

A Quick Look at Impredicativity

ALEJANDRO SERRANO, 47 Degrees, Spain and Utrecht University, The Netherlands

JURRIAN HAGE, Utrecht University, The Netherlands

SIMON PEYTON JONES, Microsoft Research, United Kingdom

DIMITRIOS VYTINIOTIS, DeepMind, United Kingdom

Type inference for parametric polymorphism is wildly successful, but has always suffered from an embarrassing flaw: polymorphic types are themselves not first class. We present Quick Look, a practical, implemented, and deployable design for impredicative type inference. To demonstrate our claims, we have modified GHC, a production-quality Haskell compiler, to support impredicativity. The changes required are modest, localised, and are fully compatible with GHC's myriad other type system extensions.

Draft March 2020. We would appreciate any comments or feedback to the email addresses below.

Additional Key Words and Phrases: Type systems, impredicative polymorphism, constraint-based inference

1 INTRODUCTION

Parametric polymorphism backed by Damas-Milner type inference was first introduced in ML [14], and has been enormously influential and widely used. But despite its this impact, it has always suffered from an embarrassing shortcoming: *Damas-Milner type inference, and its many variants, cannot instantiate a type variable with a polymorphic type*; in the jargon, the system is *predicative*.

Alas, predicativity makes polymorphism a second-class feature of the type system. The type $\forall a.[a] \rightarrow [a]$ is fine (it is the type of the list reverse function), but the type $[\forall a.a \rightarrow a]$ is not, because a \forall is not allowed inside a list. So \forall -types are not first class: they can appear in some places but not others. Much of the time that does not matter, but sometimes it matters a lot; and, tantalisingly, it is often “obvious” to the programmer what the desired impredicative instantiation should be (Section 2).

Thus motivated, a long succession of papers have tackled the problem of type inference for impredicativity [2, 11, 12, 24, 26, 27]. None has succeeded in producing a system that is simultaneously expressive enough to be useful, simple enough to support robust programmer intuition, compatible with a myriad other type system extensions, and implementable without an invasive rewrite of a type inference engine tailored to predicative type inference.

In Section 3 we introduce Quick Look, a new inference algorithm for impredicativity that, for the first time, (a) handles many “obvious” examples; (b) is expressive enough to handle all of System F; (c) requires no extension to types, constraints, or intermediate representation; and (d) is easy and non-invasive to implement in a production-scale type inference engine – indeed we have done so in GHC. We make the following contributions:

- We formalise a higher-rank baseline system (Section 4), and give the changes required for Quick Look (Section 5). A key property of Quick Look is that it requires only highly localised changes to such a specification. In particular, no new forms of types are required, and programs can be elaborated into a statically typed intermediate language based on System F. Some other approaches, such as MLF [2], require substantial changes to the intermediate language, but Quick Look does not.

Authors' addresses: Alejandro Serrano, alejandro.serrano@47deg.com, 47 Degrees, San Fernando, Spain, A.SerranoMena@uu.nl, Utrecht University, Utrecht, The Netherlands; Jurriaan Hage, J.Hage@uu.nl, Utrecht University, Utrecht, The Netherlands; Simon Peyton Jones, simonpj@microsoft.com, Microsoft Research, Cambridge, United Kingdom; Dimitrios Vytiniotis, dvytin@google.com, DeepMind, London, United Kingdom.

$head :: \forall p.[p] \rightarrow p$ $tail :: \forall p.[p] \rightarrow [p]$ $[] :: \forall p.[p]$ $(:) :: \forall p.p \rightarrow [p] \rightarrow [p]$ $single :: \forall p.p \rightarrow [p]$ $(\#) :: \forall p.[p] \rightarrow [p] \rightarrow [p]$ $id :: \forall a.a \rightarrow a$ $ids :: [\forall a.a \rightarrow a]$ $app :: \forall a b.(a \rightarrow b) \rightarrow a \rightarrow b$ $revapp :: \forall a b.a \rightarrow (a \rightarrow b) \rightarrow b$	$runST :: \forall d.(\forall s.ST s d) \rightarrow d$ $argST :: \forall s.ST s Int$ $poly :: (\forall a.a \rightarrow a) \rightarrow (Int, Bool)$ $inc :: Int \rightarrow Int$ $choose :: \forall a.a \rightarrow a \rightarrow a$ $poly :: (\forall a.a \rightarrow a) \rightarrow (Int, Bool)$ $auto :: (\forall a.a \rightarrow a) \rightarrow (\forall a.a \rightarrow a)$ $auto' :: (\forall a.a \rightarrow a) \rightarrow b \rightarrow b$ $map :: \forall p q.(p \rightarrow q) \rightarrow [p] \rightarrow [q]$
--	--

Fig. 1. Type signatures for functions used in the text

- We prove a number of theorems about our system, including about which transformations do, and do not, preserve typeability (Section 6).
- We give a type inference algorithm for Quick Look (Section 7). This algorithm is based on the now-standard approach of first *generating typing constraints* and then *solving them* [21]. As in the case of the declarative specification, no new forms of types or constraints are needed. Section 7 proves its soundness compared with the declarative specification in Section 5.¹ The implementation is in turn based very closely on this algorithm. The constraint generation judgements in Section 7 and appendix B also appear to be the first formal account of the extremely effective combination of bidirectional type inference [18] with constraint-based type inference [21, 25].
- Because Quick Look’s impact is so localised, it is simple to implement, even in a production compiler. Concretely, the implementation of Quick Look in GHC, a production compiler for Haskell, affected only 1% of GHC’s inference engine.
- To better support impredicativity, we propose to abandon contravariance of the function arrow (Section 5.4). There are independent reasons for making this change [17], but it is illuminating to see how it helps impredicativity. Moreover, we provide data on its impact (Section 9).

The paper uses a very small language, to allow us to focus on impredicativity, but Quick Look scales very well to a much larger language. Appendix B gives the details for a much richer set of features.

We present our work in the concrete setting of (a tiny subset of) Haskell, but there is nothing Haskell-specific about it. Quick Look could readily be used in any other type inference system.

We cover the rich related work in Section 10.

2 MOTIVATION

The lack of impredicativity means that polymorphism is fundamentally second class: *we cannot abstract over polymorphic types*. For example, even something as basic as function composition fails on functions with higher-rank types (types with forall in them). Suppose

$$f :: (\forall a.[a] \rightarrow [a]) \rightarrow Int \quad g :: Bool \rightarrow \forall a.[a] \rightarrow [a]$$

¹We conjecture that completeness is true as well – we are not aware of any counterexample.

Then the composition $(f \circ g)$ fails typechecking, despite the obvious compatibility of the types involved, simply because the composition requires instantiating the type of (\circ) with a polytype.

As another example, Augustsson describes an application [1] in which it was crucial to have a function var of type

```
var :: RValue a → IO (∀lr.LR lr ⇒ lr a)
```

The function var is an IO action that returns a polymorphic value. Yet in Haskell today, this is out of reach; instead you have to define a new *named* type, thus:

```
newtype LRType a = MkLR (∀lr.LR lr ⇒ lr a)
```

```
var :: RValue a → IO (LRType a)
```

Moreover, every use of var must use pattern-matching to unwrap the newtype. We call this approach “boxed impredicativity”, because the forall is wrapped in an named “box”, here *LRType*. But boxed impredicativity is tiresome at best, and declaring a new type for every polymorphic shape is gruesome.

Why not simply allow first-class polymorphism, so that $[\forall a.a \rightarrow a]$ is a valid type? The problem is in *type inference*. Consider the expression $(single\ id)$, where the type of $single$ and id are given in Figure 1. It is not clear whether to instantiate p with $\forall a.a \rightarrow a$, or with $Int \rightarrow Int$, or some other monomorphic type. Indeed $(single\ id)$ does not even have a most general (principal) type: it has two incomparable types: $\forall a.[a \rightarrow a]$ and $[\forall a.a \rightarrow a]$. Losing principal types, especially for such an innocuous program, is a heavy price to pay for first-class polymorphism.

But in many cases there is no such problem. Consider $(head\ ids)$ where, again, the types are given in Figure 1. Now there is no choice: the only possibility is to instantiate p with $\forall a.a \rightarrow a$. Our idea, just as in much previous work [24], is to exploit that special case. Our overall goals are these:

- *First class polymorphism.* We want forall-types to be first class. A function like list $reverse :: \forall a.[a] \rightarrow [a]$ should work as uniformly over $[\forall a.a \rightarrow a]$ as it does over $[Int]$ and $[Bool]$, and should do so without type annotations. No mainstream deployed language allows that; and not being able to do so is a fundamental failure of abstraction. Using boxed impredicativity is an anti-modular second best.
- *Predictable type inference:* programmers should acquire a robust mental model of what will typecheck, and what will not. Typically they do so through a process of trial and error, but our formalism in Section 5 is specifically designed to enshrine the common-sense idea that when there is clear evidence (through the argument or result type) about how to instantiate a call, type inference should take advantage of it.
- *Minimize type annotations:* “obviously typable” programs should be typeable without annotation. To substantiate this necessarily-qualitative claim we give numerous examples, especially in Figure 11.
- *Conservative extension of Damas-Milner* and its extensions to type classes, higher rank, etc etc. That is, existing programs continue to typecheck (Section 6.1).
- *Can express all of System F*, perhaps with the addition of type annotations (Section 6.1).
- *Localised, in both specification and implementation.* We seek a system that affects only a small part of the specification, and the implementation, of the type system and its inference algorithm. Modern type systems, such as that of Haskell, OCaml, or Scala, are subtle and complicated; anything that requires pervasive changes is unlikely to be implemented.

3 THE QUICK LOOK

Our new approach works as follows:

- Treat applications as a whole: a function applied to a list of arguments. The list of arguments can be empty, in which case the “function” is not necessarily a function: it can be a polymorphic value, such as the empty list $[] :: \forall p.[p]$.
- When instantiating the function, take a “quick look” at the arguments to guide that (possibly impredicative) instantiation.
- If the quick look produces a definite answer, use it; otherwise instantiate with a monotype (as usual in Hindley-Damas-Milner type inference).

In our example (*head ids*), we have to instantiate the type of $head :: \forall p.[p] \rightarrow p$. The argument $ids :: [\forall a.a \rightarrow a]$ must be compatible with the type $head$ expects, namely $[p]$. So we are forced to instantiate $p := \forall a.a \rightarrow a$.

On the other hand for (*single id*), a quick look sees that the argument $id :: \forall a.a \rightarrow a$ must be compatible with the type $single$ expects, namely p . But that does not tell us what p must be: should we instantiate that $\forall a$ or not? So the quick look produces no advice, and we revert to standard Hindley-Damas-Milner type inference by instantiating p with a monotype $\tau \rightarrow \tau$. (Operationally, the inference algorithm will instantiate p with a unification variable.)

Why is (*head ids*) easier? Because the type variable p in $head$'s type appears *guarded*, under the list type constructor; but not so for *single*. Exploiting this guardedness was the key insight of earlier work [24].

The quick-look approach scales nicely to handle multiple arguments. For example, consider the expression $(id : ids)$, where $(:)$ is Haskell's infix cons operator. How should we instantiate the type of $(:)$ given in Figure 1? A quick look at the first argument, id , yields no information; it is like the (*single id*) case. But a quick look at the second argument, ids , immediately tells us that p must be instantiated with $\forall a.a \rightarrow a$. We gain a lot from taking a quick look at *all* the arguments before committing to *any* instantiation.

3.1 “Quick-looking” the Result

So far we have concentrated on using the *arguments* of a call to guide instantiation, but we can also use the *result* type. Consider this expression, which has a user-written type signature:

$$(single\ id) :: [\forall a.a \rightarrow a]$$

When considering how to instantiate *single*, we know that it produces a result of type $[p]$, which must fit the user-specified result type $[\forall a.a \rightarrow a]$. So again there is only one possible choice of instantiation, namely $p := \forall a.a \rightarrow a$.

This same mechanism works when the “expected” type comes from an enclosing call. Suppose $foo :: [\forall a.a \rightarrow a] \rightarrow Int$, and consider $foo\ (single\ id)$. The context of the call (*single id*) specifies the result type of the call, just as the type signature did before. We need to “push down” the type expected by the context into an expression, but fortunately this ability is already well established in the form of *bidirectional type inference* [18, 19] as Section 4 discusses.

Taking a quick look at the result type is particularly important for lone variables. We can treat a lone variable as a degenerate form of call with zero arguments. Its instantiation cannot be informed by a quick look at the arguments, since it has none; but it can benefit from the result type. A ubiquitous example is the empty list $[] :: \forall p.[p]$. Consider the task of instantiating $[]$ in the context of a call $foo\ []$. Since foo expects an argument of type $[\forall a.a \rightarrow a]$, the only way to instantiate $[]$ is with $p := \forall a.a \rightarrow a$.

Finally, here is a more complicated example. Consider the call $([] \# ids)$, where the types are given in Figure 1. First we decide how to instantiate $(\#)$ and, as in the case of *head*, we can discover its instantiation $p := \forall a.a \rightarrow a$ from its second argument *ids*. Having made that decision we now

197 typecheck its first argument, $[]$, knowing that the result type must be $[\forall a.a \rightarrow a]$, and that in turn
 198 tells us the instantiation of $[]$.

200 3.2 Richer Arguments

201 So far the argument of every example call has been a simple variable. But what if it was a list
 202 comprehension? A lambda? Another call?

203 One strength of the quick-look approach is that we are free to make restrictions without affecting
 204 anything fundamental. For example, we could say (brutally) that quick look yields no advice for an
 205 argument other than a variable. The “no advice” case simply means that we will look for information
 206 in other arguments or, if none of them give advice, revert to monomorphic instantiation.

207 We have found, however, that it is both easy and beneficial to allow nested calls. For example,
 208 consider $(id : (id : ids))$. We can only learn the instantiation of the outer $(:)$ by looking at its second
 209 argument $(id : ids)$, which is a call. It would be a shame if simple call nesting broke type inference.

210 However, allowing nested calls is (currently) where we stop: if you put a list comprehension as
 211 an argument, quick look will ignore that argument. Allowing calls seems to be a sweet spot. One
 212 could go further, but the cost/benefit trade-off seems much less attractive.

213 The alert reader will note that the algorithm has complexity quadratic in the depth of function
 214 call nesting. In our example $(id : (id : ids))$ the depth was two, but if there were many elements in
 215 the list, each nested call would take a quick look into its argument, with cost linear in the depth of
 216 that argument. An easy (albeit ad-hoc) fix is to use a simple depth bound, because it is always safe
 217 to return no advice.

219 3.3 Uncurried functions

221 We have focused on exploiting n -ary calls of curried functions, but Quick Look works equally
 222 well on *uncurried* functions. For example, suppose $cons :: \forall p.(p, [p]) \rightarrow [p]$, and we have the
 223 call $cons (id, ids)$. Quick Look only has one argument to consult, namely a nested call to the pair
 224 constructor. So again, supporting nested calls is necessary, and it rapidly discovers that the only
 225 possible instantiation is $p := \forall a.a \rightarrow a$.

227 3.4 Interim Summary

228 We have described Quick Look at a high level. Its most prominent features are:

- 230 • A new “quick-look” phase is introduced, which guides instantiation of a call based on the
 231 context of the call: its arguments and expected result type.
- 232 • The quick-look phase is a modular addition. It guides at call sites, but the entire inference
 233 algorithm is otherwise undisturbed. That is in sharp contrast to earlier approaches, which
 234 have a pervasive effect throughout type inference. It seems plausible, therefore, that the
 235 quick-look approach would work equally well in other languages with very different type
 236 inference engines.

238 4 BIDIRECTIONAL, HIGHER-RANK INFERENCE

239 We begin our formalisation by giving a solid baseline, closely based on *Practical type inference for*
 240 *arbitrary-rank types* [18], which we abbreviate PTIAT. We simplify PTIAT by omitting the so called
 241 “deep skolemisation” and instantiation, and covariance and contravariance in function arrows, a
 242 choice we discuss in Section 5.4. We handle function application in an unusual way, one that will
 243 extend nicely for Quick Look, and we add *visible type application* [7].

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

Type constructors	\ni	F, G, T, \dots	Includes (\rightarrow)
Type variables	\ni	a, b, \dots	
Term variables	\ni	x, y, f, g, \dots	
Instantiation variables	\ni	κ, μ, v	
Polymorphic types	$\sigma, \phi ::=$	$\kappa \mid \forall a. \sigma \mid \rho$	
Top-level mono. types	$\rho ::=$	$\tau \mid T \bar{\sigma}$	
Fully mono. types	$\tau ::=$	$a \mid T \bar{\tau}$	
Typechecking direction	$\delta ::=$	$\uparrow \mid \downarrow$	Inference and checking respectively
Application heads	$h ::=$	x	Variable
		$\mid e :: \sigma$	Annotation
		$\mid e$	(not an application)
Arguments	$\pi ::=$	$\sigma \mid e$	
Terms / expressions	$e ::=$	$h \pi_1 \dots \pi_n$	Application ($n \geq 0$)
		$\mid \lambda x. e$	Abstraction
Match flag	$\omega ::=$	$ma \mid mnv$	Match-any and match-non-var resp.
Environments	$\Gamma ::=$	$\epsilon \mid \Gamma, x : \sigma$	
Instantiation sets	$\Delta ::=$	$\epsilon \mid \Delta, \kappa$	
Mono-substitutions	$\theta, \psi ::=$	$[\alpha := \tau]$	$\text{fiv}(\sigma)$ Free instantiation variables of σ
Poly-substitutions	$\Theta, \Psi ::=$	$[\kappa := \sigma]$	$\text{dom}(\theta)$ Domain of θ
			$\text{rng}(\theta)$ Range of θ
			$\Delta_1 \# \Delta_2$ Δ_1 is disjoint from Δ_2

Fig. 2. Syntax


4.1 Syntax

The syntax of our language is given in Figure 2.

Types. The syntax of types is unsurprising. Type constructors T include the function arrow (\rightarrow), although we usually write it infix. So $(\tau_1 \rightarrow \tau_2)$ is syntactic sugar for $((\rightarrow) \tau_1 \tau_2)$. A top-level monomorphic type, ρ , has no *top-level* forall, but may contain nested forall; while a fully monomorphic type, or *monotype*, τ , has no forall anywhere. Notice that in a polytype σ the forall can occur arbitrarily nested, including to the left or right of a function arrow. However, a *top-level monomorphic type* ρ has no forall at the top.

Terms. In order to focus on impredicativity, we restrict ourselves to a tiny term language: just the lambda calculus plus type annotations. We do not even support **let** or **case**. However, nothing essential is thereby omitted. A major feature of Quick Look is that it is completely localised to typing applications. It is fully compatible with, and leaves entirely unaffected, all other aspects of the type system, including ML-style **let**-generalisation, pattern matching, GADTs, type families, type classes, existentials, and the like (Section 5.5).

Similar to other works on type inference [5, 11, 26] our syntax uses n -ary application. The term $(h \pi_1 \dots \pi_n)$ applies a *head*, h , to a sequence of zero or more arguments $\pi_1 \dots \pi_n$. The head can be a variable x , an expression with a type annotation ($e :: \sigma$), or an expression e other than an application. The intuition is that we want to use information from the arguments to inform instantiation of the function's polymorphic variables. In fact, GHC's implementation *already* treats



application as an n -ary operation to improve error messages. Note also that a lone variable x is a valid expression e ; it is just an n -ary application with no arguments.

An argument π is either a *type argument* σ or a *value argument* e . Type arguments allow the programmer to *explicitly* instantiate the quantified variables of the function (Section 4.4).

4.2 Bidirectional Typing Rules

The typing rules for our language are given in Figures 3 and 4. Following PTIAT, to support higher rank types the typing judgment for terms is *bidirectional*, with two forms: one for checking and one for inference.

$$\Gamma \vdash_{\Downarrow} e : \rho \quad \Gamma \vdash_{\Uparrow} e : \rho$$

The first should be read “in type environment Γ , check that the term e has type ρ ”. The second should be read “in type environment Γ , the term e has inferred type ρ ”. Notice that in both cases the type ρ has no top-level quantifiers, but for checking ρ is considered as an *input* while for inference it is an *output*. When a rule has \vdash_{δ} in its conclusion, it is shorthand for two rules, one for \vdash_{\Uparrow} and one for \vdash_{\Downarrow} .

For example, rule ABS- \Uparrow deals with a lambda $(\lambda x. e)$ in inference mode. The premise extends the environment Γ with a binding $x : \tau$, for some monotype τ , and infers the type of the body e , returning its type ρ . Then the conclusion says that the type of the whole lambda is $\tau \rightarrow \rho$. Note that in inference mode the lambda-bound variable must have a monotype. A term like $\lambda x. (x \text{ True}, x \ 3)$ is ill-typed in inference mode, because x (being monomorphic) cannot be applied both to a Boolean and an integer. As is conventional, the type τ appears “out of thin air”. When constructing a typing

$$\begin{array}{c}
\boxed{\begin{array}{c} \vdash^{\text{inst}} \sigma [\bar{\pi}] \rightsquigarrow \Delta ; \bar{\phi}, \rho_r \\ \text{Invariants: } \Delta = \text{fiv}(\bar{\phi}, \rho_r) \end{array}} \\
\frac{\emptyset \vdash^i \sigma [\bar{\pi}] \rightsquigarrow \Delta ; \theta ; \bar{\phi}, \rho_r}{\vdash^{\text{inst}} \sigma [\bar{\pi}] \rightsquigarrow (\Delta \setminus \text{dom}(\theta)) ; \bar{\theta}\bar{\phi}, \rho_r} \text{INST} \\
\boxed{\begin{array}{c} \Delta \vdash^i \sigma [\bar{\pi}] \rightsquigarrow \Delta' ; \theta ; \bar{\phi}, \rho_r \\ \text{Invariants: } \text{dom}(\theta)\#\Delta', \text{dom}(\theta) \cup \Delta' = \text{fiv}(\bar{\phi}, \rho_r), \text{fiv}(\text{rng}(\theta)) \subseteq \Delta', \text{dom}(\theta)\#\text{fiv}(\rho_r) \end{array}} \\
\frac{}{\Delta \vdash^i \rho_r [] \rightsquigarrow \Delta ; \emptyset ; \rho_r} \text{IREsULT} \\
\frac{(\bar{\pi} = \epsilon \text{ or } \bar{\pi} = e \bar{\pi}') \quad \kappa \text{ fresh} \quad \Delta, \kappa \vdash^i ([a := \kappa]\rho) [\bar{\pi}] \rightsquigarrow \Delta' ; \theta ; \bar{\phi}, \rho_r}{\Delta \vdash^i (\forall a. \rho) [\bar{\pi}] \rightsquigarrow \Delta' ; \theta ; \bar{\phi}, \rho_r} \text{IALL} \\
\frac{\Delta \vdash^i (\rho[a := \sigma]) [\bar{\pi}] \rightsquigarrow \Delta' ; \theta ; \bar{\phi}, \rho_r}{\Delta \vdash^i (\forall a. \rho) [\sigma \bar{\pi}] \rightsquigarrow \Delta' ; \theta ; \bar{\phi}, \rho_r} \text{ITYARG} \\
\frac{\Delta \vdash^i \sigma_2 [\bar{\pi}] \rightsquigarrow \Delta' ; \theta ; \bar{\phi}, \rho_r}{\Delta \vdash^i (\sigma_1 \rightarrow \sigma_2) [e_1 \bar{\pi}] \rightsquigarrow \Delta' ; \theta ; \sigma_1, \bar{\phi}, \rho_r} \text{IARG} \\
\frac{\kappa \in \Delta \quad \mu, v \text{ fresh} \quad \theta_\kappa = [\kappa := (\mu \rightarrow v)] \quad (\Delta \setminus \kappa), \mu, v \vdash^i v [\bar{\pi}] \rightsquigarrow \Delta' ; \theta ; \bar{\phi}, \rho_r}{\Delta \vdash^i \kappa [e_1 \bar{\pi}] \rightsquigarrow \Delta' ; \theta \circ \theta_\kappa ; \mu, \bar{\phi}, \rho_r} \text{IVARD}
\end{array}$$

Fig. 4. Base instantiation

derivation we are free to use any τ , but of course only a suitable choice will lead to a valid typing derivation.

Rule $\text{ABS-}\Downarrow$ handles a lambda in checking mode. The type being pushed down must be a function type $\sigma_a \rightarrow \sigma_r$; we just extend the environment with $x : \sigma_a$ and check the body. Note that in checking mode the lambda-bound variable *can* have a polytype, so the lambda term in the previous paragraph *is* typeable. Notice that in $\text{ABS-}\Downarrow$ the return type σ_r may have top-level quantifiers. The judgment $\vdash_{\Downarrow}^{\forall} e : \sigma$, in Figure 3, deals with this case by adding the quantifiers to Γ before checking the expression against ρ .

The motivation for bidirectionality is that in checking mode we may push down a polytype, and thereby (as we have seen in $\text{ABS-}\Downarrow$) allow a lambda-bound variable to have a polytype. But how do we first *invoke* the checking judgment in the first place? One occasion is in rule H-ANN , where we have an explicit, user-written type signature. The second main occasion is in a function application, where we push the type expected by the function into the argument, as we show next.

4.3 Applications, Instantiation, and Subsumption

A function application with n arguments (including $n = 0$) is dealt with by rule APP , whose premises perform these five steps:

- 393 (1) Infer the (polymorphic) type σ of the function h , using \vdash_{\uparrow}^h . Usually the function is a variable
 394 x , and in that case we simply look up x in the environment Γ (rule H-VAR in Figure 3).
 395 (2) Instantiate h 's type σ with fresh *instantiation variables*, κ, μ, \dots , guided by the arguments $\bar{\pi}$
 396 to which it is applied, using the judgment \vdash^{inst} . This judgement returns
 397 • a type ϕ_i for each of the value arguments, $e_1 \dots e_n$;
 398 • the top-level-monomorphic result type ρ_r of the call; and
 399 • the *instantiation set* Δ of the freshly-instantiated variables.

400 We give some examples below.

- 401 (3) Conjure up a “magic substitution” θ that maps each of the instantiation variables in Δ to a
 402 monotype. At this point, the instantiation variables have completely disappeared.
 403 (4) Check that each value argument e_i has the expected type $\theta\phi_i$. Note that $\theta\phi_i$ can be an
 404 arbitrary polytype, which is pushed into the argument, using the checking judgment $\vdash_{\downarrow}^{\forall}$.
 405 Using the (instantiated) function type to specify the type of each argument is the essence of
 406 PTIAT.
 407 (5) Checks that the result type of the call, $\theta\phi_r$ fits the expected type ρ ; that is $\rho = \theta\phi_r$

408 The instantiation judgment, shown in Figure 4, has the form

$$409 \quad \vdash^{\text{inst}} \sigma [\bar{\pi}] \rightsquigarrow \Delta ; \bar{\phi}, \rho_r$$

410 It implements step (2) by instantiating σ , guided by the arguments $\pi_1 \dots \pi_n$. The type σ and
 411 arguments $\bar{\pi}$ should be considered inputs; the instantiation set Δ , argument types $\bar{\phi}$, and result
 412 type ρ_r are outputs. For example,

$$413 \quad \vdash^{\text{inst}} (\forall a.b \rightarrow b \rightarrow b) [\text{True}] \rightsquigarrow \{\kappa, v\} ; \kappa, (v \rightarrow v)$$

414 The instantiation judgement tells us that we can treat the type of the head of the application as
 415 $\forall \Delta. \bar{\phi} \rightarrow \rho_r$, and apply that to the value arguments of $\bar{\pi}$, ignoring type arguments which are already
 416 embodied in $\bar{\phi}$ and ρ_r . Moreover, the value arguments of $\bar{\pi}$ correspond 1-1 with the argumennt
 417 types $\bar{\phi}$.

418 During instantiation, rule IALL instantiates a leading \forall , while rule IARG decomposes a function
 419 arrow. We discuss ITYARG in Section 4.4. When the argument list is empty, IRESULT simply returns
 420 the result type ρ_r . Note that the rules deal correctly function types that have a forall nested to the
 421 right of an arrow, e.g. $f :: \text{Int} \rightarrow \forall a.[a] \rightarrow [a]$. For example,

$$422 \quad \vdash^{\text{inst}} (\text{Int} \rightarrow \forall a.[a] \rightarrow [a]) [3, x] \rightsquigarrow \{\kappa\} ; \text{Int}, [\kappa], [\kappa]$$

423 Somewhat unusually, the instantiation judgement returns an *instantiation set* Δ , the set of fresh
 424 instantiation variables it created. The set Δ starts off empty, is augmented in IALL, and returned by
 425 IRESULT. Then after instantiation, in step (3), rule APP conjures up a monomorphic substitution θ
 426 for the instantiation variables in Δ , and applies it to $\bar{\phi}$ and ϕ_r . Just like the τ in ABS- \uparrow , this θ comes
 427 “out of thin air”.

428 You may wonder why we did not simply instantiate with arbitrary monotypes in the first place,
 429 and dispense with instantiation variables, and with θ . That would be simpler, but dividing the
 430 process in two will allow us to inject Quick Look between steps 2 and 3.

431 Finally, rule IVAR deals with the case that the function type ends in a type variable, but there is
 432 another argument to come. For example, consider (*id inc* 3). We instantiate *id* with κ , so (*id inc*)
 433 has type κ ; that appears applied to 3, so we learn that κ must be $\mu \rightarrow v$. We express that knowledge
 434 by extending a local substitution θ , and adding μ, v to the instantiation set Δ . So we have

$$435 \quad \vdash^{\text{inst}} (\forall a.a \rightarrow a) [\text{inc}, 3] \rightsquigarrow \{\mu, v\} ; (\mu \rightarrow v), \mu, v$$

$$\begin{array}{c}
\boxed{\Gamma \vdash_{\uparrow} e : \rho} \quad \boxed{\Gamma \vdash_{\downarrow} e : \rho} \\
\Gamma \vdash_{\uparrow}^h h : \sigma \quad \bar{e} = \text{valargs}(\bar{\pi}) \quad \mu^{\text{inst}} \sigma [\bar{\pi}] \rightsquigarrow \Delta ; \bar{\phi}, \rho_r \\
\overline{\Gamma \vdash_{\downarrow} e_i : (\Delta, \phi_i) \rightsquigarrow \Theta_i} \\
\Theta_r = \text{if } \delta = \uparrow \text{ then } \emptyset \text{ else match}(\Delta, \rho_r, \rho, \text{mnv}) \\
\Psi = \bigoplus \Theta_i \oplus \Theta_r \quad \overline{\phi'_i = \Psi \phi_i} \quad \forall \bar{a}. \rho'_r = \Psi \rho_r \\
\overline{\text{dom}(\theta) = \bar{a} \cup \text{fiv}(\text{rng}(\Psi))} \quad \overline{\Gamma \vdash_{\downarrow}^{\forall} e_i : \theta \phi'_i} \quad \rho = \theta \rho'_r \\
\hline
\Gamma \vdash_{\delta} h \bar{\pi} : \rho \quad \text{APP}
\end{array}$$

Fig. 5. Quick look impredicativity: the new APP rule

4.4 Visible Type Application

The μ^{inst} judgement also implements visible type application (VTA) [7], a popular extension offered by GHC. The programmer can use VTA to *explicitly* instantiate a function call. For example, if $xs :: [Int]$ we could say either (*head xs*) or, using VTA, (*head @Int xs*).

Adding VTA has an immediate payoff for impredicativity: an explicit type argument can be a *polytype*, thus allowing explicit impredicative instantiation of any call. This is not particularly convenient for the programmer – the glory of Damas-Milner is that instantiation is silent – but it provides a fall-back that handles all of System F.

More precisely, rule ITYARG (Figure 4) deals with a visible type argument, by using it to instantiate the forall². The argument is a polytype σ : we allow impredicative instantiation. For example, consider the call (*map @(\forall a.a \rightarrow a) f*), where we supply one of the two type arguments that *map* expects (its type is in Figure 1). The instantiation judgement will then look like:

$$\begin{array}{c}
\mu^{\text{inst}} (\forall p q.(p \rightarrow q) \rightarrow [p] \rightarrow [q]) [@(\forall a.a \rightarrow a), f] \\
\rightsquigarrow \{ \kappa \} ; ((\forall a.a \rightarrow a) \rightarrow \kappa), [\forall a.a \rightarrow a], [\kappa]
\end{array}$$

Here we end up with just one instantiation variable κ , which instantiates q ; the other quantifier p is directly instantiated by the supplied type argument.

The attentive reader may note that our typing rules are sloppy about the lexical scoping of type variables (for example in rule GEN), so that they can appear in user-written type signatures or type arguments. Doing this properly is not hard, using the approach of Eisenberg et al. [6], but the plumbing is distracting so we omit it.

5 QUICK LOOK IMPREDICATIVITY

Building on the baseline of Section 4, we are now ready to present Quick Look, the main contribution of this paper.

The sole change to the typing rules for expressions is the shaded portion of rule APP (Figure 5). We take a quick look at each expression argument, and at the result type, each of which returns a poly-substitution Θ ; then we combine all those Θ 's with \oplus , and apply the resulting poly-substitution Ψ to the argument types $\bar{\phi}$ and result type ρ_r .

² Note that we do not address all the details of the design of Eisenberg et al. [7]. In particular, we do not account for the difference between “specified” and “inferred” quantifiers.

491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

QL argument $\Gamma \vdash_{\frac{1}{2}} e : (\Delta, \phi) \rightsquigarrow \Theta$

Invariants: $\text{dom}(\Theta) \subseteq \Delta \subseteq \text{fv}(\phi)$

$$\frac{\Gamma \vdash_{\frac{1}{2}}^h h : \sigma \quad \frac{\text{inst } \sigma[\bar{e}] \rightsquigarrow \Delta_\phi ; \bar{\phi}, \rho_r \quad \overline{\Gamma \vdash_{\frac{1}{2}} e_i : (\Delta_\phi, \phi_i) \rightsquigarrow \Theta_i}}{\Theta_\phi = \bigoplus \Theta_i \quad \omega = \text{if } \Delta_\phi = \emptyset \text{ then ma else mnv}}}{\Gamma \vdash_{\frac{1}{2}} h \bar{e} : (\Delta_\rho, \rho) \rightsquigarrow \text{match}(\Delta_\rho, \rho, \Theta_\phi \rho_r, \omega)} \text{APP-}\frac{1}{2}$$

$$\frac{e \text{ not an app. or } \phi \text{ not top-level mono.}}{\Gamma \vdash_{\frac{1}{2}} e : (\Delta, \phi) \rightsquigarrow \emptyset} \text{REST-}\frac{1}{2}$$

QL head $\Gamma \vdash_{\frac{1}{2}}^h h : \sigma$

$$\frac{x : \sigma \in \Gamma}{\Gamma \vdash_{\frac{1}{2}}^h x : \sigma} \text{H-VAR-}\frac{1}{2} \quad \frac{}{\Gamma \vdash_{\frac{1}{2}}^h (e :: \sigma) : \sigma} \text{H-ANN-}\frac{1}{2}$$

match($\Delta, \rho_1, \rho_2, \omega$) = Θ

$$\text{match}(\Delta, \rho_1, \rho_2, \omega) = \emptyset \quad \text{if } \omega = \text{mnv}, \rho_1 = \kappa \in \Delta \\ = \text{mch}(\Delta, \rho_1, \rho_2) \quad \text{otherwise}$$

mch(Δ, ρ_1, ρ_2) = Θ

$$\text{mch}(\Delta, \top \bar{\sigma}, \top \bar{\phi}) = \bigoplus (\text{mch}(\Delta, \sigma, \phi)) \\ \text{mch}(\Delta, \forall a. \sigma, \forall a. \phi) = \Theta \quad \text{if } a \notin \text{rng}(\Theta) \\ \text{where } \Theta = \text{mch}(\Delta, \sigma, \phi) \\ \text{mch}(\Delta, \kappa, \phi) = [\kappa := \phi] \quad \text{if } \kappa \in \Delta \\ \text{mch}(\Delta, \sigma, \phi) = \emptyset \quad \text{otherwise}$$

$\Theta_1 \oplus \Theta_2 = \Theta_3$

$$\Theta_1 \oplus \Theta_2 = [\kappa := \text{ljoin}(\Theta_1 \kappa, \Theta_2 \kappa) \mid \kappa \in \text{dom}(\Theta_1) \cup \text{dom}(\Theta_2)]$$

ljoin(σ_1, σ_2) = σ_3

$$\begin{aligned} \text{ljoin}(\top \bar{\sigma}, \top \bar{\phi}) &= \top \text{ljoin}(\sigma_i, \phi_i) & (1) \\ \text{ljoin}(\forall a. \sigma, \forall a. \phi) &= \forall a. \text{ljoin}(\sigma, \phi) & (2) \\ \text{ljoin}(\kappa, \phi) &= \begin{cases} \phi & \text{if } \phi \text{ contains a } \forall \\ \kappa & \text{otherwise} \end{cases} & (3) \\ \text{ljoin}(\phi, \kappa) &= \begin{cases} \phi & \text{if } \phi \text{ contains a } \forall \\ \kappa & \text{otherwise} \end{cases} & (4) \\ \text{ljoin}(\sigma, \phi) &= \sigma & \text{otherwise } (5) \end{aligned}$$

Fig. 6. Quick look impredicativity: supporting judgements

A poly-substitution Θ, Ψ maps instantiation variables (in Δ) to *polytypes* (Figure 2). It embodies the information deduced from each argument, and from the result type. It is always sound for Ψ to be empty, meaning “no useful quick-look information”.

5.1 A Quick Look at the Argument

The new APP rule is supported by the quick-look judgements in Figure 6. The main judgement takes a quick look at an argument e :

$$\Gamma \vdash_{\zeta} e : (\Delta, \phi) \rightsquigarrow \Theta$$

Here e is the argument, ϕ is the type expected by the function, and Θ is advice (in the form of a substitution for some of the variables in Δ) gained from comparing e with ϕ . The rules of this judgement are given in Figure 6. Rule $\text{REST-}\frac{\zeta}{\zeta}$ is the base case for \vdash_{ζ} . It does nothing, returning the empty substitution, for any complicated argument: list comprehensions, case expressions, lambdas, etc etc. Only in the case of a function application $h \bar{e}$ (including the case of a lone variable) does \vdash_{ζ} do anything at all (Section 3.2).

Looking now at rule $\text{APP-}\frac{\zeta}{\zeta}$, when the argument is itself a call ($h \bar{e}$) we use the same judgement again, recursively. This is perhaps confusing at first, so here is $\text{APP-}\frac{\zeta}{\zeta}$ specialised to a lone variable x :

$$\frac{x : \forall \bar{a}. \rho \in \Gamma \quad \Delta_{\phi} = \bar{\kappa} \text{ fresh} \quad \rho_r = [\bar{a} := \bar{\kappa}] \rho \quad \omega = \text{if } \Delta_{\phi} = \emptyset \text{ then ma else mnv}}{\Gamma \vdash_{\zeta} x : (\Delta, \rho) \rightsquigarrow \text{match}(\Delta, \rho, \rho_r, \omega)}$$

We look up x in Γ , instantiate any top-level quantifiers, and use *quick-look matching*, $\text{match}(\Delta, \rho, \rho_r, \omega)$, to find a substitution Θ , over variables in Δ , that makes ρ look like ρ_r .

The quick-look matching function $\text{match}(\Delta, \rho_1, \rho_2, \omega)$, defined in Figure 6, returns a poly-substitution Θ (Figure 2), from variables in Δ to polytypes σ . We discuss the first equation in Section 5.2. The second equation delegates to mch , which recursively traverses the two types to find a matching substitution, returning the empty substitution on a mis-match. It uses Θ to combine poly-substitutions from different parts of the type, just as we do in rule APP to combine poly-substitutions from different arguments of the function.

For example, when typechecking the call (*head ids*), we instantiate *head*'s type with κ , and take a quick look at the argument *ids*. For that, we look up $\text{ids} :: [\forall a.a \rightarrow a]$ in Γ , and match *head*'s expected argument type $[\kappa]$ against it. That yields $\Theta = [\kappa := \forall a.a \rightarrow a]$.

It is always sound for quick-look matching to return an empty poly-substitution, $\Theta = \emptyset$. That simply means that rule APP must find a mono-substitution θ that makes the program typecheck, just as was the case previously (Figure 3). The quick look step finds whatever polymorphic instantiation information it can; but it never fails.

5.2 Naked variables

What about the first equation of match , and the mysterious parameter ω ? Consider the call (*single id*). The argument type of $\text{single} :: \forall p.p \rightarrow [p]$ is a bare variable; after instantiation it will be, say, κ . The actual argument has type $\forall a.a \rightarrow a$. So should we instantiate κ with $\forall a.a \rightarrow a$? Or should we instantiate the argument *id*, and *then* bind its type to κ ? Since there is more than one choice, the first equation of match bails out, returning an empty substitution \emptyset . Perhaps one of the other arguments of the application will force an unambiguous choice for κ , as indeed happens in the case of (*id : ids*).

This behaviour is controlled by match 's fourth parameter ω . If $\omega = \text{mnv}$ – short for “match non-variable” – and ρ_1 is a naked variable κ , match returns the empty substitution. When do we

589 need to trigger this behaviour? Answer: when the argument is polymorphic, expressed by setting
 590 $\omega = \text{mnv}$ if $\Delta_\phi \neq \emptyset$ in APP- $\frac{1}{2}$.

591 If, on the other hand, there is no instantiation in the argument ($\Delta_\phi = \emptyset$), APP- $\frac{1}{2}$ sets $\omega = \text{ma}$,
 592 indicating that match can proceed regardless of whether ρ_1 is a naked variable. For example, in the
 593 call (*single ids*) we will perform the quick-look match

$$594 \quad \text{match}(\{\kappa\}, \kappa, [\forall a. a \rightarrow a], \text{ma})$$

595
 596 Here $\omega = \text{ma}$ disables the first equation of match, so the match returns $\Theta = [\kappa := [\forall a. a \rightarrow a]]$.

597 A very similar issue arises in rule APP (Figure 5). In checking mode ($\delta = \Downarrow$), we compute Θ_r by
 598 doing a quick-look match between the expected result type ρ and the actual result type ρ_r . For
 599 example, consider the term

$$600 \quad (\text{single id}) :: [\forall a. a \rightarrow a]$$

601 We instantiate *single* with κ , but (as we saw at the beginning of this subsection) we fail to learn
 602 anything from a quick look at the first argument. However, the call to $\text{match}(\Delta, \rho_r, \rho, \text{mnv})$ in
 603 rule APP matches the result type $\rho_r = [\kappa]$ with the type expected by the context $\rho = [\forall a. a \rightarrow a]$,
 604 yielding $\Theta_r = [\kappa := \forall a. a \rightarrow a]$, as desired.

605 Why $\omega = \text{mnv}$ in this call? Consider (*head ids*) $:: \text{Int} \rightarrow \text{Int}$. After instantiating *head* with κ , a
 606 quick look at the argument to *head* will give $\Theta_1 = [\kappa := \forall a. a \rightarrow a]$. But we will independently do
 607 a quick look at the argument

$$608 \quad \Theta_r = \text{match}(\{\kappa\}, \kappa, \text{Int} \rightarrow \text{Int}, \text{mnv})$$

609
 610 If we instead used $\omega = \text{ma}$ we would compute $\Theta_r = [\kappa := \text{Int} \rightarrow \text{Int}]$, contradicting Θ_1 . The right
 611 thing to do is to use $\omega = \text{mnv}$, which declines to quick-look match if the result type ρ_r is a naked
 612 variable.

614 5.3 Some Tricky Corners

615 In this section we draw attention to some subtler details of the typing rules.

616
 617 *Quick look does not replace “proper” typechecking.* In rule APP, we take a quick look at the
 618 arguments, with μ^{inst} , but we *also* do a proper typecheck of each argument in the final premise,
 619 using $\vdash_{\Downarrow}^{\forall} e_i : \phi_i$. This modular separation allows plenty of leeway in the quick-look part, because
 620 once the instantiation is chosen, the arguments will always be typechecked as normal.

621
 622 *Calls as arguments.* In rule APP- $\frac{1}{2}$, we generalise the rule given in Section 5.1 to work for calls as
 623 well as for lone variables. We can do that very simply by recursively invoking μ^{inst} . Our only goal is
 624 to discover the return type $\Theta_\phi \rho_r$ to use in the quick-look match in the conclusion. Rule APP- $\frac{1}{2}$ sets
 625 ω based on whether Δ_ϕ is empty, which neatly checks for instantiation *anywhere* in the argument.
 626 For example consider (*single (f 3)*) where $f :: \text{Int} \rightarrow \forall a. a \rightarrow a$, where the instantiation is hidden
 627 under an arrow.

628
 629 *Instantiation after quick look.* In rule APP while ρ_r is top-level monomorphic, Θ_{ρ_r} may not be.
 630 For example, consider (*head ids*). Instantiating *head* with $[p := \kappa]$ gives $\rho_r = \kappa$. But a quick look at
 631 the argument gives $\Theta = [\kappa := \forall a. a \rightarrow a]$, so Θ_{ρ_r} is $\forall a. a \rightarrow a$. Hence APP is careful to decompose
 632 Θ_{ρ_r} to $\forall \bar{a}. \rho'_r$, and then includes \bar{a} in the domain of the “magic substitution” θ . Note that, since θ is
 633 a mono-substitution, these \bar{a} are instantiated only with *monotypes*.

634
 635 *Quick look binds only instantiation variables.* In match (Figure 6), we find a substitution Θ that
 636 matches ρ_1 with ρ_2 , *binding only instantiation variables in Δ* . The sole goal is to find the best type
 637 with which to instantiate the call; we leave all other variables unaffected.

638 *Joining the results of independent quick looks.* The operator \oplus , defined in Figure 6, uses a function
 639 ljoin to combine the substitutions from independent quick looks. A tricky example is the call $(f\ x\ y)$
 640 where

$$641 \quad f :: \forall a. a \rightarrow a \rightarrow \text{Int} \quad x :: \forall b. (b, \sigma_{id}) \quad y :: \forall c. (\sigma_{id}, c)$$

642 and σ_{id} is shorthand for $\forall a. a \rightarrow a$. Suppose we instantiate f with $[a := \kappa]$. A quick look at the two
 643 arguments then yields

$$644 \quad \Theta_x = [\kappa := (v_b, \sigma_{id})] \quad \Theta_y = [\kappa := (\sigma_{id}, v_c)]$$

645 where v_b, v_c are the instantiation variables used to instantiate x and y respectively. Joining these
 646 leads to the desired final $\Theta = [\kappa := (\sigma_{id}, \sigma_{id})]$. The first component of the tuple is obtained from
 647 equation (3) of ljoin , the second from equation (4).

648 The same holds within a single call of match . Consider the call $(g\ xy)$ where

$$649 \quad g :: \forall a. (a, a) \rightarrow \text{Int} \quad xy :: \forall bc. ((b, \sigma_{id}), (\sigma_{id}, c))$$

650 where we want to discover the instantiation $[\kappa := (\sigma_{id}, \sigma_{id})]$.

651 Note that ljoin never fails; if it encounters a mis-match the program is un-typeable anyway, and
 652 ljoin simply returns its first argument (equation (5)). A crucial property of ljoin is that makes Quick
 653 Look stable under a mono-substitution θ ; that is $\theta(\text{ljoin}(\sigma_1, \sigma_2)) = \text{ljoin}(\theta\sigma_1, \theta\sigma_2)$; see Lemma A.2
 654 in Appendix A. This property is a key part of our soundness proof.

655 5.4 Co- and Contravariance of Function Types

656 The presentation so far treats the function arrow (\rightarrow) uniformly with other type constructors T .
 657 Suppose that

$$658 \quad f :: (\forall a. \text{Int} \rightarrow a \rightarrow a) \rightarrow \text{Bool} \quad g :: \text{Int} \rightarrow \forall b. b \rightarrow b$$

659 Then the call $(f\ g)$ is ill-typed because we use equality when comparing the expected and actual
 660 result types in rule APP . The call would also be rejected if the forall in f and g 's types were the
 661 other way around. Only if they line up will the call be accepted. The function is neither covariant
 662 nor contravariant with respect to polymorphism; it is invariant.

663 We make this choice for three reasons. First, and most important for this paper, treating the
 664 function arrow uniformly means that it acts as a guard, which in turn allows more impredicative
 665 instantiations to be inferred. For example, without an invariant function arrow (app runST argST)
 666 cannot be typed.

667 Second, such mismatches are rare (we give data in Section 9), and even when one occurs it can
 668 readily be fixed by η -expansion. For example, the call $(f\ (\lambda x. g\ x))$ is accepted regardless of the
 669 position of the forall.

670 Finally, as well as losing guardedness, co/contra-variance in the function arrow imposes other
 671 significant costs. One approach, used by GHC, is to perform automatic η -expansion, through so-
 672 called “deep skolemisation” and “deep instantiation” [18, §4.6]. But, aside from adding significant
 673 complexity to the type system, this automatic η -expansion changes the semantics of the program
 674 (both in call-by-name and call-by-need settings), which is highly questionable.

675 Instead of *actually* η -expanding, one could make the type system behave *as if* η -expansion had
 676 taken place. This would, however, impact the compiler’s intermediate language. GHC elaborates the
 677 source program to a statically-typed intermediate language based on System F; we would have to
 678 extend this along the lines of Mitchell’s System F η [15], a major change that would in turn impact
 679 GHC’s entire downstream optimisation pipeline.

680 In short, an invariant function arrow provides better impredicative inference, costs the program-
 681 mer little, and makes the type system significantly simpler. Indeed, a GHC Proposal to simplify the
 682
 683
 684
 685
 686

language by adopting an invariant function arrow has been adopted by the community, independently of impredicativity [17]. In Section 9 we quantify the impact of this change in the broader Haskell ecosystem.

5.5 Modularity

Quick Look is like many other works in that it exploits programmer-supplied type annotations to guide type inference (Section 10). But Quick Look’s truly distinctive feature is that it is *modular* and *highly localised*.

Highly localised. Through half-closed eyes the changes in Figures 5 and 6 seem substantial. However, rule APP is the sole rule of the expression judgement that is changed. When scaling up to all of Haskell, the expression judgement in Figure 3 gains dozens and dozens of rules, one for each syntactic construct. But the only change to support quick-look impredicativity remains rule APP. So Quick Look scales well to a very rich source language.

Modular. This paper has presented only a minimalistic type system, but GHC offers many, many more features, including **let**-generalisation, data types and pattern matching, GADTs, existentials, type classes, type families, higher kinds, quantified constraints, kind polymorphism, dependent kinds, and so on. GHC’s type inference engine works by generating constraints solving them separately, and elaborating the program into System F [25]. *All of these extensions, and the inference engine that supports them, are unaffected by Quick Look.* Indeed, we conjecture the Quick Look would be equally compatible with quite different type systems, such as ones involving subtyping, or dependent object types.

To substantiate these claims, Appendix B gives the extra rules for a much larger language; and we have built a full implementation in GHC (Section 8). This implementation is the first working implementation of impredicativity in GHC, despite several attempts over the last decade, each of which became mired in complexity.

6 PROPERTIES OF QUICK LOOK

In this section we give a comprehensive account of various properties of Quick Look.

6.1 Expressiveness and Backward Compatibility

Our system is able to type any program typeable in System F, maybe with additional annotations. In order to do so, we define a type-directed translation from System F into our source language, in a very similar fashion to Serrano et al. [24]. Every variable, application, and abstraction is recursively translated, and in addition:

- A System F type application ($e @\tau$) is translated as ($e' @\tau$), where e' is the translation of e .
- A System F type abstraction ($\lambda a. e$) is translated as the annotated term ($e' :: \forall a. \sigma$), where e' is the recursive translation of e , and σ is the type of e .

THEOREM 6.1 (EMBEDDING OF SYSTEM F). *Let e be a well-typed System F expression with type σ under an environment Γ , and e' the translation as defined above. Then $\Gamma \vdash_{\Downarrow} e' : \sigma$.*

The inverse translation, from our language into System F, is also simple to define. In particular, uses of the \vdash^{inst} judgment translate into type applications, and uses of $\vdash_{\Downarrow}^{\forall}$ translate into type abstractions. The fact that we elaborate to System F – a provably-sound system – means that the proposed system is sound.

One of our stated goals was to be a conservative extension of Damas-Milner polymorphism.

736 **THEOREM 6.2 (COMPATIBILITY WITH RANK-1 POLYMORPHISM).** *Let e be an expression with type τ*
 737 *under λ -calculus with predicative polymorphism, in an environment Γ whose types only have top-level*
 738 *polymorphism. Then $\Gamma \vdash e : \tau$ in the system presented in this paper.*

739 The theorem holds because, given the conditions on Γ , all the argument types $\overline{\phi}$ in rule APP will be
 740 moontypes, so the quick look will recover no useful information, and will effectively be a no-op.
 741 Technically, Damas-Milner also includes generalizing **let** bindings, but adding those in our system
 742 poses no technical challenges (Appendix B.)
 743

744 6.2 Relating Inference and Checking Mode

745 A desirable property of a type system is that if we can *infer* a type for a term then we can certainly
 746 *check* that the term can be assigned this type. The next theorem guarantees this:

747 **THEOREM 6.3.** *If $\Gamma \vdash_{\uparrow} e : \rho$ then $\Gamma \vdash_{\downarrow} e : \rho$.*

749 Although the statement of the theorem is easy, the inductive proof case for rule APP involves
 750 a delicate split on whether the QL from the arguments yields a polymorphic type or not. The
 751 subsequent QL on the result cannot cause any further instantiations.
 752

753 6.3 Stability under Transformations

754 In any type system it is desirable that simple program transformations do not make a working
 755 program fail to typecheck, or vice versa. In this section we give some such properties, using a
 756 slightly larger language than the one we presented in the paper.
 757

758 *Argument permutation.* One obviously-desirable property is that a program should be insensitive
 759 to permutation of function arguments. For example $app\ f\ x$ should be typeable iff $revapp\ x\ f$ is.

760 **CONJECTURE 6.4 (ARGUMENT PERMUTATION).** *Assume environment Γ and expressions \overline{e}_i . If:*

$$761 \Gamma, f : \forall \overline{a}. \phi_1 \rightarrow \cdots \rightarrow \phi_n \rightarrow \phi_r \vdash_{\delta} f\ \overline{e}_i : \sigma$$

762 *then for any permutation π of the indices i we have:*

$$763 \Gamma, f' : \forall \overline{a}. \phi_{\pi(1)} \rightarrow \cdots \rightarrow \phi_{\pi(n)} \rightarrow \phi_r \vdash_{\delta} f'\ \overline{e_{\pi(i)}} : \sigma$$

764 **PROOF.** (Sketch) We give a sketch of the argument, but leave a formal treatment as future work.
 765 We need two auxiliary properties:

- 766 • First, quick look never disagrees with the main typing judgement on expressions. More
 767 precisely, for any Γ, e and ϕ , let $\Delta = \text{fiv}(\phi)$, and let Θ be the quick-look substitution defined
 768 by $\Gamma \vdash_{\downarrow} e : (\Delta, \phi) \rightsquigarrow \Theta$. Then for any poly-substitution Ψ such that $\text{dom}(\Psi) = \Delta$ and
 769 $\Gamma \vdash_{\downarrow} e : \Psi(\phi)$, it must be that $\Psi = \Psi_r \cdot \Theta$, for some residual poly-substitution Ψ_r . That is, the
 770 quick-look substitution Θ is “on the way” to any valid instantiation of the function.
 771
- 772 • Second, we observe that if σ_1 and σ_2 are unifiable by a poly-substitution Ψ then $\text{ljoin}(\sigma_1, \sigma_2)$
 773 and $\text{ljoin}(\sigma_2, \sigma_1)$ can only differ in positions inside the two types where one contains an
 774 instantiation variable κ_1 and the other an instantiation variable κ_2 .
 775

776 With these two properties in mind, consider an application $(f\ e_1\ e_2)$ that is well-typed by rule APP
 777 in Figure 5. Assume we get back Θ_1 from quick-looking at e_1 , and Θ_2 from e_2 . If Ψ is the master
 778 substitution that completes the typing derivation in rule APP, then by the first property above it
 779 must be that $\Psi_r^1 \cdot \Theta_1 = \Psi = \Psi_r^2 \cdot \Theta_2$ for some residual substitutions Ψ_r^1 and Ψ_r^2 . That in turn means
 780 that the ranges of the two substitutions Θ_1 and Θ_2 must be unifiable. Now by the second property
 781 above, $\Theta_1 \oplus \Theta_2$ and $\Theta_2 \oplus \Theta_1$ can only differ by some free instantiation variables in their ranges.
 782 However these are going to be grounded to monomorphic types by the mono-substitution θ in rule
 783 APP, and so do not matter for impredicative instantiations. \square
 784

Unification variables	\ni	α, β, γ
Fully mono. types	$::=$	$\alpha \mid \kappa \mid a \mid \top \bar{\tau}$
Constraints	$::=$	$\epsilon \mid C \wedge C \mid \sigma \sim \phi \mid \forall \bar{a}. \exists \bar{a}. C$

Fig. 7. Extra syntax for inference

Let abstraction and inlining. One desirable property held by ML is let-abstraction:

$$\text{let } x = e \text{ in } b \equiv b[e/x]$$

This property does not hold in most higher-rank systems, including PTIAT, because they use the context of the call to guide the typing of the argument. For example, suppose that $f :: ((\forall a. a \rightarrow a) \rightarrow Int) \rightarrow Bool$. Then let-abstracting the argument can render the program ill-typed; but it can be fixed by adding a type signature:

$f (\lambda x. (x \text{ True}, x \ 3))$	Well typed
$\text{let } g = \lambda x. (x \text{ True}, x \ 3) \text{ in } f \ g$	Not well typed
$\text{let } g :: (\forall a. a \rightarrow a) \rightarrow Int$	
$g = \lambda x. (x \text{ True}, x \ 3)$	Well typed
$\text{in } f \ g$	

What about the opposite of let-abstraction, namely let-inlining? With PTIAT, inlining a let-binding always improves typeability, but not so for Quick Look. Suppose $f :: \forall a. (Int \rightarrow a) \rightarrow a$. Then we have

$f (\lambda x. \text{ids})$	Not well typed
$\text{let } g = \lambda x. \text{ids} \text{ in } f \ g$	Well typed
$f ((\lambda x \rightarrow \text{ids}) :: Int \rightarrow [\forall a. a \rightarrow a])$	Well typed

The trouble here is that Quick Look does not look inside lambda arguments. Again, a type signature makes the program robust to such transformation. In general,

$$\text{let } x = e :: \sigma \text{ in } b \equiv b[e :: \sigma / x]$$

η -expansion. Sometimes, as we have seen in Section 5.4, η -expansion is necessary to make a program typecheck. But sometimes the reverse is the case. For example, whereas app runST argST is accepted, $\text{app } (\lambda x. \text{runST } x) \text{ argST}$ is not. The problem is that the fact that app must be instantiated impredicatively comes solely from runST ; argST alone is not enough to know whether $\forall s. \text{ST } s$ must be left as it is or instantiated further.


7 TYPE INFERENCE

In this section we give a *type inference algorithm* that implements the specification given in Section 5, and discuss its soundness.

7.1 Inference algorithm

Typically type inference algorithms work in two stages: *generating* constraints, and then *solving* them, as described in *OutsideIn(X): Modular type inference with local assumptions* [25], which we abbreviate MTILA. To focus on impredicativity, we simplify MTILA by omitting local typing assumptions, along with data types, GADTs, and type classes – but our approach to impredicativity scales to handle all these features, as the full rules in Appendix B demonstrate.

834
835
836
837
838
839
840
841
842
843
844
845
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



Our algorithm generates *constraints* whose syntax is shown in Figure 7. Simple constraints are bags of *equality constraints* $\phi_1 \sim \phi_2$, but we will also need mixed-prefix constraints $\forall \bar{a}. \exists \bar{\alpha}. C$. These forms are not new; they are described in MTLA, and used in GHC's constraint solver. This is a key point of our approach: it requires zero changes to GHC's actual constraint language and solver.

883 *Algorithm: expressions.* Figures 8 and 9 present constraint generation for expressions. They closely
 884 follow the declarative specification in Figures 3 and 4. For example, the judgement $\Gamma \vdash_{\delta} e : \rho \rightsquigarrow C$
 885 is very similar to that in Figure 3, but in addition generates constraints C . The big difference is
 886 that instead of clairvoyantly selecting monomorphic types τ for λ -abstraction arguments and for
 887 other instantiations (e.g. the range of substitution θ in rule App) the constraint-generation rules
 888 create fresh *unification variables*, α, β, γ . Unification variables stand for monomorphic types, and
 889 are solved during a subsequent *constraint solving* pass. In contrast, instantiation variables stand
 890 for polytypes, and are solved immediately by Quick Look; the constraint solver never sees them.
 891 Hence the following invariant:

892 LEMMA 7.1. *If $\Gamma \vdash_{\delta} e : \rho \rightsquigarrow C$ (with $\text{fiv}(\Gamma) = \emptyset$) then $\text{fiv}(C) = \emptyset$.*

893 Rule GEN generates a mixed-prefix constraint (a degenerate *implication constraint* in the MTILA
 894 jargon), that encodes the fact that the unification variables $\bar{\alpha}$ generated in C are allowed to unify to
 895 types mentioning the bound variables \bar{a} .

896 Rules ABS- \Downarrow -VAR and ABS- \Uparrow generate fresh unification variables, as expected. Note that they
 897 preserve the invariant that environments and constraints only mention unification variables but
 898 never instantiation variables.

899 Rule APP follows very closely its declarative counterpart in Figure 5, with a few minor deviations.
 900 First, while rule APP in Figure 5 clairvoyantly selects a monomorphic θ to “monomorphise” any
 901 instantiation variables that are not given values by Quick Look, its algorithmic counterpart in
 902 Figure 8 generates fresh unification variables $\bar{\alpha}$. Second, rule APP in Figure 8 uses the \equiv judgement
 903 to generate further constraints; whereas the rule in Figure 5 has readily ensured that $\rho = \theta\rho_r$.
 904

905 *Algorithm: instantiation.* The algorithmic instantiation judgement in Figure 9 collects constraints
 906 that are generated in rule IVARM. In that case we have a unification variable α that we have to
 907 further unify to a function type $\beta \rightarrow \gamma$. Note how the constraint C only mentions unification *but*
 908 *no instantiation variables*. Instantiation variables κ are born and eliminated in a single quick look.
 909

910 *Algorithm: quick look, matching and joining.* These judgements remain exactly as in Figure 6, since
 911 we have carefully designed them to operate only on instantiation variables and ignore unification
 912 variables.

913 7.2 Soundness of Type Inference

914 We write $\theta \models C$ to mean that a *monomorphic* idempotent substitution θ (from any sort of variables)
 915 is a solution to constraint C ³. Due to Lemma 7.1, a solution θ to a constraint C need only refer to
 916 unification variables, but *never* instantiation variables, that must be resolved only through the QL
 917 mechanism. The main soundness theorem follows:
 918

919 THEOREM 7.2 (SOUNDNESS). *If $\Gamma \vdash_{\delta} e : \rho \rightsquigarrow C$, θ is a substitution from unification variables to*
 920 *monotypes, $\text{fiv}(\theta) = \emptyset$, and $\theta \models C$ then $\theta\Gamma \vdash_{\delta} e : \theta\rho$.*

921 The theorem relies on a chain of other lemmas for every auxiliary judgement used in our
 922 specification and the algorithm. One of the basic facts required for this chain of lemmas justifies
 923 our design of the \oplus operator on substitutions and the $\text{ljoin}(\phi_1, \phi_2)$ function:
 924


925 LEMMA 7.3. *If $\Theta^* = \bigoplus \bar{\Theta}$ and $\text{fiv}(\theta) = \emptyset$ then*

$$926 (\theta \cdot \Theta^*)|_{\text{dom}(\Theta^*)} = \bigoplus (\theta \cdot \Theta_i)|_{\text{dom}(\Theta_i)}$$

927 We additionally conjecture that completeness is true of our algorithm, but have not attempted a
 928 detailed proof.
 929

930 ³For implication constraints θ is a substitution nesting – see [24].
 931

932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980



8 IMPLEMENTATION


One of our main claims is that Quick Look can be added, in a modular and non-invasive way, to an existing, production-scale type inference engine. To substantiate the claim, we implemented Quick Look on top of the latest incarnation of GHC (ghc-8.9.0.20191101). The changes were straightforward, and were highly localised. Overall we added 800 lines, and removed 200 lines, from GHC's 90,000 line type inference engine (these figures include comment-only lines, since they are a good proxy for complexity). The implementation is publicly available.⁴

It is remarkable that there is so little to say about the implementation. By way of comparison, attempts to implement Guarded Instantiation [24] in GHC floundered in a morass of complexity.

9 THE IMPACT OF AN INVARIANT FUNCTION ARROW

In Section 5.4 we propose to make the function arrow invariant, and to drop deep instantiation and deep skolemisation. This change is independently attractive, and is the subject of a recently adopted GHC Proposal. How much effect does this change have on existing Haskell code?

⁴URL suppressed for review, but available on request.



To answer this question we used our implementation to compile a large collection of packages from Hackage; we summarize the results in Figure 10. We started from collection of packages for a recent version of GHC, obtained from Stackage (LTS 13.6), containing a total of 2,332 packages. We then selected only the 607 packages that used the extensions *RankNTypes*, *Rank2Types* or *ImpredicativeTypes*; the other 1,725 packages are certainly unaffected. We then compiled the library sections of each package. Of the 607 packages, 150 fail to compile for reasons unrelated to Quick Look – they depend on external libraries and tools we do not have available; or they do not compile with GHC 8.9 anyway. Of the remaining 457 packages, 399 compiled with no changes whatsoever; two had issues we could not solve, e.g., because of Template Haskell. We leave these two to the authors of these packages.

The remaining 56 packages could be made compilable with modest source code changes, almost all of which were a simple η -expansion on a line that was clearly identified by the error message. In total we performed 283 η -expansions in 104 of the total of 963 source files. The top three packages in this case needed 7, 7 and 9 files changed. The majority of the packages, 37, needed only one file to be changed, and 20 packages needed only a single η -expansion. In the case of two packages, *massiv* and *drinkery*, we additionally had to provide type signatures for local definitions.

In conclusion, of the 2,332 packages we started with, 74% (1,725/2,332) do not use extensions that interact with first-class polymorphism; of those that do, 87% (399/457) needed no source code changes; of those that needed changes, 97% (56/58) could be made to compile without any intimate knowledge of these packages. All but two were fixed by a handful of well-diagnosed η -expansions, two also needed some local type signatures.

10 RELATED WORK

This section explores many different strands of work on first-class polymorphism; another excellent review can be found in [3, §5]. Figure 11 compares the expressiveness of some of these systems, using examples culled from their papers. As you can see, QL performs well despite its simplicity.

Higher-rank polymorphism. Type inference for higher-rank polymorphism (in which forall can appear to the left or right of the function arrow) is a well-studied topic with successful solutions using bidirectionality [18]. Follow-up modern presentations [4] re-frame the problem within a more logical setting, and additionally describe extensions to indexed types [5].

Boxed polymorphism. Impredicativity goes beyond higher-rank, by allowing quantified types to instantiate polymorphic types and data structures. Both Haskell and OCaml have supported impredicative polymorphism, in an inconvenient form, for over a decade.

	QL	GI	MLF	HMF	FPH	HML
A POLYMORPHIC INSTANTIATION						
A1	const2 = $\lambda x y. y$	✓	✓	✓	✓	✓
	MLF infers $(b \geq \forall c. c \rightarrow c) \Rightarrow a \rightarrow b$, QL and GI infer $a \rightarrow b \rightarrow b$.					
A2	choose id	✓	✓	✓	✓	✓
	MLF and HML infer $(a \geq \forall b. b \rightarrow b) \Rightarrow a \rightarrow a$, FPH, HMF, QL, and GI infer $(a \rightarrow a) \rightarrow a \rightarrow a$.					
A3	choose [] ids	✓	✓	✓	✓	✓
A4	$\lambda(x :: \forall a. a \rightarrow a). x x$	✓	✓	✓	✓	✓
	MLF infers $(\forall a. a \rightarrow a) \rightarrow (\forall a. a \rightarrow a)$, QL and GI infer $(\forall a. a \rightarrow a) \rightarrow b \rightarrow b$.					
A5	id auto	✓	✓	✓	✓	✓
A6	id auto'	✓	✓	✓	✓	✓
A7	choose id auto	✓	✓	✓	No	No
A8	choose id auto'	No	No	✓	No	No
	QL and GI need an ann. on $id :: (\forall a. a \rightarrow a) \rightarrow (\forall a. a \rightarrow a)$.					
A9	f (choose id) ids	✓	Ext*	✓	No	✓
	where $f :: \forall a. (a \rightarrow a) \rightarrow [a] \rightarrow a$ GI needs an annotation on $id :: (\forall a. a \rightarrow a) \rightarrow (\forall a. a \rightarrow a)$.					
A10	poly id	✓	✓	✓	✓	✓
A11	poly ($\lambda x. x$)	✓	✓	✓	✓	✓
A12	id poly ($\lambda x. x$)	✓	✓	✓	✓	✓
B INFERENCE OF POLYMORPHIC ARGUMENTS						
B1	$\lambda f. (f 1, f \text{ True})$	No	No	No	No	No
	All systems require an annotation on $f :: \forall a. a \rightarrow a$.					
B2	$\lambda xs. \text{poly}(\text{head } xs)$	No	No	✓	No	No
	All systems except for MLF require annotated $xs :: [\forall a. a \rightarrow a]$.					
C FUNCTIONS ON POLYMORPHIC LISTS						
C1	length ids	✓	✓	✓	✓	✓
C2	tail ids	✓	✓	✓	✓	✓
C3	head ids	✓	✓	✓	✓	✓
C4	single id	✓	✓	✓	✓	✓
C5	id : ids	✓	✓	✓	No	✓
C6	$(\lambda x. x) : \text{ids}$	✓	✓	✓	No	✓
C7	single inc # single id	✓	✓	✓	✓	✓
C8	g (single id) ids	✓	No	✓	No	✓
	where $g :: \forall a. [a] \rightarrow [a] \rightarrow a$					
C9	map poly (single id)	✓	No	✓	✓	✓
	GI needs an ann. on $\text{single id} :: [\forall a. a \rightarrow a]$ in the previous two.					
C10	map head (single ids)	✓	✓	✓	✓	✓
D APPLICATION FUNCTIONS						
D1	app poly id	✓	✓	✓	✓	✓
D2	revapp id poly	✓	✓	✓	✓	✓
D3	runST argST	✓	✓	✓	✓	✓
D4	app runST argST	✓	✓	✓	✓	✓
D5	revapp argST runST	✓	✓	✓	✓	✓
E η-EXPANSION						
E1	$k h \text{ lst}$	No	No	No	No	No
E2	$k (\lambda x. h x) \text{ lst}$	✓	✓	✓	No	✓
	where $h :: \text{Int} \rightarrow \forall a. a \rightarrow a$, $k :: \forall a. a \rightarrow [a] \rightarrow a$, and $\text{lst} :: [\forall a. \text{Int} \rightarrow a \rightarrow a]$					
E3	$r (\lambda x y. y)$	✓	Ext*	✓	No	No
	where $r :: (\forall a. a \rightarrow \forall b. b \rightarrow b) \rightarrow \text{Int}$					

* "Ext" in GI are not accepted by the vanilla system, but are allowed with some of its extensions [24].

Fig. 11. Comparison of impredicative type systems

1079 In Haskell, one can wrap a polytype in a new, named data type or newtype, which then be-
1080 behaves like a monotype [16]. This *boxed polymorphism* mechanism is easy to implement, but the
1081 programmer has to declare new data types and explicitly box and unbox the polymorphic value.
1082 Nevertheless, boxed polymorphism is widely used in Haskell.

1083 OCaml supports *polymorphic object methods*, based on the theory of Poly-ML [9]. The programmer
1084 does not have to declare new data types, but polymorphic values must still be wrapped and
1085 unwrapped. In practice, the mechanism is little used, perhaps because it is only exposed through
1086 the object system.

1087 In FreezeML [8] the programmer chooses explicitly when to not instantiate a polymorphic type
1088 by using the freeze operator $[-]$. Another variant of this line of work is QML [23], which has
1089 two different universal quantifiers, one that can be implicitly instantiated and one that requires
1090 explicit instantiation. Introducing and eliminating the explicit quantifier is akin to the wrapping
1091 and unwrapping. We urge the reader to consult the related work section in the latter work for an
1092 overview of that line of research.

1093 Explicit wrapping and unwrapping is painful, especially since it often seems unnecessary, so our
1094 goal is to allow polymorphic functions to be implicitly instantiated with quantified types.

1095 *Guarded impredicativity.* Quick Look builds on ideas originating in previous work on Guarded
1096 Impredicative Polymorphism (GI) [24]. Specifically, GI figures out the polymorphic instantiation of
1097 variables that are “guarded” in the type of the instantiated function, or the types of the arguments;
1098 meaning they occur under some type constructor. However, GI relied on extending the constraint
1099 language with two completely new forms of constraints, one of which (delayed generalisation)
1100 was very tricky to conceptualise and implement. With Quick Look we instead eagerly figure out
1101 impredicative instantiations, quite separate from constraint solving, which can remain unmodified
1102 in GHC.

1103 HMF [11] comes the closest to QL in terms of expressiveness, a simple type vocabulary, and
1104 an equation-based unification algorithm. The key idea in HMF is that in an n -ary application we
1105 perform a subsumption between each function argument type and the type of the corresponding
1106 argument. The order in which to perform these subsumption checks is delicate: it is first performed
1107 for those arguments that correspond to an argument type that is guarded (i.e. not a naked type
1108 variable), and then on the other arguments in any order. This bears a strong similarity to our system
1109 and to GI.

1110 A key difference is that HMF is formulated to *generalize eagerly* (and solve impredicativity in one
1111 go), whereas GHC does exactly the opposite: it defers generalisation as long as possible in favour
1112 of generating constraints, only performing constraint-solving and generalisation when absolutely
1113 necessary. Switching to eager generalisation would be a deep change to GHC’s generate-and-solve
1114 approach to inference, and it is far from clear how it would work. On the other hand HMF comes
1115 with a high-level specification that enforces that polymorphism is not-guessed in the typing rules
1116 with a condition requiring that the types we assign to terms have minimal “polymorphic weights”.

1117 Another small technical difference is that our system explicitly keeps track of the variables it
1118 can unify to polytypes (with the Δ environments), whereas HMF allows arbitrary unification of
1119 any variable to a polytype when checking n -ary applications, accompanied with an after-the-fact
1120 check that no variable from the environment was unified to a polytype. This is merely a difference
1121 in the presentation and mechanics but not of fundamental nature.

1122 *Stratified inference.* The idea of a “quick-look” as a restricted pass prior to actual type inference
1123 has appeared before in the work on Stratified Type Inference (STI) [20, 22]. In the first pass, each
1124 term is annotated with a *shape*, a form of type that expresses the *quantifier structure* of the term’s
1125 eventual type, while leaving its monomorphic components as flexible unification variables that will
1126

1127

1128 be filled in by the (conventional, predicative) second pass. To avoid shortcomings with the order
 1129 in which arguments are checked this process may need to iterated [20, §7]. These works handle
 1130 higher-rank types and GADTs, but to date there has been no STI design for impredicativity. QL
 1131 has a similar flavor, but admittedly a more algorithmic specification – for example stratified type
 1132 inference does not use substitutions – rather *joins* of types in the polymorphic lattice. Nor does it
 1133 return substitutions *and* constraints. On the other hand, STI can only propagate explicitly *declared*
 1134 types whereas, by being part of type inference, QL can exploit the *inferred* types for top-level or
 1135 locally let-bound functions.

1136 *Beyond System F.* We have been reluctant to move beyond System F types, because doing so
 1137 impacts the language the programmers sees, the type inference algorithm, and the compiler’s
 1138 statically-typed intermediate language. However, once that Rubicon is crossed, there is a rich seam
 1139 of work in systems with more expressive types or more expressive unification algorithms than
 1140 first-order unification. The gold standard is MLF [2], but there are several subsequent variants,
 1141 including HML [12], and FPH [27].

1142 MLF extends type schemes with instantiation constraints, and makes the unification algorithm
 1143 aware of them. As a result it achieves the remarkable combination of: (i) typeability of the whole of
 1144 System F by only annotating function arguments that must be used polymorphically, (ii) principal
 1145 types and a sound and complete type inference algorithm, (iii) the “defining” ML property that any
 1146 sub-term can be lifted and let-bound without any further annotations without affecting typeability;
 1147 a corollary of principal types.

1148 However, MLF and variants do require intrusive modifications to a constraint solver (and in the
 1149 case of GHC, a complex constraint solver with type class constraints, implication constraints, type
 1150 families, and more) and to the type structure. Though some attempts have been made to integrate
 1151 MLF with qualified types [13], a full integration is uncharted territory. Hence, we present here a
 1152 pragmatic compromise in expressiveness for simplicity and ease of integration in the existing GHC
 1153 type inference engine.

1154 11 CONCLUSION AND FUTURE WORK

1155 In this paper we have presented Quick Look, a new approach to impredicative polymorphism which
 1156 is both modular and highly localised. We have implemented the system within GHC, and evaluated
 1157 the (scarce) changes required to a wide set of packages.

1158 Quick Look focuses on inferring impredicative *instantiations*. We plan to investigate whether a
 1159 similar approach could be used to infer polymorphic types in *argument* position. For example, being
 1160 able to accept the term $(\lambda x. x : ids)$. As in the case of visible type application with impredicative
 1161 instantiation, we expect *type variables in patterns* [6] to help with argument types.

1162 1163 ACKNOWLEDGMENTS

1164 We thank Richard Eisenberg, Stephanie Weirich, and Edward Yang for their feedback. We are deeply
 1165 grateful to Didier Rémy, for his particularly detailed review and subsequent email dialogue, going
 1166 far beyond the call of duty.

1167 1168 REFERENCES

- 1169 [1] Lennart Augustsson. 2011. Impredicative polymorphism: a use case. <http://augustss.blogspot.com/2011/07/impredicative-polymorphism-use-case-in.html>.
- 1170 [2] Didier Le Botlan and Didier Rémy. 2003. ML^F: raising ML to the power of system F. In *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming, ICFP 2003, Uppsala, Sweden, August 25-29, 2003*, Colin Runciman and Olin Shivers (Eds.). ACM, 27–38. <https://doi.org/10.1145/944705.944709>
- 1171 [3] Didier Le Botlan and Didier Rémy. 2009. Recasting MLF. *Inf. Comput.* 207, 6 (2009), 726–785.

- 1177 [4] Joshua Dunfield and Neelakantan R. Krishnaswami. 2013. Complete and easy bidirectional typechecking for higher-
1178 rank polymorphism. In *ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA -*
1179 *September 25 - 27, 2013*, Greg Morrisett and Tarmo Uustalu (Eds.). ACM, 429–442. [https://doi.org/10.1145/2500365.](https://doi.org/10.1145/2500365.2500582)
1180 2500582
- 1181 [5] Joshua Dunfield and Neelakantan R. Krishnaswami. 2019. Sound and Complete Bidirectional Typechecking for
1182 Higher-rank Polymorphism with Existentials and Indexed Types. *Proc. ACM Program. Lang.* 3, POPL, Article 9 (Jan.
1183 2019), 28 pages. <https://doi.org/10.1145/3290322>
- 1184 [6] Richard A. Eisenberg, Joachim Breitner, and Simon Peyton Jones. 2018. Type variables in patterns. In *Proceedings of*
1185 *the 11th ACM SIGPLAN International Symposium on Haskell, Haskell@ICFP 2018, St. Louis, MO, USA, September 27-17,*
1186 *2018*. 94–105. <https://doi.org/10.1145/3242744.3242753>
- 1187 [7] Richard A. Eisenberg, Stephanie Weirich, and Hamidhasan G. Ahmed. 2016. Visible Type Application. In *Proceedings*
1188 *of the 25th European Symposium on Programming Languages and Systems - Volume 9632*. Springer-Verlag New York,
1189 Inc., New York, NY, USA, 229–254. https://doi.org/10.1007/978-3-662-49498-1_10
- 1190 [8] Frank Emrich, Sam Lindley, Jan Stolarek, and James Cheney. 2019. FreezeML: Complete and Easy Type Inference for
1191 First-Class Polymorphism. Presented at TyDe 2019.
- 1192 [9] Jacques Garrigue and Didier Rémy. 1999. Semi-Explicit Higher-Order Polymorphism for ML. *Information and*
1193 *Computation* 155, 1/2 (1999), 134–169. <http://www.springerlink.com/content/m303472288241339/> A preliminary
1194 version appeared in TACS'97.
- 1195 [10] Jurriaan Hage and Bastiaan Heeren. 2009. Strategies for Solving Constraints in Type and Effect Systems. *Electronic*
1196 *Notes in Theoretical Computer Science* 236 (2009), 163 – 183. Proceedings of the 3rd International Workshop on Views
1197 On Designing Complex Architectures (VODCA 2008).
- 1198 [11] Daan Leijen. 2008. HMF: simple type inference for first-class polymorphism. In *Proceeding of the 13th ACM SIGPLAN*
1199 *international conference on Functional programming, ICFP 2008, Victoria, BC, Canada, September 20-28, 2008*. 283–294.
1200 <https://doi.org/10.1145/1411204.1411245>
- 1201 [12] Daan Leijen. 2009. Flexible types: robust type inference for first-class polymorphism. In *Proceedings of the 36th ACM*
1202 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23,*
1203 *2009*. 66–77. <https://doi.org/10.1145/1480881.1480891>
- 1204 [13] Daan Leijen and Andres Löh. 2005. Qualified types for MLF. In *Proceedings of the 10th ACM SIGPLAN International*
1205 *Conference on Functional Programming, ICFP 2005, Tallinn, Estonia, September 26-28, 2005*, Olivier Danvy and Benjamin C.
1206 Pierce (Eds.). ACM, 144–155. <https://doi.org/10.1145/1086365.1086385>
- 1207 [14] Robin Milner. 1978. A Theory of Type Polymorphism in Programming. *J. Comput. Syst. Sci.* 17, 3 (1978), 348–375.
1208 [https://doi.org/10.1016/0022-0000\(78\)90014-4](https://doi.org/10.1016/0022-0000(78)90014-4)
- 1209 [15] John C. Mitchell. 1988. Polymorphic Type Inference and Containment. *Inf. Comput.* 76, 2-3 (Feb. 1988), 211–249.
1210 [https://doi.org/10.1016/0890-5401\(88\)90009-0](https://doi.org/10.1016/0890-5401(88)90009-0)
- 1211 [16] Martin Odersky and Konstantin Läuffer. 1996. Putting type annotations to work. In *Principles of Programming Languages,*
1212 *POPL*. 54–67.
- 1213 [17] Simon Peyton Jones. 2019. GHC Proposal: “Simplify subsumption”. [https://github.com/ghc-proposals/ghc-proposals/](https://github.com/ghc-proposals/ghc-proposals/pull/287)
1214 [pull/287](https://github.com/ghc-proposals/ghc-proposals/pull/287)
- 1215 [18] Simon L. Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. 2007. Practical type inference for
1216 arbitrary-rank types. *Journal of Functional Programming* 17, 1 (2007), 1–82.
- 1217 [19] Benjamin C. Pierce and David N. Turner. 2000. Local type inference. *ACM Trans. Program. Lang. Syst.* 22, 1 (2000),
1218 1–44. <https://doi.org/10.1145/345099.345100>
- 1219 [20] François Pottier and Yann Régis-Gianas. 2006. Stratified Type Inference for Generalized Algebraic Data Types. In
1220 *Conference Record of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Charleston,
1221 South Carolina, USA) (*POPL '06*). ACM, New York, NY, USA, 232–244. <https://doi.org/10.1145/1111037.1111058>
- 1222 [21] François Pottier and Didier Rémy. 2005. The Essence of ML Type Inference. In *Advanced Topics in Types and*
1223 *Programming Languages*, Benjamin C. Pierce (Ed.). MIT Press, Chapter 10, 389–489. <http://crystal.inria.fr/attapl/>
- 1224 [22] Didier Rémy. 2005. Simple, Partial Type-inference for System F Based on Type-containment. In *Proceedings of the Tenth*
1225 *ACM SIGPLAN International Conference on Functional Programming* (Tallinn, Estonia) (*ICFP '05*). ACM, New York, NY,
1226 USA, 130–143. <https://doi.org/10.1145/1086365.1086383>
- 1227 [23] Claudio V. Russo and Dimitrios Vytiniotis. 2009. QML: Explicit First-class Polymorphism for ML. In *Proceedings of*
1228 *the 2009 ACM SIGPLAN Workshop on ML* (Edinburgh, Scotland) (*ML '09*). ACM, New York, NY, USA, 3–14. <https://doi.org/10.1145/1596627.1596630>
- 1229 [24] Alejandro Serrano, Jurriaan Hage, Dimitrios Vytiniotis, and Simon Peyton Jones. 2018. Guarded impredicative
1230 polymorphism. In *Proc ACM SIGPLAN Conference on Programming Languages Design and Implementation*. ACM.
1231 <https://www.microsoft.com/en-us/research/publication/guarded-impredicative-polymorphism/>

- 1226 [25] Dimitrios Vytiniotis, Simon L. Peyton Jones, Tom Schrijvers, and Martin Sulzmann. 2011. OUTSIDEIN(X): Modular type
1227 inference with local assumptions. *J. Funct. Program.* 21, 4-5 (2011), 333–412. <https://doi.org/10.1017/S0956796811000098>
- 1228 [26] Dimitrios Vytiniotis, Stephanie Weirich, and Simon L. Peyton Jones. 2006. Boxy types: inference for higher-rank types
1229 and impredicativity. In *Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming, ICFP 2006, Portland, Oregon, USA, September 16-21, 2006*, John H. Reppy and Julia L. Lawall (Eds.). ACM, 251–262.
1230 <https://doi.org/10.1145/1159803.1159838>
- 1231 [27] Dimitrios Vytiniotis, Stephanie Weirich, and Simon L. Peyton Jones. 2008. FPH: first-class polymorphism for Haskell.
1232 In *Proceeding of the 13th ACM SIGPLAN international conference on Functional programming, ICFP 2008, Victoria, BC, Canada, September 20-28, 2008*. 295–306. <https://doi.org/10.1145/1411204.1411246>
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274

1275 A PROOFS

1276 Section 6 already presents several properties of our constraint generation judgements. One inter-
 1277 esting lemma is that if we can infer a type, we can then check it (Theorem 6.3.) We proceed to give
 1278 the crux of the argument here.

1279 LEMMA A.1. *If $\Gamma \vdash_{\uparrow} e : \rho$ then $\Gamma \vdash_{\downarrow} e : \rho$.*

1281 PROOF. The proof is straightforward induction, with difficulties coming from rule APP. Let's
 1282 assume that $e = h \bar{\pi}$. From rule APP we'd get:

$$1283 \Gamma \vdash_{\uparrow}^h h : \sigma \quad (1)$$

$$1284 \bar{e} = \text{valargs}(\bar{\pi}) \quad (2)$$

$$1285 \text{inst} \sigma [\bar{\pi}] \rightsquigarrow \Delta ; \bar{\phi}, \rho_r \quad (3)$$

$$1286 \Gamma \vdash_{\downarrow} e_i : (\Delta, \phi_i) \rightsquigarrow \Theta_i \quad (4)$$

$$1287 \Theta = (\oplus \Theta_i) \quad \bar{\phi}'_i = \Theta \phi_i \quad \forall \bar{a}. \rho'_r = \Theta \rho_r \quad (5)$$

$$1288 \text{dom}(\theta) = \bar{a} \cup (\Delta \setminus \text{dom}(\Theta)) \quad (6)$$

$$1289 \Gamma \vdash_{\downarrow}^{\vee} e_i : \theta \phi'_i \quad (7)$$

$$1290 \rho = \theta \rho'_r \quad (8)$$

1291 Assume now that we'd like to check $\Gamma \vdash_{\downarrow} e : \theta \rho'_r$. Conditions (1), (2), and (3) directly carry over.
 1292 However now we need to perform $\text{match}(\Delta, \rho_r, \theta \rho'_r, \text{mnv})$. Let's consider the case first where $\bar{a} = \emptyset$,
 1293 for simplicity.

1294 In that case: $\theta \rho'_r = \theta \Theta \rho_r$, and we will be invoking the judgement with: $\text{match}(\Delta, \rho_r, \theta \Theta \rho_r, \text{mnv})$.
 1295 Now, the $\text{dom}(\theta)$ contains only variables in the range of Θ (it just monomorphises whatever variables
 1296 are left over uninstantiated from Θ). Hence it must be the case that matching give a substitution
 1297 which is equal to $\theta \Theta$ on the free instantiation variables of ρ that are matched to other types.
 1298 Subsequently, we'd get a $\Theta_{\downarrow} = (\oplus \Theta_i) \oplus (\theta \cdot (\oplus \Theta_i))$. Joining these two substitutions would create a
 1299 new substitution that could possibly differ by the original $\theta \Theta$ in that some bindings might not have
 1300 θ applied to them. Therefore we could always apply $\theta \Theta_{\downarrow}$ to get a substitution equivalent to $\theta \Theta$. We
 1301 can The rest follows by (6), (7), (8), above.

1302 The case when $\bar{a} \neq \emptyset$ involves a case analysis. If the original ρ_r was a naked variable mapped by
 1303 Θ to a forall type or not. If that was *not* the case then a similar analysis as before is applicable.

1304 If however, $\rho_r = \kappa$ then the matching in Figure 6 must *also* return an empty substitution, therefore
 1305 the inference and checking cases are now indistinguishable. This argument relies crucially on the
 1306 delicate design of the matching function. \square

1307 Next we give a set of auxiliary lemmas necessary to prove the main soundness results. We will
 1308 be assuming that $\text{fiv}(\Gamma) = \emptyset$, that is, the environment can only contain unification variables but not
 1309 instantiation variables. That is certainly true and preserved during constraint generation.

1310 LEMMA A.2. *Let θ be a monomorphic substitution. If $\text{ljoin}(\sigma_1, \sigma_2) = \sigma_3$ and $\text{fiv}(\theta) = \emptyset$ then*
 1311 $\text{ljoin}(\theta \sigma_1, \theta \sigma_2) = \theta \sigma_3$.

1312 PROOF. By induction on the structure of types σ_1 and σ_2 , and case analysis on the ljoin function.
 1313 \square

1314 LEMMA A.3. *If $\text{match}(\Delta, \rho_1, \rho_2, \omega) = \Theta$ and $\text{fiv}(\theta) = \emptyset$ then $\text{match}(\Delta, \theta \rho_1, \theta \rho_2, \omega) = (\theta \cdot \Theta)|_{\text{dom}(\Theta)}$.*

1315 PROOF. By induction on the structure of types, appealing to Lemma A.2. \square

1316 LEMMA A.4. *If $\Gamma \vdash_{\downarrow}^h h : \sigma$ and $\text{fiv}(\theta) = \emptyset$ then $\theta \Gamma \vdash_{\downarrow}^h h : \theta \sigma$.*

1317 PROOF. Easy case analysis. \square

Expressions	$e ::= \dots$		$\text{let } x = e \text{ in } e$
			$\text{let } x :: \sigma = e \text{ in } e$
			$\text{case } e_0 \text{ of } \{\overline{K_i x_i} \rightarrow e_i\}$
Constraints	$Q ::= \epsilon \mid Q \wedge Q$		
			$\sigma \sim \phi$
			\dots extensible
Poly. types	$\sigma, \phi ::= \dots \mid Q \Rightarrow \sigma$		
Constr. sigs	$ksig ::= K : \forall \bar{a} \bar{b}. Q \Rightarrow \bar{\sigma} \rightarrow \top \bar{a}$		
Environments	$\Gamma ::= \dots \mid \Gamma, ksig \mid \Gamma, Q$		

Fig. 12. Syntax for extended language

LEMMA A.5. *If $\vdash_{\delta} \rho_1 \equiv \rho_2 \rightsquigarrow C$ and $\theta \models C$ then $\theta\rho_1 = \theta\rho_2$.*

PROOF. Easy case analysis. □

LEMMA A.6. *If $\vdash^{\text{inst}} \sigma[\bar{e}] \rightsquigarrow \Delta$; $\bar{\phi}, \rho_r$; C and $\theta \models C$ and $\text{fiv}(\theta) = \emptyset$ then: $\vdash^{\text{inst}} \theta\sigma[\bar{e}] \rightsquigarrow \Delta$; $\theta\bar{\phi}, \theta\rho_r$.*

PROOF. Straightforward induction. The only interesting case is for rule `IVARM`, where we have to apply rule `IARG` in the declarative specification. The rest of the cases go through by directly invoking the induction hypothesis or are trivial. □

LEMMA A.7. *If $\Gamma \vdash_{\downarrow} e : (\Delta, \phi) \rightsquigarrow \Theta$ and $\text{fiv}(\theta) = \emptyset$ then $\theta\Gamma \vdash_{\downarrow} e : (\Delta, \theta\phi) \rightsquigarrow \theta \cdot \Theta|_{\Delta}$.*

As notational convenience, we write below $\theta \models C$ to mean that $\text{fiv}(\theta) = \emptyset$ and $\theta \models C$.

LEMMA A.8 (SOUNDNESS). *The following are true (proof by mutual induction on the size of the term):*

- If $\Gamma \vdash_{\uparrow}^h h : \sigma \rightsquigarrow C$ and $\theta \models C$ then $\theta\Gamma \vdash_{\uparrow}^h h : \theta\sigma$.
- If $\Gamma \vdash_{\downarrow}^v e : \sigma \rightsquigarrow C$ and $\theta \models C$ then $\theta\Gamma \vdash_{\downarrow}^v e : \theta\sigma$.
- If $\Gamma \vdash_{\uparrow} e : \rho \rightsquigarrow C$ and $\theta \models C$ then $\theta\Gamma \vdash_{\uparrow} e : \theta\rho$.
- If $\Gamma \vdash_{\downarrow} e : \rho \rightsquigarrow C$ and $\theta \models C$ then $\theta\Gamma \vdash_{\downarrow} e : \theta\rho$.

Theorem 7.2 follows directly from Lemma A.8.

B EXTENSIONS TO THE LANGUAGE

As discussed in Section 5.5, one of the salient features of Quick Look is its modularity with respect to other type system features. In this section we describe how some language extensions supported by GHC can be readily integrated in our formalization: `let` bindings, pattern matching, quantified constraints, and GADTs. For the sake of readability, Figure 12 includes all the required syntactic extensions for the foregoing developments.

B.1 Local Bindings

Our description of local bindings, given in Figure 13a, follows Vytiniotis et al. [25] closely. In particular, the type of a `let` binding is not generalized unless an explicit annotation is given. This

$$\begin{array}{c}
\boxed{\Gamma \vdash_{\uparrow} e : \rho} \quad \boxed{\Gamma \vdash_{\downarrow} e : \rho} \\
\frac{\Gamma \vdash_{\uparrow} e_1 : \rho_1 \quad \Gamma, x : \rho_1 \vdash_{\delta} e_2 : \rho_2}{\Gamma \vdash_{\delta} \text{let } x = e_1 \text{ in } e_2 : \rho_2} \text{LET} \\
\frac{\Gamma \vdash_{\downarrow}^{\forall} e_1 : \sigma_1 \quad \Gamma, x : \sigma_1 \vdash_{\delta} e_2 : \rho_2}{\Gamma \vdash_{\delta} \text{let } x :: \sigma_1 = e_1 \text{ in } e_2 : \rho_2} \text{ANNLET}
\end{array}$$

(a) No generalization

$$\begin{array}{c}
\Gamma \vdash_{\uparrow} e_1 : \rho_1 \quad \bar{a} = \text{fv}(\rho_1) - \text{fv}(\Gamma) \\
\frac{\Gamma, x : \forall \bar{a}. \rho_1 \vdash_{\delta} e_2 : \rho_2}{\Gamma \vdash_{\delta} \text{let } x = e_1 \text{ in } e_2 : \rho_2} \text{LETGEN}
\end{array}$$

(b) With generalization

Fig. 13. Local bindings

$$\begin{array}{c}
\boxed{\Gamma \vdash_{\uparrow} e : \rho} \quad \boxed{\Gamma \vdash_{\downarrow} e : \rho} \\
\Gamma \vdash_{\uparrow} e_0 : \Gamma \bar{\sigma} \\
\text{for each branch } K_i \bar{x}_i \rightarrow e_i \text{ do:} \\
K_i : \forall \bar{a}. \bar{v}_i \rightarrow \Gamma \bar{a} \in \Gamma \\
\frac{\Gamma, x_i : [\bar{a} := \bar{\sigma}] v_i \vdash_{\delta} e_i : \rho}{\Gamma \vdash_{\delta} \text{case } e_0 \text{ of } \{K_i \bar{x}_i \rightarrow e_i\} : \rho} \text{CASE}
\end{array}$$

Fig. 14. Pattern matching



design enables the most information to propagate between the definition of a local binding and its use sites when using a constraint-based formulation.

Figure 13b shows another possible design, in which generalization is performed on non-annotated lets. The main disadvantage is that when implemented in a constraint-based system, this forces us to solve the constraints obtained from e_1 before looking at e_2 . Otherwise, we cannot be sure about which type variables to generalize. Another possibility, taken by Hage and Heeren [10], Pottier and Rémy [21], is to extend the language of constraints with generalization and instantiation, making the solver aware of the order in which these constraints ought to be solved.

B.2 Pattern Matching

The integration of pattern matching in a bidirectional type system can be done in several ways, depending on the direction in which the expression being matched is type checked. In the rules given in Figure 14 we look at that expression e_0 in inference mode; the converse choice is to look at how branches are using the value to infer the instantiation.

1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470



1471 B.3 Quantified Constraints

1472 Haskell has a much richer vocabulary of polymorphic types that simply $\forall \bar{a}. \rho$. In general, in addition
 1473 to quantified type variables, a set of *constraints* may appear. Those constraints must be satisfied by
 1474 the chosen instantiation in order for the program to be accepted.

1475 Following Vytiniotis et al. [25] we leave the language of constraints open; in the case of GHC
 1476 this language includes type class constraints and equalities with type families. Figure 15 shows
 1477 that the changes required to support quantified constraints are fairly minimal:

- 1478 • Environments Γ may now also mention *local constraints*. This is a slight departure from
 1479 Vytiniotis et al. [25], in which variable environments and local constraints were kept in
 1480 separate sets; this change allows us to control better the scope of each type variable.
- 1481 • Rules GEN and IAPP now have to deal with constraints. In the former case, Q is added to
 1482 the set of local constraints. In the case of IAPP, the constraints are returned as part of the
 1483 instantiation judgment.
- 1484 • Rule APP needs to check that the constraints obtained from instantiation hold for the chosen
 1485 set of types. We use an auxiliary *constraint entailment* judgment $\Gamma \Vdash Q$ which states that
 1486 constraints Q hold in the given environment (which may contain local assumptions). This
 1487 judgment can be freely instantiated, as explained in Vytiniotis et al. [25].

1489 B.4 Generalized Algebraic Data Types

1490 GADTs extend the language by allowing local constraints and quantification also in data type
 1491 constructors. These constraints are in scope whenever pattern matching consider that case; the
 1492 updated CASE rules in Figure 16 adds those constraints to the environment.

1493 In principle, Quick Look is not affected by these changes. But we could also use some information
 1494 about the usage of types in the rest of constraints to guide the choice of impredicativity. For example,
 1495 Haskell does not allow type class instances over polymorphic types; so if we find a constraint $Eq\ a$,
 1496 we know that a should not be impredicatively instantiated.
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511
 1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519