

# Implicit Polarized F: local type inference for impredicativity

ANONYMOUS AUTHOR(S)

System F, the polymorphic lambda calculus, features the principle of *impredicativity*: polymorphic types may be (explicitly) instantiated at other types, enabling many powerful idioms such as Church encoding and data abstraction. Unfortunately, type applications need to be implicit for a language to be human-usable, and the problem of inferring all type applications in System F is undecidable. As a result, language designers have historically avoided impredicative type inference.

We reformulate System F in terms of call-by-push-value, and study type inference for it. Surprisingly, this new perspective yields a novel type inference algorithm which is extremely simple to implement (not even requiring unification), infers many types, and has a simple declarative specification. Furthermore, our approach offers type theoretic explanations of how many of the heuristics used in existing algorithms for impredicative polymorphism arise.

## 1 INTRODUCTION

System F, the polymorphic lambda calculus, was invented in the early 1970s by John Reynolds [16] and Jean-Yves Girard [6], and has remained one of the fundamental objects of study in the theory of the lambda calculus ever since.

Syntactically, it is a tiny extension of the simply-typed lambda calculus with type variables and quantification over them, with only five typing rules for the whole calculus. Despite its small size, it is capable of modeling inductive data types via Church encodings<sup>1</sup>, and supports reasoning about data abstraction and modularity via the theory of parametricity [17, 25].

Offering a combination of parsimony and expressive power, System F has been important in the study of semantics, and has also served as inspiration for much work on language design [21]. However, practical languages have historically shied away from adopting the distinctive feature of System F — full impredicative polymorphism. This is because the natural specification for type inference for full System F is undecidable.

To understand this specification, consider the following two variables:

$$\begin{aligned} \text{fst} &: \forall \alpha, \beta. (\alpha \times \beta) \rightarrow \alpha \\ \text{pair} &: \text{Int} \times \text{String} \end{aligned}$$

Here, `fst` has the polymorphic type of the first projection for pairs, and `pair` is a pair with a concrete type `Int × String`. To project out the first component in System F, we would write:

$$\text{fst} [\text{Int}] [\text{String}] \text{pair}$$

Note that each type parameter must be given explicitly. This is a syntactically heavy discipline: even in this tiny example, the number of tokens needed for type arguments equals the number of tokens which compute anything! Instead, we would prefer to write:

$$\text{fst pair}$$

---

<sup>1</sup>Indeed, any first-order function provably total in second-order arithmetic can be expressed in System F.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2022 Association for Computing Machinery.

XXXX-XXXX/2022/3-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Leaving the redundant type arguments *implicit* makes the program more readable. But how can we specify this? In the application `fst` pair, the function `fst` has the type  $\forall \alpha, \beta. (\alpha \times \beta) \rightarrow \alpha$ , but we wish to use it at a function type  $(\text{Int} \times \text{String}) \rightarrow \text{Int}$ .

The standard technical device for using one type in the place of another is *subtyping*. With subtyping, the application rule is:

$$\frac{\Theta, \Gamma \vdash f : A \quad \Theta \vdash A \leq B \rightarrow C \quad \Theta, \Gamma \vdash v : B}{\Theta, \Gamma \vdash f v : C}$$

As long our subtype relation shows that  $\forall \alpha, \beta. (\alpha \times \beta) \rightarrow \alpha \leq (\text{Int} \times \text{String}) \rightarrow \text{Int}$ , we are in the clear! Since we want subtyping to instantiate quantifiers, we can introduce subtyping rules for the universal type which reflect the *specialization order*:

$$\frac{\Theta \vdash A \text{ type} \quad \Theta \vdash [A/\alpha] B \leq C}{\Theta \vdash \forall \alpha. B \leq C} \forall L \qquad \frac{\Theta, \alpha \vdash B \leq C}{\Theta \vdash B \leq \forall \alpha. C} \forall R$$

$$\frac{\Theta \vdash X \leq A \quad \Theta \vdash B \leq Y}{\Theta \vdash A \rightarrow B \leq X \rightarrow Y} \qquad \frac{\alpha \in \Theta}{\Theta \vdash \alpha \leq \alpha}$$

Since a more general type like  $\forall \alpha, \beta. (\alpha \times \beta) \rightarrow \alpha$  can be used in the place of a more specific type like  $(\text{Int} \times \text{String}) \rightarrow \text{Int}$ , we take the standard subtyping rules for functions and base types, and add rules reflecting the principle that a universal type is a subtype of all of its instantiations.

This subtype relation is small, natural, and expressive. Thus, it has been the focus of a great deal of research in subtyping. Unfortunately, that research has revealed that this relation is *undecidable* [4, 22]. This has prompted many efforts to understand the causes of undecidability and to develop strategies which work around it. To contextualize our contributions, we briefly describe some of these efforts, but delay extended discussion of the related work until the end of the paper.

*Impredicativity.* One line of research locates the source of the undecidability in the fact that the  $\forall L$  rule is *impredicative* – it permits a quantifier to be instantiated at any type, including types with quantifiers in them.

For example, the identity function `id` can be typed in System F at  $\forall \alpha. \alpha \rightarrow \alpha$ , and impredicativity means that self-application is typeable with a polymorphic type:

$$\text{id } [\forall \alpha. \alpha \rightarrow \alpha] \text{ id } : \forall \alpha. \alpha \rightarrow \alpha$$

A perhaps more useful example is the low-precedence application function in Haskell:

$$\begin{aligned} (\$) &: \forall \alpha. \forall \beta. (\alpha \rightarrow \beta) \rightarrow \alpha \rightarrow \beta \\ (\$) \text{ f x} &= \text{f x} \end{aligned}$$

This function simply applies a function to an argument, and is used in languages like Haskell and OCaml to avoid parentheses.

In Haskell, impredicative instantiation in the apply operator, for example in expressions like `runST $ (return True)`<sup>2</sup>, is important enough that GHC, the primary Haskell compiler, has a special case typing rule for it [19]. However, this special case typing rule is non-modular and can lead to surprising results. For example, if we define an alias `app` for `$`, then `runST $ (return True)` typechecks, but `runST `app` (return True)` does not (unless the user is using the *ImpredicativeTypes* extension).

<sup>2</sup>Here, `runST` has type  $\forall \alpha. (\forall \beta. \text{ST } \beta \alpha) \rightarrow \alpha$ , which ensures that the internal state  $\beta$  is confined within the ST monad and remains inaccessible to the rest of the program. This polymorphic parameter of `runST` necessitates impredicative instantiation of  $\$$ .

99 The reason for this special case rule is that GHC’s type inference has built on a line of work [5, 11,  
100 13] which restricts the  $\forall$ L rule to be *predicative*: universally quantified types can only be instantiated  
101 at monotypes (i.e. types which do not contain any quantifiers). This recovers decidability of  
102 subsumption, but at the price of giving up any inference for uses of impredicative types.

103  
104 *Type Syntax.* Another line of work sites the difficulty of type inference in the inexpressivity of  
105 the type language of System F. The work on  $ML^F$  [1, 15] extends System F with a form of bounded  
106 quantification. These bounds are rich enough that the grammar of types can precisely record exactly  
107 which types a polymorphic quantifier might be instantiated at, which is enough to make the type  
108 inference problem tractable once more.

109 Of course, not just any language of bounds will work — the  $ML^F$  system was very carefully  
110 crafted to make inference both decidable and expressive, and is startlingly successful at this  
111 goal, only requiring type annotations on variables which are used at two different polymorphic  
112 types. Unfortunately, the  $ML^F$  algorithm is somewhat notorious for its implementation complexity,  
113 with multiple attempts to integrate it into GHC [7, 24] failing due to the difficulty of software  
114 maintenance.

115 *Heuristic Approaches.* A third line of work observes that type inference is actually easy in most  
116 cases, and that an algorithm which only does the easy things and fails otherwise may well suffice  
117 to be an algorithm which works in practice.

118 One of the oldest such approaches is that of Cardelli [2], which was invented as part of the  
119 implementation of  $F<$ , a language with impredicative polymorphism and recursive types. The  
120 heuristic algorithm Cardelli proposed was to permit impredicative instantiation without tracking  
121 bounds — his algorithm simply took any bound it saw as an equation telling it what to instantiate  
122 the quantifier to. Pierce and Turner [14], seeing how well this algorithm worked in practice,  
123 formalized it and proved it sound (naming it *local type inference*), but were not able to give a clear  
124 non-algorithmic specification.  
125

## 126 1.1 Contributions

127 Historically, approaches to impredicative type inference have been willing to pay a high cost in  
128 complexity in order to achieve high levels of expressiveness and power of type inference. In this  
129 paper, we rethink that historical approach by focusing upon the opposite tradeoff. We prioritize  
130 simplicity — of implementation, specification, and connections to existing type theories — over the  
131 expressiveness of type inference.

132 By being willing to make this sacrifice, we discover a easy-to-implement type inference algorithm  
133 for impredicative polymorphism, which has a clear declarative specification and is even easier to  
134 implement than traditional Damas-Milner type inference. Our algorithm sheds light on a classical  
135 implementation technique of Cardelli, and works over a foundational calculus, call-by-push-value.  
136

137 Call-by-push-value was invented to study the interaction of evaluation order and side-effects,  
138 and has been used extensively in semantics. In this paper, we show that it is also well-adapted to  
139 being a kernel calculus for type inference by introducing Implicit Polarized F, a polarized variant  
140 of System F.

141 Specifically, our contributions are:

- 142 • To model type inference declaratively, we first introduce a subtyping relation for Implicit  
143 Polarized F that models the specialization order. Our subtyping relation is extremely simple,  
144 and is *almost* the “off-the-shelf” relation that one would first write down. The only unusual  
145 restriction is that subtyping for shifts (the modalities connecting value and computation  
146 types) is invariant rather than covariant.  
147

We then give a type system for our variant of call-by-push-value, using subtyping as a component. Surprisingly, features of call-by-push-value invented for semantic or type-theoretic reasons turn out to be perfectly suited for specifying type inference.

For example, argument lists (sometimes called spines in the literature [3, 20]) are commonly used in practical type inference systems [7, 18, 19]. They also arise in call-by-push-value due to the shift structure mediating between functions and values. This lets us naturally scope inference to individual function applications. So we infer types for fully unambiguous function calls, and otherwise require type annotations.

This illustrates our commitment to simplicity over powerful inference, and also results in very regular and predictable annotation behavior — our language looks like ordinary (polarized) System F, except that many of the “obvious” type instantiations are implicit.

- We give algorithmic subtyping rules corresponding to the specialization order, and algorithmic typing rules corresponding to declarative typing.

Both of these algorithms are startlingly simple, as well — neither algorithm even needs unification, making them *simpler* than traditional Hindley-Damas-Milner [10] type inference! We prove that our algorithm is decidable, and that it is sound and complete with respect to the declarative specification.

- In fact, the combination of spine inference plus specialization yields an algorithm that is very similar to Cardelli’s inference algorithm for bounded System F (as well as Pierce and Turner’s local type inference). It has been known for decades that Cardelli’s algorithm works extremely well in practice. However, its type-theoretic origins have not been fully understood, nor has it been understood how to characterize the kinds of problems that Cardelli’s algorithm succeeds on using a declarative specification.

The restrictions we impose to recover decidability of specialization turn out to be natural both from the perspective of type theory and implementation, and thereby offer a theoretically motivated reconstruction of Cardelli’s algorithm. Polarity also ends up shedding light on many of the design choices underpinning a number of other type inference algorithms for System F [7, 18, 19].

An *explicit non-contribution* of this work is to propose a practical type inference algorithm for any existing language. For example, we work with a call-by-push-value calculus, which is a type structure no practical language currently has. Furthermore, we make no attempt to infer the types of arguments to functions; instead we focus solely on inferring type applications in a (polarized) variant of System F. Restricting the scope so much lets us isolate the key design issues underpinning a foundational problem in type inference.

## 2 IMPLICIT POLARIZED F

We first present Implicit Polarized F, a language with type inference that combines call-by-push-value style polarization with System F style polymorphism<sup>3</sup>.

Like call-by-push-value, Implicit Polarized F partitions terms into *values* and *computations*, ascribing *positive types* to values and *negative types* to computations. Levy [9] described the difference between values and computations in terms of the operational semantics: “a value *is*, whereas a computation *does*”. Another rule of thumb for differentiating between values and computations is to look at how we eliminate them: we tend to eliminate values by using pattern matching and computations by supplying an argument. For example, datatypes are values because they are

<sup>3</sup>Adding polarization does not result in us losing any expressiveness compared to typed System F: we demonstrate in the appendix that typeable System F can be embedded within Implicit Polarized F.

197	Values	$v ::= x \mid \{t\}$
198	Computations	$t ::= \lambda x : P. t \mid \Lambda \alpha. t \mid \text{return } v \mid$
199		$\text{let } x = v(s); t \mid \text{let } x : P = v(s); t$
200		
201	Argument lists	$s ::= \epsilon \mid v, s$
202	Positive types	$P ::= \alpha \mid \downarrow N$
203	Negative types	$N ::= P \rightarrow N \mid \forall \alpha. N \mid \uparrow P$
204		
205	Typing contexts	$\Theta ::= \cdot \mid \Theta, \alpha$
206	Typing environments	$\Gamma ::= \cdot \mid \Gamma, x : P$
207		
208		
209		
210		
211		
212		
213		
214		
215		
216		
217		
218		
219		
220		
221		
222		
223		
224		
225		
226		
227		
228		
229		
230		
231		
232		
233		
234		
235		
236		
237		
238		
239		
240		
241		
242		
243		
244		
245		

Fig. 1. Implicit Polarized F

eliminated by pattern matching, whereas functions are computations because they are eliminated by supplying arguments.

We present the term language for Implicit Polarized F in Figure 1. Terms are split between values and computations as follows:

- Variables  $x$  are values, since this is one of the invariants of call-by-push-value.
- Lambda abstractions  $\lambda x : P. t$  are computations, not values, since they are eliminated by passing an argument. We can think of a lambda abstraction as a computation  $t$  of type  $N$  with holes in it representing the (positive) bound variable, so lambda abstractions have the negative type  $P \rightarrow N$ .
- Thunks  $\{t\}$  are values that suspend arbitrary computations. Surrounding a computation  $t$  of type  $N$  with braces suspends that computation, giving us a thunk  $\{t\}$  of type  $\downarrow N$ . To create a traditional function value, we make a lambda abstraction  $\lambda x : P. t$  into a thunk  $\{\lambda x : P. t\}$ . Traditional thunks therefore have the type  $\downarrow(P \rightarrow N)$ .
- System-F style type abstractions  $\Lambda \alpha. t$  are computations of type  $\forall \alpha. N$ . Our choice of polarities for the types of polymorphic terms is motivated by the correspondence between function abstraction  $\lambda x : P. t$  and type abstraction  $\Lambda \alpha. t$ .
- The let forms  $\text{let } x = v(s); t$  and  $\text{let } x : P = v(s); t$  are sequencing computations that operate much like the bind operation in a monad. Each form takes the thunked computation  $v$ , passes it the arguments  $s$  it needs, binds the result to the variable  $x$  and continues the computation  $t$ .
- $\text{return } v$  is a trivial computation that just returns a value, completing the sequencing monad. We can also use  $\text{return}$  to return values from a function – in fact since all the terminal symbols for terms live in the grammar of values, the syntax mandates that functions eventually return a value. The type structure of returned values is symmetric to that of thunks: returned values have type  $\uparrow P$ .

Besides the addition of type annotations to guide inference and some differences to the syntax, our language is nothing more than call-by-push-value plus polymorphism. Our use of argument lists serves only to reverse the order of argument lists from Levy’s original work so that our syntax matches traditional languages with n-ary functions like JavaScript and C. Just like those languages, all of a function’s arguments must be passed at once in Implicit Polarized F.

246  $\boxed{\Theta; \Gamma \vdash e : A}$  The term  $e$  synthesizes the type  $A$

247  $\boxed{\Theta; \Gamma \vdash s : N \gg M}$  When passed to a head of type  $N$ , the argument list  $s$  synthesizes

248 the type  $M$

249

250 
$$\frac{x : P \in \Gamma}{\Theta; \Gamma \vdash x : P} \text{Dvar} \qquad \frac{\Theta; \Gamma \vdash t : N}{\Theta; \Gamma \vdash \{t\} : \downarrow N} \text{Dthink}$$

251

252

253 
$$\frac{\Theta; \Gamma, x : P \vdash t : N}{\Theta; \Gamma \vdash \lambda x : P. t : P \rightarrow N} \text{Dlabs} \qquad \frac{\Theta, \alpha; \Gamma \vdash t : N}{\Theta; \Gamma \vdash \Lambda \alpha. t : \forall \alpha. N} \text{Dgen}$$

254

255

256

257 
$$\frac{\Theta; \Gamma \vdash v : P}{\Theta; \Gamma \vdash \text{return } v : \uparrow P} \text{Dreturn}$$

258

259

260 
$$\frac{\Theta; \Gamma \vdash v : \downarrow M \quad \Theta; \Gamma \vdash s : M \gg \uparrow Q \quad \Theta \vdash \uparrow Q \leq^- \uparrow P \quad \Theta; \Gamma, x : P \vdash t : N}{\Theta; \Gamma \vdash \text{let } x : P = v(s); t : N} \text{Dambiguouslet}$$

261

262

263 
$$\frac{\Theta; \Gamma \vdash v : \downarrow M \quad \Theta; \Gamma \vdash s : M \gg \uparrow Q \quad \Theta; \Gamma, x : Q \vdash t : N \quad \forall P. \text{if } \Theta; \Gamma \vdash s : M \gg \uparrow P \text{ then } \Theta \vdash Q \cong^+ P}{\Theta; \Gamma \vdash \text{let } x = v(s); t : N} \text{Dunambiguouslet}$$

264

265

266

267

268 
$$\frac{}{\Theta; \Gamma \vdash \epsilon : N \gg N} \text{Dspinenil} \qquad \frac{\Theta; \Gamma \vdash v : P \quad \Theta \vdash P \leq^+ Q \quad \Theta; \Gamma \vdash s : N \gg M}{\Theta; \Gamma \vdash v, s : (Q \rightarrow N) \gg M} \text{Dspinecons}$$

269

270

271 
$$\frac{\Theta \vdash P \text{type}^+ \quad \Theta; \Gamma \vdash s : [P/\alpha]N \gg M}{\Theta; \Gamma \vdash s : (\forall \alpha. N) \gg M} \text{Dspintypeabs}$$

272

273

274

Fig. 2. Declarative type system

## 2.1 Declarative typing

We present in Figure 2 a declarative type system for Implicit Polarized F. Our system has simple, mostly syntax directed rules, with its main complexity lying in the unusual premise to Dunambiguouslet and the rules for typing argument lists.

The system has five main judgments:

- $\Theta; \Gamma \vdash v : P$ : In the context  $\Theta$  and typing environment  $\Gamma$ , the value  $v$  synthesizes the positive type  $P$ .
- $\Theta; \Gamma \vdash t : N$ : In the context  $\Theta$  and typing environment  $\Gamma$ , the computation  $t$  synthesizes the negative type  $N$ .
- $\Theta; \Gamma \vdash s : N \gg M$ : In the context  $\Theta$  and typing environment  $\Gamma$ , and when passed to a head of type  $N$ , the argument list  $s$  synthesizes the type  $M$ .
- $\Theta \vdash P \leq^+ Q$ : In the context  $\Theta$ ,  $P$  is a positive subtype of  $Q$ .
- $\Theta \vdash N \leq^- M$ : In the context  $\Theta$ ,  $N$  is a negative subtype of  $M$ . We reverse the alphabetical order of  $M$  and  $N$  in the negative judgment to better indicate the symmetry between positive and negative subtyping — we will see what this symmetry is in Section 4.1.

When we want to be ambiguous between values and computations, we write  $\Theta; \Gamma \vdash e : A$  and  $\Theta \vdash A \leq^{\pm} B$ .

Our inference rules are the following:

- Dvar lets us reference variables within the type environment  $\Gamma$ .
- Dthunk states that if we know that a computation has type  $N$ , then thunking it produces a value of type  $\downarrow N$ .
- Dlabs is the standard typing rule for lambda abstraction. The value hypothesis  $x : P$  in this rule maintains the invariant that the type environment only contains bindings to values.
- Dgen is to Dlabs as type abstraction is to function abstraction. We add a type variable to the typing context, as opposed to adding a variable binding to the typing environment.
- Dreturn complements Dthunk: if a value has type  $P$ , then returning it produces a computation of type  $\uparrow P$ .
- Dambiguouslet allows us to let-bind the results of function applications.
  - In the first premise,  $\Theta; \Gamma \vdash v : \downarrow M$ , we require  $v$  to be a thunked computation.
  - We then take the type  $M$  of this computation and use the spine judgment  $\Theta; \Gamma \vdash s : M \gg \uparrow Q$  to type the function application  $v(s)$ . The requirement that the argument list produces an upshifted type  $\uparrow Q$  means that the type has to be maximally applied. This encodes the fact that partial application is forbidden. Typically  $v$  will be a downshifted function and  $s$  will be the argument list to pass to that function. For example,  $v$  could be a function stored within the type environment, e.g.  $f$  when  $f : \downarrow(P \rightarrow N) \in \Gamma$ , or one written inline  $\{\lambda x : P. t\}$ . However we can also use an embedded value  $\{\text{return } v\}$  as a head along with an empty argument list to let-bind values.
  - Now that we have the type  $\uparrow Q$  that the argument list takes the head to, we use the subtyping judgment  $\Theta \vdash \uparrow Q \leq^{\sim} \uparrow P$  to check that the type synthesized by the function application is compatible with the annotation  $P$ . For our algorithm to work, we need to check the stronger compatibility constraint that  $\uparrow Q$  is a subtype of  $\uparrow P$ , rather than just checking that  $Q$  is a subtype of  $P$ .
  - Finally,  $\Theta; \Gamma, x : P \vdash t : N$  tells us that the type of the let term is given by the type of  $t$  in the type environment  $\Gamma$  extended with  $x$  bound to  $P$ .
- Dunambiguouslet is a variant of Dambiguouslet that lets us infer the return types of function applications when they are unambiguous. The first and second premises are identical to those of Dambiguouslet, and the third premise is the same as the last premise of Dambiguouslet except that it binds  $x$  directly to the type synthesized by the function application. However the final premise is unusual, and one might wonder whether it is in fact well founded. This premise encodes the condition that in order to omit the annotation, the return type of the function application must be unambiguous. By quantifying over all possible inferred types, we check that there is only one type (modulo isomorphism) that can be inferred for the return type of the function application. If every inferred type is equivalent to  $Q$ , then we can arbitrarily choose  $Q$  to bind  $x$  to, since any other choice would synthesize the same type for the let-binding. The well-foundedness of our rule stems from the use of a syntactically smaller subterm in the premise of Dunambiguouslet compared to the conclusion. In  $\Theta; \Gamma \vdash s : M \gg \uparrow P$ ,  $s$  is a smaller subterm than  $\text{let } x = v(s); t$ , and every subderivation of that premise acts on a smaller subterm too.

A consequence of this unusual premise is that the soundness and completeness theorems end up being mutually recursive.

344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392

$$\boxed{\Theta \vdash A \text{ type}^\pm} \quad \text{In the context } \Theta, A \text{ is a well-formed positive/negative type}$$

$$\frac{\alpha \in \Theta}{\Theta \vdash \alpha \text{ type}^+} \text{ Twfivar} \quad \frac{\Theta \vdash N \text{ type}^-}{\Theta \vdash \downarrow N \text{ type}^+} \text{ Twfshift}\downarrow \quad \frac{\Theta, \alpha \vdash N \text{ type}^-}{\Theta \vdash \forall \alpha. N \text{ type}^-} \text{ Twffforall}$$

$$\frac{\Theta \vdash P \text{ type}^+ \quad \Theta \vdash N \text{ type}^-}{\Theta \vdash P \rightarrow N \text{ type}^-} \text{ Twfarrow} \quad \frac{\Theta \vdash P \text{ type}^+}{\Theta \vdash \uparrow P \text{ type}^-} \text{ Twfshift}\uparrow$$

Fig. 3. Well-formedness of declarative types

$$\boxed{\Theta \vdash A \leq^\pm B} \quad \text{In the context } \Theta, A \text{ is a positive/negative declarative subtype of } B$$

$$\frac{\Theta \vdash \alpha \text{ type}^+}{\Theta \vdash \alpha \leq^+ \alpha} \leq^\pm \text{Drefl} \quad \frac{\Theta \vdash M \leq^- N \quad \Theta \vdash N \leq^- M}{\Theta \vdash \downarrow N \leq^+ \downarrow M} \leq^\pm \text{Dshift}\downarrow$$

$$\frac{\Theta, \alpha \vdash N \leq^- M}{\Theta \vdash N \leq^- \forall \alpha. M} \leq^\pm \text{Dforallr} \quad \frac{\Theta \vdash P \text{ type}^+ \quad \Theta \vdash [P/\alpha]N \leq^- M}{\Theta \vdash \forall \alpha. N \leq^- M} \leq^\pm \text{Dforalll}$$

$$\frac{\Theta \vdash Q \leq^+ P \quad \Theta \vdash N \leq^- M}{\Theta \vdash P \rightarrow N \leq^- Q \rightarrow M} \leq^\pm \text{Darrow} \quad \frac{\Theta \vdash Q \leq^+ P \quad \Theta \vdash P \leq^+ Q}{\Theta \vdash \uparrow P \leq^- \uparrow Q} \leq^\pm \text{Dshift}\uparrow$$

Fig. 4. Declarative subtyping

The last three rules teach us how to type function applications:

- $\text{Dspineil}$  states that an empty argument list does not change the type of the head of a function application.
- $\text{Dspinecons}$  tells us how to type a non-empty argument list. The argument list  $v, s$  takes  $Q \rightarrow N$  to  $M$  if the value  $v$  synthesizes a type  $P$ , this type  $P$  is either  $Q$  or a subtype of it, and the remaining argument list  $s$  takes  $N$  to  $M$ .
- $\text{Dspinetypeabs}$  lets us instantiate the head that an argument list can be passed to by replacing a type variable  $\alpha$  with a well-formed (see Figure 3) type  $P$ . Note that the instantiation is implicit — there is no annotation for  $P$ , so we are doing inference here.

## 2.2 Subtyping

To enable implicit type application, our type system uses a subtyping relation  $\Theta \vdash A \leq^\pm B$  which reifies the specialization order, expressing when a term of type  $A$  can be safely used where a term of type  $B$  is expected. This relation is defined in Figure 4 and consists of the following rules:

- $\leq^\pm \text{Drefl}$  is the standard rule for type variables: if we are expecting a term of type  $\alpha$ , then the only terms we can safely use in its place are those that also have type  $\alpha$ .
- $\leq^\pm \text{Darrow}$  is the classic rule for function subtyping: we can replace a term of type  $Q \rightarrow M$  with a function that takes inputs that are at least as general as  $Q$  and produces outputs that are at least as specific as  $M$ .



393  $\Theta \vdash A \cong^\pm B$  In the context  $\Theta$ , the types  $A$  and  $B$  are isomorphic  
 394

$$395 \quad \Theta \vdash A \cong^\pm B \text{ iff } \Theta \vdash A \leq^\pm B \text{ and } \Theta \vdash B \leq^\pm A.$$

396  
397  
398 Fig. 5. Isomorphic types  
399

- 400
- 401 •  $\leq^\pm D\text{forallll}$  and  $\leq^\pm D\text{forallr}$  correspond to the left and right rules governing  $\forall$  quantifiers
  - 402 in the LK sequent calculus system. Surprisingly, the  $\leq^\pm D\text{forallll}$  rule does not have any
  - 403 restrictions beyond well-formedness:  $P$  is an arbitrary well-formed type, so full impredicative
  - 404 instantiation is possible.
  - 405 •  $\leq^\pm D\text{shift}\downarrow$  and  $\leq^\pm D\text{shift}\uparrow$  are unusual since the standard denotational and operational
  - 406 semantics of shifts both give rise to covariant rules. However using these covariant rules
  - 407 would lead to an undecidable system [4, 22]. To restore decidability, we use more restrictive
  - 408 invariant shift rules.

409 This idea is motivated by a similar restriction that Cardelli [2] introduced as a heuristic  
 410 for type inference in an implementation of System  $F_{<}$ . Instead of using subtyping to infer  
 411 the types assigned to type parameters, Cardelli used first-order unification to “synthesize”  
 412 these types.

413 Since the goal of unification is to make expressions equal, an equivalent notion of Cardelli’s  
 414 restriction in our system is requiring the types to be *isomorphic*. We say that types  $A$  and  
 415  $B$  are isomorphic, denoted as  $\Theta \vdash A \cong^\pm B$ , if both  $\Theta \vdash A \leq^\pm B$  and  $\Theta \vdash B \leq^\pm A$  hold (see  
 416 Figure 5). Indeed, this requirement corresponds exactly to the premises of the shift rules. To  
 417 our knowledge, this is the first time that Cardelli’s restriction has been presented as part of  
 418 a declarative specification of typing.

### 419 3 EXAMPLES

420 In this section we characterize the behavior of our system by giving some examples of terms and  
 421 subtyping relationships that are permitted and some that are not.

#### 422 3.1 Subtyping

- 423
- 424 • Our subtyping relation allows us to swap quantifiers at identical depths. As indicated by the
  - 425 subtyping in both directions below, we can substitute a term of either of the types below
  - 426 for a term of the other type:

$$427 \quad \forall\alpha, \beta. \downarrow(\alpha \rightarrow \uparrow\beta) \rightarrow [\alpha] \rightarrow \uparrow[\beta] \quad \stackrel{\leq^-}{\leq^+} \quad \forall\beta, \alpha. \downarrow(\alpha \rightarrow \uparrow\beta) \rightarrow [\alpha] \rightarrow \uparrow[\beta]$$

- 428
- 429 • Subtypes can push in quantifiers as long as those quantifiers do not cross a shift boundary:

$$430 \quad \forall\alpha. \alpha \rightarrow (\forall\beta. \beta \rightarrow \uparrow(\alpha \times \beta)) \quad \leq^- \quad \forall\alpha, \beta. \alpha \rightarrow \beta \rightarrow \uparrow(\alpha \times \beta)$$

- 431
- 432 • Unusually, our subtyping relation allows the instantiations of type variables to be im-
  - 433 predicated. For example, here we can instantiate the  $\alpha$  in  $[\alpha]$  to be an identity function
  - 434  $\forall\beta. \beta \rightarrow \uparrow\beta$ , giving us a list of *polymorphic* identity functions. Predicative systems would
  - 435 allow us to produce a list of identity functions that are all parameterized over the same type,
  - 436 i.e.  $\forall\beta. [\beta \rightarrow \uparrow\beta]$ , but the list produced by our system  $[\forall\beta. \beta \rightarrow \uparrow\beta]$  is truly polymorphic
  - 437 with each element of the list having its own type variable.
- 438  
439  
440  
441

442	head	: $\downarrow(\forall\alpha. [\alpha] \rightarrow \uparrow\alpha)$	inc	: $\downarrow(\text{Int} \rightarrow \uparrow\text{Int})$
443	tail	: $\downarrow(\forall\alpha. [\alpha] \rightarrow \uparrow[\alpha])$	choose	: $\downarrow(\forall\alpha. \alpha \rightarrow \alpha \rightarrow \uparrow\alpha)$
444	[]	: $\downarrow(\forall\alpha. \uparrow[\alpha])$	poly	: $\downarrow(\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha) \rightarrow \uparrow(\text{Int} \times \text{Bool}))$
445	(::)	: $\downarrow(\forall\alpha. \alpha \rightarrow [\alpha] \rightarrow \uparrow[\alpha])$	auto	: $\downarrow(\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha) \rightarrow (\forall\alpha. \alpha \rightarrow \uparrow\alpha))$
446	single	: $\downarrow(\forall\alpha. \alpha \rightarrow \uparrow[\alpha])$	auto'	: $\downarrow(\forall\alpha. \downarrow(\forall\beta. \beta \rightarrow \uparrow\beta) \rightarrow \alpha \rightarrow \uparrow\alpha)$
447	append	: $\downarrow(\forall\alpha. [\alpha] \rightarrow [\alpha] \rightarrow \uparrow[\alpha])$	app	: $\downarrow(\forall\alpha, \beta. \downarrow(\alpha \rightarrow \uparrow\beta) \rightarrow \alpha \rightarrow \uparrow\beta)$
448	length	: $\downarrow(\forall\alpha. [\alpha] \rightarrow \uparrow\text{Int})$	revapp	: $\downarrow(\forall\alpha, \beta. \alpha \rightarrow \downarrow(\alpha \rightarrow \uparrow\beta) \rightarrow \uparrow\beta)$
449	map	: $\downarrow(\forall\alpha, \beta. \downarrow(\alpha \rightarrow \beta) \rightarrow [\alpha] \rightarrow \uparrow[\beta])$	flip	: $\downarrow(\forall\alpha, \beta, \gamma. \downarrow(\alpha \rightarrow \beta \rightarrow \uparrow\gamma) \rightarrow \beta \rightarrow \alpha \rightarrow \uparrow\gamma)$
450	id	: $\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)$	runST	: $\downarrow(\forall\alpha. \downarrow(\forall\beta. \text{ST } \beta \ \alpha) \rightarrow \uparrow\alpha)$
451	ids	: $\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)$	argST	: $\downarrow(\forall\alpha. \text{ST } \alpha \ \text{Int})$
452	a9	: $\downarrow(\forall\alpha. \downarrow(\alpha \rightarrow \uparrow\alpha) \rightarrow [\alpha] \rightarrow \uparrow\alpha)$	c8	: $\downarrow(\forall\alpha. [\alpha] \rightarrow [\alpha] \rightarrow \uparrow\alpha)$
453	h	: $\downarrow(\text{Int} \rightarrow (\forall\alpha. \alpha \rightarrow \uparrow\alpha))$	k	: $\downarrow(\forall\alpha. \alpha \rightarrow [\alpha] \rightarrow \uparrow\alpha)$
454	lst	: $\downarrow(\forall\alpha. \text{Int} \rightarrow \alpha \rightarrow \uparrow\alpha)$	r	: $\downarrow(\downarrow(\forall\alpha. \alpha \rightarrow \forall\beta. \beta \rightarrow \uparrow\beta) \rightarrow \uparrow\text{Int})$

Fig. 6. Environment for the examples from Serrano et al. [19], translated into Implicit Polarized F

$$\forall\alpha. \uparrow[\alpha] \leq^- \uparrow[\downarrow(\forall\beta. \beta \rightarrow \uparrow\beta)]$$

- An important restriction of our system that is essential to decidability is that types underneath shifts must be isomorphic. For example, we must establish both  $M \leq^- N$  and  $N \leq^- M$  before we can infer that  $\downarrow M \leq^+ \downarrow N$ .

This means that terms of either of the types below may not be substituted for terms of the other type. While  $\forall\alpha, \beta. \alpha \rightarrow \beta \rightarrow \uparrow(\alpha \times \beta)$  is indeed a subtype of  $\text{Int} \rightarrow \text{String} \rightarrow \uparrow(\text{Int} \times \text{String})$ , the subtyping does not hold the other way around. Therefore in this example where we zip a list of integers and a list of strings together, the type of the pairing function  $\downarrow(\text{Int} \rightarrow \text{String} \rightarrow \uparrow(\text{Int} \times \text{String}))$  is neither a subtype nor a supertype of the type of the generic pairing function  $\downarrow(\forall\alpha, \beta. \alpha \rightarrow \beta \rightarrow \uparrow(\alpha \times \beta))$ .

$$\downarrow(\text{Int} \rightarrow \text{String} \rightarrow \uparrow(\text{Int} \times \text{String})) \rightarrow [\text{Int}] \rightarrow [\text{String}] \rightarrow \uparrow[\text{Int} \times \text{String}]$$

$$\downarrow(\forall\alpha, \beta. \alpha \rightarrow \beta \rightarrow \uparrow(\alpha \times \beta)) \rightarrow [\text{Int}] \rightarrow [\text{String}] \rightarrow \uparrow[\text{Int} \times \text{String}]$$

### 3.2 Typing

To illustrate Implicit Polarized F's type system, we translate the examples used in the comparison of type systems in Serrano et al. [19] (GI) to Implicit Polarized F and investigate which of these translations (a) typecheck without needing additional annotations beyond the GI example, (b) typecheck but require additional annotations beyond the GI example, and (c) do not typecheck. These examples make use of a few common functions whose types we give in Figure 6.

Despite the extreme bias towards simplicity that is reflected in our straightforward local heuristic, surprisingly many of these examples typecheck. We discuss a few characteristic ones below; the full set are in Figure 7.

	GI example	Implicit Polarized F translation		
491	A1	const2 = $\lambda x y. y$	let const2 = {return { $\Lambda\alpha. \Lambda\beta. \lambda x : \alpha. \lambda y : \beta. \text{return } y$ }}; ...	Ann
492	A2	choose id	$\lambda x : \Downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha). (\text{let } t = \text{choose id } x; \text{return } t)$	Ann
493	A3	choose [] ids	let n : [ $\Downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)$ ] = []; let t = choose n ids; return t	Ann
494	A4	$\lambda x : (\forall\alpha. \alpha \rightarrow \alpha). x x$	$\lambda x : (\forall\alpha. \alpha \rightarrow \uparrow\alpha). (\text{let } y = x x; \text{return } y)$	✓
495	A5	id auto	let t = id auto; return t	✓
496	A6	id auto'	let t = id auto'; return t	✓
497	A7	choose id auto	let t = choose id auto; return t	×
498	A8	choose id auto'	let t = choose id auto'; return t	×
499	A9	a9 (choose id) ids	let f = {return { $\lambda x : \Downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha).$ let y = choose id x; return y}}};	Ann
500			let t = a9 f ids; return t	
501			let t = poly id; return t	✓
502	A10	poly id	let t = poly { $\Lambda\alpha. \lambda x : \alpha. \text{return } x$ }; return t	Ann
503	A11	poly ( $\lambda x. x$ )	let t = poly { $\Lambda\alpha. \lambda x : \alpha. \text{return } x$ }; return t	Ann
504	A12	id poly ( $\lambda x. x$ )	let x = id poly; let t = x { $\Lambda\alpha. \lambda y : \alpha. \text{return } y$ }; return t	Ann
505	B1	$\lambda f. (f 1, f \text{True})$	$\lambda f : \Downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha). (\text{let } l = f 1; \text{let } r = f \text{true}; \text{return } (l, r))$	Ann
506	B2	$\lambda xs. \text{poly } (\text{head } xs)$	$\lambda xs : [\Downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)]. (\text{let } x = \text{head } xs; \text{let } y = \text{poly } x; \text{return } y)$	Ann
507	C1	length ids	let t = length ids; return t	✓
508	C2	tail ids	let t = tail ids; return t	✓
509	C3	head ids	let t = head ids; return t	✓
510	C4	single id	let t = single id; return t	✓
511	C5	cons id ids	let t = cons id ids; return t	✓
512	C6	cons ( $\lambda x. x$ ) ids	let t = cons { $\Lambda\alpha. \lambda x : \alpha. \text{return } x$ } ids; return t	Ann
513	C7	append (single inc) (single ids)	let x = single inc; let y = single id; let t = append x y; return t	×
514	C8	c8 (single id) ids	let x = single id; let t = c8 x ids; return t	✓
515	C9	map poly (single id)	let x = single id; let t = map poly x; return t	✓
516	C10	map head (single ids)	let x = single ids; let t = map head x; return t	×
517	D1	app poly id	let t = app poly id; return t	✓
518	D2	revapp id poly	let t = revapp id poly; return t	✓
519	D3	runST argST	let t = runST argST; return t	✓
520	D4	app runST argST	let t = app runST argST; return t	×
521	D5	revapp argST runST	let t = revapp argST runST; return t	×
522	E1	k h lst	let t = k h lst; return t	×
523	E2	k ( $\lambda x. h x$ ) lst	let f = {return { $\Lambda\alpha. \lambda y : \text{Int}.$ return { $\lambda x : \alpha. \text{let } z = h y x; \text{return } z$ }}}}};	×
524			let t = k f lst; return t	
525	E3	r ( $\lambda x y. y$ )	let t = r { $\Lambda\alpha. \Lambda\beta. \lambda x : \alpha. \lambda y : \beta. \text{return } y$ }; return t	×
526				

Fig. 7. Examples used in the comparison of type systems in Serrano et al. [19] along with their translations to Implicit Polarized F. In the right hand column:

- ✓ indicates that an example typechecks without needing additional annotations beyond the GI example.
- Ann indicates that the example typechecks, but additional annotations are required beyond the GI example.
- × indicates that the example does not typecheck in Implicit Polarized F.

540 *Shifts make types much more fine-grained.* The shift structure in Implicit Polarized F types makes  
 541 them much more fine-grained compared to systems like GI. For example, in Implicit Polarized F  
 542 the number of arguments we can pass to a function is encoded in its type. If we instantiate `id` at  
 543  $\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)$ , it still takes only a single argument. In contrast, in systems like GI, instantiating `id`  
 544 at  $\forall\alpha. \alpha \rightarrow \alpha$  gives us a function that can take two arguments: first `id` and then an arbitrary value.

545 As a result, while `id` instantiated at  $\forall\alpha. \alpha \rightarrow \alpha$  and `auto` represent the same thing in systems like  
 546 GI, `id` instantiated at  $\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)$  and `auto` represent fundamentally different things in Implicit  
 547 Polarized F. Therefore in A7 we can not synthesize a consistent return type for `choose id auto` and  
 548 so it does not typecheck.

549 Examples A8, C7, and E2 do not work for similar reasons. For instance, in C7 our use of invariant  
 550 shifts prevents us from using single `ids` at a less polymorphic type.

551 Another example of the shift structure in Implicit Polarized F making types much more fine-  
 552 grained is A3: a direct translation `let t = choose [] ids; return t` will not typecheck since `[]` has a  
 553 shift at its top-level while `ids` does not. As a result, we need to unwrap the top-level shift in the  
 554 type of `[]` using a `let`-binding like so: `let n : [ $\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)$ ] = []`.  
 555

556 *Hyperlocal inference.* Inference in Implicit Polarized F refuses to incorporate far away information.  
 557 As encoded in the last premise of the `Dunambiguouslet` rule, we must annotate a `let`-bound variable  
 558 unless its type is completely determined by the application  $v(s)$ .

559 For example, in A3 there are many types we can infer for the empty list constructor `[]`, for instance  
 560 `[Int]`, `[ $\forall\alpha. \alpha$ ]`, and  `$\forall\alpha. [\alpha]$` . We can see from the later application `choose n ids` that `[ $\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)$ ]`  
 561 is the right choice for this type. However, Implicit Polarized F only considers the immediate context  
 562 of the function application. In this context the type for `n` is ambiguous, so we need to annotate the  
 563 `let`-binding.

564 While A3 needs an annotation, examples like C4 do not. While the type `[ $\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha)$ ]` that is  
 565 synthesized for the `let`-bound variable in C4 contains universal quantification, this quantification  
 566 occurs underneath a shift. Since shifts are invariant, we can not perform subtyping underneath  
 567 them, so this type is the only one modulo isomorphism that we can infer for the `let`-bound variable.  
 568

569 *We sometimes need to eta-expand terms to get them to type check.* An example of this is C10  
 570 `map head (single ids)`. Here `map` expects an argument of type  $\downarrow(\alpha \rightarrow \uparrow\beta)$ , but `head` has type  
 571  $\downarrow(\forall\alpha. [\alpha] \rightarrow \uparrow\alpha)$ . Since shifts are invariant in Implicit Polarized F, these types are incompatible.  
 572 We can remedy this by eta-expanding the term, for instance for C10 we can create a new term  
 573  $\lambda x : \downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha). \text{let } y = \text{head } x; \text{return } y$  that specializes `head` to the type  $\downarrow([\downarrow\forall\alpha. \alpha \rightarrow \uparrow\alpha] \rightarrow$   
 574  $\uparrow\downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha))$ . Examples D4, D5, E1, and E3 fail to typecheck for similar reasons.

575 Another situation where eta-expansion is necessary is when replicating partial application. Since  
 576 Implicit Polarized F does not allow partial application, partial applications like `choose id` in GI  
 577 must be represented as  $\lambda x : \downarrow(\forall\alpha. \alpha \rightarrow \uparrow\alpha). \text{let } y = \text{choose id } x; \text{return } y$  in Implicit Polarized F.  
 578 Since we syntactically require users to annotate lambda-bound variables, these examples all require  
 579 annotations in Implicit Polarized F. This affects examples A2, A9, and E2.  
 580

581 *Implicit Polarized F tends to prefer inferring impredicative types to inferring type schemes.* Adding  
 582 impredicativity to Hindley-Milner style systems means that, without changes to the type language  
 583 like the introduction of bounded quantification, terms are no longer guaranteed to have a most  
 584 general type. For instance, when typing `single id` we find that `[ $\forall\alpha. \alpha \rightarrow \alpha$ ]` and  `$\forall\alpha. [\alpha \rightarrow \alpha]$`  are  
 585 equally general. Due to its simple local heuristic, Implicit Polarized F will instantiate the type  
 586 variable in `single` impredicatively, thereby typing `single` as `[ $\forall\alpha. \alpha \rightarrow \alpha$ ]`. In GI, the choice we make  
 587 depends on whether the type variable  $\alpha$  in the definition of `single` appears under a type constructor  
 588

Positive types	$P ::= \dots \mid \hat{\alpha}$
Contexts	$\Theta ::= \dots \mid \Theta, \hat{\alpha} \mid \Theta, \hat{\alpha} = P$

Fig. 8. Additions to declarative types and contexts to form their algorithmic counterparts

in any of the arguments to it. Since it does not, GI assigns  $\alpha$  a “top-level monomorphic” type, thereby typing single as  $\forall\alpha. [\alpha \rightarrow \alpha]$ .

Implicit Polarized F’s preference for inferring impredicative types is not an unquestionable pro or con: while it allows us to infer examples like C8 and C9 that systems that prefer type schemes can not, systems that prefer type schemes can deal better with eta-reduced terms.

## 4 ALGORITHMIC TYPING

In this section, we discuss how to implement our declarative system algorithmically.

There are two rules that prevent us from directly implementing the declarative system. These are  $\text{Dspinetypeabs}$  and  $\leq^{\pm}\text{Dforall}$ , both of which have a type  $P$  that appears in the premises but not in the conclusion of each rule. As a result, a direct implementation of type inference would need to somehow invent a type  $P$  out of thin air.

To get an algorithm, we follow Dunfield and Krishnaswami [5] in replacing these types  $P$  with *existential type variables*  $\hat{\alpha}$ . These represent unsolved types that the algorithm will solve at a later stage. Since type variables and their solutions are positive, existential type variables must also be positive.

The introduction of existential variables has some knock-ons on our system: each judgment needs to be adapted to manage existential variables and their instantiation. For instance, as shown in Figure 8, the contexts of algorithmic judgments will contain not only type variables, but also solved and unsolved existential type variables.

### 4.1 Algorithmic subtyping

We first describe the algorithmic judgments for subtyping. These will be used to solve all relevant existential type variables.

Our algorithm for subtyping consists of two mutually recursive judgments  $\Theta \vdash P \leq^+ Q \dashv \Theta'$  and  $\Theta \vdash N \leq^- M \dashv \Theta'$ . Each judgment takes an input context  $\Theta$ , checks that the left-hand side type is a subtype of the right-hand side type, and produces an output context  $\Theta'$ . The shape of this output context is identical to that of the corresponding input context: the only difference between them is that the output context contains the solutions to any newly solved existential variables. Since existential type variables can only have a single solution, output contexts can only solve existential variables that are unsolved in the input context.

As a result of polarization, existential type variables only appear on the left-hand side of negative subtyping judgments and the right-hand side of positive subtyping judgments. This *groundness invariant* is critical to the workings of our algorithm. Since we will frequently make use of this invariant, we briefly introduce some terminology related to it. We call types that do not contain any existential variables *ground*. The side of a judgment that can not contain any existential type variables is the *ground side*, and the opposite side is the *non-ground side*. We will observe symmetries between properties of the positive and negative judgments based on the groundness of the types (see Lemma 6.1 for instance). This is why we swap the alphabetical order of  $M$  and  $N$  when writing negative judgments.

$\Theta \vdash A \leq^{\pm} B \dashv \Theta'$

In the context  $\Theta$ ,  $A$  checks algorithmically as a subtype of  $B$ , producing the context  $\Theta'$

$$\frac{}{\Theta_L, \alpha, \Theta_R \vdash \alpha \leq^+ \alpha \dashv \Theta_L, \alpha, \Theta_R} \leq^{\pm} \text{Arefl} \quad \frac{\Theta_L \vdash P \text{ type}^+ \quad P \text{ ground}}{\Theta_L, \hat{\alpha}, \Theta_R \vdash P \leq^+ \hat{\alpha} \dashv \Theta_L, \hat{\alpha} = P, \Theta_R} \leq^{\pm} \text{Ainst}$$

$$\frac{\Theta \vdash M \leq^- N \dashv \Theta' \quad \Theta' \vdash N \leq^- [\Theta']M \dashv \Theta''}{\Theta \vdash \downarrow N \leq^+ \downarrow M \dashv \Theta''} \leq^{\pm} \text{Ashift}\downarrow$$

$$\frac{\Theta, \hat{\alpha} \vdash [\hat{\alpha}/\alpha]N \leq^- M \dashv \Theta', \hat{\alpha} [= P] \quad M \neq \forall\beta. M'}{\Theta \vdash \forall\alpha. N \leq^- M \dashv \Theta'} \leq^{\pm} \text{Aforallll}$$

$$\frac{\Theta, \alpha \vdash N \leq^- M \dashv \Theta', \alpha}{\Theta \vdash N \leq^- \forall\alpha. M \dashv \Theta'} \leq^{\pm} \text{Aforallr} \quad \frac{\Theta \vdash Q \leq^+ P \dashv \Theta' \quad \Theta' \vdash [\Theta']N \leq^- M \dashv \Theta''}{\Theta \vdash P \rightarrow N \leq^- Q \rightarrow M \dashv \Theta''} \leq^{\pm} \text{Aarrow}$$

$$\frac{\Theta \vdash Q \leq^+ P \dashv \Theta' \quad \Theta' \vdash [\Theta']P \leq^+ Q \dashv \Theta''}{\Theta \vdash \uparrow P \leq^- \uparrow Q \dashv \Theta''} \leq^{\pm} \text{Ashift}\uparrow$$

Fig. 9. Algorithmic subtyping

We present our algorithmic subtyping rules in Figure 9.

- $\leq^{\pm} \text{Arefl}$  is identical to our declarative rule for type variables, however we have moved the  $\Theta \vdash \alpha \text{ type}^+$  judgment into an equivalent requirement that the input and output contexts have the form  $\Theta_L, \alpha, \Theta_R$ .
- $\leq^{\pm} \text{Ainst}$  is a new rule that describes how to solve existential type variables. The idea is that if we reach a subtyping judgment like  $P \leq^+ \hat{\alpha}$ , then there is no further work to be done, and we can just set  $\hat{\alpha}$  to be equal to  $P$ . Since this rule is the only way we solve existential variables, all solutions to existential variables are ground.

The premises of  $\leq^{\pm} \text{Ainst}$  ensure that the rule preserves well-formedness: as shown in Figure 14, our contexts are ordered, and a solution  $P$  to an existential variable  $\hat{\alpha}$  must be well-formed with respect to the context items prior to  $\hat{\alpha}$ , i.e.  $\Theta_L$ . This means that the solution for  $\hat{\alpha}$  can only contain type variables that were in scope when  $\hat{\alpha}$  was introduced: an important property for the soundness of the system.

- $\leq^{\pm} \text{Aforallll}$  implements the strategy we described to make the  $\leq^{\pm} \text{Dforallll}$  rule algorithmic: rather than replacing the type variable  $\alpha$  with a ground type  $P$ , we procrastinate deciding what  $P$  should be and instead replace  $\alpha$  with a fresh existential type variable  $\hat{\alpha}$ . The bracketed  $\hat{\alpha} [= P]$  in the output context indicates that the algorithm doesn't have to solve  $\hat{\alpha}$ : we will later see in Section 6.1.1 that it solves  $\hat{\alpha}$  if and only if the type variable  $\alpha$  appears in the term  $N$ .

The second premise of  $\leq^{\pm} \text{Aforallll}$  states that the right-hand side must not be a prenex polymorphic form. We introduce this premise to ensure that the system is strictly syntax-oriented: we have to eliminate all the quantifiers on the right with  $\leq^{\pm} \text{Aforallr}$  before eliminating any quantifiers on the left with  $\leq^{\pm} \text{Aforallll}$ . In general, introducing a premise like this could be problematic for completeness, as in completeness (see Theorem 7.2) we want to be able to take apart the declarative derivation in the same order that the algorithmic

687  $\boxed{[\Theta]A}$  Applying a context  $\Theta$ , as a substitution, to a type  $A$

$$\begin{array}{ll}
 688 & [\cdot]A = A & [\Theta, \hat{\alpha}]A = [\Theta]A \\
 689 & [\Theta, \alpha]A = [\Theta]A & [\Theta, \hat{\alpha} = P]A = [\Theta]([P/\hat{\alpha}]A)
 \end{array}$$

692 Fig. 10. Applying a context to a type

693  
694  
695 system would derive the term. Thankfully, the invertibility of  $\leq^{\pm}D\text{forallr}$  ensures that we  
696 can transform declarative derivations concluding with  $\leq^{\pm}D\text{forallll}$  into equivalent ones that,  
697 like the algorithmic systems, use  $\leq^{\pm}D\text{forallr}$  as much as possible to eliminate all the prenex  
698 quantifiers on the right-hand side before using  $\leq^{\pm}D\text{forallll}$ .

- 699 •  $\leq^{\pm}A\text{forallr}$  is almost an exact copy of  $\leq^{\pm}D\text{forallll}$ . The forall rules must preserve the ground-  
700 ness invariant, so while  $\leq^{\pm}A\text{forallll}$  can introduce existential type variables when eliminating  
701 a quantifier on the left-hand side,  $\leq^{\pm}A\text{forallr}$  can not introduce any existential variables  
702 while eliminating a quantifier on the right-hand side. This ensures that  $M$  remains ground  
703 from the conclusion of the rule to its premise.
- 704 •  $\leq^{\pm}A\text{arrow}$  describes how to check function types. This rule is interesting since it has  
705 multiple premises. After checking the first premise  $\Theta \vdash Q \leq^+ P \dashv \Theta'$ , we have an output  
706 context  $\Theta'$  that might solve some existential variables that appear in  $P$ . Since some of these  
707 existential variables might also appear in  $N$ , we need to ensure that we propagate their  
708 solutions when we check  $N \leq^- M$ . We do this by substituting all the existential variables  
709 that appear in the possibly non-ground  $N$  by their solutions in  $\Theta'$ . We denote this operation  
710 as  $[\Theta']N$  and define it formally in Figure 10.
- 711 •  $\leq^{\pm}A\text{shift}\downarrow$  and the symmetric  $\leq^{\pm}A\text{shift}\uparrow$  are our shift rules. In these rules the groundness  
712 invariant defines the order in which the premises need to be checked. For instance, in  
713  $\leq^{\pm}A\text{shift}\downarrow$  only  $N$  is ground in the conclusion, therefore we need to check  $M \leq^- N$  first. As  
714 we will see in Section 6.1.1, checking this gives us an algorithmic context  $\Theta'$  that solves  
715 all the existential variables that appear in  $M$ . We can use this context to *complete*  $M$  by  
716 substituting all the existential variables in  $M$  by their solutions in  $\Theta'$ , giving us the ground  
717 type  $[\Theta']M$ . Now that we have a ground type  $[\Theta']M$ , we can check  $N \leq^- M$  by verifying  
718  $\Theta' \vdash N \leq^- [\Theta']M \dashv \Theta''$ .

## 720 4.2 Algorithmic type system

721 With a mechanism to solve existential variables in hand, we can now construct an algorithmic  
722 system to implement our declarative typing rules. We present our algorithmic type system in  
723 Figure 11.

- 724 • Most of the rules ( $A\text{var}$ ,  $A\lambda\text{abs}$ ,  $A\text{gen}$ ,  $A\text{thunk}$ ,  $A\text{return}$ ,  $A\text{spinenil}$ , and  $A\text{spinecons}$ ) are  
725 identical to their declarative counterparts, modulo adding output contexts and applying  
726 output contexts to subsequent premises.
- 727 • We split the declarative  $D\text{spinetypeabs}$  rule into two rules depending on whether the new  
728 universal variable  $\alpha$  appears in the type we are quantifying over  $N$ . If it does not, as in  
729  $A\text{spinetypeabsnotin}$ , then we do not need to introduce a new existential variable. If it does,  
730 as in  $A\text{spinetypeabsin}$ , then as with  $\leq^{\pm}A\text{forallll}$  we introduce a fresh existential variable  $\hat{\alpha}$   
731 to replace  $\alpha$ . Note that unlike most rules,  $A\text{spinetypeabsin}$  adds a new existential variable  
732 to the output context of the conclusion *that does not appear in its input context*. This will  
733 prove to be important in the  $A\text{unambiguouslet}$  and  $A\text{ambiguouslet}$  rules.

736  $\boxed{\Theta; \Gamma \vdash e : A \dashv \Theta'}$  The term  $e$  synthesizes the type  $A$ , producing the context  $\Theta'$

737

738  $\boxed{\Theta; \Gamma \vdash s : N \gg M \dashv \Theta'}$  When passed to a head of type  $N$ , the argument list  $s$  synthesizes

739 the type  $M$ , producing the context  $\Theta'$

740

741 
$$\frac{x : P \in \Gamma}{\Theta; \Gamma \vdash x : P \dashv \Theta} \text{Avar}$$

742

743 
$$\frac{\Theta; \Gamma, x : P \vdash t : N \dashv \Theta'}{\Theta; \Gamma \vdash \lambda x : P. t : P \rightarrow N \dashv \Theta'} \text{A}\lambda\text{abs} \qquad \frac{\Theta, \alpha; \Gamma \vdash t : N \dashv \Theta', \alpha}{\Theta; \Gamma \vdash \Lambda \alpha. t : \forall \alpha. N \dashv \Theta'} \text{A}\text{gen}$$

744

745

746 
$$\frac{\Theta; \Gamma \vdash t : N \dashv \Theta'}{\Theta; \Gamma \vdash \{t\} : \downarrow N \dashv \Theta'} \text{A}\text{thunk} \qquad \frac{\Theta; \Gamma \vdash v : P \dashv \Theta'}{\Theta; \Gamma \vdash \text{return } v : \uparrow P \dashv \Theta'} \text{A}\text{return}$$

747

748

749

750 
$$\frac{\Theta; \Gamma \vdash v : \downarrow M \dashv \Theta' \quad \Theta'; \Gamma \vdash s : M \gg \uparrow Q \dashv \Theta'' \quad \Theta'' \vdash P \leq^+ Q \dashv \Theta''' \quad \Theta''' \vdash [\Theta''']Q \leq^+ P \dashv \Theta^{(4)} \quad \Theta^{(5)} = \Theta^{(4)} \upharpoonright \Theta \quad \Theta^{(5)}; \Gamma, x : P \vdash t : N \dashv \Theta^{(6)}}{\Theta; \Gamma \vdash \text{let } x : P = v(s); t : N \dashv \Theta^{(6)}} \text{Aambiguouslet}$$

751

752

753

754 
$$\frac{\Theta; \Gamma \vdash v : \downarrow M \dashv \Theta' \quad \Theta'; \Gamma \vdash s : M \gg \uparrow Q \dashv \Theta'' \quad \text{FEV}(Q) = \emptyset \quad \Theta''' = \Theta'' \upharpoonright \Theta \quad \Theta'''; \Gamma, x : Q \vdash t : N \dashv \Theta^{(4)}}{\Theta; \Gamma \vdash \text{let } x = v(s); t : N \dashv \Theta^{(4)}} \text{Aunambiguouslet}$$

755

756

757

758

759

760 
$$\frac{}{\Theta; \Gamma \vdash \epsilon : N \gg N \dashv \Theta} \text{Aspinenil}$$

761

762 
$$\frac{\Theta; \Gamma \vdash v : P \dashv \Theta' \quad \Theta' \vdash P \leq^+ [\Theta']Q \dashv \Theta'' \quad \Theta''; \Gamma \vdash s : [\Theta'']N \gg M \dashv \Theta'''}{\Theta; \Gamma \vdash v, s : Q \rightarrow N \gg M \dashv \Theta'''} \text{Aspinecons}$$

763

764

765 
$$\frac{\Theta; \Gamma \vdash s : N \gg M \dashv \Theta' \quad \alpha \notin \text{FUV}(N)}{\Theta; \Gamma \vdash s : (\forall \alpha. N) \gg M \dashv \Theta'} \text{Aspinetypeabsnotin}$$

766

767

768 
$$\frac{\Theta, \hat{\alpha}; \Gamma \vdash s : [\hat{\alpha}/\alpha]N \gg M \dashv \Theta', \hat{\alpha} [= P] \quad \alpha \in \text{FUV}(N)}{\Theta; \Gamma \vdash s : (\forall \alpha. N) \gg M \dashv \Theta', \hat{\alpha} [= P]} \text{Aspinetypeabsin}$$

769

770

771

772

773

774

775

Fig. 11. Algorithmic type system

- The first, second, and last premises of `Aambiguouslet` are the same as their declarative counterparts. The third and fourth just inline the algorithmic  $\leq^\pm \text{Ashift} \uparrow$  rule corresponding to the declarative premise  $\Theta \vdash \uparrow Q \leq^- \uparrow P$  in `Dambiguouslet`. We saw in the `Aspinetypeabsin` rule that performing type inference on spines can introduce new existential variables. To simplify our proofs, we do not want these to leak. Therefore we introduce a notion of *context restriction* in Figure 12. The fifth premise  $\Theta^{(5)} = \Theta^{(4)} \upharpoonright \Theta$  creates a new context that restricts the context  $\Theta^{(4)}$  to contain only the existential variables in  $\Theta$ . So if  $\Theta^{(4)}$  contains new existential variables compared to  $\Theta$ , then these will not be



$$\boxed{\Theta' \upharpoonright \Theta} \quad \Theta' \text{ restricted to only contain existential variables which appear in } \Theta$$

$$\frac{}{\cdot \upharpoonright \cdot = \cdot} \upharpoonright_{\text{empty}} \qquad \frac{\Theta' \upharpoonright \Theta = \Theta''}{\Theta', \alpha \upharpoonright \Theta, \alpha = \Theta'', \alpha} \upharpoonright_{\text{uvar}}$$

$$\frac{\Theta' \upharpoonright \Theta = \Theta''}{\Theta', \hat{\alpha} [= P] \upharpoonright \Theta, \hat{\alpha} [= Q] = \Theta'', \hat{\alpha} [= P]} \upharpoonright_{\text{guessin}} \qquad \frac{\Theta' \upharpoonright \Theta = \Theta'' \quad \hat{\alpha} [= Q] \notin \Theta}{\Theta', \hat{\alpha} [= P] \upharpoonright \Theta = \Theta''} \upharpoonright_{\text{guessnotin}}$$

Fig. 12. Definition of context restriction

present in  $\Theta^{(5)}$ . However any new solutions in  $\Theta^{(4)}$  to existential variables already in  $\Theta$  will be present in  $\Theta^{(5)}$ , so the algorithm will not solve any existential variable as two incompatible solutions.

- The first, second, and last premises of `Aunambiguouslet` are lifted from the declarative rule `Dunambiguouslet`. We replace the quantification in the last premise of `Dunambiguouslet` with a statement that there are no existential variables left in the type  $Q$ , i.e.  $\text{FEV}(Q) = \emptyset$ . Since no instantiations are possible, there are not any other types left that  $Q$  could have. Like `Aambiguouslet`, we restrict the output context of the spine judgment to remove any existential variables newly introduced by the spine.

Our algorithm is really easy to implement – we wrote a bare-bones implementation in 250 lines of OCaml.

## 5 PROPERTIES OF DECLARATIVE TYPING

Implicit Polarized F should support the full complement of metatheoretic properties, but in this paper we focus on the properties needed to establish our theorems about type inference. For example, we will only prove substitution at the type level, and totally ignore term-level substitution.

### 5.1 Subtyping

Because we use subtyping to model type instantiation, we need to know quite a few properties about how subtyping works. For example, we are able to show that the subtyping relation admits both reflexivity and transitivity:

LEMMA 5.1 (DECLARATIVE SUBTYPING IS REFLEXIVE). *If  $\Theta \vdash A \text{ type}^\pm$  then  $\Theta \vdash A \leq^\pm A$ .*

LEMMA 5.2 (DECLARATIVE SUBTYPING IS TRANSITIVE). *If  $\Theta \vdash A \text{ type}^\pm$ ,  $\Theta \vdash B \text{ type}^\pm$ ,  $\Theta \vdash C \text{ type}^\pm$ ,  $\Theta \vdash A \leq^\pm B$ , and  $\Theta \vdash B \leq^\pm C$ , then  $\Theta \vdash A \leq^\pm C$ .*

Because of impredicativity, transitivity is surprisingly subtle to get right. We discuss the needed metric in Section 6.2 as this metric is also used to show the decidability of algorithmic subtyping.

We also show that subtyping is stable under substitution, which is useful for proving properties of type instantiation:

LEMMA 5.3 (DECLARATIVE SUBTYPING IS STABLE UNDER SUBSTITUTION). *If  $\Theta_L, \Theta_R \vdash P \text{ type}^+$ , then:*

- *If  $\Theta_L, \alpha, \Theta_R \vdash Q \text{ type}^+$ ,  $\Theta_L, \alpha, \Theta_R \vdash R \text{ type}^+$ , and  $\Theta_L, \alpha, \Theta_R \vdash Q \leq^+ R$ , then  $\Theta_L, \Theta_R \vdash [P/\alpha]Q \leq^+ [P/\alpha]R$ .*
- *If  $\Theta_L, \alpha, \Theta_R \vdash N \text{ type}^-$ ,  $\Theta_L, \alpha, \Theta_R \vdash M \text{ type}^-$ , and  $\Theta_L, \alpha, \Theta_R \vdash N \leq^- M$ , then  $\Theta_L, \Theta_R \vdash [P/\alpha]N \leq^- [P/\alpha]M$ .*

$$\boxed{\Theta \vdash \Gamma_1 \cong \Gamma_2} \text{ In the context } \Theta, \text{ the environments } \Gamma_1 \text{ and } \Gamma_2 \text{ are isomorphic}$$

$$\frac{}{\Theta \vdash \cdot \cong \cdot} \text{Eisoempty} \qquad \frac{\Theta \vdash \Gamma_1 \cong \Gamma_2 \quad \Theta \vdash P \cong^+ Q}{\Theta \vdash \Gamma_1, x : P \cong \Gamma_2, x : Q} \text{Eisovar}$$

Fig. 13. Isomorphic environments

$$\boxed{\Theta \text{ ctx}} \text{ } \Theta \text{ is a well-formed context}$$

$$\frac{}{\cdot \text{ ctx}} \text{Cwfempty} \qquad \frac{\Theta \text{ ctx}}{\Theta, \alpha \text{ ctx}} \text{Cwfuvar} \qquad \frac{\Theta \text{ ctx}}{\Theta, \hat{\alpha} \text{ ctx}} \text{Cwfunolvedguess}$$

$$\frac{\Theta \text{ ctx} \quad \Theta \vdash P \text{ type}^+ \quad P \text{ ground}}{\Theta, \hat{\alpha} = P \text{ ctx}} \text{Cwfsolvedguess}$$

Fig. 14. Well-formedness of contexts

## 5.2 Typing

We also show an important property about our typing judgment. Suppose we have two types  $A$  and  $B$  that are isomorphic, i.e.  $\Theta \vdash A \cong^\pm B$ . Then, by the intuitive definition of subtyping, it should be possible to safely use any term of type  $A$  where a term of type  $B$  is expected, and vice versa. Concretely, the map function can be given both the type  $P = \downarrow(\forall \alpha, \beta. \downarrow(\alpha \rightarrow \beta) \rightarrow [\alpha] \rightarrow \uparrow[\beta])$  and also the type  $Q = \downarrow(\forall \beta, \alpha. \downarrow(\alpha \rightarrow \beta) \rightarrow [\alpha] \rightarrow \uparrow[\beta])$ , and there is no reason to prefer one to the other. Therefore if we have an environment  $\Gamma$  where  $\text{map} : P$  and an environment  $\Gamma'$  where  $\text{map} : Q$ , the inference we get within each environment should be the same.

We formalize this idea with the following lemma, which makes use of the isomorphic environment judgment we define in Figure 13:

LEMMA 5.4 (ISOMORPHIC ENVIRONMENTS TYPE THE SAME TERMS). *If  $\Theta \vdash \Gamma \cong \Gamma'$ , then:*

- *If  $\Theta; \Gamma \vdash v : P$  then  $\exists P'$  such that  $\Theta \vdash P \cong^- P'$  and  $\Theta; \Gamma' \vdash v : P'$ .*
- *If  $\Theta; \Gamma \vdash t : N$  then  $\exists N'$  such that  $\Theta \vdash N \cong^- N'$  and  $\Theta; \Gamma' \vdash t : N'$ .*
- *If  $\Theta; \Gamma \vdash s : N \gg M$  and  $\Theta \vdash N \cong^- N'$ , then  $\exists M'$  such that  $\Theta \vdash M \cong^- M'$  and  $\Theta; \Gamma \vdash s : N' \gg M'$ .*

This lemma tells us that regardless of which type we give the variable  $\text{map}$ , any term which uses it will have the same type (up to isomorphism).

## 6 PROPERTIES OF ALGORITHMIC TYPING

### 6.1 Well-formedness

We first establish some of the invariants our type system maintains.

#### 6.1.1 Subtyping.

*Context well-formedness.* A simple property that our algorithm maintains is that any solutions the algorithm chooses are well-formed: they are ground and only contain universal variables in

883  $\boxed{\Theta \vdash A \text{ type}^\pm}$  In the context  $\Theta$ ,  $A$  is a well-formed positive/negative type

$$884 \frac{\hat{\alpha} \in \text{EV}(\Theta)}{\Theta \vdash \hat{\alpha} \text{ type}^+} \text{Twfguess}$$

885  
886  
887  
888 Fig. 15. Additional well-formedness rules for algorithmic types.  $\text{EV}(\Theta)$  contains all the existential type  
889 variables in  $\Theta$ , independently of whether they are solved or unsolved.  
890

891  $\boxed{\Theta \longrightarrow \Theta'}$   $\Theta$  extends to  $\Theta'$

$$892 \frac{}{\cdot \longrightarrow \cdot} \text{Cempty} \quad \frac{\Theta \longrightarrow \Theta'}{\Theta, \alpha \longrightarrow \Theta', \alpha} \text{Cuvar} \quad \frac{\Theta \longrightarrow \Theta'}{\Theta, \hat{\alpha} \longrightarrow \Theta', \hat{\alpha}} \text{Cunsolvedguess}$$

$$893$$

$$894 \frac{\Theta \longrightarrow \Theta'}{\Theta, \hat{\alpha} \longrightarrow \Theta', \hat{\alpha} = P} \text{Csolveguess} \quad \frac{\Theta \longrightarrow \Theta' \quad \|\Theta\| \vdash P \cong^+ Q}{\Theta, \hat{\alpha} = P \longrightarrow \Theta', \hat{\alpha} = Q} \text{Csolvedguess}$$

$$895$$

$$896$$

$$897$$

$$898$$

$$899$$

900  
901 Fig. 16. Context extension  
902

903  $\boxed{\|\Theta\|}$  The declarative context corresponding to the algorithmic context  $\Theta$

$$904 \|\cdot\| = \cdot \quad \|\Theta, \alpha\| = \|\Theta\|, \alpha \quad \|\Theta, \hat{\alpha}\| = \|\Theta\| \quad \|\Theta, \hat{\alpha} = P\| = \|\Theta\|$$

$$905$$

$$906$$

$$907$$

908 Fig. 17. Producing a declarative context from an algorithmic context  
909  
910

911 scope at the time of the corresponding existential variable's creation. We formalize this notion of  
912 the well-formedness of contexts in Figures 14 and 15.

913 *Output context solves all existential variables.* Another property of our algorithm is that given  
914 well-formed inputs, the algorithm will solve all the necessary solutions: in other words, the output  
915 context contains solutions to all existential variables appearing in the non-ground type.  
916

917 *Context extension.* Finally, we want to show that the algorithm produces an output context that  
918 only adds solutions to the input context. For instance, our algorithm should not change an existing  
919 solution to an incompatible one. We formalize this idea with a context extension judgment based  
920 on the earlier work of Dunfield and Krishnaswami [5]. This judgment  $\Theta \longrightarrow \Theta'$ , as defined in  
921 Figure 16, indicates information gain:  $\Theta'$  must contain at least as much information as  $\Theta$ .

922 Most of the rules defining context extension are homomorphic. The exceptions are *Csolveguess*,  
923 which allows the algorithm to solve an existential variable, and *Csolvedguess*, which allows the  
924 algorithm to change the solution  $P$  for an existential variable  $\hat{\alpha}$  to an isomorphic type  $Q$ . This  
925 ability is not crucial to the algorithm; in fact the algorithm will maintain syntactically identical  
926 solutions. However since the declarative system can invent any solutions it likes when instantiating  
927 quantified types, we need this little bit of flex to prove completeness.

928 Solutions to existential types are ground, so we do not need to introduce a new notion of  
929 isomorphism over algorithmic types. However since the context  $\Theta$  is algorithmic rather than  
930 declarative, we do need to introduce a simple operation  $\|\Theta\|$  that converts an algorithmic context  
931

932  $\boxed{\Theta \Longrightarrow \Theta'}$  The context  $\Theta'$  weakly extends the context  $\Theta$

933

934  $\frac{}{\cdot \Longrightarrow \cdot} \text{Wcempty} \quad \frac{\Theta \Longrightarrow \Theta'}{\Theta, \alpha \Longrightarrow \Theta', \alpha} \text{Wcuvar} \quad \frac{\Theta \Longrightarrow \Theta'}{\Theta, \hat{\alpha} \Longrightarrow \Theta', \hat{\alpha}} \text{Wcunsolvedguess}$

935

936

937  $\frac{\Theta \Longrightarrow \Theta'}{\Theta, \hat{\alpha} \Longrightarrow \Theta', \hat{\alpha} = P} \text{Wcsolveguess} \quad \frac{\Theta \Longrightarrow \Theta' \quad \|\Theta\| \vdash P \cong^+ Q}{\Theta, \hat{\alpha} = P \Longrightarrow \Theta', \hat{\alpha} = Q} \text{Wcsolvedguess}$

938

939

940

941  $\frac{\Theta \Longrightarrow \Theta'}{\Theta \Longrightarrow \Theta', \hat{\alpha}} \text{Wcnewunsolvedguess} \quad \frac{\Theta \Longrightarrow \Theta'}{\Theta \Longrightarrow \Theta', \hat{\alpha} = P} \text{Wcnewsolvedguess}$

942

943

944

945 Fig. 18. Weak context extension. We highlight the rules that are “new” compared with context extension.

947  $\boxed{\Theta \vdash \Gamma \text{ env}}$  The environment  $\Gamma$  is well-formed with respect to the context  $\Theta$

948

949

950  $\frac{}{\Theta \vdash \cdot \text{ env}} \text{Ewfempty} \quad \frac{\Theta \vdash \Gamma \text{ env} \quad \Theta \vdash P \text{ type}^+ \quad P \text{ ground}}{\Theta \vdash \Gamma, x : P \text{ env}} \text{Ewfvvar}$

951

952

953

954 Fig. 19. Well-formedness of typing environments

955

956 into a declarative one by dropping all the context items involving existential variables. This

957 operation is defined in Figure 17.

958 We formalize all these properties in the well-formedness statement about the algorithmic sub-

959 typing relation below.

960

961 LEMMA 6.1 (ALGORITHMIC SUBTYPING IS W.F.).

- 962 • If  $\Theta \vdash P \leq^+ Q \dashv \Theta'$ ,  $\Theta \text{ ctx}$ ,  $P \text{ ground}$ , and  $[\Theta]Q = Q$ , then  $\Theta' \text{ ctx}$ ,  $\Theta \longrightarrow \Theta'$ , and  $[\Theta']Q$
- 963 *ground*.
- 964 • If  $\Theta \vdash N \leq^- M \dashv \Theta'$ ,  $\Theta \text{ ctx}$ ,  $M \text{ ground}$ , and  $[\Theta]N = N$ , then  $\Theta' \text{ ctx}$ ,  $\Theta \longrightarrow \Theta'$ , and  $[\Theta']N$
- 965 *ground*.

966 The premises of each case encode our expectations about the inputs to the subtyping algorithm:

967 the input context should be well-formed, the type on the ground side should be ground, and the

968 type on the non-ground side should not contain any existential variables that have already been

969 solved. We prove that when given these inputs, the subtyping algorithm produces an output context

970 that is well-formed, is compatible with the input context, and that solves all the existential variables

971 that appear in the input types.

972

973 **6.1.2 Typing.** The well-formedness statement for typing is very similar to the one for subtyping.

974 However stating well-formedness for the spine judgment is more complex because this judgment

975 (specifically the *Aspinetypeabsin* rule) can introduce new existential variables that do not appear

976 within its input context.

977 To tackle this, we introduce a new notion of context extension, *weak context extension*, in

978 Figure 18. This is identical to normal context extension, except it has two additional rules to

979 permit adding new existential variables: *Wcnewunsolvedguess* lets it add an unsolved variable,

981	$ A _{\text{NQ}}$ The size of a type $A$ , ignoring quantification		
982			
983	$ \alpha _{\text{NQ}} = 1$	$ \downarrow N _{\text{NQ}} =  N _{\text{NQ}} + 1$	$ P \rightarrow N _{\text{NQ}} =  P _{\text{NQ}} +  N _{\text{NQ}} + 1$
984	$ \hat{\alpha} _{\text{NQ}} = 1$	$ \forall \alpha. N _{\text{NQ}} =  N _{\text{NQ}}$	$ \uparrow P _{\text{NQ}} =  P _{\text{NQ}} + 1$
985			
986	$\text{NPQ}(A)$ The number of prenex quantifiers in a type $A$		
987			
988	$\text{NPQ}(\alpha) = 0$	$\text{NPQ}(\downarrow N) = 0$	$\text{NPQ}(P \rightarrow N) = 0$
989	$\text{NPQ}(\hat{\alpha}) = 0$	$\text{NPQ}(\forall \alpha. N) = 1 + \text{NPQ}(N)$	$\text{NPQ}(\uparrow P) = 0$
990			

Fig. 20. Definitions used in the decidability metric for Implicit Polarized F

and `Wcnewsolvedguess` lets it add a solved variable. We also extend in Figure 19 our notion of the well-formedness of algorithmic contexts to typing environments in the obvious way, with `Ewfvar` paralleling `Cwfsolvedguess`.

With this weaker notion of context extension we can state well-formedness of typing as follows:

LEMMA 6.2 (ALGORITHMIC TYPING IS W.F.). *Given a typing context  $\Theta$  and typing environment  $\Gamma$  such that  $\Theta$  ctx and  $\Theta \vdash \Gamma$  env:*

- *If  $\Theta; \Gamma \vdash v : P \dashv \Theta'$ , then  $\Theta'$  ctx,  $\Theta \longrightarrow \Theta'$ ,  $\Theta' \vdash P$  type<sup>+</sup>, and  $P$  ground.*
- *If  $\Theta; \Gamma \vdash t : N \dashv \Theta'$ , then  $\Theta'$  ctx,  $\Theta \longrightarrow \Theta'$ ,  $\Theta' \vdash N$  type<sup>-</sup>, and  $N$  ground.*
- *If  $\Theta; \Gamma \vdash s : N \gg M \dashv \Theta'$ ,  $\Theta \vdash N$  type<sup>-</sup>, and  $[\Theta]N = N$ , then  $\Theta'$  ctx,  $\Theta \Longrightarrow \Theta'$ ,  $\Theta' \vdash M$  type<sup>-</sup>,  $[\Theta']M = M$ , and  $\text{FEV}(M) \subseteq \text{FEV}(N) \cup (\text{FEV}(\Theta') \setminus \text{FEV}(\Theta))$ .*

In addition to the standard postconditions, we also prove in the spine judgment case that the free existential variables in the output type  $M$  either come from the input type  $N$  or were added to the context by `Aspinetypeabsin` while instantiating universal quantifiers.

## 6.2 Determinism and decidability

Since our system is syntax-directed, determinism of algorithmic typing follows from straightforward rule inductions on the typing and subtyping rules. However, decidability is more intricate.

**6.2.1 Subtyping.** Our goal in proving the decidability of subtyping is finding a metric that decreases from the conclusion to each premise of the algorithmic rules. However the obvious metrics do not work. For instance, a metric based on the size of the type will not work because types can get bigger when we instantiate type variables. Predicative systems might use the lexicographic ordering of the number of prenex quantifiers followed by the size of the type as a metric. But while instantiations will not increase the number of quantifiers in predicative systems, they can do in impredicative systems.

In order to find a metric that works, we go back to the declarative system. Our declarative subtyping relation has the property that (using a size metric  $|\_ |_{\text{NQ}}$  for terms that ignores quantifiers), a well-formed<sup>4</sup>  $\Theta \vdash P \leq^+ Q$  judgment implies that  $|Q|_{\text{NQ}} \leq |P|_{\text{NQ}}$ , and similarly a well-formed  $\Theta \vdash N \leq^- M$  judgment implies that  $|N|_{\text{NQ}} \leq |M|_{\text{NQ}}$ . The intuition behind this property is that subtyping reflects the specialization order, and when a polymorphic type is instantiated, occurrences of type variables  $\alpha$  get replaced with instantiations  $P$ . As a result, the size of a subtype has to be smaller than its supertypes. In ordinary System F, this intuition is actually false due to the contravariance of the function type. However, in Implicit Polarized F, universal quantification is

<sup>4</sup>One for which the well-formedness assumptions in Lemma 6.1 (Algorithmic subtyping is w.f.) hold.

a *negative type* of the form  $\forall\alpha. N$ , and ranges over *positive types*. Since positive type variables  $\alpha$  are occurring within a negative type  $N$ , each occurrence of  $\alpha$  must be underneath a shift, which is invariant in our system. As a result, the intuitive property is true.

Now, note that our algorithm always solves all of the existential problems in a well-formed subtyping problem. In particular, given an algorithmic derivation  $\Theta \vdash P \leq^+ Q \dashv \Theta'$ , we know that  $[\Theta']Q$  will be ground. So we therefore expect  $[\Theta']Q$  to be the same size as it would be in a declarative derivation. This gives us a hint about what to try for the metric.

Since the size of the ground side (e.g.  $P$ ) bounds the size of the completed non-ground side (e.g.  $[\Theta']Q$ ), we will incorporate the size of the ground side into our metric. This size decreases between the conclusions and the premises of all of the rules apart from  $\leq^\pm\text{Aforalll}$  and  $\leq^\pm\text{Aforallr}$ , where it stays the same.  $\leq^\pm\text{Aforalll}$  and  $\leq^\pm\text{Aforallr}$  both however pick off a prenex quantifier, so we can take the lexicographic ordering of this size followed by the total number of prenex quantifiers in the subtyping judgment. This is  $(|P|_{\text{NQ}}, \text{NPQ}(P) + \text{NPQ}(Q))$  for positive judgments  $\Theta \vdash P \leq^+ Q \dashv \Theta'$  and  $(|M|_{\text{NQ}}, \text{NPQ}(M) + \text{NPQ}(N))$  for negative judgments  $\Theta \vdash N \leq^- M \dashv \Theta'$ . We give formal definitions of  $|\_|\_{\text{NQ}}$  and  $\text{NPQ}(\_)$  in Figure 20.

Note that this is somewhat backwards from a typical metric for predicative systems where we count quantifiers first, then the size.

We prove that this metric assigns a total ordering to derivations in the algorithmic subtyping system. In each rule with multiple hypotheses, we invoke the bounding property described above, which we formalize in the following lemma:

LEMMA 6.3 (COMPLETED NON-GROUND SIZE BOUNDED BY GROUND SIZE).

- If  $\Theta \vdash P \leq^+ Q \dashv \Theta'$ ,  $\Theta \text{ ctx}$ ,  $P$  ground, and  $[\Theta]Q = Q$ , then  $||[\Theta']Q|_{\text{NQ}} \leq |P|_{\text{NQ}}$ .
- If  $\Theta \vdash N \leq^- M \dashv \Theta'$ ,  $\Theta \text{ ctx}$ ,  $M$  ground, and  $[\Theta]N = N$ , then  $||[\Theta']N|_{\text{NQ}} \leq |M|_{\text{NQ}}$ .

As a result of our key tactic for deriving this metric coming from the declarative system, it turns out that this metric assigns total orderings to derivations both in the algorithmic and the declarative subtyping systems. This allows us to reuse it in both the proof of completeness and the proof of transitivity, where in both cases using the height of the derivation as an induction metric is too weak.

**6.2.2 Typing.** The decidability of algorithmic typing ends up being relatively straightforward. Almost every rule decreases in the standard structural notions of size  $|\_|\_$  on terms and spines. The  $\text{Dspinetypeabs}$  rule however preserves the size of the spine involved, so taking this size alone is insufficient. Therefore when comparing two spine judgments, we take the lexicographic ordering of  $(|s|, \text{NPQ}(N))$ , where  $s$  is the spine and  $N$  is the input type. We prove that this metric assigns a total ordering to algorithmic typing derivations, thus demonstrating that our typing algorithm is decidable. Alongside our proof, we include an exact statement of this metric in the appendix.

Note that, as with subtyping, this metric assigns a total ordering not only to algorithmic typing derivations but also to declarative typing derivations. Therefore we reuse this metric in the proofs of soundness, completeness, and the behavior of isomorphic types. In each of these cases, using the height of the derivation as an induction metric is too weak.

## 7 SOUNDNESS AND COMPLETENESS

We have now set out declarative subtyping and typing systems for Implicit Polarized F as well as algorithms to implement them. In this section, we will demonstrate that the algorithms are sound and complete with respect to their declarative counterparts. Proofs of all of these results are in the appendix.

## 7.1 Subtyping

7.1.1 *Soundness.* Consider the algorithmic subtyping judgment  $\Theta \vdash P \leq^+ Q \dashv \Theta'$ . The non-ground input  $Q$  might contain some unsolved existential variables, which are solved in  $\Theta'$ . For soundness, we want to say that no matter what solutions are ever made for these existential variables, there will be a corresponding declarative derivation. To state this, we introduce a notion of a *complete context*. These are algorithmic contexts which contain no unsolved existential variables:

$$\text{Complete contexts } \Omega ::= \cdot \mid \Omega, \alpha \mid \Omega, \hat{\alpha} = P$$

We can now use the context extension judgment  $\Theta' \longrightarrow \Omega$  to state that  $\Omega$  is a complete context that contains solutions to all the unsolved existential variables in  $\Theta'$ . Applying this complete context to  $Q$  gives us our hoped-for declarative subtyping judgment  $\|\Theta\| \vdash P \leq^+ [\Omega]Q$ . (The  $\|\Theta\|$  judgment (defined in Figure 17) drops the existential variables to create a declarative context.)

This gives us the following statement of soundness, which we prove in the appendix:

**THEOREM 7.1 (SOUNDNESS OF ALGORITHMIC SUBTYPING).** *Given a well-formed algorithmic context  $\Theta$  and a well-formed complete context  $\Omega$ :*

- *If  $\Theta \vdash P \leq^+ Q \dashv \Theta'$ ,  $\Theta' \longrightarrow \Omega$ ,  $P$  ground,  $[\Theta]Q = Q$ ,  $\Theta \vdash P \text{ type}^+$ , and  $\Theta \vdash Q \text{ type}^+$ , then  $\|\Theta\| \vdash P \leq^+ [\Omega]Q$ .*
- *If  $\Theta \vdash N \leq^- M \dashv \Theta'$ ,  $\Theta' \longrightarrow \Omega$ ,  $M$  ground,  $[\Theta]N = N$ ,  $\Theta \vdash N \text{ type}^-$ , and  $\Theta \vdash M \text{ type}^-$ , then  $\|\Theta\| \vdash [\Omega]N \leq^- M$ .*

The side conditions for soundness are similar to those of well-formedness (Lemma 6.1), except we also require the complete context and the types in the algorithmic judgment to be well-formed.

7.1.2 *Completeness.* Completeness is effectively the reverse of soundness. For each positive declarative judgment  $\|\Theta\| \vdash P \leq^+ [\Omega]Q$  where  $\Theta \longrightarrow \Omega$  we want to find a corresponding algorithmic derivation  $\Theta \vdash P \leq^+ Q \dashv \Theta'$ , and likewise for negative judgments. No matter what declarative solutions we had for a solved type, the algorithmic system infers compatible solutions.

Since context extension  $\Theta \longrightarrow \Omega$  now appears in the premise, we need to strengthen our induction hypothesis to prove completeness for rules with multiple premises. To do this, we prove that the algorithm's output context  $\Theta'$  is also compatible with  $\Omega$ , i.e.  $\Theta' \longrightarrow \Omega$ . This gives us the following statement of completeness, in which the side conditions are identical to soundness.

**THEOREM 7.2 (COMPLETENESS OF ALGORITHMIC SUBTYPING).** *If  $\Theta$  ctx,  $\Theta \longrightarrow \Omega$ , and  $\Omega$  ctx, then:*

- *If  $\|\Theta\| \vdash P \leq^+ [\Omega]Q$ ,  $\Theta \vdash P \text{ type}^+$ ,  $\Theta \vdash Q \text{ type}^+$ ,  $P$  ground, and  $[\Theta]Q = Q$ , then  $\exists \Theta'$  such that  $\Theta \vdash P \leq^+ Q \dashv \Theta'$  and  $\Theta' \longrightarrow \Omega$ .*
- *If  $\|\Theta\| \vdash [\Omega]N \leq^- M$ ,  $\Theta \vdash M \text{ type}^-$ ,  $\Theta \vdash N \text{ type}^-$ ,  $M$  ground, and  $[\Theta]N = N$ , then  $\exists \Theta'$  such that  $\Theta \vdash N \leq^- M \dashv \Theta'$  and  $\Theta' \longrightarrow \Omega$ .*

## 7.2 Typing

Due to the unusual implication within `Dunambiguouslet`, soundness and completeness of typing are mutually recursive in our system. In soundness we use completeness while proving this implication, and in completeness we use soundness while unpacking the implication. We justify each of these uses by only applying either soundness or completeness to a judgment involving a subterm.

7.2.1 *Soundness.* Soundness of typing is formulated in the same way as soundness of subtyping. We introducing a complete context  $\Omega$  which the output context extends to. For the spine judgment we allow the completed output type  $[\Omega]M$  to be isomorphic to the declarative output type  $M'$ .

1128 THEOREM 7.3 (SOUNDNESS OF ALGORITHMIC TYPING). *If  $\Theta$  ctx,  $\Theta \vdash \Gamma$  env,  $\Theta' \longrightarrow \Omega$ , and  $\Omega$  ctx,*  
 1129 *then:*

- 1130 • *If  $\Theta; \Gamma \vdash v : P \dashv \Theta'$ , then  $\|\Theta\|; \Gamma \vdash v : [\Omega]P$ .*
- 1131 • *If  $\Theta; \Gamma \vdash t : N \dashv \Theta'$ , then  $\|\Theta\|; \Gamma \vdash t : [\Omega]N$ .*
- 1132 • *If  $\Theta; \Gamma \vdash s : N \gg M \dashv \Theta'$ ,  $\Theta \vdash N$  type<sup>-</sup>, and  $[\Theta]N = N$ , then  $\exists M'$  such that  $\|\Theta\| \vdash [\Omega]M \cong^-$*   
 1133  *$M'$  and  $\|\Theta\|; \Gamma \vdash s : [\Omega]N \gg M'$ .*

### 1135 7.3 Completeness

1136 The main challenge in stating completeness of typing is taking into account the fact that spine  
 1137 judgments can introduce new existential variables. Practically, the introduction of new existential  
 1138 variables means that the output context of the algorithmic spine judgment  $\Theta'$  will not necessarily  
 1139 extend to the complete context  $\Omega$ . To deal with this, we make use of weak context extension  
 1140 and instead require there to exist some new complete context  $\Omega'$  containing the new existential  
 1141 variables. The original complete context  $\Omega$  should then weakly extend to  $\Omega'$ , and the output context  
 1142 of the spine judgment  $\Theta'$  should strongly extend to  $\Omega'$ .

1143 This gives us the following statement of the completeness of typing:  
 1144

1145 THEOREM 7.4 (COMPLETENESS OF ALGORITHMIC TYPING). *If  $\Theta$  ctx,  $\Theta \vdash \Gamma$  env,  $\Theta \longrightarrow \Omega$ , and  $\Omega$  ctx,*  
 1146 *then:*

- 1147 • *If  $\|\Theta\|; \Gamma \vdash v : P$  then  $\exists \Theta'$  such that  $\Theta; \Gamma \vdash v : P \dashv \Theta'$  and  $\Theta' \longrightarrow \Omega$ .*
- 1148 • *If  $\|\Theta\|; \Gamma \vdash t : N$  then  $\exists \Theta'$  such that  $\Theta; \Gamma \vdash t : N \dashv \Theta'$  and  $\Theta' \longrightarrow \Omega$ .*
- 1149 • *If  $\|\Theta\|; \Gamma \vdash s : [\Omega]N \gg M$ ,  $\Theta \vdash N$  type<sup>-</sup>, and  $[\Theta]N = N$ , then  $\exists \Theta', \Omega'$  and  $M'$  such that*  
 1150  *$\Theta; \Gamma \vdash s : N \gg M' \dashv \Theta'$ ,  $\Omega \Longrightarrow \Omega'$ ,  $\Theta' \longrightarrow \Omega'$ ,  $\|\Theta\| \vdash [\Omega']M' \cong^- M$ ,  $[\Theta']M' = M'$ , and*  
 1151  *$\Omega'$  ctx.*

## 1153 8 RELATED WORK

1154 There has been considerable research into working around the undecidability of type inference  
 1155 for System F. Broadly, it falls into three main categories: enriching the language of types to make  
 1156 inference possible, restricting the subtype relation, and using heuristics to knock off the easy cases.  
 1157

1158 *Enriching the Type Language.* The “gold standard” for System F type inference is the  $\text{ML}^F$  system  
 1159 of Botlan and Rémy [1]. They observe that type inference for System F is complicated by the lack of  
 1160 principal types. To create principal types, they extend the type language of System F with bounded  
 1161 type constraints, and show that this admits a type inference algorithm with a simple and powerful  
 1162 specification: only arguments to functions used at multiple parametric types need annotation.

1163 Unfortunately,  $\text{ML}^F$  is notoriously difficult to implement (see Rémy and Jakobowski [15] for the  
 1164 state-of-the-art), and there have been multiple attempts to find simpler subsystems. FPH (Vytiniotis  
 1165 et al. [24]) attempted to simplify  $\text{ML}^F$  by limiting the user-visible language of types to the standard  
 1166 System F types and only using  $\text{ML}^F$ -style constrained types internally within the type checker.  
 1167 Another attempt to simplify  $\text{ML}^F$  was HML (Leijen [8]). Unlike FPH, HML exposes part of the  $\text{ML}^F$   
 1168 machinery (flexible types) to the user, which changes the language of types but also types more  
 1169 programs. Both FPH and HML require annotations only for polymorphic arguments, however  
 1170 unfortunately, both approaches proved too intricate to adapt to GHC.

1171 *Restricting the subtype relation.* The most widely used approach for System F type inference  
 1172 simply abandons impredicativity. This line of work was originated by Odersky and Läufer [11], who  
 1173 proposed restricting type instantiation in the subtype relation to monotypes. This made subtyping  
 1174 decidable, and forms the basis for type inference in Haskell [13]. Dunfield and Krishnaswami [5]  
 1175 give a simple variant of this algorithm based on bidirectional typechecking.  
 1176



1177 This style of inference omits impredicativity from the subtype relation altogether. *Boxy types* [23]  
 1178 combine predicative subtyping with a generalization of bidirectional typechecking to support  
 1179 impredicativity. Vytiniotis et al. introduce a new type system feature, *boxes*, which merge the  
 1180 synthesis and checking judgments from bidirectional typechecking into marks on types (i.e. the  
 1181 boxes) which indicate whether part of a type came from inference or annotation. (This is somewhat  
 1182 reminiscent of “colored local type inference” [12].) Unfortunately, boxy types lack a clear non-  
 1183 algorithmic specification of when type annotations are required.

1184 HMF [7] introduced another approach to restricting the specialization relation. It restricted  
 1185 subtyping for the function type constructor — instead of being contravariant, function types  
 1186 were invariant. This meant that inference could be done with only a modest modification to the  
 1187 Damas-Milner algorithm. Our work retains function contravariance, and only imposes invariance  
 1188 at shifts.

1189 *Heuristic Approaches.* Many of the type inference problems which arise in practice are actually  
 1190 easy to solve. Cardelli [2] invented one of the oldest such techniques while constructing a type  
 1191 inference algorithm for a language with F-bounded impredicative polymorphism. Rather than  
 1192 doing anything difficult, Cardelli’s algorithm simply instantiated each quantifier with the first type  
 1193 constraint it ran into. Pierce and Turner [14] formalized Cardelli’s approach, and noticed that it did  
 1194 not need unification to implement. To make their algorithm work, they needed to use a seemingly  
 1195 ad-hoc representation of types, with a single type constructor that was simultaneously an  $n$ -ary  
 1196 small and big lambda. Furthermore, while they proved that inference was sound, they did not offer  
 1197 a declarative characterization of inference.

1198 Serrano et al. [19] recently revisited the idea of controlling inference with heuristics with their  
 1199 system GI. They restrict impredicative instantiation to *guarded instantiations*, which are (roughly  
 1200 speaking) the cases when the type variable being instantiated is underneath a type constructor.  
 1201 This restriction is automatically achieved in our setting via the presence of shifts, which suggests  
 1202 that this syntactic restriction actually arises for deeper type-theoretic reasons. In follow up work,  
 1203 Serrano et al. [18] further simplify their approach, making some dramatic simplifications to the type  
 1204 theory (e.g., giving up function type contravariance) in order to achieve a simpler implementation.

1205 *Conclusions.* Many researchers have noticed that type inference algorithms benefit from being able  
 1206 to look at the entire argument list to a function. From the perspective of plain lambda calculus, this  
 1207 looks ad-hoc and non-compositional. For example, HMF was originally described in two variants,  
 1208 one using argument lists and one without, and only the weaker algorithm without argument lists  
 1209 has a correctness proof. However, from the perspective of polarized type theory, argument lists are  
 1210 entirely type-theoretically natural: they mark the change of a polarity boundary!

1211 This explains why Pierce and Turner’s use of “jumbo” function types which combine multiple  
 1212 quantifiers and arguments makes sense: the merged connectives are all negative, with no interposed  
 1213 shifts. Making shifts explicit means that small connectives can have the same effect, which makes  
 1214 it possible to give a clear specification for the system.

1215 We have also seen that many algorithms omit function contravariance from the specialization  
 1216 order to support impredicative inference. Our work clarifies that contravariance *per se* is not  
 1217 problematic, but rather that the benefits for inference arise from controlling the crossing of polarity  
 1218 boundaries. This again permits a simpler and more regular specification of subtyping.

1219

## 1220 REFERENCES

- 1221 [1] Didier Le Botlan and Didier Rémy. 2003. MLF Raising ML to the Power of System F. In *ICFP '03*. ACM Press, Uppsala,  
 1222 Sweden, 52–63.
- 1223 [2] Luca Cardelli. 1993. *An Implementation of F<:*. Technical Report. Systems Research Center, Digital Equipment  
 1224 Corporation.

1225

- 1226 [3] I. Cervesato. 2003. A Linear Spine Calculus. *Journal of Logic and Computation* 13, 5 (Oct. 2003), 639–688. <https://doi.org/10.1093/logcom/13.5.639>
- 1227
- 1228 [4] Jacek Chrzaszcz. 1998. Polymorphic Subtyping without Distributivity. In *Proceedings of the 23rd International Symposium on Mathematical Foundations of Computer Science (MFCS '98)*. Springer-Verlag, Berlin, Heidelberg, 346–355.
- 1229
- 1230 [5] Jana Dunfield and Neel Krishnaswami. 2013. Complete and Easy Bidirectional Typechecking for Higher-Rank Polymorphism. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming - ICFP '13*. ACM Press, Boston, Massachusetts, USA, 429. <https://doi.org/10.1145/2500365.2500582>
- 1231
- 1232 [6] Jean-Yves Girard. 1971. Une Extension de L'Interpretation de Gödel à L'Analyse, et Son Application à L'Elimination Des Coupures Dans L'Analyse et La Theorie Des Types. In *Proceedings of the Second Scandinavian Logic Symposium*, J.E. Fenstad (Ed.). Studies in Logic and the Foundations of Mathematics, Vol. 63. Elsevier, North Holland, 63–92. [https://doi.org/10.1016/S0049-237X\(08\)70843-7](https://doi.org/10.1016/S0049-237X(08)70843-7)
- 1233
- 1234 [7] Daan Leijen. 2008. *HMF: Simple Type Inference for First-Class Polymorphism*. Technical Report MSR-TR-2008-65. Microsoft Research. 15 pages.
- 1235
- 1236 [8] Daan Leijen. 2008. *Robust Type Inference for First-Class Polymorphism*. Technical Report MSR-TR-2008-55. Microsoft Research. 10 pages.
- 1237
- 1238 [9] Paul Blain Levy. 2006. Call-by-Push-Value: Decomposing Call-by-Value and Call-by-Name. *Higher-Order and Symbolic Computation* 19, 4 (Dec. 2006), 377–414. <https://doi.org/10.1007/s10990-006-0480-6>
- 1239
- 1240 [10] Robin Milner. 1978. A Theory of Type Polymorphism in Programming. *J. Comput. System Sci.* 17, 3 (Dec. 1978), 348–375. [https://doi.org/10.1016/0022-0000\(78\)90014-4](https://doi.org/10.1016/0022-0000(78)90014-4)
- 1241
- 1242 [11] Martin Odersky and Konstantin Läuffer. 1996. Putting Type Annotations to Work. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages - POPL '96*. ACM Press, St. Petersburg Beach, Florida, United States, 54–67. <https://doi.org/10.1145/237721.237729>
- 1243
- 1244 [12] Martin Odersky, Christoph Zenger, and Matthias Zenger. 2001. Colored Local Type Inference. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages - POPL '01*. ACM Press, London, United Kingdom, 41–53. <https://doi.org/10.1145/360204.360207>
- 1245
- 1246 [13] Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. 2007. Practical Type Inference for Arbitrary-Rank Types. *Journal of Functional Programming* 17, 01 (Jan. 2007), 1. <https://doi.org/10.1017/S0956796806006034>
- 1247
- 1248 [14] Benjamin Pierce and David Turner. 2000. Local Type Inference. *ACM Transactions on Programming Languages and Systems* 22, 1 (Jan. 2000), 1–44. <https://doi.org/10.1145/345099.345100>
- 1249
- 1250 [15] Didier Rémy and Boris Yakobowski. 2008. From ML to MLF: Graphic Type Constraints with Efficient Type Inference. In *ICFP*. ACM, Victoria, BC, Canada, 63–74.
- 1251
- 1252 [16] John C Reynolds. 1974. Towards a Theory of Type Structure. In *Programming Symposium*. Springer, Springer, Paris, France, 408–425.
- 1253
- 1254 [17] John C. Reynolds. 1983. Types, Abstraction and Parametric Polymorphism. In *IFIP Congress*. North-Holland/IFIP, Paris, France, 513–523.
- 1255
- 1256 [18] Alejandro Serrano, Jurriaan Hage, Simon Peyton Jones, and Dimitrios Vytiniotis. 2020. A Quick Look at Impredicativity. *Proceedings of the ACM on Programming Languages* 4, ICFP (Aug. 2020), 1–29. <https://doi.org/10.1145/3408971>
- 1257
- 1258 [19] Alejandro Serrano, Jurriaan Hage, Dimitrios Vytiniotis, and Simon Peyton Jones. 2018. Guarded Impredicative Polymorphism. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation - PLDI 2018*. ACM Press, Philadelphia, PA, USA, 783–796. <https://doi.org/10.1145/3192366.3192389>
- 1259
- 1260 [20] Arnaud Spiwack. 2014. A Dissection of L. (2014).
- 1261
- 1262 [21] Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones, and Kevin Donnelly. 2007. System F with Type Equality Coercions. In *Proceedings of the 2007 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation (TLDI '07)*. Association for Computing Machinery, New York, NY, USA, 53–66. <https://doi.org/10.1145/1190315.1190324>
- 1263
- 1264 [22] Jerzy Tiuryn and Pawel Urzyczyn. 1996. The Subtyping Problem for Second-Order Types Is Undecidable. In *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science (LICS '96)*. IEEE Computer Society, USA, 74.
- 1265
- 1266 [23] Dimitrios Vytiniotis, Stephanie Weirich, and Simon Peyton Jones. 2006. Boxy Types: Inference for Higher-Rank Types and Impredicativity. In *ICFP '06*. ACM Press, Portland, Oregon, USA, 251–262.
- 1267
- 1268 [24] Dimitrios Vytiniotis, Stephanie Weirich, and Simon Peyton Jones. 2008. FPH: First-class Polymorphism for Haskell. In *ICFP '08*. ACM Press, Victoria, BC, Canada, 295–306.
- 1269
- 1270 [25] Philip Wadler. 1989. Theorems for Free!. In *Proceedings of the Fourth International Conference on Functional Programming Languages and Computer Architecture - FPCA '89*. ACM Press, Imperial College, London, United Kingdom, 347–359. <https://doi.org/10.1145/99370.99404>
- 1271
- 1272
- 1273
- 1274