
FORMALISING GHC'S TYPE SYSTEM

A RIGOROUS AND MACHINE CHECKED FORMULATION OF $\text{OUTSIDEIN}(X)$
IN A DEPENDENTLY TYPED PROGRAMMING LANGUAGE

A THESIS BY

LIAM O'CONNOR-DAVIS

SUPERVISED BY

MANUEL M. T. CHAKRAVARTY

SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE HONOURS DEGREE OF
Bachelor of Science (Computer Science)

*School of Computer Science and Engineering
University of New South Wales*

THE UNIVERSITY OF
NEW SOUTH WALES



SYDNEY • AUSTRALIA

OCTOBER 16TH 2012

Abstract

GHC, a state of the art Haskell compiler (Marlow (Ed.), 2010), offers numerous extensions to the standard Haskell type system (Schrijvers *et al.*, 2009; Yorgey *et al.*, 2012; Kiselyov *et al.*, 2010; Peyton Jones *et al.*, 2007). Each of these extensions is usually specified only semi-formally, and only in isolation. Very little work has been done examining type system properties when multiple type system extensions are combined, which is the scenario actually being faced by GHC developers. To address this, the GHC team published `OUTSIDEIN(X)`, a mostly-rigorous formulation of GHC's type inference system (Vytiniotis *et al.*, 2011), which encompasses every type system extension developed for GHC to date.

We formalise `OUTSIDEIN(X)` in a mechanical proof assistant, in order to provide a body of formal work upon which future extensions can be developed. By using a mechanical proof assistant we not only ensure correctness of our proofs and complete rigour in our definitions, but also make possible the incremental development of the formal work alongside the more practically-minded type checker implementation in GHC. This additional accessibility will hopefully prevent further extensions from being developed without regard to the effect such an extension may have on other parts of the type system.

Our formalisation is developed in Agda (Norell, 2008). As a dependently typed programming language which enforces totality, Agda doubles as a proof assistant (Howard, 1980). It is still under heavy development, and is quite experimental. By formalising `OUTSIDEIN(X)` in Agda, we demonstrate its readiness for type system work, and also provide an example to encourage further type systems research in Agda.

Contents

1	Introduction	5
1.1	Related Work	5
1.2	Why mechanise <code>OUTSIDEIN(X)</code> ?	5
1.3	Why Agda?	6
1.4	A Brief Introduction to Agda	6
1.5	Overview	7
2	Making <code>OutsideIn(X)</code> rigorous	8
2.1	Regarding <code>let</code> -generalisation	8
2.2	Phases of Inference	9
2.3	The Extended Constraint Language	10
2.4	Fresh Variables	11
2.5	The Prenexer	13
2.6	Solver and Separator	14
2.7	Top Level Rules	15
3	Encoding Types and Terms	16
3.1	Names	16
3.2	Indexed Terms	17
3.3	Nested Data Types	18
3.4	Separator and Prenexer	22
3.5	Expressions	23
3.6	Constraint Generator and Solver	25

4 Simple Instantiation	28
4.1 Unification	28
4.2 Solving	29
4.3 Results	30
5 Conclusions and Future Work	31
5.1 Proofs	31
5.2 Instantiating X for Haskell	32
5.3 Summary of Contributions	32
References	33

“There are many ways of trying to understand programs. People often rely too much on one way, which is called ‘debugging’ and consists of running a partly-understood program to see if it does what you expected. Another way, which ML advocates, is to install some means of understanding in the very programs themselves.”

— Robin Milner, *Foreword to The Little MLer*

1 Introduction

Haskell is a purely functional programming language, with a type system that supports algebraic data types, type inference, parametric (higher-kinded) polymorphism, and type class constraints (Marlow (Ed.), 2010). In recent years, the developers of GHC, a prominent Haskell compiler, have implemented a variety of extensions to this type system, with the aim of providing greater expressiveness, ease of use, or static verification capabilities. Some are straightforward, such as generalising type classes to type relations via multi-parameter type classes. Some require more significant extensions to the type system, such as type families (Kiselyov *et al.*, 2010) and the earlier functional dependencies extension; some make type inference significantly more difficult and are major extensions, such as GADTs (Schrijvers *et al.*, 2009), impredicative polymorphism and arbitrary-rank types (Peyton Jones *et al.*, 2007).

While GHC accommodates all of these extensions simultaneously, the papers that introduce each one discuss type inference and type checking only in isolation, and sometimes quite informally. This makes the properties of GHC’s type reconstruction algorithm difficult to determine when multiple extensions are combined.

As a first step towards solving this problem, the GHC team (specifically Vytiniotis, Peyton Jones, Schrijvers and Suzmann) published `OUTSIDEIN(X)`, a modular type inference system that accommodates all of these extensions (and possibly more), along with soundness and principality proofs (Vytiniotis *et al.*, 2011).

Our work aims to more rigorously formalise `OUTSIDEIN(X)` in the dependently-typed programming language *cum* proof assistant Agda 2 (Norell, 2008).

1.1 Related Work

The `OUTSIDEIN(X)` system itself is similar in presentation to the `HM(X)` system, a parameterised formalisation of ML’s type inference, first presented in (Odersky *et al.*, 1997) and more rigorously formalised in (Pottier & Rémy, 2005). It is the culmination of years of work and is the latest in a series of type inference systems, starting with the original inference system for GADTs in Haskell based on “wobbly types” (Schrijvers *et al.*, 2009), right up to the very similar `LHM(X)` system presented in (Vytiniotis *et al.*, 2010).

Very little work has been done on formalising type inference in a proof assistant, and what little work that has been done is primarily focused on type inference for ML. Nipkow and Narachewski have formalised Milner’s original \mathcal{W} algorithm for the HM calculus in Isabelle/HOL and proven soundness (Naraschewski & Nipkow, 1999), and Dubois *et al.* simultaneously performed a similar verification in Coq (Dubois & Ménéssier-Morain, 1999) for the purposes of developing a certified ML compiler (Dubois, 2000). Surprisingly, no type inference algorithms have been formalised in Agda before. Even more surprisingly, no algorithm which includes support for GADTs or other advanced type system features has ever been formalised in a proof assistant, to our knowledge.

The techniques we use for term representation are drawn heavily from Bird’s observation of monadic structure in syntax trees (Bird & Paterson, 1999), a concept first explored in (Bellegard & Hook, 1994). We also draw several representational tricks from works in generic programming using dependent types (Morris *et al.*, 2004) and McBride’s work on structurally recursive unification (McBride, 2003), which is used directly in our simple instantiation.

1.2 Why mechanise `OutsideIn(X)`?

Any handwritten formalisation or proof will likely lack the amount of rigour necessary to be accepted by a proof assistant, much like a handwritten algorithm is unlikely to be accepted by a programming language compiler. Formalising `OUTSIDEIN(X)` in a proof assistant is therefore more difficult than it appears at first glance. The use of a proof assistant requires us to redesign those parts of the system that are not amenable to automated checking, and to make rigorous all those parts of the original formalisation that are left to the reader’s intuition.

By formalising `OUTSIDEIN(X)` in a proof assistant, we achieve two main goals. Firstly, we make explicit that which was implicit, and to prove that which was assumed in the original `OUTSIDEIN(X)` paper, ensuring that our

formal work stands on solid ground; and secondly, we encourage those developing extensions for the type system to use our work as a foundation for their formalisation, along with the necessary practical implementation of the new extension in GHC’s type checker. By making our formalisation available as code, we hope to mitigate the social problem of formal work on a type system being published in a long paper and subsequently ignored¹.

1.3 Why Agda?

Agda is an interesting choice of proof assistant for this task. This choice was made not simply because Agda is the most familiar to the author, but also because Agda is still quite experimental, and the subject of a great deal of new research. By formalising `OUTSIDEIN(X)` in Agda, we show Agda’s readiness for type systems work, and provide an example for others researching type systems and considering Agda. Similar work has been done in Coq (Dubois & Ménéssier-Morain, 1999) and Isabelle (Naraschewski & Nipkow, 1999), but our formalisation is the first such work in Agda.

1.4 A Brief Introduction to Agda

Agda is a programming language with a concrete syntax similar to Haskell, based on the dependent intuitionistic type theory of Per Martin-Löf (Martin-Löf, 1984). Agda enforces totality by mandating that all functions be structurally recursive², meaning that programs correspond to proofs in a higher order intuitionistic logic. Features include (co)-inductive data types and families, “mix-fix” syntax (Danielsson & Norell, 2011), parameterised modules, “View from the left” style pattern matching (McBride & McKinna, 2004) and compile time proof irrelevance annotations. For a complete tutorial in Agda programming, we defer to the experts (Norell, 2008); the Agda examples in this thesis only require a rudimentary knowledge of Agda’s syntax.

1.4.1 A Key Point of Difference

Unlike other dependently typed theorem provers such as Coq, when working in Agda one does not write a proof script consisting of a series of proof *tactics* which transform or generate a proof *object* (i.e. the dependently typed program); the program or proof is written directly. This has two main implications:

1. *Proof terms are explicit, automation is not available.* There is limited opportunity for automated generation and manipulation of the proof object (i.e. complicated proof tactics) when one writes the proof object code directly. Recently, a new reflection interface was added in Agda version 2.3.0, which allows Agda programs to inspect the current goal and generate solutions for it. It offers a kind of automatic generation of proof objects using Agda itself as a tactic language, however it remains highly experimental and few tactics are available. Our formalisation does not make use of this feature.
2. *Great care must be taken to keep representations manageable.* In most theorem provers, the formal properties we want to prove and the definitions they describe tend to be quite distinct. Because Agda uses the same language to talk about both, we can combine them in ways that would be impossible in Isabelle/HOL or unusual in Coq³. Key properties about our definitions are implied by their structure, rather than independently proven as lemmas, using a variety of definitional tricks. We employ these tricks extensively for our representation of names, substitutions, and abstract syntax trees.

¹ See the new `DataKinds` extension (Yorgey *et al.*, 2012)

² Or, in the case of coinduction, structurally corecursive.

³ Such tricks are certainly possible in Coq, but less commonly employed.

1.5 Overview

There are three main components to our work, each discussed in a separate chapter:

1. Our approach to `OUTSIDEIN(X)` is discussed in chapter two, where we rework various parts of `OUTSIDEIN(X)` that were informally presented originally, providing a more rigorous formalisation.
2. Our encoding of the system in Agda, including our method of term representation, is explained in chapter three. As discussed in the previous chapter, the choice of representation is very important for Agda work.
3. Chapter four introduces a simple instantiation of the X parameter, which allows us not only to sanity-check our definitions, but also to provide a testbed for experimentation with the system.

Lastly, chapter five discusses future directions for this work, including proof work and instantiating the system for Haskell itself.

2 Making OutsideIn(X) rigorous

OUTSIDEIN(X), published in (Vytiniotis *et al.*, 2011), is an approach to type inference approach that supports modular type inference and, most interestingly, *local assumptions*, such as those introduced by pattern matching on a *generalised algebraic data type*, or GADT. A GADT is more general than a regular algebraic data type because its constructors can have substantially more flexible types, including constraints and type variables not mentioned in the constructed type. This allows for a variety of useful features, such as existential quantification and type indexing (Schrijvers *et al.*, 2009).

For example, suppose we had the following GADT.

```
data EqOrShow :: * → * where
  IsEq    :: (Eq τ) ⇒ EqOrShow τ
  IsShow  :: (Show τ) ⇒ EqOrShow τ
```

When we pattern match on a GADT such as this, we introduce a *local assumption*, in this case about the type variable α :

```
f :: EqOrShow α → α → Either String Bool
f IsEq x = Right (x ≡ x)
f IsShow x = Left (show x)
```

Here, the local assumption $Eq\ \alpha$ allows the first alternative to type check, while the local assumption $Show\ \alpha$ allows the second alternative to type check. It is important to note that these assumptions must *remain local*. This is especially apparent when the constraints are contradictory, for example using equality constraints⁴:

```
data IntOrBool :: * → * where
  IsInt    :: (τ ~ Int) ⇒ IntOrBool τ
  IsBool   :: (τ ~ Bool) ⇒ IntOrBool τ
```

A function which pattern matches on this GADT must obviously localise each constraint assumption to each alternative — it is impossible for $\alpha \sim Bool$ and $\alpha \sim Int$ unless $Int \sim Bool$, which is of course not the case⁵.

These local assumptions have historically been difficult to deal with, resulting in lack of principal types. While the general typing rules in a language may allow terms which lack principal types, OUTSIDEIN(X) only infers principal types. If a term lacks a principal type, then type inference will fail — it is not complete. Most similar systems, for example HM(X), are complete but lack support for GADTs and type classes (Pottier & Rémy, 2005).

2.1 Regarding let-generalisation

When it was first published, OUTSIDEIN(X) was not the inference system used by GHC. Some significant changes had to be made to the static semantics of GHC Haskell in order to accommodate it; specifically, generalisation of inferred types in local **let**-expressions was removed.

⁴ More conventionally, these constructors would not use type variables at all, instead each constructor would be parameterised by the concrete types *Int* and *Bool* respectively, however this is desugared into a type variable with an explicit equality constraint (Schrijvers *et al.*, 2009)

⁵ At least, it is not the case in languages outside those whose names contain the letter C and not much else.

Continuing from our earlier example, if we had the following definition:

$$\begin{aligned}
 f &:: \text{IntOrBool } \alpha \rightarrow \alpha \rightarrow \text{Bool} \\
 f \ x \ y &= \mathbf{let} \ g \ z = \mathit{not} \ y \ \mathbf{in} \\
 &\quad \mathbf{case} \ x \ \mathbf{of} \\
 &\quad \quad \mathit{IsInt} \quad \rightarrow \ \text{True} \\
 &\quad \quad \mathit{IsBool} \quad \rightarrow \ g \ ()
 \end{aligned}$$

Most programmers would expect the binding g to give a type error, as it requires y to be of type Bool , *before* that has been established via pattern matching on x ⁶. A principal type does exist, however, for g :

$$g :: \forall \beta. (\alpha \sim \text{Bool}) \Rightarrow \beta \rightarrow \text{Bool}$$

That is, the constraint $\alpha \sim \text{Bool}$, rather than being rejected, is *abstracted over* in the inferred type. Then, at any call site, we must provide evidence that $\alpha \sim \text{Bool}$, which can be done in this example thanks to the pattern matching.

The authors of `OUTSIDEIN(X)` state that abstracting over all inferred constraints imposes a significant complexity cost on the implementation of the type checker, for two main reasons. Firstly, at each call site of such a generalised expression, the (potentially large) set of constraints that have been abstracted over must be shown satisfiable, which, if necessary for every locally bound expression, would be difficult to perform efficiently.

Secondly, it becomes impossible to solve constraints “on-the-fly” in a similar manner to Milner’s \mathcal{W} algorithm (Milner, 1978); instead the compiler must generate all constraints and then solve them in discrete phases. GHC relies on this on-the-fly solving to resolve equality constraints efficiently using mutable type variables (Peyton Jones *et al.*, 2007). They also observe that the principal types inferred by such a method are often highly confusing and result in the type checker accepting almost-certainly erroneous code, such as the above example.

In `OUTSIDEIN(X)`, this generalisation and abstraction of inferred constraints is only performed *at the top level*, i.e. where the type environment is empty. Therefore, the difficulties above disappear. At the local level, *no generalisation is performed at all*. That is, an (unannotated) `let`-binding `let x = y in e` is equivalent to `(λx. e) y`. While a cherished feature of most type inference algorithms since the original \mathcal{W} , `let` generalisation at a local level turns out to be relied upon very rarely in practice, with a total of only 127 lines needing to be modified in the `base` Haskell library consisting of 94,954 lines after this change was made (Vytiniotis *et al.*, 2010).

2.2 Phases of Inference

As mentioned in the previous section, GHC resolves equality constraints as soon as they can be solved for efficiency reasons. Specifically, only those constraints which rely on information not available when they are generated are deferred until after constraint generation for solving. All other constraints are solved immediately. `OUTSIDEIN(X)`, however, presents inference as two separate phases:

1. *Generate* constraints (according to a set of syntax-directed rules), combining them into one large overall constraint.
2. *Solve* the constraint, using unification and standard constraint-solving methods.

The method used by GHC then can be viewed then as interleaving these two phases, whereas `OUTSIDEIN(X)` (at least in theory) keeps them distinct. Our formal work is derived from `OUTSIDEIN(X)`, not GHC, and for such formal work we do not particularly care about the *efficiency* of type inference, but rather the various properties we can prove about it. For this reason our formalisation also separates these phases. Indeed, we add *additional* phases to this process in order to more clearly separate the solver phase; as well as to deal with our extended constraint language, discussed in the next section.

⁶ Note that `not :: Bool → Bool`

Term variables	\in	x, y, z, f, g, h
Type variables	\in	a, b, c
Unification type variables	\in	$\alpha, \beta, \gamma, \delta$
Data constructors	\in	K
Type constructors	\in	\mathbf{T}
	v	$::= K \mid x$
Expressions	e	$::= v \mid \lambda x. e \mid e_1 e_2$ $\mid \text{case } e \text{ of } \{ \overline{K \bar{x} \rightarrow e} \}$ $\mid \text{let } x = e_1 \text{ in } e_2 \mid \text{let } x :: \sigma = e_1 \text{ in } e_2$
Type schemes	σ	$::= \forall \bar{a}. Q \Rightarrow \tau$
Programs	$prog$	$::= f :: \sigma = e, prog \mid f = e, prog \mid \epsilon$
Constraints*	Q	$::= \epsilon \mid Q_1 \wedge Q_2 \mid \tau_1 \sim \tau_2 \mid \dots$
Extended constraints	C	$::= Q \mid C_1 \wedge' C_2 \mid \exists \bar{\beta}. Q \supset C \mid \exists \alpha. C$
Monotypes*	τ	$::= tv \mid \mathbf{T} \bar{\tau} \mid \tau_1 \rightarrow \tau_2 \mid \dots$
	tv	$::= a$
Environments	Γ	$::= \epsilon \mid (v : \sigma), \Gamma$
Axiom Schema*	\mathcal{Q}	$::= \dots$

Γ_0 : Types of data constructors*
 $K : \forall \bar{a} \bar{b}. Q \Rightarrow \bar{\tau} \rightarrow \mathbf{T} \bar{a}$ * — part of the parameter X

Fig. 1: Our updated syntax, slightly extended from $\text{OUTSIDEIN}(X)$

2.3 The Extended Constraint Language

Constraint generation in $\text{OUTSIDEIN}(X)$ is specified independently of the exact constraint system or type system used — the constraint and type terms are part of the X parameter. Clearly, some intuitive conditions must be met by these components, summarised in two additional parameters: an *entailment relation* of global constraint schema to locally inferred constraints, and certain *simplifier conditions* that ensure that the provided solver behaves consistently with this entailment relation. Specifically, the constraint system must include constraint conjunction and equality constraints (see Fig. 1) that behave as one would expect, and any locally inferred constraints must be resolved trivially if they restate an axiom in a global constraint scheme (see Fig. 2). This parameterisation of the system is similar to $\text{HM}(X)$, the parameterised extension of ML’s type inference presented in (Odersky *et al.*, 1997)⁷, with the addition of global *axiom schema*, which are designed to accommodate Haskell’s type classes. Specifically, type class instances can generate top-level implication constraints such as *Show* $\alpha \Rightarrow \text{Show } [\alpha]$. In $\text{OUTSIDEIN}(X)$, these top-level constraints are expressed in axiom schema and do not form part of the main constraint language.

In order to deal with local assumptions, $\text{OUTSIDEIN}(X)$ extends Q , the original constraint language in X , to an *algorithmic* constraint language C . C is just Q with an additional form, $\exists \bar{\beta}. (Q \supset C)$, where Q is a local assumption, C is a constraint, and $\bar{\beta}$ are the *only* variables that can be unified while solving the constraint $Q \supset C$. The local assumption is defined as a constraint in the *original* constraint language Q specifically, rather than the larger C , as all local assumptions come from the constraint clause of a generalised type signature, either provided by the user or in the type of a generalised data constructor; they are not generated locally by the algorithm itself.

Our presentation of this extended language differs from the presentation in $\text{OUTSIDEIN}(X)$ in several respects. One minor change is that we make the separation between the languages C , Q and \mathcal{Q} much more clear. In particular, we add a new form of conjunction to the extended language C , written $\phi \wedge' \psi$, as the conjunction inherited from Q can, technically, only contain Q -constraints. In addition, we have changed \mathcal{Q} , the language of axiom schema, so that it no longer includes all of Q . Instead, we make no assumptions about the forms that \mathcal{Q} schema may take, and have reformulated the entailment relation to require *both* a \mathcal{Q} -constraint *and* a Q -constraint as context (see Fig. 2). We have also added to the entailment relation requirements for conjunction elimination rules as well as a rule to deal with ϵ -constraints. These rules were inexplicably absent from the original presentation, despite being implicitly used repeatedly in the soundness proof presented in the same work.

⁷ or the more rigorous formalisation presented by Pottier and Remy in their chapter of *Advanced Types and Programming Languages* (Pottier & Rémy, 2005)

$\boxed{Q; Q \Vdash Q}$		
part of the parameter X , subject to the following requirements:		
Tautology	$Q; Q \Vdash \epsilon$	(R ₁)
Reflexivity	$Q; Q \Vdash Q$	(R ₂)
Transitivity	$Q; Q_1 \Vdash Q_2$ and $Q; Q_2 \Vdash Q_3$ implies $Q; Q_1 \Vdash Q_3$	(R ₃)
Substitution	$Q; Q_1 \Vdash Q_2$ implies $\theta Q; \theta Q_1 \Vdash \theta Q_2$ where θ is a type substitution	(R ₄)
Conjunction intro.	$Q; Q \Vdash Q_1$ and $Q; Q \Vdash Q_2$ implies $Q; Q \Vdash Q_1 \wedge Q_2$	(R ₅)
Conjunction elim.	$Q; Q_1 \wedge Q_2 \Vdash Q_1$	(R ₆)
	$Q; Q_1 \wedge Q_2 \Vdash Q_2$	(R ₇)
Type eq. reflexivity	$Q; Q \Vdash \tau \sim \tau$	(R ₈)
Type eq. symmetry	$Q; Q \Vdash \tau_1 \sim \tau_2$ implies $Q; Q \Vdash \tau_2 \sim \tau_1$	(R ₉)
Type eq. transitivity	$Q; Q \Vdash \tau_1 \sim \tau_2$ and $Q; Q \Vdash \tau_2 \sim \tau_3$ implies $Q; Q \Vdash \tau_1 \sim \tau_3$	(R ₁₀)
Type eq. substitutivity	$Q; Q \Vdash \tau_1 \sim \tau_2$ implies $Q; Q \Vdash [\alpha \mapsto \tau_1]\tau \sim [\alpha \mapsto \tau_2]\tau$	(R ₁₁)
$\boxed{Q; Q_{\text{given}}; \overline{\alpha_{\text{tch}}} \xrightarrow{\text{simp}} Q_{\text{wanted}} \rightsquigarrow Q_{\text{residual}}; \theta}$		
part of the parameter X , subject to the following requirements:		
Touchable-aware:	$\text{dom}(\theta) \subseteq \overline{\alpha_{\text{tch}}}$	
Soundness:	$Q; Q_g \xrightarrow{\text{simp}} Q_w \rightsquigarrow Q_r; \theta$ implies $Q; Q_g \wedge Q_r \Vdash \theta Q_w$	
Principality: (Guess-freedom)	$Q; Q_g \xrightarrow{\text{simp}} Q_w \rightsquigarrow Q_r; \theta$ implies $Q; Q_g \wedge Q_w \Vdash Q_r \wedge \mathcal{E}_\theta$ where $\mathcal{E}_\theta = \{(\alpha \sim \tau) \mid [\alpha \mapsto \tau] \in \theta\}$.	

Fig. 2: Entailment relation and simplifier conditions

The constraint solver is also part of the parameter X and therefore can only act on constraints in Q , not the extended constraint language C . Therefore, $\text{OUTSIDEIN}(X)$ includes additional machinery to solve implication constraints given a solver for Q -constraints. In $\text{OUTSIDEIN}(X)$, and in this thesis, the term *simplifier* is used to describe the Q -solver, whereas the term *solver* is reserved for the top-level C -solver.

2.4 Fresh Variables

A common sin against mathematical rigour often committed in type inference literature is that of the magically fresh variable. This is an example taken from a constraint generation rule (for lambda abstractions) in $\text{OUTSIDEIN}(X)$:

$$\frac{\alpha \text{ fresh} \quad \Gamma, (x : \alpha) \vdash e : \tau \rightsquigarrow C}{\Gamma \vdash \lambda x. e : (\alpha \rightarrow \tau) \rightsquigarrow C}$$

These fresh variables must be globally unique and in scope throughout the entire program, despite being summoned *ad-hoc* as constraints are generated; they must be completely unused variable names before being introduced here. Narachewski and Nipkow’s approach to this problem, when they verified \mathcal{W} in Isabelle, was to thread an infinite source of known globally unique variable names (i.e. a natural number n for which all names in $\{N_i \mid i \geq n\}$ are unique and unused) as state through the program, removing a name from the source when a fresh variable was introduced (i.e. incrementing n) (Naraschewski & Nipkow, 1999). While this approach is perhaps closer to how \mathcal{W} would be implemented in a compiler, it has a certain inelegance that complicates Agda definitions of these rules considerably. Specifically, the rules would need to live within a state monad, introducing needless dependency between rule invocations which would otherwise be independent.

Our approach is instead to reuse some machinery that is already present in $\text{OUTSIDEIN}(X)$ for local assumptions. We shall extend the constraint language slightly while generating constraints, and simplify it again before solving them. Specifically, we add another form to the extended constraint language C : an (existential) quantifier for unification variables, which we denote with $\exists\alpha$. C (\exists is used here rather than \exists to distinguish between the two existential quantifiers in C ; \exists is for local assumptions, and \exists is for unification variables). This approach is similar to the existential quantifiers used in (Pottier & Rémy, 2005).

To use these quantifiers, we must first rearrange the constraint generation rules so that the type is viewed as

$$\boxed{
\begin{array}{c}
\Delta; \Gamma \vdash e : \tau \rightsquigarrow C \\
\\
\frac{(v : \forall \bar{a}. Q_1 \Rightarrow \tau_1) \in \Gamma}{\Delta; \Gamma \vdash v : \tau \rightsquigarrow \exists \bar{a}. [\bar{a} \mapsto \bar{\alpha}] Q_1 \wedge' (\tau \sim [\bar{a} \mapsto \bar{\alpha}] \tau_1)} \text{VARCON} \\
\\
\frac{\alpha_1, \alpha_2, \alpha_3, \Delta; \Gamma \vdash e_1 : \alpha_1 \rightsquigarrow C_1 \quad \alpha_1, \alpha_2, \alpha_3, \Delta; \Gamma \vdash e_2 : \alpha_2 \rightsquigarrow C_2}{\Delta; \Gamma \vdash e_1 e_2 : \tau \rightsquigarrow \exists \alpha_1. \exists \alpha_2. \exists \alpha_3. C_1 \wedge' C_2 \wedge' (\alpha_1 \sim (\alpha_2 \rightarrow \alpha_3)) \wedge' (\tau \sim \alpha_3)} \text{APP} \\
\\
\frac{\alpha, \beta, \Delta; \Gamma, (x : \alpha) \vdash e : \beta \rightsquigarrow C}{\Delta; \Gamma \vdash \lambda x. e : \tau \rightsquigarrow \exists \alpha. \exists \beta. C \wedge' (\tau \sim (\alpha \rightarrow \beta))} \text{ABS} \\
\\
\frac{\alpha_1, \alpha_2, \Delta; \Gamma \vdash e_1 : \alpha_1 \rightsquigarrow C_1 \quad \alpha_1, \alpha_2, \Delta; \Gamma, (x : \alpha_1) \vdash e_2 : \alpha_2 \rightsquigarrow C_2}{\Delta; \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \tau \rightsquigarrow \exists \alpha_1. \exists \alpha_2. C_1 \wedge' C_2 \wedge' (\tau \sim \alpha_2)} \text{LET} \\
\\
\frac{\alpha_1, \alpha_2, \Delta; \Gamma \vdash e_1 : \alpha_1 \rightsquigarrow C_1 \quad \alpha_1, \alpha_2, \Delta; \Gamma, (x : \alpha_1) \vdash e_2 : \alpha_2 \rightsquigarrow C_2}{\Delta; \Gamma \vdash \text{let } x :: \tau' = e_1 \text{ in } e_2 : \tau \rightsquigarrow \exists \alpha_1. \exists \alpha_2. C_1 \wedge' C_2 \wedge' (\tau \sim \alpha_2) \wedge' (\alpha_1 \sim \tau')} \text{LETA} \\
\\
\frac{\sigma_1 = \forall \bar{a}. Q \Rightarrow \tau' \quad Q \neq \epsilon \text{ or } \bar{a} \neq \epsilon \quad \bar{\alpha}, \beta_1, \beta_2, \Delta; \Gamma \vdash e_1 : \beta_1 \rightsquigarrow C \quad C_1 = \exists \bar{\alpha}. \exists \epsilon. ([\bar{a} \mapsto \bar{\alpha}] Q \supset C \wedge' \beta_1 \sim [\bar{a} \mapsto \bar{\alpha}] \tau') \quad \beta_1, \beta_2, \Delta; \Gamma, (x : \sigma_1) \vdash e_2 : \beta_2 \rightsquigarrow C_2}{\Delta; \Gamma \vdash \text{let } x :: \sigma_1 = e_1 \text{ in } e_2 : \tau \rightsquigarrow \exists \beta_1. \exists \beta_2. C_1 \wedge' C_2 \wedge' (\tau \sim \beta_2)} \text{GLETA} \\
\\
\frac{\alpha, \beta, \bar{\gamma}, \bar{\delta}, \Delta; \Gamma \vdash e : \alpha \rightsquigarrow C \quad (K_i : \forall \bar{a} \bar{b}. Q_i \Rightarrow \bar{\tau}_i \rightarrow \mathbf{T} \bar{a}) \in \Gamma \quad \alpha, \beta, \bar{\gamma}, \bar{\delta}, \bar{\rho}, \Delta; \Gamma, (x_i : [b \mapsto \rho][a \mapsto \gamma] \tau_i) \vdash e_i : \delta_i \rightsquigarrow C_i}{\Delta; \Gamma \vdash \text{case } e \text{ of } \{\bar{K}_i \bar{x}_i \rightarrow e_i\} : \tau \rightsquigarrow \exists \alpha. \exists \beta. \exists \bar{\gamma}. \exists \bar{\delta}. C \wedge' (\mathbf{T} \bar{\gamma} \sim \alpha) \wedge' (\bigwedge C'_i) \wedge' (\tau \sim \beta)} \text{CASE} \\
\text{where } C'_i = \begin{cases} C_i \wedge' \delta_i \sim \beta & \text{if } \bar{b}_i = \epsilon \text{ and } Q_i = \epsilon \\ \exists \bar{\rho}. \exists \epsilon. ([\bar{b} \mapsto \rho][\bar{a} \mapsto \bar{\gamma}] Q_i) \supset C_i \wedge' \delta_i \sim \beta & \text{otherwise} \end{cases}
\end{array}
}$$

Fig. 3: Constraint Generation Rules (using our new quantifier)

input, rather than output. This does not affect the algorithm significantly — types are always just a single metavariable⁸, and the exact form of the inferred type is instead indicated by an explicit equality constraint. This means that the only place where these new unification variables are mentioned is within the generated constraint, and not within the type:

$$\frac{\alpha \text{ fresh} \quad \beta \text{ fresh} \quad \Gamma, (x : \alpha) \vdash e : \beta \rightsquigarrow C}{\Gamma \vdash \lambda x. e : \tau \rightsquigarrow C \wedge (\tau \sim (\alpha \rightarrow \beta))}$$

Then, we can simply add an environment of available type variables Δ to the constraint generation judgement, and replace the fresh variable introductions with quantifiers (where α and β are not in the environment Δ):

$$\frac{\alpha, \beta, \Delta; \Gamma, (x : \alpha) \vdash e : \beta \rightsquigarrow C}{\Delta; \Gamma \vdash \lambda x. e : \tau \rightsquigarrow \exists \alpha. \exists \beta. C \wedge (\tau \sim (\alpha \rightarrow \beta))}$$

The changes required to most of the other rules are in a similar vein (See Fig. 3 for a full set). This more rigorous formulation of constraint generation is equivalent to the original presentation, with the added benefit of a formalized notion of fresh variables, and less ambiguity in the presentation of the constraint languages.

⁸ This does not lead to ambiguity dangers, as the system is still syntax-directed

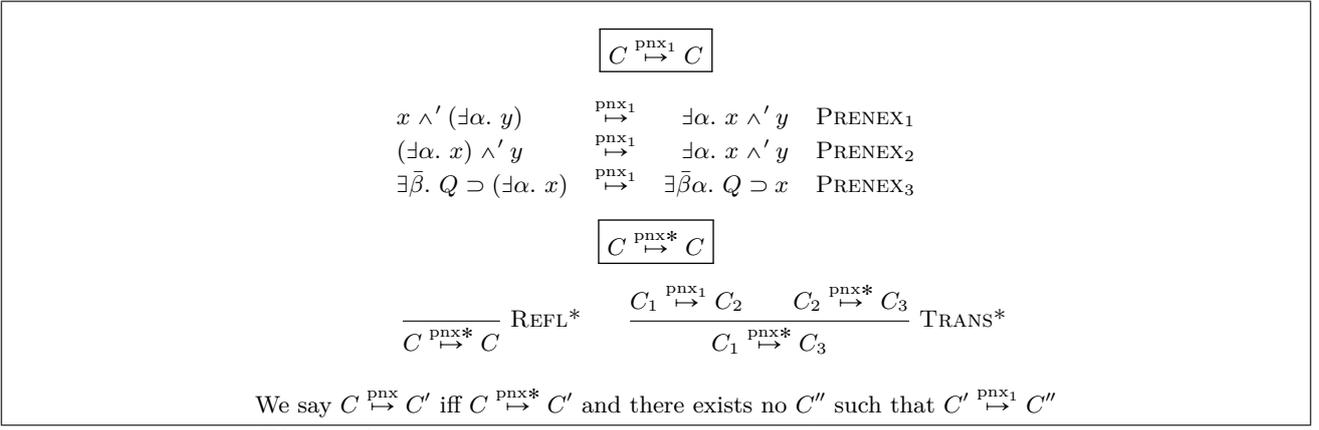


Fig. 4: Prenexer rewrite rules, and their reflexive transitive closure

2.5 The Prenexer

Of particular interest is the interaction between local assumption forms and these fresh variable quantifiers. Here is a rule from the original $\text{OUTSIDEIN}(X)$ formulation where a user has specified a general type for a local `let` binding:

$$\frac{\sigma_1 = \forall \bar{a}. Q_1 \Rightarrow \tau_1 \quad Q_1 \neq \epsilon \text{ or } \bar{a} \neq \epsilon \quad \Gamma \vdash e_1 : \tau \rightsquigarrow C \quad \bar{\beta} = \text{fuv}(\tau, C) - \text{fuv}(\Gamma) \quad C_1 = \exists \bar{\beta}. (Q_1 \supset C \wedge \tau \sim \tau_1) \quad \Gamma, (x : \sigma_1) \vdash e_2 : \tau_2 \rightsquigarrow C_2}{\Gamma \vdash \text{let } x :: \sigma_1 = e_1 \text{ in } e_2 : \tau_2 \rightsquigarrow C_1 \wedge C_2} \text{ ORIGINAL}$$

This rule introduces a local assumption constraint where the variables bound by the quantifier $(\exists \bar{\beta})$ are all free unification variables introduced when generating the constraint for e_1 and τ . With our new quantifiers, there will never be any free unification variables under any circumstances (as any unification variables introduced would have been bound within the constraint C), which makes the updated rule look somewhat odd — the existential quantifier in the local assumption form is empty⁹:

$$\frac{\sigma_1 = \forall \bar{a}. Q \Rightarrow \tau' \quad Q \neq \epsilon \text{ or } \bar{a} \neq \epsilon \quad \bar{\alpha}, \beta_1, \beta_2, \Delta; \Gamma \vdash e_1 : \beta_1 \rightsquigarrow C \quad C_1 = \exists \bar{\alpha}. \exists \epsilon. ([\bar{a} \mapsto \bar{\alpha}]Q \supset C \wedge \beta_1 \sim [\bar{a} \mapsto \bar{\alpha}]\tau') \quad \beta_1, \beta_2, \Delta; \Gamma, (x : \sigma_1) \vdash e_2 : \beta_2 \rightsquigarrow C_2}{\Delta; \Gamma \vdash \text{let } x :: \sigma_1 = e_1 \text{ in } e_2 : \tau \rightsquigarrow \exists \beta_1. \exists \beta_2. C_1 \wedge C_2 \wedge (\tau \sim \beta_2)} \text{ GLETA}$$

We resolve this oddity by introducing a new constraint simplification phase, which is run before solving, called the *prenexer*. The prenexer is responsible for eliminating all of the new \exists -quantifiers introduced by constraint generation, so that solving can proceed unchanged from the original $\text{OUTSIDEIN}(X)$ formulation. This phase is called the *prenexer* because it moves the \exists -quantifiers leftward, towards the front of the constraint expression — a logical formula is considered to be in *prenex normal form* if all quantifiers are moved to the leftmost, outermost position possible.

Despite the name, the prenexer does not necessarily leave constraints in such a normal form due to the presence of other quantifiers in the formula, specifically the existential quantifiers introduced by local assumption constraints. Observe the rule PRENEX_3 , where a new type variable is bound within an implication constraint. Here, the variable bound by the $\exists \alpha$ quantifier is added to the set of variables bound by the existential quantifier $\exists \bar{\beta}$. This rule therefore ensures that local assumption forms again bind all unification variables introduced within the implication, as they do in the original $\text{OUTSIDEIN}(X)$ formulation.

Ordinarily with an explicitly named representation of type variables, we would have to take care when performing these rewrites that no name conflicts occur. As our representation is based on de Bruijn indices, this is quite

⁹ We also have to introduce fresh unification variables for all those bound in the user's general type (\bar{a}). As they are bound *outside* the implication constraint, they are treated as skolem variables within it. Therefore, this change does not affect the semantics of the original rule.

implication constraints	$I ::= \epsilon \mid \exists^I \bar{\alpha}. Q \supset C \mid I \wedge^I I$
separated constraints	$\mathcal{C} ::= Q \cdot I$
implic $[C_1 \wedge^I C_2]$	$=$ implic $[C_1] \wedge^I$ implic $[C_2]$
implic $[Q]$	$= \epsilon$
implic $[\exists \bar{\beta}. Q \supset C]$	$= \exists^I \bar{\beta}. Q \supset$ sep $[C]$
simple $[C_1 \wedge^I C_2]$	$=$ simple $[C_1] \wedge$ simple $[C_2]$
simple $[Q]$	$= Q$
simple $[\exists \bar{\beta}. Q \supset C]$	$= \epsilon$
sep $[C]$	$=$ simple $[C] \cdot$ implic $[C]$

$$\mathcal{Q}; Q_{\text{given}}; \bar{\alpha}_{\text{tch}} \xrightarrow{\text{solv}} \mathcal{C}_{\text{wanted}} \rightsquigarrow Q_{\text{residual}}; \theta$$

$$\frac{\mathcal{Q}; Q_g; \bar{\alpha} \xrightarrow{\text{simp}} Q \rightsquigarrow Q_r; \theta \quad \forall ((\exists^I \bar{\beta}_i. Q_i \supset C_i) \in I). \mathcal{Q}; Q_g \wedge Q_r \wedge Q_i; \bar{\beta}_i \xrightarrow{\text{solv}} C_i \rightsquigarrow \epsilon; \theta_i}{\mathcal{Q}; Q_g; \bar{\alpha} \xrightarrow{\text{solv}} Q \cdot I \rightsquigarrow Q_r; \theta}$$

Fig. 5: Solver, separated constraints and separator functions.

mechanical. As each quantifier is moved out, we adjust the indices within each expression accordingly, thus ensuring that names remain unique (See chapter 3 for details).

Once all \exists quantifiers have been moved by the prenexer, we will have an expression of the form $\exists \bar{\alpha}. C$ where C consists only of the extended syntax used originally in $\text{OUTSIDEIN}(X)$, and all fresh names introduced by the constraint generation have been made globally unique — exactly the semantics of the informal **fresh** constructor used previously. Therefore, after rewriting, we can simply pass C to the subsequent solver phase, using $\bar{\alpha}$ as the set of variables the solver may unify (see the top level rules in Fig. 6).

2.6 Solver and Separator

Our presentation of the solver infrastructure is somewhat different to the solver infrastructure in the original paper. Here we introduce a new form of *separated* C -constraint, denoted \mathcal{C} , which contains a vanilla Q -constraint, already solvable by the provided simplifier, and a special I constraint, which contains only local assumption implications (whose bodies are, in turn, separated \mathcal{C} constraints, see Fig. 5). Our solver operates on these separated constraints rather than the C -constraints directly, so we introduce a function **sep** : $C \rightarrow \mathcal{C}$, defined in terms of helper functions **simple** : $C \rightarrow Q$ and **implic** : $C \rightarrow I$, which forms an additional phase in the inference algorithm.

In the original presentation of the $\text{OUTSIDEIN}(X)$ solver, the separation of constraints into simple and implication components was performed inline with the main solver rule, rather than in a separate phase. This makes the termination argument for the main solving rule slightly less clear, as it must be established that, for all C , **simple** $[C]$ and **implic** $[C]$ are no larger than C itself. This, while obvious to a human observer, is not so obvious to a proof assistant. We simply sidestep any termination trouble by encoding the separation of constraints as a separate phase, rather than interleaving them as in the original presentation.

Our main solving judgement is of the following form:

$$\mathcal{Q}; Q_{\text{given}}; \bar{\alpha}_{\text{tch}} \xrightarrow{\text{solv}} \mathcal{C}_{\text{wanted}} \rightsquigarrow Q_{\text{residual}}; \theta$$

This can be read as, “Given the axiom scheme \mathcal{Q} and constraint Q_{given} , the constraint $\mathcal{C}_{\text{wanted}}$ is simplified to Q_{residual} producing the substitution θ where $\text{dom}(\theta) \subseteq \bar{\alpha}_{\text{tch}}$.” Note that while the solver is not obliged to solve *all* constraints, the remaining residual constraint is a Q -constraint, not a \mathcal{C} -constraint, which means that, at the very least, all implication constraints must be resolved. This is important, as the constraints left residual from the solver are abstracted over when generalising on top-level definitions (see Fig. 6), and C constraints do not form part of the constraint language available to the user of the language — It would be highly unusual for a type inference algorithm to provide type signatures that the user of the language could not express themselves!

$$\boxed{
\begin{array}{c}
\mathcal{Q}; \Delta; \Gamma \vdash \text{prog} \\
\\
\frac{}{\mathcal{Q}; \Delta; \Gamma \vdash \epsilon} \text{EMPTY} \quad \frac{\bar{a}, \Delta; \Gamma \vdash e : \tau \rightsquigarrow C \quad C \xrightarrow{\text{pnx}} \exists \bar{\beta}. C' \quad \mathcal{Q}; \mathcal{Q}; \bar{\beta} \xrightarrow{\text{solv}} \text{sep}[C'] \rightsquigarrow \epsilon; \theta \quad \mathcal{Q}; \Delta; \Gamma, (f : \forall \bar{a}. Q \Rightarrow \tau) \vdash \text{prog}}{\mathcal{Q}; \Delta; \Gamma \vdash f :: (\forall \bar{a}. Q \Rightarrow \tau) = e, \text{prog}} \text{BINDA} \\
\\
\frac{\gamma, \Delta; \Gamma \vdash e : \gamma \rightsquigarrow C \quad C \xrightarrow{\text{pnx}} \exists \bar{\beta}. C' \quad \mathcal{Q}; \mathcal{Q}; \gamma, \bar{\beta} \xrightarrow{\text{solv}} \text{sep}[C'] \rightsquigarrow Q_r; \theta \quad \bar{\alpha} = \text{free unification variables in } Q_r, \theta \gamma \quad \mathcal{Q}; \Delta; \Gamma, (f : \forall \bar{\alpha}. Q_r \Rightarrow \theta \gamma) \vdash \text{prog}}{\mathcal{Q}; \Delta; \Gamma \vdash f = e, \text{prog}} \text{BIND}
\end{array}
}$$

Fig. 6: Top level algorithmic rules.

2.7 Top Level Rules

Our presentation of the top-level rules differ significantly from the original $\text{OUTSIDEIN}(X)$ presentation, in order to accommodate the other changes we have made to the system. In particular:

- As we view the type in the constraint generation rule as *input* rather than *output*, we do not need to add an equality constraint in the rule BINDA, reconciling the provided type signature with the generated type. Instead, we simply pass the provided type directly in to constraint generation. Similarly, we cannot simply use the type returned by constraint generation in the rule BIND, but must instead introduce a new “fresh” type variable γ , use it for constraint generation, then apply the solution substitution θ to it in order to determine the type of f .
- As we now have an explicit notion of available type variables in the environment Δ for constraint generation, we add a similar environment here.
- Constraints must be prenexed *and* separated before being solved, which necessitates some additions to the BIND rules.

Despite the added rigour in our version of the $\text{OUTSIDEIN}(X)$ system, some informality still remains in these definitions - in particular, the possibility of name conflicts is ignored, assertions are made informally about substitution domains, and the rule BIND in Figure 6 relies on an informally specified “free unification variables” operation. Naturally, in order to formalise this system in a proof assistant, we still must resolve these issues. All of these problematic elements are eliminated via our representation of terms and names, discussed in the next chapter.

3 Encoding Types and Terms

Constraints and expressions cannot be expressed in Agda simply as a series of data types, because the exact structure of these terms is dependent on the parameter X . C -constraints, which are not part of X , can contain Q -constraints, which are part of it. Similarly, expressions, not part of X , can contain types, which are.

We use Agda’s module system to represent this parameterisation, providing a record type X containing all definitions within the parameter as an argument to our modules. Here is a sketch of the Agda code used to represent such a parameterisation:

```
record X : Set1 where
  field Type : ...
  field QConstraint : ...
  ...

module OutsideIn (x : X) where
  open X(x)
  ...
  data CONSTRAINT ...
  data EXPRESSION ...
```

(In all future code, we will use sans-serif font to refer to elements of the parameter X)

Note that the type of X is Set_1 , rather than the usual type-of-types Set , because the X parameter sits one meta-level higher than the definitions themselves — the type Set cannot contain Set without making Agda inconsistent¹⁰.

Over the course of this chapter, we will gradually refine this sketch into the concrete definitions we use in our formalisation.

3.1 Names

One of the most commonly examined facets of term representation is how to represent variable names. Much literature has been published on the subject, and a wide range of techniques exist. Perhaps the most common is that of the *de Bruijn index* (de Bruijn, 1972), a simple system of assigning numerical indices to binders instead of names. For example, the term $\lambda x. \lambda y. x y$ can be restated with (stack-based) de Bruijn indices as $\lambda. \lambda. 1 0$. These indices are sometimes presented the other way around, where the innermost binder is referenced by the *highest* available index, but for our purposes, this orientation is easier.

These indices make reasoning and manipulating terms substantially easier in many cases: avoiding name clashes when rewriting constraints is simply a matter of small arithmetic operations on indices, and α -equivalent terms are propositionally equal.

Nicolas Pouillard has generalised de Bruijn indices in a series of systems (implemented in Agda, no less), starting with “Nameless, Painless”, published in (Pouillard, 2011). Based on the notion of an *abstract world* of variable names, these systems are designed chiefly to avoid programming errors when working with de Bruijn indices (which, over the years, have established some notoriety for being somewhat difficult beasts to tame). While his approach is certainly not without merit, we feel that using such a library to represent terms in $\text{OUTSIDEIN}(X)$ may needlessly complicate our definitions. The approach we have taken, based on de Bruijn indices, allows us to exploit type-indexing techniques to generate a number of “theorems for free” via parametricity (Wadler, 1989), and expose an elegant categorical structure in our representation. It is not clear that we could retain this simplicity and generality were we to rely on Pouillard’s work.

¹⁰ The inconsistency arises from Russell’s paradox.

3.2 Indexed Terms

The simplest possible representation which uses de Bruijn indices simply uses the full set \mathbb{N} to represent type variables:¹¹

```

record X : Set1 where
  field Type      : Set
        Var       :  $\mathbb{N} \rightarrow \text{Type}$ 
         $\rightarrow'$  :  $\text{Type} \rightarrow \text{Type} \rightarrow \text{Type}$ 
  field QConstraint : Set
         $\epsilon$       : QConstraint
         $\wedge$       :  $\text{QConstraint} \rightarrow \text{QConstraint} \rightarrow \text{QConstraint}$ 
         $\sim$        :  $\text{Type} \rightarrow \text{Type} \rightarrow \text{QConstraint}$ 

module OutsideIn (x : X) where
  open X(x)
  data CONSTRAINT : Set where
    QC          :  $\text{QConstraint} \rightarrow \text{CONSTRAINT}$ 
     $\wedge'$        :  $\text{CONSTRAINT} \rightarrow \text{CONSTRAINT} \rightarrow \text{CONSTRAINT}$ 
     $\exists$           :  $\text{CONSTRAINT} \rightarrow \text{CONSTRAINT}$ 
     $\exists \_ \_ \supset \_ \_$  :  $\mathbb{N} \rightarrow \text{QConstraint} \rightarrow \text{CONSTRAINT} \rightarrow \text{CONSTRAINT}$ 

```

Fig. 7: Naïve constraint representation

This encoding has a number of obvious problems. For example, all terms have an infinite number of free variables available, as the full type \mathbb{N} is used for variable names. This makes it impossible to determine instantly whether a term is closed or if a term contains free variables; one must instead analyse the term to extract this information. A common technique used to solve this problem when encoding binders in dependently typed languages is to index the type of terms by the number of available variables in the term. This technique is used often in generic programming literature, such as (Morris *et al.*, 2004), and was also shown by McBride to provide a convenient termination measure that can be used to phrase first-order unification as structural recursion (McBride, 2003). Reworking the above term definition to include such indexing, we get:

```

record X : Set1 where
  field Type      :  $\mathbb{N} \rightarrow \text{Set}$ 
        Var       :  $\forall \{n : \mathbb{N}\} \rightarrow \text{FIN } n \rightarrow \text{Type } n$ 
         $\rightarrow'$     :  $\forall \{n : \mathbb{N}\} \rightarrow \text{Type } n \rightarrow \text{Type } n \rightarrow \text{Type } n$ 
  field QConstraint :  $\mathbb{N} \rightarrow \text{Set}$ 
         $\epsilon$        :  $\forall \{n : \mathbb{N}\} \rightarrow \text{QConstraint } n$ 
         $\wedge$       :  $\forall \{n : \mathbb{N}\} \rightarrow \text{QConstraint } n \rightarrow \text{QConstraint } n \rightarrow \text{QConstraint } n$ 
         $\sim$        :  $\forall \{n : \mathbb{N}\} \rightarrow \text{Type } n \rightarrow \text{Type } n \rightarrow \text{QConstraint } n$ 

module OutsideIn (x : X) where
  open X(x)
  data CONSTRAINT (n :  $\mathbb{N}$ ) : Set where
    QC          :  $\text{QConstraint } n \rightarrow \text{CONSTRAINT } n$ 
     $\wedge'$        :  $\text{CONSTRAINT } n \rightarrow \text{CONSTRAINT } n \rightarrow \text{CONSTRAINT } n$ 
     $\exists$           :  $\text{CONSTRAINT } (\text{suc } n) \rightarrow \text{CONSTRAINT } n$ 
     $\exists \_ \_ \supset \_ \_$  :  $(m : \mathbb{N}) \rightarrow \text{QConstraint } n \rightarrow \text{CONSTRAINT } (n + m) \rightarrow \text{CONSTRAINT } n$ 

```

Fig. 8: FIN-named constraint representation

¹¹ Where \mathbb{N} is the type of the standard Peano naturals with zero : \mathbb{N} and suc : $\mathbb{N} \rightarrow \mathbb{N}$

This definition provides a type-level distinction between closed terms and terms that may contain some free variables, eliminating the problems with the earlier encoding. It enforces this by demanding that type variables be of type $\text{FIN } n$, where n is the number of available variables in the term. $\text{FIN } n$ is a type containing exactly n inhabitants — a finite set of natural numbers $[0, n)$ — defined as follows:

data $\text{FIN} : \mathbb{N} \rightarrow \text{Set}$ **where**
 $\text{zero} : \forall \{n : \mathbb{N}\} \rightarrow \text{FIN } (\text{suc } n)$
 $\text{suc} : \forall \{n : \mathbb{N}\} \rightarrow \text{FIN } n \rightarrow \text{FIN } (\text{suc } n)$

With this definition, the previously infinite number of available variables is now restricted to a finite number described by the type index. For example, the term $\exists. \exists. \text{QC } (\text{Var } \text{zero} \sim \text{Var } (\text{suc } \text{zero}))$ is closed and could therefore be of type $\text{CONSTRAINT } n$ for any $n : \mathbb{N}$ — that is, it could appear in a context with any number of available variables (including zero). The body of that constraint, $\text{QC } (\text{Var } \text{zero} \sim \text{Var } (\text{suc } \text{zero}))$, is by contrast typed most generally as $\text{CONSTRAINT } (\text{suc } (\text{suc } n))$ for any $n : \mathbb{N}$, which means that the term can only validly appear in a context with two or more available variables. Therefore, our form for the \exists quantifier *introduces* a new type variable by incrementing the index:

$$\exists : \text{CONSTRAINT } (\text{suc } n) \rightarrow \text{CONSTRAINT } n$$

The local assumption constraint, unlike the \exists quantifier, introduces more than one variable at a time:

$$\exists _ _ \supset _ : (m : \mathbb{N}) \rightarrow \text{QConstraint } n \rightarrow \text{CONSTRAINT } (n + m) \rightarrow \text{CONSTRAINT } n$$

Note that the antecedent QConstraint has only n available variables, not $n + m$, as it is impossible for the antecedent to mention any unification variables introduced in the succedent C -constraint.

3.3 Nested Data Types

While this indexing provides us a very nice way to handle de Bruijn indices for bound variables, we have not introduced an elegant way to handle type constructors, such as *Int* or *Maybe*, which exist in the top-level environment.

From the perspective of the constraint generation and solver infrastructure, type constructors are no different from type variables. One possible solution, therefore, is to simply assign indices to these top level type constructors. This makes type constructors difficult to distinguish from type variables, however, which becomes a serious problem when instantiating the X parameter — type constructors, which are treated as *rigid* and cannot be unified, must be distinguished from unification type variables, which can be substituted. In addition, assigning indices to top-level global definitions is aesthetically unpleasing.

A common alternative is to introduce a separate *Con* introduction form for **Type**, which refers to type constructors, as opposed to the *Var* form for type variables. This solution becomes unsatisfactory when we examine the solver infrastructure (see Fig. 5). Note that, when solving each implication constraint, *all* variables bound outside the implication constraint, *including* variables that are unifiable outside the implication, are treated as rigid, skolem variables. This means that a variable previously treated as unifiable could, in a different context, be treated as rigid like a constructor. Using a separate *Con* form therefore does not bring any advantage — we still have some subset of the available type variables being treated as skolem.

Our approach allows us to treat type constructors and type variables identically in the constraint generation and solver infrastructure of the $\text{OUTSIDEIN}(X)$ system itself, but retain the ability to separate skolem variables from unification variables when implementing the simplifier. In addition, type constructors can be represented by *any* type¹², and therefore do not need to be assigned indices, which makes working with the system a great deal more convenient.

¹² Provided such a type has decidable equality.

Our approach is to index terms not by the *size* of the set of available type variables, but by the *set itself*, as shown below:

```

record X : Set1 where
  field Type      : Set → Set
         Var       : ∀{n : Set} → n → Type n
         →'        : ∀{n : Set} → Type n → Type n → Type n
  field QConstraint : Set → Set
         ε         : ∀{n : Set} → QConstraint n
         ∧         : ∀{n : Set} → QConstraint n → QConstraint n → QConstraint n
         ~         : ∀{n : Set} → Type n → Type n → QConstraint n

module OutsideIn (x : X) where
  open X(x)
  data CONSTRAINT (n : Set) : Set where
    QC          : QConstraint n → CONSTRAINT n
    ∧'          : CONSTRAINT n → CONSTRAINT n → CONSTRAINT n
    ⊥           : CONSTRAINT (S n) → CONSTRAINT n
    ∃_._ ⊃ _    : (m : ℕ) → QConstraint n → CONSTRAINT (n ⊕ m) → CONSTRAINT n

```

Fig. 9: Constraint representation with nested data types

The secret to this representation lies in the \mathcal{S} data type, used when new type variables are made available by quantifiers. As an additional bound variable is now available, the type $\mathcal{S} \tau$ must be isomorphic to $\tau + 1$ ¹³, and is therefore implemented as follows:

```

data S (τ : Set) : Set where
  zero  : S τ
  suc   : τ → S τ

```

We also introduce another operation \oplus , for local assumptions, which is essentially repeated application of \mathcal{S} :

```

_ ⊕ _      : Set → ℕ → Set
x ⊕ zero   = x
x ⊕ suc n  = (S x) ⊕ n

```

Then, a simplifier can take as arguments constraints of type $\text{QConstraint } (x \oplus n)$, where x is the type for rigid variables and constructors, and n is the number of additional variables that can be unified. This obviates the need for an explicit set $\overline{\alpha_{\text{tch}}}$ to be passed to the solver — instead this information can be gleaned purely from the type of constraints passed in. This approach is advantageous because it allows us to handle this reinterpretation of variables in a unification or skolem context without changing the structure of the term.

3.3.1 Monads

Another elegant observation about this representation is that type terms can now form a categorical abstraction familiar to every Haskell programmer — a *monad*. This monadic structure is not an original discovery; it has been demonstrated by others in the past, doing similar work on term representation (Bird & Paterson, 1999; Bellegard & Hook, 1994). Originally from category theory, a *monad* is understood by functional programmers to be a type constructor m , a function $\text{unit} : \alpha \rightarrow m \alpha$ and a *Kleisli composition* operator $\circ_m : (\beta \rightarrow m \gamma) \rightarrow (\alpha \rightarrow m \beta) \rightarrow (\alpha \rightarrow m \gamma)$, such that the following laws hold:

```

Left Identity:   unit ◦m f      ≐ f
Right Identity: g ◦m unit     ≐ g
Associativity:  f ◦m (g ◦m h) ≐ (f ◦m g) ◦m h

```

¹³ i.e. Haskell's `Maybe` type.

A *category* is comprised of a class of *objects*, a class of *arrows* or *morphisms* between those objects, including an identity morphism for each object, and a morphism composition operation \circ which must be an associative and respect identity morphisms. From this definition, it is obvious that the above *monad laws* are just restatements of the *category laws* for a specific category, called the *Kleisli category* for the monad m .

We introduce informally the notion of a category **Agda**, consisting of Agda types as objects, Agda functions as morphisms, $(\lambda x \rightarrow x)$ as every identity morphism and function composition \circ as morphism composition.¹⁴ Then, the Kleisli category for **Type** is a *subcategory* of **Agda** with all the same objects, but only those functions with types of the form $\alpha \rightarrow \mathbf{Type} \beta$ as morphisms. Note that the type $\alpha \rightarrow \mathbf{Type} \beta$ is that of a substitution; Kleisli composition is substitution composition, where **Var** is the identity substitution. **Var** is therefore the *unit* operation for our monad and the identity morphism for our Kleisli category.

By requiring in our X parameter that **Type** be a monad, we are able to assume that substitution is well-behaved with respect to the structure of the type terms.

3.3.2 Functors

We can simply derive the familiar *bind* function from Kleisli composition and *unit*:

$$\begin{aligned} \mathit{bind} &: \forall\{\alpha \beta\} \rightarrow (\alpha \rightarrow m \beta) \rightarrow m \alpha \rightarrow m \beta \\ \mathit{bind} \ f \ a &= (f \circ_m (\lambda x \rightarrow a)) \ \mathit{it} \end{aligned}$$

(Where “it” is a constructor for a unit type)

When applied to types, *bind* is clearly application of a substitution to a term:

$$\mathit{bind} : \forall\{\alpha \beta\} \rightarrow (\alpha \rightarrow \mathbf{Type} \beta) \rightarrow \mathbf{Type} \alpha \rightarrow \mathbf{Type} \beta$$

Viewed categorically, *bind* could be viewed here as a mapping from morphisms in the Kleisli category of **Type** (i.e functions with types of the form $\alpha \rightarrow \mathbf{Type} \beta$) to morphisms in a new subcategory of **Agda** in the image of **Type**, which we call the **Type**-subcategory. This category has morphisms consisting of Agda functions, and objects consisting of types of the form **Type** τ for any τ . Note that the following two properties hold for *bind*:

$$\begin{aligned} \mathbf{Identity}: \quad & \mathit{bind} \ \mathit{unit} \quad \doteq \quad \mathit{id} \\ \mathbf{Composition}: \quad & \mathit{bind} \ f \circ \mathit{bind} \ g \quad \doteq \quad \mathit{bind} \ (f \circ_{\mathbf{Type}} g) \end{aligned}$$

If we have a mapping f from a category A to a category B , such that identity maps to identity, and composition maps to composition, then f is called a *functor* from A to B . As this is true of *bind*, we can say that *bind* is a functor from the Kleisli category of **Type** to the **Type**-subcategory.

We would like to be able to perform substitution on more than just **Type** terms. Specifically, it is also necessary to perform substitution on **QConstraints** and, ultimately, **CONSTRAINTS**.

To achieve this we require in the parameter X *another* functor, also from the Kleisli category of **Type**, but this time to the subcategory of **Agda** in the image of **QConstraint** — the **QConstraint**-subcategory:

$$Q\text{-subst} : \forall\{\alpha \beta\} \rightarrow (\alpha \rightarrow \mathbf{Type} \beta) \rightarrow (\mathbf{QConstraint} \alpha \rightarrow \mathbf{QConstraint} \beta)$$

The functor laws tell us that this, once again, respects substitution operations:

$$\begin{aligned} \mathbf{Identity}: \quad & Q\text{-subst} \ \mathit{unit} \quad \doteq \quad \mathit{id} \\ \mathbf{Composition}: \quad & Q\text{-subst} \ f \circ Q\text{-subst} \ g \quad \doteq \quad Q\text{-subst} \ (f \circ_{\mathbf{Type}} g) \end{aligned}$$

Given this functor, similar functors can be produced for **CONSTRAINT** and, indeed, any type that similarly contains **QCONSTRAINTS**.

Some difficulty arises, however, when defining the functor for terms which introduce variables: The form $\perp. C$ contains a subexpression of type **CONSTRAINT** ($\mathcal{S} \alpha$). Showing that \mathcal{S} is a monad¹⁵ produces the requisite

¹⁴ This is similar to the imaginary category **Hask** for Haskell, however unlike **Hask**, Agda types and functions *do* actually form a category

¹⁵ The familiar **Maybe** monad, no less

functors to transform our substitution $\alpha \rightarrow \text{Type } \beta$ to a substitution $\mathcal{S} \alpha \rightarrow \text{Type } (\mathcal{S} \beta)$, as required for this case. The form $\exists n. Q \supset C$ contains a subexpression of type `CONSTRAINT` $(\alpha \oplus n)$, which is a slightly more complicated beast to tame. Here, we must show not only that \mathcal{S} is a monad, but that $\mathcal{S} \circ \mathcal{S}$ is also, and $\mathcal{S} \circ \mathcal{S} \circ \mathcal{S}$, and so on. In other words, we have to show that for all $n, \lambda \alpha. \alpha \oplus n$ is a monad. Towards this end, we borrow the concept of a *monad transformer*, i.e. a monad homomorphism *lift* from any monad M to $M \circ \mathcal{S}$. By showing that *lift* is monad homomorphism for the \mathcal{S} -transformer, and that the resultant type $M \circ \mathcal{S}$ is a monad if M is a monad, it can be trivially shown that any number of \mathcal{S} 's form a monad, and thus that $\lambda \alpha. \alpha \oplus n$ is a monad as required.

Using the above, we have defined similar “substitution functors” for all types that are indexed by the set of type variables.

3.3.3 Renaming

Another common operation for `Types`, `QConstraints` and so on is *renaming*. Renaming can be implemented via functor composition in terms of substitution:

$$\begin{aligned} \tau\text{-rename} &: \forall\{\alpha \beta\} \rightarrow (\alpha \rightarrow \beta) \rightarrow (\text{Type } \alpha \rightarrow \text{Type } \beta) \\ \tau\text{-rename } f &= \text{bind } (\text{rename } f) \\ \\ Q\text{-rename} &: \forall\{\alpha \beta\} \rightarrow (\alpha \rightarrow \beta) \rightarrow (\text{QConstraint } \alpha \rightarrow \text{QConstraint } \beta) \\ Q\text{-rename } f &= Q\text{-subst } (\text{rename } f) \\ \\ &\dots \end{aligned}$$

Where *rename* is a functor from the base `Agda` category to the Kleisli category of `Type`:

$$\begin{aligned} \text{rename} &: \forall\{\alpha \beta\} \rightarrow (\alpha \rightarrow \beta) \rightarrow (\alpha \rightarrow \text{Type } \beta) \\ \text{rename } f &= \text{unit} \circ f \end{aligned}$$

With these definitions, the common de Bruijn index “up-shift” substitution which increments every index is trivially *rename suc*. Up-shifting by multiple variables at a time can be done simply by *rename unit_n* where *unit_n* is the *unit* morphism of the monad $\lambda \alpha. \alpha \oplus n$.

3.3.4 A Note on Equality

Propositional equality is defined simply in `Agda`¹⁶, as a data type that reifies definitional equality:

$$\begin{aligned} \mathbf{data} \quad \equiv & \{A : \text{Set}\} : A \rightarrow A \rightarrow \text{Set} \mathbf{where} \\ \text{refl} &: \forall\{x : A\} \rightarrow x \equiv x \end{aligned}$$

This allows us to introduce definitional equality constraints to the local context by pattern matching, in order to prove basic theorems like transitivity and symmetry of equality:

$$\begin{aligned} \text{trans} &: \forall\{A : \text{Set}\}\{x \ y \ z : A\} \rightarrow x \equiv y \rightarrow y \equiv z \rightarrow x \equiv z \\ \text{trans refl refl} &= \text{refl} \\ \\ \text{sym} &: \forall\{A : \text{Set}\}\{x \ y : A\} \rightarrow x \equiv y \rightarrow y \equiv x \\ \text{sym refl} &= \text{refl} \end{aligned}$$

The monad and functor laws shown above, however, are in terms of extensional equality, not this propositional equality. Unfortunately, this equality is not extensional — that is, the statement that two functions f and g are

¹⁶ This definition is somewhat simplified - in particular, universe polymorphism is removed.

propositionally equal, if, for all x , $f(x)$ is propositionally equal to $g(x)$ is not provable in Agda. Any value x can only be said to be propositionally equal to some other value y if x and y both normalise to the same result. If the definitions of two functions are (intensionally) different, they will not normalise to the same definition, even if they give the same result for all inputs.

The general approach for proving lemmas which require extensionality is to prove them within Altenkirch's setoid-based¹⁷ model (Altenkirch, 1999). As this becomes quite tedious in practice, we opt for the perhaps more inelegant approach of simply postulating functional extensionality:

postulate *extensionality* : $\{A B : \text{Set}\} \{f g : A \rightarrow B\} \rightarrow (\forall x \rightarrow f\ x \equiv g\ x) \rightarrow f \equiv g$

Agda's logic is consistent with this postulate, however we do lose the sometimes-valuable property of *canonicity* for equality proofs – that is, not all equality proofs normalise to a canonical closed term (i.e *refl*), as the postulate prevents normalisation. In practice, this means that programs which depend on this postulate will not give meaningful results (they crash the program much like Haskell's `error`), however the postulate can be freely used for proof work - that is, to show that a particular type is inhabited.

We do not need extensionality to model the `OUTSIDEIN(X)` algorithm itself. We require it only to prove monad laws and similar identities about our term representation. The loss of canonicity in exchange for simpler proofs is a trade we deem acceptable.

3.4 Separator and Prenexer

The definition of `SEPARATEDCONSTRAINT`, our representation of \mathcal{C} , is uninteresting save that it too is indexed by the set of available type variables, and is similarly equipped with functors for substitution.

It seems correct, then, that our separator should be a total function of the following type:

$sep : \forall \{tv : \text{Set}\} \rightarrow \text{CONSTRAINT } tv \rightarrow \text{SEPARATEDCONSTRAINT } tv$

However our *sep* function is not total, as it is. The `sep` function defined in Figure 5 is not defined for constraints of the form $\perp . C$. To solve this problem, we adjust our constraint representation to be indexed by a *strata*:

```

record X : Set1 where
  field Type      : Set → Set
         Var       : ∀{n : Set} → n → Type n
         →'        : ∀{n : Set} → Type n → Type n → Type n
  field QConstraint : Set → Set
         ε         : ∀{n : Set} → QConstraint n
         ∧         : ∀{n : Set} → QConstraint n → QConstraint n → QConstraint n
         ~         : ∀{n : Set} → Type n → Type n → QConstraint n
module OutsideIn (x : X) where
  open X(x)
  data STRATA : Set where
    Hi : STRATA
    Lo : STRATA
  data CONSTRAINT (n : Set) : STRATA → Set where
    QC      : ∀{σ : STRATA} → QConstraint n → CONSTRAINT n σ
    ∧'     : ∀{σ : STRATA} → CONSTRAINT n σ → CONSTRAINT n σ → CONSTRAINT n σ
    ⊥       : CONSTRAINT (S n) Hi → CONSTRAINT n Hi
    ∃_._ ⊃ _ : ∀{σ} → (m : ℕ) → QConstraint n → CONSTRAINT (n ⊕ m) σ
              → CONSTRAINT n σ

```

Fig. 10: Final Constraint Representation

¹⁷ A *setoid* being a dependent product (τ, \approx) where \approx is an equivalence relation on the type τ .

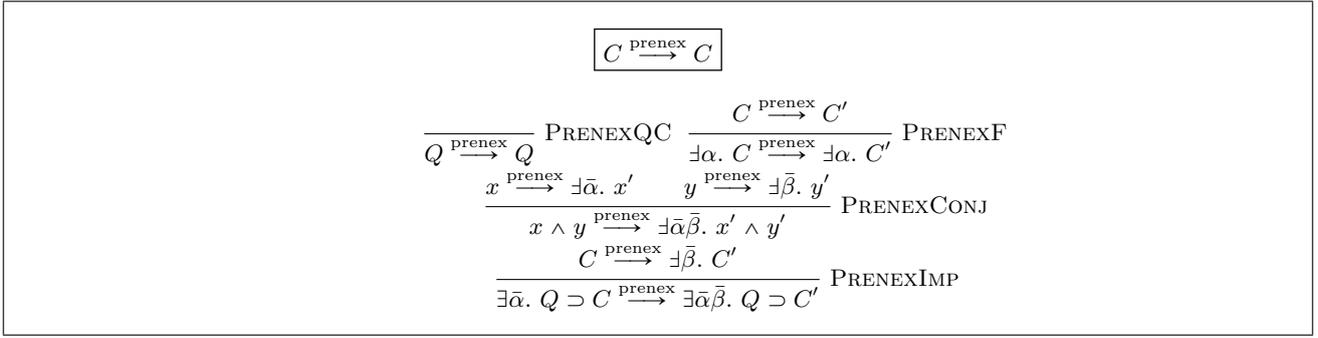


Fig. 11: A structurally recursive prenexer

Note that the type `CONSTRAINT n Hi` contains all forms of constraints, whereas the type `CONSTRAINT n Lo` contains all constraints except for constraints of the form $\exists. C$. This allows us to achieve a limited kind of subtyping via type indexing.

Now we can more accurately specify the domain of the `sep` function, so as to make it total:

$$\text{sep} : \forall \{tv : \text{Set}\} \rightarrow \text{CONSTRAINT } tv \text{ Lo} \rightarrow \text{SEPARATEDCONSTRAINT } tv$$

We can also give a more detailed account of the type of the prenexer:

$$\text{prenex} : \forall \{tv : \text{Set}\} \rightarrow \text{CONSTRAINT } tv \text{ Hi} \rightarrow \exists (\lambda n \rightarrow \text{CONSTRAINT } (tv \oplus n) \text{ Lo})$$

Implementing the prenexer as presented in Figure 4 is somewhat difficult, as it is presented as a rewrite system rather than as a neat structural recursion. We instead implement a structurally recursive prenexer as shown in Figure 11. The equivalence between these two formalisations can be proven via an easy induction.

3.5 Expressions

Our representation of `EXPRESSION` merits some attention as it is indexed by not one, but *two* types — one for available type variable names, and one for available variables on the *value* level:

data `EXPRESSION` (*ev* *tv* : `Set`) : ...

This approach brings a number of advantages. For example, type environments like Γ can be represented as a total function — it is impossible for an environment lookup to fail. This also ensures that we update the type environment whenever new value-level variables are available.

It is problematic, however, when dealing with pattern matching. Inside a pattern alternative, some number of new value variables are introduced into scope — this number should be reflected in the *ev* index. The exact number of variables, however, is dependent on the *type* (or, at least, the arity) of the data constructor in the pattern, which is not information we have in our possession when constructing syntax trees. Furthermore, it is already clear that data constructors require separate treatment from regular variables, as patterns contain only names of data constructors, not arbitrary variables.

This calls for a more sophisticated representation of names in expressions. Instead of simply using *ev* to refer to value-level variables, we will use `NAME ev`, defined as follows:

data `NAME` *TYPE* : `Set` **where**
`Binding` : `NAME` *TYPE*
`Datacon` : $\mathbb{N} \rightarrow \text{NAME } \textit{TYPE}$

data `NAME` (*n* : `Set`) : `NAME` *TYPE* \rightarrow `Set` **where**
`N` : *n* \rightarrow `NAME` *n* `Binding`
`DC` : $\forall \{x\} \rightarrow \text{dc } x \rightarrow \text{NAME } n \text{ (Datacon } x)$

(Where `dc` is a type for datacon names, indexed by their arity, provided in the parameter *X*)

A NAME can signify, therefore, either a straightforward binding *or* a data constructor of some arity. As the arity is presented in the type NAME, we can easily produce the desired type for a pattern matching alternative:

data ALTERNATIVE ($ev\ tv : \text{Set}$) : SHAPE \rightarrow Set **where**
 $\xrightarrow{\text{alt}}$: $\forall\{n : \mathbb{N}\} \rightarrow \text{NAME } ev\ (Datacon\ n) \rightarrow \text{EXPRESSION } (ev \oplus n)\ tv \rightarrow$
 ALTERNATIVE $ev\ tv$

This technique works greatly to our advantage as there is another instance where data constructors require special treatment; this time in constraint generation. Observe the two rules in constraint generation where the environment is consulted; the rules VARCON and CASE:

$$\frac{(v : \forall \bar{a}. Q_1 \Rightarrow \tau_1) \in \Gamma}{\Delta; \Gamma \vdash v : \tau \rightsquigarrow \exists \bar{\alpha}. [\bar{a} \mapsto \bar{\alpha}] Q_1 \wedge' (\tau \sim [\bar{a} \mapsto \bar{\alpha}] \tau_1)} \text{VARCON}$$

$$\frac{\begin{array}{l} \alpha, \beta, \bar{\gamma}, \bar{\delta}, \Delta; \Gamma \vdash e : \alpha \rightsquigarrow C \\ (K_i : \forall \bar{a} \bar{b}. Q_i \Rightarrow \bar{\tau}_i \rightarrow \mathbb{T} \bar{a}) \in \Gamma \quad \alpha, \beta, \bar{\gamma}, \bar{\delta}, \bar{\rho}, \Delta; \Gamma, (\bar{x}_i : [\bar{b} \mapsto \bar{\rho}] [\bar{a} \mapsto \bar{\gamma}] \tau_i) \vdash e_i : \delta_i \rightsquigarrow C_i \\ C'_i = \begin{cases} C_i \wedge' \delta_i \sim \beta & \text{if } \bar{b}_i = \epsilon \text{ and } Q_i = \epsilon \\ \exists \bar{\rho}. \exists \epsilon. ([\bar{b} \mapsto \bar{\rho}] [\bar{a} \mapsto \bar{\gamma}] Q_i) \supset C_i \wedge' \delta_i \sim \beta & \text{otherwise} \end{cases} \end{array}}{\Delta; \Gamma \vdash \text{case } e \text{ of } \{K_i \bar{x}_i \rightarrow e_i\} : \tau \rightsquigarrow \exists \alpha. \exists \beta. \exists \bar{\gamma}. \exists \bar{\delta}. C \wedge' (\mathbb{T} \bar{\gamma} \sim \alpha) \wedge' (\bigwedge C'_i) \wedge' (\tau \sim \beta)} \text{CASE}$$

Note that in each instance, the form of the type schema retrieved from the environment is different! One form for data constructors ($\forall \bar{a} \bar{b}. Q \Rightarrow \bar{\tau} \rightarrow \mathbb{T} \bar{a}$), and another for bindings ($\forall \bar{a}. Q \Rightarrow \tau$). As Type is part of the X parameter, we do not have the luxury of simply pattern-matching on the structure of the type schema to determine if it is of the correct form. Instead we adjust our encoding of type schema, in order to force the user to provide correctly-structured types in the environment for data constructors. Specifically, we index our representation of type schema by a NAME TYPE, providing a general type schema form for types of regular Bindings but more structured forms for types of data constructors:

data TYPE SCHEMA ($tv : \text{Set}$) : NAME TYPE \rightarrow Set **where**
 $\forall' _ . _ \Rightarrow _$: $(n : \mathbb{N}) \rightarrow \text{QConstraint } (tv \oplus n) \rightarrow \text{Type } (tv \oplus n) \rightarrow \text{TYPE SCHEMA } tv \text{ Regular}$
 $\forall' _ . _ \rightarrow _$: $(a : \mathbb{N}) \{r : \mathbb{N}\} \rightarrow \text{VEC } (\text{Type } (tv \oplus a))\ r \rightarrow tv \rightarrow \text{TYPE SCHEMA } tv\ (Datacon\ r)$
 $\forall' _ . _ . _ \Rightarrow _ \rightarrow _$: $(a\ b : \mathbb{N}) \{r : \mathbb{N}\} \rightarrow \text{QConstraint } (tv \oplus a \oplus b) \rightarrow$
 $\text{VEC } (\text{Type } (tv \oplus a \oplus b))\ r \rightarrow tv \rightarrow \text{TYPE SCHEMA } tv\ (Datacon\ r)$

Here, we use a length-indexed vector VEC to ensure that the arity of the function type for data constructors matches the arity mentioned in their NAME.

This allows us to define environments as a total mapping of any NAME of type n to TYPE SCHEMA for the same name type n .

Environment : Set \rightarrow Set \rightarrow Set
Environment $ev\ tv = \forall\{n : \text{NAME TYPE}\} \rightarrow \text{NAME } ev\ n \rightarrow \text{TYPE SCHEMA } tv\ n$

3.5.1 Shape Indexing

One of the issues that results from indexing expressions by the set of available type variables is that constraint generation is no longer quite as trivially structurally recursive. For example, take the constraint generation rule for abstraction:

$$\frac{\alpha, \beta, \Delta; \Gamma, (x : \alpha) \vdash e : \beta \rightsquigarrow C}{\Delta; \Gamma \vdash \lambda x. e : \tau \rightsquigarrow \exists \alpha. \exists \beta. C \wedge' (\tau \sim (\alpha \rightarrow \beta))} \text{ABS}$$

In this rule, two new type variables are introduced into the set of available variables Δ . This corresponds in our representation to an index \mathcal{S} ($\mathcal{S}\ tv$), however the expression e 's index is merely tv . The obvious way to

```

data NAMEType : Set where
  Binding   : NAMEType
  Datacon   : ℕ → NAMEType

data NAME (n : Set) : NAMEType → Set where
  N         : n → NAME n Binding
  DC        : ∀ {x} → dc x → NAME n (Datacon x)

data ALTERNATIVE (ev tv : Set) : SHAPE → Set where
   $\xrightarrow{\text{alt}}$  : ∀ {n : ℕ} {σ : SHAPE} → NAME ev (Datacon n) → EXPRESSION (ev ⊕ n) tv σ →
    ALTERNATIVE ev tv (Unary σ)

data ALTERNATIVES (ev tv : Set) : SHAPE → Set where
  esac      : ALTERNATIVES ev tv Nullary
  ||        : ∀ {σ1 σ2 : SHAPE} → ALTERNATIVE ev tv σ1 → ALTERNATIVES ev tv σ2 →
    ALTERNATIVES ev tv (Binary σ1 σ2)

data EXPRESSION (ev tv : Set) : SHAPE → Set where
  EVar      : ∀ {x : NAMEType} → NAME ev x → EXPRESSION ev tv Nullary
  λ'        : ∀ {σ : SHAPE} → EXPRESSION (S ev) tv σ → EXPRESSION ev tv (Unary σ)
  app       : ∀ {σ1 σ2 : SHAPE} → EXPRESSION ev tv σ1 → EXPRESSION ev tv σ2 →
    EXPRESSION ev tv (Binary σ1 σ2)
  let       : ∀ {σ1 σ2 : SHAPE} → EXPRESSION ev tv σ1 → EXPRESSION (S ev) tv σ2 →
    EXPRESSION ev tv (Binary σ1 σ2)
  leta    : ∀ {σ1 σ2 : SHAPE} → EXPRESSION ev tv σ1 → Type tv → EXPRESSION (S ev) tv σ2 →
    EXPRESSION ev tv (Binary σ1 σ2)
  letga   : ∀ {σ1 σ2 : SHAPE} → (n : ℕ) → EXPRESSION ev (tv ⊕ n) σ1 →
    QConstraint (tv ⊕ n) → Type (tv ⊕ n) → EXPRESSION (S ev) tv σ2 →
    EXPRESSION ev tv (Binary σ1 σ2)
  case      : ∀ {σ1 σ2 : SHAPE} → EXPRESSION ev tv σ1 → ALTERNATIVES ev tv σ2 →
    EXPRESSION ev tv (Binary σ1 σ2)

```

Fig. 12: The complete EXPRESSION representation

remedy this is to simply “up-shift” the de Bruijn indices in e twice using the appropriate substitution functor, however this causes Agda’s termination checker to complain — the recursive invocation is no longer on the obviously smaller e ; instead it is on the not-so-obviously smaller $e\text{-subst}$ ($\text{rename } (\text{succ} \circ \text{succ})$) e . In order to show termination, we must show that substitution on types (which change the structure of *type* terms) do not change the recursion pattern used by the constraint generation (the structure of *expression* terms). We achieve this using a simple technique: we add a *third* index to the EXPRESSION type, which represents the structure or the *shape* of the expression tree (see Fig. 12). The SHAPE type itself is a straightforward tree structure which is sufficient to encode the recursion pattern of constraint generation:

```

data SHAPE : Set where
  Nullary   : SHAPE
  Unary     : SHAPE → SHAPE
  Binary    : SHAPE → SHAPE → SHAPE

```

With this addition, the type of the substitution functor for expressions now becomes:

$$e\text{-subst} : \forall \{\sigma : \text{SHAPE}\} \{ev \alpha \beta : \text{Set}\} \rightarrow (\alpha \rightarrow \text{Type } \beta) \rightarrow \text{EXPRESSION } ev \alpha \sigma \rightarrow \text{EXPRESSION } ev \beta \sigma$$

The fact that the shape of the expression is not changed by type substitution is now made clear in the type of the type substitution operation. In constraint generation, recursion once again becomes structural, not on the expressions themselves, but on their shape parameters.

3.6 Constraint Generator and Solver

Now that we have established our choice of representation for terms and types, we must still choose a representation for the *computations* themselves, specifically constraint generation, and solving.

Broadly speaking, there are two main ways to encode this computation. The first is to directly encode the computation as an Agda function (using constraint generation as an example):

```

genConstraint : {ev tv : Set} {σ : SHAPE}
  → Environment ev tv
  → EXPRESSION ev tv σ
  → Type tv
  → CONSTRAINT tv
...

```

This is certainly the most straightforward encoding of the algorithm, and was the method we had used during development of our formalisation. Ultimately, we chose not to use this simple representation, however, opting instead for a *propositional* approach, defining a propositional data type for the constraint generation judgement:

```

syntax CONSTRAINTGEN Γ τ e C = Γ ⊢ e : τ ∼ C
data CONSTRAINTGEN {ev tv : Set} (Γ : Environment ev tv) (τ : Type tv) :
  {σ : SHAPE} → EXPRESSION ev tv σ → CONSTRAINT tv Hi → Set where
...

```

This type contains constructors corresponding to each rule in the constraint generation (see Fig. 3). Then, to provide an algorithmic interpretation of these rules, and therefore provide a way to actually *run* our constraint generator and infer types, we write a function that generates a proof of this proposition given the requisite inputs¹⁸:

```

genConstraint' : {ev tv : Set} {σ : SHAPE} (Γ : Environment ev tv) (e : EXPRESSION ev tv σ) (τ : Type tv)
  → ∃ (λ C → Γ ⊢ e : τ ∼ C)

```

This approach brings a number of advantages. For one, it brings a substantial performance improvement — the constraint generator takes one tenth the time and one fifth the memory to type check. It also gives us an Agda proposition exactly the same in form to the rules presented in the previous chapter. Not only a nice aesthetic consequence, this also liberates us from certain petty annoyances that arise from encoding the algorithm directly. For example, the solver is partial — not all constraints are satisfiable. In a direct functional encoding, this would necessitate **Maybe** types or similar mechanisms to deal with partiality. The propositional encoding of the solver, however, is only ever constructible in the happy case; the partiality is relegated to the algorithmic interpretation of the proposition.

If we desire to prove things about these operations (for example, inference soundness and principality), then the propositional encoding of these algorithms is much more preferable. Instead of evaluating expressions and pattern matching on their results, we can more directly match on the various rules used to justify the proposition. This makes proofs done on paper about these systems much more straightforward to transfer to our formalisation.

Another advantage of the propositional encoding comes from the solver. Observe the rule for the solver infrastructure:

$$\frac{\mathcal{Q}; Q_g; \bar{\alpha} \vdash^{\text{simp}} Q \rightsquigarrow Q_r; \theta \quad \forall ((\exists^I \bar{\beta}_i. Q_i \supset C_i) \in I). \mathcal{Q}; Q_g \wedge Q_r \wedge Q_i; \bar{\beta}_i \vdash^{\text{solv}} C_i \rightsquigarrow \epsilon; \theta_i}{\mathcal{Q}; Q_g; \bar{\alpha} \vdash^{\text{solv}} Q \cdot I \rightsquigarrow Q_r; \theta}$$

In particular, note that each implication constraint body must resolve to ϵ . In a propositional encoding, this is a non-issue, but an algorithmic encoding must check the return value of the solver function given the implication constraint, to *ensure* that the result is ϵ . As **QConstraints** are part of the parameter and therefore abstract, we cannot pattern match to make these assurances, and must therefore demand a proof of the decidability of equality to ϵ in the X constraint:

```

field dec-ε : (x : QConstraint) → DEC (x ≡ ε)

```

¹⁸ Interestingly, as this function must be total and terminating, this can be viewed as a proof of the decidability of the **CONSTRAINTGEN** proposition

Where DEC is defined (in the standard library) as follows:

```
data DEC ( $P$  : Set) : Set where  
  yes  :  $P \rightarrow$  DEC  $P$   
  no   :  $(P \rightarrow \perp) \rightarrow$  DEC  $P$ 
```

By adopting a propositional encoding for the solver, we do not eliminate this issue, but at least we are able to confine it merely to the algorithmic interpretation of the proposition rather than have such nonsense pollute the rules themselves.

4 Simple Instantiation

This chapter describes a simple instantiation of the X parameter of the `OUTSIDEIN(X)` system. We have formalised this instantiation in Agda and used it to infer types for basic programs.

By developing this simple instantiation, we not only provide a way to actually *use* the type checker we have formalised, but we also ensure that it is in fact *possible* to instantiate. Otherwise, it would be possible that the requirements in the X parameter were too onerous to actually be instantiable¹⁹.

Our instantiation of `QConstraint` is the smallest that meets the requirements of the X parameter. It consists only of ϵ , constraint conjunction and type equality constraints:

```
data QCONSTRAINT (x : Set) : Set where
  ~  :  TYPE x → TYPE x → QCONSTRAINT x
  ^  :  QCONSTRAINT x → QCONSTRAINT x → QCONSTRAINT x
   $\epsilon$  :  QCONSTRAINT x
```

`TYPE` consists of type variables (or constructors), type application and function types — the typical assortment of types from System F.

```
data TYPE (x : Set) : Set where
  Var  :  n → TYPE n
  →'   :  TYPE n → TYPE n → TYPE n
  app  :  TYPE n → TYPE n → TYPE n
```

`TYPE` is also a monad, as required by the X parameter, but this proof is entirely uneventful and uninteresting, as is the proof of the functor laws for the Q -*subst* functor. Indeed, the only truly interesting component of the simple instantiation is the simplifier.

The simplest possible simplifier is simply the no-op:

$$\mathcal{Q} ; Q_g ; \overline{\alpha_{\text{tch}}} \mapsto Q_w \rightsquigarrow Q_w ; \text{unit}_{\text{Type}}$$

While this simplifier meets the simplifier soundness and principality criteria in Figure 2, it is not a particularly interesting simplifier to work with. Instead, we define a simplifier that is actually capable of solving constraints.

Broadly speaking, the simplifier is divided into two main components:

1. **Unification**, for solving equality constraints.
2. **Solving**, which uses the results from unification to solve constraint terms.

4.1 Unification

Unification is the process of producing a substitution θ given two terms τ_1 and τ_2 such that $\theta\tau_1 \equiv \theta\tau_2$. Unification for first-order terms is decidable via Robinson’s algorithm (Robinson, 1965), but this has a nontrivial termination argument, based on the fact that each variable substituted reduces the number of available variables in the term. This gives a termination measure, but it is not one that is immediately visible as a structural recursion. This is problematic in Agda, which mandates that all functions be structurally recursive.

McBride more recently demonstrated a structurally recursive presentation of first-order unification in a dependently typed setting, indexing terms by the number of available unification variables (McBride, 2003). For the purposes of our unification, we shamelessly reuse his presentation with little modification.

One slight difference is that our `TYPE` terms do not come indexed by the number of available variables they contain. Rather, they come in the form `TYPE ($\tau \oplus n$)`, with skolem variables and constructor names represented by the type τ , and n unification variables available. We define our unification algorithm in terms of a special

¹⁹ As a most extreme example: if we had, by some accident, required a proof of \perp in our X parameter, this would not only be impossible to instantiate, it would also render any proof about the system entirely useless!

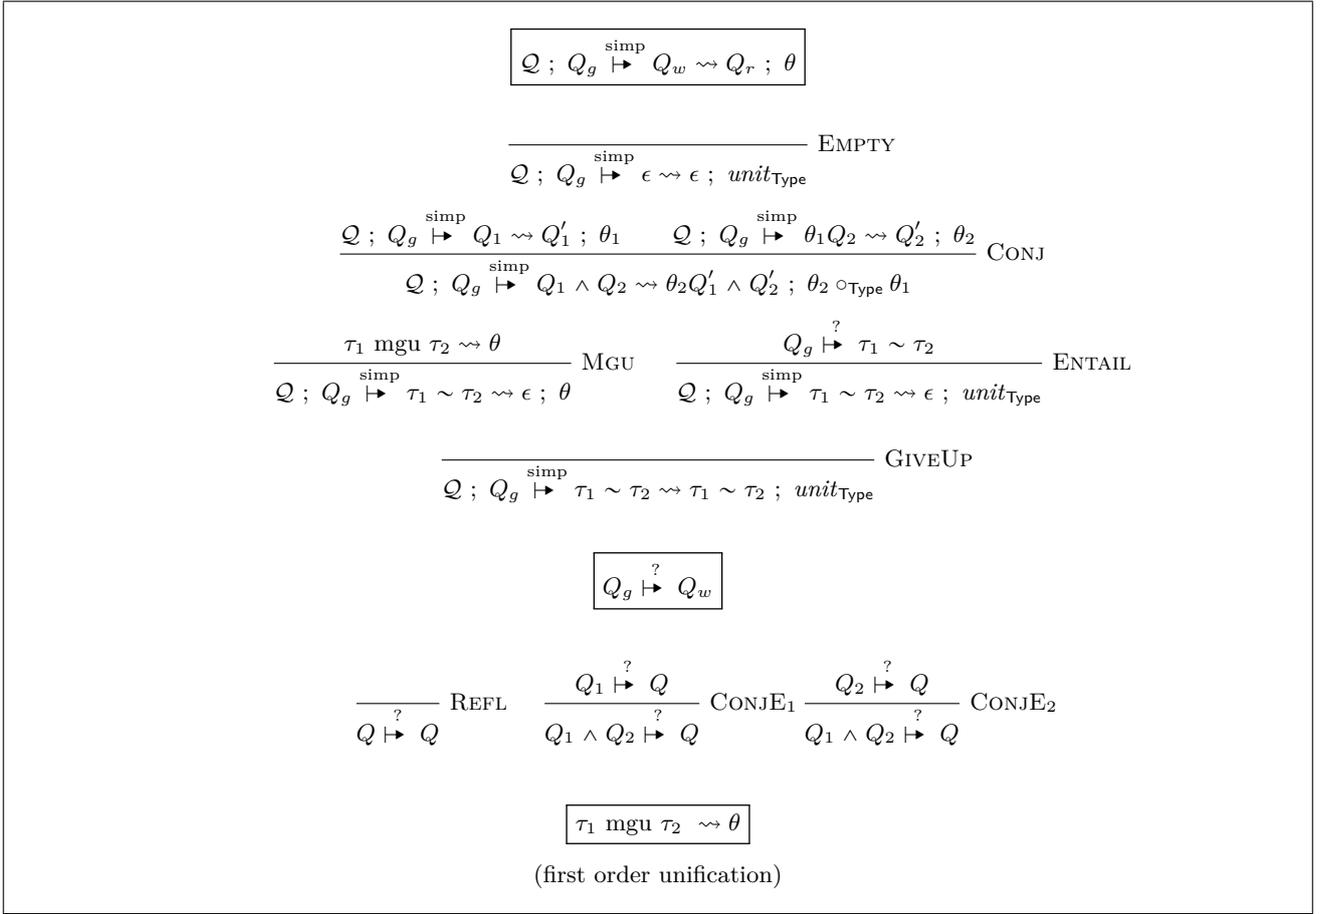


Fig. 13: Simple instantiation of the simplifier

type used for names, which explicitly separates skolem from unification variables — after all, they are treated very differently when performing unification:

```

data SNAME (sk : Set) (un : ℕ) : Set where
  unification : FIN un → SNAME sk un
  rigid       : sk → SNAME sk un

```

Then, we define the straightforward isomorphism between the type used for names in the simplifier, and this special SNAME type used in the unification:

$$\begin{aligned}
iso_1 &: \forall \{m : \mathbb{N}\} \{t : \text{Set}\} \rightarrow t \oplus m \rightarrow \text{SNAME } t \ m \\
iso_2 &: \forall \{m : \mathbb{N}\} \{t : \text{Set}\} \rightarrow \text{SNAME } t \ m \rightarrow t \oplus m
\end{aligned}$$

This isomorphism can be applied to types and constraints using the standard *rename* functors, which allows us to use a convenient representation for unification without affecting any other component of the instantiation or of the system.

As this representation trivially ensures that the domain of any substitution is restricted only to the unification variables $\overline{\alpha_{\text{tch}}}$, we omit mention of $\overline{\alpha_{\text{tch}}}$ in Figure 13.

4.2 Solving

The overall simplifier has to deal with the three possible constraint forms (see Fig. 13). Empty ϵ constraints resolve trivially to the identity substitution, conjunctions resolve to a composition of the substitutions resulting from each of the two conjuncts, and type equality resolves to the most general unifier of the two types. There are two other rules for type equality: ENTAIL, which searches the given context for an identical constraint to

the wanted equality; and GIVEUP, which simply gives up on solving the constraint entirely, returning it as a residual constraint. The presence of these rules makes the system not syntax-directed, and therefore it does not directly correspond to a deterministic algorithm. We resolve this nondeterminism by introducing an ordering on the rules: MGU is attempted first, then ENTAIL, then GIVEUP.

Another problem presents itself when one attempts to encode this simplifier in Agda. In the rule CONJ, the first substitution θ is applied to the second conjunct Q_2 , and recursion occurs on the substituted constraint θQ_2 . This recursion is *non-structural* — the dreaded termination checker has once again reared its head.

In order to convince Agda that applying a type substitution to a constraint does not affect the shape of the constraint, we employ a similar tactic to the technique used for representing EXPRESSIONS previously. We shall define a new form of constraint, indexed by its *shape*:

```
data SCONSTRAINT (x : Set) : SHAPE → Set where
  ~  :  TYPE x → TYPE x → SCONSTRAINT x Nullary
  ^  :  ∀{s1 s2} → SCONSTRAINT x r1 → SCONSTRAINT x r2 → SCONSTRAINT x (Binary s1 s2)
  ε  :  SCONSTRAINT x Nullary
```

With this indexed data type, substitution now has the following type:

$$SC\text{-subst} : \forall\{\alpha \beta\}\{s\} \rightarrow (\alpha \rightarrow \text{TYPE } \beta) \rightarrow SConstraint \alpha s \rightarrow SConstraint \beta s$$

As with EXPRESSION, we can now glean from the type that the *shape* of the constraint term does not change with substitution. Therefore, structural recursion can occur on the SHAPE index of the SCONSTRAINT term, rather than on the term itself. As SCONSTRAINT and QCONSTRAINT are structurally identical, it is trivial to convert between them.

4.3 Results

We have used this simple instantiation to successfully infer types, including polymorphic types and types involving GADTs. In particular, recall the example used earlier to illustrate the necessity of removing **let**-generalisation:

```
data IntOrBool :: * → * where
  IsInt  :: (τ ~ Int) ⇒ IntOrBool τ
  IsBool :: (τ ~ Bool) ⇒ IntOrBool τ

f :: IntOrBool α → α → Bool
f x y = let g z = not y in
  case x of
    IsInt  → True
    IsBool → g ()
```

This expression rightly fails to type-check in our simple instantiation, however if we add a type signature to g :

```
f :: IntOrBool α → α → Bool
f x y = let g : ∀β. (α ~ Bool) → β → Bool
  g z = not y in
  case x of
    IsInt  → True
    IsBool → g ()
```

The program now type-checks in our simple instantiation of OUTSIDEIN(X), as desired.

5 Conclusions and Future Work

Our formalisation comes to approximately three thousand lines of Agda definitions and proofs, excluding comments, which makes it one of the largest single formalisations in Agda to date, approximately one fifth the size of the entire Agda standard library. It takes approximately 46 seconds to type check, using the latest development version of Agda 2.3.2 on a mid-2011 MacBook Pro. Individual modules take no longer than five seconds to check.

During our development, constraint generation and other phases of the algorithm were written predominantly in a direct, functional style, rather than the final propositional encoding which we present here. Constraint generation, written this way, took approximately five minutes to type check, and required approximately six gigabytes of RAM — just for the constraint generation module alone. Exactly what causes this is still unclear, however we have raised the issue with the Agda developers.

While in development, we attempted to make extensive use of several experimental Agda features, such as instance arguments (Devriese & Piessens, 2011), however this frequently resulted in strangely unsolved metavariables, confusing error messages, or worse, internal errors from Agda’s type checker.

One of our goals in this endeavour is to show Agda’s readiness for type systems work. Despite all of the above problems, we believe we have succeeded in this respect. Agda has, for the most part, provided a very expressive language for us to encode our formalisation. Certain tricks, such as our encoding of expression names or our indexing of type terms, would be highly uncommon in proof assistants such as Coq and impossible in theorem provers such as Isabelle. On the other hand, Agda requires that all termination arguments be phrased in terms of structural recursion, which did result in some minor contortions to our formalisation that would not be necessary if a separate termination proof could be provided. Conor McBride has made the argument that such contortions are merely exposing pre-existing structure of our algorithms that was previously hidden, and therefore a structurally recursive presentation is clearer than a presentation with a separate termination proof (McBride, 2003). We are not certain whether we agree with this argument, but offer it as justification for the abovementioned contortions.

Agda is most definitely a programming language *first* and a proof assistant *second*. For this reason, it could be argued that we have been playing to Agda’s strengths — having encoded the structures, definitions and algorithms of the system, we have formalised the part of the system to which Agda is best suited. It remains to be seen whether Agda is suitable for serious proof work (see section 5.1 for a discussion of a possible soundness proof for our formalisation).

Another goal of our research was to provide a base upon which other type system formalisations can be developed, and to create a testbed for experimenting with $\text{OUTSIDEIN}(X)$ on solid formal ground. By including a simple instantiation of the system, we have demonstrated the capability of our formalisation to represent real type systems and infer types for real programs. An obvious direction for future work is to instantiate X for all of Haskell, a prospect we discuss in section 5.2, as well as expanding the simple instantiation to include proofs of the simplifier soundness and principality conditions.

5.1 Proofs

We have completed a preliminary investigation into a proof of *soundness* for the inference system we have presented. *Soundness* here refers the notion that if a particular type τ is inferred for an expression e , then there exists a *proof* according to some natural, permissive typing rules that e is well typed by τ . Also of interest is a *principality* proof, which suggests that any inferred type of an expression e is the principal type for e . The original $\text{OUTSIDEIN}(X)$ paper presents typing rules as well as a proof of soundness and principality for their system, assuming simplifier soundness and guess-freedom conditions. A number of obstacles make converting those proofs to our formalisation less than straightforward:

1. The proofs presented in the $\text{OUTSIDEIN}(X)$ paper are not very detailed, mostly due to space constraints, and as a result some steps in the proof are not very clear. For example, both proofs mention induction, however the exact structure of the induction is not clarified; and in both proofs, several cases are dismissed as similar to a previous case, without clarification as to exactly what is different or how they are similar.

2. Some properties of entailment upon which the soundness proof relies are not even formalised in the entailment relation presented in the paper. For example, it is not required that ϵ is entailed by any context, nor is conjunction elimination included in the original presentation of entailment. Nevertheless, these properties are relied upon by the soundness proof.
3. Our formalisation introduces a few new layers of indirection between constraint generation and constraint solving (such as the prenexer, separator and so on). This added indirection means that the constraint produced by the constraint generator and the constraint given to the solver are not the same, which makes reasoning more difficult.
4. The original proof takes certain properties of conjunctions as automatic and invisible; specifically associativity, commutativity, and identity with ϵ . This means that goals might be perfectly provable but not fitting to the exact structure of the constraints required, which can make proving much more difficult.

The first two points here mostly concern problems with the original `OUTSIDEIN(X)` proof. These problems are not insurmountable, and indeed they are to be expected — all proofs become sloppy when viewed through the harsh, clear lens of a proof assistant. The final two points, however, reveal some weaknesses in Agda’s approach to theorem proving. It is common to find oneself *encumbered by structure* in languages such as Agda, where proofs of even simple properties can be complicated by the wrong choice of representation. In Isabelle, for example, correspondences across layers of indirection and simple lemmas about entailment (such as commutativity of conjunction) can be added to an automatic simplifier that would make these problems trivially soluble. Agda, on the other hand, places on us the onerous burden of manually rewriting terms to be of the correct structure, and manually transporting properties across isomorphisms.

While these obstacles are irritating, they are not impossible to overcome. Therefore, our subsequent research effort shall be devoted to developing a soundness proof for our formalisation.

5.2 Instantiating X for Haskell

Another obvious progression from this point is to investigate an instantiation of the X parameter with a sophisticated type system that supports the many extensions available in Haskell, such as type classes (and their corresponding top-level axiom schemes), type families, and so on. A surface evaluation reveals that this would likely be quite difficult, chiefly because the simplifier for this Haskell-like type system does not have a straightforward termination measure — Indeed, none is presented in the original `OUTSIDEIN(X)` paper, which presents this simplifier as a sequence of rewrite rules. The GHC implementation of the simplifier also has no obvious termination argument. While a proof of termination could no doubt be found (the system does, after all, appear to terminate!), it may well be difficult to phrase such a sophisticated simplifier as a structural recursion.

One method towards circumventing this problem (which does not involve determining a complicated termination argument for the rewrite system) is to simply pass an arbitrary large natural number to the rewrite system which reduces this number by one on each rule application, failing if the number reaches zero. A proof of termination would then be rephrased as a proof that there exists a number for any term which, when given to this system along with the term, will result in an expression in normal form (i.e. not fail). By rephrasing it this way, Agda will view the rewrite system as structurally recursive in the natural number, and a separate proof of termination can be provided²⁰.

5.3 Summary of Contributions

- A more rigorous and clear formulation of the `OUTSIDEIN(X)` inference system.
- An Agda formalisation of this system, instantiable for real type systems, that can be used for experimentation and development of new type system extensions on solid formal ground.
- A simple example instantiation of the above Agda formalisation, which can be used to infer types for basic programs, and can also serve as a basis for more sophisticated instantiations.
- A demonstration of a variety of representational techniques for names and terms, suitable for similar work in both dependently-typed languages like Agda, and almost-dependently-typed languages like Haskell.

²⁰ Or simply assumed, if no proof is found.

Through our work on `OUTSIDEIN(X)`, we have shown that Agda gives us through its type system a very useful and powerful language to express structure and maintain invariants about our data. The addition of a soundness and principality proof, an obvious direction for future work, would make our formalisation highly compelling for those seeking to extend or experiment with Haskell’s type inference system. Even without such a proof, our formalisation already serves as a platform for experimentation, as it is the only implementation of `OUTSIDEIN(X)` that is actually abstracted over its X parameter. Ultimately, this project has brought us one small step closer to a rigorously formalised type system for GHC Haskell.

Acknowledgements. The author would like to thank the members of the Agda community, specifically Arseniy Alekseyev, Conor McBride, Daniel Peebles, Andrea Vezzosi and others, all of whom were of great assistance when learning Agda and developing this formalisation. Edward Kmett and Ashley Yakely also offered their advice and assistance during the development of the category theory components of the formalisation. Thanks are also due to Manuel Chakravarty, who supervised this project, and the entire PLS group at UNSW.

On a personal note, I would like to thank my partner Tran for her support and understanding in this busy year.

References

- 1 Thorsten Altenkirch. (1999). Extensional Equality in Intensional Type Theory. *Pages 412–412 of: Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science*. Washington, DC, USA: IEEE Computer Society.
- 2 Françoise Bellegard, and James Hook. (1994). Substitution: a Formal Methods Case Study Using Monads and Transformations. *Science of computer programming*, **23**(2-3), 287–311.
- 3 Richard S. Bird, and Ross Paterson. (1999). de Bruijn Notation as a Nested Datatype. *Journal of functional programming*, **9**(1), 77–91.
- 4 Nils Danielsson, and Ulf Norell. (2011). Parsing Mixfix Operators. *Pages 80–99 of: Implementation and Application of Functional Languages*. Springer Berlin / Heidelberg.
- 5 Nicolaas Govert de Bruijn. (1972). Lambda Calculus Notation with Nameless Dummies: a Tool for Automatic Formula Manipulation with Application to the Church-Rosser Theorem. *Indagationes mathematicae (elsevier)*, **34**, 381–392.
- 6 Dominique Devriese, and Frank Piessens. (2011). On the bright side of type classes: instance arguments in Agda. *Pages 143–155 of: ICFP ’11: Proceedings of the 16th ACM SIGPLAN international conference on Functional programming*. New York, New York, USA: ACM.
- 7 Catherine Dubois. (2000). Proving ML Type Soundness Within Coq. *Pages 126–144 of: TPHOLs ’00: Proceedings of "Theorem Proving with Higher Order Logics" 2000*. Springer.
- 8 Catherine Dubois, and Valérie Ménessier-Morain. (1999). Certification of a type inference tool for ML: Damas-Milner within Coq. *Journal of automated reasoning*, **23**, 3–4.
- 9 W.A. Howard. (1980). *The Formulae-as-Types Notion of Construction*. Letters To H.B. Curry: Essays on Combinatory Logic.
- 10 Oleg Kiselyov, Simon Peyton Jones, and Chung-chieh Shan. (2010). Fun with Type Functions. *Pages 301–331 of: Reflections on the Work of C.A.R. Hoare*. London: Springer London.
- 11 Simon Marlow (Ed.). (2010). *Haskell 2010 Language Report*. Tech. rept.
- 12 Per Martin-Löf. (1984). *Intuitionistic Type Theory*. Bibliopolis.
- 13 Conor McBride. (2003). First-Order Unification by Structural Recursion. *Journal of functional programming*, **13**(6), 1061–1075.
- 14 Conor McBride, and James McKinna. (2004). The View From The Left. *Journal of functional programming*, **14**(1), 69–111.
- 15 Robin Milner. (1978). A Theory of Type Polymorphism in Programming. *Journal of computer and system sciences*, **17**, 348–375.
- 16 Peter Morris, Thorsten Altenkirch, and Conor McBride. (2004). Exploring the Regular Tree Types. *Page 30 of: Types for Proofs and Programs*. SpringerVerlag.
- 17 Wolfgang Naraschewski, and Tobias Nipkow. (1999). Type Inference Verified: Algorithm W in Isabelle/HOL. *Journal of automated reasoning*, **23**(3), 299–318.
- 18 Ulf Norell. 2008 (May). Dependently Typed Programming In Agda. *AFP 08: Sixth International Summer School on Advanced Functional Programming*.

- 19 Martin Odersky, Martin Sulzmann, and Martin Wehr. (1997). Type Inference with Constrained Types. *Fourth International Workshop on Foundations of Object-Oriented Programming (FOOL)*.
- 20 Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. (2007). Practical Type Inference for Arbitrary-Rank Types. *Journal of functional programming*, **17**(01), 1–82.
- 21 François Pottier, and Didier Rémy. (2005). Essence of ML Type Inference. Benjamin Pierce (ed), *Advanced Topics in Types and Programming Languages*. Cambridge, Massachusetts: MIT Press.
- 22 Nicolas Pouillard. (2011). Nameless, Painless. *ICFP '11: Proceedings of the 16th ACM SIGPLAN international conference on Functional programming*. ACM Request Permissions.
- 23 John Alan Robinson. (1965). A Machine-Oriented Logic Based on the Resolution Principle. *Journal of the acm*, **12**(1), 23–41.
- 24 Tom Schrijvers, Simon Peyton Jones, Martin Sulzmann, and Dimitrios Vytiniotis. (2009). Complete and Decidable Type Inference for GADTs. *ICFP '09: Proceedings of the 14th ACM SIGPLAN international conference on Functional programming*. ACM Request Permissions.
- 25 Dimitrios Vytiniotis, Simon Peyton Jones, and Tom Schrijvers. (2010). Let Should Not Be Generalized. *TLDI '10: Proceedings of the 5th ACM SIGPLAN workshop on Types in language design and implementation*. ACM Request Permissions.
- 26 Dimitrios Vytiniotis, Simon Peyton Jones, Tom Schrijvers, and Martin Suzmann. (2011). OutsideIn(X): Modular Type Inference with Local Assumptions. *Journal of functional programming*, **21**(4-5), 333–412.
- 27 Philip Wadler. (1989). Theorems for free! *Proceedings of the Fourth International Conference on Functional Programming and Computer Architecture (FPCA '89)*.
- 28 Brent Yorgey, Stephanie Weirich, Simon Peyton Jones, Julian Cretin, Dimitrios Vytiniotis, and José Pedro Magalhães. (2012). Giving Haskell a Promotion. *Pages 53–66 of: Proceedings of the 8th ACM SIGPLAN workshop on Types in language design and implementation*. New York, NY, USA: ACM.