



An Existential Crisis Resolved

Type Inference for First-Class Existential Types

RICHARD A. EISENBERG, Tweag, France

GUILLAUME DUBOC, ENS Lyon, France and Tweag, France

STEPHANIE WEIRICH, University of Pennsylvania, USA

DANIEL LEE, University of Pennsylvania, USA

Despite the great success of inferring and programming with universal types, their dual—existential types—are much harder to work with. Existential types are useful in building abstract types, working with indexed types, and providing first-class support for refinement types. This paper, set in the context of Haskell, presents a bidirectional type-inference algorithm that infers where to introduce and eliminate existentials without any annotations in terms, along with an explicitly typed, type-safe core language usable as a compilation target. This approach is backward compatible. The key ingredient is to use *strong* existentials, which support (lazily) projecting out the encapsulated data, not weak existentials accessible only by pattern-matching.

CCS Concepts: • **Software and its engineering** → **Abstract data types**; *Functional languages*.

Additional Key Words and Phrases: existential types, type inference, Haskell

ACM Reference Format:

Richard A. Eisenberg, Guillaume Duboc, Stephanie Weirich, and Daniel Lee. 2021. An Existential Crisis Resolved: Type Inference for First-Class Existential Types. *Proc. ACM Program. Lang.* 5, ICFP, Article 64 (August 2021), 29 pages. <https://doi.org/10.1145/3473569>

1 INTRODUCTION

Parametric polymorphism through the use of universally quantified type variables is pervasive in functional programming. Given its overloaded numbers, a beginning Haskell programmer literally cannot ask for the type of $1 + 1$ without seeing a universally quantified type variable.

However, universal quantification has a dual: existentials. While universals claim the spotlight, with support for automatic elimination (that is, instantiation) in all non-toy typed functional languages we know and automatic introduction (frequently, **let**-generalization) in some, existentials are underserved and impoverished. In every functional language we know, both elimination and introduction must be done explicitly every time, and languages otherwise renowned for their type inference—such as Haskell—require that users define a new top-level datatype for every existential.

While not as widely useful as universals, existential quantification comes up frequently in richly typed programming. Further examples are in Section 2, but consider writing a *dropWhile* function on everyone’s favorite example datatype, the length-indexed vector:

Authors’ addresses: Richard A. Eisenberg, rae@richarde.dev, Tweag, 207 Rue de Bercy, 75012, Paris, France; Guillaume Duboc, guillaume.duboc@ens-lyon.fr, ENS Lyon, Lyon, France and Tweag, France; Stephanie Weirich, sweirich@seas.upenn.edu, University of Pennsylvania, 3330 Walnut St., Philadelphia, PA, 19104, USA; Daniel Lee, laniel@seas.upenn.edu, University of Pennsylvania, 3330 Walnut St., Philadelphia, PA, 19104, USA.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/8-ART64

<https://doi.org/10.1145/3473569>

-- *dropWhile predicate vec* drops the longest prefix of *vec* such that all elements in the prefix satisfy *predicate*. In this type, *n* is the vector's length, while *a* is the type of elements.

dropWhile :: (*a* → *Bool*) → *Vec n a* → *Vec ??? a*

How can we fill in the question marks? Without knowing the contents of the vector and the predicate we are passing, we cannot know the length of the output. Furthermore, returning an ordinary, unindexed list would require copying a suffix of the input vector, an unacceptable performance degradation.

Existentials come to our rescue: *dropWhile* :: (*a* → *Bool*) → *Vec n a* → ∃*m. Vec m a*. Though this example can be written today in a number of languages, all require annotations in terms both to pack (introduce) the existential and unpack (eliminate) it through the application or pattern-matching of a data constructor.

This paper describes a type-inference algorithm that supports implicit introduction and elimination of existentials, with a concrete setting in Haskell. We offer the following contributions:

- Section 4 presents our type-inference algorithm, the primary contribution of this paper. The algorithm is a small extension to an algorithm that accepts a Hindley-Milner language; our language, \mathbb{X} , is thus a superset of Hindley-Milner (Theorem 7.3). In addition, it supports several *stability properties* [Bottu and Eisenberg 2021]; a language is *stable* if small, seemingly innocuous changes to the input program (such as **let**-inlining) do not cause a change in the type or acceptability of a program (Theorems 7.4–7.6). Our algorithm is easily integrable with the latest inference algorithm [Serrano et al. 2020] in the Glasgow Haskell Compiler (GHC) (Section 8).
- Section 5 presents a core language based on System F, \mathbb{FX} , that is a suitable target of compilation (Section 6) for \mathbb{X} . We prove \mathbb{FX} is type-safe (Theorems 5.1 and 5.2) and supports type erasure (Theorem 5.3). It is designed in a way that is compatible with the existing System FC [Sulzmann et al. 2007] language used internally within GHC. All programs accepted by our algorithm elaborate to well-typed programs in \mathbb{FX} (Theorem 7.1). In addition, elaboration preserves the semantics of the source program, as we can observe by examining the result of type erasure (Theorem 7.2).

We normally desire type-inference algorithms to come with a declarative specification, where automatic introduction and elimination of quantifiers can happen anywhere, in the style of the Hindley-Milner type system [Hindley 1969; Milner 1978]. These specifications come alongside syntax-directed algorithms that are sound and complete with respect to the specification [Clément et al. 1986; Damas and Milner 1982]. However, we do not believe such a system is possible with existentials; while negative results are hard to prove conclusively, we lay out our arguments against this approach in Section 9.1. Instead, we present just our algorithm, though we avoid the complication and distraction of unification variables by allowing our algorithm to non-deterministically guess monotypes τ in the style of a declarative specification.

There is a good deal of literature in this area; much of it is focused on module systems, which often wish to hide the nature of a type using an existential package. We review some important prior work in Section 10.

The concrete examples in this paper are set in Haskell, but the fundamental ideas in our inference algorithm are fully portable to other settings, including in languages without **let**-generalization.

2 MOTIVATION AND EXAMPLES

Though not as prevalent as examples showing the benefits of universal polymorphism, easy existential polymorphism smooths out some of the wrinkles currently inherent in programming with indexed types such as GADTs [Xi et al. 2003].

$$\begin{array}{ll}
\text{filter} :: (a \rightarrow \text{Bool}) \rightarrow \text{Vec } n \ a \rightarrow \text{ExVec } a & \text{filter} :: (a \rightarrow \text{Bool}) \rightarrow \text{Vec } n \ a \rightarrow \exists m. \text{Vec } m \ a \\
\text{filter } _ \text{Nil} & = \text{MkEV Nil} & \text{filter } _ \text{Nil} & = \text{Nil} \\
\text{filter } p \ (x \text{:>} \ xs) \mid p \ x & , \text{MkEV } v \leftarrow \text{filter } p \ xs & \text{filter } p \ (x \text{:>} \ xs) \mid p \ x & = x \text{:>} \text{filter } p \ xs \\
& = \text{MkEV } (x \text{:>} \ v) & & \mid \text{otherwise} = \text{filter } p \ xs \\
& \mid \text{otherwise} = \text{filter } p \ xs & & \\
\text{(a)} & & \text{(b)} &
\end{array}$$

Fig. 1. Implementations of *filter* over vectors (a) in today's Haskell, and (b) with our extensions

2.1 Unknown Output Indices

We first return to the example from the introduction, writing an operation that drops an indeterminate number of elements from a length-indexed vector:

```

data Nat = Zero | Succ Nat
type Vec :: Nat → Type → Type -- -XStandaloneKindSignatures, new in GHC 8.10
data Vec n a where
  Nil :: Vec Zero a
  (:>) :: a → Vec n a → Vec (Succ n) a
infixr 5 :>

```

In today's Haskell, the way to write *dropWhile* over vectors is like this:

```

type ExVec :: Type → Type
data ExVec a where
  MkEV :: ∀(n :: Nat) (a :: Type). Vec n a → ExVec a
dropWhile :: (a → Bool) → Vec n a → ExVec a
dropWhile _ Nil = MkEV Nil
dropWhile p (x :> xs) | p x = dropWhile p xs
                       | otherwise = MkEV (x :> xs)

```

However, with our inference of existential introduction and elimination, we can simplify to this:

```

dropWhile :: (a → Bool) → Vec n a → ∃m. Vec m a
dropWhile _ Nil = Nil
dropWhile p (x :> xs) | p x = dropWhile p xs
                       | otherwise = x :> xs

```

There are two key differences: we no longer need to define the *ExVec* type, instead using $\exists m. \text{Vec } m \ a$; and we can omit any notion of packing in the body of *dropWhile*. Similarly, clients of *dropWhile* would not need to unpack the result, allowing the result of *dropWhile* to be immediately consumed by a *map*, for example.

2.2 Increased Laziness

Another function that produces an output of indeterminate length is *filter*. It is enlightening to compare the implementation of *filter* using today's existentials and the version possible with our new ideas; see Figure 1.

What if **unpack** were simply lazy? The problem is that this is not simple! A straightforward typed operational semantics would not suffice, because there is no way to, say, reduce an **unpack** into a substitution (the way we would handle a lazy **let**). We could imagine an untyped operational semantics that did not require **unpack** to evaluate the existential package, binding its variable with a lazy binding. Without types, though, we would be unable to prove safety. In order to keep a typed operational semantics with a lazy **unpack**, we must model a set of heap bindings and an evaluation stack in our semantics. While this is possible, such an operational semantics is unsuitable for a (dependently typed) language where we also might wish to evaluate in types, which is our eventual goal for Haskell. The claim here is not that a lazy **unpack** is impossible, but that it is not obviously superior to the approach we advocate for here.

Relatedly, one could wonder whether we should just use a lazy Haskell pattern in *filter*. Alas, Haskell does not allow a lazy pattern to bind existential variables: writing $\sim(MkEV\ v) \leftarrow \text{filter } p\ xs$ in Figure 1(a) would cause a compile-time error. This restriction in today's Haskell is not incidental, because the internal language would require exactly the power of the **open** approach we propose here in order to support such a lazy pattern.

Aside 1. Why lazy **unpack** is no easy answer

Beyond just the change to the types and the disappearance of terms to pack and unpack existentials, we can observe that the *laziness* of the function has changed. (See Aside 1 for why we cannot easily make **unpack** bind lazily.) In Figure 1(a), we see that the recursive call to *filter* must be made before the use of the cons operator $:>$. This means that, say, computing *take 2 (filter p vec)* (assuming *take* is clever enough to expect an *ExVec*) requires computing the result of the entire *filter*, even though the analogous expression on lists would only require filtering enough of *vec* to get the first two elements that satisfy *p*. The implementation of *filter* also requires enough stack space to store all the recursive calls, requiring an amount of space linear in the length of the input vector.

By contrast, the implementation in Figure 1(b) is lazy in the tail of the vector. Computing *take 2 (filter p vec)* really would only process enough elements of *vec* to find the first 2 that satisfy *p*. In addition, the computation requires only constant stack space, because *filter* will immediately return a cons cell storing a thunk for filtering the tail. If a bounded number of elements satisfy *p*, this is an asymptotic improvement in space requirements.

We can support the behavior evident in Figure 1(b) only because we use *strong* existential packages, where the existentially packed type can be projected out from the existential package, instead of relying on the use of a pattern-match. Furthermore, projection of the packed type is requires no evaluation of any expression. We return to explain more about this key innovation in Section 3.

2.3 Object Encoding

Suppose we have a pretty-printer feature in our application, making use of the following class:

```
class Pretty a where
  pretty :: a → Doc
```

There are *Pretty* instances defined for all relevant types. Now, suppose we have $order :: Order$, $client :: Client$, and $status :: OrderStatus$; we wish to create a message concatenating these three details. Today, we might say $vcat [pretty\ order, pretty\ client, pretty\ status]$, where $vcat :: [Doc] \rightarrow Doc$. However, equipped with lightweight existentials, we could instead write $vcat [order, client, status]$, where $vcat :: [\exists a. Pretty\ a \wedge a] \rightarrow Doc$. Here, the \wedge type constructor allows us to pack a witness for a constraint (such as a type class dictionary [Hall et al. 1996]) inside an existential package. Each element of the list is checked against the type $\exists a. Pretty\ a \wedge a$. Choosing one, checking $order$ against $\exists a. Pretty\ a \wedge a$ uses unification to determine that the choice of a should be $Order$, and we will then need to satisfy a *Pretty Order* constraint. In the implementation of $vcat$, elements of type $\exists a. Pretty\ a \wedge a$ will be available as arguments to *pretty*:

```
vcat :: [\exists a. Pretty a \wedge a] \rightarrow Doc
vcat []      = empty
vcat (x : xs) = pretty x $$ vcat xs
```

While the code simplification at call sites is modest, the ability to abstract over a constraint in forming a list makes it easier to avoid the types from preventing users from expressing their thoughts more directly.

Our main formal presentation in this paper does not include the packed constraints required here, but Section 9.2 considers an extension to our work that would support this example.

2.4 Richly Typed Data Structures

Suppose we wish to design a datatype whose inhabitants meet certain invariants by construction. If the invariants are complex enough, this can be done only by designing the datatype as a generalized algebraic datatype (GADT) [Xi et al. 2003]. Though other examples in this space abound (for example, encoding binary trees [McBride 2014] and regular expressions [Weirich 2018]), we will use the idea of a well-typed expression language, perhaps familiar to our readers.¹

The idea is encapsulated in these definitions:

```
data Ty = Ty :> Ty | ... -- base types elided
type Exp :: [Ty] -- types of in-scope variables
         \> Ty -- type of expression
         \> Type
data Exp ctx ty where
  App :: Exp ctx (arg :> result) \> Exp ctx arg \> Exp ctx result
  ...
```

An expression of type $Exp\ ctx\ ty$ is guaranteed to be well-typed in our object language: note that a function application requires the function to have a function type $arg :> result$ and the argument to have type arg . (The ctx is a list of the types of in-scope variables; using de Bruijn indices means we do not need to map names.) We are thus unable to represent the syntax tree applying, say, the number 5 to an argument *True*.

However, if we are to use Exp in a running interpreter, we have a problem: users might not type well-typed expressions. How can we take a user-written program and represent it in Exp ? We must type-check it.

¹This well-worn idea perhaps originates in a paper by Pfenning and Lee [1989], though that paper does not use an indexed datatype. Augustsson and Carlsson [1999] extend the idea to use a datatype, much as we have done here. A more in-depth treatment of this example is the subject of a functional pearl by Eisenberg [2020].

Assuming a type $UExp$ (“unchecked expression”) that is like Exp but without its indices, we would write the following:²

```

typecheck :: (ctx :: [ Ty ]) → UExp → Maybe (∃ty. Exp ctx ty)
typecheck ctx (UApp fun arg) = do -- using the Maybe monad
  fun' ← typecheck ctx fun
  arg' ← typecheck ctx arg
  -- decompose the type of fun' into expectedArgTy :→ _resultTy:
  (expectedArgTy, _resultTy) ← checkFunctionTy (typeOf fun')
  -- Check whether expectedArgTy and the type of arg' are the same (failing if not)
  -- Refl is a proof the types coincide; matching on it reveals this fact to the type-checker:
  Refl ← checkEqual expectedArgTy (typeOf arg')
  return (App fun' arg')

```

The use of an existential type is critical here. There is no way to know what the type of an expression is before checking it, and yet we need this type available for compile-time reasoning to be able to accept the final use of App . An example such as this one can be written today, but with extra awkward packing and unpacking of existentials, or through the use of a continuation-passing encoding. With the use of lightweight existentials, an example like this is easier to write, lowering the barrier to writing richly typed, finely specified programs.

3 KEY IDEA: EXISTENTIAL PROJECTIONS

In our envisioned source language, introduction and elimination of existential types are implicit. Precise locations are determined by type inference (as pinned down in Section 4)—accordingly, these locations may be hard to predict. Once these locations have been identified, the compiler must produce a fully annotated, typed core language that makes these introductions and eliminations explicit. We provide a precise account of this core language in Section 5. But before we do that, we use this section to informally justify why we need new forms in the first place. Why can we no longer use the existing encoding of existential types (based on Mitchell and Plotkin [1988] and Läufer [1996]) internally?

The key observation is that, since the locations of introductions and eliminations are hard to predict, they must not affect evaluation. Any other design would mean that programmers lose the ability to reason about when their expressions are reduced.

The existing datatype-based approach requires an existential-typed expression to be evaluated to head normal form to access the *type* packed in the existential. This is silly, however: types are completely erased, and yet this rule means that we must perform runtime evaluation simply to access an erased component of a some data.

To illustrate the problem, consider this Haskell datatype:

```

data Exists (f :: Type → Type) = ∀(a :: Type). Ex !(f a)

```

With this construct, we can introduce existential types using the data constructor Ex and eliminate them by pattern matching on Ex . Note the presence of the strictness annotation, written with $!$. A use of the Ex data constructor, if it is automatically inserted by the type inferencer, must not block reduction.³

²This rendering of the example assumes the ability to write using dependent types, to avoid clutter. However, do not be distracted: the dependent types could easily be encoded using singletons [Eisenberg and Weirich 2012; Monnier and Haguenaer 2010], while we focus here on the use of existential types.

³Similarly, our choice of explicit introduction form for the core language must be strict in its argument if it is to be unobservable.

The difficult issue, however, is elimination. To access the value carried by *Exists*, we must use pattern matching. We cannot use a straightforward projection function $\text{unExists} :: \text{Exists } f \rightarrow f ???$: it would allow the abstracted type variable to escape its scope—exactly why we cannot write a well-scoped type signature for *unExists*. As a result, we cannot use this value without weak-head evaluation of the term. As Section 3.2 shows, this forcing can decrease the laziness of our program.

While perhaps not as fundamental as our desire for introduction and elimination to be transparent to evaluation, another design goal is to allow arbitrary **let**-inlining. In other words, if $\text{let } x = e1 \text{ in } e2$ type-checks, then $e2 [e1 / x]$ should also type-check. This property gives flexibility to users: they (and their IDEs) can confidently refactor their program without fear of type errors.

Taken together, these design requirements—transparency to evaluation and support for **let**-inlining—drive us to enhance our core language with *strong* existentials [Howard 1969]: existentials that allow projection of both the type witness and the packed value, without pattern-matching.⁴

3.1 Strong Existentials via pack and open

Our core language $\mathbb{F}\mathbb{X}$ adopts the following constructs for introducing and eliminating existential types:⁵

$$\frac{\text{PACK} \quad \Gamma \vdash e : \tau_2[\tau_1/a]}{\Gamma \vdash \text{pack } \tau_1, e \text{ as } \exists a. \tau_2 : \exists a. \tau_2} \quad \frac{\text{OPEN} \quad \Gamma \vdash e : \exists a. \tau}{\Gamma \vdash \text{open } e : \tau[[e : \exists a. \tau] / a]}$$

The **pack** typing rule is fairly standard [Pierce 2002, Chapter 24]. This term creates an existential package, hiding a type τ_1 in the package with an expression e . Our operational semantics (Figure 7) includes a rule that makes this construct strict.

To eliminate existential types, we use the **open** construct (from Cardelli and Leroy [1990]) instead of pattern matching. The **open** construct eliminates an existential without forcing it, as **opens** are simply erased during compilation. The type of **open** e is interesting: we substitute away the bound variable a , replacing it with $[e : \exists a. \tau]$. This type is an *existential projection*. The idea is that we can think of an existential package $\exists a. \tau$ as a (dependent) pair, combining the choice for a (say, τ_0) with an expression of type $\tau[\tau_0 / a]$. The type $[e : \exists a. \tau]$ projects out the type τ_0 from the pair.

A key aspect of **open** is that the type form $[e : \exists a. \tau]$ is a completely opaque type. In our surface language, $[e : \exists a. \tau]$ is equal to itself and no other type. Computation is not necessary in types. One way to think of this is to imagine that $[e : \exists a. \tau]$ is like a fresh type variable whose name is long—not as a construct that, say, accesses a type within e .

The simple idea of **open** is very powerful. It means that we can talk about the type in an existential package without unpacking the package. It would even be valid to project out the type of an existential package that will never be computed. Because types can be erased in our semantics, even projecting out the type from a bottoming expression (of existential type) is harmless.⁶

Note that the type of the existential package expression is included in the syntax for projections $[e : \exists a. \tau]$: this annotation is necessary because expressions in our surface language \mathbb{X} might have multiple, different types. (For example, $\lambda x \rightarrow x$ has both type $\text{Int} \rightarrow \text{Int}$ and type $\text{Bool} \rightarrow \text{Bool}$.) Including the type annotation fixes our interpretation of e , but see Section 6 for more on this point.

⁴Strong existentials stand in contrast to *weak* existentials. A strong existential package supports operators that access the encapsulated type and datum, while a weak existential requires pattern-matching in order to extract the datum and bring its type into scope. In a lazy language, strong existentials thus have greater expressive power, as we can use a lazy projection, as we do here.

⁵These rules are slightly simplified. The full rules appear in Section 5.

⁶Readers may be alarmed at that sentence: how could $[\perp : \exists a. a]$ be a valid type? Perhaps a more elaborate system might want to reject such a type, but there is no need to. As all types are erased and have no impact on evaluation, an exotic type like this is no threat to type safety.

3.2 The **unpack** Trap

Adding the **open** term to the language comes at a cost to complexity. Let us take a moment to reflect on why a more traditional elimination form (called **unpack**) is insufficient.

A frequent presentation of existentials in a language based on System F uses the **unpack** primitive. Pierce [2002, Chapter 24] presents the idea with this typing rule:

$$\text{UNPACK} \frac{\Gamma \vdash e_1 : \exists a. \tau_2 \quad \Gamma, a, x:\tau_2 \vdash e_2 : \tau \quad a \notin \text{fv}(\tau)}{\Gamma \vdash \text{unpack } e_1 \text{ as } a, x \text{ in } e_2 : \tau}$$

The idea is that **unpack** extracts out the packed expression in a variable x , also binding a type variable a to represent the hidden type. The typing rule corresponds to the pattern-match in **case** e_1 of $Ex (x :: _ a) \rightarrow e_2$, where x and a are brought into scope in e_2 .⁷

This approach is attractive because it is simple to add to a language like System F. It does not require the presence of terms in types and the necessary machinery that we describe in Section 5. However, it is also not powerful enough to accommodate some of the examples we would like to support.

*The **unpack** term impacts evaluation.* Because it is based on pattern matching, the **unpack** term must reduce its argument to a weak-head normal form before providing access to the hidden type. The standard reduction rule looks like this:

$$\text{unpack } (\text{pack } \tau_1, e_1 \text{ as } \exists a. \tau_2) \text{ as } a, x \text{ in } e_2 \longrightarrow e_1[e_1/x][\tau_1/a]$$

What this rule means is that the only parts of the term that have access to the abstract type are the ones that are evaluated after the existential has been weak-head normalized. Without weak-head normalizing the argument to a **pack**, we have nothing to substitute for x and a .

Let us rewrite the *filter* example from Section 2.2, making more details explicit so that we can see why this is an issue.

```
filter :: ∀n a. (a → Bool) → Vec n a → ∃m. Vec m a
filter = λn a → λ(p :: a → Bool) (vec :: Vec n a) →
  case vec of
    (:>) n1 (x :: a) (xs :: Vec n1 a)      -- vec is x :> xs
    | p x      → ...
    | otherwise → filter n1 a p xs
  Nil → pack Zero, Nil as ∃m. Vec m a  -- vec is Nil
```

The treatment above makes all type abstraction and application explicit. Note that the pattern-match for the cons operator $:>$ includes a compile-time (or type-level) binding for the length of the tail, $n1$.

The question here is: what do we put in the ... in the case where $p x$ holds? One possibility is to apply the $(:>)$ operator to build the result. However, right away, we are stymied: what do we pass to that operator as the length of the resulting vector? It depends on the length of the result of the recursive call. A use of **unpack** cannot help us here, as **unpack** is used in a term, not in a type index; even if we could use it, we would have to return the packed type, not something we can ordinarily do.

⁷See Eisenberg et al. [2018] for more details on how Haskell treats that type annotation.

Instead, we must use **unpack** (and **pack**) *before* calling the (:>) operator. Specifically, we can write

```
unpack filter n1 a p xs as n2, ys in pack n2, ( $\text{:>}$ ) n2 x ys as  $\exists m. \text{Vec } m \ a$ 
```

This use of **unpack** is type-correct, but we have lost the laziness of *filter* we so prized in Section 2.2.

On the other hand, **open** allows us to fill in the ... with the following code, using the the existential projection to access the new (type-level) length for the arguments to **pack** and to :> .

```
let ys ::  $\exists m. \text{Vec } m \ a$  -- usual lazy let
```

```
    ys = filter n1 a p xs
```

```
in pack [ys ::  $\exists m. \text{Vec } m \ a$ ], ( $\text{:>}$ ) [ys ::  $\exists m. \text{Vec } m \ a$ ] x (open ys) as  $\exists m. \text{Vec } m \ a$ 
```

As we expand on in the next subsection, we do not have to **let**-bind *ys*; instead, we could just repeat the sub-expression *filter n1 a p xs*.

3.3 The Importance of Strength

Beyond the peculiarities of the *filter* example, having a lazy construct that accesses the abstracted type in an existential package is essential to supporting inferrable existential types.

Here is a somewhat contrived example to illustrate this point:

```
data Counter a = Counter { zero :: a, succ :: a  $\rightarrow$  a, toInt :: a  $\rightarrow$  Int }
```

```
mkCounter :: String  $\rightarrow$   $\exists a. \text{Counter } a$  -- a counter with a hidden representation
```

```
mkCounter = ...
```

```
initial1 :: Int
```

```
initial1 = let c = mkCounter "hello" in (toInt c) (zero c)
```

```
initial2 :: Int
```

```
initial2 = (toInt (mkCounter "hello")) (zero (mkCounter "hello"))
```

We would like our language to accept both *initial1* and *initial2*. After all, one of the benefits of working in a pure, lazy language is referential transparency: programmers (and tools, such as IDEs) should be able to perform expression inlining with no change in behavior. In both *initial1* and *initial2*, the compiler must automatically eliminate the existential that results from each use of *mkCounter*. In the definition *initial1*, elaboration is not difficult, even if we only have the weak **unpack** elimination form to work with.

However, supporting *initial2* is more problematic. Maintaining the order of evaluation of the source language requires two separate uses of the elimination form.

To type-check the application of *toInt (mkCounter "hello")* to *zero (mkCounter "hello")*, we must first know the type packed into the package returned from *mkCounter "hello"*. Accessing this type should not evaluate *mkCounter "hello"*, however: a programmer rightly expects that *toInt* is evaluated before any call to *mkCounter* is, which may have performance or termination implications. More generally, we can imagine the need for a hidden type arbitrarily far away from the call site of a function (such as *mkCounter*) that returns an existential; eager evaluation of the function would be most unexpected for programmers.

Note that, critically, both calls to *mkCounter* in *initial2* contain the *same* argument. Since we are working in a pure context, we know that the result of the two calls to *mkCounter "hello"* in *initial2* must be the same, and thus that the program is well-typed.

In sum, if the compiler is to produce the elimination form for existentials, that elimination form must be *nonstrict*, allowing the packed witness type to be accessed without evaluation. Any other

choice means that programmers must expect hard-to-predict changes to the evaluation order of their program. In addition, if we wish to allow users to inline their **let**-bound identifiers, this projection form must also be *strong*, and remember the existentially typed expression in its type.

Note that we are taking advantage of Haskell’s purity in this part of the design. We can soundly support a strong elimination form like **open** only because we know that the expressions which appear in types are pure. All projections of the type witness from the same expression will be equal. In a language without this property, such as ML, we would need to enforce a value restriction on the type projections. Such a value restriction would prevent us from injecting, say, a non-deterministic expression into types; as there is no notion of evaluating a type, it would be unclear when and how often to evaluate the expression which could yield different results at each evaluation.

4 INFERRING EXISTENTIALS

In this section we present the surface language, \mathbb{X} , that we use to manipulate existentials, and the bidirectional type system that infers them. As our concrete setting is in Haskell, our starting point is the surface language described by Serrano et al. [2020], modified to add support for existentials. We add a syntax for existential quantifiers $\exists a.\epsilon$ and existential projections $[e : \epsilon]$. An important part of our type system is the type instantiation mechanism, which implicitly handles the opening of existentials (Section 4.3).

4.1 Language Syntax

The syntax of our types is given in Figure 2.

σ	$::= \epsilon \mid \forall a.\sigma$	universally quantified type
ϵ	$::= \rho \mid \exists b.\epsilon$	existentially quantified type
ρ	$::= \tau \mid \sigma_1 \rightarrow \sigma_2$	top-level monomorphic type
τ	$::= a \mid \text{Int} \mid \tau_1 \rightarrow \tau_2 \mid [e : \epsilon]$	monomorphic type
a, b	$::= \dots$	type variable
Γ	$::= \emptyset \mid \Gamma, a \mid \Gamma, x:\sigma$	typing context

Fig. 2. Type stratification

Polytypes σ can quantify an arbitrary number (including 0) universal variables and, within the universal quantification, an arbitrary number (including 0) existential variables. This stratification is enforced through the distinction between σ -types and ϵ -types. Note that the type $\exists a.\forall b.\tau$ is ruled out.⁸ Top-level monotypes ρ have no top-level quantification. Monotypes τ include a projection form $[e : \epsilon]$ that occurs every time an existential is opened, as described in Section 3.1. Universal and existential variables draw from the same set of variable names, denoted with a or b .

The expressions of \mathbb{X} are defined in Figure 3. This language is a fairly small λ -calculus, with type annotations and n -ary application (including type application). The expression $h \pi_1 \dots \pi_n$ applies a head to a sequence of arguments π_i that can be expressions or types. The head is either a variable x , an annotated expression $e :: \sigma$, or an expression e that is not an application.⁹

⁸As usual, stratifying the grammar of types simplifies type inference. In our case, this choice drastically simplifies the challenge of comparing types with mixed quantifiers. Dunfield and Krishnaswami [2019, Section 2] have an in-depth discussion of this challenge.

⁹Our grammar does not force a head expression h to be something other than an application, but we will consistently assume this restriction is in force. It would add clutter and obscure our point to bake this restriction in the grammar.

x	$::= \dots$	term variable
n	$::= \dots$	integer literal
e	$::= h\bar{\pi} \mid \lambda x. e \mid \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 \mid n$	expression
h	$::= x \mid e \mid e :: \sigma$	expression head
π	$::= e \mid \sigma$	argument

Fig. 3. Our surface language, \mathbb{X}

An important complication of our type system is that expressions may appear in types: this happens in the projection form $[e : \epsilon]$. We thus must address how to treat type equality. For example, suppose term variable x (of type Int) is free in a type τ ; is $\tau[(\lambda y. y) 1 / x]$ equal to $\tau[1 / x]$? That is, does type equality respect β -reduction? Our answer is “no”: we restrict type equality in our language to be syntactic equality (modulo α -equivalence, as usual). We can imagine a richer type equality relation—which would accept more programs—but this simplest, least expressive version satisfies our needs. (However, see [Aside 2](#) in [Section 7.3](#) for a wrinkle here.) Adding such an equality relation is largely orthogonal to the concerns around existential types that draw our focus.¹⁰

4.2 Type System

The typing rules of our language appear in [Figure 4](#). This bidirectional type system uses two forms for typing judgments: $\Gamma \vdash e \Rightarrow \rho$ means that, in the type environment Γ , the program e has the inferred type ρ , while $\Gamma \vdash e \Leftarrow \rho$ means that, in the type environment Γ , e is checked to have type ρ . We also use a third form to simplify the presentation of the rules: $\Gamma \vdash e \Leftrightarrow \rho$, which means that the rule can be read by replacing \Leftrightarrow with either \Rightarrow or \Leftarrow in both the conclusion and premises. Although the rules are fairly close to the standard rules of a typed λ -calculus, handling existentials through **packing** and **opening** has an impact on the rules **LET** and **GEN**.

We review the rules in [Figure 4](#) here, deferring the most involved rule, **APP**, until after we discuss the instantiation judgment \vdash^{inst} , in [Section 4.3](#).

4.2.1 Simple Subsumption. Bidirectional type systems typically rely on a reflexive, transitive *subsumption relation* \leq , where we expect that if $e : \sigma_1$ and $\sigma_1 \leq \sigma_2$, then $e : \sigma_2$ is also derivable. For example, we would expect that $\forall a. a \rightarrow a \leq \text{Int} \rightarrow \text{Int}$. This subsumption relation is then used when “switching modes”; that is, if we are checking an expression e against a type σ_2 where e has a form resistant to type propagation (the case when e is a function call), we infer a type σ_1 for e and then check that $\sigma_1 \leq \sigma_2$.

However, our type system refers to no such \leq relation: we essentially use equality as our subsumption relation, invoking it implicitly in our rules through the use of a repeated metavariable. (Though hard to see, the repeated metavariable is the ρ_r in rule **APP**, when replacing the \Leftrightarrow in the conclusion with a \Leftarrow .) We get away with this because our bidirectional type-checking algorithm works over top-level monotypes ρ , not the more general polytype σ . A type ρ has no top-level quantification at all. Furthermore, our type system treats all types as invariant—including \rightarrow . This treatment follows on from the ideas in [Serrano et al. \[2020, Section 5.8\]](#), which describes how Haskell recently made its arrow type similarly invariant.

We adopt this simpler approach toward subsumption both to connect our presentation with the state-of-the-art for type inference in Haskell [[Serrano et al. 2020](#)] and also because this approach

¹⁰Our core language \mathbb{FX} does need to think harder about this question, in order to prove type safety. See [Section 5.1](#).

$\Gamma \vdash^{\forall} e \Leftarrow \sigma$	<i>(Universal type checking)</i>									
$\text{GEN} \quad \frac{\Gamma, \bar{a} \vdash e \Leftarrow \rho[\bar{\tau} / \bar{b}] \quad f\nu(\bar{\tau}) \subseteq \text{dom}(\Gamma, \bar{a})}{\Gamma \vdash^{\forall} e \Leftarrow \forall \bar{a}. \exists \bar{b}. \rho}$										
$\Gamma \vdash e \Rightarrow \rho \quad \Gamma \vdash e \Leftarrow \rho$	<i>(Type synthesis and type checking)</i>									
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%; vertical-align: top; padding: 5px;"> $\text{APP} \quad \frac{\Gamma \vdash_h h \Rightarrow \sigma \quad \Gamma \vdash^{\text{inst}} h : \sigma ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r \quad \bar{e} = \text{exprargs}(\bar{\pi}) \quad \Gamma \vdash^{\forall} e_i \Leftarrow \sigma_i}{\Gamma \vdash h \bar{\pi} \Leftrightarrow \rho_r}$ </td> <td style="width: 33%; vertical-align: top; padding: 5px;"> $\text{LABS} \quad \frac{\Gamma, x:\tau \vdash e \Rightarrow \rho \quad f\nu(\tau) \subseteq \text{dom}(\Gamma) \quad \bar{a} \text{ fresh} \quad \rho' = \rho[\bar{a} / \lfloor \rho \rfloor_x] \quad (\text{see } \S 4.2.3)}{\Gamma \vdash \lambda x. e \Rightarrow \tau \rightarrow \exists \bar{a}. \rho'}$ </td> <td style="width: 33%; vertical-align: top; padding: 5px;"> $\text{CABS} \quad \frac{\Gamma, x:\sigma_1 \vdash^{\forall} e \Leftarrow \sigma_2 \quad f\nu(\sigma_1) \subseteq \text{dom}(\Gamma)}{\Gamma \vdash \lambda x. e \Leftarrow \sigma_1 \rightarrow \sigma_2}$ </td> </tr> <tr> <td colspan="3" style="text-align: center; padding: 10px 0;"> $\text{LET} \quad \frac{\Gamma \vdash e_1 \Rightarrow \rho_1 \quad \bar{a} = f\nu(\rho_1) \setminus \text{dom}(\Gamma) \quad \Gamma, x:\forall \bar{a}. \rho_1 \vdash e_2 \Leftrightarrow \rho_2}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \Leftrightarrow \rho_2[e_1 / x]}$ </td> </tr> <tr> <td colspan="3" style="text-align: center; padding: 10px 0;"> $\text{INT} \quad \frac{}{\Gamma \vdash n \Leftrightarrow \text{Int}}$ </td> </tr> </table>		$\text{APP} \quad \frac{\Gamma \vdash_h h \Rightarrow \sigma \quad \Gamma \vdash^{\text{inst}} h : \sigma ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r \quad \bar{e} = \text{exprargs}(\bar{\pi}) \quad \Gamma \vdash^{\forall} e_i \Leftarrow \sigma_i}{\Gamma \vdash h \bar{\pi} \Leftrightarrow \rho_r}$	$\text{LABS} \quad \frac{\Gamma, x:\tau \vdash e \Rightarrow \rho \quad f\nu(\tau) \subseteq \text{dom}(\Gamma) \quad \bar{a} \text{ fresh} \quad \rho' = \rho[\bar{a} / \lfloor \rho \rfloor_x] \quad (\text{see } \S 4.2.3)}{\Gamma \vdash \lambda x. e \Rightarrow \tau \rightarrow \exists \bar{a}. \rho'}$	$\text{CABS} \quad \frac{\Gamma, x:\sigma_1 \vdash^{\forall} e \Leftarrow \sigma_2 \quad f\nu(\sigma_1) \subseteq \text{dom}(\Gamma)}{\Gamma \vdash \lambda x. e \Leftarrow \sigma_1 \rightarrow \sigma_2}$	$\text{LET} \quad \frac{\Gamma \vdash e_1 \Rightarrow \rho_1 \quad \bar{a} = f\nu(\rho_1) \setminus \text{dom}(\Gamma) \quad \Gamma, x:\forall \bar{a}. \rho_1 \vdash e_2 \Leftrightarrow \rho_2}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \Leftrightarrow \rho_2[e_1 / x]}$			$\text{INT} \quad \frac{}{\Gamma \vdash n \Leftrightarrow \text{Int}}$		
$\text{APP} \quad \frac{\Gamma \vdash_h h \Rightarrow \sigma \quad \Gamma \vdash^{\text{inst}} h : \sigma ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r \quad \bar{e} = \text{exprargs}(\bar{\pi}) \quad \Gamma \vdash^{\forall} e_i \Leftarrow \sigma_i}{\Gamma \vdash h \bar{\pi} \Leftrightarrow \rho_r}$	$\text{LABS} \quad \frac{\Gamma, x:\tau \vdash e \Rightarrow \rho \quad f\nu(\tau) \subseteq \text{dom}(\Gamma) \quad \bar{a} \text{ fresh} \quad \rho' = \rho[\bar{a} / \lfloor \rho \rfloor_x] \quad (\text{see } \S 4.2.3)}{\Gamma \vdash \lambda x. e \Rightarrow \tau \rightarrow \exists \bar{a}. \rho'}$	$\text{CABS} \quad \frac{\Gamma, x:\sigma_1 \vdash^{\forall} e \Leftarrow \sigma_2 \quad f\nu(\sigma_1) \subseteq \text{dom}(\Gamma)}{\Gamma \vdash \lambda x. e \Leftarrow \sigma_1 \rightarrow \sigma_2}$								
$\text{LET} \quad \frac{\Gamma \vdash e_1 \Rightarrow \rho_1 \quad \bar{a} = f\nu(\rho_1) \setminus \text{dom}(\Gamma) \quad \Gamma, x:\forall \bar{a}. \rho_1 \vdash e_2 \Leftrightarrow \rho_2}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \Leftrightarrow \rho_2[e_1 / x]}$										
$\text{INT} \quad \frac{}{\Gamma \vdash n \Leftrightarrow \text{Int}}$										
$\Gamma \vdash_h h \Rightarrow \sigma$	<i>(Head synthesis)</i>									
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%; vertical-align: top; padding: 5px;"> $\text{H-VAR} \quad \frac{x:\sigma \in \Gamma}{\Gamma \vdash_h x \Rightarrow \sigma}$ </td> <td style="width: 33%; vertical-align: top; padding: 5px;"> $\text{H-ANN} \quad \frac{\Gamma \vdash^{\forall} e \Leftarrow \sigma \quad f\nu(\sigma) \subseteq \text{dom}(\Gamma)}{\Gamma \vdash_h (e :: \sigma) \Rightarrow \sigma}$ </td> <td style="width: 33%; vertical-align: top; padding: 5px;"> $\text{H-INFER} \quad \frac{\Gamma \vdash e \Rightarrow \rho}{\Gamma \vdash_h e \Rightarrow \rho}$ </td> </tr> </table>		$\text{H-VAR} \quad \frac{x:\sigma \in \Gamma}{\Gamma \vdash_h x \Rightarrow \sigma}$	$\text{H-ANN} \quad \frac{\Gamma \vdash^{\forall} e \Leftarrow \sigma \quad f\nu(\sigma) \subseteq \text{dom}(\Gamma)}{\Gamma \vdash_h (e :: \sigma) \Rightarrow \sigma}$	$\text{H-INFER} \quad \frac{\Gamma \vdash e \Rightarrow \rho}{\Gamma \vdash_h e \Rightarrow \rho}$						
$\text{H-VAR} \quad \frac{x:\sigma \in \Gamma}{\Gamma \vdash_h x \Rightarrow \sigma}$	$\text{H-ANN} \quad \frac{\Gamma \vdash^{\forall} e \Leftarrow \sigma \quad f\nu(\sigma) \subseteq \text{dom}(\Gamma)}{\Gamma \vdash_h (e :: \sigma) \Rightarrow \sigma}$	$\text{H-INFER} \quad \frac{\Gamma \vdash e \Rightarrow \rho}{\Gamma \vdash_h e \Rightarrow \rho}$								

Fig. 4. Type inference for \mathbb{X}

simplifies our typing rules. We see no obstacle to incorporating our ideas with a more powerful subsumption judgment, such as the deep-skolemization judgment of Peyton Jones et al. [2007, Section 4.6.2] or the slightly simpler co- and contravariant judgment of Odersky and Läufer [1996, Figure 2].

4.2.2 Checking against a Polytype. Rule **GEN**, the sole rule for the $\Gamma \vdash^{\forall} e \Leftarrow \sigma$ judgment, deals with the case when we are checking against a polytype σ . If we want to ensure that e has type σ , then we must *skolemize* any universal variables bound in σ : these variables behave essentially as fresh constants while type-checking e . Rule **GEN** thus just brings them into scope.

On the other hand, if there are existential variables bound in σ , then we must *instantiate* these. If we are checking that e has some type $\exists a.\tau_0$, that means we must find some type τ such that e has type $\tau_0[\tau / a]$. This is very different than the skolemization of a universal variable, where we must keep the variable abstract. Instead, when checking against $\exists a.e$, we guess a monotype τ and check e against the type $e[\tau / a]$. Rule **GEN** simply does this for nested existential quantification over variables \bar{b} . A real implementation might use unification variables, but we here rely on the rich body of literature [e.g., Dunfield and Krishnaswami 2013] that allows us to guess monotypes

during type inference, knowing how to translate this convention into an implementation using unification variables.

4.2.3 Abstractions. Rule **iAbs** synthesizes the type of a λ -abstraction, by guessing the (mono)type τ of the bound variable and then inferring the type of the body e to be ρ . However, rule **iAbs** also can pack existentials. This is necessary to avoid skolem escape: it is possible that the type ρ contains x free. However, it would be disastrous if $\lambda x.e$ was assigned a type mentioning x , as x is no longer in scope.

For example, suppose we have $\Gamma = f:\text{Int} \rightarrow \exists a.a \rightarrow \text{Bool}$. Now, consider inferring the type ρ in $\Gamma \vdash \lambda x.f x \Rightarrow \rho$. Guessing $x:\text{Int}$, we will infer $\Gamma, x:\text{Int} \vdash f x \Rightarrow [f x : \exists a.a \rightarrow \text{Bool}] \rightarrow \text{Bool}$. It is tempting, then, to say $\Gamma \vdash \lambda x.f x \Rightarrow \text{Int} \rightarrow [f x : \exists a.a \rightarrow \text{Bool}] \rightarrow \text{Bool}$, but this is wrong: the type mentions x free, but Γ does not bind x . Instead, rule **iAbs** infers $\Gamma \vdash \lambda x.f x \Rightarrow \exists a.\text{Int} \rightarrow a \rightarrow \text{Bool}$, by using a instead of the ill-scoped $[f x : \exists a.a \rightarrow \text{Bool}]$.

More generally, we must identify all existential projections within ρ that have x free. These are replaced with fresh variables \bar{a} . We use the notation $[\rho]_x$ to denote the list of projections in ρ ; multiple projections of the same expression (that is, multiple occurrences of $[e_0 : \epsilon_0]$ for some e_0 and ϵ_0) are commoned up in this list. Formally,

$$[\rho]_x = \{[e : \epsilon] \mid ([e : \epsilon] \text{ is a sub-expression of } \rho) \wedge (x \text{ is a free variable in } e)\}.$$

The notation $\rho[\bar{a} / [\rho]_x]$ denotes the type ρ where the \bar{a} are written in place of these projections. Note that this notation is set up *backward* from the way it usually works, where we substitute some type for a variable. Here, instead, we are replacing the type with a fresh variable.

In the conclusion of the rule, we existentially quantify the \bar{a} , to finally obtain a function type of the form $\tau \rightarrow \exists \bar{a}.\rho'$.¹¹

The checking rule **cAbs** is much simpler. We know the type of the bound variable by decomposing the known expected type $\sigma_1 \rightarrow \sigma_2$. We also need not worry about skolem escape because we have been provided with a well-scoped σ_2 result type for our function. The only small wrinkle is the need to use \vdash^\forall in order to invoke rule **Gen** to remove any quantifiers on the type σ_2 .

4.2.4 Let Skolem-escape. Rule **Let** deals with **let**-expressions, both in synthesis and in checking modes. It performs standard **let**-generalization, computing generalized variables \bar{a} by finding the free variables in ρ_1 and removing any variables additionally free in Γ . Indeed, all that is unexpected in this rule is the type substitution in the conclusion.

The problem, like with rule **iAbs** is the potential for skolem-escape. The variable x might appear in the type ρ_2 . However, x is out of scope in the conclusion, and thus it cannot appear in the overall type of the **let**-expression. One solution to this problem would be to **pack** all the existentials that fall out of scope, much like we do in rule **iAbs**. However, doing so would mean that our bidirectional type system now infers existential types ϵ instead of top-level monomorphic types ρ ; keeping with the simpler ρ is important to avoid the complications of a non-trivial subsumption judgment. Hence we choose to replace all occurrences of x inside of projections by the expression e_1 . This does not pose a problem since e_1 is well-typed according to the premises of the **Let** rule.

4.2.5 Inferring the Types of Heads. Following [Serrano et al. \[2020\]](#), our system treats n -ary applications directly, instead of recurring down a chain of binary applications $e_1 e_2$. The head of an n -ary application is denoted with h ; heads' types are inferred with the $\Gamma \vdash_h h \Rightarrow \sigma$ judgment. Variables simply perform a context lookup, annotated expressions check the contained expression against the

¹¹Our language works well without this special substitution. Instead, we could have a check that the final inferred type in rule **iAbs** is well scoped. However, this extra existential packing is easy enough to add, and so we have.

$$\boxed{\Gamma \vdash^{\text{inst}} e : \sigma ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r} \quad (\text{Instantiation judgment})$$

$$\begin{array}{c}
\text{ITYARG} \\
\frac{\Gamma \vdash^{\text{inst}} e \sigma' : \sigma[\sigma' / a] ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r \quad f\nu(\sigma') \subseteq \text{dom}(\Gamma)}{\Gamma \vdash^{\text{inst}} e : \forall a. \sigma ; \sigma', \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r}
\end{array}
\quad
\begin{array}{c}
\text{IARG} \\
\frac{\Gamma \vdash^{\text{inst}} e e' : \sigma_2 ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r}{\Gamma \vdash^{\text{inst}} e : (\sigma_1 \rightarrow \sigma_2) ; e', \bar{\pi} \rightsquigarrow \sigma_1, \bar{\sigma} ; \rho_r}
\end{array}$$

$$\begin{array}{c}
\text{IALL} \\
\frac{\Gamma \vdash^{\text{inst}} e : \sigma[\tau / a] ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r \quad \bar{\pi} \neq \sigma', \bar{\pi}' \quad f\nu(\tau) \subseteq \text{dom}(\Gamma)}{\Gamma \vdash^{\text{inst}} e : \forall a. \sigma ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r}
\end{array}
\quad
\begin{array}{c}
\text{IEXIST} \\
\frac{\Gamma \vdash^{\text{inst}} e : \epsilon[[e : \exists a. \epsilon] / a] ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r}{\Gamma \vdash^{\text{inst}} e : \exists a. \epsilon ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r}
\end{array}$$

$$\begin{array}{c}
\text{IRESET} \\
\frac{}{\Gamma \vdash^{\text{inst}} e : \rho_r ; [] \rightsquigarrow [] ; \rho_r}
\end{array}$$

Fig. 5. Instantiation

provided type, and other expressions infer a ρ -type. It is understood here that we use rule **H-INFER** only when the other rules do not apply, for example, for λ -abstractions.

4.3 Instantiation Semantics

The instantiation rules of Figure 5 present an auxiliary judgment used in type-checking applications. The judgment $\Gamma \vdash^{\text{inst}} e : \sigma ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r$ means: with in-scope variables Γ , apply function e of type σ to arguments $\bar{\pi}$ requires $\text{exprargs}(\bar{\pi})$ (the value arguments) to have types $\bar{\sigma}$, resulting in an expression $e \bar{\pi}$ of type ρ_r . This judgment is directly inspired by Serrano et al. [2020, Figure 4].

The idea is that we use \vdash^{inst} to figure out the types of term-level arguments to a function in a pre-pass that examines only type arguments. Having determined the expected types of the term-level arguments $\bar{\sigma}$, rule **APP** (in Figure 4) actually checks that the arguments have the correct types. This pre-pass is not necessary in order to infer the types for existentials, but it sets the stage for Section 8, where we integrate our design with the current implementation in GHC.

Application. Rule **ITYARG** handles type application by instantiating the bound variable a with the supplied type argument σ' . Rule **IARG** handles routine expression application simply by remembering that the argument should have type σ_1 . Note that we do *not* check that the argument e' has type σ_1 here.

Quantifiers. Rule **IALL** deals with universal quantifiers in the function's type by instantiating with a guessed monotype τ . The first premise is to avoid ambiguity with rule **ITYARG**; we do not wish to guess an instantiation when the user provides it explicitly with a type argument.

Rule **IEXIST** eagerly opens existentials by substituting a projection in place of the bound variable a . This is the only place in the judgment where we need the function expression e : whenever we open an existential type, we must remember what expression has that type, so that we do not confuse two different existentially packed types.

For example, if f has type $\text{Bool} \rightarrow \exists b.(b, b \rightarrow \text{Int})$, then the function application $f \text{True}$ will be given the opened pair type:

$$([f \text{True} : \exists b.(b, b \rightarrow \text{Int})], [f \text{True} : \exists b.(b, b \rightarrow \text{Int})] \rightarrow \text{Int})$$

Rule **IREsULT** concludes computing the instantiation in a function application by copying the function type to be the result type.

The APP rule. Having now understood the instantiation judgment, we turn our attention to rule **APP**. After inferring the type σ for an application head h , σ gets instantiated, revealing argument types $\bar{\sigma}$. Each argument e_i is checked against its corresponding type σ_i , where the entire function application expression has type ρ_r . Rule **APP** operates in both synthesis and checking modes. When synthesizing, it simply returns ρ_r from the instantiation judgment; when checking, it ensures that the instantiated type ρ_r matches what was expected. We need do no further instantiation or skolemization because we have a simple subsumption relation.

5 CORE LANGUAGE

Perhaps we can infer existential types using existential projections $[e : \epsilon]$, but how do we know such an approach is sound? We show that it is by elaborating our surface expressions into a core language $\mathbb{F}\mathbb{X}$, inspired by a similar language described by [Cardelli and Leroy \[1990, Section 4\]](#), and we prove the standard progress and preservation theorems of this language. This section presents $\mathbb{F}\mathbb{X}$ and states key metatheory results; the following section connects \mathbb{X} to $\mathbb{F}\mathbb{X}$ by presenting our elaboration algorithm.

The syntax of $\mathbb{F}\mathbb{X}$ is in Figure 6 and selected typing rules are in Figure 7; full typing rules appear in the appendix.¹² Note that we use upright Latin letters to denote $\mathbb{F}\mathbb{X}$ expressions and types; when we mix \mathbb{X} and $\mathbb{F}\mathbb{X}$ in close proximity, we additionally use colors.

B	$::= \rightarrow \text{Int} \dots$	base type
t, r, s	$::= a B\bar{t} \forall a.t \exists a.t [e]$	type
e, h	$::= x n \lambda x:t.e e_1 e_2 \Lambda a.e e t \mathbf{pack} t, e \mathbf{as} t_2$ $\mathbf{open} e \mathbf{let} x = e_1 \mathbf{in} e_2 e \triangleright \gamma$	expression
v	$::= n \lambda x:t.e \Lambda a.v \mathbf{pack} t, v \mathbf{as} t_2$	value
γ	$::= \langle t \rangle \mathbf{sym} \gamma \gamma_1 ;; \gamma_2 [\eta] \gamma_1 @ \gamma_2 \mathbf{projpack} t, e \mathbf{as} t_2 \dots$	type coercion
η	$::= e \triangleright \gamma \mathbf{step} e$	expression coercion
G	$::= \emptyset G, x : t G, a$	typing context

Fig. 6. Syntax of the core language, $\mathbb{F}\mathbb{X}$

The nub of $\mathbb{F}\mathbb{X}$ is System F, with fully applied base types B (because they are fully applied, we do not need to have a kind system) and ordinary universal quantification. We thus omit typing rules from this presentation that are standard. The inclusion of existential types, **pack** and **open** is fitting for a core language supporting existentials. This language necessarily has mutually recursive grammars for types and expressions, but the typing rules are not mutually recursive: rule **CT-PROJ** shows that a projection in a type is well-formed when the expression is well-scoped. (The $\vdash G \mathbf{ok}$ premise refers to a routine context-well-formedness judgment, omitted.) We do not require the existential package to be well-typed (though it would be, in practice).

5.1 Coercions

The biggest surprise in $\mathbb{F}\mathbb{X}$ is its need for type and expression *coercions*. The motivation for these can be seen in rule **CS-OPENPACK**. If we are stepping an expression **open** ($\mathbf{pack} t, v \mathbf{as} \exists a.t_2$), we want to extract the value v from the existential package. The problem is that v has the wrong type.

¹²<https://richarde.dev/papers/2021/exists/exists-extended.pdf>

$G \vdash e : t$	<p style="text-align: right;"><i>(Expression typing)</i></p> <div style="display: flex; justify-content: space-between;"> <div style="width: 30%;"> <p>CE-ABS</p> $\frac{G, x : t_1 \vdash e : t_2 \quad x \notin \text{fv}(t_2)}{G \vdash \lambda x : t_1. e : t_1 \rightarrow t_2}$ </div> <div style="width: 30%;"> <p>CE-LET</p> $\frac{G \vdash e_1 : t_1 \quad G, x : t_1 \vdash e_2 : t_2}{G \vdash \text{let } x = e_1 \text{ in } e_2 : t_2[e_1 / x]}$ </div> <div style="width: 30%;"> <p>CE-PACK</p> $\frac{G \vdash t : \text{type} \quad G \vdash \exists a. t_2 : \text{type} \quad G \vdash e : t_2[t / a]}{G \vdash \text{pack } t, e \text{ as } \exists a. t_2 : \exists a. t_2}$ </div> </div> <div style="display: flex; justify-content: space-around; margin-top: 10px;"> <div style="width: 45%;"> <p>CE-OPEN</p> $\frac{G \vdash e : \exists a. t}{G \vdash \text{open } e : t[[e] / a]}$ </div> <div style="width: 45%;"> <p>CE-CAST</p> $\frac{G \vdash e : t_1 \quad G \vdash \gamma : t_1 \sim t_2}{G \vdash e \triangleright \gamma : t_2}$ </div> </div>
$G \vdash t : \text{type}$	<p style="text-align: right;"><i>(Type well-formedness)</i></p> <p>CT-PROJ</p> $\frac{\vdash G \text{ ok} \quad \text{fv}(e) \subseteq \text{dom}(G)}{G \vdash [e] : \text{type}}$
$G \vdash \gamma : t_1 \sim t_2$	<p style="text-align: right;"><i>(Type coercion typing)</i></p> <div style="display: flex; justify-content: space-between;"> <div style="width: 20%;"> <p>CG-REFL</p> $\frac{G \vdash t : \text{type}}{G \vdash \langle t \rangle : t \sim t}$ </div> <div style="width: 20%;"> <p>CG-SYM</p> $\frac{G \vdash \gamma : t_1 \sim t_2}{G \vdash \text{sym } \gamma : t_2 \sim t_1}$ </div> <div style="width: 20%;"> <p>CG-TRANS</p> $\frac{G \vdash \gamma_1 : t_1 \sim t_2 \quad G \vdash \gamma_2 : t_2 \sim t_3}{G \vdash \gamma_1 ; \gamma_2 : t_1 \sim t_3}$ </div> <div style="width: 20%;"> <p>CG-PROJ</p> $\frac{G \vdash \eta : e_1 \sim e_2}{G \vdash [\eta] : [e_1] \sim [e_2]}$ </div> </div> <div style="display: flex; justify-content: space-around; margin-top: 10px;"> <div style="width: 45%;"> <p>CG-INSTEXISTS</p> $\frac{G \vdash \gamma_1 : (\exists a. t_1) \sim (\exists a. t_2) \quad G \vdash \gamma_2 : t_3 \sim t_4}{G \vdash \gamma_1 @ \gamma_2 : t_1[t_3 / a] \sim t_2[t_4 / a]}$ </div> <div style="width: 45%;"> <p>CG-PROJPACK</p> $\frac{G \vdash \text{pack } t, e \text{ as } t_2 : t_2}{G \vdash \text{projpack } t, e \text{ as } t_2 : [\text{pack } t, e \text{ as } t_2] \sim t_2}$ </div> </div>
$G \vdash \eta : e_1 \sim e_2$	<p style="text-align: right;"><i>(Expression coercion typing)</i></p> <p>CH-STEP</p> $\frac{G \vdash e : t \quad G \vdash e' : t \quad G \vdash e \longrightarrow e'}{G \vdash \text{step } e : e \sim e'}$
$G \vdash e \longrightarrow e'$	<p style="text-align: right;"><i>(Small-step operational semantics)</i></p> <p>CS-PACKCONG</p> $\frac{G \vdash e \longrightarrow e'}{G \vdash \text{pack } t, e \text{ as } t_2 \longrightarrow \text{pack } t, e' \text{ as } t_2}$ <p>CS-OPENPACK</p> $\frac{}{G \vdash \text{open } (\text{pack } t, v \text{ as } t_2) \longrightarrow v \triangleright \langle t_2 \rangle @ (\text{sym } (\text{projpack } t, v \text{ as } t_2))}$ <p>CS-OPENCONG</p> $\frac{G \vdash e : t \quad G \vdash e \longrightarrow e'}{G \vdash \text{open } e \longrightarrow \text{open } e' \triangleright \langle t \rangle @ (\text{sym } [\text{step } e])}$

Fig. 7. Selected typing and reduction rules of the core language, $\mathbb{F}\mathbb{X}$

Suppose that v has type t_0 . Then, we have $\text{pack } t, v \text{ as } \exists a. t_2 : \exists a. t_2$ and $\text{open } (\text{pack } t, v \text{ as } \exists a. t_2) : t_2 [\text{pack } t, v \text{ as } \exists a. t_2 / a]$, according to rule **CE-OPEN**. This last type is not syntactically the same as t_0 , although it must be that $t_0 = t_2 [t / a]$ to satisfy the premises of rule **CE-PACK**. Because the type of the opened existential does not match the type of the packed value, a naïve reduction rule like $G \vdash \text{open } (\text{pack } t, v \text{ as } t_2) \longrightarrow v$ would not preserve types.

There are, in general, two ways to build a type system when encountering such a problem. We could have a non-trivial type equality relation, where we say that $[\text{pack } t, e \text{ as } t_2] \equiv t$. Doing so would simplify the reduction rules, but this simplification comes at a cost: our language would now have a conversion rule that allows an expression of one type t_1 to have another type t_2 as long as $t_1 \equiv t_2$. This rule is not syntax-directed; accordingly, it is hard to determine whether type-checking remains decidable. Furthermore, a non-trivial type equality relation makes proofs considerably more involved. In effect, we are just moving the complexity we see in the right-hand side of a rule like rule **CS-OPENPACK** into the proofs.

The alternative approach to a non-trivial equality relation is to use explicit coercions, as we have here. The cost is clutter. Casts sully our reduction steps, and we need to explicitly shunt coercions in several (omitted, unenlightening) reduction rules—for example, when reducing $((\lambda x:t. e_1) \triangleright \gamma) e_2$ where the cast intervenes between a λ -abstraction and its argument. Despite the presence of these rules in our operational semantics, coercions can be fully erased: we can write an alternative, untyped operational semantics that omits coercions entirely. Theorem 7.2 shows that erasure preserves program behavior.

Both approaches—an enriched definitional equality vs. explicit coercions—are essentially equivalent: we can view explicit coercions simply as an encoding of the derivation of an equality judgment.¹³ We choose explicit coercions both because FX is a purely internal language (and thus clutter is less noisome) and because it allows for an easy connection to the implementation of the core language in GHC, based on System FC [Sulzmann et al. 2007], with similar explicit coercions.

The coercion language for FX includes constructors witnessing that they encode an equivalence relation (rules **CG-REFL**, **CG-SYM**, and **CG-TRANS**), along with several omitted forms showing that the equivalence is also a congruence over types. Coercions also include several decomposition operations; rule **CG-INSTEXISTS** shows one, used in our reduction rules. The two forms of interest to use are $[\eta]$ (rule **CG-PROJ**) and **projpack** (rule **CG-PROJPACK**). The former injects the equivalence relation on expressions (witnessed by expression coercions η) into the type equivalence relation, and the latter witnesses the equivalence between $[\text{pack } t, e \text{ as } t]$ and its packed type t .

The equivalence relation on expressions is surprisingly simple: we need only the two rules in Figure 7. These rules allow us to drop casts (supporting a coherence property which states that the presence of casts is essentially unimportant) and to reduce expressions.

5.2 Metatheory

We prove (almost) standard progress and preservation theorems for this language:

THEOREM 5.1 (PROGRESS). *If $G \vdash e : t$, where G contains only type variable bindings, then one of the following is true:*

- (1) *there exists e' such that $G \vdash e \longrightarrow e'$;*
- (2) *e is a value v ; or*
- (3) *e is a casted value $v \triangleright \gamma$.*

THEOREM 5.2 (PRESERVATION). *If $G \vdash e : t$ and $G \vdash e \longrightarrow e'$, then $G \vdash e' : t$.*

¹³Weirich et al. [2017] makes this equivalence even clearer by presenting two proved-equivalent versions of a language, one with a non-trivial, undecidable type equality relation and another with explicit coercions.

$\Gamma \vdash^{\forall} e \Leftarrow \sigma \Rightarrow \mathbf{e}$ $\Gamma \vdash e \Leftarrow \rho \Rightarrow \mathbf{e}$ $\Gamma \vdash_h h \Rightarrow \sigma \Rightarrow \mathbf{h}$ $\Gamma \vdash^{\text{inst}} e : \sigma \Rightarrow \mathbf{e}; \bar{\pi} \rightsquigarrow \bar{\sigma}; \rho_r \Rightarrow \mathbf{e}_r$ $\sigma \Rightarrow \mathbf{s}$ $\Gamma \Rightarrow \mathbf{G}$	elaboration of polymorphic expressions elaboration of expressions elaboration of application heads elaboration of application spines elaboration of types elaboration of typing contexts
ELAB-GEN $\frac{\Gamma, \bar{a} \vdash e \Leftarrow \rho[\bar{\tau}/\bar{b}] \Rightarrow \mathbf{e} \quad \tau \Rightarrow \mathbf{t} \quad \rho \Rightarrow \mathbf{r} \quad \text{fv}(\bar{\tau}) \subseteq \text{dom}(\Gamma, \bar{a})}{\Gamma \vdash^{\forall} e \Leftarrow \forall \bar{a}. \exists \bar{b}. \rho \Rightarrow \Lambda \bar{a}. \text{pack } \bar{t}, \mathbf{e} \text{ as } \exists \bar{b}. \mathbf{r}}$	
ELAB-IABS	
$\bar{a} \text{ fresh}$	
$\Gamma, x:\tau \vdash e \Rightarrow \rho \Rightarrow \mathbf{e}$ $\text{fv}(\tau) \subseteq \text{dom}(\Gamma)$ $\rho' = \rho[\bar{a}/[\rho]_x] \quad \tau \Rightarrow \mathbf{t}$ $\rho \Rightarrow \mathbf{r} \quad \rho' \Rightarrow \mathbf{r}'$ <hr style="width: 100%;"/> $\Gamma \vdash \lambda x. e \Rightarrow \tau \rightarrow \exists \bar{a}. \rho' \Rightarrow \lambda x:\mathbf{t}. \text{pack } [\mathbf{r}]_{x}, \mathbf{e} \text{ as } \exists \bar{a}. \mathbf{r}'$	ELAB-APP $\Gamma \vdash_h h \Rightarrow \sigma \Rightarrow \mathbf{h}$ $\Gamma \vdash^{\text{inst}} h : \sigma \Rightarrow \mathbf{h}; \bar{\pi} \rightsquigarrow \bar{\sigma}; \rho_r \Rightarrow \mathbf{e}_r$ <hr style="width: 100%;"/> $\Gamma \vdash h \bar{\pi} \Leftarrow \rho_r \Rightarrow \mathbf{e}_r$
ELAB-IARG $\frac{\Gamma \vdash^{\forall} e' \Leftarrow \sigma_1 \Rightarrow \mathbf{e}' \quad \Gamma \vdash^{\text{inst}} e' : \sigma_2 \Rightarrow \mathbf{e} \mathbf{e}'; \bar{\pi} \rightsquigarrow \bar{\sigma}; \rho_r \Rightarrow \mathbf{e}_r}{\Gamma \vdash^{\text{inst}} e : (\sigma_1 \rightarrow \sigma_2) \Rightarrow \mathbf{e}; \mathbf{e}', \bar{\pi} \rightsquigarrow \sigma_1, \bar{\sigma}; \rho_r \Rightarrow \mathbf{e}_r}$	
ELAB-IEEXIST $\frac{\Gamma \vdash^{\text{inst}} e : \epsilon[[e : \exists a. \epsilon] / a] \Rightarrow \text{open } \mathbf{e}; \bar{\pi} \rightsquigarrow \bar{\sigma}; \rho_r \Rightarrow \mathbf{e}_r}{\Gamma \vdash^{\text{inst}} e : \exists a. \epsilon \Rightarrow \mathbf{e}; \bar{\pi} \rightsquigarrow \bar{\sigma}; \rho_r \Rightarrow \mathbf{e}_r}$	
ELAB-IRESULT $\frac{}{\Gamma \vdash^{\text{inst}} e : \rho_r \Rightarrow \mathbf{e}_r; [] \rightsquigarrow []; \rho_r \Rightarrow \mathbf{e}_r}$	

Fig. 8. Judgments and selected rules for elaborating from \mathbb{X} into \mathbb{FX} .

In addition, we prove that types can still be erased in this language. Let $|e|$ denote the expression e with all type abstractions, type applications, **packs**, **opens** and casts dropped. Furthermore, overload \longrightarrow to mean the reduction relation over the erased language.

THEOREM 5.3 (ERASURE). *If $G \vdash e \longrightarrow^* e'$, then $|e| \longrightarrow^* |e'|$.*

The proofs largely follow the pattern set by previous papers on languages with explicit coercions and are unenlightening. They appear, in full, in the appendix.

6 ELABORATION

We now augment our inference rules from Section 4 to describe the elaboration from the surface language \mathbb{X} into our core \mathbb{FX} . The notation \Rightarrow denotes elaboration of a surface term, type or context

into its core equivalent. Some of our rules appear in Figure 8. The rest appear in the appendix. In order to aid understanding, we use **blue** for \mathbb{X} terms and **red** for \mathbb{FX} terms.

The rules in Figure 8 allow packing multiple existentials at once, when given a list of types as the first argument to **pack**; see rules **ELAB-GEN** and **ELAB-IABS**. Rule **ELAB-GEN** checks a surface expression e against an expected type $\forall \bar{a}. \exists \bar{b}. \rho$. We see that the result of elaboration uses nested Λ -abstractions and our nested **pack** notation to produce an \mathbb{FX} expression that has the desired type. Rule **ELAB-IABS** echoes rule **IABS**, producing an \mathbb{FX} expression with **packs** necessary to accommodate any projections that mention the bound variable x ; recall the special treatment of such projections from Section 4.2.3.

Rule **ELAB-APP** elaborates the head h to \mathbf{h} , and then calls the \vdash^{inst} judgment. This judgment takes the elaborated \mathbf{h} as an *input* (despite its appearance on the right of a \Rightarrow). This input of an elaborated expression is built up as the application spine is checked, to be returned in rule **ELAB-IRESULT**. In order to build this elaborated expression as we go, rule **ELAB-IARG** elaborates arguments, in contrast to our original rule **IARG**; rule **ELAB-APP** then no longer needs to check these arguments in a second pass.¹⁴ Rule **ELAB-IEEXIST** is the place where **open** is introduced, as it open an expression with an existential type.

The omitted rules all appear in the appendix and broadly follow the pattern set here.

6.1 Tweaking the IEEXIST Rule

In the instantiation judgment for the surface language (Figure 5), rule **IEEXIST** opens existentials. That is, given an expression e with an existential type $\exists a. \epsilon$, it infers for e the type resulting from replacing the type variable with the projection $[e : \exists a. \epsilon]$. However, these projections pose a problem during the elaboration process. Specifically, if we have an application $e_1 e_2$ such that e_1 expects an argument whose type mentions $[e_0 : \epsilon]$ —and e_2 indeed has a type mentioning $[e_0 : \epsilon]$ —we cannot be sure that the application remains well-typed after elaboration. After all, type-checking in \mathbb{X} is non-deterministic, given the way it guesses instantiations and the types of λ -bound variables. Another wrinkle is that $[e_0 : \epsilon]$ might appear under binders, making it even easier for type inference to come to two different conclusions when computing $\Gamma \vdash^{\vee} e_0 \Leftarrow \epsilon$.

There are two approaches to fix this problem: we can require our elaboration process to be deterministic, or we can modify rule **IEEXIST** to make sure that projections in the surface language actually use pre-elaborated core expressions. We take the latter approach, as it is simpler and more direct. However, we discuss later in this section the possible disadvantages of this choice, and a route to consider the first one.

Accordingly, we now introduce the following new **IEEXISTCORE** and rule **LETCORE** rules, replacing rules **IEEXIST** and rule **LET**:

$$\begin{array}{c}
 \text{IEEXISTCORE} \\
 \Gamma \vdash^{\vee} e \Leftarrow \exists a. \epsilon \Rightarrow \mathbf{e} \\
 \hline
 \Gamma \vdash^{\text{inst}} e : \epsilon[[e] / a] ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r \\
 \hline
 \Gamma \vdash^{\text{inst}} e : \exists a. \epsilon ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r
 \end{array}
 \qquad
 \begin{array}{c}
 \text{LETCORE} \\
 \Gamma \vdash e_1 \Rightarrow \rho_1 \Rightarrow \mathbf{e}_1 \\
 \bar{a} = \text{fv}(\rho_1) \setminus \text{dom}(\Gamma) \\
 \Gamma, x : \forall \bar{a}. \rho_1 \vdash e_2 \Leftarrow \rho_2 \\
 \hline
 \Gamma \vdash \text{let } x = \mathbf{e}_1 \text{ in } \mathbf{e}_2 \Leftarrow \rho_2[\Lambda \bar{a}. \mathbf{e}_1 / x]
 \end{array}$$

Fig. 9. Updated rules to support \mathbb{FX} expressions in \mathbb{X} types

¹⁴Knowledgeable readers will wonder how this new treatment interacts with the Quick-Look algorithm, which critically depends on waiting to type-check arguments *after* a quick look at the entire argument spine. The solution is to be lazy: the elaborated is not needed until after all arguments have been checked. Accordingly, we could, for example, use a mutable cell to hold the elaborated expression, and then fill in this cell only during the second pass. Our formal presentation here need not worry about this technicality, however.

Now, the elaboration process $\tau \Rightarrow \mathfrak{t}$ is indeed deterministic, making \Rightarrow a function on types τ and contexts Γ . Having surmounted this hurdle, elaboration largely very straightforward.

6.2 A Different Approach

We may want to refrain from using core expressions inside of projections, because doing so introduces complexity for the programmer who is not otherwise exposed to the core language. To wit, \mathbb{X} would keep using projections of the form $[e : \epsilon]$, where we understand that $\Gamma \vdash^{\forall} e \Leftarrow \epsilon$ in the ambient context Γ , while \mathbb{FX} uses the form $[e]$.

It is vitally important that, if our surface-language typing rules accept a program, the elaborated version of that program is type-correct. (We call this property *soundness*; it is Theorem 7.1.) Yet, if elaboration of types is non-deterministic, we will lose this property, as explained above.

This alternative approach is simply to *assume* that elaboration is deterministic. Doing so is warranted because, in practice, a type-checker implementation will proceed deterministically—it seems far-fetched to think that a real type-checker would choose different types for the same expression and expected type, if any. In essence, a deterministic elaborator means that we can consider $[e : \epsilon]$ as a proxy for $[e]$. The first is preferable to programmers because it is written in the language they program in. However, a type-checker implementation may choose to use the latter, and thus avoid the possibility of unsoundness from arising out of a non-deterministic elaborator.

7 ANALYSIS

The surface language \mathbb{X} allows us to easily manipulate existentials in a λ -calculus while delegating type consistency to an explicit core language \mathbb{FX} . The following theorems establish the soundness of this approach, via the elaboration transformation \Rightarrow , as well as the general expressivity and consistency of our bidirectional type system.

7.1 Soundness

If our surface language is to be type safe, we must know that any term accepted in the surface language corresponds to a well-typed term in the core language:

THEOREM 7.1 (SOUNDNESS).

- (1) If $\Gamma \vdash^{\forall} e \Leftarrow \sigma \Rightarrow e$, then $G \vdash e : s$, where $\Gamma \Rightarrow G$ and $\sigma \Rightarrow s$.
- (2) If $\Gamma \vdash e \Rightarrow \rho \Rightarrow e$, then $G \vdash e : r$, where $\Gamma \Rightarrow G$ and $\rho \Rightarrow r$.
- (3) If $\Gamma \vdash e \Leftarrow \rho \Rightarrow e$, then $G \vdash e : r$, where $\Gamma \Rightarrow G$ and $\rho \Rightarrow r$.

Furthermore, in order to eliminate the possibility of a trivial elaboration scheme, we would want the elaborated term to behave like the surface-language one. We capture this property in this theorem:

THEOREM 7.2 (ELABORATION ERASURE).

- (1) If $\Gamma \vdash^{\forall} e \Leftarrow \sigma \Rightarrow e$, then $|e| = |e|$.
- (2) If $\Gamma \vdash e \Rightarrow \rho \Rightarrow e$, then $|e| = |e|$.
- (3) If $\Gamma \vdash e \Leftarrow \rho \Rightarrow e$, then $|e| = |e|$.

This theorem asserts that, if we remove all type annotations and applications, the \mathbb{X} expression is the same as the \mathbb{FX} one.

7.2 Conservativity

Not only do we want our \mathbb{X} programs to be sound, but we also want \mathbb{X} to be a comfortable language to program in. As our language is an extension of Hindley-Milner, we know that all the conveniences programmers are used to in that setting carry over here.

Theorem 7.4 tells us that expanding out a well-typed **let** remains well typed. However, if we selectively expand a repeated **let**, a larger expression may become ill typed. Suppose we have $f :: \text{Int} \rightarrow \exists a. (a, a \rightarrow \text{Int})$ and write $(\text{snd } (f (\text{let } x = 5 \text{ in } x+x))) (\text{fst } (f (\text{let } x = 5 \text{ in } x+x)))$. That expression is a well-typed Int . However, if we inline only one of the **lets**, to $(\text{snd } (f (5 + 5))) (\text{fst } (f (\text{let } x = 5 \text{ in } x + x)))$, we now have an ill-typed expression. The problem is that our language uses a very fine-grained expression equality relation: just α -equivalence. Accordingly, $\text{let } x = 5 \text{ in } x + x$ and $5 + 5$ are considered distinct, and when these expressions appear in types (via existential projections), the types are different.

The solution is straightforward, if not entirely lightweight: extend the expression equality relation. Doing so would require a more explicit treatment of equality in our type inference algorithm (in particular, rule **APP** of Figure 4 would need to invoke the equality relation), as well as additions to FX to accommodate this new development. It is not clear whether the added expressiveness are worth the complexity cost, and so we kept our equivalence relationship simple for ease of presentation.

Aside 2. Selective **let**-inlining sometimes causes trouble

THEOREM 7.3 (CONSERVATIVE EXTENSION OF HINDLEY-MILNER). *If e has no type arguments or type annotations, and Γ, e, τ, σ contain no existentials, then:*

- (1) $(\Gamma \vdash_{HM} e : \tau)$ implies $(\Gamma \vdash e \Rightarrow \tau)$
- (2) $(\Gamma \vdash_{HM} e : \sigma)$ implies $(\Gamma \vdash^{\forall} e \Leftarrow \sigma)$

where \vdash_{HM} denotes typing in the Hindley-Milner type system, as described by Clément et al. [1986, Figure 3].

7.3 Stability

The following theorems denote stability properties [Bottu and Eisenberg 2021]. In other words, they ensure that small user-written transformations do not change drastically the static semantics of our programs. The **let**-inlining property is specifically permitted by our approach to existentials, and it is a major feature of our type system.

THEOREM 7.4 (LET-INLINING). *If x is free in e_2 then:*

$$\begin{array}{ll} (\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \Rightarrow \rho) & \text{implies } (\Gamma \vdash e_2[e_1 / x] \Rightarrow \rho) \\ (\Gamma \vdash^{\forall} \text{let } x = e_1 \text{ in } e_2 \Leftarrow \sigma) & \text{implies } (\Gamma \vdash^{\forall} e_2[e_1 / x] \Leftarrow \sigma) \end{array}$$

Interestingly, the system we present here does not support a small generalization of the **let**-inlining property, as we explore in Aside 2.

This next theorem tells us that the order variables appear in an existential quantification does not affect usage sites:

THEOREM 7.5 (ORDER OF QUANTIFICATION DOES NOT MATTER). *Let ρ' (resp. σ') be two types that differ from ρ (resp. σ) only by the ordering of quantified type variables in their existential types. Then:*

- (1) $(\Gamma \vdash e \Rightarrow \rho)$ if and only if $(\Gamma \vdash e \Rightarrow \rho')$
- (2) $(\Gamma \vdash^{\forall} e \Leftarrow \sigma)$ if and only if $(\Gamma \vdash^{\forall} e \Leftarrow \sigma')$

Lastly, this theorem tells us that extra, redundant type annotations do not disrupt typability:

THEOREM 7.6 (SYNTHESIS IMPLIES CHECKING). *If $\Gamma \vdash e \Rightarrow \rho$ then $\Gamma \vdash e \Leftarrow \rho$.*

$$\boxed{\Gamma \vdash^{\forall} e \Leftarrow \sigma}$$

(Universal type checking)

$$\begin{array}{c} \text{GENIMPREDICATIVE} \\ \bar{\kappa} \text{ fresh} \quad \rho' = \rho[\bar{\kappa} / \bar{b}] \\ \Gamma, \bar{a} \vdash_{\zeta} e : \rho' \rightsquigarrow \Theta \\ \rho'' = \Theta \rho' \\ \text{dom}(\theta) = \text{fiv}(\rho'') \\ \Gamma, \bar{a} \vdash e \Leftarrow \theta \rho'' \\ \hline \Gamma \vdash^{\forall} e \Leftarrow \forall \bar{a}. \exists \bar{b}. \rho \end{array}$$

Fig. 10. Allowing impredicative instantiation in the \vdash^{\forall} judgment

8 INTEGRATING WITH TODAY'S GHC AND QUICK LOOK

We envision integrating our design into GHC, allowing Haskell programmers to use existential types in their programs. Accordingly, we must consider how our work fits with GHC's latest type-inference algorithm, dubbed Quick Look [Serrano et al. 2020]. The structure behind our inference algorithm—with heads applied to lists of arguments instead of nested applications—is based directly on Quick Look, and it is straightforward to extend our work to be fully backward-compatible with that design. Indeed, our extension is essentially orthogonal to the innovations of impredicative type inference in the Quick Look algorithm.

It would take us too far afield from our primary goal—describing type inference for existential types—to explain the details of Quick Look here. We thus build on the text already written by Serrano et al. [2020]; readers uninterested in the details may safely skip the rest of this section.

Serrano et al. [2020] explains their algorithm progressively, by stating in their Figures 3 and 4 a baseline system. That baseline also effectively serves as our baseline here. Then, in their Figure 5, the authors add a few new premises to specific rules, along with judgments those premises refer to. Given this modular presentation, we can adopt the same changes: their rule **IARG** is our rule **IARG**, and their rule **APP-↓** is our rule **APP**. The only wrinkle in merging these systems is that their presentation uses a notion of *instantiation variable*, which Serrano et al. write as κ . An instantiation variable is allowed to unify with a polytype, in contrast to an ordinary unification variable, which must unify with a monotype. Given that impredicative instantiation is not a primary goal of our work, we choose not to use this approach in our main formal presentation, instead preferring the more conventional idiom of using guessed τ -types. However, in order to integrate inferred existentials with Quick Look impredicativity, we must explicitly use instantiation variables in the rule below.

Since we have a more elaborate notion of polytype, one rule needs adjustment in our system: the rule implementing the $\Gamma \vdash^{\forall} e \Leftarrow \sigma$ judgment, rule **GEN**. That rule skolemizes (makes fresh constants out of) the variables universally quantified in σ and guesses $\bar{\tau}$ to instantiate the existentially quantified variables. In order to allow these instantiations to be impredicative, we must modify the rule, as in Figure 10.

This rule follows broadly the pattern from rule **GEN**, but using instantiation variables $\bar{\kappa}$ instead of guessing $\bar{\tau}$. The third premise invokes the Quick Look judgment \vdash_{ζ} [Serrano et al. 2020, Figure 5] to generate a substitution Θ . Such a substitution Θ maps instantiation variables κ to polytypes σ ; by contrast, a substitution θ includes only monotypes τ in its codomain. The next two premises of rule **GENIMPREDICATIVE** apply the Θ substitution, and then use θ to eliminate any remaining instantiation variables κ : the $\text{fiv}(\rho'')$ extracts all the *free* instantiation variables in ρ'' . Note that

the range of θ appears unconstrained here; the types in its range are guessed, just like the $\bar{\tau}$ in rule **GEN**.

With this one new rule—along with the changes evident in Figure 5 of Serrano et al.—our system supports impredicative type inference, and is a conservative extension of their algorithm.

9 DISCUSSION

We have described how our inference algorithm allows users to program with existentials while avoiding the need to thinking about packing and unpacking. Here, we review some subtleties that arise as our approach encounters more practical settings.

9.1 No Declarative (Non-syntax-directed) System with Existentials

When we first set out to understand type inference with existentials better, our goal was to develop a type system with existential types, unguided type inference (no additional annotation obligations for the programmer), and principal types. Our assumption was that if this is possible with universal quantification [Hindley 1969; Milner 1978], it should also be possible for existential quantification. Unfortunately, it seems such a design is out of reach.

To see why, consider $f\ b = \mathbf{if}\ b\ \mathbf{then}\ (1, \lambda y \rightarrow y + 1)\ \mathbf{else}\ (True, \lambda z \rightarrow 1)$. We can see that f can be assigned one of two different types:

- (1) $Bool \rightarrow \exists a. (a, Int \rightarrow Int)$
- (2) $Bool \rightarrow \exists a. (a, a \rightarrow Int)$

Neither of these types is more general than the other, and neither seems likely to be ruled out by straightforward syntactic restrictions (such as the Hindley-Milner type system's requirement that all universal quantification be in prenex form).

One possible approach to inference for a definition like f is to use an *anti-unification* [Pfenning 1991] algorithm to relate the types of $(1, \lambda y \rightarrow y + 1)$ and $(True, \lambda z \rightarrow 1)$: infer the former to have type $(Int, Int \rightarrow Int)$ and the latter to have type $(Bool, \alpha \rightarrow Int)$ for some unknown type α . The goal then is to find some type τ such that τ can instantiate to either of these two types: this is anti-unification. The problem is, in this case, α : we get different results depending on whether α becomes Int or $Bool$.

We might imagine a way of choosing between the two hypothetical types for f , above, but any such restriction would break the desired symmetry and elegance of a declarative system that allows arbitrary generalization and specialization. Instead, we settle for the practical, predictable bidirectional algorithm presented in this paper, leaving the search for a more declarative approach as an open problem—one we think unlikely to have a satisfying solution.

9.2 Class Constraints on Existentials

The algorithm we present in this paper works with a typing context storing the types of bound variables. In full Haskell, however, we also have a set of constraint assumptions, and accepting some expressions requires proving certain constraints. A type system with these assumptions and obligations is often called a *qualified type system* [Jones 1992]. Our extension to support both universal and existential qualified types is in Figure 11.

This extension introduces type classes C and constraints Q . Constraints are applied type classes (like $Show\ Int$), and perhaps others; the details are immaterial. Instead, we refer to an abstract logical entailment relation \Vdash , which relates assumptions and the constraints they entail. Universally quantified types σ can now require proving a constraint: to use $e : Q \Rightarrow \sigma$, the constraint Q must

$C ::= \dots$		type class
$Q ::= C \bar{c} \mid \dots$		constraint
$\sigma ::= \epsilon \mid \forall a. \sigma \mid Q \Rightarrow \sigma$		universally quantified type
$\epsilon ::= \rho \mid \exists b. \epsilon \mid Q \wedge \epsilon$		existentially quantified type
$\Gamma ::= \emptyset \mid \Gamma, a \mid \Gamma, x : \sigma \mid \Gamma, Q \mid \Gamma, [e : \epsilon]$		typing context
$\Gamma \Vdash Q$		logical entailment
GENQUALIFIED		
$\frac{\Gamma' = \Gamma, \bar{a}, Q_1, \overline{[e : Q \wedge \epsilon]}}{\Gamma' \vdash^V e \Leftarrow Q \wedge \epsilon} \quad \frac{e \in e_0}{\Gamma' \vdash e_0 \Leftarrow \rho[\bar{c}/\bar{b}]}$	$\text{IGIVEN} \quad \frac{\Gamma \vdash^{\text{inst}} e : \epsilon ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r}{[e : Q \wedge \epsilon] \in \Gamma}$	$\text{IWANTED} \quad \frac{\Gamma \vdash^{\text{inst}} e : \sigma ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r}{\Gamma \Vdash Q}$
$\Gamma \vdash^V e_0 \Leftarrow (\forall \bar{a}. Q_1 \Rightarrow \exists \bar{b}. Q_2 \wedge \rho)$	$\Gamma \vdash^{\text{inst}} e : Q \wedge \epsilon ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r$	$\Gamma \vdash^{\text{inst}} e : Q \Rightarrow \sigma ; \bar{\pi} \rightsquigarrow \bar{\sigma} ; \rho_r$

Fig. 11. Type system extension to support existentially packed class constraints

hold. Existentially quantified types ϵ can now provide the proof of a constraint: the expression $e : Q \wedge \epsilon$ contains evidence that Q holds. Assumed constraints appear in contexts Γ .¹⁵

The surprising feature here is that we have a new form of assumption, $[e : \epsilon]$. This assumption is allowed only when ϵ has the form $Q \wedge \epsilon'$; the assumed constraint is Q . However, by including the expression e that proves Q in the context, we remember how to compute Q when it is required.

9.2.1 Static Semantics. Examining the typing rules, we see rule **GENQUALIFIED** assumes Q_1 as a given (following the usual treatment of givens in qualified type systems) and also assumes an arbitrary list of projections $\overline{[e : \epsilon]}$. This arbitrary assumption is quite like how rule **GEN** assumes types \bar{c} to replace the existential variables \bar{b} . To prevent the type system from working in an unbounded search space for assumptions to make, the expressions e must be sub-expressions of our checked expression e_0 .

The instantiation judgment \vdash^{inst} must also accommodate constraints. When, in rule **IGIVEN**, it comes across an expression whose type includes a packed assumption $Q \wedge \epsilon$, it checks to make sure that assumption was included in Γ . The design here requiring an arbitrary guess of assumptions, only to validate the guess later, is merely because our presentation is somewhat declarative. By contrast, an implementation would work by emitting constraints and solving them (that is, computing \Vdash) later [Pottier and Rémy 2005]; when the constraint-generation pass encounters an expression of type $Q \wedge \epsilon$, it simply emits the constraint as a given. Rule **IWANTED** is a straightforward encoding of the usual behavior of qualified types, where the usage of an expression of type $Q \Rightarrow \sigma$ requires proving Q .

9.2.2 Dynamic Semantics. An interesting new challenge with packed class constraints is that class constraints are not erasable. In practice, a function *pretty* of type $\text{Pretty } a \Rightarrow a \rightarrow \text{String}$ (§2.3) takes *two* runtime arguments: a *dictionary* [Hall et al. 1996] containing implementations of the methods in *Pretty*, as well as the actual, visible argument of type a . When this dictionary comes from an existential projection, the expression producing the existential will have to be evaluated.

¹⁵Other presentations of qualified type systems frequently have a judgment that looks like $P \mid \Gamma \vdash e : \rho$, or similar, with a separate set of logical assumptions P . Because our assumptions may include expressions, we must mix the logical assumptions with variable assumptions right in the same context Γ .

For example, suppose we have $mk :: Bool \rightarrow \exists a. Pretty\ a \wedge a$ and call $pretty\ (mk\ True)$. Calling $pretty$ requires passing the dictionary giving the the implementation of the function at the specific type $pretty$ is instantiated at $(\lfloor mk\ True :: \exists a. Pretty\ a \wedge a \rfloor$, in this case). Getting this dictionary requires evaluating $mk\ True$. Naïvely, this means $mk\ True$ would be evaluated *twice*. This makes some sense if we think of $Q \wedge \epsilon$ as the type of pairs of a dictionary for Q and the inhabitant of ϵ : the naïve interpretation of $pretty\ (mk\ True)$ thus is like calling $pretty\ (fst\ (mk\ True))\ (snd\ (mk\ True))$. We do not address how to do better here, as standard optimization techniques can apply to improve the potential repeated work. Once again, purity works to our advantage here, in that we can be assured that commoning up the calls to $mk\ True$ does not introduce (or eliminate) effects.

9.3 Relevance and Existentials

One of the primary motivations for this work is to set the stage for an eventual connection between Liquid Haskell [Vazou et al. 2014] and the rest of Haskell’s type system. A Liquid Haskell refinement type is exemplified by $\{v :: Int \mid v \geq 0\}$; any element of such a type is guaranteed to be non-negative. Yet what would it mean to have a function $return$ such a type? To be concrete, let us imagine $mk :: Bool \rightarrow \{v :: Int \mid v \geq 0\}$. This function would return a value v of type Int , along with a proof that $v \geq 0$: this is a dependent pair, or an existential package. Thus, we can rephrase the type of mk to be $Bool \rightarrow \exists (v :: Int). Proof\ (v \geq 0)$, where $Proof\ q$ encodes a proof of the logical property q .

However, our new form of existential is different than the others considered in this paper. Here, the relevant part is the *first* component, not the second. That is, we want to be able to project out $v :: Int$ at runtime, discarding the compile-time proof that $v \geq 0$.

The core language presented in this paper cannot, without embellishment, support relevant first components of existentials. In other words, $\lfloor e : \epsilon \rfloor$ is always a compile-time type, never a runtime term. Nevertheless, existing approaches to deal with relevance will work in this new setting. Haskell’s \forall construct universally quantifies over an irrelevant type. Yet, work on dependent Haskell [Eisenberg 2016; Gundry 2013; Weirich et al. 2017] shows how we can make a similar, relevant construct. Similar approaches could work in a core language modeled on FX . Indeed, other dependently typed languages, such as Coq, Agda, and Idris support existential packages with relevant dependent components.

The big step our current work brings to this story is type inference. Whether relevant or not, we would still want existential packages to be packed and unpacked without explicit user direction, and we would still want type inference to have the properties of the algorithm presented in this paper. In effect, the choice of relevance of the dependent component is orthogonal to the concerns in this paper. We are thus confident that our approach would work in a setting with relevant types.

10 RELATED WORK

There is a long and rich body of literature informing our knowledge of existential types. We review some of the more prominent work here.

History. Existential types were present from the beginning in the design of polymorphic programming languages, present in Girard’s System F [Girard 1972] and independently discovered by Reynolds [1974], though in a less expressive form. Mitchell and Plotkin [1988] recognized the ability of existential types to model abstract datatypes and remarked on their connection with the Σ -types of Martin-Löf type theory [Martin-Löf 1975]. They proposed an elimination form, called *abstype*, that is equivalent to the now standard **unpack**.

Cardelli and Leroy [1990] compared Mitchell and Plotkin’s **unpack** based approach to various calculi with projection-based existentials. Their “calculus with a dot notation” includes the ability

for the type language to project the type component from term variables of an existential type. At the end of the report (Section 4), they generalize to allow arbitrary expressions in projections. It is this language that is most similar to our core language. They also note a number of examples that are expressible only in this language.

Integration with type inference. Full type checking and type inference for domain-free System F with existential types is known to be undecidable [Nakazawa and Tatsuta 2009; Nakazawa et al. 2008]. As a result, several language designers have used explicit forms such as datatype declarations or type annotations to extend their languages with existential types.

The datatype-based version of existentials found in GHC was first suggested by Perry [1991] and implemented in Hope+. It was formalized by Läufer and Odersky [1994] and implemented in the Caml Light compiler for ML, along with the Haskell B compiler [Augustsson 1994].

The Utrecht Haskell Compiler (UHC) also supports a version of existential type [Dijkstra 2005], in a form that does not require the explicit connection to datatypes found in GHC. As in this work, values of existential types can be opened in place, without the use of an **unpack** term. However, unlike here, UHC generates a fresh type variable for the abstracted type with each use of **open**. As a result, UHC does not need the form of dependent types that we propose, but also cannot express some of the examples allowed by our system (§3.3).

Leijen [2006] describes an extension of MLF [Le Botlan and Rémy 2003] with first-class existential types. Like this work, programmers never needed to add explicit **pack** or **unpack** expressions. However, because the type system was based on MLF, polymorphic types include instantiation constraints and the type-inference algorithm is very different from that used by GHC. In contrast, our work requires only a small extension of GHC's most recent implementation of first-class polymorphism. Furthermore, Leijen does not describe a translation from his source language to an explicitly typed core language; a necessary implementation step for GHC.

Dunfield and Krishnaswami [2019] extend a bidirectional type system with indexes in existential types in order to support GADTs. As in this work, the introduction and elimination of existentials is implicit and determined by type annotations. Existentials are introduced via subsumption and eliminated via pattern matching. As a result, this type system has the same scoping limitations as one based on **unpack**.

In other contexts, if the domain of types that existentials are allowed to quantify over is restricted, more aggressive type inference is possible. For example, Tate et al. [2008] restrict existentials to hide only class types and develop a type-inference framework for a small object-oriented typed assembly language.

Module systems. This paper also relates to work on ML-style module systems. We do not summarize that field here but mention some papers that are particularly inspirational or relevant.

MacQueen [1986] noted the deficiencies of Mitchell and Plotkin [1988] with respect to expressing modular structure. This work proposed the original form of the ML module system as a dependent type system based on strong Σ -types. As in our system, modules support projections of the abstracted type and values. However, unlike this work, the ML module language supports additional type system features: a phase separation between the compile-time and runtime parts of the language, a treatment of generativity which determines when module expressions should and should not define new types, etc, as described in Harper and Pierce [2005]. We do not intend to use this type system to express modular structure.

F-ing modules [Rossberg et al. 2014] present a formalization of ML modules using existential types and a translation of a module language into System F_ω augmented with **pack** and **unpack**. Our approach is similar to theirs, in that we also use a translation of a surface language into our FX . However, because the ML module system includes a phase separation, our concerns about

strictness do not apply in that setting. As a result they can target the non-dependent language F_ω and use **unpack** as their elimination form. Rossberg [2015] extends the source language to a more uniform design while still retaining the translation to a non-dependent core calculus.

Montagu and Rémy [2009] present an extension of System F to compute *open* existential types. They introduce the idea of decomposing the usual explicit **pack** and **unpack** constructs of System F, and we were inspired by those ideas to design the type system of our implicit surface language with opened existentials. Interestingly, for a long time, it was unknown whether full abstraction could be achieved with strong existentials. Crary [2017] plugged this hole, proving Reynold’s abstraction theorem for a module calculus based on strong Σ -types.

11 CONCLUSION

By leveraging strong existential types, we have presented a type-inference algorithm that can infer introduction and elimination sites for existential packages. Users can freely create and consume existentials with no term-level annotations. The type annotation burden is small, and it dovetails with programmers’ current expectations around bidirectional type inference. The algorithm we present is designed to integrate well with GHC/Haskell’s state-of-the-art approach to type inference, the Quick Look algorithm [Serrano et al. 2020].

In order to prove our approach sound, we include an elaboration into a type-safe core language, inspired by Cardelli and Leroy [1990] and supporting the usual progress and preservation proofs. This core language is a small extension on System FC, the current core language implemented within GHC, and thus is suitable for implementation.

Beyond just soundness, we prove that inlining a **let**-binding preserves types, a non-trivial property in a type system with inferred existential types. We also prove that our type-inference algorithm is a conservative extension of a basic Hindley-Milner type system.

We believe and hope that our forthcoming implementation within GHC—in active development at the time of writing—will enable programmers to verify more aspects of their programs, even when that verification requires the use of existential types. We also hope that this new feature will provide a way forward to integrate the user-facing success of Liquid Haskell with GHC’s internal language and optimizer.

ACKNOWLEDGMENTS

The authors thank Neel Krishnaswami and Simon Peyton Jones for their collaboration and review, along with our anonymous reviewers. This material is based upon work supported by the National Science Foundation under Grant No. 1703835 and Grant No. 1704041. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- Lennart Augustsson. 1994. *Haskell B. user’s manual*. <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.50.5800&rep=rep1&type=pdf>
- Lennart Augustsson and Magnus Carlsson. 1999. An exercise in dependent types: A well-typed interpreter. (1999). <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.39.2895&rep=rep1&type=pdf> Unpublished manuscript.
- Gert-Jan Bottu and Richard A. Eisenberg. 2021. Seeking Stability by being Lazy and Shallow: Lazy and shallow instantiation is user friendly. In *ACM SIGPLAN Haskell Symposium*.
- Luca Cardelli and Xavier Leroy. 1990. Abstract types and the dot notation. In *IFIP TC2 working conference on programming concepts and methods*. North-Holland, 479–504.
- Dominique Clément, Thierry Despeyroux, Gilles Kahn, and Joëlle Despeyroux. 1986. A Simple Applicative Language: Mini-ML. In *Conference on LISP and Functional Programming* (Cambridge, Massachusetts, USA) (*LFP ’86*). ACM.

- Karl Cray. 2017. Modules, Abstraction, and Parametric Polymorphism. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages* (Paris, France) (POPL 2017). Association for Computing Machinery, New York, NY, USA, 100–113. <https://doi.org/10.1145/3009837.3009892>
- Luis Damas and Robin Milner. 1982. Principal Type-schemes for Functional Programs. In *Symposium on Principles of Programming Languages* (Albuquerque, New Mexico) (POPL '82). ACM.
- Atze Dijkstra. 2005. *Stepping through Haskell*. Ph.D. Dissertation. Universiteit Utrecht.
- Jana Dunfield and Neelakantan R. Krishnaswami. 2013. Complete and Easy Bidirectional Typechecking for Higher-rank Polymorphism. In *International Conference on Functional Programming (ICFP '13)*. ACM.
- Jana Dunfield and Neelakantan R. Krishnaswami. 2019. Sound and Complete Bidirectional Typechecking for Higher-Rank Polymorphism with Existentials and Indexed Types. *Proc. ACM Program. Lang.* 3, POPL, Article 9 (Jan. 2019), 28 pages. <https://doi.org/10.1145/3290322>
- Richard A. Eisenberg. 2016. *Dependent Types in Haskell: Theory and Practice*. Ph.D. Dissertation. University of Pennsylvania.
- Richard A. Eisenberg. 2020. Stitch: The Sound Type-Indexed Type Checker (Functional Pearl). In *Proceedings of the 13th ACM SIGPLAN International Symposium on Haskell* (Virtual Event, USA) (Haskell 2020). Association for Computing Machinery, New York, NY, USA, 39–53. <https://doi.org/10.1145/3406088.3409015>
- Richard A. Eisenberg, Joachim Breitner, and Simon Peyton Jones. 2018. Type Variables in Patterns. In *Proceedings of the 11th ACM SIGPLAN International Symposium on Haskell* (St. Louis, MO, USA) (Haskell 2018). Association for Computing Machinery, New York, NY, USA, 94–105. <https://doi.org/10.1145/3242744.3242753>
- Richard A. Eisenberg and Stephanie Weirich. 2012. Dependently Typed Programming with Singletons. In *ACM SIGPLAN Haskell Symposium*.
- Jean-Yves Girard. 1972. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Ph.D. Dissertation. Université Paris 7.
- Adam Gundry. 2013. *Type Inference, Haskell and Dependent Types*. Ph.D. Dissertation. University of Strathclyde.
- Cordelia V. Hall, Kevin Hammond, Simon L. Peyton Jones, and Philip L. Wadler. 1996. Type Classes in Haskell. *ACM Trans. Program. Lang. Syst.* 18, 2 (March 1996).
- Robert Harper and Benjamin C. Pierce. 2005. Design Considerations for ML-Style Module Systems. In *Advanced Topics in Types and Programming Languages*, Benjamin C. Pierce (Ed.). The MIT Press, 293–346.
- J. Roger Hindley. 1969. The Principal Type-Scheme of an Object in Combinatory Logic. *Trans. Amer. Math. Soc.* 146 (1969).
- William Alvin Howard. 1969. The Formulae-as-types Notion of Construction. (1969). <https://www.dcc.fc.up.pt/~acm/howard2.pdf> Dedicated to H. B. Curry on the occasion of his 80th birthday.
- Mark P. Jones. 1992. A Theory of Qualified Types. In *Proceedings of the 4th European Symposium on Programming (ESOP '92)*. Springer-Verlag, Berlin, Heidelberg, 287–306.
- Konstantin Läufer. 1996. Type classes with existential types. *Journal of Functional Programming* 6, 3 (1996), 485–518. <https://doi.org/10.1017/S0956796800001817>
- Konstantin Läufer and Martin Odersky. 1994. Polymorphic type inference and abstract data types. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 16, 5 (1994), 1411–1430.
- Didier Le Botlan and Didier Rémy. 2003. ML^F: Raising ML to the power of System F. In *International Conference on Functional Programming*. ACM.
- Daan Leijen. 2006. First-class polymorphism with existential types. (2006). Unpublished.
- David B MacQueen. 1986. Using dependent types to express modular structure. In *Proceedings of the 13th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*. 277–286.
- Per Martin-Löf. 1975. An intuitionistic theory of types: Predicative part. In *Studies in Logic and the Foundations of Mathematics*. Vol. 80. Elsevier, 73–118.
- Conor Thomas McBride. 2014. How to Keep Your Neighbours in Order. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming* (Gothenburg, Sweden) (ICFP '14). ACM.
- Robin Milner. 1978. A Theory of Type Polymorphism in Programming. *J. Comput. System Sci.* 17 (1978).
- John C Mitchell and Gordon D Plotkin. 1988. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 10, 3 (1988), 470–502.
- Stefan Monnier and David Haguenaer. 2010. Singleton types here, singleton types there, singleton types everywhere. In *Programming languages meets program verification (PLPV '10)*. ACM.
- Benoît Montagu and Didier Rémy. 2009. Modeling Abstract Types in Modules with Open Existential Types. In *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Savannah, GA, USA) (POPL '09). Association for Computing Machinery, New York, NY, USA, 354–365. <https://doi.org/10.1145/1480881.1480926>
- Koji Nakazawa and Makoto Tatsuta. 2009. Type Checking and Inference for Polymorphic and Existential Types. In *Proceedings of the Fifteenth Australasian Symposium on Computing: The Australasian Theory - Volume 94* (Wellington, New Zealand) (CATS '09). Australian Computer Society, Inc., AUS, 63–72.

- Koji Nakazawa, Makoto Tatsuta, Yukiyooshi Kameyama, and Hiroshi Nakano. 2008. Undecidability of Type-Checking in Domain-Free Typed Lambda-Calculi with Existence. In *Computer Science Logic*, Michael Kaminski and Simone Martini (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 478–492.
- Martin Odersky and Konstantin Läuffer. 1996. Putting Type Annotations to Work. In *Symposium on Principles of Programming Languages (POPL '96)*. ACM.
- Nigel Perry. 1991. *The implementation of practical functional programming languages*. Ph.D. Dissertation. Imperial College of Science, Technology and Medicine, University of London.
- Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. 2007. Practical type inference for arbitrary-rank types. *Journal of Functional Programming* 17, 1 (Jan. 2007).
- F. Pfenning. 1991. Unification and anti-unification in the calculus of constructions. In *Proceedings 1991 Sixth Annual IEEE Symposium on Logic in Computer Science*. IEEE Computer Society, Los Alamitos, CA, USA, 74,75,76,77,78,79,80,81,82,83,84,85. <https://doi.org/10.1109/LICS.1991.151632>
- Frank Pfenning and Peter Lee. 1989. LEAP: A language with eval and polymorphism. In *TAPSOFT '89*, J. Díaz and F. Orejas (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 345–359.
- Benjamin C. Pierce. 2002. *Types and Programming Languages*. MIT Press, Cambridge, MA.
- François Pottier and Didier Rémy. 2005. The Essence of ML Type Inference. In *Advanced Topics in Types and Programming Languages*, Benjamin C. Pierce (Ed.). The MIT Press, 387–489.
- John C. Reynolds. 1974. Towards a theory of type structure. In *Programming Symposium*, B. Robinet (Ed.). Lecture Notes in Computer Science, Vol. 19. Springer Berlin Heidelberg, 408–425.
- Andreas Rossberg. 2015. 1ML – Core and Modules United (F-Ing First-Class Modules). In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming (Vancouver, BC, Canada) (ICFP 2015)*. Association for Computing Machinery, New York, NY, USA, 35–47. <https://doi.org/10.1145/2784731.2784738>
- Andreas Rossberg, Claudio Russo, and Derek Dreyer. 2014. F-ing modules. *Journal of functional programming* 24, 5 (2014), 529–607.
- Alejandro Serrano, Jurriaan Hage, Simon Peyton Jones, and Dimitrios Vytiniotis. 2020. A Quick Look at Impredicativity. *Proc. ACM Program. Lang.* 4, ICFP, Article 89 (Aug. 2020), 29 pages. <https://doi.org/10.1145/3408971>
- Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones, and Kevin Donnelly. 2007. System F with type equality coercions. In *Types in languages design and implementation (Nice, Nice, France) (TLDI '07)*. ACM.
- Ross Tate, Juan Chen, and Chris Hawblitzel. 2008. *A Flexible Framework for Type Inference with Existential Quantification*. Technical Report MSR-TR-2008-184. <https://www.microsoft.com/en-us/research/publication/a-flexible-framework-for-type-inference-with-existential-quantification/>
- Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton Jones. 2014. Refinement Types for Haskell. In *International Conference on Functional Programming (Gothenburg, Sweden) (ICFP '14)*. ACM.
- Stephanie Weirich. 2018. Dependent Types in Haskell. Haskell eXchange keynote.
- Stephanie Weirich, Antoine Voizard, Pedro Henrique Azevedo de Amorim, and Richard A. Eisenberg. 2017. A Specification for Dependent Types in Haskell. *Proc. ACM Program. Lang.* 1, ICFP, Article 31 (Aug. 2017), 29 pages. <https://doi.org/10.1145/3110275>
- Hongwei Xi, Chiyang Chen, and Gang Chen. 2003. Guarded recursive datatype constructors. In *Principles of Programming Languages (New Orleans, Louisiana, USA) (POPL '03)*. ACM.