *not* do any kind of shifting or type-change of the result: if the inner parser is of type, say, $Parser\ (Succ\ n)\ (Fin\ (Succ\ n))$, then the outer result will have type $Parser\ n\ (Fin\ (Succ\ n))$. Note that the index to the Fin does not change.

The $bind$ function is implemented using a new combinator $hlocal$, inspired by the $local$ method of the $MonadReader$ class from the mtl (monad transformer library). The relevant part of this class is

```
class Monad m ⇒ MonadReader r m | m → r where
  local :: (r → r) → m a → m a
  ...
```

The $local$ method allows a computation to assume a local value of the environment for some smaller computation. This is exactly what we want here. The only problem is that the type of the local environment is *different* than the type of the outer environment: the outer environment has type $Vec\ String\ n$ while the local one has type $Vec\ String\ (Succ\ n)$.

We must accordingly define a heterogeneous reader monad, which allows a type change for the local environment. Here is the class definition:

```

736 class Monad m => MonadHReader r1 m | m -> r1 where
737   type SetEnv r2 m :: Type -> Type
738   hlocal :: (r1 -> r2) -> (Monad (SetEnv r2 m) => SetEnv r2 m a) -> m a
739

```

The `MonadHReader` class allows for the possibility that the environment (denoted with the r variables here) in a local computation is different than the environment in the outer computation. Because there may be many types that have `MonadHReader` instances, we must use the associated type family `SetEnv` to update the monad type with the new environment type.

In the inner computation, we need to know that the underlying monad, with the updated environment, is still a member of the `Monad` type class. This fact is assumed by putting the constraint `Monad (SetEnv r2 m)` on the inner computation, leveraging Haskell's support for higher-rank types [Peyton Jones et al. 2007].⁷

Returning to our indexed parser, we need these two instances:

```

749 instance Monad m => MonadHReader r1 (ReaderT r1 m) where
750   type SetEnv r2 (ReaderT r1 m) = ReaderT r2 m
751   hlocal f thing_inside = ...
752
753 instance MonadHReader r1 m => MonadHReader r1 (ParsecT s u m) where
754   type SetEnv r2 (ParsecT s u m) = ParsecT s u (SetEnv r2 m)
755   hlocal f thing_inside = ...

```

Here, `ReaderT` is the monad-transformer form of the `Reader` monad we saw earlier in the definition of `Parser`. (`Reader` is just defined to be a `ReaderT` based on the `Identity` monad.) The first instance says that the environment associated with a `ReaderT r1 m` is r_1 ; that is why the r_1 is the first parameter in the `MonadHReader` instance. It then describes that to update the environment from r_1 to r_2 , we just replace the type parameter to `ReaderT`. The implementation is straightforward and elided here.

The `ParsecT` instance lifts a `MonadHReader` instance through the `ParsecT` monad transformer, propagating the action of `SetEnv`. The implementation requires the usual type chasing characteristic of monad-transformer code, but offered no particular coding challenge.

With all this in place, it is straightforward to use the `hlocal` method in the `bind` function, giving us exactly the behavior that we want.

6 THE TYPE-INDEXED EXPRESSION AST

We now are ready to greet the `Exp` type, the type-indexed AST for expressions. Its definition appears in Figure 6. The `Exp` type is indexed by two parameters: a typing context of kind `Ctx n`, where n is the number of bound variables; and a type of kind `Type`.

Compare the definition of `Exp` with the typing rules in Figure 1. Each constructor corresponds with precisely one rule; the types of the constructor arguments correspond precisely with the premises of the rule; and the type of the constructor result corresponds precisely with the rule conclusion. Take function application as an example. The `T_APP` rule has two premises: one gives expression e_1 type $\tau_1 \rightarrow \tau_2$, and the other checks to see that e_2 has the argument type τ_1 . In the same way, the first argument to the constructor `App` takes an expression in some context `ctx` and with some type `arg` \rightarrow `res`. The second argument to `App` then has type `arg`. Furthermore, just as the conclusion to the `T_APP` rule says that the overall $e_1 e_2$ expression has type τ_2 , the result type

⁷A reader informed about recent updates to GHC might wonder why we do not use *quantified constraints* [Bottu et al. 2017] here. While this approach would seem to work, the current implementation fails us, because the head of a quantified constraint cannot be a type family, as described at <https://ghc.haskell.org/trac/ghc/ticket/14860>.


```

785 type Ctx n = Vec Type n
786 data Exp :: ∀n. Ctx n → Type → Type where
787   Var  :: Elem ctx ty → Exp ctx ty
788   Lam  :: TypeRep arg → Exp (arg :> ctx) res → Exp ctx (arg → res)
789   App  :: Exp ctx (arg → res) → Exp ctx arg → Exp ctx res
790   Arith :: Exp ctx Int → ArithOp ty → Exp ctx Int → Exp ctx ty
791   IntE :: Int → Exp ctx Int
792   ...
793
794   -- An encoding of (\x: Int. x + 1) 5, as an example
795 example :: Exp VNil Int
796 example = App (Lam (typeRep @Int) (Arith (Var EZ) Plus (IntE 1))) (IntE 5)
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833

```

Fig. 6. The type-indexed *Exp* expression AST

of the *App* constructor is an expression of type *res*. An easier example is for the constructor *IntE*, where the resulting type is simply *Int*, regardless of the context.

It is for this reason that modeling a typed language is such a perfect fit for GADTs—the information in the typing rules is directly expressed in the AST type definition.

6.1 The *Elem* type and type-indexed de Bruijn indices

Perhaps the most distinctive aspect of *Exp*—other than its indices—is the choice of representation for variables. *Exp* continues our use of de Bruijn indices, but we must be careful here: we need the type of a variable to be expressed in the return index to the *Var* constructor. While it is conceivable to do this via some *Lookup* type family, the *Elem* type is a much more direct approach:

```

812 data Elem :: ∀a n. Vec a n → a → Type where
813   EZ :: Elem (x :> xs) x
814   ES :: Elem xs x → Elem (y :> xs) x
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833

```

The *Elem* type is indexed by a vector (of any element type *a*) and a distinguished element of that vector. An *Elem* value, when viewed as a Peano natural number, is simply the index into the vector that selects that distinguished element. Equivalently, a value of type *Elem xs x* is a proof that *x* is an element of the vector *xs*; the computational content of the proof is *x*'s location in *xs*.

The definitions of the two constructors support this description. The *EZ* constructor has type *Elem (x :> xs) x*—we can see plainly that the distinguished element *x* is the first element in the vector. The *ES* constructor takes a proof that *x* is in a vector *xs* and produces a proof that *x* is in the vector *y :> xs* (for any *y*). Naturally, *x*'s index in *y :> xs* is one greater than *x*'s index in *xs*, thus underpinning the interpretation of *ES* as a Peano successor operator.

In the case of our use of *Elem* within the *Exp* type, the vectors at hand are contexts (vectors of *Types*) and the elements are types of Stitch variables. The *Elem* type gives us exactly what we need: a type-level relationship between a context and a type, along with the term-level information (the de Bruijn index) to locate that type within that context.

6.2 *Lam* requires the indexed *TypeRep*

Note the *Lam* constructor for building λ -abstractions. The first argument is *TypeRep arg*. This argument contains both a runtime type representation, suitable for runtime comparisons and

```

834 data ArithOp ty where
835   Plus, Minus, Times, Divide, Mod      :: ArithOp Int
836   Less, LessE, Greater, GreaterE, Equals :: ArithOp Bool
837
838   -- Like Ex, but includes a Typeable constraint for the existentially bound index
839   -- This is declared in the Data.Exists module with the Ex type
840 data TypeableEx :: (k → Type) → Type where
841   TypeableEx :: Typeable i ⇒ a i → TypeableEx a
842   -- UArithOp ("unchecked ArithOp") is an existential package for an ArithOp
843 type UArithOp = TypeableEx ArithOp
844 pattern UArithOp op = TypeableEx op
845 {-# complete UArithOp #-}
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882

```

Fig. 7. Arithmetic operators, from the Op module.

pretty-printing, and also a compile-time type index *arg*, used later in the type of *Lam*. Like *replicate*, this is a place where a dependent type is called for. Happily, the *TypeRep* singleton works well here.

It may be interesting to note that this *TypeRep* argument was actually not required in an early (but fully working) version of *Stitch*. Lacking the *TypeRep* meant that the pretty-printer was unable to annotate type-checked λ -expressions, but that was the only drawback. The *arg* type index was (and still is) an existential type, packed by the *Lam* constructor. Because the choice of *arg* was never needed at runtime, no runtime witness was necessary. The addition of *TypeRep* was forced, however, when implementing common-subexpression elimination, as the argument is necessary in order to write *Exp*'s *TestEquality* instance. See Section 9.1.

6.3 Arithmetic operators

The *Arith* constructor contains two subexpressions and the choice of arithmetic operator. All binary operators in *Stitch* operate on two *Ints*, so the subexpressions are constrained each to have type *Int*. The return type, on the other hand, varies with the operator. For example, *+* produces an *Int* while *<* produces a *Bool*. We thus need another indexed type, *ArithOp*, indexed by the return type of the operation.⁸ The definition appears in Figure 7.

Given the introduction above, this definition should be very unsurprising. Additionally, Figure 7 includes definitions for *UArithOp*, the unindexed variant of *ArithOp*, used before type checking. A *UArithOp* must store the singleton associated with the existentially bound type index so that the *Stitch* type checker can compare this type with the expected type of an expression.

7 THE SOUND TYPE-INDEXED TYPE CHECKER

We are ready now for the part we have all been waiting for: the sound type-indexed type checker. Many cases appear in Figure 8; these cases illustrate the points of interest.

At its core, the *check* function takes an unchecked expression of type *UExp* and converts it into a checked expression of type *Exp*. Already we see an unexpected twist in the type of *check*: it is written in continuation-passing style (CPS). The reason for this is that there is naturally no way to know what indices should be placed on the output *Exp*. What we would like to write, ideally, is $check :: UExp \text{ Zero} \rightarrow \exists ty. Exp \text{ VNil } ty$ (ignoring the monadic context). However, Haskell does not support such a convenient construct. While we could use the *Ex* existential package here quite

⁸It would be easy to generalize this also to be indexed by argument types, if they varied among operators.

```

883 check :: (MonadError Doc m, MonadReader Globals m)
884         ⇒ UExp Zero → (∀(t :: Type). TypeRep t → Exp VNil t → m r) → m r
885 check = go SCNil
886
887 where
888 go :: (MonadError Doc m, MonadReader Globals m, SNat1 n)
889       ⇒ SCtx (ctx :: Ctx n) → UExp n → (∀t. TypeRep t → Exp ctx t → m r) → m r
890 go ctx (UVar n) k = check_var n ctx $ λty elem →
891                   k ty (Var elem)
892 where check_var :: Fin n → SCtx (ctx :: Ctx n)
893         → (∀t. TypeRep t → Elem ctx t → m r) → m r
894 check_var FZ (ty :%> _) k0 = k0 ty EZ
895 check_var (FS n0) ( _ :%> ctx0) k0 = check_var n0 ctx0 $ λty elem →
896                                         k0 ty (ES elem)
897
898 go _ (UGlobal n) k = do globals ← ask
899                       lookupGlobal globals n $ λty exp →
900                           k ty (shifts0 exp)
901
902 go ctx (ULam (Ty arg_ty) body) k = go (arg_ty :%> ctx) body $ λres_ty body' →
903                                       k (arg_ty :→ res_ty) (Lam arg_ty body')
904
905 go ctx e@(UApp e1 e2) k = go ctx e1 $ λfun_ty e'1 →
906                               go ctx e2 $ λarg_ty e'2 →
907                                   case fun_ty of arg_ty' :→ res_ty
908                                       | Just HRefl ← eqTypeRep arg_ty arg_ty'
909                                       → k res_ty (App e'1 e'2)
910                                       _ → typeError e ...
911
912 go ctx e@(UArith e1 (UArithOp op) e2) k = go ctx e1 $ λty1 e'1 →
913                                               go ctx e2 $ λty2 e'2 →
914                                                   case (isTypeRep @Int ty1, isTypeRep @Int ty2) of
915                                                       (Just HRefl, Just HRefl)
916                                                       → k typeRep (Arith e'1 op e'2)
917                                                       _ → typeError e ...
918
919 go _ (UIntE n) k = k typeRep (IntE n)

```

Fig. 8. The sound type-indexed type checker (excerpts)

profitably, I found that CPS was easier and made for code with a better flow. With CPS, we can pass the type index t to the continuation using a higher-rank type for *check*. We also must pass *TypeRep t* to the continuation, so that runtime comparisons can be performed.

The *check* function works over closed expressions, as we always call it on a top-level expression. However, it must recur into open expressions, and so we define the more-general *go* local helper function. The *go* function's type mimics that of *check* but allows for the possibility of open expressions, quantifying over the context length, n , and context *ctx*. Because we will need to look up variable types at runtime, we need the context to be available both at compile-time (to use as an

index to *Exp*) and at runtime. This means that we need a singleton for the context, as embodied by this definition:

```

932 data SCtx :: ∀n. Ctx n → Type where
933   SCNil :: SCtx VNil
934   (:%>) :: TypeRep t → SCtx ts → SCtx (t :> ts)

```

An *SCtx* operates analogously to a *SNat*, forcing the runtime value to match exactly the compile-time type index.

The other small curiosity in the type of *go* is that it adds a *SNat!* *n* constraint, where *n* is the length of the typing context. This constraint is not needed for type checking but instead is needed only for pretty-printing. In the text produced for type errors (elided here), we often want to print parts of expressions. Recall that the pretty-printer colors the de Bruijn indices in the output to indicate the indices' provenance (i.e., which binder they refer to). While the numeral to output can be read directly from the *Fin* or *Elem* datatype, the color cannot—the color is computed by subtracting the value of the de Bruijn index from the number of in-scope variables.⁹ For example, suppose the second bound variable is rendered in purple (as it is in the examples in Section 2.2). When two variables are in scope, index 0 should be purple. But if three variables are in scope, index 1 should be purple: the invariant here is that the index two less than the number of in-scope variables is purple. Accordingly, the pretty-printer needs to know the number of in-scope variables at runtime. This number is the type index *n*, and thus we need the singleton for *n*; in this case, it is convenient to pass it implicitly, leading to the *SNat!* *n* constraint.

7.1 Checking variables

The variable case is handled by the helper function *check_var*. The *check_var* function uses the *Fin* *n* stored by the *UVar* constructor to index into the typing context, stored as a the singleton *SCtx*. When *check_var* finds the type it is looking for, it passes that type to the continuation, along with an *Elem* value which will store the de Bruijn index in the *Exp* type. GHC's type checker is working hard here to make sure this function definition is correct, using the definition of *Fin* to ensure that our pattern-match is complete,¹⁰ and that the *Elem* we build really does show that the type *t* is in the context *ctx*. Note that there is no possibility of errors here: the use of *Fin* in the *UExp* type guarantees that the variable is in scope.

7.2 Inlining globals

Stitch allows its users to declare global variables in the REPL, as demonstrated in Section 2.2. Expressions to be stored in globals are parsed and type-checked, with the type-checked *Exp* stored for later retrieval. Of course, a global can have any type, and so the data structure used to store the globals must use an existential. All globals are closed, and so we already know that the context must be empty. The definition of the *Globals* datatype appears in Figure 9.

Globals is a newtype wrapping a finite map from strings (global variable names) to existential-packed expressions. We pack these expressions along with a *Typeable* constraint containing the expressions' types, for retrieval during type checking. As the type checking algorithm uses an explicit *TypeRep* for types, we use the *unpackTypeRepEx* function, which unpacks a *TypeableEx* existential package, converting the implicit *Typeable* type representation to an explicit *TypeRep*. (Recall that the *typeRep* function used in *unpackTypeRepEx* has type *Typeable* *a* ⇒ *TypeRep* *a*.)

⁹That is, the color is a representation of a de Bruijn *level*, not an index.

¹⁰Note that we match the *Fin* before the vector, as we did in Section 3.2.2.

```

981 newtype Globals = Globals (M.Map String (TypeableEx (Exp VNil)))
982
983 lookupGlobal :: MonadError Doc m
984             => Globals → String → (∀ty. TypeRep ty → Exp VNil ty → m r) → m r
985 lookupGlobal (Globals globals) var k = case M.lookup var globals of
986             Just exp → unpackTypeRepEx exp k
987             Nothing → throwError ...
988
989 -- From Data.Exists; unpacks a TypeableEx, providing an explicit TypeRep
990 unpackTypeRepEx :: TypeableEx a → (∀i. TypeRep i → a i → r) → r
991 unpackTypeRepEx (TypeableEx x) k = k typeRep x
992

```

Fig. 9. Storing and retrieving global variables; module *M* refers to *Data.Map* from the containers package

A further complication arises in the fact that we inline the value of a global variable into an expression with potentially a non-empty context. Globals have an empty context, and so we must be careful to shift de Bruijn indices when inlining the global. I defer discussion of the *shifts0* function until we have talked about evaluation, where de Bruijn shifting is more at home. See Section 8.3. Regardless of the details, however, we can see already that the strongly typed discipline within Stitch prevents us from forgetting about this shifting: the continuation in *go* expects a *Exp ctx t*, where the context is provided as a parameter to *go*. That is, the context is known ahead of time. Since *lookupGlobal* passes an expression in an empty context to its continuation, that expression cannot be directly passed to the continuation of *go*: GHC would issue an error saying that it cannot prove that *ctx* is *VNil*. This error is spot on, pointing out that we have confused an expression in an empty context with one in a potentially non-empty one, necessitating shifting.

7.3 Checking a λ -abstraction

The *Lam* case is remarkably straightforward. We check the abstraction body, learning its result type *res_ty* and getting the type-checked expression *body'*. We then continue with a function type composed from *arg_ty* (as unpacked from the *Ty* stored by the *ULam* constructor) and *res_ty*, using our \rightarrow pattern synonym (Section 4.1). Note that if we did not store the *arg_ty* indexed *TypeRep* in the *ULam*, we would be stuck here.

7.4 Checking an application

Checking function applications is really the heart of any type checker: this is the principal place where two types may be in conflict. In our case, we check the two expressions separately, getting their types and type-checked expression trees. We then must ensure that *fun_ty*, the type of the applied function, is indeed a function type. This is done by matching against the \rightarrow pattern synonym. We then must ensure that the actual argument type *arg_ty* matches the function's expected argument type *arg_ty'*. We use the *eqTypeRep* function, exported from GHC's *Type.Reflection* module and explained in Section 4.2. If successful, this function returns a proof to the type checker that *arg_ty* equals *arg_ty'*, and we are then allowed to build the application with *App*. If either check fails, we issue an error.

The type discipline in Stitch is working hard to keep us correct here. If we skipped the type checks, the *App* application would be ill-typed, as *App* expects its first argument to be a function and its second argument to have the argument type of that function. The checks ensure this to GHC, which then allows our use of *App* to succeed.

1030 7.5 Arithmetic expressions

1031 Arithmetic expressions are straightforward to check, following broadly the pattern we saw in
 1032 the function application case: simply check all the *TypeReps*. We make use here of the *isTypeRep*
 1033 function we defined in Section 4.2 to check that both arguments are indeed *Ints*. Upon success, we
 1034 can retrieve the result type of the expression by using *typeRep*; recall that the *UArithOp* type (Section
 1035 6.3) stores a *Typeable* constraint for the operation type via its definition in terms of *TypeableEx*.
 1036 Type inference figures out that the use of *typeRep* here should correspond to the result type of
 1037 *Arith*, in turn set by the use of *op* as an argument to *Arith*.

1038 We conclude with the case for integer literals. In the call of the continuation, we can once again
 1039 use *typeRep*, as the use of *IntE* tells us we need the representation for the type *Int*.
 1040

1041 There are several more cases in the type checker, all similar to those presented here. In all, this
 1042 type checker was remarkably easy to write, given the groundwork in setting up the types correctly.
 1043 GHC’s type checker stops us from making mistakes here—the whole point of using an indexed
 1044 expression AST—and GHC’s type inference allows us the convenience to pass type representations
 1045 implicitly. Furthermore, the type errors I encountered during implementation were indeed helpful,
 1046 pointing out any missing type equality checks.

1047 Beyond these observations, I wish to note simply that such a type checker is possible to write at
 1048 all. In conversations with experienced functional programmers, some have been surprised that the
 1049 type-indexed expression AST has any practical use at all, despite the fact that this technique is not
 1050 new [e.g., Pašalić et al. 2002]. After all, how could you guarantee that expressions are well typed?
 1051 The answer is, of course, by checking them, as *check* does for us here.
 1052

1053 8 EVALUATION WITH AN INDEXED AST

1054 Writing evaluators is where the indexed AST really shines: we essentially can not get it wrong.

1055 A type-indexed AST allows us to easily write a *tagless* interpreter, where a value does not need
 1056 to be stored with a runtime tag that indicates the value’s type. To see the problem, imagine an
 1057 unindexed AST and a function *eval* :: *Exp* → *Value*. The *Value* type would have to be a sum type
 1058 with several constructors, say, for integer, Boolean, and function values. This means that every
 1059 time we extract a value, we have to check the tag, a potentially costly step at runtime. With our
 1060 indexed expression type, we can evaluate to a type *Value ty*, where *Value* is this type family:
 1061

1062 **type family** *Value t where*

1063 *Value Int* = *Int*

1064 *Value Bool* = *Bool*

1065 *Value (a → b)* = *Exp VNil a → Exp VNil b*
 1066

1067 Values are accordingly tagless—no runtime check needs to be performed when inspecting one.
 1068 Tagless interpreters have been studied at some length [Carette et al. 2009; Pašalić et al. 2002; Taha
 1069 et al. 2001], and we will not explore this aspect of Stitch further.

1070 The two evaluators for Stitch are straightforward transcriptions of Stitch’s operational semantics
 1071 (Figure 1). There is only one small hitch: encoding values. We sometimes need to translate a value
 1072 back into an expression—for example, when we substitute that value in for a variable during
 1073 β -reduction. We thus define a type *ValuePair* :: *Type* → *Type* that stores closed expressions along
 1074 with the untagged values. As there is only one constructor for the *ValuePair* type, its tag need not
 1075 be checked at runtime. Its definition, along with the big-step evaluator, appear in Figure 10.

1076 The helper functions *apply* and *arith* are routine and elided. Note, however, the *impossibleVar*
 1077 function, which eliminates the possibility of encountering a variable in an empty context. It is
 1078

```

1079     data ValuePair ty = ValuePair { expr :: Exp VNil ty, val :: Value ty }
1080     eval :: Exp VNil t → ValuePair t
1081     eval (Var v)           = impossibleVar v
1082     eval e@(Lam _ body) = ValuePair e $ λarg → subst arg body
1083     eval (App e1 e2)   = eval (apply (eval e1) (eval e2))
1084     eval (Arith e1 op e2) = eval (arith (val $ eval e1) op (val $ eval e2))
1085     eval e@(IntE n)      = ValuePair e n
1086     ...
1087
1088     impossibleVar :: Elem VNil x → a
1089     impossibleVar = λcase { }
1090

```

Fig. 10. Implementation of big-step operational semantics

```

1094     data Length :: ∀a n. Vec a n → Type where
1095     LZ :: Length VNil
1096     LS :: Length xs → Length (x :> xs)
1097
1098     subst :: ∀ctx s t. Exp ctx s → Exp (s :> ctx) t → Exp ctx t
1099     subst e = go LZ
1100     where
1101     go :: Length (locals :: Ctx n) → Exp (locals ++ s :> ctx) t0 → Exp (locals ++ ctx) t0
1102     go len (Var v)           = subst_var len v
1103     go len (Lam ty body) = Lam ty (go (LS len) body)
1104     ... -- other forms are treated homomorphically
1105
1106     subst_var :: Length (locals :: Ctx n) → Elem (locals ++ s :> ctx) t0 → Exp (locals ++ ctx) t0
1107     subst_var LZ     EZ     = e           -- no locals; substitute
1108     subst_var LZ     (ES v) = Var v       -- no locals; decrement index
1109     subst_var (LS _)  EZ     = Var EZ     -- variable is local; no change
1110     subst_var (LS len) (ES v) = shift (subst_var len v) -- recur
1111

```

Fig. 11. Indexed substitution, from the Eval module

implemented via an empty **case** expression. Empty **case** expressions are strict in Haskell, in contrast to non-empty **cases**. When the *Elem VNil x* is evaluated, it must be *ES* or *EZ*, both of which cannot be indexed by an empty context. GHC thus discovers that *Elem VNil x* is an empty type, and the empty **case** is accepted as a complete pattern match.

8.1 Substitution

We are left to discuss the bane of implementors using de Bruijn indices: substitution. Once again, the type indices save us from making errors—there seems to be no real way to go wrong, and the type errors that we encounter gently guide us to the right answer. The final result is in Figure 11.

The *subst* function takes an expression *e* of type *s* and another expression with a free variable of type *s* and substitutes *e* into the latter expression. The *subst* function’s type requires that the variable to be substituted have a de Bruijn index of 0, as is needed during β -reduction. However,

1128 as anyone who has proved a substitution lemma knows, we must generalize this type to get a
 1129 powerful enough recursive function to do the job.

1130 Note that the type of *subst* is precisely the shape of a substitution lemma: that if $\Gamma \vdash e_1 : \sigma$
 1131 and $\Gamma, x:\sigma \vdash e_2 : \tau$, then $\Gamma \vdash e_2[e_1/x] : \tau$. A proof of this lemma must strengthen the induction
 1132 hypothesis to allow bound local variables, leading to a proof of this stronger claim: if $\Gamma \vdash e_1 : \sigma$
 1133 and $\Gamma, x:\sigma, \Gamma' \vdash e_2 : \tau$, then $\Gamma, \Gamma' \vdash e_2[e_1/x] : \tau$. If we call Γ' *locals* and Γ *ctx*, this strengthened
 1134 induction hypothesis matches up with the type of the helper function *go*. (Recall that contexts
 1135 in the implementation are in reverse order to those in the formalism.) As one implements such a
 1136 function, this correspondence is a strong hint that the function type is correct.

1137 The *go* function takes one additional argument: a value of type *Length locals*. The *Length* type
 1138 is included in Figure 11; values are Peano naturals that describe the length of a vector.¹¹ This
 1139 extra piece is necessary as local variables get treated differently in a substitution than do variables
 1140 from the outer context. The number of locals informs the *subst_var* function when to substitute,
 1141 when to shift, and when to leave well enough alone. Pierce [2002, Chapter 6] offers an accessible
 1142 introduction to the delicate operation of substitution in the presence of de Bruijn indices, and a
 1143 full exploration of this algorithm would take us too far afield; suffice it to say that any misstep in
 1144 *subst_var* would be caught by GHC’s type checker.

1145

1146 8.2 Shifting

1147 As hinted at previously, substitution with de Bruijn indices is subtle not only because it is hard to
 1148 keep track of which variable one is substituting, but also because the expression being substituted
 1149 suddenly appears in a new context and accordingly may require adjustments to its indices. This
 1150 process is called *shifting*.¹² If we have an expression *#1 #0* (where both variables are free) and wish
 1151 to substitute into an expression with an additional bound variable, we must shift to *#2 #1*. I have
 1152 intentionally kept the colors consistent during the shift, as the identity of these variables does *not*
 1153 change—just the index does.

1154 Shifting is an operation that makes sense both on full expressions *Exp* and also on indices
 1155 *Elem* directly. We will discover that both of these are sometimes necessary when performing
 1156 common-subexpression elimination (CSE, Section 9), and so we generalize the notion of shifting by
 1157 introducing a type class. The relevant definitions are in Figure 12.

1158 The first detail to notice here is that *Shiftable* classifies a polykinded type variable *a*—note the $\forall n$
 1159 in *a*’s kind. This gives *Shiftable* a *higher-rank kind*. GHC deals with this exotic species in stride; the
 1160 only challenge is that GHC will never infer a variable to have a polykind, and so all introductions
 1161 of *a* must be written with a kind annotation. We see this in the type of *shift*. The polymorphism in
 1162 the kind of *a* is essential here because, as a stand-in for *Exp* or *Elem*, *a* must be able to be applied to
 1163 contexts of any length. Without this polymorphism, it would be impossible to write the *Shiftable*
 1164 class.

1165 As before, the implementation of these functions is straightforward, once we have written down
 1166 the types and can be guided by GHC’s type checker. The types themselves come straight from
 1167 standard type theory, where they correspond to the weakening and strengthening lemmas.

1168

1169 8.3 Using *shifts0* in the type checker

1170 Part of the discussion about the *UGlobal* case in the type checker (Section 7.2) was deferred until
 1171 after we have introduced shifting. We return to this case here. The code is in Figure 8.

1172

1173 ¹¹Although vectors are indexed by their length, that index is a compile-time natural only. To get the length of a vector at
 1174 runtime, it is still necessary to recur down the length of the vector.

1175 ¹²In a call-by-value λ -calculus, this shifting will never affect a substituted expression, as all such expressions are closed.
 1176 However, the definition of substitution is general and must take this shifting into account.

1176


```

1177 class Shiftable (a ::  $\forall n. \text{Ctx } n \rightarrow \text{Type} \rightarrow \text{Type}$ ) where
1178   shifts    ::  $\text{Length prefix} \rightarrow a \text{ ctx } ty \rightarrow a (\text{prefix} \text{++} \text{ctx}) ty$   -- multishifts are needed in CSE
1179   shifts0   ::  $a \text{ VNil } ty \rightarrow a \text{ prefix } ty$ 
1180   unshifts ::  $\text{Length prefix} \rightarrow a (\text{prefix} \text{++} \text{ctx}) ty \rightarrow \text{Maybe } (a \text{ ctx } ty)$   -- needed for CSE
1181
1182 instance Shiftable Exp where
1183   shifts    = shiftsExp
1184   shifts0   = shifts0Exp  -- see Section 8.3
1185   unshifts = unshiftsExp  -- elided
1186
1187 instance Shiftable Elem where . . .
1188   -- Convenient abbreviation for the common case of shifting by only one index
1189   shift ::  $\forall (a :: \forall n. \text{Ctx } n \rightarrow \text{Type} \rightarrow \text{Type}) \text{ ctx } t \text{ ty}. \text{Shiftable } a \Rightarrow a \text{ ctx } ty \rightarrow a (t \text{:>} \text{ctx}) ty$ 
1190   shift = shifts (LS LZ)
1191   shiftsExp ::  $\forall \text{prefix ctx ty}. \text{Length prefix} \rightarrow \text{Exp ctx ty} \rightarrow \text{Exp } (\text{prefix} \text{++} \text{ctx}) ty$ 
1192   shiftsExp prefix = go LZ
1193   where
1194     go ::  $\text{Length } (\text{locals} :: \text{Ctx } n) \rightarrow \text{Exp } (\text{locals} \text{++} \text{ctx}) ty_0 \rightarrow \text{Exp } (\text{locals} \text{++} \text{prefix} \text{++} \text{ctx}) ty_0$ 
1195     go len (Var v)          = Var (shifts_var len v)
1196     go len (Lam ty body) = Lam ty (go (LS len) body)
1197     . . .  -- other forms are treated homomorphically
1198
1199     shifts_var ::  $\text{Length } (\text{locs} :: \text{Ctx } n) \rightarrow \text{Elem } (\text{locs} \text{++} \text{ctx}) ty_0 \rightarrow \text{Elem } (\text{locs} \text{++} \text{prefix} \text{++} \text{ctx}) ty_0$ 
1200     shifts_var LZ v          = weakenElem prefix v
1201     shifts_var (LS _) EZ   = EZ
1202     shifts_var (LS l) (ES e) = ES (shifts_var l e)
1203
1204     -- Weaken an Elem to work against a larger vector.
1205     weakenElem ::  $\text{Length prefix} \rightarrow \text{Elem } xs \ x \rightarrow \text{Elem } (\text{prefix} \text{++} xs) \ x$ 
1206     weakenElem LZ e         = e
1207     weakenElem (LS len) e = ES (weakenElem len e)

```

Fig. 12. De Bruijn index shifting, from the Shift module

The challenge is that globals all refer to *closed* expressions, and yet the global might be used in a context with several bound variables. We must, therefore, adjust the context of the expression stored in the global. However, the usual shifting logic surely is overkill here: a global variable expression is closed, after all. There is no way shifting can possibly make a difference!

While we *could* use the general shifting mechanism, we instead prefer to use a specialization of shifting, tailored for closed expressions, *shifts0*. See Figure 13, which defines *shifts0Exp*, the definition of *shifts0* in the *Shiftable* instance for *Exp*. This function tiresomely walks the entire structure of its argument in order to do nothing. The problem is that the type of the output really is different than the type of the input; the only way to convince GHC that no action needs to be taken is a full recursive traversal.

This is disappointing. We want our types to help prevent errors, not require extra runtime work. It is conceivable that a language with full dependent types would support a proof that *shifts0Exp*

```

1226 shifts0Exp :: ∀ prefix ty. Exp VNil ty → Exp prefix ty
1227 shifts0Exp = go LZ
1228 where
1229   go :: Length (locals :: Ctx n) → Exp locals ty0 → Exp (locals ++ prefix) ty0
1230   go len (Var v)      = Var (shifts0_var v len)
1231   go len (Lam ty body) = Lam ty (go (LS len) body)
1232   ... -- other forms are treated homomorphically
1233
1234   shifts0_var :: Elem locals ty0 → Length (locals :: Ctx n) → Elem (locals ++ prefix) ty0
1235   shifts0_var EZ _      = EZ
1236   shifts0_var (ES v) (LS len) = ES (shifts0_var v len)
1237
1238 -- Because shifts0Exp provably does nothing, we can short-circuit it:
1239 {-# ninline shifts0Exp #-}
1240 {-# rules "shifts0Exp" shifts0Exp = unsafeCoerce #-}
1241
1242
1243
1244

```

Fig. 13. Shifting closed expressions should be trivial

has no runtime effect, but this is still hard to imagine, given that the output of *shifts0Exp* has a different type than its input.

The fullness of GHC’s feature set comes to the rescue here. GHC supports *rewrite rules* [Peyton Jones et al. 2001], which allow a programmer to provide arbitrary term rewriting rules that GHC applies during its optimization passes. These rules are type-checked to make sure both sides have the same type, but no checking is done for semantic consistency. It is just the ticket for us here: we can fix the types up with an *unsafeCoerce* and trust our by-hand analysis that *shifts0Exp* really does nothing at runtime. The *noinline* is necessary because GHC might observe that *shifts0Exp* is a short function (because it is defined almost immediately in terms of *go*) and decide to inline it. The *noinline* tells GHC not to, and that way the rewrite rule can trigger.

Is this design a win or a loss? I am not sure. It surely has aspects of a loss because the compiler can not figure out that *shifts0Exp* is pointless. On the other hand, the workaround is very easy and fully effective. And, even in a language with a richer type system than GHC’s Haskell, it is not clear we can do better.

9 COMMON-SUBEXPRESSION ELIMINATION

Having covered the basic necessities of an interpreter, we now explore an extension, as evidence that we can still implement non-trivial transformations over an indexed AST. Common-subexpression elimination is a standard optimization pass, which identifies expressions with common subexpressions, transforming these to use a let-bound variable instead. A full description of the CSE algorithm is unnecessary here but is well documented in the CSE module; instead, we will focus on the (indexed) data structures used to power the CSE algorithm.

The key data structure needed for CSE is a finite map that uses expressions as keys. Using such a map, we can store what expressions we have seen so far in order to find duplicates, and we can map expressions to fresh let-bound variables. The challenge here is that we need to make sure an expression of type *ty* maps to a variable of type *ty*; failing to do so would lead the CSE algorithm not to pass GHC’s type checker.

Naturally, we want the CSE algorithm to be reasonably efficient. Instead of creating our own mapping structure, we would like to use the existing optimized *HashMap* structure from the

```

1275 -- from GHC's Data.Type.Equality module
1276 class TestEquality (t :: k → Type) where testEquality :: t a → t b → Maybe (a ~: b)
1277 class IHashable (t :: k → Type) where ihashWithSalt :: Int → t a → Int -- in Data.IHashable
1278 instance TestEquality (Elem xs) where ... -- in Data.Vec
1279 -- in Exp
1280 type KnownLength (ctx :: Ctx n) = SNatI n -- "a context's length is available at runtime"
1281 instance TestEquality (Exp ctx) where ...
1282 instance KnownLength ctx ⇒ IHashable (Exp ctx) where ...
1283 instance KnownLength ctx ⇒ IHashable (Elem ctx) where ...
1284 -- In Data.IHashMap.Base:
1285 data IHashMap :: ∀k. (k → Type) → (k → Type) → Type where ...
1286 insert :: (TestEquality k, IHashable k) ⇒ k i → v i → IHashMap k v → IHashMap k v
1287 lookup :: (TestEquality k, IHashable k) ⇒ k i → IHashMap k v → Maybe (v i)
1288 map :: (∀i. v1 i → v2 i) → IHashMap k v1 → IHashMap k v2
1289 type ExpMap ctx a = IHashMap (Exp ctx) a -- In CSE

```

Fig. 14. Key definitions for indexed *HashMap*s

unordered-containers library, a widely-used containers implementation. However, a *HashMap* requires that all the keys in the map have the same type. This is usually a desired property, but not in our case here: the different keys will all be *Exps*, but they may have different type indices. The solution is to alter *HashMap* to work with indexed types. To implement this idea, I took the source code from unordered-containers, made a few small changes to the types, and then simply fixed the errors that GHC reported. Some key definitions are in Figure 14.

9.1 Indexed maps

Just as a traditional mapping structure must depend on a key's *Eq* instance, an indexed mapping structure must depend on a key's *TestEquality* instance. The *TestEquality* class includes indexed types where an equality test can inform the equality of the indices. In our case, this clearly includes *Exp ctx*, because we can compare two expressions; if they are equal (in the shared context), then surely their types are the same. As *Exp* is indexed by its type, a comparison between the values gives us an equality between their type indices—exactly the contract *TestEquality* requires.

We also must generalize the *Hashable* class used for traditional *HashMap*s so that we can state that *Exp* has a hash, no matter its type. This is straightforward to do; see *IHashable*.

In the definition of *IHashMap*, we must index the map by the type constructors, not the concrete types. Note that in the definition for *ExpMap*, the key is *Exp ctx*, not *Exp ctx ty*. In this way, a map can contain expressions of many types. Accordingly, the *insert* and *lookup* functions work by applying the key type *k* and value type *v* to an index *i*. (Note: the *k* in the definition of *IHashMap* is the kind of the index, not the key.) The magic here is that *IHashMap* is not itself indexed by *i*, so we can look up *k i*, for any *i*, in a *IHashMap k v*, retrieving (perhaps) a *v i*.

Though not used in CSE, I have included here the type of the *map* function. Its function argument must be polymorphic in the index *i*. This is because the function must work over all values stored in the map; these values, of course, may have different indices. With a higher-rank type, however, *map* (and other functions) are straightforward to adapt to the indexed setting.

9.2 Experience report

The adaptation of *HashMap* into an indexed setting was shockingly easy. Once I had committed to adapting the existing implementation, it took me roughly 2 hours to update the 2.5k lines of code implementing lazy *HashMaps* and *HashSets*. The process flowed as we all imagine typed refactoring should: I changed the datatype definitions and just followed the errors. It all worked splendidly once it compiled. I was aided by the fact that *TestEquality* is already exported from GHC's set of libraries and that this class has just the right shape for usage in a finite map structure.

Many functions, such as *map* require higher-rank types. Interestingly, several class instance definitions also require a higher rank, but these require a higher-rank *constraint*, also known as a quantified constraint [Bottu et al. 2017]. For example, here are the instance heads for two instances of *IHashMap*:

```
instance (TestEquality k, IHashable k,  $\forall i. (Read (k\ i), Read (v\ i)) \Rightarrow Read (IHashMap\ k\ v)$ )
instance ( $\forall i. (Show (k\ i), Show (v\ i)) \Rightarrow Show (IHashMap\ k\ v)$ )
```

In order to parse the contents of a *IHashMap k v*, we need to be able to read elements of type *k i* and *v i*, for any *i*, and similarly for pretty-printing. With quantified constraints, we can express this fact directly, and type-checking proceeds without a hiccup.

The CSE implementation overall was also agreeably easy. While the design of the algorithm took some careful thought, working with indexed types was an aid to the process, not an obstacle. The way *Exp*'s indices track contexts, in particular, was critical, because any recursive algorithm over *Exps* must occasionally change contexts; it would have been very easy to forget a shift or unshift during this process without GHC's type checker helping me get it right.

10 DISCUSSION

10.1 Polymorphic recursion in types

It is well known that polymorphic recursion is impossible with Damas-Milner type inference [Henglein 1993; Mycroft 1984]. If we want to write a polymorphic recursive function, we must supply a type signature.

However, what if a *type* is polymorphic recursive? That is, a recursive occurrence in a type definition might have a parameter of a different kind than the outer definition. A handy example is the *Length* type, repeated here:

```
data Length ::  $\forall a\ n. Vec\ a\ n \rightarrow Type$  where
  LZ :: Length VNil
  LS :: Length xs  $\rightarrow Length (x\ :>\ xs)$ 
```

This type is polymorphic recursive because the recursive occurrence in the *LS* constructor takes a parameter *xs* which has a different kind (*Vec a n*) than the kind of the parameter of the return type of *LS*, which is *Vec a (Succ n)*. When should GHC accept such a definition? In other words, when does a type have a *kind signature*?

Given the syntax of GHC, this is not an easy question to answer. For example, the *Length* type as written above still requires a small amount of kind inference: I have not written the kinds of *a* or *n*. Other forms of type declarations have other confounding details. Worse, the decision whether or not a type has a kind signature must be made very early, before doing any kind inference on the type: the signal must be purely syntactic.

Accordingly, GHC defines a set of rules describing when types have a so-called *complete user-specified kind signature*, or CUSK. These rules, as documented in the GHC manual, say that a datatype declaration has a CUSK when any kind variables mentioned in its explicit kind are

1373 explicitly quantified (among other rules). This means that the $\forall a\ n$ above is compulsory—if I omit
 1374 this, the type does not have a CUSK and thus cannot be polymorphic recursive.

1375 This leads to an unpleasant user experience. Leaving out the explicit quantification induces an
 1376 error message about mismatched kinds. It is not hard to work out that GHC is struggling to infer
 1377 polymorphic recursion from this message, but nothing suggests to add explicit quantification to
 1378 solve the problem. Instead, the programmer has to already be familiar with the vagaries of CUSKs
 1379 to figure out what to do.

1380 Happily, there is already an accepted GHC proposal [Eisenberg 2017] to fix this problem by
 1381 allowing users to write kind signatures distinct from type declarations, much as we do with
 1382 term-level functions.

1383

1384 10.2 *let should sometimes be generalized*

1385 Type inference in the presence of GADTs is hard [Chen and Erwig 2016; Peyton Jones et al. 2006,
 1386 2004; Vytiniotis et al. 2011]. One of the confounding effects of GADTs is that GHC does not
 1387 generalize local **let**-bound variables in a module with the `MonoLocalBinds` language flag enabled,
 1388 which is implied by the GADTs extension [Vytiniotis et al. 2010].¹³ However, in two separate places,
 1389 this lack of generalization stymied my implementation:

1390

1391 *Generalizing type signatures.* If a function’s type signature can be kind-generalized, GHC will
 1392 automatically generalize it. For example, if we declare `typeRepShow :: TypeRep a → String`, GHC
 1393 will infer that we really mean `typeRepShow :: ∀ k (a :: k). TypeRep a → String`. This implicit
 1394 generalization is useful and rarely gets in the way.

1395 However, if I am declaring a local function whose type mentions in-scope variables from an
 1396 outer scope, GHC does not kind-generalize, for exactly the same reasons that it does not generalize
 1397 term-level **let**-definitions. (Vytiniotis et al. [2010] lay out these motivations in great detail.) This
 1398 means that my type signature must explicitly mention any kind variables I wish to generalize over.
 1399 This restriction bit me in the `go` helper functions to `subst` and `shiftsExp`, where the functions must
 1400 be generalized over the length of the local context. I had not explicitly done so at first, and it took
 1401 me some time to figure out what was going wrong. It might be helpful for GHC to alert a user
 1402 when a **let** or type signature has been prevented from generalization.

1403

1404 *Generalizing polymorphic traversal functions.* In the adaptation of `HashMap` to `IHashMap`, it was
 1405 necessary to make many traversal functions have higher-rank types, like `map` in Section 9.1. Other
 1406 functions in the `HashMap` library use these traversals with locally defined helper functions, which
 1407 generally lacked type signatures. However, because **lets** were not generalized in the module, the
 1408 type of the **let**-bound function was not polymorphic enough to be used as the argument to the
 1409 higher-rank traversal function. While adding the type signatures to the local functions was not
 1410 terribly difficult, it was tedious, and I opted instead to specify `NoMonoLocalBinds`, to good effect.

1411

1412 10.3 *Dependent types*

1413 To my surprise, this project did *not* strongly want for full dependent types. As we have seen, we
 1414 needed a few singletons. A language with support for dependent types would naturally not need
 1415 these singletons. However, one of the real pain points for singletons—costly runtime conversions
 1416 between singletons and unrefined types—arose in only one place: the calculation of what color is
 1417 used to render a de Bruijn index. Another big pain point is code duplication, but that problem, too,

1418

1419 ¹³More precisely, GHC does not generalize local **let**-bound variables whose right-hand side mentions a variable bound from
 1420 an outer scope. In other words, if the local definition can be easily lifted out to top-level, GHC still *does* generalize it.

1421

was almost entirely absent from Stitch. Despite being the author of the singletons library [Eisenberg and Weirich 2012] that automates working with them, I was not tempted to use it here.

10.4 Type errors and editor integration

One aspect in which GHC/Haskell lags behind other dependently typed languages is in its editor integration. Idris, for example, supports interactive type errors, allowing a user to explore typing contexts and other auxiliary information in reading an error [Christiansen 2015]. Idris, Agda, and Coq all allow a programmer to focus on one goal at a time. The closest feature in GHC is its support for typed holes [Gissurarson 2018], where a programmer can replace an expression with an underscore and GHC will tell you the desired type of the expression and suggest type-correct replacements.

The extra features in other language systems would have been helpful, but their lack did not bite in this development. I used typed holes a few times, and I had to comment out code in order to focus on smaller sections, but these were not burdens. Type errors were often screen-filling, but it was easy enough to discern the key details without being overwhelmed. So, while I agree that GHC has room to improve in this regard, its current state is still quite usable.

10.5 Related work

The basic idea embodied in Stitch is not new. Though written before the invention of indexed data types, Pfenning and Lee [1989] consider an encoding of System F in a third-order polymorphic λ -calculus (F_3); only well-typed programs are representable. Their encoding is very much a foreshadowing of more recent papers. Perhaps the first elucidation of the technique of using an indexed AST is by Augustsson and Carlsson [1999], who implemented their interpreter in Cayenne [Augustsson 1998]. The idea was picked up by Pašalić et al. [2002], who use the example of an indexed AST to power the introduction of Meta-D, a language useful for writing indexed ASTs. Other work principally focusing on an index AST includes that by Chen and Xi [2003], which includes an indexed CPS transform, implemented in ATS [Xi 2004]. An implementation of this idea in Haskell is described by Guillemette and Monnier [2008], who embed System F; their encoding is limited by the lack of, e.g., rich kinds in Haskell at the time, and their focus is more on compiler transformations than on type checking. More recently, an indexed AST has been encoded in Agda [Allais et al. 2018, 2017]; the authors' focus in both works cited is in generating correct definitions and proofs without boilerplate. Going beyond just embedding the λ -calculus, Weirich [2017] embeds a richly typed AST for regular expressions in Haskell. The indexed AST idea comes up, in passing or with focus, in many more works beyond these, both in the folklore and in published literature.

The real focus of this paper is not an indexed AST, however; it is to serve as a tutorial to the advanced features of Haskell. In this space, this paper's contribution is indeed novel: to my knowledge, this is the first peer-reviewed tutorial paper aiming to cover these techniques. There is educational material in the folklore and posted online [Ishii 2014; Le 2018]. A tutorial focusing on an indexed AST embedding in Idris [Brady 2013] is part of that language's online documentation [The Idris Team 2017], and Benton et al. [2012] use an indexed AST to explore intrinsic-verification features of Coq. In contrast to those materials, this paper presents its tutorial in the context of a complete software artifact that is a practical tool for teaching the operation of the λ -calculus, with a user-oriented executable. The goal in doing so is to demonstrate that it is indeed possible to build relatively mundane software components, such as a REPL or parser, using fancy types in Haskell—a fact not necessarily yet appreciated by the broader programming language community.

10.6 Conclusion

I have presented Stitch, a simply typed λ -calculus interpreter, amenable for pedagogic use and implemented using an indexed AST. This paper has explored the implementation and described the features of modern Haskell that power the encoding and enable Stitch to be written. I have reported on Haskell's support for richly typed work such as Stitch, concluding that Haskell is ready as a host language for serious work with fancy types.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1704041. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- Guillaume Allais, Robert Atkey, James Chapman, Conor McBride, and James McKinna. 2018. A Type and Scope Safe Universe of Syntaxes with Binding: Their Semantics and Proofs. *Proc. ACM Program. Lang.* 2, ICFP, Article 90 (July 2018), 30 pages. <https://doi.org/10.1145/3236785>
- Guillaume Allais, James Chapman, Conor McBride, and James McKinna. 2017. Type-and-scope Safe Programs and Their Proofs. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs (CPP 2017)*. ACM, New York, NY, USA, 195–207. <https://doi.org/10.1145/3018610.3018613>
- Lennart Augustsson. 1998. Cayenne—a language with dependent types. In *Proc. ACM SIGPLAN International Conference on Functional Programming (ICFP '98)*. ACM, 239–250.
- Lennart Augustsson and Magnus Carlsson. 1999. An exercise in dependent types: A well-typed interpreter. (1999). <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.39.2895&rep=rep1&type=pdf> Unpublished manuscript.
- Nick Benton, Chung-Kil Hur, Andrew J. Kennedy, and Conor McBride. 2012. Strongly Typed Term Representations in Coq. *Journal of Automated Reasoning* 49, 2 (01 Aug 2012), 141–159. <https://doi.org/10.1007/s10817-011-9219-0>
- Gert-Jan Bottu, Georgios Karachalias, Tom Schrijvers, Bruno C. d. S. Oliveira, and Philip Wadler. 2017. Quantified Class Constraints. In *Proceedings of the 10th ACM SIGPLAN International Symposium on Haskell (Haskell 2017)*. ACM, New York, NY, USA, 148–161. <https://doi.org/10.1145/3122955.3122967>
- Edwin Brady. 2013. Idris, a general-purpose dependently typed programming language: Design and implementation. *J. Funct. Prog.* 23 (2013).
- Jacques Carette, Oleg Kiselyov, and Chung-chieh Shan. 2009. Finally tagless, partially evaluated: Tagless staged interpreters for simpler typed languages. *J. Funct. Program.* 19, 5 (Sept. 2009), 509–543.
- Manuel M. T. Chakravarty, Gabriele Keller, and Simon Peyton Jones. 2005. Associated Type Synonyms. In *International Conference on Functional Programming (ICFP '05)*. ACM.
- Chiyan Chen and Hongwei Xi. 2003. Implementing Typeful Program Transformations. In *Proceedings of the 2003 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-based Program Manipulation (PEPM '03)*. ACM, New York, NY, USA, 20–28. <https://doi.org/10.1145/777388.777392>
- Sheng Chen and Martin Erwig. 2016. Principal Type Inference for GADTs. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '16)*. ACM, New York, NY, USA, 416–428. <https://doi.org/10.1145/2837614.2837665>
- James Cheney and Ralf Hinze. 2003. *First-Class Phantom Types*. Technical Report. Cornell University.
- Adam Chlipala. 2008. Parametric Higher-order Abstract Syntax for Mechanized Semantics. In *Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming (ICFP '08)*. ACM, New York, NY, USA, 143–156. <https://doi.org/10.1145/1411204.1411226>
- David Raymond Christiansen. 2015. A Pretty Printer that Says What it Means. Talk, Haskell Implementors Workshop, Vancouver, BC, Canada. (2015). <https://www.youtube.com/watch?v=m7BBccIDXSg>
- Nicolaas Govert de Bruijn. 1972. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *Indagationes Mathematicae (Proceedings)* 75, 5 (1972), 381–392. [https://doi.org/10.1016/1385-7258\(72\)90034-0](https://doi.org/10.1016/1385-7258(72)90034-0)
- Iavor Diatchki. 2014. Type-level literals. GHC User's Manual. (2014). https://downloads.haskell.org/~ghc/latest/docs/html/users_guide/glasgow_exts.html#type-level-literals
- Richard A. Eisenberg. 2016. *Dependent Types in Haskell: Theory and Practice*. Ph.D. Dissertation. University of Pennsylvania.
- Richard A. Eisenberg. 2017. Top-level kind signatures. A GHC Proposal. (2017). <https://github.com/ghc-proposals/ghc-proposals/blob/master/proposals/0036-kind-signatures.rst>

- 1520 Richard A. Eisenberg, Dimitrios Vytiniotis, Simon Peyton Jones, and Stephanie Weirich. 2014. Closed Type Families with
 1521 Overlapping Equations. In *Principles of Programming Languages (POPL '14)*. ACM.
- 1522 Richard A. Eisenberg and Stephanie Weirich. 2012. Dependently Typed Programming with Singletons. In *ACM SIGPLAN
 1523 Haskell Symposium*.
- 1524 Richard A. Eisenberg, Stephanie Weirich, and Hamidhasan Ahmed. 2016. Visible Type Application. In *European Symposium
 1525 on Programming (ESOP) (LNCS)*. Springer-Verlag.
- 1526 Matthias Páll Gissurarson. 2018. *Suggesting Valid Hole Fits for Typed-Holes in Haskell*. Master's thesis. Chalmers University
 1527 of Technology, University of Gothenburg. <https://mpg.is/papers/gissurarson2018suggesting-msc.pdf>
- 1528 Andrew D. Gordon. 1994. A mechanisation of name-carrying syntax up to alpha-conversion. In *Higher Order Logic Theorem
 1529 Proving and Its Applications*, Jeffrey J. Joyce and Carl-Johan H. Seger (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg,
 413–425.
- 1530 Louis-Julien Guillemette and Stefan Monnier. 2008. A type-preserving compiler in Haskell. In *Proc. 13th ACM SIGPLAN
 1531 International Conference on Functional Programming (ICFP '08)*. ACM, 75–86.
- 1532 Adam Gundry. 2013. *Type Inference, Haskell and Dependent Types*. Ph.D. Dissertation. University of Strathclyde.
- 1533 Fritz Henglein. 1993. Type Inference with Polymorphic Recursion. *ACM Trans. Program. Lang. Syst.* 15, 2 (April 1993),
 253–289. <https://doi.org/10.1145/169701.169692>
- 1534 Hiromi Ishii. 2014. Dependent Types in Haskell. School of Haskell blog. (2014). [https://www.schoolofhaskell.com/user/
 1535 konn/prove-your-haskell-for-great-safety/dependent-types-in-haskell](https://www.schoolofhaskell.com/user/konn/prove-your-haskell-for-great-safety/dependent-types-in-haskell)
- 1536 Mark P. Jones. 1995. Functional Programming with Overloading and Higher-Order Polymorphism. In *Advanced Functional
 1537 Programming*, Johan Jeuring and Erik Meijer (Eds.). LNCS, Vol. 925. Springer Verlag.
- 1538 Georgios Karachalias, Tom Schrijvers, Dimitrios Vytiniotis, and Simon Peyton Jones. 2015. GADTs meet their match. In
 1539 *International Conference on Functional Programming (ICFP '15)*. ACM.
- 1540 Edward Kmett. 2012. bound. Haskell package. (2012). <https://github.com/ekmett/bound/>
- 1541 Justin Le. 2018. Introduction to Singletons. (2018). <https://blog.jle.im/entry/introduction-to-singletons-3.html>
- 1542 Daan Leijen. 2001. *Parsec: a fast combinator parser*. Technical Report UU-CS-2001-26. University of Utrecht.
- 1543 Sam Lindley and Conor McBride. 2013. Hasochism: the pleasure and pain of dependently typed Haskell programming. In
 1544 *ACM SIGPLAN Haskell Symposium*.
- 1545 Andres Löh. 2012. lhs2TeX. Haskell package. (2012). <https://www.andres-loeh.de/lhs2tex/>
- 1546 Conor McBride. 2011. The Strathclyde Haskell Enhancement. <https://personal.cis.strath.ac.uk/conor.mcbride/pub/she/>.
 1547 (2011).
- 1548 Stefan Monnier and David Haguenaer. 2010. Singleton types here, singleton types there, singleton types everywhere. In
 1549 *Programming languages meets program verification (PLPV '10)*. ACM.
- 1550 Alan Mycroft. 1984. *Polymorphic type schemes and recursive definitions*. Springer, Berlin, Heidelberg.
- 1551 Emir Pašalić, Walid Taha, and Tim Sheard. 2002. Tagless Staged Interpreters for Typed Languages. In *Proceedings of the
 1552 Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP '02)*. ACM, New York, NY, USA, 218–229.
 1553 <https://doi.org/10.1145/581478.581499>
- 1554 Simon Peyton Jones. 2003. Wearing the Hair Shirt: A Retrospective on Haskell. Invited talk at POPL. (2003).
- 1555 Simon Peyton Jones, Andrew Tolmach, and Tony Hoare. 2001. Playing by the Rules: Rewriting as a practical optimisation
 1556 technique in GHC. In *Proceedings of the Haskell Workshop*.
- 1557 Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. 2007. Practical type inference for arbitrary-
 1558 rank types. *Journal of Functional Programming* 17, 1 (Jan. 2007).
- 1559 Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Geoffrey Washburn. 2006. Simple unification-based type
 1560 inference for GADTs. In *International Conference on Functional Programming (ICFP '06)*. ACM.
- 1561 Simon Peyton Jones, Geoffrey Washburn, and Stephanie Weirich. 2004. *Wobbly types: type inference for generalised algebraic
 1562 data types*. Technical Report MS-CIS-05-26. University of Pennsylvania.
- 1563 Simon Peyton Jones, Stephanie Weirich, Richard A. Eisenberg, and Dimitrios Vytiniotis. 2016. A reflection on types. In *A
 1564 list of successes that can change the world*. Springer. A festschrift in honor of Phil Wadler.
- 1565 Frank Pfenning and Conal Elliott. 1988. Higher-order Abstract Syntax. In *Proceedings of the ACM SIGPLAN 1988 Conference
 1566 on Programming Language Design and Implementation (PLDI '88)*. ACM, New York, NY, USA, 199–208. [https://doi.org/10.
 1567 1145/53990.54010](https://doi.org/10.1145/53990.54010)
- 1568 Frank Pfenning and Peter Lee. 1989. LEAP: A language with eval and polymorphism. In *TAPSOFT '89*, J. Diaz and F. Orejas
 (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 345–359.
- 1569 Matthew Pickering, Gergő Erdi, Simon Peyton Jones, and Richard A. Eisenberg. 2016. Pattern Synonyms. In *ACM SIGPLAN
 1570 Haskell Symposium*. ACM.
- 1571 Benjamin C. Pierce. 2002. *Types and Programming Languages*. MIT Press, Cambridge, MA.
- 1572 Andrew M. Pitts. 2003. Nominal logic, a first order theory of names and binding. *Information and Computation* 186, 2 (2003),
 165 – 193. [https://doi.org/10.1016/S0890-5401\(03\)00138-X](https://doi.org/10.1016/S0890-5401(03)00138-X) Theoretical Aspects of Computer Software (TACS 2001).

- 1569 Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Gilles Peskine, Thomas Ridge, Susmit Sarkar, and Rok Strniša. 2010.
1570 Ott: Effective tool support for the working semanticist. *Journal of Functional Programming* 20, 1 (Jan. 2010).
- 1571 Walid Taha, Henning Makhholm, and John Hughes. 2001. Tag elimination and Jones-optimality. In *Programs as Data Objects*,
1572 Olivier Danvy and Andrzej Filinski (Eds.). LNCS, Vol. 2053. Springer Verlag.
- 1573 The Idris Team. 2017. Example: The Well-Typed Interpreter. The Idris Tutorial. (2017). <http://docs.idris-lang.org/en/latest/tutorial/interp.html>
- 1574 Dimitrios Vytiniotis, Simon Peyton Jones, and Tom Schrijvers. 2010. Let Should Not Be Generalized. In *Types in Language*
1575 *Design and Implementation (TLDI '10)*. ACM.
- 1576 Dimitrios Vytiniotis, Simon Peyton Jones, Tom Schrijvers, and Martin Sulzmann. 2011. OUTSIDEIN(X): Modular Type
1577 Inference with Local Assumptions. *Journal of Functional Programming* 21, 4-5 (Sept. 2011).
- 1578 Philip Wadler. 2003. A Prettier Printer. In *The Fun of Programming: Essays in Honor of Richard Bird*, Jeremy Gibbons and
1579 Oege De Moor (Eds.). Palgrave.
- 1580 Stephanie Weirich, Justin Hsu, and Richard A. Eisenberg. 2013. System FC with Explicit Kind Equality. In *International*
1581 *Conference on Functional Programming (ICFP '13)*. ACM.
- 1582 Stephanie Weirich. 2017. The Influence of Dependent Types. Keynote, POPL '17. (2017). https://www.youtube.com/watch?v=rflCw9bT4_0
- 1583 Stephanie Weirich, Antoine Voizard, Pedro Henrique Azevedo de Amorim, and Richard A. Eisenberg. 2017. A Specification
1584 for Dependent Types in Haskell. *Proc. ACM Program. Lang.* 1, ICFP, Article 31 (Aug. 2017), 29 pages. <https://doi.org/10.1145/31110275>
- 1585 Stephanie Weirich, Brent A. Yorgey, and Tim Sheard. 2011. Binders Unbound. In *Proceedings of the 16th ACM SIGPLAN*
1586 *International Conference on Functional Programming (ICFP '11)*. ACM, New York, NY, USA, 333–345. <https://doi.org/10.1145/2034773.2034818>
- 1587 Hongwei Xi. 2004. Applied Type System. In *Types for Proofs and Programs*, Stefano Berardi, Mario Coppo, and Ferruccio
1588 Damiani (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 394–408.
- 1589 Hongwei Xi, Chiyan Chen, and Gang Chen. 2003. Guarded recursive datatype constructors. In *Principles of Programming*
1590 *Languages (POPL '03)*. ACM.
- 1591 Brent A. Yorgey, Stephanie Weirich, Julien Cretin, Simon Peyton Jones, Dimitrios Vytiniotis, and José Pedro Magalhães.
1592 2012. Giving Haskell a promotion. In *Types in Language Design and Implementation (TLDI '12)*. ACM.
- 1593 Vladislav Zavialov. 2018. Remove the * kind syntax. A GHC Proposal. (2018). <https://github.com/ghc-proposals/ghc-proposals/blob/master/proposals/0030-remove-star-kind.rst>
- 1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617