Elaboration with First-Class Implicit Function Types*

ANDRÁS KOVÁCS, Eötvös Loránd University, Hungary

Implicit functions are dependently typed functions, such that arguments are provided (by default) by inference machinery instead of programmers of the surface language. Implicit functions in Agda are an archetypal example. In the Haskell language as implemented by the Glasgow Haskell Compiler (GHC), polymorphic types are another example. Implicit function types are *first-class* if they are treated as any other type in the surface language. This holds in Agda and partially holds in GHC. Inference and elaboration in the presence of first-class implicit functions poses a challenge; in the context of Haskell and ML-like languages, this has been dubbed “impredicative instantiation” or “impredicative inference”. We propose a new solution for elaborating first-class implicit functions, which is applicable to full dependent type theories and compares favorably to prior solutions in terms of power, generality and simplicity. We build atop Norell’s bidirectional elaboration algorithm for Agda, and we note that the key issue is incomplete information about insertions of implicit abstractions and applications. We make it possible to track and refine information related to such insertions, by adding a function type to a core Martin-Löf type theory, which supports strict (definitional) currying. This allows us to represent undetermined domain arities of implicit function types, and we can decide at any point during elaboration whether implicit abstractions should be inserted.

CCS Concepts: • **Theory of computation** → **Type theory**; • **Software and its engineering** → **Polymorphism**.

Additional Key Words and Phrases: impredicative polymorphism, type theory, elaboration, type inference

ACM Reference Format:

András Kovács. 2020. Elaboration with First-Class Implicit Function Types. *Proc. ACM Program. Lang.* 4, ICFP, Article 101 (August 2020), 29 pages. <https://doi.org/10.1145/3408983>

1 INTRODUCTION

Programmers and users of proof assistants do not like to write out obvious things. Type inference and elaboration serve the purpose of filling in tedious details, translating terse surface-level languages to explicit core languages. Modern systems such as Agda have gotten quite adept at this task. However, in practice, programmers still have to tell the compiler where to try filling in details on its own.

Implicit function types are a common mechanism for conveying to the compiler that particular function arguments should be inferred by default. In Agda and Coq, one can use bracketed function domains for this purpose:

$id : \{A : \text{Set}\} \rightarrow A \rightarrow A$ Definition $id \{A : \text{Type}\}(x : A) := x.$
 $id \ x = x$

*The author was supported by the European Union, co-financed by the European Social Fund (EFOP-3.6.3-VEKOP-16-2017-00002).

Author’s address: András Kovács, Department of Programming Languages and Compilers, Eötvös Loránd University, Budapest, Hungary, kovacsandras@inf.elte.hu.

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

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/8-ART101

<https://doi.org/10.1145/3408983>

In GHC, one can use `forall` to define implicit function types¹

$$\begin{aligned} id &:: \text{forall } (a :: *) . a \rightarrow a \\ id \ x &= x \end{aligned}$$

In all of the above cases, if we apply `id` to an argument, the implicit type argument is provided by elaboration. For example, in Agda, `id true` is elaborated to `id {Bool} true`, and analogously in GHC and Coq. In all three systems, there is also a way to explicitly specify implicit arguments: in Agda we may put arguments in brackets as we have seen, in Coq we can prefix a name with `@` to make every implicit argument explicit, as in `@id bool true`, and in GHC we can enable the language extension `TypeApplications` [Eisenberg et al. 2016] and write `id @Bool True`.

Implicit function types are **first-class** if they can be manipulated like any other type. Coq is an example for a system where this is *not* the case. In Coq, the core language does not have an actual implicit function type, instead, implicitness is tied to particular *names*, and while we can write `list (forall {A : Type}, A → A)` for a list type with polymorphic elements, the brackets here are simply ignored by Coq. For example, Coq accepts the following definition:

$$\begin{aligned} \text{Definition } poly &: \text{forall } (f : \text{forall } \{A : \text{Type}\}, A \rightarrow A), \text{bool} * \text{nat} := \\ &\text{fun } f \Rightarrow (f \text{ bool true}, f \text{ nat } 0) \end{aligned}$$

This is a higher-rank polymorphic function which returns a pair. Note that `f` is applied to two arguments, because the implicitness in `forall {A : Type}, A → A` is silently dropped.

In GHC Haskell, `forall` types are more flexible. We can write the following, with `RankNTypes` enabled:

$$\begin{aligned} poly &:: (\text{forall } a . a \rightarrow a) \rightarrow (\text{Bool}, \text{Int}) \\ poly \ f &= (f \ \text{True}, f \ 0) \end{aligned}$$

However, polymorphic types are only supported in function domains and as fields of algebraic data constructors. We cannot instantiate an arbitrary type parameter to a `forall`, as in `[forall a. a → a]` for a list type with polymorphic elements. While this type is technically allowed by the `ImpredicativeTypes` language extension, as of GHC 8.8 this extension is deprecated and is not particularly usable in practice.

In Agda, implicit functions are truly a first-class notion, and we may have `List ({A : Set} → A → A)` without issues. However, Agda's elaboration still has limitations when it comes to handling implicit functions. Assume that we have `[]` for the empty list and `- :: -` for list extension, and consider the following code:

$$\begin{aligned} polyList &: \text{List } (\{A : \text{Set}\} \rightarrow A \rightarrow A) \\ polyList &= (\lambda x \rightarrow x) :: [] \end{aligned}$$

Agda 2.6.0.1 does not accept this. However, it does accept `polyList = (λ {A} x → x) :: []`. The issue is the following. Agda first infers a type for `(λ x → x) :: []`, then tries to unify the inferred type with the given `List ({A : Set} → A → A)` annotation. However, when Agda elaborates `λ x → x`, it does not yet know anything about the element type of the list; it is an undetermined unification variable. Hence, Agda does not know whether it should insert an extra `λ {A}` or not. If the element type is later found to be an implicit function, then it should, otherwise it should not. To solve this conundrum, Agda simply assumes that any unknown type is *not* an implicit function type, and

¹This notation requires language extensions `KindSignatures` and `RankNTypes`; one could also write the type `a → a` and GHC would silently insert the quantification.

elects to not insert a lambda. This assumption is often correct, but sometimes — as in the current case — it is not.

There is significant literature on type inference in the presence of first-class polymorphic types, mainly in relation to GHC and ML-like languages; see e.g. [Leijen 2008, 2009; Serrano et al. 2018; Vytiniotis et al. 2006]. The above issue in Agda is a specific instance of the challenges described in the mentioned works. Currently none of the above solutions are supported in production compilers, for reasons of complexity, fragility and interaction with other language features. A recent GHC development [Serrano et al. 2020] offers a solution which is relatively simple, and which is likely to land in an official GHC release. However, none of these works support dependent types, which is a key point in our work.

The solution presented in this paper is to gradually accumulate information about implicit insertions, and to have a setup where insertions can be refined and performed at any time after a particular expression is elaborated. In the current example, our algorithm wraps $\lambda x \rightarrow x$ in an implicit lambda with unknown arity, whose domain is later refined to be $A : \text{Set}$ when the inferred type is unified with the annotation.

1.1 Contributions

- We propose an elaboration algorithm which translates from a small Agda-like surface language to a small Martin-Löf type theory extended with implicit function types, telescopes and *strictly curried function types* with telescope domain. We use these extensions to accumulate information about implicit insertions. Our algorithm is based on Norell’s bidirectional elaborator for Agda [Norell 2007, Chapter 3].
- In the System F fragment, the presented elaborator is comparable or superior to previous solutions for impredicative inference. However, it also supports full dependent type theory. Our inference is also global, i.e. it can consider the whole program and not just particular n-ary applications.
- We provide an executable implementation of the elaborator described in this paper.
- Our solution is simple: we implemented elaboration, evaluation and unification in about 680 lines of Haskell, of which about 230 lines implement the novel enhancements on the top of Norell’s basic elaborator.
- The target theory of elaboration serves as a general platform for elaborating implicit function types: our concrete elaborator is a relatively simple one, and there is room to further develop it.

1.2 Note on Terminology

We prefer to avoid the term “impredicative inference” in order to avoid confusion with impredicativity in type theory. The two notions sometimes coincided historically, but currently they are largely orthogonal. In type theory, impredicativity is a property of a universe, i.e. closure of a universe under arbitrary products. In the type inference literature, impredicativity means the ability to instantiate type variables and metavariables to polymorphic types. In particular, we have that

- Agda has type-theory-predicative universes, but implements type-inference-impredicative elaboration with first-class implicit function types.
- Coq has type-theory-impredicative Prop universe (and optionally also Set), but implements type-inference-predicative elaboration, because of the lack of implicit function types.
- GHC is type-theory-impredicative with RankNTypes enabled and ImpredicativeTypes *disabled*, as we have $(\text{forall } (a :: *) . a \rightarrow a) :: *$.

2 BIDIRECTIONAL ELABORATION

First, we present a variant of Norell’s bidirectional elaborator [Norell 2007, Chapter 3]. Compared to *ibid.* we make some extensions and simplifications; what we end up with can be viewed as a toy version of the actual Agda elaborator. Our main goal with the elaborator is to have a setup which is as simple as possible, but sufficient to illustrate the enhancements in Section 5, which could be used to conservatively extend elaboration in existing dependently typed languages. Hence, we do not have more fine-grained control over metavariable insertion, nor the advanced unification algorithms used in Agda [Abel and Pientka 2011]; these are orthogonal improvements to the presented system.

$t, u, v, A, B, C ::=$	x	variable
	$(x : A) \rightarrow B$	function type
	$\{x : A\} \rightarrow B$	implicit function type
	$t u$	application
	$t \{u\}$	implicit application
	$\lambda x. t$	lambda abstraction
	$\lambda \{x\}. t$	implicit abstraction
	U	universe
	let $x : A = t$ in u	let-definition
	$_$	hole for inferred term

Fig. 1. Syntax of the surface language.

2.1 Surface Syntax

Figure 1 shows the possible constructs in the surface language. We only have terms, as we have Russell-style universe in the core, and we can conflate types and terms for convenience. The surface syntax does not have semantics or any well-formedness relations attached; its sole purpose is to serve as input to elaboration. Hence, the surface syntax can be also viewed as a small untyped tactic language which is interpreted by the elaborator.

The syntactic constructs are almost the same in the surface language as in the core syntax. The difference is that $_$ holes only appear in surface syntax. The $_$ can be used to request a term to be inferred by elaboration, the same way as in Agda. This can be used to give let-definitions without type annotation, as in **let** $x : _ = U$ **in** x .

2.2 Core Syntax

Figure 2 lists selected rules of the core language.

Metacontexts are used to record metavariables which are created during elaboration. In our case, metacontexts are simply a context prefix, and we have variables pointing into it. This corresponds to a particularly simple variant of *crisp type theory* [Licata et al. 2018], where we do not have modal type operators or functions with crisp (“meta”) domain. The non-meta typing context additionally supports *defined variables*, which is used in the typing rule for **let**-definitions, and we have that

$\boxed{\Theta \vdash}$	<i>metacontext formation</i>			
$\boxed{\Theta \Gamma \vdash}$	<i>context formation</i>			
$\boxed{\Theta \Gamma \vdash t : A}$	<i>typing</i>			
$\boxed{\Theta \Gamma \vdash t \equiv u : A}$	<i>term equality</i>			
METACON/EMPTY $\frac{}{\bullet \vdash}$	METACON/BIND $\frac{\Theta \vdash \quad \Theta \bullet \vdash A : U}{\Theta, \alpha : A \vdash}$	CON/EMPTY $\frac{\Theta \vdash}{\Theta \bullet \vdash}$	CON/BIND $\frac{\Theta \Gamma \vdash \quad \Theta \Gamma \vdash A : U}{\Theta \Gamma, x : A \vdash}$	
		CON/DEFINE $\frac{\Theta \Gamma \vdash \quad \Theta \Gamma \vdash t : A}{\Theta \Gamma, x : A = t \vdash}$	METAVAR $\frac{}{\Theta_0, \alpha : A, \Theta_1 \Gamma \vdash \alpha : A}$	BOUND-VAR $\frac{}{\Theta \Gamma, x : A, \Delta \vdash x : A}$
DEFINED-VAR $\frac{}{\Theta \Gamma, x : A = t, \Delta \vdash x : A}$	UNIVERSE $\frac{}{\Theta \Gamma \vdash U : U}$	LET $\frac{\Theta \Gamma \vdash t : A \quad \Theta \Gamma, x : A = t \vdash u : B}{\Theta \Gamma \vdash \mathbf{let} x : A = t \mathbf{in} u : B[x \mapsto t]}$		
FUN $\frac{\Theta \Gamma \vdash A : U \quad \Theta \Gamma, x : A \vdash B : U}{\Theta \Gamma \vdash (x : A) \rightarrow B : U}$		IMPLICIT-FUN $\frac{\Theta \Gamma \vdash A : U \quad \Theta \Gamma, x : A \vdash B : U}{\Theta \Gamma \vdash \{x : A\} \rightarrow B : U}$		
APP $\frac{\Theta \Gamma \vdash t : (x : A) \rightarrow B \quad \Theta \Gamma \vdash u : A}{\Theta \Gamma \vdash t u : B[x \mapsto u]}$		IMPLICIT-APP $\frac{\Theta \Gamma \vdash t : \{x : A\} \rightarrow B \quad \Theta \Gamma \vdash u : A}{\Theta \Gamma \vdash t \{u\} : B[x \mapsto u]}$		
LAM $\frac{\Theta \Gamma, x : A \vdash t : B}{\Theta \Gamma \vdash \lambda x. t : (x : A) \rightarrow B}$		IMPLICIT-LAM $\frac{\Theta \Gamma, x : A \vdash t : B}{\Theta \Gamma \vdash \lambda \{x\}. t : \{x : A\} \rightarrow B}$		
FUN-β $\frac{\Theta \Gamma, x : A \vdash t : B \quad \Theta \Gamma \vdash u : A}{\Theta \Gamma \vdash (\lambda x. t) u \equiv t[x \mapsto u] : B[x \mapsto u]}$		IMPLICIT-FUN-β $\frac{\Theta \Gamma, x : A \vdash t : B \quad \Theta \Gamma \vdash u : A}{\Theta \Gamma \vdash (\lambda \{x\}. t) \{u\} \equiv t[x \mapsto u] : B[x \mapsto u]}$		
FUN-η $\frac{\Theta \Gamma \vdash t : (x : A) \rightarrow B}{\Theta \Gamma \vdash (\lambda x. t x) \equiv t : (x : A) \rightarrow B}$		IMPLICIT-FUN-η $\frac{\Theta \Gamma \vdash t : \{x : A\} \rightarrow B}{\Theta \Gamma \vdash (\lambda \{x\}. t \{x\}) \equiv t : \{x : A\} \rightarrow B}$		
DEFINITION $\frac{}{\Theta \Gamma, x : A = t, \Delta \vdash x \equiv t : A}$				

Fig. 2. Selected rules of the core language.

any defined variable is equal to its definition. We support **let** as a convenience feature which is important for defining more involved example programs².

The universe U is Russell-style, and we have the type-in-type rule. This causes our core syntax to be non-total, and our elaboration algorithm to be possibly non-terminating. Consistent universe setups are straightforward to devise, but they are orthogonal to the focus of this work, and we choose to avoid needless noise.

Function types only differ from each other in notation: implicit functions have the same rules as “explicit” functions. The primary purpose of implicit function types is to *guide elaboration*: the elaborator will at times compute a type and branch on whether it is an implicit function type.

We avoid a fully formal presentation in this paper. For examples of what such a presentation looks like, see [Abel et al. 2018] or [Wieczorek and Biernacki 2018]. Some notes on what is elided:

- We use nameful notation and implicit weakening, i.e. whenever a term is well-formed in some context, it is assumed to be well-formed (as it is) in extended contexts. We also assume that any specifically mentioned name is fresh, e.g. when we write $\Theta, \alpha : A$, we assume that α is fresh in Θ . Formally, we would use de Bruijn indices for variables, and define variable renaming and parallel substitution by recursion on presyntax, e.g. as in [Schäfer et al. 2015].
- Fixing any Θ metacontext, parallel substitutions of bound and defined variables form morphisms of a category, where the identity substitution id maps each variable to itself and composition $- \circ -$ is given by pointwise substitution. The action of parallel substitution on terms is functorial, i.e. $t[\sigma \circ \delta] \equiv t[\sigma][\delta]$ and $t[\text{id}] \equiv t$, and typing is stable under substitution.
- Definitional equality is understood to be a congruence and an equivalence relation, which is respected by substitution and typing.
- We elide a number of well-formedness assumptions in rules. For instance, whenever a context appears in a rule, it is assumed to be well-formed. Likewise, whenever we have $\Theta | \Gamma \vdash t : A$, we assume that $\Theta | \Gamma \vdash A : U$.

From now on, we will only consider well-formed core syntax, and unless otherwise mentioned, presented constructions on the core syntax respect definitional equality.

Alternatively, one could present the syntax as a generalized algebraic theory [Sterling 2019] or a quotient inductive-inductive type [Altenkirch and Kaposi 2016], in which case we would get congruences and quotienting for free, and we would also get a rich model theory for our syntax. However, it seems that there are a number of possible choices for giving an algebraic presentation of metacontexts, and existing works on algebraic presentations of dependent modal contexts (e.g. [Birkedal et al. 2020]) do not precisely cover the current use case. We leave this to future work, along with the investigation of elaboration from an algebraic perspective.

Notations. We use Agda-like syntactic sugar both in the surface syntax and in core syntax.

- We use $A \rightarrow B$ to refer to non-dependent functions.
- We group domain types together in functions, and omit function arrows, as in $\{A B : U\}(x : A) \rightarrow B \rightarrow A$.
- We group multiple λ -s, as in $\lambda \{A\} \{B\} x y. x$.

Definition 2.1 (Spines). We use a spine notation for neutral terms. A spine is a list of terms, noted as $\bar{t}$, where terms may be wrapped in brackets to signal implicit application. For example, if $\bar{u} \equiv (\{A\}, \{B\}, x)$, then $t \bar{u}$ denotes $t \{A\} \{B\} x$. In $t \bar{u}$, we call t the *head* of the neutral term. In particular, if t is a metavariable, the neutral term is *meta-headed*.

²In presentations of dependent type theories without explicit substitutions, the **let** rule is not derivable from function application, unlike in simple type theories.

$$\boxed{\theta : \Theta_0 \Rightarrow \Theta_1} \quad \text{metasubstitution formation}$$

$$\begin{array}{c}
\text{METASUB/EMPTY} \\
\frac{\Theta \vdash}{() : \Theta \Rightarrow \bullet}
\end{array}
\quad
\begin{array}{c}
\text{METASUB/EXTENDED} \\
\frac{\theta : \Theta_0 \Rightarrow \Theta_1 \quad \Theta_0 | \bullet \vdash t : A[\theta]}{(\theta, \alpha \mapsto t) : \Theta_0 \Rightarrow (\Theta_1, \alpha : A)}
\end{array}
\quad
\begin{array}{c}
\text{METASUB/CON-ACTION} \\
\frac{\theta : \Theta_0 \Rightarrow \Theta_1 \quad \Theta_1 | \Gamma \vdash}{\Theta_0 | \Gamma[\theta] \vdash}
\end{array}$$

$$\begin{array}{c}
\text{METASUB/TM-ACTION} \\
\frac{\theta : \Theta_0 \Rightarrow \Theta_1 \quad \Theta_1 | \Gamma \vdash t : A}{\Theta_0 | \Gamma[\theta] \vdash t[\theta] : A[\theta]}
\end{array}
\quad
\begin{array}{c}
\text{METASUB/IDENTITY} \\
\frac{}{\text{id} : \Theta \Rightarrow \Theta}
\end{array}
\quad
\begin{array}{c}
\text{METASUB/COMPOSITION} \\
\frac{\theta_0 : \Theta_1 \Rightarrow \Theta_2 \quad \theta_1 : \Theta_0 \Rightarrow \Theta_1}{\theta_0 \circ \theta_1 : \Theta_0 \Rightarrow \Theta_2}
\end{array}$$

$$\begin{array}{c}
\text{METASUB/WEAKENING} \\
\frac{}{p : (\Theta, \alpha : A) \Rightarrow \Theta}
\end{array}$$

Fig. 3. Admissible rules for metasubstitutions.

Example 2.2. The core syntax is quite expressive as a programming language, thanks to **let**-definitions and the type-in-type rule which allows Church encodings for a large class of inductive types. For example, the following term computes a list of types by mapping:

```

let List : U → U
  = λ A. (L : U) → (A → L → L) → L → L in
let map : {AB : U} → (A → B) → List A → List B
  = λ {A} {B} f as L cons nil. as L (λ a. cons (f a)) nil in
  map {U} {U} (λ A. A → A) (λ L cons nil. cons U (cons U nil))

```

2.3 Metasubstitutions

Before we can move on to the description of the elaborator, we need to specify metasubstitutions. These are essentially just parallel substitutions of metacontexts, and their purpose is to keep track of meta-operations (e.g. fresh meta creation or solution of a meta).

- A metasubstitution $\theta : \Theta_0 \Rightarrow \Theta_1$ assigns to each variable in Θ_1 a term in Θ_0 , hence it is represented as a list of terms $(\alpha_1 \mapsto t_1, \dots, \alpha_i \mapsto t_i)$.
- We define the action of a metasubstitution on contexts and terms by recursion; we write action on contexts as $\Gamma[\theta]$ and action on terms as $t[\theta]$. We remark that there is no abstraction for metavariables in the core syntax, so we do not have to handle variable capture (or index shifting).
- We list some of the admissible rules in Figure 3, which are relevant in the rest of the paper.

The identity substitution id maps each variable to itself. Composition is given by pointwise term substitution, id and $- \circ -$ yields a category, and the action of metasubstitution on contexts and terms is functorial. The weakening substitution p (the naming comes from categories-with-families terminology [Dybjer 1995]) can be defined as dropping the last entry from $\text{id} : (\Theta, \alpha : A) \Rightarrow (\Theta, \alpha : A)$.

2.4 Fresh Metavariables

Using *contextual metavariables* is a standard practice in the implementation of dependently typed languages. This means that every “hole” in the surface language is represented as an unknown function which abstracts over all bound variables in the scope of a hole. Unlike [Nanevski et al. \[2008\]](#) and similarly to [Gundry \[2013\]](#), we do not have a first-class notion of contextual types, and instead reuse the standard dependent function type to abstract over enclosing contexts.

Definition 2.3 (Closing type). For each $\Theta | \Gamma \vdash A : \mathcal{U}$, we define a term $\Gamma \Rightarrow A$ by recursion on Γ , such that $\Theta | \bullet \vdash \Gamma \Rightarrow A : \mathcal{U}$.

$$\begin{aligned} ((\Gamma, x : A) \Rightarrow B) & \quad \equiv (\Gamma \Rightarrow ((x : A) \rightarrow B)) \\ ((\Gamma, x : A = t) \Rightarrow B) & \quad \equiv (\Gamma \Rightarrow B[x \mapsto t]) \\ (\bullet \Rightarrow B) & \quad \equiv B \end{aligned}$$

Definition 2.4 (Contextualization). For each $\Theta | \bullet \vdash t : \Gamma \Rightarrow A$, we define the spine $\overline{\text{vars}}_\Gamma$ such that that $\Theta | \Gamma \vdash t \overline{\text{vars}}_\Gamma : A$, which is t applied to all bound variables in Γ .

$$\begin{aligned} (t \overline{\text{vars}}_{\Gamma, x:B}) & \quad \equiv (t \overline{\text{vars}}_\Gamma) x \\ (t \overline{\text{vars}}_{\Gamma, x:B=u}) & \quad \equiv (t \overline{\text{vars}}_\Gamma) \\ (t \overline{\text{vars}}_\bullet) & \quad \equiv t \end{aligned}$$

Example 2.5. If we have $\Gamma \equiv (\bullet, A : \mathcal{U}, B : A \rightarrow \mathcal{U})$, then $(\Gamma \Rightarrow \mathcal{U}) \equiv ((A : \mathcal{U})(B : A \rightarrow \mathcal{U}) \rightarrow \mathcal{U})$ and $t \overline{\text{vars}}_\Gamma \equiv t A B$.

Definition 2.6 (Fresh meta creation). We define $\text{freshMeta}_{\Theta | \Gamma} A$ as follows:

$$\frac{\Theta | \Gamma \vdash A : \mathcal{U}}{\text{freshMeta}_{\Theta | \Gamma} A \in \{(\Theta', \theta, t) \mid (\theta : \Theta' \Rightarrow \Theta) \wedge (\Theta' | \Gamma[\theta] \vdash t : A[\theta])\}} \\ \text{freshMeta}_{\Theta | \Gamma} A \equiv ((\Theta, \alpha : \Gamma \Rightarrow A), p, \alpha \overline{\text{vars}}_\Gamma)$$

In the definition above, we extend Θ with a fresh meta, which has the closing type $\Gamma \Rightarrow A$. The p weakening relates the new metacontext to the old one, by “dropping” the new entry. Lastly, $\alpha \overline{\text{vars}}_\Gamma$ is the fresh meta applied to all bound variables.

2.5 Unification

We assume that there is a unification procedure, which returns a unifying metasubstitution on success. We only have *homogeneous* unification, i.e. the two terms to be unified must have the same type. The specification is as follows:

$$\frac{\Theta | \Gamma \vdash t : A \quad \Theta | \Gamma \vdash u : A}{\text{unify } t u \in \{(\Theta', \theta) \mid (\theta : \Theta' \Rightarrow \Theta) \wedge (\Theta' | \Gamma[\theta] \vdash t[\theta] \equiv u[\theta] : A[\theta])\} \cup \{\text{fail}\}}$$

For a simple example, assuming $\Theta \equiv (\bullet, \alpha : \mathcal{U}, \beta : \mathcal{U})$, $\text{unify } \alpha (\beta \rightarrow \beta)$ yields $\Theta' \equiv (\bullet, \beta : \mathcal{U})$ and the substitution $\theta \equiv (\alpha \mapsto (\beta \rightarrow \beta), \beta \mapsto \beta)$, where $\theta : \Theta' \Rightarrow \Theta$.

Here, we do not require that unification returns most general unifiers, nor do we go into the details of how unification is implemented. [Gundry \[2013, Chapter 4\]](#) describes unification in detail for a similar syntax, with a similar (though more featureful) setup for metacontexts. See also [\[Abel and Pientka 2011\]](#) for a reference on unification. Note that our unification algorithm does not support *constraint postponing*, as we have not talked about constraints at all. In our concrete prototype implementation, unification supports basic pattern unification and metavariable pruning.

2.6 Elaboration

Elaboration consists of two (partial) functions, checking and inferring, which are defined by mutual induction on surface syntax. We also have helper functions for implicit argument insertion, which can be used as a post-processing step after inference. First, about the used notations:

- We use a Haskell-like monadic pseudocode notation, where the side effect is failure via `fail`.
- We use pattern matching notation on core terms; e.g. we may match on whether a type is a function type. This assumes an evaluation/normalization procedure on core terms; but note that we already assume this in unification. Patterns are matched in top-down order.
- We abbreviate $\theta_1 \circ \theta_2$ as θ_{12} and $\theta_1 \circ \theta_2 \circ \theta_3$ as θ_{123} and analogously in other cases. We do this to reduce the visual noise caused by threading composed metasubstitutions everywhere in the elaboration algorithm.

We present the specifications and definitions below, then we describe them in order.

INSERTION

$$\frac{\Theta_0 | \Gamma \vdash (\Theta, \theta, t, A) \in \{(\Theta, \theta, t, A) \mid (\theta : \Theta \Rightarrow \Theta_0) \wedge (\Theta | \Gamma[\theta] \vdash t : A)\} \cup \{\text{fail}\}}{\begin{array}{l} \text{insert}'(\Theta, \theta, t, A) \in \{(\Theta, \theta, t, A) \mid (\theta : \Theta \Rightarrow \Theta_0) \wedge (\Theta | \Gamma[\theta] \vdash t : A)\} \cup \{\text{fail}\} \\ \text{insert}(\Theta, \theta, t, A) \in \{(\Theta, \theta, t, A) \mid (\theta : \Theta \Rightarrow \Theta_0) \wedge (\Theta | \Gamma[\theta] \vdash t : A)\} \cup \{\text{fail}\} \end{array}}$$

$$\begin{aligned} \text{insert}'(\Theta_1, \theta_1, t, \{x : A\} \rightarrow B) &\equiv \mathbf{let} (\Theta_2, \theta_2, u) = \text{freshMeta}_{\Theta_1 | \Gamma[\theta_1]} A \\ &\quad \mathbf{in} \text{insert}'(\Theta_2, \theta_{12}, (t[\theta_2]) \{u\}, B[\theta_2][x \mapsto u]) \\ \text{insert}'(\Theta_1, \theta_1, t, A) &\equiv \mathbf{return} (\Theta_1, \theta_1, t, A) \\ \text{insert}' \text{ fail} &\equiv \text{fail} \\ \text{insert}(\Theta, \theta, (\lambda\{x\}.t), A) &\equiv \mathbf{return} (\Theta, \theta, (\lambda\{x\}.t), A) \\ \text{insert}(\Theta, \theta, t, A) &\equiv \text{insert}'(\Theta, \theta, t, A) \\ \text{insert fail} &\equiv \text{fail} \end{aligned}$$

CHECK

$$\frac{t \text{ is a surface expression} \quad \Theta | \Gamma \vdash A : \mathbf{U}}{\llbracket t \rrbracket \Downarrow_{\Theta | \Gamma} A \in \{(\Theta', \theta, t') \mid (\theta : \Theta' \Rightarrow \Theta) \wedge (\Theta' | \Gamma[\theta] \vdash t' : A[\theta])\} \cup \{\text{fail}\}}$$

INFER

$$\frac{t \text{ is a surface expression} \quad \Theta | \Gamma \vdash}{\llbracket t \rrbracket \Uparrow_{\Theta | \Gamma} \in \{(\Theta', \theta, t', A) \mid (\theta : \Theta' \Rightarrow \Theta) \wedge (\Theta' | \Gamma[\theta] \vdash t' : A)\} \cup \{\text{fail}\}}$$

$$\begin{aligned} \llbracket \lambda x. t \rrbracket \Downarrow_{\Theta | \Gamma} ((x : A) \rightarrow B) &\equiv \mathbf{do} \\ &\quad (\Theta', \theta, t') \leftarrow \llbracket t \rrbracket \Downarrow_{\Theta | \Gamma, x:A} B \\ &\quad \mathbf{return} (\Theta', \theta, \lambda x. t') \\ \llbracket \lambda \{x\}. t \rrbracket \Downarrow_{\Theta | \Gamma} (\{x : A\} \rightarrow B) &\equiv \mathbf{do} \\ &\quad (\Theta', \theta, t') \leftarrow \llbracket t \rrbracket \Downarrow_{\Theta | \Gamma, x:A} B \\ &\quad \mathbf{return} (\Theta', \theta, \lambda \{x\}. t') \\ \llbracket t \rrbracket \Downarrow_{\Theta | \Gamma} (\{x : A\} \rightarrow B) &\equiv \mathbf{do} \\ &\quad (\Theta', \theta, t') \leftarrow \llbracket t \rrbracket \Downarrow_{\Theta | \Gamma, x:A} B \end{aligned}$$

```

return ( $\Theta', \theta, \lambda \{x\}. t'$ )
 $\llbracket \text{let } x : A = t \text{ in } u \rrbracket \Downarrow_{\Theta_0 | \Gamma} B \equiv \text{do}$ 
  ( $\Theta_1, \theta_1, A'$ )  $\leftarrow \llbracket A \rrbracket \Downarrow_{\Theta_0 | \Gamma} \cup$ 
  ( $\Theta_2, \theta_2, t'$ )  $\leftarrow \llbracket t \rrbracket \Downarrow_{\Theta_1 | \Gamma[\theta_1]} A'$ 
  ( $\Theta_3, \theta_3, u'$ )  $\leftarrow \llbracket u \rrbracket \Downarrow_{\Theta_2 | \Gamma[\theta_{12}], x:A'[\theta_2]=t'} (B[\theta_{12}])$ 
  return ( $\Theta_3, \theta_{123}, \text{let } x : A'[\theta_{23}] = t'[\theta_3] \text{ in } u'$ )
 $\llbracket \_ \rrbracket \Downarrow_{\Theta | \Gamma} A \equiv \text{do}$ 
  return ( $\text{freshMeta}_{\Theta | \Gamma} A$ )
 $\llbracket t \rrbracket \Downarrow_{\Theta_0 | \Gamma} A \equiv \text{do}$ 
  ( $\Theta_1, \theta_1, t', B$ )  $\leftarrow \text{insert} (\llbracket t \rrbracket \Uparrow_{\Theta_0 | \Gamma})$ 
  ( $\Theta_2, \theta_2$ )  $\leftarrow \text{unify} (A[\theta_1]) B$ 
  return ( $\Theta_2, \theta_{12}, t'[\theta_2]$ )

 $\llbracket x \rrbracket \Uparrow_{\Theta | \Gamma} \equiv \text{do}$ 
  if ( $\Gamma \equiv (\Gamma_0, x : A, \Gamma_1)$ )  $\vee$  ( $\Gamma \equiv (\Gamma_0, x : A = t, \Gamma_1)$ )
    then return ( $\Theta, \text{id}, x, A$ )
    else fail
 $\llbracket U \rrbracket \Uparrow_{\Theta | \Gamma} \equiv \text{do}$ 
  return ( $\Theta, \text{id}, U, U$ )
 $\llbracket (x : A) \rightarrow B \rrbracket \Uparrow_{\Theta_0 | \Gamma} \equiv \text{do}$ 
  ( $\Theta_1, \theta_1, A'$ )  $\leftarrow \llbracket A \rrbracket \Downarrow_{\Theta_0 | \Gamma} \cup$ 
  ( $\Theta_2, \theta_2, B'$ )  $\leftarrow \llbracket B \rrbracket \Downarrow_{\Theta_1 | \Gamma[\theta_1], x:A'} \cup$ 
  return ( $\Theta_2, \theta_{12}, ((x : A'[\theta_2]) \rightarrow B'), U$ )
 $\llbracket \{x : A\} \rightarrow B \rrbracket \Uparrow_{\Theta_0 | \Gamma} \equiv \text{do}$ 
  ( $\Theta_1, \theta_1, A'$ )  $\leftarrow \llbracket A \rrbracket \Downarrow_{\Theta_1 | \Gamma} \cup$ 
  ( $\Theta_2, \theta_2, B'$ )  $\leftarrow \llbracket B \rrbracket \Downarrow_{\Theta_2 | \Gamma[\theta_1], x:A'} \cup$ 
  return ( $\Theta_2, \theta_{12}, (\{x : A'[\theta_2]\} \rightarrow B'), U$ )
 $\llbracket t u \rrbracket \Uparrow_{\Theta_0 | \Gamma} \equiv \text{do}$ 
  ( $\Theta_1, \theta_1, t', A$ )  $\leftarrow \text{insert}' (\llbracket t \rrbracket \Uparrow_{\Theta_0 | \Gamma})$ 
  let ( $\Theta_2, \theta_2, A_0$ ) =  $\text{freshMeta}_{\Theta_1 | \Gamma[\theta_1]} U$ 
  let ( $\Theta_3, \theta_3, A_1$ ) =  $\text{freshMeta}_{\Theta_2 | \Gamma[\theta_{12}], x:A_0} U$ 
  ( $\Theta_4, \theta_4$ )  $\leftarrow \text{unify} (A[\theta_{23}]) ((x : A_0[\theta_3]) \rightarrow A_1)$ 
  ( $\Theta_5, \theta_5, u'$ )  $\leftarrow \llbracket u \rrbracket \Downarrow_{\Theta_4 | \Gamma[\theta_{1234}]} (A_0[\theta_{34}])$ 
  return ( $\Theta_5, \theta_{12345}, (t'[\theta_{2345}]) u', A_1[\theta_{45}][x \mapsto u']$ )
 $\llbracket t \{u\} \rrbracket \Uparrow_{\Theta_0 | \Gamma} \equiv \text{do}$ 
  ( $\Theta_1, \theta_1, t', A$ )  $\leftarrow \llbracket t \rrbracket \Uparrow_{\Theta_0 | \Gamma}$ 
  let ( $\Theta_2, \theta_2, A_0$ ) =  $\text{freshMeta}_{\Theta_1 | \Gamma[\theta_1]} U$ 

```

```

let ( $\Theta_3, \theta_3, A_1$ ) = freshMeta  $\Theta_2 | \Gamma[\theta_{12}], x:A_0 \cup$ 
( $\Theta_4, \theta_4$ )  $\leftarrow$  unify ( $A[\theta_{23}]$ ) ( $\{x : A_0[\theta_3]\} \rightarrow A_1$ )
( $\Theta_5, \theta_5, u'$ )  $\leftarrow$   $\llbracket u \rrbracket \Downarrow_{\Theta_4 | \Gamma[\theta_{1234}]} (A_0[\theta_{34}])$ 
return ( $\Theta_5, \theta_{12345}, (t'[\theta_{2345}]) \{u'\}, A_1[\theta_{45}][x \mapsto u']$ )

 $\llbracket \lambda x. t \rrbracket \Uparrow_{\Theta_0 | \Gamma} \equiv$  do
  let ( $\Theta_1, \theta_1, A$ ) = freshMeta  $\Theta_0 | \Gamma \cup$ 
  ( $\Theta_2, \theta_2, t', B$ )  $\leftarrow$  insert ( $\llbracket t \rrbracket \Uparrow_{\Theta_1 | \Gamma[\theta_1], x:A}$ )
  return ( $\Theta_2, \theta_{12}, \lambda x. t', (x : A[\theta_2]) \rightarrow B$ )

 $\llbracket \lambda \{x\}. t \rrbracket \Uparrow_{\Theta_0 | \Gamma} \equiv$  do
  let ( $\Theta_1, \theta_1, A$ ) = freshMeta  $\Theta_0 | \Gamma \cup$ 
  ( $\Theta_2, \theta_2, t', B$ )  $\leftarrow$  insert ( $\llbracket t \rrbracket \Uparrow_{\Theta_1 | \Gamma[\theta_1], x:A}$ )
  return ( $\Theta_2, \theta_{12}, \lambda \{x\}. t', \{x : A[\theta_2]\} \rightarrow B$ )

 $\llbracket \text{let } x : A = t \text{ in } u \rrbracket \Uparrow_{\Theta_0 | \Gamma} \equiv$  do
  ( $\Theta_1, \theta_1, A'$ )  $\leftarrow$   $\llbracket A \rrbracket \Downarrow_{\Theta_0 | \Gamma} \cup$ 
  ( $\Theta_2, \theta_2, t'$ )  $\leftarrow$   $\llbracket t \rrbracket \Downarrow_{\Theta_1 | \Gamma[\theta_1]} A'$ 
  ( $\Theta_3, \theta_3, u', B$ )  $\leftarrow$   $\llbracket u \rrbracket \Uparrow_{\Theta_2 | \Gamma[\theta_{12}], x:A'[\theta_2] = t'}$ 
  return ( $\Theta_3, \theta_{123}, (\text{let } x : A'[\theta_{23}] = t'[\theta_3] \text{ in } u'), B$ )

 $\llbracket \_ \rrbracket \Uparrow_{\Theta | \Gamma} \equiv$  do
  let ( $\Theta', \theta, A$ ) = freshMeta  $\Theta | \Gamma \cup$ 
  return (freshMeta  $\Theta' | \Gamma[\theta] A$ )

```

2.6.1 Implicit Argument Insertion. We define two functions, **insert'** and **insert** which insert implicit applications around a core term. For example, if we have a defined name *id* with type $\{A : \mathcal{U}\} \rightarrow A \rightarrow A$ in a surface program, we usually want to expand *id* occurrences to *id* $\{\alpha\}$, where α is a fresh metavariable. We define the insertion functions to take as input the output of $\llbracket - \rrbracket \Uparrow$, so that it is more convenient to use as an optional post-processing step after inference.

The difference between **insert** and **insert'** is that the former avoids insertion to implicit λ -terms. We usually want to avoid that, as it creates a new implicit β -redex where the fresh metavariable is often under-constrained. We only create such β -redex when the surface program explicitly asks for it; we shall shortly see this in the discussion of the $\llbracket - \rrbracket \Uparrow$ case for explicit function application. With this, we largely follow the insertion behavior of Agda.

2.6.2 Checking. The first two clauses are checking λ -s, where the expected type exactly matches the λ binders. Hence, we simply check under binders with $\llbracket t \rrbracket \Downarrow_{\Theta | (\Gamma, x:A)} B$, and wrap the resulting term in the appropriate (implicit or explicit) λ .

The third clause for $\llbracket t \rrbracket \Downarrow_{\Theta | \Gamma} (\{x : A\} \rightarrow B)$ is more interesting. Here, we are checking a surface term which is *not* a λ (this follows from our top-down pattern matching notation), with an implicit function expected type. Here, we check *t* in the extended $\Gamma, x : A$ context, and we insert a new implicit λ in the elaboration output. This is the only point where implicit λ -s are introduced by elaboration. Practically, this rule is commonly useful whenever we have a higher-order function where some arguments have implicit function type. For example, in the surface syntax, assume

natural numbers, and an induction principle for them:

$$\text{NatInd} : (P : \text{Nat} \rightarrow \mathcal{U}) \rightarrow P \text{ zero} \rightarrow (\{n : \text{Nat}\} \rightarrow P n \rightarrow P (\text{suc } n)) \rightarrow (n : \text{Nat}) \rightarrow P n$$

Then, define addition using induction:

$$\begin{aligned} \text{let NatPlus} : \text{Nat} \rightarrow \text{Nat} \rightarrow \text{Nat} \\ = \text{NatInd } (\lambda n. \text{Nat} \rightarrow \text{Nat}) (\lambda m. m) (\lambda f m. \text{suc } (f m)) \text{ in } \dots \end{aligned}$$

When the above is elaborated, the $\lambda f m. \text{suc } (f m)$ function is checked with the expected type $\{n : \text{Nat}\} \rightarrow (\text{Nat} \rightarrow \text{Nat}) \rightarrow (\text{Nat} \rightarrow \text{Nat})$, and the elaboration output is $\lambda \{n\} f m. \text{suc } (f m)$. Hence, in this case we do not have to write implicit λ in the surface syntax.

For $\llbracket \text{let } x : A = t \text{ in } u \rrbracket \Downarrow_{\Theta_0|\Gamma} B$, we simply let checking fall through. For $\llbracket _ \rrbracket \Downarrow_{\Theta|\Gamma} A$, we return a fresh metavariable with the expected type. In any other $\llbracket t \rrbracket \Downarrow_{\Theta_0|\Gamma} A$ case, we have a *change of direction*: we infer a type for t (with implicit insertions) then unify the expected and inferred types.

2.6.3 Inferring. For $\llbracket x \rrbracket \Uparrow_{\Theta|\Gamma}$, we look up the type of x in Γ . In the case of $\mathcal{U}$, we always succeed and infer $\mathcal{U}$ as type. In the cases for function types, we check that the domains and codomains have type $\mathcal{U}$.

The $\llbracket t u \rrbracket \Uparrow_{\Theta_0|\Gamma}$ and $\llbracket t \{u\} \rrbracket \Uparrow_{\Theta_0|\Gamma}$ cases are again interesting. Here, we first infer a type for t , then refine the type to a function type, and lastly check the argument with the domain type. Note the difference between the explicit and implicit case. In the former case we use $\text{insert}' (\llbracket t \rrbracket \Uparrow_{\Theta_0|\Gamma})$, which inserts implicit applications. In the latter case we do no insertion. This ensures that implicit applications in the surface syntax behave similarly as in Agda. For example, given $\text{id} : \{A : \mathcal{U}\} \rightarrow A \rightarrow A$ in scope, we elaborate $\text{id } U$ as follows:

- (1) The expression is an explicit application, so we infer id and insert implicit arguments, returning $\text{id } \{\alpha\}$, where α is a fresh meta.
- (2) We check that U has type α . Here we immediately change direction, inferring U as type for U and unifying α with U .
- (3) Hence, the resulting output is $\text{id } \{U\} U$.

On the other hand, we elaborate $\text{id } \{U\}$ as follows:

- (1) This is an implicit application, so we infer a type for id without inserting implicit arguments. This yields the inferred type $\{A : \mathcal{U}\} \rightarrow A \rightarrow A$, and we check that U has type $\mathcal{U}$.
- (2) We change direction and infer U as type for U , and successfully unify U with U .

Also note that we use insert' here instead of insert . Thus, we get expected (Agda-like) insertions with β -redexes as well. For example $(\lambda \{A\} x. \text{id } \{A\} x) U$ is elaborated to $(\lambda \{A\} x. \text{id } \{A\} x) \{U\} U$.

It would be more efficient (and also allow more user-friendly error messages) to not immediately refine the inferred type of t to a function type, but rather match on the inferred type, and only perform refining when the type is meta-headed. We present the unoptimized version here for the sake of brevity.

For λ -s, we create a fresh meta for the domain type (since our surface λ -s are not annotated), and infer types for the bodies.

In the case of **let**, inference again just falls through, and we infer a type for the **let** body. For $\llbracket _ \rrbracket \Uparrow_{\Theta|\Gamma}$, we create a fresh meta for the type of the hole, and another fresh meta for the hole itself.

2.6.4 Starting and Finishing Elaboration. Given a surface term t , we initiate elaboration by computing $\llbracket t \rrbracket \Uparrow_{\bullet|\bullet}$. If this succeeds, we get a (Θ, θ, t') result. Elaboration is successful overall if $\Theta = \bullet$ and $\theta = \text{id}$, i.e. no unsolved metas remain.

2.6.5 Properties of Elaboration. First, elaboration is *sound* in the sense that it only produces well-formed output.

THEOREM 2.7 (SOUNDNESS). *The definitions of $\llbracket - \rrbracket \Downarrow$ and $\llbracket - \rrbracket \Uparrow$ conform to the *CHECK* and *INFER* specifications. This follows by induction on surface syntax, while also relying on the properties of substitution, metasubstitution, insert, unify and freshMeta. $\square$*

We remark that this notion of soundness is only a “sanity” or well-typing statement for elaboration. In fact, we could define elaboration as a constantly failing partial function, and it would also conform to the specification. One way to view this is that $\llbracket - \rrbracket \Downarrow$ and $\llbracket - \rrbracket \Uparrow$ together with their specification constitute the semantics of surface syntax. We do not give any other semantics to the surface syntax, nor does it support any other operation.

We do not present any *completeness* result for elaboration in this paper. For an example of what this would entail, in [Dunfield and Krishnaswami 2013] completeness means that whenever there is a way to fill in missing details in the surface syntax, algorithmic typechecking *always* finds it. In *ibid.* this means figuring out domain types for λ -s and inserting all implicit applications. However, we target a far stronger theory, and to our knowledge no prior work has attempted to characterize inferable terms or show completeness for a comparably strong elaborator. Our experience in Agda is that it is not tractable in general to figure out which arguments are inferable, by looking at function types, and often we have to run elaboration to see what works. This does not imply that type inference with dependent types must necessarily be an unwieldy affair. While we lose completeness, we gain the ability to infer more parts of programs by following richer type dependencies, as well as more power to refine and synthesize terms when types are known beforehand.

We can still say something about the behavior of our elaborator. For this, we consider a translation from core terms to surface terms, the evident forgetful translation, which maps core terms to surface counterparts. Now, this is an “evil” construction on core terms, since it does not preserve definitional equality, but we shall only use this evil notion in the following statement.

THEOREM 2.8 (CONSERVATIVITY). *Elaboration is conservative over the surface syntax, in the sense that for any surface term t , if checking or inference outputs t' , then the forgetful translation of t' differs from t only by*

- *Having all $_$ holes filled with core expressions.*
- *Having extra implicit λ -s and implicit applications inserted.*

This follows by straightforward induction on surface syntax. $\square$

Remark. It is *not* the case that for every term in the core syntax there exists a surface term which elaborates to it. The main reason is that λ -s in the surface syntax are not annotated, and it is easy to find core λ expressions with uninferable domain types. We skip optional surface λ domain type annotations for the sake of brevity.

2.6.6 Omitted Features.

- *Let-generalization.* This is an open research topic in settings with dependent types, and we make no attempt at covering it. See [Eisenberg 2016] for a treatment in a proposed dependent version of Haskell.
- *Polymorphic subtyping.* In some prior works, e.g. in [Dunfield and Krishnaswami 2013] and [Vytiñiotis et al. 2008], there is a subtyping relation arising along instantiations of polymorphic types. In GHC 8 polymorphic subtyping is implemented for function types only. Polymorphic subtyping complicates type inference, and to our knowledge it has not been implemented in any dependently typed setting. We also believe that it is undesirable in

dependent settings, because elaboration of subtyping must insert coercions which significantly change the intensional character of programs. For example, if we have covariant list types, then coercing $t : \text{List} (\{A : \mathbf{U}\} \rightarrow A \rightarrow A)$ to $t : \text{List} (\text{Bool} \rightarrow \text{Bool})$ requires mapping over t and inserting implicit applications to Bool for each list element. In System F, all such coercions are erasable, since types are computationally irrelevant, but in our core syntax we have implicit functions with arbitrary (relevant) domains. In GHC, subtyping coercions for functions change operational semantics; this is a reason for abandoning subtyping in recent developments of impredicative inference for GHC [Serrano et al. 2020].

3 ISSUES WITH FIRST-CLASS IMPLICIT FUNCTIONS

We revisit now the *polyList* example from Section 1. We assume the following:

$$\text{List} : \mathbf{U} \rightarrow \mathbf{U} \quad \text{nil} : \{A : \mathbf{U}\} \rightarrow \text{List } A \quad \text{cons} : \{A : \mathbf{U}\} \rightarrow A \rightarrow \text{List } A \rightarrow \text{List } A$$

Below, we present a trace of checking $\text{cons } (\lambda x. x) \text{ nil}$ at type $\text{List} (\{A : \mathbf{U}\} \rightarrow A \rightarrow A)$. We omit context and metacontext parameters everywhere, and mark recursive calls by indentation. We also omit some checking, inference, implicit insertion and unification calls which are not essential for illustration.

```

0       $\llbracket \text{cons } (\lambda x. x) \text{ nil} \rrbracket \Downarrow (\text{List } (\{A : \mathbf{U}\} \rightarrow A \rightarrow A))$ 
1       $\llbracket \text{cons } (\lambda x. x) \text{ nil} \rrbracket \Uparrow$ 
2       $\llbracket \text{cons } (\lambda x. x) \rrbracket \Uparrow$ 
3       $\llbracket \text{cons} \rrbracket \Uparrow$ 
4       $= \text{cons } \{\alpha_0\} : \alpha_0 \rightarrow \text{List } \alpha_0 \rightarrow \text{List } \alpha_0$ 
5       $\llbracket \lambda x. x \rrbracket \Downarrow \alpha_0$ 
6       $= \lambda x. x$ 
7       $= \text{cons } \{\alpha_1 \rightarrow \alpha_1\} (\lambda x. x) : \text{List } (\alpha_1 \rightarrow \alpha_1) \rightarrow \text{List } (\alpha_1 \rightarrow \alpha_1)$ 
8       $\llbracket \text{nil} \rrbracket \Downarrow (\text{List } (\alpha_1 \rightarrow \alpha_1))$ 
9       $= \text{nil } \{\alpha_1 \rightarrow \alpha_1\}$ 
10      $= \text{cons } \{\alpha_1 \rightarrow \alpha_1\} (\lambda x. x) (\text{nil } \{\alpha_1 \rightarrow \alpha_1\}) : \text{List } (\alpha_1 \rightarrow \alpha_1)$ 
11      $\text{unify } (\text{List } (\{A : \mathbf{U}\} \rightarrow A \rightarrow A)) (\text{List } (\alpha_1 \rightarrow \alpha_1))$ 
12      $\text{unify } (\{A : \mathbf{U}\} \rightarrow A \rightarrow A) (\alpha_1 \rightarrow \alpha_1)$ 
13      $= \text{fail}$ 

```

Above, we first infer $\text{cons } (\lambda x. x) \text{ nil}$, which inserts implicit applications to fresh metas in cons and nil , and returns $\text{cons } \{\alpha_1 \rightarrow \alpha_1\} (\lambda x. x) (\text{nil } \{\alpha_1 \rightarrow \alpha_1\}) : \text{List } (\alpha_1 \rightarrow \alpha_1)$. Here, the α_0 meta is refined to $\alpha_1 \rightarrow \alpha_1$ when we check $\lambda x. x$. In the end, we need to unify the expected and inferred types, which fails, since we have an implicit function type on one side and an explicit function on the other side.

Why does this fail? The culprit is line 5, where we call $\llbracket \lambda x. x \rrbracket \Downarrow \alpha_0$. At this point, the checking type is not an implicit function type (it is a meta), so we do not insert an implicit λ . At the heart of the issue is that elaboration makes insertion choices based on core types.

- (1) $\llbracket t \rrbracket \Downarrow_{\Theta \mid \Gamma} A$ can insert a λ only if A is an implicit function type.
- (2) $\text{insert } (\Theta, \theta, t, A)$ can insert an application only if A is an implicit function type.

In both of these cases, if A is of the form $\alpha \bar{u}$ (i.e. meta-headed), then it is possible that α is later refined to an implicit function type, but at that point we have already missed our shot at implicit insertion.

There is a potential naive solution: just *postpone* checking a term until the shape of the checking type is known for sure. This was included as part of a proposed solution for smarter λ -insertions by [Johansson and Lloyd 2015]. This means that checking with a meta-headed type returns a “guarded constant” [Norell 2007, Chapter 3], an opaque stand-in which only computes to an actual core term when the checking type becomes known. In practice, this solution has a painful drawback: *we get no information at all from checked terms before the guard is unblocked*.

For an example for unexpected behavior with this solution, let us assume $Bool : U$ and $true : Bool$, and try to infer type for the surface term **let** $x : _ = true$ **in** x . We first insert a fresh meta α for the hole, and then check $true$ with α . We postpone this checking, returning a guarded constant, and then infer a type for x , which is α . Hence, this small example yields an unsolved meta and a guarded constant in the output.

Now, this particular example can be repaired by special-casing the elaboration of a **let**-definition without an explicit type annotation. However, our experience from playing with an implementation of this solution, is that we are missing too much information by postponing, and this cascades in an unfortunate way: postponing yields more unsolved metas, which cause more postponing.

4 TELESCOPES AND STRICTLY CURRIED FUNCTIONS

As part of the proposed solution, we extend the core theory with telescopes and strictly curried functions. Figure 4 lists the typing rules and definitional equalities.

4.1 Telescopes

Telescopes can be viewed as a generic implementation of record types. We have Tel as the type of telescopes. Elements of Tel are right-nested telescopes of types, with ϵ denoting the empty telescope, and $- \triangleright -$ telescope extension. For example, we can define the signature of natural number algebras as follows:

$$\text{let } NatAlgSig : Tel = (N : U) \triangleright (zero : N) \triangleright (suc : N \rightarrow N) \triangleright \epsilon \text{ in } \dots$$

We interpret an $A : Tel$ as a record type as $Rec\ A$, which behaves as the evident iterated Σ -type corresponding to the telescope. Hence, $Rec\ \epsilon$ is isomorphic to the unit type, with inhabitant $[]$, and $Rec\ ((x : A) \triangleright B)$ behaves as a Σ -type, with pairing constructor $- :: -$ and projections π_1 and π_2 . We also have the β and η rules for record constructors and projections in Figure 4. We present definitional equalities in a compact form, but note that they still stand for $\boxed{\Theta | \Gamma \vdash t \equiv u : A}$ judgments. Hence, the sides of the equations must have the same types, and in particular the left side of the $[]$ - η rule has type $Rec\ \epsilon$.

We naturally want to make sure that the extended core syntax is still sensible. In relation to type theories, the following three metatheoretical properties are commonly considered.

- A theory is *consistent* if not every closed type has a closed inhabitant. An inconsistent theory cannot be used as a *logic*, from the propositions-as-types perspective [Univalent Foundations Program 2013, Chapter 1], because there every proposition is provable, so proofs do not distinguish truth from falsehood.
- A theory has the *canonicity* property, if every closed term is definitionally equal to a canonical term. For example, every closed natural number is a numeral, every $Bool$ is $true$ or $false$ and so on. A canonical theory is useful as a *total programming language*; if canonicity fails, then some closed programs diverge or compute to stuck terms which cannot be viewed as

values. In non-total languages, progress and type preservation constitute a weaker analogue of canonicity [Harper 2016, Chapter 6].

- A theory has *decidable conversion*, if definitional equality is decidable. Undecidable conversion usually implies undecidable type checking for type theories. Often, decidable conversion follows from a normalization procedure for open terms: we first normalize, then decide equality of normal forms.

The properties are listed roughly in the order of increasing strength. Canonicity often implies consistency (because empty types have no canonical terms), and normalization implies canonicity. Now, our current core theory has *none* of the above properties, because the $U : U$ rule precludes them. However, as we mentioned in Section 2.2, this is easily fixed by stratified universes.

In any case, a *reduction* of new features to old features is a common and convenient strategy to reduce metatheoretical properties of extended theories to that of smaller or better-known theories. See Boulier et al. [2017] for a collection of examples for this. In the simplest case, new features are definable from old features. Here, we can either assume that new features are merely definitional shorthands, or give an evident injective translation from the new theory to the old one. Both imply that the desired properties transfer from the old theory to the new one.

Telescopes and records are derivable from natural numbers, the unit type and Σ -types. We write $\top$ for the unit type, and we have $\text{tt} : \top$. We use Agda-like pattern matching notation in the following. First, we define length-indexed telescopes.

$$\begin{aligned} \text{Tel}' &: \text{Nat} \rightarrow U \\ \text{Tel}' \text{ zero} &:= \top \\ \text{Tel}' (\text{suc } n) &:= \Sigma(A : U). (A \rightarrow \text{Tel}' n) \end{aligned}$$

Then, we have $\text{Tel} := \Sigma(n : \text{Nat}). (\text{Tel}' n)$, and define records:

$$\begin{aligned} \text{Rec} &: \text{Tel} \rightarrow U \\ \text{Rec} (\text{zero}, _) &:= \top \\ \text{Rec} (\text{suc } n, (A, B)) &:= \Sigma(a : A). (\text{Rec } (n, B a)) \end{aligned}$$

From the above, ϵ , $- \triangleright -$, $- :: -$ and $[]$ are evident, and all expected equalities hold definitionally. Hence, it does not matter whether telescopes are primitive or derived. In the rest of the paper we never need to look inside the representation of telescopes, and we only use the “interface” from Figure 4.

4.2 Strictly Curried Functions

These are function types whose domains are telescopes, and they are immediately computed to iterated implicit function types when the domain telescope is canonical. See $\text{FUN-}\epsilon$ and $\text{FUN-}\triangleright$: a curried function with empty domain computes to simply the codomain, while a function with a non-empty domain computes to an implicit function type. We explicitly write telescopes in both λ -abstractions and applications for strictly curried functions, since they are relevant in the computation rules.

Curried function types tend to be computed away, but they can persist if the domain telescope is neutral, and in particular when it is meta-headed. For example, assuming a meta $\alpha : \text{Tel}$, the type $\{x : \bar{\alpha}\} \rightarrow B$ cannot be computed further. During elaboration, we will use strictly curried function types to represent unknown insertions, but these types are eventually computed away if a surface expression can be successfully elaborated. Since the surface language remains unchanged, telescopes and curried functions are merely an internal implementation detail from the perspective of programmers.

TEL	EMPTY-TEL	NONEMPTY-TEL	RECORD-TYPE
$\frac{}{\Theta \Gamma \vdash \text{Tel} : \mathbb{U}}$	$\frac{}{\Theta \Gamma \vdash \epsilon : \text{Tel}}$	$\frac{\Theta \Gamma \vdash A : \mathbb{U} \quad \Theta \Gamma, x : A \vdash B : \text{Tel}}{\Theta \Gamma \vdash (x : A) \triangleright B : \text{Tel}}$	$\frac{\Theta \Gamma \vdash A : \text{Tel}}{\Theta \Gamma \vdash \text{Rec } A : \mathbb{U}}$
	EMPTY-RECORD	NONEMPTY-RECORD	
	$\frac{}{\Theta \Gamma \vdash [] : \text{Rec } \epsilon}$	$\frac{\Theta \Gamma \vdash t : A \quad \Theta \Gamma \vdash u : \text{Rec } (B[x \mapsto t])}{\Theta \Gamma \vdash t :: u : \text{Rec } ((x : A) \triangleright B)}$	
RECORD-PROJECTION		CURRIED-FUN	
$\frac{\Theta \Gamma \vdash t : \text{Rec } ((x : A) \triangleright B)}{\Theta \Gamma \vdash \pi_1 t : A \quad \Theta \Gamma \vdash \pi_2 t : \text{Rec } (B[x \mapsto \pi_1 t])}$		$\frac{\Theta \Gamma \vdash A : \text{Tel} \quad \Theta \Gamma, x : \text{Rec } A \vdash B : \mathbb{U}}{\Theta \Gamma \vdash \{x : \bar{A}\} \rightarrow B : \mathbb{U}}$	
CURRIED-LAM		CURRIED-APP	
$\frac{\Theta \Gamma, x : \text{Rec } A \vdash t : B}{\Theta \Gamma \vdash \lambda \{x : \bar{A}\}. t : \{x : \bar{A}\} \rightarrow B}$		$\frac{\Theta \Gamma \vdash t : \{x : \bar{A}\} \rightarrow B \quad \Theta \Gamma \vdash u : \text{Rec } A}{\Theta \Gamma \vdash t \{u : \bar{A}\} : B[x \mapsto u]}$	
$\pi_1\text{-}\beta$	$\pi_1 (t :: u)$	$\equiv t$	
$\pi_2\text{-}\beta$	$\pi_2 (t :: u)$	$\equiv u$	
$::\text{-}\eta$	$(\pi_1 t :: \pi_2 t)$	$\equiv t$	
$[]\text{-}\eta$	t	$\equiv []$	
FUN- ϵ	$\{x : \bar{\epsilon}\} \rightarrow B$	$\equiv B[x \mapsto []]$	
FUN- $\triangleright$	$\{x : \overline{(y : A) \triangleright B}\} \rightarrow C$	$\equiv \{y : A\} \rightarrow (\{b : \bar{B}\} \rightarrow C[x \mapsto (y :: b)])$	
LAM- ϵ	$\lambda \{x : \bar{\epsilon}\}. t$	$\equiv t[x \mapsto []]$	
LAM- $\triangleright$	$\lambda \{x : \overline{(y : A) \triangleright B}\}. t$	$\equiv \lambda \{y\}. \lambda \{b : \bar{B}\}. t[x \mapsto (y :: b)]$	
APP- ϵ	$t \{u : \bar{\epsilon}\}$	$\equiv t$	
APP- $\triangleright$	$t \{u : \overline{(x : A) \triangleright B}\}$	$\equiv t \{\pi_1 u\} \{\pi_2 u : \overline{B[x \mapsto \pi_1 u]}\}$	
CURRIED- β	$\lambda (\{x : \bar{A}\}. t) \{u : \bar{A}\}$	$\equiv t[x \mapsto u]$	
CURRIED- η	$\lambda \{x : \bar{A}\}. t \{x : \bar{A}\}$	$\equiv t$	

Fig. 4. Rules for telescopes and strictly curried functions.

Curried functions are *mostly* derivable from Nat , $\top$ and Σ . The type former is defined as follows:

$$\begin{aligned}
\Pi^C : (A : \text{Tel}) &\rightarrow (\text{Rec } A \rightarrow \mathbb{U}) \rightarrow \mathbb{U} \\
\Pi^C(\text{zero}, _) B &\quad \equiv B \text{ tt} \\
\Pi^C(\text{suc } n, (A, B)) C &\equiv \{a : A\} \rightarrow \Pi^C(n, B a)(\lambda b. C(a, b))
\end{aligned}$$

With this, we can also define $\text{app} : \Pi^C A B \rightarrow (a : \text{Rec } A) \rightarrow B a$ and $\text{lam} : ((a : \text{Rec } A) \rightarrow B a) \rightarrow \Pi^C A B$, and all equations in Figure 4 hold definitionally, except $\text{CURRIED-}\beta$ and $\text{CURRIED-}\eta$. These do not hold strictly, because Π^C , app and lam are all defined by recursion on the A telescope, but the $\beta\eta$ rules are specified generically for arbitrary (possibly neutral) telescopes. $\text{CURRIED-}\beta$ is still provable as a propositional equality, and assuming function extensionality $\text{CURRIED-}\eta$ is provable as

well. For details, see our Agda formalization of these definitions, which is included alongside the prototype implementation.

Hence, we can derive a somewhat weaker version of curried functions, with propositional β and η . From this, we still get consistency and canonicity very cheaply. This is because models proving consistency or canonicity for the base theory with Nat , $\top$ and Σ usually support equality reflection, i.e. that propositional equality implies definitional equality. For consistency, standard set-theoretical models and models in extensional type theory have this property, for canonicity, glued models (e.g. as in [Kaposi et al. 2019] and [Sterling 2019]) also have this property.

In contrast, showing normalization and decidability of conversion would require some extra work. We leave this to future work, but we expect that it is not difficult to extend previous proofs to cover strict β and η for curried functions.

5 EXTENDING ELABORATION

We shall utilize the extended core theory to implement smarter elaboration. Recall from Section 3 that the old elaborator makes two kinds of unforced insertion choices:

- (1) $\llbracket t \rrbracket \Downarrow_{\Theta \vdash \Gamma} A$ does not insert an implicit λ when A is meta-headed.
- (2) $\text{insert}(\Theta, \theta, t, A)$ does not insert an implicit application when A is meta-headed.

In the following, we shall only enhance λ -insertions (1). This allows a simple implementation which only requires small changes to unification, and which is already remarkably powerful. It seems that enhancing implicit application insertions in (2) requires extending unification; we discuss this in Section 7.3. First, we modify closing types and contextualization to take advantage of telescopes.

Definition 5.1 (Closing types). We use curried function types to close over record types in the scope. If a bound variable does not have a record type, then we do as before³. We prepend the following clause to Definition 2.3:

$$((\Gamma, x : \text{Rec } A) \Rightarrow B) \equiv (\Gamma \Rightarrow (\{x : \bar{A}\} \rightarrow B))$$

Definition 5.2 (Contextualization). We extend spine notation to applications of curried functions. For example, we may have a spine $\bar{t} \equiv (\{x : \bar{A}\}, \{y : \bar{B}\})$. We accordingly revise Definition 2.4 for $\overline{\text{vars}}_{\Gamma}$ so that we use curried function application for each record type in Γ .

5.1 Handling Superfluous Implicit Functions

Before we can move on to unification and elaboration, we have to address a curious issue. Assuming $\text{Bool} : \mathcal{U}$, $\text{true} : \text{Bool}$ and $\text{false} : \text{Bool}$, consider the following surface expression:

let $x : _ = \text{true}$ **in** x

What should this expression elaborate to? We would expect the result to be simply

let $x : \text{Bool} = \text{true}$ **in** x

However, there are infinitely many core terms which are conservative over the surface expression in the sense of Theorem 2.8. That is, we can wrap definitions with any number of implicit λ -s, and add implicit applications accordingly to usage sites of the defined name. For example, we could have

let $x : \{y : \text{Bool}\} \rightarrow \text{Bool} = \lambda \{y\}. \text{true}$ **in** $x \{ \text{true} \}$

³This implies that we close over meta-headed types using plain functions. In theory, this causes a higher-order version of the basic implicit insertion problem: we are uncertain about whether we should be uncertain about implicit insertions. This issue does not seem to come up in realistic programs in the prototype implementation; the first-order treatment of uncertain insertions seem to leave little remaining “higher-order” uncertainty on the table.

$$\begin{array}{c}
\text{METAACON/CONSTANCY} \\
\frac{\Theta|\Gamma, x : \text{Rec } A \vdash B : \mathbf{U}}{\Theta, \text{constancy}_{\Gamma, x:\text{Rec } A} B \vdash} \\
\\
\text{CONSTANCY-}\equiv \\
\frac{x \notin \text{FreeVars}(B)}{\Theta_0, \text{constancy}_{\Gamma_0, x:\text{Rec } A} B, \Theta_1|\Gamma_1 \vdash A \equiv \epsilon : \text{Tel}} \\
\\
\text{METASUB/CONSTANCY} \\
\frac{\theta : \Theta_0 \Rightarrow \Theta_1 \quad x \notin \text{FreeVars}(B[\theta]) \text{ implies } \Theta_0|\Gamma[\theta] \vdash A[\theta] \equiv \epsilon : \text{Tel}}{(\theta, \text{solve}_{\Gamma, x:\text{Rec } A} B) : \Theta_0 \Rightarrow (\Theta_1, \text{constancy}_{\Gamma, x:\text{Rec } A} B)} \\
\\
\text{METASUB/WEAKEN-CONSTANCY} \\
\frac{}{p : (\Theta, \text{constancy}_{\Gamma, x:\text{Rec } A} B) \Rightarrow \Theta}
\end{array}$$

Fig. 5. Rules for constancy constraints

This is clearly undesirable. With the type $\{y : \text{Bool}\} \rightarrow \text{Bool}$, the implicit argument y is never inferable, because the codomain type does not depend on the domain, and the argument is never constrained. Hence, with the above definition, we always have to write $x \{ \text{true} \}$ or $x \{ \text{false} \}$ when we want to use x . In order to avoid such nonsense, we adopt the following principle: *elaboration should never invent non-dependent implicit function types*.

This was a non-issue in the old elaborator, because it was not able to invent implicit function types; it was only utilizing the type annotations present in the surface input. In the case of `let x : _ = true`, the old elaborator checks `true` with a fresh meta, and just assumes that the meta does not stand for an implicit function type.

5.1.1 Constancy Constraints. We use these constraints to get rid of curried function types as soon as we learn that they are non-dependent. They are constraints in the usual sense in unification algorithms (e.g. as in [Abel and Pientka 2011] or [Vytiniotis et al. 2011]). We formalize them in a compact way, by adding a new kind of context extension for metacontexts. The rules are given in Figure 5.

In the rule METAACON/CONSTANCY we specify extension of a metacontext with a constraint. The CONSTANCY- $\equiv$ rule expresses that, assuming we have a constancy constraint for A and B in context, if B does not depend on the $x : \text{Rec } A$ domain variable, then A is equal to the empty telescope ϵ .

The METASUB/CONSTANCY rule defines metasubstitutions whose codomains are extended with constraints. Intuitively, while the METASUB/EXTENDED rule from Section 2.3 can be used to solve a metavariable (by mapping it to a term), METASUB/CONSTANCY solves a constraint. We can only extend $\theta : \Theta_0 \Rightarrow \Theta_1$ to map into an additional constraint if θ forces the constraint to hold. In METASUB/WEAKEN-CONSTANCY, we overload p for the weakening substitution which drops a constraint.

Definition 5.3 (Creating a new constraint). We do this similarly to Definition 2.6, by simply returning a weakening substitution.

$$\text{newConstancy}_{\Theta|\Gamma, x:\text{Rec } A} B \equiv ((\Theta, \text{constancy}_{\Gamma, x:\text{Rec } A} B), p)$$

5.1.2 Algorithmic Implementation of Constraint Solving. The above specification for constancy constraints is compact but not particularly algorithmic: we just magically get new definitional equalities whenever we have constraints in contexts. In our prototype implementation, we implement eager removal of solvable constraints.

After solving a meta α during unification, which yields a unifying θ substitution, we review all $(\text{constancy}_{\Gamma, x:\text{Rec } A} B)$ constraints in the context, such that x occurs in B inside a $\bar{t}$ spine of some $\bar{\alpha} \bar{t}$ term. In other words, we review constraints where the new meta solution might make a difference.

- (1) If we have $x \in \text{FreeVars}(B[\theta])$, where x occurs rigidly in $B[\theta]$, i.e. the occurrence is not inside a spine of a meta, then no metasubstitution can possibly remove this occurrence. In this

case the constraint holds vacuously, so we can use the `METASUB/CONSTANCY` rule to return a θ' substitution which also solves the constraint.

- (2) If we have $x \notin \text{FreeVars}(B[\theta])$, we recursively unify $A[\theta]$ with ϵ . If that succeeds, we get a θ' which unifies $A[\theta]$ and ϵ and thus forces the constraint to hold, so we can again use `METASUB/CONSTANCY` to solve the constraint.
- (3) In any other case we simply return θ and keep the constraint around.

Also, when we create a new constancy constraint, we immediately review it as described above.

Remark. In the case with $x \in \text{FreeVars}(B[\theta])$, it would be also sound to solve the constraint when the occurrence is not rigid. However, this way we could lose potential non- ϵ solutions of $A[\theta]$.

5.2 Unification For Strictly Curried Function Types

Although we omit most details of unification, we shall discuss it for curried function types, as it is essential in the extended elaboration algorithm. We seek to solve meta-telescopes in a way such that iterated implicit function types on the two sides of a unification problem become definitionally equal. First we define a helper function which counts leading implicit function types:

$$\begin{aligned} \text{implArity } (\{x : A\} \rightarrow B) &::= 1 + \text{implArity } B \\ \text{implArity } A &::= 0 \end{aligned}$$

Then, we extend unification. The most interesting case is when we unify a curried function type with an implicit function type. In this case, if there are “missing” implicit arguments on the side of the curried function type (note the `implArity` condition on the pattern), we refine the A telescope domain to an extended $(x_0 : A_0) \triangleright A_1$ telescope. Since we invent a fresh A_1 domain for a curried function type, we need to add a constancy constraint for it as well.

```
unifyΘ0|Γ ({x :  $\bar{A}$ } → B) ({x0 : A0} → B') | implArity B < implArity B' + 1 ::= do
  let (Θ1, θ1, A1) = freshMetaΘ0|Γ, x0:A0 Tel
  (Θ2, θ2) ← unifyΘ1|Γ[θ1] (A[θ1]) ((x : A0[θ1]) ▷ A1)
  let (Θ3, θ3) = newConstancyΘ2|Γ[θ1], x0:A0[θ1], x1:Rec(A1[θ2]) (B[θ1][x ↦ (x0 :: x1)])
  unifyΘ3|Γ[θ1], x0:A0[θ1] ({x1 : A1[θ2]} → B[θ1][x ↦ (x0 :: x1)]) B'
```

We have the symmetric `unifyΘ0|Γ ({x0 : A0} → B') ({x :  $\bar{A}$ } → B)` case the same way as above.

Below, let us assume that B' is not meta-headed, and it also does not match the pattern above. In this case we solve the domain telescope to be empty.

```
unifyΘ0|Γ ({x :  $\bar{A}$ } → B) B' ::= do
  (Θ1, θ1) ← unifyΘ0|Γ A  $\epsilon$ 
  unifyΘ1|Γ[θ1] (B[θ1][x ↦ []]) (B'[θ1])
```

Again, we also have the symmetric case. Some examples for unification behavior:

- If B is not meta-headed, `unify ({x :  $\bar{\alpha}$ } → B) B` causes α to be solved as ϵ , for the following reason: if B is an implicit function type, the arity condition fails, otherwise the pattern match fails, so either way we fall through the first rule and solve α to ϵ .
- If B is again not meta-headed, `unify ({x :  $\bar{\alpha}$ } → B) ({x : A} → B)` causes α to be solved as $(x : A) \triangleright \epsilon$.

Note that this algorithm matches one curried function type at a time, and does not recognize that multiple possible solutions exist e.g. in the case of `unify ({x :  $\bar{\alpha}$ } {y :  $\bar{\beta}$ } → C) ({x : A} {y : B} → C)`, where we could solve α to ϵ , or $(x : A) \triangleright \epsilon$, or $(x : A) \triangleright (y : B) \triangleright \epsilon$. We leave such more sophisticated n-ary matching to future work.

For telescopes, Rec, inhabitants of Rec, curried abstraction and application, unification is structural, and other cases remain the same as in the basic elaborator of Section 2.

5.3 Elaboration

In the definition of checking, we insert a new clause after $\llbracket t \rrbracket \Downarrow_{\Theta_0|\Gamma} (\{x : A\} \rightarrow B)$:

```

 $\llbracket t \rrbracket \Downarrow_{\Theta_0|\Gamma} (\alpha \bar{u}) \equiv \mathbf{do}$ 
   $\mathbf{let} (\Theta_1, \theta_1, A) = \mathbf{freshMeta}_{\Theta_0|\Gamma} \mathbf{Tel}$ 
   $(\Theta_2, \theta_2, t', B) \leftarrow \mathbf{insert} (\llbracket t \rrbracket \Uparrow_{\Theta_1|\Gamma[\theta_1], x:\mathbf{Rec} A})$ 
   $\mathbf{let} (\Theta_3, \theta_3) = \mathbf{newConstancy}_{\Theta_2|\Gamma[\theta_{12}], x:\mathbf{Rec} (A[\theta_2])} B$ 
   $(\Theta_4, \theta_4) \leftarrow \mathbf{unify} ((\alpha \bar{u})[\theta_{123}]) (\{x : \overline{A[\theta_{23}]}\} \rightarrow B[\theta_3])$ 
   $\mathbf{return} (\Theta_4, \theta_{1234}, (\lambda \{x : \overline{A[\theta_{234}]}\}. t'[\theta_{34}]))$ 

```

Hence, when checking a term with a meta-headed type, we create a fresh meta with Tel type, infer a type for the term and wrap the result in a strictly curried λ . We again need to create a new constancy constraint. This way, if B does not depend on $x : \mathbf{Rec} A$, the strictly curried λ in the output immediately computes away, since A is solved to ϵ .

This concludes the definition of the extended elaborator. The new algorithm is sound with respect to the extended core syntax, and it also has the conservativity property from Theorem 2.8 if we additionally allow elaboration to insert λ -s for strictly curried functions. We present some examples of the algorithm in action.

Example 5.4. We return to the *polyList* example from Section 3. We trace elaboration using the extended algorithm.

```

0       $\llbracket \mathbf{cons} (\lambda x. x) \mathbf{nil} \rrbracket \Downarrow (\mathbf{List} (\{A : \mathbf{U}\} \rightarrow A \rightarrow A))$ 
1       $\llbracket \mathbf{cons} (\lambda x. x) \mathbf{nil} \rrbracket \Uparrow$ 
2       $\llbracket \mathbf{cons} (\lambda x. x) \rrbracket \Uparrow$ 
3       $\llbracket \mathbf{cons} \rrbracket \Uparrow$ 
4       $= \mathbf{cons} \{\alpha_0\} : \alpha_0 \rightarrow \mathbf{List} \alpha_0 \rightarrow \mathbf{List} \alpha_0$ 
5       $\llbracket \lambda x. x \rrbracket \Downarrow \alpha_0$ 
6       $= \lambda \{y : \overline{\alpha_1}\}. \lambda x. x$ 
7       $= \mathbf{cons} \{\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\}\} (\lambda \{y : \overline{\alpha_1}\}. \lambda x. x)$ 
8       $\llbracket \mathbf{nil} \rrbracket \Downarrow (\mathbf{List} (\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\}))$ 
9       $= \mathbf{nil} \{\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\}\}$ 
10      $= \mathbf{cons} \{\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\}\} (\lambda \{y : \overline{\alpha_1}\}. \lambda x. x)$ 
         $(\mathbf{nil} \{\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\}\})$ 
         $: \mathbf{List} (\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\})$ 
11      $\mathbf{unify} (\mathbf{List} (\{A : \mathbf{U}\} \rightarrow A \rightarrow A)) (\mathbf{List} (\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\}))$ 
12      $\mathbf{unify} (\{A : \mathbf{U}\} \rightarrow A \rightarrow A) (\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\})$ 
13      $\mathbf{unify} \alpha_1 ((A : \mathbf{U}) \triangleright (\alpha_3 A))$ 
14      $\mathbf{unify} (A \rightarrow A) (\{z : \overline{\alpha_3 A}\} \rightarrow \alpha_2 \{A\} \{z : \overline{\alpha_3 A}\} \rightarrow \alpha_2 \{A\} \{z : \overline{\alpha_3 A}\})$ 
15      $\mathbf{unify} (\alpha_3 A) \epsilon$ 

```

```

16      unify (A → A) (α2 {A} → α2 {A})
17      unify A (α2 {A})
18      unify A A
19  = cons {{A : U} → A → A} (λ {A} x. x) (nil {{A : U} → A → A})

```

We diverge from the previous attempt at line 5. Here, we check $\lambda x. x$ with the meta α_0 , so we create a fresh $\alpha_1 : \text{Tel}$ meta and wrap the result as $\lambda \{y : \overline{\alpha_1}\}. \lambda x. x$. This result has the inferred type $\{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\} \rightarrow \alpha_2 \{y : \overline{\alpha_1}\}$, and we promptly unify α_0 (which stands for the list element type) with it. On line 10, we return the elaborated *cons* expression, where the list element type is made explicit. It only remains to unify the inferred and expected types. On line 12 we have the case where a curried function type is matched with an implicit function type, so on line 13 we refine α_1 to a non-empty telescope, and proceed unifying the domains. Note that on line 14 we have $\alpha_2 \{A\} \{z : \overline{\alpha_3 A}\}$, which results from applying the substitution $y \mapsto (A :: z)$, and computing $\alpha_2 \{(A :: z) : (A : U) \triangleright (\alpha_3 A)\}$ further using the *APP-▷* rule from Figure 4.

On line 14 we have the case when a curried function type is matched with a type which is not a curried function, implicit function or a meta-headed type. We accordingly unify domain telescope $\alpha_3 A$ with ϵ , which causes α_3 to be solved as $\lambda A. \epsilon$ according to standard pattern unification. Now, $\alpha_2 \{A\} \{z : \alpha_3 A\}$ computes to $\alpha_2 \{A\} \{z : \epsilon\}$, which further computes to $\alpha_2 \{A\}$ by *APP-ε*. From here, unification finishes with success. On the last line, the expected output is returned: since the α_1 telescope meta is now solved, curried function types and abstractions are computed away.

Computing $\llbracket \text{cons } (\lambda \{A\} x. x) \text{ nil} \rrbracket \Downarrow (\text{List } (\{A : U\} \rightarrow A \rightarrow A))$ also succeeds, as in Agda and our basic elaborator. In this case elaboration also inserts a curried λ around $(\lambda \{A\} x. x)$, but its domain is then solved to ϵ according to the rules for curried function type unification.

Example 5.5. We illustrate now the action of constancy constraints, using the example from Section 5.1. Again, assume $\text{Bool} : U$ and $\text{true} : \text{Bool}$.

```

0       $\llbracket \text{let } x : \_ = \text{true in } x \rrbracket \Uparrow$ 
1       $\llbracket \_ \rrbracket \Downarrow U = \alpha_0$ 
2       $\llbracket \text{true} \rrbracket \Downarrow \alpha_0$ 
3      freshMeta Tel = α1
4       $\llbracket \text{true} \rrbracket \Uparrow = \text{true} : \text{Bool}$ 
5      newConstancy•, x : \text{Rec } \alpha_1 Bool
6      unify α1 ε
7      unify α0 Bool
8      = true
9       $\llbracket x \rrbracket \Uparrow \text{true} = \text{true} : \text{Bool}$ 
10     = (let x : Bool = true in x) : Bool

```

First we create a fresh meta α_0 for the hole on line 1, then we check *true* with it. Since we are checking with a meta-headed type, we create a fresh telescope meta α_1 , then infer *Bool* for *true*. On line 5 we create a constancy constraint and immediately try to solve it, as described in Section 5.1.2. Since *Bool* does not depend on the domain, we solve the constraint and thereby solve α_1 as ϵ . Hence, we simply return *true* on line 8, since $\lambda \{x : \overline{\epsilon}\}. \text{true}$ computes to that, and elaboration succeeds with the expected output.

6 IMPLEMENTATION

We provide an implementation of the elaborator from Section 5. It is a standalone Haskell program which reads a surface expression from standard input, and outputs the result of elaboration, or optionally the type or the normal form of the result. It is implemented in about 1100 lines of Haskell. Of this, 680 lines constitute the core syntax, evaluation, unification and elaboration. Of these 680 lines, about 450 lines implement the basic elaborator of Section 2 and 230 lines implement the extended elaborator.

Elaboration is implemented in the style of Coquand’s algorithm [Coquand 1996], where elaboration is interleaved with normalization-by-evaluation. Hence, we do not perform any substitution operations on core syntax, instead we evaluate core terms into the semantic domain, and perform unification, strengthening and occurrence checking on semantic values. There is also a quoting (or readback) operation which yields normalized core terms from values.

Metacontexts are a key point in the implementation. We avoid the tedious (and inefficient) threading of composed metasubstitutions, instead we have a mutable reference which stores the current metasubstitution. We have a “forcing” operation which computes a semantic value to a head normal form with respect to the current metasubstitution. Hence, instead of constantly updating every term and type by performing metasubstitution, we only force them whenever we need to pattern match on the shape of types, for example when we are inserting implicit arguments based on types. In this paper we still stick with the “threaded metasubstitution” presentation because it is more conventional, and also more straightforward to formalize.

We use normalization-by-evaluation in Abel’s style (see e.g. [Abel 2013, Chapter 3]), where semantic values use de Bruijn levels, and the core syntax uses de Bruijn indices. This is practically very favorable, because the evaluator never has to perform weakening on values. The implementation of curried functions presents a bit of a complication, because the computation rules are type-directed, so we have to annotate applications and abstractions with telescopes (just as in our notation for the core syntax). We implement constraints in a fairly optimized way: we keep track of relevant “blocking” metas for each constraint and upon solving a meta we only review constraints which were blocked on the meta.

7 RELATED WORKS AND EVALUATION

In this section we examine how the elaborator fares in practice, its limitations, and how it compares to related works. We refer to our elaborator as FCIF when comparing it to others.

7.1 Related Works

MLF [Botlan and Rémy 2003] extends System F with polymorphic subtyping bounds, and supports strong inference for first class polymorphism. HML [Leijen 2009] is a simplified variant of MLF. These systems rely critically on subtyping and thus diverge markedly from (fragments of) Martin-Löf type theory. HMF [Leijen 2008] is a relatively simple system with first-class polymorphism; however, it is also weaker than others in terms of inferable annotations, and its more powerful variant (extended with n-ary application handling) is more complex.

There have been several attempts in the context of GHC. Boxy types [Vytiniotis et al. 2006] and FPH [Vytiniotis et al. 2008] were two early iterations which suffered from complex specification, complex implementation or fragility. More recent works are guarded impredicativity (GI) [Serrano et al. 2018] and quick look impredicativity (QL) [Serrano et al. 2020]. They both work by examining n-ary applications to find metavariable instantiations which arise from occurrences guarded by rigid type constructors. GI’s use of *generalization constraints* anticipates our use of telescopes, but overall GI is also burdened by a great deal of complexity. QL streamlines GI by eschewing

$IdTy \equiv \{A : U\} \rightarrow A \rightarrow A$	$inc : Int \rightarrow Int$
$single : \{A : U\} \rightarrow A \rightarrow List A$	$auto : IdTy \rightarrow IdTy$
$id : IdTy$	$auto' : \{B : U\} \rightarrow IdTy \rightarrow B \rightarrow B$
$ids : List IdTy$	$choose : \{A : U\} \rightarrow A \rightarrow A \rightarrow A$
$nil : \{A : U\} \rightarrow List A$	$app : \{AB : U\} \rightarrow (A \rightarrow B) \rightarrow A \rightarrow B$
$cons : \{A : U\} \rightarrow A \rightarrow List A \rightarrow List A$	$revapp : \{AB : U\} \rightarrow A \rightarrow (A \rightarrow B) \rightarrow B$
$head : \{A : U\} \rightarrow List A \rightarrow A$	$runST : \{A : U\} \rightarrow (\{S : U\} \rightarrow ST S A) \rightarrow A$
$tail : \{A : U\} \rightarrow List A \rightarrow List A$	$argST : \{S : U\} \rightarrow ST S Int$
$map : \{AB : U\} \rightarrow (A \rightarrow B) \rightarrow List A \rightarrow List B$	$poly : IdTy \rightarrow Pair Int Bool$

Fig. 6. Types used in Figure 7.

constraints in favor of an eager preprocessing pass on neutral expressions which finds polymorphic instantiations. QL also takes advantage of bidirectional type propagation. Although QL seems to be practically the most favorable so far, it is far from being elegant: it rechecks expressions multiple times, and the restriction of preprocessing to (nested) neutral applications is rather ad-hoc.

7.2 Evaluation

We borrow a very useful collection of inference benchmarks from the GI [Serrano et al. 2018] and QL [Serrano et al. 2020] papers: see Figure 6 and Figure 7. We also include a source file in the implementation which reproduces these results. In *ibid.* a comparison is presented between multiple systems, but here we only include FCIF. The relative performance of FCIF in this benchmark is easy to remember: it handles exactly the same cases as QL (which is slightly more than what GI covers). We mark some cases as “Yes*”, where FCIF successfully infers a type, but since FCIF does no let-generalization, the inferred types are not fully constrained without extra contextual information.

In the A1 case, FCIF inserts only a single curried λ on the outside, which is why the inferred type is not unifiable with $\{A : U\} \rightarrow A \rightarrow \{B : U\} \rightarrow B \rightarrow A$. However, this could be easily remedied by adding an extra curried λ insertion in the definition of $\llbracket \lambda x. t \rrbracket \uparrow \uparrow \ominus \Gamma$.

In A8, the failure is intentional: the types of id and $auto'$ are not unifiable because of the mismatched order of implicit and explicit arguments. We do not float out or reorder implicit arguments in any way, because we want to support arbitrary mixing of implicit and explicit arguments in the surface language, and such reordering would be problematic anyway in the presence of dependent types. The same situation arises in E1, where we cannot unify $\{A : U\} \rightarrow Int \rightarrow A \rightarrow A$ with $Int \rightarrow \{A : U\} \rightarrow A \rightarrow A$.

Some general comments on the comparison of FCIF to prior works. First, FCIF supports dependent types and prior solutions do not. It is also unclear whether prior solutions can scale to dependent types. MLF and HML rely on subtyping, and solutions which work one neutral spine at a time (HMF, GI, QL) also face issues with dependently typed spines. With such spines, we cannot simply process later arguments when previous ones are not yet elaborated, as return types depend on argument values, and skipping over arguments may clog up type computation in an unacceptable way.

Secondly, FCIF supports global inference. For an example for this, **let** $x : _ = single\ id$ **in** $cons\ \{List\ IdTy\}\ x\ nil$ works in FCIF, MLF and HML but does not work in GI, QL and HMF. MLF and HML work here by immediately giving a principal type to $single\ id$ which involves generalization

A POLYMORPHIC INSTANTIATION		
A1	$\lambda x y. y$ we infer a type which may be solved to $\{A B : U\} \rightarrow A \rightarrow B \rightarrow A$ but not to $\{A : U\} \rightarrow A \rightarrow \{B : U\} \rightarrow B \rightarrow A$	Yes*
A2	<i>choose id</i>	Yes*
A3	<i>choose nil ids</i>	Yes
A4	$\lambda (x : \{A : U\} \rightarrow A \rightarrow A). x x$	Yes*
A5	<i>id auto</i>	Yes
A6	<i>id auto'</i>	Yes*
A7	<i>choose id auto</i>	Yes
A8	<i>choose id auto'</i>	No
A9	$\lambda (f : \{A : U\} \rightarrow (A \rightarrow A) \rightarrow List A \rightarrow A). f (choose id) ids$	Yes*
A10	<i>poly id</i>	Yes
A11	<i>poly</i> ($\lambda x. x$)	Yes
B INFERENCE OF POLYMORPHIC ARGUMENTS		
B1	$\lambda f. pair (f zero) (f true)$	No
B2	$\lambda xs. poly (head xs)$	No
C FUNCTIONS ON POLYMORPHIC LISTS		
C1	<i>length ids</i>	Yes
C2	<i>tail ids</i>	Yes
C3	<i>head ids</i>	Yes*
C4	<i>single id</i>	Yes*
C5	<i>cons id ids</i>	Yes
C6	<i>cons</i> ($\lambda x. x$) <i>ids</i>	Yes
C7	<i>append</i> (<i>single inc</i>) (<i>single id</i>)	Yes
C8	$\lambda (g : \{A : U\} \rightarrow List A \rightarrow List A \rightarrow A). g (single id) ids$	Yes*
C9	<i>map poly</i> (<i>single id</i>)	Yes
C10	<i>map head</i> (<i>single ids</i>)	Yes
D APPLICATION FUNCTIONS		
D1	<i>app poly id</i>	Yes
D2	<i>revapp id poly</i>	Yes
D3	<i>runST argST</i>	Yes
D4	<i>app runST argST</i>	Yes
D5	<i>revapp argST runST</i>	Yes
E η -EXPANSION		
	assuming $k : \{A : U\} \rightarrow A \rightarrow List A \rightarrow A$, $h : Int \rightarrow \{A : U\} \rightarrow A \rightarrow A$ and $lst : List (\{A : U\} \rightarrow Int \rightarrow A \rightarrow A)$	
E1	$k h lst$	No
E2	$k (\lambda x. h x) lst$	Yes*
E3	$\lambda (r : (\{A : U\} \rightarrow A \rightarrow \{B : U\} \rightarrow B \rightarrow B) \rightarrow Int). r (\lambda x y. y)$	Yes

“Yes*” means that our system can infer a type for the expression, but the lack of let-generalization yields unsolved metas in the type.

Fig. 7. Elaboration benchmark from [Serrano et al. 2018].

with a subtyping bound. In contrast, FCIF makes no promises about principal typing — which is not feasible with dependent types — and does no generalization, but it can still infer a type for *id* which can be later constrained to *IdTy*.

7.3 Polymorphic Argument Inference and Curried Application Insertion

FCIF cannot infer the $B1$ and $B2$ cases on Figure 7, which would require inferring polymorphic types for function arguments. In very simple cases, FCIF is able to infer (ambiguous) polymorphic argument types. For example, it can infer a type for $(\lambda f. f\{zero\})$. Assuming $Pair : U \rightarrow U \rightarrow U$ and $pair : \{AB : U\} \rightarrow A \rightarrow B \rightarrow Pair\ A\ B$, it can also infer type for the following:

$$\lambda f\ x. pair\ (f\ \{x\}\ x)\ (f\ x)$$

Here, FCIF can proceed because it first encounters $(f\ \{x\}\ x)$, where it learns that f should have two arguments, where the first one is implicit. Armed with this knowledge, it applies f to a fresh meta when elaborating $(f\ x)$. The following fails though:

$$\lambda f\ x. pair\ (f\ x)\ (f\ \{x\}\ x)$$

Here, FCIF decides upon elaborating $(f\ x)$ that f is a function whose first argument is explicit, and fails at $f\ \{x\}$. This is similar to the implicit λ insertion issue, but here we would have to refine implicit applications, i.e. cover issue (2).

We can try inserting an unknown curried application after the first f , which could be later refined. Below we present an extension of insert which accomplishes this. We omit metacontexts and metasubstitutions for brevity.

```

insert (t,  $\alpha\ \bar{u}$ )  $\equiv$  do
  A  $\leftarrow$  freshMeta $_{\Gamma}$  Tel
  B  $\leftarrow$  freshMeta $_{\Gamma, x:\text{Rec}\ A}\ U$ 
  unify ( $\alpha\ \bar{u}$ ) ( $\{x : \bar{A}\} \rightarrow B$ )
  newConstancy $_{\Gamma, x:\text{Rec}\ A}\ B$ 
  u  $\leftarrow$  freshMeta $_{\Gamma}$  (Rec A)
  return (t {u :  $\bar{A}$ }, B[x  $\mapsto$  u])

```

However, this alone does not work: in $t\ \{u : \bar{A}\}$, we return a term applied to a record with *unknown telescope type*, which later may give rise to unification problems where metas are applied to such unknown records. Our current unification algorithm does not handle these cases meaningfully, and we are not aware of any applicable algorithm in prior literature.

A robust system for this would give a solution for polymorphic argument inference. We do not expect that this is easy. The only prior work which supports this feature is MLF, which does not support dependent types, but whose type inference algorithms are already very complex.

We nevertheless sketch a solution which could be fleshed out in future work. We work with the following example:

let $f : IdTy \rightarrow Bool = \lambda _ . true$ **in** $\lambda x. f\ x$

Plain FCIF will attempt to elaborate $f\ x$ to $f\ (\lambda\{A\}. x)$, but fails to infer a type for x , because A is not in the scope of x 's type. FCIF with curried application insertion instead returns $f\ (\lambda\{A\}. x\ \{\alpha_2\ x\ A : \bar{\alpha}_0\})$ from the same expression, and we also have $x : \{y : \bar{\alpha}_0\} \rightarrow \alpha_1\ \{y : \bar{\alpha}_0\}$. However, at this point FCIF still needs to unify expected and inferred types, so it needs to compute

$$\text{unify}\ (\alpha_1\ \{\alpha_2\ x\ A : \bar{\alpha}_0\})\ (A \rightarrow A)$$

This does not fall into any pattern fragment, and has no most general solution. However, we may make several assumptions about unknown telescopes in spines. We can view this as an extension of Miller's pattern unification [Miller 1991] with unknown spine slices: $\alpha_2\ x\ A$ stands for a part of the spine which may contain bound variables x and A . We make the following assumption: α_2

may be only solved with *order-preserving embeddings*, so the four potential solutions are $(\lambda x A. [])$, $(\lambda x A. (A :: []))$, $(\lambda x A. (x :: []))$ and $(\lambda x A. (x :: A :: []))$. On the one hand, this rules out solutions which contain arbitrary terms or non-linear variable occurrences, which would yield unsolvable spines⁴. On the other hand, we probably do not throw away many useful solutions, because all curried applications arise from elaboration, and we do not have to deal with arbitrary such applications coming from the surface language. In short, we try to keep the invariant that metas returning in record types are only ever solved with order-preserving embeddings.

Out of the four possible embeddings, $(\lambda x A. (x :: A :: []))$ and $(\lambda x A. (x :: []))$ are ruled out because they would require a cyclic solution for α_0 , and $(\lambda x A. [])$ is ruled out because it would yield the unsolvable unify $\alpha_1 (A \rightarrow A)$ problem. Thus, $(\lambda x A. (A :: []))$ is the unique order-preserving embedding solution in this case, which does yield the expected elaboration output. We note though that “unique order-preserving embedding” does not yet provide an efficient algorithmic implementation, and we would need to be smarter in considering particular unification problems, to avoid backtracking through possible solutions.

8 CONCLUSION AND FUTURE WORK

In type theory, it is a common endeavor to search for theories and features which confer the greatest amount of expressive power for the least amount of formal and conceptual complexity. In relation to elaboration and inference algorithms, we would similarly like to use tools and concepts with high power-to-weight ratio. Sometimes core theories need to be extended to allow more powerful elaboration. Certainly, type inference would be cumbersome without the notion of metavariables. Likewise, contextual metavariables are essential for elaboration in the presence of type dependencies.

In this paper we propose the concept of strictly curried functions, which supports elaboration of first-class implicit functions. It seems that trying to solve this problem by fiddling with postponing and heuristics, without extending the core theory, does not really cut it. In a dependently typed spine $t u v$, when elaborating v we already want to have an elaborated version of u with computational behavior. While contextual metavariables yield a nice modal type theory which computes in the presence of *unknown terms*, our curried functions extend this to a system which also computes in the presence of unknown *implicit insertions*. Since implicit insertions are in a way also just unknown terms, we can reuse most of the infrastructure of contextual/crisp type theory, and only add modest extensions.

Often it is the case that constructions on dependent type theory are forced to be more principled and structured than constructions on less powerful theories, because the many interactions and intricacies leave less room for whims. In our case, we believe that focusing on dependent type theory helped us home in on the essential parts of the problem, and it is likely that restrictions of our solution would be also favorably simple and powerful in simpler settings such as System F.

In future work, we would like to

- Investigate extending elaboration with curried application insertions, along the lines of Section 7.3.
- Formalize elaboration and unification from an algebraic/categorical perspective, in particular give algebraic definitions for the core theories.
- Investigate implementing features of our elaborator in production systems such as Agda. Also, investigate simplifying our algorithm to non-dependent (e.g. System F) settings.
- Investigate completeness and properties related to inferability.

⁴Technically, we could consider permutations such as $\lambda x A. (A :: x :: [])$ too, but these are superfluous, since nothing in our system depends on the argument ordering of metavariable solutions.

REFERENCES

- Andreas Abel. 2013. *Normalization by Evaluation: Dependent Types and Impredicativity*. Ph.D. Dissertation. Ludwig-Maximilians-Universität München. Habilitation thesis.
- Andreas Abel, Joakim Öhman, and Andrea Vezzosi. 2018. Decidability of conversion for type theory in type theory. *Proc. ACM Program. Lang.* 2, POPL (2018), 23:1–23:29. <https://doi.org/10.1145/3158111>
- Andreas Abel and Brigitte Pientka. 2011. Higher-Order Dynamic Pattern Unification for Dependent Types and Records. In *Typed Lambda Calculi and Applications - 10th International Conference, TLCA 2011, Novi Sad, Serbia, June 1-3, 2011. Proceedings (Lecture Notes in Computer Science, Vol. 6690)*, C.-H. Luke Ong (Ed.). Springer, 10–26. https://doi.org/10.1007/978-3-642-21691-6_5
- Thorsten Altenkirch and Ambrus Kaposi. 2016. Type theory in type theory using quotient inductive types. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav Bodik and Rupak Majumdar (Eds.). ACM, 18–29. <https://doi.org/10.1145/2837614.2837638>
- Lars Birkedal, Ranald Clouston, Bassel Mannaa, Rasmus Ejlers Møgelberg, Andrew M. Pitts, and Bas Spitters. 2020. Modal dependent type theory and dependent right adjoints. *Math. Struct. Comput. Sci.* 30, 2 (2020), 118–138. <https://doi.org/10.1017/S0960129519000197>
- Didier Le Botlan and Didier Rémy. 2003. ML^F : raising ML to the power of system F. *SIGPLAN Notices* 38, 9 (2003), 27–38. <https://doi.org/10.1145/944746.944709>
- Simon Boulrier, Pierre-Marie Pédro, and Nicolas Tabareau. 2017. The Next 700 Syntactical Models of Type Theory. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs (Paris, France) (CPP 2017)*. ACM, New York, NY, USA, 182–194. <https://doi.org/10.1145/3018610.3018620>
- Thierry Coquand. 1996. An Algorithm for Type-Checking Dependent Types. *Sci. Comput. Program.* 26, 1-3 (1996), 167–177. [https://doi.org/10.1016/0167-6423\(95\)00021-6](https://doi.org/10.1016/0167-6423(95)00021-6)
- Joshua Dunfield and Neelakantan R. Krishnaswami. 2013. Complete and easy bidirectional typechecking for higher-rank polymorphism. In *ACM SIGPLAN International Conference on Functional Programming, ICFP’13, Boston, MA, USA - September 25 - 27, 2013*, Greg Morrisett and Tarmo Uustalu (Eds.). ACM, 429–442. <https://doi.org/10.1145/2500365.2500582>
- Peter Dybjer. 1995. Internal Type Theory. In *Types for Proofs and Programs, International Workshop TYPES’95, Torino, Italy, June 5-8, 1995, Selected Papers (Lecture Notes in Computer Science, Vol. 1158)*, Stefano Berardi and Mario Coppo (Eds.). Springer, 120–134. https://doi.org/10.1007/3-540-61780-9_66
- Richard A. Eisenberg. 2016. Dependent Types in Haskell: Theory and Practice. *CoRR abs/1610.07978* (2016). arXiv:1610.07978 <http://arxiv.org/abs/1610.07978>
- Richard A. Eisenberg, Stephanie Weirich, and Hamidhasan G. Ahmed. 2016. Visible Type Application. In *Programming Languages and Systems - 25th European Symposium on Programming, ESOP 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings (Lecture Notes in Computer Science, Vol. 9632)*, Peter Thiemann (Ed.). Springer, 229–254. https://doi.org/10.1007/978-3-662-49498-1_10
- Adam Michael Gundry. 2013. *Type inference, Haskell and dependent types*. Ph.D. Dissertation. University of Strathclyde, Glasgow, UK. http://oleg.lib.strath.ac.uk/R/?func=dbin-jump-full&object_id=22728
- Robert Harper. 2016. *Practical foundations for programming languages*. Cambridge University Press.
- Marcus Johansson and Jesper Lloyd. 2015. *Eliminating the problems of hidden-lambda insertion-Restricting implicit arguments for increased predictability of type checking in a functional programming language with depending types*. Master’s thesis. Chalmers University of Technology.
- Ambrus Kaposi, Simon Huber, and Christian Sattler. 2019. Gluing for Type Theory. In *4th International Conference on Formal Structures for Computation and Deduction, FSCD 2019, June 24-30, 2019, Dortmund, Germany (LIPIcs, Vol. 131)*, Herman Geuvers (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 25:1–25:19. <https://doi.org/10.4230/LIPIcs.FSCD.2019.25>
- Daan Leijen. 2008. HMF: simple type inference for first-class polymorphism. In *Proceeding of the 13th ACM SIGPLAN international conference on Functional programming, ICFP 2008, Victoria, BC, Canada, September 20-28, 2008*, James Hook and Peter Thiemann (Eds.). ACM, 283–294. <https://doi.org/10.1145/1411204.1411245>
- Daan Leijen. 2009. Flexible types: robust type inference for first-class polymorphism. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, Zhong Shao and Benjamin C. Pierce (Eds.). ACM, 66–77. <https://doi.org/10.1145/1480881.1480891>
- Daniel R. Licata, Ian Orton, Andrew M. Pitts, and Bas Spitters. 2018. Internal Universes in Models of Homotopy Type Theory. In *3rd International Conference on Formal Structures for Computation and Deduction, FSCD 2018, July 9-12, 2018, Oxford, UK (LIPIcs, Vol. 108)*, Hélène Kirchner (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 22:1–22:17. <https://doi.org/10.4230/LIPIcs.FSCD.2018.22>
- Dale Miller. 1991. A Logic Programming Language with Lambda-Abstraction, Function Variables, and Simple Unification. *J. Log. Comput.* 1, 4 (1991), 497–536. <https://doi.org/10.1093/logcom/1.4.497>

- Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. 2008. Contextual modal type theory. *ACM Trans. Comput. Log.* 9, 3 (2008), 23:1–23:49. <https://doi.org/10.1145/1352582.1352591>
- Ulf Norell. 2007. *Towards a practical programming language based on dependent type theory*. Ph.D. Dissertation. Chalmers University of Technology.
- Steven Schäfer, Tobias Tebbi, and Gert Smolka. 2015. Autosubst: Reasoning with de Bruijn Terms and Parallel Substitutions. In *Interactive Theorem Proving - 6th International Conference, ITP 2015, Nanjing, China, August 24-27, 2015, Proceedings (Lecture Notes in Computer Science, Vol. 9236)*, Christian Urban and Xingyuan Zhang (Eds.). Springer, 359–374. https://doi.org/10.1007/978-3-319-22102-1_24
- Alejandro Serrano, Jurriaan Hage, Simon Peyton Jones, and Dimitrios Vytiniotis. 2020. A quick look at impredicativity. (January 2020). <https://www.microsoft.com/en-us/research/publication/a-quick-look-at-impredicativity/> In submission.
- Alejandro Serrano, Jurriaan Hage, Dimitrios Vytiniotis, and Simon Peyton Jones. 2018. Guarded impredicative polymorphism. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2018, Philadelphia, PA, USA, June 18-22, 2018*, Jeffrey S. Foster and Dan Grossman (Eds.). ACM, 783–796. <https://doi.org/10.1145/3192366.3192389>
- Jonathan Sterling. 2019. Algebraic Type Theory and Universe Hierarchies. *CoRR* abs/1902.08848 (2019). arXiv:1902.08848 <http://arxiv.org/abs/1902.08848>
- The Univalent Foundations Program. 2013. *Homotopy Type Theory: Univalent Foundations of Mathematics*. <https://homotopytypetheory.org/book>, Institute for Advanced Study.
- Dimitrios Vytiniotis, Simon L. Peyton Jones, Tom Schrijvers, and Martin Sulzmann. 2011. OutsideIn(X) Modular type inference with local assumptions. *J. Funct. Program.* 21, 4-5 (2011), 333–412. <https://doi.org/10.1017/S0956796811000098>
- Dimitrios Vytiniotis, Stephanie Weirich, and Simon L. Peyton Jones. 2006. Boxy types: inference for higher-rank types and impredicativity. In *Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming, ICFP 2006, Portland, Oregon, USA, September 16-21, 2006*, John H. Reppy and Julia L. Lawall (Eds.). ACM, 251–262. <https://doi.org/10.1145/1159803.1159838>
- Dimitrios Vytiniotis, Stephanie Weirich, and Simon L. Peyton Jones. 2008. FPH: first-class polymorphism for Haskell. In *Proceeding of the 13th ACM SIGPLAN international conference on Functional programming, ICFP 2008, Victoria, BC, Canada, September 20-28, 2008*, James Hook and Peter Thiemann (Eds.). ACM, 295–306. <https://doi.org/10.1145/1411204.1411246>
- Paweł Wieczorek and Dariusz Biernacki. 2018. A Coq formalization of normalization by evaluation for Martin-Löf type theory. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2018, Los Angeles, CA, USA, January 8-9, 2018*, June Andronick and Amy P. Felty (Eds.). ACM, 266–279. <https://doi.org/10.1145/3167091>