

Structural Refinement Types

David Binder

Department of Computer Science
University of Tübingen
Tübingen, Germany
david.binder@uni-tuebingen.de

Klaus Ostermann

Department of Computer Science
University of Tübingen
Tübingen, Germany
klaus.ostermann@uni-tuebingen.de

Ingo Skupin

Department of Computer Science
University of Tübingen
Tübingen, Germany
skupin@informatik.uni-tuebingen.de

David Läwen

Department of Computer Science
University of Tübingen
Tübingen, Germany
david.laewen@gmx.de

Abstract

Static types are a great form of lightweight static analysis. But sometimes a type like `List` is too coarse – we would also like to work with its *refinements* like non-empty lists, or lists containing exactly 42 elements. Dependent types allow for this, but they impose a heavy proof burden on the programmer. We want the checking and inference of refinements to be fully automatic.

Existing approaches to refinement types support extensive automation, but they make various trade-offs. In the original system of Freeman and Pfenning [6], the programmer was required to manually specify all refinements they want to use. We only require the user to specify the original data type, but not its refinements, which are discovered automatically. With liquid refinement types [17, 18], an external solver is used to solve the proof obligations for the programmer. By contrast, we use only familiar techniques from constraint-based type inference. The system of Jones and Ramsay [12] is very similar in spirit to our system, but the refinements they allow are less expressive. They support refinements which can be expressed as the removal of constructors from the definition of a type, but require these removals to be hereditary. As such, they cannot express the type of non-empty lists, which requires the removal of the `Nil` constructor only at the top level. Our system lifts this limitation, supporting all refinements that can be expressed as a regular sublanguage of the original type.

In this article we present a simple refinement type system and inference algorithm which uses only variants of familiar concepts from constraint-based type inference. Concretely, we build on the algebraic subtyping approach [3, 4, 14] and extend it with typing rules which combine properties of nominal and structural type systems in a novel way. Despite the simplicity of our approach, the resulting type system is very expressive and allows to specify and infer non-trivial properties of programs.

Keywords: Structural Types, Nominal Types, Refinement Types, Algebraic Subtyping

ACM Reference Format:

David Binder, Ingo Skupin, Klaus Ostermann, and David Läwen. 2022. Structural Refinement Types. In *Proceedings of Workshop on Type-Driven Development (TyDe’22)*. ACM, New York, NY, USA, 13 pages.

1 Introduction

Consider the following simple predecessor function on Peano numbers:

```
def pred := λx.case x of {S(n) ⇒ n}
```

According to most type systems and implementations, this function is partial. How programming languages handle this partiality may differ; some implementations will crash at runtime while others generate a warning or an error at compile-time. But there is another alternative: This function is a perfectly well-behaved total function on the domain of *non-zero* Peano numbers. Expressing and inferring such *refinements* of data types is the subject of this article.

We are certainly not the first to have made this observation. More fine-grained types of this sort are commonly known as *refinement types* [6]. In distinction to dependent types, which can express almost arbitrary subsets of types, refinement types are much more lightweight. A heavy emphasis on automation makes it so that little to no additional annotations are required by the programmer to benefit from their expressiveness.

As soon as we introduce refinement types into a type system, some notion of subtyping comes into play: If a function expects a number as an argument, it is clearly also valid to pass a subtype such as a non-zero number. So we need subtyping, but the design of a global type inference algorithm with nice properties such as principal types for a system with both subtyping and parametric polymorphism is known to be a hairy problem. Luckily for us, we don’t have to solve this problem ourselves, since Dolan and Mycroft [3, 4] as well as Parreaux [14] have already done so for us. The *algebraic subtyping* approach they have developed can be used

111 as a blueprint to design a type system which fulfills the de-
 112 sired properties, similar to how the basic Hindley-Milner
 113 algorithm can be used as a blueprint for further type system
 114 extensions.

115 While experimenting with an implementation of the al-
 116 gebraic subtyping type inference algorithm, we made the
 117 following discovery, which we present as our central contri-
 118 bution in this paper: We implemented both ordinary nominal
 119 algebraic data types, as well as purely structural data types
 120 in the form of polymorphic variants [7]. We then realized
 121 that if we combine the typing rules for polymorphic vari-
 122 ants with the typing rules for nominal data types, we get a
 123 form of refinement types for free! Before we go into further
 124 detail, let us briefly review the spectrum of nominal and
 125 structural type systems, one of the principal ways in which
 126 type systems can be classified.

127 A nominal type system can be recognized by the pres-
 128 ence of a sublanguage for specifying types and the relation-
 129 ships between them. In inheritance-based, object-oriented
 130 languages, for example, this sublanguage is used to spec-
 131 ify classes and their subtyping relationships, while in data-
 132 oriented languages, it is used to specify a wide variety of
 133 algebraic data types, such as enums and structs. By contrast,
 134 structural type systems consider the *structure* of types to
 135 reason about type compatibility. In systems for structural
 136 records and polymorphic variants [7, 8], for example, the
 137 types reflect which fields are present in records, or which
 138 variants may be present in a sum type. Nominal subtyping
 139 can be seen as a subset of structural subtyping in which a
 140 nominal subtyping declaration must be present in addition
 141 to a structural subtyping relation, which on one hand facili-
 142 tates unplanned extensibility but on the other hand leads to
 143 undesirable subtype relationships [13].

144 In practice, a lot of systems combine nominal and struc-
 145 tural aspects within the same system. To illustrate this, con-
 146 sider the following examples:

$$\{ a = 2, b = \text{True} \} : \{ a : \mathbb{N}, b : \mathbb{B} \}$$

$$\text{\`IoError}(5) : \langle \text{\`IoError}(\mathbb{N}) \rangle$$

153 In the first line, both fields a and b are reflected in the struc-
 154 tural record type. The fields themselves, on the other hand,
 155 are typed using nominal types \mathbb{N} and \mathbb{B} . The second example
 156 illustrates the use of polymorphic variants (or structural sum
 157 types), using syntax inspired by OCaml. The constructor
 158 `\`IoError` is not part of some algebraic data type declared
 159 elsewhere in the program or the standard library. It is intro-
 160 duced in this term, so its name and argument type is reflected
 161 in the inferred type $\langle \text{\`IoError}(\mathbb{N}) \rangle$.

162 The combination of such structural variants and records
 163 with equi-recursive types is surprisingly expressive. Equi-re-
 164 cursive types have the form $\mu\alpha.\tau$ and are for typechecking

166 purposes indistinguishable from their unfolding $\tau[\mu\alpha.\tau/\alpha]$ ¹.
 167 As an example for this expressivity, consider the following
 168 function:

```
169 def isEven := λx.case x of {Z ⇒ True;
170                               `S(x) ⇒ not(isEven(x))}
```

171 In a system which can infer equi-recursive types, this func-
 172 tion has the inferred type $\mu\alpha.\langle \text{\`Z}, \text{\`S}(\alpha) \rangle \rightarrow \mathbb{B}$. This type
 173 correctly reflects both the constructors matched against and
 174 the recursive nature of the function. But while this type is
 175 hardly more informative than the type $\mathbb{N} \rightarrow \mathbb{B}$, the following
 176 example shows how structural types in combination with
 177 equi-recursive types can give more information than the
 178 corresponding nominal types can. Consider the following
 179 function which doubles its argument:

```
182 def double := λx.case x of {Z ⇒ `Z;
183                               `S(x) ⇒ `S(`S(double(x)))}
```

184 The inferred type of this function is $\mu\alpha.\langle \text{\`Z} \mid \text{\`S}(\alpha) \rangle \rightarrow$
 185 $\mu\beta.\langle \text{\`Z} \mid \text{\`S}(\text{\`S}(\beta)) \rangle$. This type reveals that the function
 186 only returns even numbers!

187 But, alas, not everything is fine with these structural types.
 188 Consider the predecessor function which we used as our first
 189 example, but this time with a polymorphic variant:

```
192 def pred := λx.case x of {`S(n) ⇒ n}
```

193 Its inferred type scheme is the rather surprisingly polymor-
 194 phic $\forall\alpha.\langle \text{\`S}(\alpha) \rangle \rightarrow \alpha$. This type states that the predecessor
 195 function can be applied to any typeable term as long as it is
 196 wrapped in the successor constructor, even such nonsensical
 197 terms as `\`S(True)`. Under the typing rules of a structural
 198 system this is correct and doesn't violate the soundness of
 199 the type theory. But it is clearly confusing for a programmer
 200 who expects that the application of the successor constructor
 201 to a boolean should result in a type error.

202 The main idea presented in this paper is that if we combine
 203 the ordinary typing rules for data types with the typing rules
 204 for polymorphic variants, we obtain an interesting refine-
 205 ment type system. We call these types *structural refinement*
 206 *types*. In this system, structural types are used as the refine-
 207 ments of nominal types which are declared by the user. For
 208 instance, given the nominal type \mathbb{N} of Peano numbers, we
 209 can use various structural refinement types, such as

```
211 all natural numbers : μα.⟨ ℕ | Z, S(α) ⟩
212 non-zero numbers : ⟨ ℕ | S(μα.⟨ ℕ | Z, S(α) ⟩) ⟩
213 even numbers : μα.⟨ ℕ | Z, S(⟨ ℕ | S(a) ⟩) ⟩
```

214
 215
 216
 217
 218
 219
 220
 221
 222
 223
 224
 225
 226
 227
 228
 229
 230
 231
 232
 233
 234
 235
 236
 237
 238
 239
 240
 241
 242
 243
 244
 245
 246
 247
 248
 249
 250
 251
 252
 253
 254
 255
 256
 257
 258
 259
 260
 261
 262
 263
 264
 265
 266
 267
 268
 269
 270
 271
 272
 273
 274
 275
 276
 277
 278
 279
 280
 281
 282
 283
 284
 285
 286
 287
 288
 289
 290
 291
 292
 293
 294
 295
 296
 297
 298
 299
 300
 301
 302
 303
 304
 305
 306
 307
 308
 309
 310
 311
 312
 313
 314
 315
 316
 317
 318
 319
 320
 321
 322
 323
 324
 325
 326
 327
 328
 329
 330
 331
 332
 333
 334
 335
 336
 337
 338
 339
 340
 341
 342
 343
 344
 345
 346
 347
 348
 349
 350
 351
 352
 353
 354
 355
 356
 357
 358
 359
 360
 361
 362
 363
 364
 365
 366
 367
 368
 369
 370
 371
 372
 373
 374
 375
 376
 377
 378
 379
 380
 381
 382
 383
 384
 385
 386
 387
 388
 389
 390
 391
 392
 393
 394
 395
 396
 397
 398
 399
 400
 401
 402
 403
 404
 405
 406
 407
 408
 409
 410
 411
 412
 413
 414
 415
 416
 417
 418
 419
 420
 421
 422
 423
 424
 425
 426
 427
 428
 429
 430
 431
 432
 433
 434
 435
 436
 437
 438
 439
 440
 441
 442
 443
 444
 445
 446
 447
 448
 449
 450
 451
 452
 453
 454
 455
 456
 457
 458
 459
 460
 461
 462
 463
 464
 465
 466
 467
 468
 469
 470
 471
 472
 473
 474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485
 486
 487
 488
 489
 490
 491
 492
 493
 494
 495
 496
 497
 498
 499
 500

With these refinement types, the predecessor function can be typed with $\langle \mathbb{N} \mid S(\mu\alpha.\langle \mathbb{N} \mid Z, S(\alpha) \rangle) \rangle \rightarrow \mu\alpha.\langle \mathbb{N} \mid Z, S(\alpha) \rangle^2$. That is, the function expects a non-zero number as an argument and returns some natural number.

1.1 Overview

The rest of this paper is structured as follows:

- In [Section 2](#) we present our central ideas using a simplified version of our type system.
- In [Section 3](#) we present the full declarative type system with user defined and parameterized data types.
- In [Section 4](#) we present the type inference algorithm, and the algorithms used to simplify types. The type inference algorithm uses a variant of the biunification algorithm introduced by Dolan and Mycroft [3, 4] with some modifications inspired by Parreaux [14]. Type simplification is achieved by encoding types in finite automata and using familiar simplification techniques from automata theory.
- We discuss related work in [Section 5](#), future work in [Section 6](#) and conclude in [Section 7](#).

2 The Main Idea

Our main idea is to build upon the algebraic subtyping approach [3, 4, 14] and to extend it with types which combine the typing rules for nominal and structural types. In [Section 3](#) and [Section 4](#) we will present these rules with all the gory details. Since we show how to implement structural refinement types for arbitrary user-defined data types, the resulting rules are quite complex. In order to make them more palatable, we specialize the rules in this section to two examples: Peano numbers and lists. In [Section 2.1](#) we show how to work with refinements of Peano numbers, since they are the simplest example involving recursive types. We show how to deal with refinements of lists in [Section 2.2](#). Parameterized types like lists pose an interesting design question. Should we compute refinements of the spine of the list and its elements separately, or together? We chose to refine them separately, and motivate that choice in that subsection.

2.1 Peano Numbers

Natural numbers \mathbb{N} can be represented by a data type with two constructors: zero (Z) and successor (S). The typing rules for the constructors and the pattern match are familiar:

$$\frac{}{\Gamma \vdash Z : \mathbb{N}} Z_{Nominal} \quad \frac{\Gamma \vdash e : \mathbb{N}}{\Gamma \vdash S(e) : \mathbb{N}} S_{Nominal}$$

$$\frac{\Gamma \vdash e : \mathbb{N} \quad \Gamma \vdash e_Z : \tau \quad \Gamma, x : \mathbb{N} \vdash e_S : \tau}{\Gamma \vdash \text{case } e \text{ of } \{Z \Rightarrow e_Z, S(x) \Rightarrow e_S\} : \tau} \text{CASE}_{Nominal}^{\mathbb{N}}$$

²There are many ways in which the presentation of these refinement types can be improved when shown to the user. We use a presentation which is close to the form used in the formalization.

Nothing about these rules is surprising. But note that even in these very simple rules one of the essential characteristics of all static analyses is already present. The types that we use are a conservative approximation, as the single type \mathbb{N} is used for all natural numbers.

In contrast to this, let us consider what would happen if we used polymorphic variants [7, 8] for natural numbers. The typing rule for zero and the successor would then be the following:

$$\frac{}{\Gamma \vdash `Z : \langle `Z \rangle} Z_{Structural} \quad \frac{\Gamma \vdash e : \sigma}{\Gamma \vdash `S(e) : \langle `S(\sigma) \rangle} S_{Structural}$$

If we compare them to the previous rules, we can make the following observations: The inferred type $\langle `Z \rangle$ for zero is much more informative, it is even specific enough to deduce the only possible inhabitant, namely the number 0! And the same holds true for the rule for the successor; from the type in the conclusion of the rule we can deduce that the outermost constructor must be $`S$. But the rule $S_{Structural}$ is also somewhat disappointing: It does not impose any restriction on the type of the term that $`S$ is applied to. This is why a term like $`S(\text{True})$ is typeable with this rule.

Similar to how the typing rules for constructors reflect the constructor used to construct the term in the type, the rule for pattern matches reflects the constructors which are matched against. For example, a pattern match which only matches against the constructor $`S$ is typed as follows:

$$\frac{\Gamma \vdash e : \langle `S(\tau) \rangle \quad \Gamma, x : \tau \vdash e_S : \rho}{\Gamma \vdash \text{case } e \text{ of } \{`S(x) \Rightarrow e_S\} : \rho} \text{CASE}_{Structural}^S$$

In this rule we don't require e to be a natural number. We only place the minimal requirement on the type e , namely to be wrapped with the constructor $`S$.

So what is the result when we combine the nominal and structural rules above? In our system, the typing rules for zero and the successor look as follows:

$$\frac{}{\Gamma \vdash Z : \langle \mathbb{N} \mid Z \rangle} Z_{Refinement} \quad \frac{\Gamma \vdash e : \sigma \quad \sigma <: \mu\alpha.\langle \mathbb{N} \mid Z, S(\alpha) \rangle}{\Gamma \vdash S(e) : \langle \mathbb{N} \mid S(\sigma) \rangle} S_{Refinement}$$

In the typing rule $S_{Refinement}$ we have combined the properties of both systems. In the conclusion of the rule, we can still deduce from the type $\langle \mathbb{N} \mid S(\sigma) \rangle$ which constructor was used to build the term. And in the second premise we still put a constraint on the permissible arguments of S , namely that the argument to S must be some subtype of the Peano natural numbers, i.e. a refinement of \mathbb{N} . The situation is similar in the case of pattern matches:

$$\frac{\langle \mathbb{N} \mid \emptyset \rangle <: \tau <: \mu\alpha.\langle \mathbb{N} \mid Z, S(\alpha) \rangle}{\Gamma \vdash e : \langle \mathbb{N} \mid S(\tau) \rangle} \quad \frac{\Gamma, x : \tau \vdash e_S : \rho}{\Gamma \vdash \text{case } e \text{ of } \{S(x) \Rightarrow e_S\} : \rho} \text{CASE}_{Refinement}^S$$

We require the term e to be of type $\langle \mathbb{N} \mid S(\tau) \rangle$, since we only match against the successor constructor. But we also require the argument of S to be *some* natural number, which we express here with a lower and an upper bound on τ . We might learn more about the requirements that τ must satisfy in the body e_S , for example that τ must be non-zero, too.

2.2 Lists

In the previous section we used the example of natural numbers, whose definition in the formal syntax of Section 3 will look like this:

```
data Nat : ( ; ) → * { Z() : Nat( ; ), S(rec@( ; )) : Nat( ; ) }
```

After the name of the type constructor `Nat` we have to specify its kind. Since natural numbers are not parameterized, the kind is simply $(;) \rightarrow *$, which is isomorphic to the kind $*$ of inhabited types. The recursive occurrence in the argument of the constructor S is written using the special `rec` symbol, which has to be applied to zero type arguments, using the syntax `rec@(;)`. This syntax makes more sense once we see the definition of a parameterized type. The type of lists of some element type α is defined as follows:

```
data List : (* ; ) → * {
  ∀α. Nil() : List(α ; ),
  ∀α. Cons(α, rec@(α ; )) : List(α ; ) }
```

This definition introduces the type constructor `List` of kind $(* ;) \rightarrow *$. In the algebraic subtyping system we have to keep covariant and contravariant parameters of a type strictly separate. The element type of a list is a covariant argument, and covariant arguments are written to the left of the semicolon in the kind of `List`.

There is one central design question concerning parameterized types like lists. The problem can be illustrated with the list `Cons(True, Cons(False, Nil))`. What should be the type of this list? As we see it, there are two possible choices:

1. We infer the refinement as if the programmer had written a data type declaration of lists of booleans. We call this the *monomorphising* approach. With this approach, we infer the type of a two-element list which contains the term `True` at the first position and the term `False` at the second position.
2. We infer separate refinements for the spine of the list and the elements of the list. Using this approach, we infer the type of a two element list (that is, a refinement on lists) whose elements are among the set containing both `True` and `False` (that is, a refinement of booleans). On a technical level, this corresponds to introducing a single unification variable for all of the elements of the list.

The types that are inferred using the first approach are more precise, but this precision comes at a cost. Type inference has to keep track of a lot more information, making

it much harder to scale to realistic programs. Secondly, the inferred types are much harder to decipher for the user, and the approach is overall less modular. For these reasons we have decided to specify and implement the second approach.

An example for functions on lists can be found in Figure 1. The `mapNonEmpty` function operates on the type `NonEmpty` of non-empty lists, which is a subtype of the type `FullList` of list of arbitrary length. This illustrates how we can keep reasoning about the *shape* of the list separate from reasoning about the lists *contents*: `mapNonEmpty` keeps track of the shape and allows us to recover the shape of the input in the output, while at the same time remaining parametric over the types of the elements of the lists involved. This allows us to use its output as an input to the `max` function. At the same time, we would still be able to use its output in any parametric function on lists like a `listLength` function since `NonEmpty` is a subtype of `FullList`.

2.3 The Expressivity of Structural Refinement Types

How expressive are the structural refinement types that we present here? In their paper, Freeman and Pfenning [6] already made the observation that they can only express refinements that correspond to a regular sublanguage of the original type, and we have the same restriction in our system. For example, it is possible to express the refinement type of even natural numbers, but it is not possible to express the refinement type of natural numbers that are prime.

This restriction of the expressive power is essential in several different respects. The favorable closure properties of regular languages are crucial when we compute and simplify unions and intersections of refinement types. The close correspondence between types and finite automata, which explains the restriction to regular sublanguages, is also the basis of the simplification algorithms described by Dolan [3], which we modify for our purposes in Section 4.4.

3 Formalization

In this section we describe the declarative type system; the type inference algorithm will then be described in Section 4. We introduce terms in Section 3.1, kinds in Section 3.2 and most types in Section 3.3. We introduce the syntax and rules of structural refinement types in Section 3.4. We use the notation \bar{e} to denote a (possibly empty) list of elements e , following the conventions of Igarashi et al. [10].

3.1 Terms

The term system, presented in Figure 2, is entirely standard for a language with functions and data types. Terms e can be variables x , which are taken from the set of term variables `VAR`. The two syntactic forms for functions are lambda abstractions $\lambda x.e$ and function applications $e e$. Elements of an algebraic data type can be constructed with constructors applied to arguments $C(\bar{e})$, where the names of constructors

```

441 data List : (*;) → * {
442   ∀α.Nil() : List(α;),
443   ∀α.Cons(α, rec@(α;)) : List(α; ) }
444
445 type FullList : (*;) → * :=
446   μδ.⟨ List(γ; ) | Nil, Cons(γ, δ) ⟩
447
448 map : (α → β) →
449   FullList@(α; ) → FullList@(β; )
450
451 map := λf xs. case xs of {
452   Nil ⇒ Nil;
453   Cons(y, ys) ⇒ Cons(f y, map f ys) }
454
455 type NonEmpty : (*;) → * :=
456   ⟨ List(α; ) | Cons(α, FullList@(α;)) ⟩
457
458 mapNonEmpty : (α → β) →
459   NonEmpty@(α; ) → NonEmpty@(β; )
460
461 mapNonEmpty := λf xs. case xs of {
462   Cons(y, ys) ⇒ Cons(f y, map f ys) }
463
464 max : NonEmpty@(ℕ; ) → ℕ
465 max := λns. case ns of {
466   Cons(m, ms) ⇒ case ms of {
467     Nil ⇒ m
468     Cons(o, os) ⇒ case (max Cons(o, os)) ≤ m of {
469       True ⇒ m
470       False ⇒ max Cons(o, os) } } }
471
472 maxLength : NonEmpty@(String; ) → ℕ
473 maxLength := λss. max (mapNonEmpty length ss)
474
475
476
477

```

Figure 1. Mapping on non-empty lists

are from the set `CTORNAME`. In order to deconstruct an element of an algebraic data type one uses a pattern match `case e of { $\overline{C(\bar{x})} \Rightarrow e$ }`.

```

484   x, y ∈ VAR   C ∈ CTORNAME
485
486 e, ei ::= x | λx.e | e e
487         |  $\overline{C(\bar{e})}$  | case e of {  $\overline{C(\bar{x})} \Rightarrow e$  }
488
489

```

Figure 2. The syntax of terms.

We do not specify any operational semantics for the terms of Figure 2. The type theory we present is equally suitable for both strict and non-strict evaluation orders.

3.2 Kinds

We formalize a system with higher kinds whose syntax is given in Figure 3:

$$\kappa ::= * \mid (\bar{\kappa}; \bar{\kappa}') \rightarrow * \quad \text{Kinds}$$

Figure 3. The syntax of kinds.

The kind $*$ classifies inhabited types. Higher kinds have the form $(\bar{\kappa}; \bar{\kappa}') \rightarrow *$. They have two lists of arguments $\bar{\kappa}$ and $\bar{\kappa}'$ and always return the kind $*$. The two argument lists are needed to keep track of variance; the first list stands for covariant arguments and the second list stands for contravariant arguments. As an example, the list type constructor has kind $(*;) \rightarrow *$, taking one covariant argument of kind $*$ and no contravariant arguments.

3.3 Types

We will first discuss the types of the core system, which are given in Figure 4. Type variables are taken from a set `TYVAR` and occur as both skolem variables α and unification variables $\alpha^?$ once we consider type inference in Section 4. The function type $\sigma \rightarrow \tau$ is the only built-in concrete type, all other concrete types are specified by the programmer. We postpone the discussion of user-specified types until Section 3.4. The application of a higher-kinded type constructor to type arguments is written $\tau@(\bar{\sigma}; \bar{\rho})$; in this example the type constructor τ is applied to the covariant arguments $\bar{\sigma}$ and the contravariant arguments $\bar{\rho}$. Since we are in a subtyping system we also have union and intersection types $\sigma \sqcup \tau$ and $\sigma \sqcap \tau$, as well as top and bottom types \top and \perp . Equi-recursive types are written $\mu\alpha.\tau$ and are indistinguishable from their unfolding $\tau[\mu\alpha.\tau/\alpha]$.

$$\alpha, \beta, \gamma, \delta, \rho \in \text{TYVAR}$$

σ, τ	::= α	Skolem variable
	$\alpha^?$	Unification variable
	$\tau \rightarrow \tau$	Function type
	$\tau@(\bar{\tau}; \bar{\tau}')$	Type Application
	$\tau \sqcup \tau$ $\tau \sqcap \tau$	Union, Intersection
	\top \perp	Top and Bottom
	$\mu\alpha.\tau$	Recursive type
Δ	::= ϵ $\alpha : \kappa, \Delta$	Kind environments
Σ	::= ϵ $\sigma <: \tau, \Sigma$ $\triangleright (\sigma <: \tau), \Sigma$	Hypotheses context

Figure 4. Syntax of types for the core system.

Since we have type constructors with higher kinds in our system, we have to specify rules for checking the kinds of types. These rules are given in Figure 5. The judgement

$\Delta \vdash \tau : \kappa$ states that the type τ has kind κ in a kind environment Δ which assigns kinds to type variables. The syntax of kind environments is specified in Figure 4.

$$\begin{array}{c}
 \boxed{\text{Kinding rules: } \Delta \vdash \tau : \kappa} \\
 \frac{\Delta(\alpha) : \kappa}{\Delta \vdash \alpha : \kappa} \text{K-VAR} \quad \frac{\Delta, \alpha : \kappa \vdash \tau : \kappa}{\Delta \vdash \mu\alpha.\tau : \kappa} \text{K-MU} \\
 \frac{}{\Delta \vdash \top : \kappa} \text{K-TOP} \quad \frac{}{\Delta \vdash \perp : \kappa} \text{K-BOT} \\
 \frac{\Delta \vdash \sigma : \kappa \quad \Delta \vdash \tau : \kappa}{\Delta \vdash \sigma \sqcup \tau : \kappa} \text{K-UNION} \\
 \frac{\Delta \vdash \sigma : \kappa \quad \Delta \vdash \tau : \kappa}{\Delta \vdash \sigma \sqcap \tau : \kappa} \text{K-INTER} \\
 \frac{\Delta \vdash \sigma : * \quad \Delta \vdash \tau : *}{\Delta \vdash \sigma \rightarrow \tau : *} \text{K-FUN} \\
 \frac{\Delta \vdash \sigma : (\overline{\kappa}; \overline{\kappa}') \rightarrow * \quad \Delta \vdash \tau : \kappa \quad \Delta \vdash \tau' : \kappa'}{\Delta \vdash \sigma @ (\overline{\tau}; \overline{\tau}') : *} \text{K-TYAPP} \\
 \frac{\Delta \vdash \tau : (;) \rightarrow *}{\Delta \vdash \tau : *} \text{K-STAR}
 \end{array}$$

Figure 5. Kinding rules.

Note that in the kinding rules in Figure 5 the function type can only be applied to types of the inhabited kind $*$, but all other constructions like unions, intersections, the top and bottom type and recursive types can occur at arbitrary kinds. This is an essential prerequisite for the correct treatment of structural refinement types: For example, we need to be able to talk about the union of the type constructors of lists and non-empty lists. The K-STAR rule witnesses the fact that a higher order kind with no parameters is the same as kind $*$.

The declarative typing rules for the core system are presented in Figure 6. The presentation of these rules is adapted from that in [14]. The judgement $\Gamma \vdash e : \tau$ states that the expression e has type τ in the variable context Γ which assigns types to term variables. Note that in this form the rule T-SUB is not syntax directed, since it allows to change the type of an expression at an arbitrary point in a typing derivation.

$$\begin{array}{c}
 \boxed{\text{Typing rules: } \Gamma \vdash e : \tau} \\
 \frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau} \text{T-VAR} \quad \frac{\Gamma \vdash e_1 : \sigma \rightarrow \tau \quad \Gamma \vdash e_2 : \sigma}{\Gamma \vdash e_1 e_2 : \tau} \text{T-APP} \\
 \frac{\Gamma, x : \sigma \vdash e : \tau}{\Gamma \vdash \lambda x.e : \sigma \rightarrow \tau} \text{T-LAM} \quad \frac{\Gamma \vdash e : \sigma \quad \sigma <: \tau}{\Gamma \vdash e : \tau} \text{T-SUB}
 \end{array}$$

Figure 6. Declarative typing rules

The last set of rules concern the formalization of the subtyping lattice, and are presented in Figure 7. The judgement $\Sigma \vdash \sigma <: \tau$ says that σ is a subtype of τ under the assumption that the hypotheses in Σ hold. There are two different

kinds of hypotheses, guarded hypotheses $\triangleright(\sigma <: \tau)$ and unguarded hypotheses $\sigma <: \tau$. Unguarded hypotheses can be used directly using the rule S-HYP, but guarded hypotheses must first be unlocked with the \triangleleft operation before they can be used. Hypotheses are unguarded every time we pass a type constructor in the rule S-FUN. The additional complication of hypotheses contexts is necessary for the correct treatment of recursive types. Without them, we would not be able to determine that $\mu\alpha.\mathbf{Bool} \rightarrow \alpha$ is a subtype of $\mu\alpha.\mathbf{Bool} \rightarrow \mathbf{Bool} \rightarrow \alpha$, even though they have the same infinite unfolding. This example is well-explained in [14].

$$\begin{array}{c}
 \boxed{\text{Subtyping rules: } \Sigma \vdash \sigma <: \tau} \\
 \frac{}{\vdash \tau <: \tau} \text{S-REFL} \quad \frac{H \in \Sigma}{\Sigma \vdash H} \text{S-HYP} \quad \frac{\Sigma, \triangleright H \vdash H}{\Sigma \vdash H} \text{S-ASSUM} \\
 \frac{\vdash H}{\Sigma \vdash H} \text{S-WEAKEN} \quad \frac{}{\vdash \tau <: \top} \text{S-TOP} \quad \frac{}{\vdash \perp <: \tau} \text{S-BOT} \\
 \frac{\Sigma \vdash \tau <: \tau' \quad \Sigma \vdash \tau' <: \tau''}{\Sigma \vdash \tau <: \tau''} \text{S-TRANS} \\
 \frac{\triangleleft \Sigma \vdash \sigma' <: \sigma \quad \triangleleft \Sigma \vdash \tau <: \tau'}{\Sigma \vdash \sigma \rightarrow \tau <: \sigma' \rightarrow \tau'} \text{S-FUN} \\
 \frac{\Sigma \vdash \tau <: \tau' \quad \Sigma \vdash \overline{\sigma} <: \overline{\sigma'} \quad \Sigma \vdash \overline{\rho} <: \overline{\rho'}}{\Sigma \vdash \tau @ (\overline{\sigma}; \overline{\rho}) <: \tau' @ (\overline{\sigma}'; \overline{\rho}')} \text{S-APP} \\
 \frac{\forall j \exists i : \Sigma \vdash \tau_i <: \sigma_j}{\Sigma \vdash \prod_i \tau_i <: \prod_j \sigma_j} \text{S-MEET} \\
 \frac{\forall i \exists j : \Sigma \vdash \tau_i <: \sigma_j}{\Sigma \vdash \sqcup_i \tau_i <: \sqcup_j \sigma_j} \text{S-JOIN} \\
 \frac{}{\Sigma \vdash \tau [\mu\rho.\tau/\rho] <: \mu\rho.\tau} \text{S-}\mu\text{-R} \\
 \frac{}{\Sigma \vdash \mu\rho.\tau <: \tau [\mu\rho.\tau/\rho]} \text{S-}\mu\text{-L}
 \end{array}$$

where

$$\begin{array}{ll}
 \triangleleft \overline{\Sigma} = \overline{\triangleleft \Sigma} & \triangleleft (\triangleright H) = H \\
 & \triangleleft (\tau <: \sigma) = \tau <: \sigma
 \end{array}$$

Figure 7. Declarative subtyping rules

3.4 Structural refinement types

We now extend the system with user defined data types, which are specified in Figure 8. A program contains a set of data type declarations of the form $\mathbf{data} N : (\overline{\kappa}; \overline{\kappa}') \rightarrow * \{ \overline{ct} \}$, which introduce a new algebraic data type N with kind $(\overline{\kappa}; \overline{\kappa}') \rightarrow *$ whose name is taken from the set TYNAME, together with its constructors \overline{ct} . For example, the type of

lists is defined with the following declaration:

$$\begin{aligned} \text{data List} &: (*;) \rightarrow * \{ \\ &\forall \alpha. \text{Nil}() : \text{List}(\alpha;), \\ &\forall \alpha. \text{Cons}(\alpha, \text{rec}@(\alpha;)) : \text{List}(\alpha;) \} \end{aligned}$$

The types of the arguments inside data type declarations are restricted to *non-polar types* v , which are the subset of all types that do not contain any lattice constructors \sqcap, \sqcup, \top or \perp . We also require all occurrences of **rec** to be in positive positions.

$N, M \in \text{TYNAME}$

$$\begin{aligned} d &::= \text{data } N : (\bar{\kappa}; \bar{\kappa}) \rightarrow * \{ \bar{ct} \} && \text{Type declaration} \\ ct &::= \forall \bar{\alpha}, \bar{\alpha}'. C(\bar{v}) : N(\bar{\alpha}; \bar{\alpha}') && \text{Constructor declaration} \\ v &::= \alpha && \text{Skolem variable} \\ &| v \rightarrow v && \text{Function type} \\ &| v@(\bar{v}; \bar{v}) && \text{Type application} \\ &| \mu \alpha. v && \text{Recursive type} \\ &| \langle N(\bar{\alpha}; \bar{\alpha}) \mid \overline{C(\bar{v})} \rangle && \text{Refinement type} \\ &| \text{rec} && \text{Recursive occurrence} \\ \sigma, \tau &::= \dots && \text{(From Figure 4)} \\ &| \langle N(\bar{\alpha}; \bar{\alpha}) \mid \overline{C(\bar{\tau})} \rangle && \text{Refinement type} \end{aligned}$$

Figure 8. The syntax of refinement types and data declarations.

The additional kinding, typing and subtyping rules are presented in Figure 9.

4 Type Inference

In this section we present the algorithm which we use to infer structural refinement types. The overall architecture of the type inference algorithm follows Pottier and Rémy [15]. In their approach, one introduces a language of constraints (Section 4.1) and then implements separate phases for first generating constraints (Section 4.2) and then solving them (Section 4.3). The resulting types might be very verbose and can be compressed, which is why a type simplification pass is needed. We will present automata-based type simplification for structural refinement types in Section 4.4.

4.1 Constraints and polar types

At the heart of Hindley-Milner based type inference is the unification algorithm which solves a set of type equalities of the form $\sigma \sim \tau$. Starting from the initial set of generated constraints, the unification algorithm successively applies one of the following three operations, until no further constraints remain: (1) Remove equations of the form $\tau \sim \tau$. (2) Decompose constraints like $\sigma \rightarrow \tau \sim \sigma' \rightarrow \tau'$ into simpler constraints $\sigma \sim \sigma'$ and $\tau \sim \tau'$. (3) Solve equations between a unification variable and another type (like $\alpha^? \sim \tau$)

Kinding rules: $\Delta \vdash \tau : \kappa$

$$\frac{\text{data } N : \kappa \{ \bar{ct} \} \text{ OK} \quad \forall \bar{\beta}, \bar{\beta}'. C(\bar{\tau}[\bar{\beta}/\alpha, \bar{\beta}'/\alpha']) : N(\bar{\beta}; \bar{\beta}') \subseteq \bar{ct}}{\Delta \vdash \langle N(\bar{\alpha}; \bar{\alpha}') \mid \overline{C(\bar{\tau})} \rangle : \kappa} \text{K-REFINE}$$

Well-formedness rules: $d \text{ OK}$

$$\frac{\overline{\bar{\alpha} : \bar{\kappa}, \bar{\alpha}' : \bar{\kappa}' \vdash \bar{\tau} : *}}{\text{data } N : (\bar{\kappa}; \bar{\kappa}') \rightarrow * \{ \sqrt{\bar{\alpha}, \bar{\alpha}'. C(\bar{\tau}) : N(\bar{\alpha}; \bar{\alpha}')} \} \text{ OK}} \text{K-DATA}$$

Typing rules: $\Gamma \vdash e : \tau$

$$\frac{\forall i : \Gamma \vdash e_i : \sigma_i[\bar{\alpha}'/\bar{\alpha}, \bar{\beta}'/\bar{\beta}] \quad \forall \bar{\alpha}, \bar{\beta}. C(\bar{\sigma}_i) : N(\bar{\alpha}; \bar{\beta}) \in \text{Ctors}}{\Gamma \vdash C(\bar{e}_i) : \langle N(\bar{\alpha}; \bar{\beta}) \mid \overline{C(\bar{\sigma}_i)} \rangle @(\bar{\alpha}'; \bar{\beta}')} \text{T-CTOR}$$

$$\frac{\Gamma \vdash e : \langle N(\bar{\alpha}; \bar{\beta}) \mid \overline{C_i(\bar{\sigma}_{ij})} \rangle @(\bar{\alpha}'; \bar{\beta}') \quad \forall i : \Gamma, x_{ij} : \sigma_{ij}[\bar{\alpha}'/\bar{\alpha}, \bar{\beta}'/\bar{\beta}] \vdash e_i : \tau[\bar{\alpha}'/\bar{\alpha}, \bar{\beta}'/\bar{\beta}] \quad \left\{ \sqrt{\bar{\alpha}, \bar{\beta}. C_i(\bar{\sigma}_{ij}) : N(\bar{\alpha}; \bar{\beta})} \right\} = \text{Ctors}(N)}{\Gamma \vdash \text{case } e \text{ of } \{ C_i(\bar{x}_{ij}) \Rightarrow e_i \} : \tau} \text{T-CASE}$$

Subtyping rules: $\Sigma \vdash \sigma <: \tau$

$$\frac{\triangleleft \Sigma \vdash s <: t}{\Sigma \vdash \langle N(\bar{\alpha}; \bar{\beta}) \mid \overline{C(\bar{s})} \rangle <: \langle N(\bar{\alpha}; \bar{\beta}) \mid \overline{C(\bar{t}), \varphi} \rangle} \text{S-REFINE}$$

Figure 9. Extension with structural refinement types.

by substituting τ for all occurrences of $\alpha^?$ in the remaining constraints, possibly after performing an occurs-check. The result of the unification algorithm is a type substitution which maps unification variables to types, such that all constraints are satisfied.

In the algebraic subtyping approach [3, 4, 14] this scheme is modified in order to work with type inequality constraints $\sigma <: \tau$ instead of type equality constraints $\sigma \sim \tau$. There are two important problems one encounters when trying to adapt Hindley-Milner based type inference to type inequality constraints. First, the presence of equi-recursive types (and the corresponding absence of an occurs-check) make special precautions necessary in order to guarantee the termination of the inference algorithm. This is solved by adding a cache of solved constraints, which we will discuss in Section 4.3. The second problem is that unions and intersections as well as top and bottom types can now occur in constraints. To see why this can be a problem, consider the two constraints

$\sigma_1 \sqcup \sigma_2 <: \tau$ and $\sigma <: \tau_1 \sqcup \tau_2$. The first constraint is equivalent to the conjunction of the two constraints $\sigma_1 <: \tau$ and $\sigma_2 <: \tau$, and can easily be handled by the constraint solver. The second constraint $\sigma <: \tau_1 \sqcup \tau_2$, on the other hand, is equivalent to a *disjunction* of constraints which cannot be easily handled. The solution to this problem, proposed by [3], is to exclude constraints of the second kind by introducing a notion of *polar types* which restrict where unions and intersections, top and bottom types can occur within types.

$$\begin{array}{ll}
p & ::= + \mid - & \text{Polarities} \\
\Pi & ::= \epsilon \mid \alpha : p, \Pi & \text{Polarity context} \\
q & ::= \sigma^+ <: \tau^- & \text{Constraints} \\
\neg(+)= - & \quad \neg(-)= + & \text{(Polarity switch)}
\end{array}$$

$$\begin{array}{c}
\boxed{\text{Polarity check: } \Pi \vdash \tau : p} \\
\frac{\Pi(\alpha) = p}{\Pi \vdash \alpha : p} \text{P-BOUND} \qquad \frac{\alpha \notin \Pi}{\Pi \vdash \alpha : p} \text{P-FREE} \\
\frac{}{\Pi \vdash \top : -} \text{P-TOP} \qquad \frac{}{\Pi \vdash \perp : +} \text{P-BOT} \\
\frac{\Pi \vdash \sigma : + \quad \Pi \vdash \tau : +}{\Pi \vdash \sigma \sqcup \tau : +} \text{P-UNION} \\
\frac{\Pi \vdash \sigma : - \quad \Pi \vdash \tau : -}{\Pi \vdash \sigma \sqcap \tau : -} \text{P-INTER} \\
\frac{\Pi \vdash \sigma : \neg(p) \quad \Pi \vdash \tau : p}{\Pi \vdash \sigma \rightarrow \tau : p} \text{P-FUN} \\
\frac{\Pi, \alpha : p \vdash \tau : p}{\Pi \vdash \mu\alpha.\tau : p} \text{P-MU} \\
\frac{\Pi \vdash \tau : p \quad \overline{\Pi \vdash \sigma : p} \quad \overline{\Pi \vdash \rho : \neg(p)}}{\Pi \vdash \tau @ (\overline{\sigma}; \overline{\rho}) : p} \text{P-TYAPP} \\
\frac{\overline{\Pi, \alpha : \overline{p}, \beta : \neg(\overline{p}) \vdash \tau : p}}{\Pi \vdash \langle N(\overline{\alpha}; \overline{\beta}) \mid \overline{C(\overline{\tau})} \rangle : p} \text{P-REFINE}
\end{array}$$

Figure 10. The restriction to polar types.

The sub-syntax of polar types is formalized in Figure 10. We distinguish two different polar types, positive types τ^+ and negative types τ^- . The polarity of a type is checked with the help of the judgement $\Pi \vdash \tau : p$, where a type τ is checked to have polarity p in the polarity context Π which assigns polarities to type variables. The only rules which require types to have concrete polarities are the rules P-UNION, P-INTER, P-TOP and P-BOT; this is the essential restriction of polar types. The rules P-FUN, P-TYAPP and P-REFINE implement the usual rules for type constructors which deal with co- and contravariance, and make use of the helper function $\neg(p)$ for switching the polarity of contravariant arguments. The rule P-MU extends the context with the recursive variable; this is needed to guarantee that the unfolding of a polar

recursive type is always a polar type with the same polarity. Type variables which occur in the context Π must have the polarity assigned to them in the context (rule P-BOUND), whereas type variables not contained in Π can be given any polarity (rule P-FREE). Without the latter rule, we would not be able to assign a polarity to the sensible type $\alpha \rightarrow \alpha$. Note that types which don't contain any unions and intersections or top and bottom types, can always be checked with both polarities.

Constraints q must always have the form $\sigma^+ <: \tau^-$. This restriction excludes the problematic example $\sigma <: \tau_1 \sqcup \tau_2$ we considered earlier in this section.

4.2 Constraint Generation

$$\begin{array}{ll}
\llbracket \alpha \rrbracket_{\tau}^{\perp} := \alpha & \llbracket \alpha^? \rrbracket_{\tau}^{\perp} := \alpha^? \\
\llbracket \alpha \rrbracket_{\tau}^{\top} := \alpha & \llbracket \alpha^? \rrbracket_{\tau}^{\top} := \alpha^? \\
\llbracket v \rightarrow v' \rrbracket_{\tau}^{\perp} := \llbracket v \rrbracket_{\tau}^{\top} \rightarrow \llbracket v' \rrbracket_{\tau}^{\perp} & \\
\llbracket v \rightarrow v' \rrbracket_{\tau}^{\top} := \llbracket v \rrbracket_{\tau}^{\perp} \rightarrow \llbracket v' \rrbracket_{\tau}^{\top} & \\
\llbracket v @ (\overline{v'}; \overline{v''}) \rrbracket_{\tau}^{\perp} := \llbracket v \rrbracket_{\tau}^{\perp} @ (\overline{\llbracket v' \rrbracket_{\tau}^{\perp}}; \overline{\llbracket v'' \rrbracket_{\tau}^{\top}}) & \\
\llbracket v @ (\overline{v'}; \overline{v''}) \rrbracket_{\tau}^{\top} := \llbracket v \rrbracket_{\tau}^{\top} @ (\overline{\llbracket v' \rrbracket_{\tau}^{\top}}; \overline{\llbracket v'' \rrbracket_{\tau}^{\perp}}) & \\
\llbracket \mu\alpha.v \rrbracket_{\tau}^{\perp} := \mu\alpha.\llbracket v \rrbracket_{\tau}^{\perp} & \\
\llbracket \mu\alpha.v \rrbracket_{\tau}^{\top} := \mu\alpha.\llbracket v \rrbracket_{\tau}^{\top} & \\
\llbracket \langle M(\overline{\alpha}; \overline{\alpha'}) \mid \overline{C(\overline{v})} \rangle \rrbracket_{\tau}^{\perp} := \langle M(\overline{\alpha}; \overline{\alpha'}) \mid \overline{C(\llbracket v \rrbracket_{\tau}^{\perp})} \rangle & \\
\llbracket \langle M(\overline{\alpha}; \overline{\alpha'}) \mid \overline{C(\overline{v})} \rangle \rrbracket_{\tau}^{\top} := \langle M(\overline{\alpha}; \overline{\alpha'}) \mid \overline{C(\llbracket v \rrbracket_{\tau}^{\top})} \rangle &
\end{array}$$

And, given **data** $N : (\overline{\kappa}, \overline{\kappa'}) \rightarrow * \{ \overline{C(\overline{v})} \}$:

$$\begin{array}{ll}
\llbracket \mathbf{rec} \rrbracket_N^{\perp} := \langle N(\overline{\alpha}; \overline{\alpha'}) \mid \emptyset \rangle & \\
\llbracket \mathbf{rec} \rrbracket_N^{\top} := \mu\rho.\langle N(\overline{\alpha}; \overline{\alpha'}) \mid \overline{C(\llbracket v \rrbracket_{\rho}^{\perp})} \rangle & \\
\llbracket \mathbf{rec} \rrbracket_{\rho}^{\perp} := \rho & \\
\llbracket \mathbf{rec} \rrbracket_{\rho}^{\top} := \rho &
\end{array}$$

Figure 11. Definitions of lower and upper bound translation

During constraint generation we take a term e (from Figure 2) and a typing context Γ , and return a positive type τ^+ , together with a set of constraints Ξ . This is expressed in the corresponding judgement form $\Gamma \vdash e : \tau \rightsquigarrow \Xi$. The type τ that is generated contains unification variables, whose precise type will be determined in the subsequent constraint solving step. Constraint generation uses the *upper* and *lower bound translation* functions, which are given in Figure 11. These functions translate the types used inside a data type declaration to the *greatest* and respectively *least* types with respect respect to the surrounding declaration.

The constraint generation rules are given in Figure 12a, we will now discuss the rules in detail.

881 The type of a variable x is determined in rule G-VAR by
882 looking up the type assigned to x in the variable context Γ .

883 The terms for function abstraction and application are han-
884 dled in rules G-LAM and G-APP. In order to generate a type
885 for a lambda abstraction $\lambda x.e$, a fresh unification variable $\beta^?$
886 is generated for the variable x . The body of the function e
887 is then checked in the extended context $\Gamma, x : \beta^?$, resulting
888 in the type τ , and the type $\beta^? \rightarrow \tau$ is returned. For function
889 applications $e_1 e_2$ the types σ_1 and σ_2 are inferred for the
890 function and argument, respectively. We generate a fresh uni-
891 fication variable $\beta^?$ for the result of the function application,
892 and add the constraint $\sigma_1 <: \sigma_2 \rightarrow \beta^?$.

893 Constructors are handled in rule G-COR. We first look up
894 the signature of the constructor C in the program; assume
895 that this results in $\forall \bar{\alpha}, \bar{\alpha}'. C(\bar{\sigma}) : N(\bar{\alpha}; \bar{\alpha}')$. We then generate
896 fresh unification variables $\bar{\beta}^?$ and $\bar{\beta}'^?$ for the covariant and
897 contravariant arguments of the type N to which the con-
898 structor C belongs. We return $\langle N(\bar{\alpha}; \bar{\alpha}') \mid C(\bar{\sigma}) \rangle @ (\bar{\beta}^?; \bar{\beta}'^?)$,
899 that is the refinement type of N containing only the construc-
900 tor C , applied to the freshly generated unification variables.
901 But we still have to generate constraints for the arguments
902 of the constructor! For every term e in the argument list of
903 the constructor, we infer a type τ . We then add a constraint
904 between the inferred type τ and the type σ declared in the
905 data type declaration. Before we do this, we have to replace
906 the type variables α and α' in the type declaration with the
907 fresh unification variables. We also have to replace the re-
908 cursive occurrences of N in σ by the fully refined type N .
909 This results in the constraint $\tau <: \llbracket \sigma \rrbracket_N^{\top} [\bar{\beta}^? / \bar{\alpha}, \bar{\beta}'^? / \bar{\alpha}']$ which
910 is ultimately emitted.

911 Pattern matches are handled in rule G-CASE. First, we
912 typecheck the term on which the pattern matching is done.
913 The result of this typecheck is then constrained from above
914 by the refinement type which contains all the constructors
915 that appear in the patterns. Then we generate new unifica-
916 tion variables $\beta_{ij}^?$ for all variables in the patterns and check
917 the cases e_i of the pattern match using these new unifica-
918 tion variables, resulting in types τ_i . The resulting type of the
919 whole pattern match is a new unification variable $\gamma^?$ which
920 is an upper bound of the types of all the cases, thus we add
921 constraints $\tau_i <: \gamma$ for all i . Finally, we need to ensure that
922 the pattern variables are only used in a way that is *compat-*
923 *ible* with the refinement type of the pattern match. This is
924 done via the G-COMPAT rule which constrains the unification
925 variables $\beta_{ij}^?$ to be above the *lower bound translation* $\llbracket \cdot \rrbracket_N^{\perp}$
926 of the corresponding type in declaration of the current refine-
927 ment type N , and below the *upper bound translation* $\llbracket \cdot \rrbracket_N^{\top}$
928 of the same type. The compatibility rule is also indexed by
929 the name of the current refinement type to allow translating
930 occurrences of **rec**. Adding constraints involving the lower
931 bound translation is necessary to allow rejection of terms
932 like $\lambda n. \text{case } n \text{ of } \{S(m) \Rightarrow \text{case } m \text{ of } \{\text{True} \Rightarrow \text{True}\}\}$. In
933 this instance, the constraint $\llbracket \mathbb{N} \rrbracket_N^{\perp} <: \beta^?$ for the unification
934
935

936 variable $\beta^?$ of m which is generated in the pattern match for
937 n is required to later reject this program when this constraint
938 clashes with the constraint $\beta^? <: \langle \text{Bool} \mid \text{True} \rangle$. This leads
939 to the constraint $\langle \mathbb{N} \mid \emptyset \rangle <: \langle \text{Bool} \mid \text{True} \rangle$ by transitivity.
940

941 4.3 Constraint Solving

942 The constraint solver takes a list of constraints qs as input,
943 and produces a list of variable bounds B as output, or else fails
944 if the set of constraints is not satisfiable. A variable bound B
945 assigns to a unification variable $\alpha^?$ a list of lower and upper
946 bounds. Constraint solving is defined as a state transition
947 system on constraint solver states S . A constraint solver
948 state S consists of a cache of already solved constraints, a list
949 of constraints to be solved, and a list of computed variable
950 bounds. If a constraint is encountered that cannot be solved,
951 the constraint solver transitions into the special **Fail** state.

$$\begin{aligned} B &:= \overline{\sigma^+} <: \alpha^? <: \overline{\sigma^-} && \text{Variable Bounds} \\ S &:= \{Q\}; \overline{Q} \vdash \overline{B} \mid \text{Fail} && \text{Constraint Solver State} \end{aligned}$$

952
953
954
955 The algorithm terminates when there are no constraints
956 left to be processed. The cache is necessary to ensure ter-
957 mination and in order to avoid exponential runtime, as it
958 is possible for the algorithm to add an already processed
959 constraint back to its list of constraints. Initially, the cache
960 is empty and the list of constraints to be processed is given
961 by the output of the constraint generator. The initial config-
962 uration of the constraint solver for a list of constraints qs is
963 therefore:

$$964 \emptyset; qs \vdash \text{empty}$$

965
966 As long as the list of constraints to be processed is not
967 empty, one of the rules of Figure 12b will determine the
968 next state of the constraint solver. These rules are applied
969 repeatedly until the list of constraints is empty. In the end,
970 this gives us a constraint solver state that correctly reflects
971 all the upper and lower bounds implied by the set of initial
972 constraints.

973 The rule CACHEHIT will remove the constraint if it is found
974 in the cache, and has therefore already been processed.

975 The rules UPPERBOUND and LOWERBOUND solve *atomic*
976 constraints whose left or the right hand side consist of a uni-
977 fication variable. If both σ_1 and σ_2 are unification variables,
978 the rule UPPERBOUND will fire, ensuring determinism³. We
979 only discuss the rule UPPERBOUND, since the rule LOWER-
980 BOUND works similarly. The constraint $\alpha^? <: \sigma$ is solved by
981 adding σ to the list of upper bounds of $\alpha^?$. But, as discussed
982 in [14], we must also make sure that this new bound is con-
983 sistent with the existing lower bounds of $\alpha^?$. We therefore
984 add one new constraint between σ and each of the existing
985 lower bounds to the constraints to be processed. These are
986 the subconstraints that are implied by transitivity of the sub-
987 typing relation. Note that no occurs check is performed; it is

988 ³This is analogous to the situation in Hindley-Milner, where to solve a
989 constraint like $\alpha \sim \beta$, we are free to either substitute α for β or β for α .
990

Constraint generation: $\Gamma \vdash e : \tau \rightsquigarrow \Xi$

$$\begin{array}{c}
\frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau \rightsquigarrow \emptyset} \text{G-VAR} \qquad \frac{\Gamma, x : \beta^? \vdash e : \tau \rightsquigarrow \Xi \quad \text{Fresh}(\beta^?)}{\Gamma \vdash \lambda x. e : \beta^? \rightarrow \tau \rightsquigarrow \Xi} \text{G-LAM} \\
\frac{\Gamma \vdash e_1 : \sigma_1 \rightsquigarrow \Xi_1 \quad \Gamma \vdash e_2 : \sigma_2 \rightsquigarrow \Xi_2 \quad \text{Fresh}(\beta^?)}{\Gamma \vdash e_1 e_2 : \beta^? \rightsquigarrow \{\sigma_1 <: \sigma_2 \rightarrow \beta^?\} \cup \Xi_1 \cup \Xi_2} \text{G-APP} \\
\frac{\Gamma \vdash e : \tau \rightsquigarrow \Xi \quad \forall \bar{\alpha}, \bar{\alpha}'. C(\bar{\sigma}) : N(\bar{\alpha}; \bar{\alpha}') \in \text{Ctors} \quad \text{Fresh}(\bar{\beta}^?, \bar{\beta}'^?)}{\Gamma \vdash C(\bar{e}) : \langle N(\bar{\alpha}; \bar{\alpha}') \mid C(\bar{\sigma}) \rangle @(\bar{\beta}^?, \bar{\beta}'^?) \rightsquigarrow \left\{ \tau <: [\sigma]_N^T[\bar{\beta}^?/\bar{\alpha}, \bar{\beta}'^?/\bar{\alpha}'] \right\} \cup (\cup_i \Xi_i)} \text{G-CTOR} \\
\frac{\frac{\Gamma, x : \beta^? \vdash e : \tau \rightsquigarrow \Xi}{\Gamma \vdash e : \sigma \rightsquigarrow \Xi} \quad \text{Fresh}(\bar{\beta}^?, \bar{\gamma}^?, \bar{\delta}^?, \bar{\delta}'^?) \quad \frac{C(\bar{\beta}^?) \heartsuit_N C(\bar{\sigma}[\bar{\delta}^?/\bar{\alpha}, \bar{\delta}'^?/\bar{\alpha}']) \rightsquigarrow \Xi_\heartsuit}{\forall \bar{\alpha}, \bar{\alpha}'. C(\bar{\sigma}) : N(\bar{\alpha}; \bar{\alpha}') \subseteq \text{Ctors}}}{\Gamma \vdash \text{case } e \text{ of } \{C(\bar{x}) \Rightarrow e\} : \gamma^? \rightsquigarrow \left\{ \sigma <: \langle N(\bar{\alpha}; \bar{\alpha}') \mid C(\bar{\sigma}) \rangle @(\bar{\delta}^?, \bar{\delta}'^?), \tau <: \gamma^?, \right\} \cup \Xi \cup \{\cup_i \Xi_i\} \cup \{\cup_i \Xi_{\heartsuit, i}\}} \text{G-CASE} \\
\frac{C(\bar{\beta}^?) \heartsuit_N C(\bar{\tau}) \rightsquigarrow \left\{ [\tau]_N^T <: \beta^? <: [\tau]_N^T \right\}}{C(\bar{\beta}^?) \heartsuit_N C(\bar{\tau}) \rightsquigarrow \left\{ [\tau]_N^T <: \beta^? <: [\tau]_N^T \right\}} \text{G-COMPAT}
\end{array}$$

(a) Constraint generation rules. Inputs are **contexts** and **terms**, outputs are **types** and **constraint sets**.

Constraint solver step: $S \rightarrow S'$

$$\begin{array}{c}
\frac{q \in \text{ca}}{ca; q, qs \vdash bs \rightarrow ca; qs \vdash bs} \text{CACHEHIT} \\
\frac{q \notin \text{ca} \quad q = \alpha^? <: \sigma \quad bs(\alpha^?) = lbs <: \alpha^? <: \text{ubs}}{ca; q, qs \vdash bs \rightarrow ca, q; \{lb <: \sigma\}_{lb \in lbs}, qs \vdash bs[\alpha^? \mapsto lbs <: \alpha^? <: \{\text{ubs}, \sigma\}]} \text{UPPERBOUND} \\
\frac{q \notin \text{ca} \quad q = \sigma <: \alpha^? \quad \sigma \notin \text{TYVAR} \quad bs(\alpha^?) = lbs <: \alpha^? <: \text{ubs}}{ca; q, qs \vdash bs \rightarrow ca, q; \{\sigma <: \text{ub}\}_{ub \in \text{ubs}}, qs \vdash bs[\alpha^? \mapsto \{\sigma, lbs\} <: \alpha^? <: \text{ubs}]} \text{LOWERBOUND} \\
\frac{q \notin \text{ca} \quad q = \sigma_1 <: \sigma_2 \quad \sigma_1 \notin \text{TYVAR} \quad \sigma_2 \notin \text{TYVAR} \quad \text{Sub}(q) = qs'}{ca; q, qs \vdash bs \rightarrow ca, q; qs' \# qs \vdash bs} \text{SUBOK} \\
\frac{q \notin \text{ca} \quad q = \sigma_1 <: \sigma_2 \quad \sigma_1 \notin \text{TYVAR} \quad \sigma_2 \notin \text{TYVAR} \quad \text{Sub}(q) = \text{Fail}}{ca; q, qs \vdash bs \rightarrow \text{Fail}} \text{SUBFAIL}
\end{array}$$

(b) The biunification algorithm.

Decomposing non-atomic constraints: $\text{Sub}(-) : q \rightarrow \bar{q}/\text{Fail}$

$$\begin{array}{l}
\text{Sub}(\tau <: \top) := \emptyset \qquad \text{Sub}(\perp <: \sigma) := \emptyset \\
\text{Sub}(\tau_1 \sqcup \tau_2 <: \sigma) := \{\tau_1 <: \sigma, \tau_2 <: \sigma\} \quad \text{Sub}(\tau <: \sigma_1 \sqcap \sigma_2) := \{\tau <: \sigma_1, \tau <: \sigma_2\} \\
\text{Sub}(\tau <: \mu\alpha.\sigma) := \{\tau <: \sigma[\mu\alpha.\sigma/\alpha]\} \quad \text{Sub}(\mu\alpha.\tau <: \sigma) := \{\tau[\mu\alpha.\tau/\alpha] <: \sigma\} \\
\text{Sub}(\alpha <: \alpha) := \emptyset \\
\text{Sub}(\sigma_1 \rightarrow \tau_1 <: \sigma_2 \rightarrow \tau_2) := \{\sigma_2 <: \sigma_1, \tau_1 <: \tau_2\} \\
\text{Sub}(\tau @(\bar{\sigma}; \bar{\rho}) <: \tau' @(\bar{\sigma}'; \bar{\rho}')) := \{\tau <: \tau', \bar{\sigma} <: \bar{\sigma}', \bar{\rho}' <: \bar{\rho}\} \\
\text{Sub}(\langle N(\bar{\alpha}; \bar{\alpha}') \mid \emptyset \rangle <: \langle N(\bar{\beta}; \bar{\beta}') \mid \psi \rangle) := \emptyset \\
\text{Sub}(\langle N(\bar{\alpha}; \bar{\alpha}') \mid C(\bar{\tau}), \varphi \rangle <: \langle N(\bar{\beta}; \bar{\beta}') \mid C(\bar{\sigma}), \psi \rangle) := \left\{ \tau_i <: \sigma_i[\bar{\alpha}/\bar{\beta}] \right\} \cup \text{Sub}(\langle N(\bar{\alpha}; \bar{\alpha}') \mid \varphi \rangle <: \langle N(\bar{\beta}; \bar{\beta}') \mid \psi \rangle)
\end{array}$$

(c) Decomposing subtyping constraints.

Figure 12. Type inference.

possible that type variables refer to themselves recursively through their lower or upper bounds.

If neither σ_1 nor σ_2 are type variables, we invoke the $\text{Sub}(-)$ function to decompose the constraint into a list of subconstraints. This function is specified in Figure 12c. If none of the rules in that figure match, then the function $\text{Sub}(-)$ returns Fail.

4.4 Type Simplification

The principal types in vanilla Hindley Milner type inference have the useful property that they are syntactically unique⁴. For this reason, it is not necessary to simplify or normalize the principal types that have been inferred in such a system. The situation is different in a subtyping system, where there are usually many equivalent but syntactically distinct principal types for any typeable term. For example, the term $\lambda x y z. \text{case } x \text{ of } \{\text{True} \Rightarrow y, \text{False} \Rightarrow z\}$ can be given the syntactically distinct but equivalent types $\forall \alpha \beta. \text{Bool} \rightarrow \alpha \rightarrow \beta \rightarrow \alpha \sqcap \beta$ and $\forall \alpha. \text{Bool} \rightarrow \alpha \rightarrow \alpha \rightarrow \alpha$. In this case, we would prefer to display the second type to the user, since it is syntactically shorter and easier to understand.

There are two main approaches to simplify types. The direct approach presented by Parreaux [14, Section 4] works directly on types and has a simpler implementation. The alternative is the automata-based approach of Dolan [3]. It uses a close correspondence between types and their encoding in finite automata, witnessed by a *representation theorem* [3, Section 7]. This theorem states that any two type automata represent the same type *iff* they accept the same language. Therefore, any algorithm for simplifying finite automata can be used to simplify types, which permits the use of generic, well-known algorithms from automata theory.

For our purposes, we assume the same equivalence for our type automata, using them to simplify our inferred types, though we do not prove the corresponding theorem.

As an example, let us observe how the type of the Boolean negation function is simplified. The negation function is given by

```
def not := λb. case b of {True ⇒ False, False ⇒ True}.
```


The inferred type is somewhat convoluted. Figure 13 depicts the type automata as generated initially along with the type automata after the determinisation step. The simplified inferred type of not is—as one would expect—

$$\langle \text{Bool} \mid \text{True}, \text{False} \rangle \rightarrow \langle \text{Bool} \mid \text{True}, \text{False} \rangle.$$

Meet and join types correspond to non-deterministic transitions, while equi-recursive types correspond to cyclical type automata, as shown in Figure 14. Here we can see the necessity of the polarity restriction for recursive types discussed in Section 4.1: Occurrences of μ -bound variables are

⁴Unique up to α -equivalence, of course.

Before determinisation:



After determinisation:




Figure 13. Determinisation of type automata for type $\langle \text{Bool} \mid \text{True}, \text{False} \rangle \rightarrow \langle \text{Bool} \mid \text{True}, \text{False} \rangle$. The shading of the states signifies polarity, where gray is negative and white is positive.

represented by cyclical transitions back to the state corresponding to the recursive type; as such, they must possess the same polarity.




Figure 14. Type automata for refinement type of even Peano numbers $\mu \alpha. \langle \mathbb{N} \mid Z, S(\langle \mathbb{N} \mid S(\alpha) \rangle) \rangle$

Higher-kinded types are represented in the same way as types of kind $*$, except that the kind annotations (not visible in the automata figures) are different. For meets and joins, only nodes with the same kind may be merged during the determinisation. Type application is handled analogously to the representation of the function type. An example for an automaton involving higher-kinded types can be found in Figure 15 which depicts the type of a singleton list containing only the value True.

5 Related Work

We have identified the following areas of related work:

Algebraic Subtyping: The elegant combination of type inference, parametric polymorphism and subtyping has been a longstanding problem. The algebraic subtyping approach was developed as a solution to this problem by Dolan [3]




Figure 15. Type automaton for type of one-element **Bool** list $\langle \text{List}(\alpha;) \mid \text{Cons}(\alpha, \langle \text{List}(\beta;) \mid \text{Nil} \rangle @(\alpha;)) \rangle @(\langle \text{Bool} \mid \text{True} \rangle;)$

and Dolan and Mycroft [4]. Algebraic subtyping is the central foundation on which our paper builds. We have also profited from the presentation of the underlying algorithm and ideas of the algebraic subtyping approach by Parreaux [14].

Polymorphic variants: The algebraic subtyping approach is, in general, not particularly concerned with the concrete types available in the system, only with the algebraic properties of the subtyping lattice. For example, [3, 4] illustrate their approach with a system which only contains booleans, functions and records. In this paper we study the combination of ordinary algebraic data types, which all languages in the ML lineage support, with polymorphic variants. Polymorphic variants [7, 8] allow to program with the familiar tools of functional programming: building up data with constructors and decomposing with pattern matching. But they don't require the programmer to declare the data types and their constructors in advance. Structural refinement types combine the typing rules for algebraic data types with the typing rules of polymorphic variants.

Record subtyping and codata types: There is a well-known duality between data types and codata types [1, 5, 9, 16]. A special instance of this duality is the duality of polymorphic variants and extensible records [8]. In fact, we have developed and implemented the ideas presented in this paper in the more general context of algebraic data and codata types. We have specialized the formalization to data types and polymorphic variants in order to simplify the presentation, but the extension to the more general setting is straightforward.

Refinement types: Refinement types in our sense were introduced in a seminal paper by Freeman and Pfenning [6]. They conceive of refinement types as abstract domains, which have to be manually specified by the programmer in advance. They require a finite number of refinement types in order to keep type inference decidable within the intersection type system they use.

Another recent system with a very similar aim are the intensional refinement types of Jones and Ramsay [12]. In distinction to our system, they can only express the complete absence of a constructor in a refinement type. For example, they cannot express the type of nonempty lists, since that type requires to exclude the `nil` constructor only as the outermost constructor. A huge benefit of their paper, on the other hand, is the detailed study of the computational cost of tracking type refinements. We have not yet attempted to characterize the computational cost of our refinements theoretically or empirically.

Variance of type parameters The idea of co- and contravariant type parameters has been explored in-depth in the literature on object-oriented subtyping [11]. A significant trade-off in the design of a variance mechanism is the difference between use-site and definition-site variance [2]. With regard to this distinction, structural refinement types as presented here offer definition-site variance.

6 Future Work

We have implemented the refinement types described in this paper in the `DUALSUB` language, and we plan to provide both a theoretical and empirical evaluation of that system. We plan to provide proofs of the soundness of the system as well as the principal types property in the future. We conjecture that it is relatively straight-forward to modify the existing proofs provided by Dolan [3] and Parreaux [14]. For an empirical validation we plan to implement larger case studies which can help evaluate whether the expressivity of the refinement types is useful in practice, and whether the type inference algorithm scales well with larger programs.

7 Discussion

We have presented a type inference algorithm for *structural refinement types*, which combines properties of both algebraic data types and polymorphic variants. These novel types allow us to express refinements on algebraic data types which correspond to regular sublanguages. In contrast to other refinement type systems, our approach uses only familiar techniques from constraint based type inference. As a result, the addition of these types to a language which already implements the algebraic subtyping approach is comparatively simple. We have implemented *structural refinement types* in the `DUALSUB` language.

References

- [1] Andreas Abel, Brigitte Pientka, David Thibodeau, and Anton Setzer. 2013. Copatterns: Programming infinite structures by observations. In *Proceedings of the Symposium on Principles of Programming Languages* (Rome, Italy). ACM, New York, 27–38.
- [2] John Altidor, Shan Shan Huang, and Yannis Smaragdakis. 2011. Taming the Wildcards: Combining Definition- and Use-Site Variance. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation* (San Jose, California, USA) (*PLDI*

- 1321 '11). Association for Computing Machinery, New York, NY, USA, 602–
1322 613. <https://doi.org/10.1145/1993498.1993569>
- 1323 [3] Stephen Dolan. 2017. *Algebraic Subtyping: Distinguished Dissertation*
1324 2017. BCS, Swindon, GBR.
- 1325 [4] Stephen Dolan and Alan Mycroft. 2017. Polymorphism, Subtyping,
1326 and Type Inference in MLsub. In *Proceedings of the 44th ACM SIGPLAN*
1327 *Symposium on Principles of Programming Languages (Paris, France)*
1328 *(POPL 2017)*. Association for Computing Machinery, New York, NY,
1329 USA, 60–72.
- 1330 [5] Paul Downen, Zachary Sullivan, Zena M Ariola, and Simon Peyton
1331 Jones. 2019. Codata in Action. In *European Symposium on Programming*.
1332 Springer, 119–146.
- 1333 [6] Tim Freeman and Frank Pfenning. 1991. Refinement Types for ML.
1334 In *Proceedings of the ACM SIGPLAN 1991 Conference on Programming*
1335 *Language Design and Implementation (Toronto, Ontario, Canada) (PLDI*
1336 *'91)*. Association for Computing Machinery, New York, NY, USA, 268–
1337 277. <https://doi.org/10.1145/113445.113468>
- 1338 [7] Jacques Garrigue. 1998. Programming with polymorphic variants. In
1339 *ML Workshop*, Vol. 13. Baltimore.
- 1340 [8] Benedict R Gaster and Mark P Jones. 1996. *A polymorphic type system*
1341 *for extensible records and variants*. Technical Report. Citeseer.
- 1342 [9] Tatsuya Hagino. 1989. Codatatypes in ML. *Journal of Symbolic Com-*
1343 *putation* 8, 6 (1989), 629–650.
- 1344 [10] Atsushi Igarashi, Benjamin C Pierce, and Philip Wadler. 2001. Feather-
1345 weight Java: a minimal core calculus for Java and GJ. *ACM Transactions*
1346 *on Programming Languages and Systems (TOPLAS)* 23, 3 (2001), 396–
1347 450.
- 1348 [11] Atsushi Igarashi and Mirko Viroli. 2006. Variant Parametric Types:
1349 A Flexible Subtyping Scheme for Generics. *ACM Trans. Program.*
1350 *Lang. Syst.* 28, 5 (sep 2006), 795–847. [https://doi.org/10.1145/1152649.](https://doi.org/10.1145/1152649.1152650)
1351 [1152650](https://doi.org/10.1145/1152650)
- 1352 [12] Eddie C Jones and Steven J. Ramsay. 2021. Intensional datatype refine-
1353 ment: with application to scalable verification of pattern-match safety.
1354 *Proceedings of the ACM on Programming Languages* 5 (2021), 1 – 29.
1355 1377
- 1356 [13] Klaus Ostermann. 2008. Nominal and Structural Subtyping in
1357 Component-Based Programming. *Journal of Object Technology* 7,
1358 1 (2008), 121 – 145. http://www.jot.fm/issues/issue_2008_01/article4/
1359 1379
- 1360 [14] Lionel Parreaux. 2020. The Simple Essence of Algebraic Subtyping:
1361 Principal Type Inference with Subtyping Made Easy (Functional Pearl).
1362 *Proc. ACM Program. Lang.* 4, ICFP, Article 124 (Aug. 2020), 28 pages.
1363 <https://doi.org/10.1145/3409006> 1380
- 1364 [15] François Pottier and Didier Rémy. 2005. The Essence of ML Type
1365 Inference. In *Advanced Topics in Types and Programming Languages*,
1366 Benjamin C. Pierce (Ed.). MIT Press, Chapter 10, 389–489. [http:](http://crystal.inria.fr/attapl/)
1367 [//crystal.inria.fr/attapl/](http://crystal.inria.fr/attapl/) 1381
- 1368 [16] Tillmann Rendel, Julia Trieflinger, and Klaus Ostermann. 2015. Auto-
1369 matic Refunctionalization to a Language with Copattern Matching:
1370 With Applications to the Expression Problem. In *Proceedings of the*
1371 *20th ACM SIGPLAN International Conference on Functional Program-*
1372 *ming (Vancouver, BC, Canada) (ICFP 2015)*. Association for Computing
1373 Machinery, New York, NY, USA, 269–279. [https://doi.org/10.1145/](https://doi.org/10.1145/2784731.2784763)
1374 [2784731.2784763](https://doi.org/10.1145/2784731.2784763) 1382
- 1375 [17] Patrick M. Rondon, Ming Kawaguci, and Ranjit Jhala. 2008. Liquid
1376 Types. In *Proceedings of the 29th ACM SIGPLAN Conference on Pro-*
1377 *gramming Language Design and Implementation (Tucson, AZ, USA)*
1378 *(PLDI '08)*. Association for Computing Machinery, New York, NY, USA,
1379 159–169. <https://doi.org/10.1145/1375581.1375602> 1383
- 1380 [18] Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and
1381 Simon Peyton-Jones. 2014. Refinement Types for Haskell. In *Pro-*
1382 *ceedings of the 19th ACM SIGPLAN International Conference on Func-*
1383 *tional Programming (Gothenburg, Sweden) (ICFP '14)*. Association
1384 for Computing Machinery, New York, NY, USA, 269–282. [https:](https://doi.org/10.1145/2628136.2628161)
1385 [//doi.org/10.1145/2628136.2628161](https://doi.org/10.1145/2628136.2628161) 1384
- 1386 1385
- 1387 1386
- 1388 1387
- 1389 1388
- 1390 1389
- 1391 1390
- 1392 1391
- 1393 1392
- 1394 1393
- 1395 1394
- 1396 1395
- 1397 1396
- 1398 1397
- 1399 1398
- 1400 1399
- 1401 1400
- 1402 1401
- 1403 1402
- 1404 1403
- 1405 1404
- 1406 1405
- 1407 1406
- 1408 1407
- 1409 1408
- 1410 1409
- 1411 1410
- 1412 1411
- 1413 1412
- 1414 1413
- 1415 1414
- 1416 1415
- 1417 1416
- 1418 1417
- 1419 1418
- 1420 1419
- 1421 1420
- 1422 1421
- 1423 1422
- 1424 1423
- 1425 1424
- 1426 1425
- 1427 1426
- 1428 1427
- 1429 1428
- 1430 1429
- 1431 1430