

Partial Type Signatures for Haskell

Extended Version with Proofs of the Theorems

Thomas Winant¹, Dominique Devriese¹, Frank Piessens¹, and Tom Schrijvers²

¹ KU Leuven `firstname.lastname@cs.kuleuven.be`

² UGent `tom.schrijvers@ugent.be`

Abstract. Strong type systems can be used to increase the reliability and performance of programs. In combination with type inference the overhead for the programmer can be kept small. Nevertheless, explicit type signatures often remain needed or useful. In languages with standard Hindley-Milner-based type systems, programmers have a binary choice between omitting the type of a value or function (and rely on type inference) or explicitly providing the type entirely; there are no intermediate options. Some proposals for partial type signatures exist, but none support features like local constraints and GHC’s non-generalisation of local bindings. Therefore we propose and motivate a practical form of partial type signatures that can be used with present-day Haskell. We formally describe our proposal as an extension of the `OUTSIDEIN(X)` system and we prove some of its properties. We have developed a (not yet complete) implementation for the GHC Haskell compiler. We find that our design fits naturally in both the `OUTSIDEIN(X)` formalism and the compiler. Our wildcards map nicely to HM unification variables and our generalisation rules are carefully designed to align with existing behaviour. This is the extended version of a paper submitted to the PADL symposium. It is identical to that paper except for two appendices containing proofs of our main theorems.

1 Introduction

Static type checking can help catch errors at compile-time and provide useful information for compiler optimisations. Through the use of type inference, programmers are not required to provide explicit type signatures for all values in a program. Nevertheless, explicit signatures can still be needed or useful: type signatures provide a form of machine-checked documentation, they can be used to make general inferred types more specific, and help to verify whether the program corresponds to a programmer’s intentions.

Haskell’s overloaded math operators exemplify the need for type signatures:

```
let harmonic  $x\ y = \frac{2}{\frac{1}{x} + \frac{1}{y}}$  in print (harmonic 3 2)
```

Under Haskell’s defaulting, x and y are interpreted as floating point numbers leading to the inexact output `2.4000000000000004`. The exact output $\frac{12}{5}$ is produced with the signature `harmonic :: Rational -> Rational -> Rational`. Without defaulting, an ambiguous type variable would make a type signature mandatory.

Additionally, type inference is fundamentally limited. It is impossible to infer types for all programs that are typeable in more complex type systems (see e.g. [1, §23.6]). Consider the following Haskell program.

```
foo x = (x [True, False], x ['a', 'b'])
test = foo reverse
```

This program is rejected by Haskell's type checker, because of the Damas-Milner rule that a lambda-bound argument (like x) must have a monomorphic type. x could be assigned the type $[Bool] \rightarrow [Bool]$, or $[Char] \rightarrow [Char]$, but not $\forall a.[a] \rightarrow [a]$, see e.g. [2]. With a correct signature, the program is accepted:

```
foo :: ( $\forall a.[a] \rightarrow [a]$ )  $\rightarrow$  ([Bool], [Char])
foo x = (x [True, False], x ['a', 'b'])
```

Haskell, like many other programming languages provides a binary, all-or-nothing choice when it comes to type signatures: either the programmer writes the whole signature or none at all. Nevertheless, in many of the situations where type signatures are needed or useful, it suffices to pin down certain parts of the type. Providing the full type is unneeded and sometimes tedious or distracting. For example, when types are intended to document the code or to make its inferred type more specific, this is often only needed for one argument of a function or for the monad in which a computation runs, but not its result type. For example, only the type of foo 's argument cannot be inferred, but its result type can. In cases where we want or need to specify only a part of a type, it can be beneficial to *not* specify the rest. That remainder can be boilerplate, tedious or obscure the intention of the type signature. Not providing this information can save the programmer some thought and work, especially if the uninteresting bits of the type are unknown or prone to frequent change during development.

For such cases, *partial type signatures* can specify a type only partially and leave the rest for the type inferencer to decide. For foo , we could use:

```
foo :: ( $\forall a.[a] \rightarrow [a]$ )  $\rightarrow$  _
foo x = (x [True, False], x ['a', 'b'])
```

This partial signature specifies that foo is a function and defines the polymorphic type of foo 's first argument. The result type is unspecified, as indicated by a type wildcard (written $_$). Similarly, for the *harmonic* example, it would suffice to write the shorter signature $harmonic :: Rational \rightarrow _$.

At this point, we should mention some partial workarounds for the lack of partial signatures in Haskell. foo could for example use a pattern type signature:

```
foo (x ::  $\forall a.[a] \rightarrow [a]$ ) = ...
```

However, such a solution only applies if the parts of the type we want/do not want to specify happen to coincide with the type of one or more input variables and the rest. Expression type signatures similarly provide a partial solution. Another way to simulate partial type signatures uses explicitly typed identity functions in the implementation; we could write foo for example as follows:

```
foo = (id ::  $\forall b.(\forall a.[a] \rightarrow [a]) \rightarrow b$ ) ( $\lambda x \rightarrow$  (x [True, False], x ['a', 'b']))
```

A downside is that *foo*'s implementation is obscured with computationally insignificant code and Kiselyov has proposed to place such code in fake clauses [3]:

```
foo x | False = (id :: (∀a.[a] → [a]) → b) x
foo x = (x [True, False], x ['a', 'b'])
```

A combinator library supports this technique [4]. These workarounds are generally poorly legible, cumbersome to use (e.g. requiring lambda functions instead of left-hand-side patterns) and limited (e.g. only a lower bound on type constraints). Their existence does prove the need for actual partial type signatures.

We propose and study a form of partial type signatures in the context of a language with HM-based type inference. Our partial type signatures extend normal signatures with type wildcards ($_$). During type inference, such wildcards can be instantiated to arbitrary types, e.g. the type $_ \rightarrow _$ can be instantiated to $Int \rightarrow (Bool \rightarrow Int)$ or $(Int \rightarrow Bool) \rightarrow String$. They map nicely to the unification variables used internally by most type inferencers.

In the context of HM-based type inference, we take care to properly interact with the type generalisation that is performed to achieve let-polymorphism. If (part of) the type instantiating a wildcard is not restricted by type inference, a HM-style type inferencer will quantify over it. Consider the following program.

```
bar :: _ → _
bar _ = True
```

From the return value *True*, the type checker learns that the second wildcard in the partial signature of *bar* must be instantiated to *Bool*. However, the first wildcard remains open. In this case, type generalisation will infer *bar*'s principal type $\forall a.a \rightarrow Bool$, like when the type signature is omitted entirely.

A second, related challenge is dealing with constraints, for example type class constraints (e.g. $\forall a.Num\ a \Rightarrow a$) and equality constraints (e.g. $(Fun1\ a \sim (b \rightarrow b)) \Rightarrow a \rightarrow b$) supported by GHC. Our partial signatures allow the inference of additional constraints if and only if the type contains an *extra-constraints wildcard*, written as an underscore just before the double arrow: $_ \Rightarrow a \rightarrow b$. For example, the signature $_ \rightarrow b$ (without an extra-constraints wildcard) forbids types with additional constraints like $Num\ b \Rightarrow Int \rightarrow b$. That type can be allowed explicitly with the signature $_ \Rightarrow _ \rightarrow b$. Only one extra-constraints wildcard can be present and allows any number of constraints to be added.

In a GHC ticket discussion, Peyton Jones has argued the usefulness of an extra-constraints wildcard based on the following example (where the wildcard would be instantiated to $(Num\ a, Show\ a)$) [5]:

```
f :: _ ⇒ [a] → String
f xs = show (sum xs)
```

We also allow multiple references to a wildcard within a signature using *named wildcards* (written as e.g. $_a$). In the following example, we use them to shorten a tedious type signature.

```
isMeltdown :: NukeMonad param1 param2 Bool
unlessMeltdown :: _nm () → _nm ()
```

```
unlessMeltdown c = do m ← isMeltdown  
                if m then return () else c
```

To make our proposal precise, we give a formal account based on Vytiniotis et al.’s OUTSIDEIN(X) formalism [6]. We define natural and algorithmic typing rules and prove their correspondence. Additionally, we prove that our new rules generalise the old ones for signatures without wildcards and that a partial signature $f :: _ \Rightarrow _$ has the same effect as no signature at all. Such correspondences are important for consistency and to align with users’ expectations.

We have an implementation of our proposal in the Glasgow Haskell Compiler, but it is not yet complete at the time of writing. Our current version correctly unifies wildcards with concrete types, but unifying with open types, generalisation and the extra-constraints wildcard are not yet working as we intend. We hope to finish our modifications in the coming months.

Contributions The idea of partial type signatures is not novel. Several languages support them in some form or other [7,8] and they have been proposed for Haskell several times before [9,10]. Dijkstra [11] and Sulzmann and Wazny [12,13] have detailed proposals for Haskell-like languages. Still, we believe that ours is the first rigorous formalisation of partial type signatures for a HM-style inference that supports all the features of present-day Haskell. Specifically, we support local constraints (that arise e.g. from pattern matching on GADTs) and align with GHC’s non-generalisation of local bindings.

More specifically, our contributions are the following:

- A formalised proposal for partial type signatures, including generalisation, in a Hindley-Milner-style type inference system. Our work plugs into the constraint-based type inference approach OUTSIDEIN(X) [6], currently employed by the de facto standard Haskell compiler GHC.
- We align our partial type signatures with the OUTSIDEIN(X) policy that *let should not be generalised*.
- We formally show that the new typing rules generalise the existing rules for signatures without wildcards and for omitted signatures.
- A (not yet complete) implementation in the GHC Haskell compiler.

Outline In Section 2, we describe our additional syntax, both informally and formally. Formal rules for handling wildcard syntax are listed in Section 3. We extend OUTSIDEIN(X) typing rules to support wildcards in Section 4. Local bindings with partial type signatures are described in Section 5. We prove the correspondence of our rules to the standard ones for the uninformative signature $_ \Rightarrow _$ and for signatures without wildcards in section 6. We discuss our implementation in Section 7, related work in Section 8 and conclude in Section 9.

This extended version is identical to the submitted version except for the two appendices which contain the proofs of our formal results, which are omitted from the submitted version for space reasons.

Term variables	$\in x, y, z, f, g, h$
Type variables	$\in a, b, c$
Named wildcards	$\in \underline{a}, \underline{b}, \underline{c}$
Data constructors	$\in K$
	$\nu ::= K \mid x$
Programs	$prog ::= \epsilon \mid f = e, prog \mid$ $f ::= \underline{\sigma} = e, prog$
Expressions	$e ::= \nu \mid \lambda x. e \mid e_1 e_2 \mid$ $\text{case } e \text{ of } \{K \bar{x} \rightarrow e\}$
Type schemes	$\sigma ::= \forall \bar{a}. Q \Rightarrow \tau$
Type schemes with wildcards	$\underline{\sigma} ::= \forall \bar{a}. \underline{Q} \Rightarrow \underline{\tau}$
Constraints	$Q ::= \epsilon \mid Q_1 \wedge Q_2 \mid \tau_1 \sim \tau_2 \mid \mathbf{D} \bar{\tau} \mid \dots$
Constraints with wildcards	$\underline{Q}^w ::= \underline{Q} \mid \underline{Q}_1^w \wedge \underline{Q}_2^w \mid \underline{\tau}_1 \sim \underline{\tau}_2 \mid \mathbf{D} \underline{\tau} \mid \dots$
Constraints with extra constraints wildcard	$\underline{Q} ::= \underline{Q}^w \mid \underline{Q}^w \wedge _$
Monotypes	$\tau, \nu ::= tv \mid \mathbf{Int} \mid \mathbf{Bool} \mid [\tau] \mid \mathbf{T} \bar{\tau} \mid \dots$
Monotypes with wildcards	$\underline{\tau}, \underline{\nu} ::= \tau \mid _ \mid \underline{a} \mid [\underline{\tau}] \mid \mathbf{T} \underline{\tau}$
Type environments	$\Gamma ::= \epsilon \mid (\nu : \sigma), \Gamma$
Free type variables	$ftv(\cdot)$
Top-level axiom schemes	$\mathcal{Q} ::= \epsilon \mid \mathcal{Q} \wedge \mathcal{Q} \mid \forall \bar{a}. Q \Rightarrow Q$
Unification variables	$\in \alpha, \beta, \gamma, \omega, \dots$
Unifiers	$\theta, \varphi ::= [\bar{\alpha} \mapsto \bar{\tau}]$
Unification or rigid (skolem) variables	$tv ::= \alpha \mid a$
Algorithm-generated constraints	$C ::= Q \mid C_1 \wedge C_2 \mid \exists \bar{\alpha}. (Q \supset C)$
Free unification variables	$fuw(\cdot)$
Named wildcards	$nwc(\cdot)$

Fig. 1. Wildcard syntax extension of [6, Fig. 1, page 12 and Fig. 5, page 17]

2 Wildcard syntax

In the introduction we already gave an informal account of the wildcard syntax we support. We quickly reiterate and formalise the syntax of wildcards as an extension of the syntax in `OUTSIDEIN(X)` [6]. Figure 1 contains the formal definitions with additions and changes highlighted in grey.

First of all, *type wildcards* can take the place of monotypes, e.g. $f :: _ \rightarrow _$. For type inference, they are translated to unification variables (see Section 3.2). By convention, we write unification variables that arise from wildcards as $\omega_1, \omega_2, \dots$.

A wildcard *in* a constraint is called a *constraint wildcard*, e.g. $Eq _ \Rightarrow _$. A wildcard occurring *as* a constraint is an *extra-constraints wildcard*, e.g. $_ \Rightarrow _$. When it is present, any number of constraints may be added to the type during inference. Because one extra-constraints wildcard can be instantiated to any number of constraints, more than one such wildcard would be pointless. For clarity, we allow only one and require that it comes last in the list of constraints.

Additionally, we support *named wildcards*, e.g. $_a \rightarrow _a$. All instances of a named wildcard within a partial type signature must unify with the same type. Named wildcards are particularly useful to express constraints on wildcard

types, e.g. $Eq _a \Rightarrow _a$ or $(_a \sim b) \Rightarrow _a \rightarrow [b]$. Although syntactically similar, named wildcards should not be confused with type variables: they can unify with concrete types. Only when not unified with concrete types, they are generalised over and behave like type variables.

In Figure 1 we provide variants of type schemes (σ), constraints (Q), and monotypes (τ) that *can* contain wildcards, respectively $\underline{\sigma}$, \underline{Q} , and $\underline{\tau}$. A distinction between *constraints with wildcards* (\underline{Q}^w) and *constraints with [an] extra-constraints wildcard* (\underline{Q}) is made to enforce that the extra-constraints wildcard can occur at most once and must come last.

3 Wildcard instantiation and desugaring

Before we introduce the adapted typing rules, we formalise the relation between wildcards and types. To this end, we define two judgments: the *wildcard instantiation judgment* and the *wildcard desugaring judgment*. They are employed in Section 4 by the natural and algorithmic typing rules respectively and the latter should be understood as an algorithmic version of the former.

3.1 Wildcard instantiation

The wildcard instantiation judgment $\underline{Q}; \underline{\tau} \Rightarrow Q; \tau$ can be read as “The wildcards in constraints \underline{Q} and monotype $\underline{\tau}$ can be instantiated to obtain constraints Q and monotype τ ”. Each wildcard in \underline{Q} and $\underline{\tau}$ corresponds to a concrete type or a type variable in Q and τ . Remember that \underline{Q} and $\underline{\tau}$ can contain wildcards, whereas Q and τ cannot. This judgment will be used by the adapted typing rules to instantiate a partial type signature to a type signature without wildcards.

The rules of the judgment are shown in Figure 2. The rule NAMEDWC requires monotypes \bar{v} that are substituted by the named wildcards in \underline{Q} and $\underline{\tau}$. We then delegate to two subjudgments that instantiate the unnamed wildcards in respectively \underline{Q}^w and τ . The rule EXTRAWC states that an extra-constraints wildcard can be instantiated to an arbitrary conjunction of constraints Q_{TES} , which can consist of zero or more constraints. Remember that \underline{Q} can contain an extra-constraints wildcard and \underline{Q}^w cannot.

The first subjudgment $\underline{\tau} \Rightarrow^f \tau$ instantiates wildcards in a monotype to concrete types or type variables. The rule TYWC states that a type wildcard can be instantiated to any monotype τ . A monotype without wildcards is instantiated to itself (TYNOWC) and there is a congruence rule for type constructor applications (TYAPP). Note that function types: (\rightarrow) , tuples: $(,)$, lists: $[, \dots]$ are all treated as type constructor applications.

The second subjudgment $\underline{Q}^w \Rightarrow^c Q$ instantiates wildcards in constraints to concrete types or type variables. Constraints without wildcards need no further wildcard instantiation (CONNOWC). A conjunction of constraints is handled recursively in CONCONJ. A type-class constraint can also contain wildcards (CONTC), which will be instantiated using the previously described subjudgment. Type wildcards in equality constraints are handled in CONEQ.

$$\boxed{Q; \tau \Rightarrow Q; \tau}$$

$$\frac{\overline{a} = nwc(\underline{\tau}) \cup nwc(Q^w) \quad [\underline{a} \mapsto v] Q^w \Rightarrow^c Q \quad [\underline{a} \mapsto v] \underline{\tau} \Rightarrow^t \tau}{\underline{Q}^w; \underline{\tau} \Rightarrow Q; \tau} \text{NAMEDWC}$$

$$\frac{\underline{Q}^w; \underline{\tau} \Rightarrow Q; \tau}{\underline{Q}^w \wedge \underline{\tau}; \underline{\tau} \Rightarrow Q \wedge Q_{res}; \tau} \text{EXTRAWC}$$

$$\frac{}{\underline{\tau} \Rightarrow^t \tau} \text{TYWC} \quad \frac{}{\tau \Rightarrow^t \tau} \text{TYNOWC} \quad \frac{\forall i. \underline{\tau}_i \Rightarrow^t \tau_i}{\mathbf{T} \underline{\tau}_i \Rightarrow^t \mathbf{T} \overline{\tau}_i} \text{TYAPP}$$

$$\frac{}{Q \Rightarrow^c Q} \text{CONNOWC} \quad \frac{\underline{Q}_1^w \Rightarrow^c Q_1 \quad \underline{Q}_2^w \Rightarrow^c Q_2}{\underline{Q}_1^w \wedge \underline{Q}_2^w \Rightarrow^c Q_1 \wedge Q_2} \text{CONCONJ}$$

$$\frac{\forall i. \underline{\tau}_i \Rightarrow^t \tau_i}{\mathbf{D} \underline{\tau}_i \Rightarrow^c \mathbf{D} \overline{\tau}_i} \text{CONTC} \quad \frac{\underline{\tau}_1 \Rightarrow^t \tau_1 \quad \underline{\tau}_2 \Rightarrow^t \tau_2}{\underline{\tau}_1 \sim \underline{\tau}_2 \Rightarrow^c \tau_1 \sim \tau_2} \text{CONEQ}$$

Fig. 2. Natural wildcard instantiation judgment rules

3.2 Wildcard desugaring

We also define an algorithmic variant of the wildcard instantiation judgment, the *wildcard desugaring judgment*. Instead of instantiating wildcards to concrete types or type variables as the wildcard instantiation judgment does, the wildcard desugaring judgment replaces them by fresh unification variables in order to participate in $\text{OUTSIDEIN}(X)$'s type inference.

The wildcard desugaring judgment $\underline{Q}; \underline{\tau} \Rightarrow_a Q; \tau; extra$ can be read as: replacing all the wildcards in \underline{Q} and $\underline{\tau}$ with fresh unification variables, gives us Q , τ , and *extra*. This last boolean output parameter indicates whether the constraints contained an extra-constraints wildcard or not, e.g. the first underscore in $\underline{\tau} \Rightarrow \underline{\tau}$. If and only if *extra* = true, extra constraints can be generated.

The rules of this judgment are shown in Figure 3. As their structure resembles the structure of the rules of the wildcard instantiation judgment, we shall only highlight the differences.

If \underline{Q} contains an extra-constraints wildcard, *extra* will be true (AEXTRAWC). Subsequently, or if it did not contain such a wildcard, the named wildcards in \underline{Q}^w and $\underline{\tau}$ are collected and replaced with fresh unification variables $\omega_1, \omega_2, \dots$ in the rule ANAMEDWC. Note that multiple occurrences of a named wildcard are replaced with the same unification variable. Unnamed wildcards in $\underline{\tau}$ and \underline{Q}^w are desugared separately by two subjudgments $\underline{\tau} \Rightarrow_a^t \tau$ and $\underline{Q}^w \Rightarrow_a^c Q$ respectively. The only difference with the corresponding wildcard instantiation subjudgments is that in the rule ATYWC, a wildcard is replaced with a fresh unification variable instead of a monotype τ .

$$\begin{array}{c}
\boxed{Q; \underline{\tau} \Rightarrow_a Q; \tau; \text{extra}} \\
\\
\frac{\overline{a} = \text{nwc}(\underline{\tau}) \cup \text{nwc}(Q^w) \quad \overline{w} \text{ fresh} \quad \frac{[\underline{a} \mapsto \omega] Q^w \Rightarrow_a^c Q \quad [\underline{a} \mapsto \omega] \underline{\tau} \Rightarrow_a^t \tau}{Q^w; \underline{\tau} \Rightarrow_a Q; \tau; \text{false}} \text{ANAMEDWC} \quad \frac{Q^w; \underline{\tau} \Rightarrow_a Q; \tau; \text{false}}{Q^w \wedge _; \underline{\tau} \Rightarrow_a Q; \tau; \text{true}} \text{AEXTRAWC}}{Q^w; \underline{\tau} \Rightarrow_a Q; \tau; \text{false}} \\
\\
\frac{\omega \text{ fresh} \quad \text{ATYWC} \quad \frac{\underline{\tau} \Rightarrow_a^t \tau}{\tau \Rightarrow_a^t \tau} \text{ATYNOWC} \quad \frac{\forall i. \underline{\tau}_i \Rightarrow_a^t \tau_i}{\text{T } \underline{\tau}_i \Rightarrow_a^t \text{T } \tau_i} \text{ATYAPP}}{_ \Rightarrow_a^t \omega} \\
\\
\frac{\frac{Q^w \Rightarrow_a^c Q}{Q \Rightarrow_a^c Q} \text{ACONNOWC} \quad \frac{Q_1^w \Rightarrow_a^c Q_1 \quad Q_2^w \Rightarrow_a^c Q_2}{Q_1^w \wedge Q_2^w \Rightarrow_a^c Q_1 \wedge Q_2} \text{ACONCONJ}}{\frac{\forall i. \underline{\tau}_i \Rightarrow_a^t \tau_i}{\text{D } \underline{\tau}_i \Rightarrow_a^c \text{D } \tau_i} \text{ACONTC} \quad \frac{\underline{\tau}_1 \Rightarrow_a^t \tau_1 \quad \underline{\tau}_2 \Rightarrow_a^t \tau_2}{\underline{\tau}_1 \sim \underline{\tau}_2 \Rightarrow_a^c \tau_1 \sim \tau_2} \text{ACONTC}}{Q^w \Rightarrow_a^c Q}
\end{array}$$

Fig. 3. Algorithmic wildcard desugaring judgment rules

4 Typing rules

When checking a partial type signature, the wildcards are unified with concrete types if necessary, otherwise they are replaced with fresh universally quantified type variables, i.e. the type is generalised. If an extra-constraints wildcard is present, additional constraints may be generated and added to the annotated constraints. We formalise this by adapting the `OUTSIDEIN(X)` typing rules [6].

4.1 Natural typing rules

Figure 4 shows the three top-level natural typing rules in [6]: `EMPTY`, the base case, `BIND`, for definitions without a type signature, and `BINDA`, for definitions with a signature. It also shows the new rule `BINDPA` which replaces `BINDA`. Changes in `BINDPA` w.r.t. `BINDA` are greyed. The rules refer to the *constraint entailment* judgment $\mathcal{Q} \Vdash Q$, which should be read as: “the axioms \mathcal{Q} imply Q ”.

Compared to `BINDA`, `BINDPA` supports partial type signatures. It is extended with the premise $\underline{Q}; \underline{\tau} \Rightarrow Q; \tau$, i.e. \underline{Q} and $\underline{\tau}$ are instantiated to Q and τ (see Section 3.1). Additional type variables that were not present in the partial type signature but arose from the generalisation of the type, are captured in \overline{b} , and also universally quantified over in the final type of the top-level definition.

4.2 Constraint solver

Before discussing the new top-level algorithmic typing rules, which make use of `OUTSIDEIN(X)`’s constraint solver, we shall briefly describe the constraint

$$\begin{array}{c}
\boxed{\mathcal{Q}; \Gamma \vdash prog} \quad \frac{ftv(\Gamma) = fuv(\mathcal{Q}) = \emptyset}{\mathcal{Q}; \Gamma \vdash \epsilon} \text{EMPTY} \\
\\
\frac{Q_1; \Gamma \vdash e : \tau \quad \bar{a} = ftv(Q) \cup fuv(\tau) \quad \mathcal{Q} \wedge Q \Vdash Q_1}{\mathcal{Q}; \Gamma, (f : \forall \bar{a}. Q \Rightarrow \tau) \vdash prog} \text{BIND} \\
\\
\frac{Q_1; \Gamma \vdash e : \tau \quad \bar{a} = ftv(Q) \cup fuv(\tau) \quad \mathcal{Q} \wedge Q \Vdash Q_1}{\mathcal{Q}; \Gamma, (f : \forall \bar{a}. Q \Rightarrow \tau) \vdash prog} \text{BINDA} \\
\\
\frac{\boxed{Q; \underline{\tau} \Rightarrow Q; \tau} \quad Q_1; \Gamma \vdash e : \tau \quad \bar{a} \boxplus \bar{b} = ftv(Q) \cup fuv(\tau)}{\mathcal{Q} \wedge Q \Vdash Q_1 \quad \mathcal{Q}; \Gamma, (f : \forall \bar{a}\bar{b}. Q \Rightarrow \tau) \vdash prog} \text{BINDPA} \\
\\
\mathcal{Q}; \Gamma \vdash f : \forall \bar{a}. \underline{Q} \Rightarrow \underline{\tau} = e, prog
\end{array}$$

Fig. 4. Natural top-level typing rules, adapted from [6, Fig. 4, p. 15]

solver [6, Section 5.5, p. 41]. The $\text{OUTSIDEIN}(X)$ type inference system is in fact parameterised by a constraint domain X . For present-day Haskell, X would be instantiated to a constraint domain that contains type-class and equality constraints (and Vytiniotis et al. present a concrete solver for this X [6]), but the $\text{OUTSIDEIN}(X)$ typing rules and algorithms are designed to support alternative domains as well. In this text, we keep X abstract. We will only describe the form of the constraint solver, not the implementation, which is specific to X .

We have already seen the natural constraint entailment relation $\mathcal{Q} \Vdash Q$. On the algorithmic side, the constraint solver (Figure 5) has the following signature.

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

The inputs in this signature are:

- \mathcal{Q} : the top-level axiom scheme. In a concrete setting, it will contain for example class instances or reduction rules of type functions, but we will leave it abstract. It does not change during type-checking.
- Q_{given} : the given constraints that arise from type annotations (or pattern matching),
- $\bar{\alpha}_{tch}$: the *touchable* unification variables that the solver is allowed to instantiate, and
- C_{wanted} : the constraints to be solved.

The outputs are:

- $Q_{residual}$: residual constraints that the solver has not been able to solve, and
- θ : a substitution mapping unification variables to types, with $\text{dom}(\theta) \subseteq \bar{\alpha}_{tch}$.

Vytiniotis et al. keep the constraint solver abstract, but require certain properties of it. It is required to be *sound* and yield *guess-free solutions*, two formal properties (specified in terms of the natural constraint entailment relation \Vdash) that we do not go into further. We will however require the solver to support a somewhat larger form of inputs. In the next section, we explain this further.

4.3 Wildcards in constraints

We have chosen to allow both named and unnamed wildcards in constraints. Nevertheless, it is important to point out a limitation of such wildcards in our system. The `OUTSIDEIN(X)` infrastructure will never apply unification to two constraints. Consider the following example

$$\begin{aligned} h &:: Eq _ \Rightarrow a \rightarrow a \rightarrow Bool \\ h &= (\equiv) \end{aligned}$$

In this case, h 's implementation generates the wanted constraint $Eq \ a$, which one might expect to be unified with $Eq _$, so that the wildcard is instantiated with type a , but this is not what happens. The `OUTSIDEIN(X)` constraint solver does not unify the given constraint $Eq _$ with the wanted constraint $Eq \ a$. In general, it will never unify one constraint with another; the algorithm will only instantiate wildcards $_a$ in constraints C if

- $_a$ is a named wildcard also mentioned in the non-constraint part of the signature and it is instantiated during unification with the inferred type.
- The instantiation follows semantically from the constraint, i.e. $C \supset _a \sim \dots$

In `OUTSIDEIN(X)`, unifying the non-constraint part of a signature with the inferred part happens through the generation of equality constraints, so in this sense the first case is comprised in the second. As a result, for h we get an error that the constraint $Eq \ a$ cannot be solved from given constraints $Eq _$.

Nevertheless, this limitation does not mean that wildcards in constraints are useless. Consider the following example:

$$f :: Monad _m \Rightarrow _m Bool$$

For this signature, $_m$ can either be unified with a concrete type constructor like `Maybe` for which there is a `Monad` instance or be generalised to a universally quantified monad m . Similarly, we can say something like

$$g :: (_a, _) \sim F _b \Rightarrow _b \rightarrow _a$$

This signature states that g is a function whose domain type is mapped by type function `F` to a tuple whose first element is its range type.

Formally, the choice to allow wildcards in constraints implies that we have to drop an invariant of the constraint solver. For the constraint entailment relation $\mathcal{Q}; Q_{given}; \bar{\alpha}_{tch} \xrightarrow{solv} C_{wanted} \rightsquigarrow Q_{residual}; \theta$, Vytiniotis et al. mention two invariants that should hold: $\bar{\alpha}_{tch} \#_{fvv}(Q_{given})$ and $dom(\theta) \#_{fvv}(Q_{given})$, i.e. the free unification variables in Q_{given} should not be unified. This means that the given constraints Q_{given} are not allowed to contain unification variables that the solver can instantiate. In order to support wildcards in constraints, it is required to remove this restriction and we propose to do so. Although the proofs about the concrete constraint solver in the second part of Vytiniotis et al.'s paper [6] rely upon this restriction, we conjecture that this is a technical restriction that can be remedied.

Contrary to the behaviour of wildcards in the non-constraint part of a signature, some of the behaviour of wildcards in constraints we just discussed could be

$$\begin{array}{c}
\boxed{\mathcal{Q}; \Gamma \vdash \text{prog}} \quad \frac{}{\mathcal{Q}; \Gamma \vdash \epsilon} \text{EMPTY} \\
\\
\frac{\Gamma \vdash e : \tau \rightsquigarrow C \quad \mathcal{Q}; \epsilon; \text{fuv}(\tau) \cup \text{fuv}(C) \vdash^{\text{solv}} C \rightsquigarrow \mathcal{Q}; \theta \quad \bar{a} \text{ fresh} \\
\bar{\alpha} = \text{fuv}(\theta\tau) \cup \text{fuv}(\mathcal{Q}) \quad \mathcal{Q}; \Gamma, (f : \forall \bar{a}. [\bar{\alpha} \mapsto \bar{a}] (Q \Rightarrow \theta\tau)) \vdash \text{prog}}{\mathcal{Q}; \Gamma \vdash f = e, \text{prog}} \text{BIND} \\
\\
\frac{\Gamma \vdash e : v \rightsquigarrow C \quad \mathcal{Q}; \mathcal{Q}; \text{fuv}(v) \cup \text{fuv}(C) \vdash^{\text{solv}} C \wedge v \sim \tau \rightsquigarrow \epsilon; \theta \\
\mathcal{Q}; \Gamma, (f : \forall \bar{a}. Q \Rightarrow \tau) \vdash \text{prog}}{\mathcal{Q}; \Gamma \vdash f :: \forall \bar{a}. Q \Rightarrow \tau = e, \text{prog}} \text{BINDA} \\
\\
\boxed{\mathcal{Q}; \mathcal{Q}_{\text{given}}; \bar{\alpha}_{\text{tch}} \vdash^{\text{solv}} C_{\text{wanted}} \rightsquigarrow \mathcal{Q}_{\text{residual}}; \theta}
\end{array}$$

Fig. 5. Top-level algorithmic rules, taken from [6, Fig. 12, page 39]

$$\begin{array}{c}
\Gamma \vdash e : v \rightsquigarrow C \quad \boxed{\mathcal{Q}; \underline{\tau} \Rightarrow_a \mathcal{Q}; \tau; \text{extra}} \\
\mathcal{Q}; \mathcal{Q}; \text{fuv}(v) \cup \text{fuv}(C) \cup \text{fuv}(\tau, \mathcal{Q}) \vdash^{\text{solv}} C \wedge v \sim \tau \rightsquigarrow \mathcal{Q}_{\text{res}}; \theta \\
\boxed{\text{extra} \vee (\mathcal{Q}_{\text{res}} = \epsilon)} \quad \beta = \text{fuv}(\theta\tau) \cup \text{fuv}(\theta\mathcal{Q} \wedge \mathcal{Q}_{\text{res}}) \quad \bar{b} \text{ fresh} \\
\frac{\mathcal{Q}; \Gamma, (f : \forall \bar{a} \bar{b}. [\bar{\beta} \mapsto \bar{b}] (\theta\mathcal{Q} \wedge \mathcal{Q}_{\text{res}} \Rightarrow \theta\tau)) \vdash \text{prog}}{\mathcal{Q}; \Gamma \vdash f :: \forall \bar{a}. \mathcal{Q} \Rightarrow \underline{\tau} = e, \text{prog}} \text{BINDPA}
\end{array}$$

Fig. 6. New top-level algorithmic rule, adapted from Fig. 5

unexpected by programmers. Because of this, one might consider the possibility to disallow both named and unnamed type wildcards in constraints, but we have chosen not to do this. We think the limitations of wildcards in constraints can be explained to the user and our examples show that they can be useful despite the limitations.

4.4 Algorithmic typing rules

In addition to the new top-level natural typing rule BINDPA, we also adapt the top-level algorithmic typing rules. The original top-level algorithmic typing rules are shown in Figure 5. As wildcards can only occur in a type signature, only the rule BINDA that handles declarations with a type annotation, has to be adapted. The adapted rule is presented in Figure 6, with changes w.r.t. BINDA highlighted in grey.

The BINDPA rule works as follows. We start with the first expression: the type v of e is inferred using the constraint generation judgment from [6] while generating the constraints C . The wildcards in \mathcal{Q} and $\underline{\tau}$ are replaced with fresh unification variables with the wildcard desugaring judgment we defined earlier. The *extra* output parameter indicates whether we are allowed to infer extra constraints.

On the second line, the invocation of the constraint solver has been slightly modified. The free unification variables in τ and \mathcal{Q} , introduced during the wildcard desugaring, are added to the set of touchable unification variables that

$$\begin{array}{c}
\frac{\Gamma \vdash e_1 : \tau_1 \rightsquigarrow C_1 \quad \Gamma, (x : \tau_1) \vdash e_2 : \tau_2 \rightsquigarrow C_2}{\Gamma \vdash \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 : \tau_2 \rightsquigarrow C_1 \wedge C_2} \text{LET} \\
\\
\frac{\Gamma \vdash e_1 : \tau \rightsquigarrow C_1 \quad \Gamma, (x : \tau_1) \vdash e_2 : \tau_2 \rightsquigarrow C_2}{\Gamma \vdash \mathbf{let} \ x :: \tau_1 = e_1 \ \mathbf{in} \ e_2 : \tau_2 \rightsquigarrow C_1 \wedge C_2 \wedge \tau \sim \tau_1} \text{LETA} \\
\\
\frac{\begin{array}{l} \sigma_1 = \forall \bar{a}. Q_1 \Rightarrow \tau_1 \quad Q_1 \neq \epsilon \text{ or } \bar{a} \neq \epsilon \quad \Gamma \vdash e_1 : \tau \rightsquigarrow C \\ \bar{\beta} = (fuv(\tau) \cup fuv(C)) - fuv(\Gamma) \quad C_1 = \exists \bar{\beta}. (Q_1 \supset C \wedge \tau \sim \tau_1) \\ \Gamma, (x : \sigma_1) \vdash e_2 : \tau_2 \rightsquigarrow C_2 \end{array}}{\Gamma \vdash \mathbf{let} \ x :: \sigma_1 = e_1 \ \mathbf{in} \ e_2 : \tau_2 \rightsquigarrow C_1 \wedge C_2} \text{GLETA} \\
\\
\boxed{\Gamma \vdash e : \tau \rightsquigarrow C} \\
\\
\frac{\begin{array}{l} \underline{\sigma}_1 = \forall \bar{a}. \underline{Q}_1 \Rightarrow \underline{\tau}_1 \quad \Gamma \vdash e_1 : \tau \rightsquigarrow C \quad \underline{Q}_1 ; \underline{\tau}_1 \Rightarrow_a Q_1 ; \tau_1 ; \text{false} \\ Q_1 \neq \epsilon \text{ or } \bar{a} \neq \epsilon \text{ or } fuv(\tau_1) \cup fuv(Q_1) \neq \emptyset \\ \bar{\beta} = ((fuv(\tau) \cup fuv(C)) - fuv(\Gamma)) \cup fuv(\tau_1) \cup fuv(Q_1) \\ C_1 = \exists \bar{\beta}. (Q_1 \supset C \wedge \tau \sim \tau_1) \quad \Gamma, (x : \forall \bar{a}. Q_1 \Rightarrow \tau_1) \vdash e_2 : \tau_2 \rightsquigarrow C_2 \end{array}}{\Gamma \vdash \mathbf{let} \ x :: \underline{\sigma}_1 = e_1 \ \mathbf{in} \ e_2 : \tau_2 \rightsquigarrow C_1 \wedge C_2} \text{GLETPA}
\end{array}$$

Fig. 7. Constraint generation for local **let**-bound definitions, taken and adapted from [6, Fig. 13, page 40]

the constraint solver is allowed to instantiate. We also capture the residual constraints, which were not allowed in the previous version of the rule, in Q_{res} . Now they are allowed, but only if *extra* is true.

In the next step, we collect the remaining free unification variables in $\theta\tau$ and $\theta Q \wedge Q_{res}$. These unification variables were not instantiated to concrete types while solving the constraints and so we generalise over them. They are replaced with fresh, universally quantified type variables, \bar{b} . The residual constraints, i.e. the extra constraints that have not been solved by the constraint solver, are added to the annotated constraints.

Exercise to the reader: apply the rule BINDPA to the following example.

$$\begin{array}{l}
f :: _ \Rightarrow _ \rightarrow _ \rightarrow Bool \rightarrow _ \\
f \ x \ y \ b = x \equiv y \wedge b
\end{array}$$

Theorem 1 (Algorithm soundness). *If $\mathcal{Q}; \Gamma \vdash prog$ then $\mathcal{Q}; \Gamma \vdash prog$ in a closed top-level Γ .*

5 Typing of local definitions

Advanced type system features like GADTs have a profound impact on a type system. Crucially, the clean and simple principal typing property that the HM system satisfies is no longer valid (see e.g. [6]). This makes type inference a harder problem and Vytiniotis et al. present one possible way out. They advocate the policy that the types of local (unannotated) definitions should not be generalised, with the slogan “Let should not be generalised”.

For partial type signatures of local definitions, we align with the policy to not generalise local definitions. Next, we present the adapted typing rules for local definitions, but we omit natural typing rules as the required changes are minimal. The existing algorithmic rules and our adapted rule are shown in Figure 7.

The rule LETA applies to definitions with an annotated monomorphic type, GLETA for polymorphic type signatures and LET for definitions without a signature. The rule LET is remarkably simple, because it applies the NOGEN policy of not generalising the inferred type at all. Our adapted typing rule GLETPA extends this policy to partial type signatures.

The GLETPA rule applies to local bindings with a partial type signature, either polymorphic or monomorphic. It first desugars the partial type signature. The *extra* parameter must be false, i.e. we forbid an extra-constraints wildcard, since the NOGEN policy forbids additional constraints. We verify that the type signature was indeed partial by requiring free unification variables in the desugared type and constraints. Next, the set of unification variables allowed to unify, i.e. the *touchables*, is extended with those resulting from the wildcards in the desugared τ_1 and Q_1 . Solving the implication constraint should unify them, fixing the definition’s actual type. The local binding, annotated with the desugared type, is added to the environment to type check the body e_2 . Following the NOGEN policy, no generalisation is performed.

The example *foo* shows the effect of not generalising in GLETPA.

```
foo = let g :: _ → _
      g x = x
      h :: Eq _a ⇒ _a → _a → Bool
      h x y = x ≡ y
      in (g True, g 'v', h True True, h 'a' 'b')
```

Instead of being quantified over, the free unification variables in the type of g unify with the *Bool* type at the first call of g . Thus, g ’s type is $Bool \rightarrow Bool$. As g is also called with a *Char* argument, the program will be rejected. Similarly, the unification variable for the named wildcard $_a$ in h ’s type is not generalised. Instead, it unifies with the *Bool* type, producing the type $Eq\ Bool \Rightarrow Bool \rightarrow Bool \rightarrow Bool$ for h .

6 Alignment with existing rules

Partial type signatures are a generalisation of the binary choice between a full signature or none at all. Using wildcards, partial type signatures can mix annotated and inferred types. To demonstrate that partial type signatures truly are a generalisation of the existing inference, we prove two properties.

First, partial type signatures are a conservative extension: the adapted typing rules are equivalent to the original rules for signatures without wildcards.

Second, definitions without a type signature are equivalent with definitions with a partial type signature of the form $_ \Rightarrow _$. More formally: the BINDPA rule (Figure 6) can be used to type check a definition $f = e$ without a type

signature by treating it as if it had the partial type signature $f :: _ \Rightarrow _ = e$. The following `ALTBIND` rule simply transforms definitions without a type signature into definitions with the equivalent partial type signature:

$$\boxed{\frac{\mathcal{Q}; \Gamma \vdash f :: _ \Rightarrow _ = e, \text{prog}}{\mathcal{Q}; \Gamma \vdash f = e, \text{prog}} \text{ALTBIND}}$$

Theorem 2. *Given a program prog in which every definition f has either a type signature without wildcards, i.e. $f :: \forall \bar{a}. Q \Rightarrow \tau = e$, or no type signature at all, i.e. $f = e$. If $\mathcal{Q}; \Gamma \vdash \text{prog}$, using `BIND`, `BINDA`, and `EMPTY` (Figure 5), then $\mathcal{Q}; \Gamma \vdash \text{prog}$, using `ALTBIND`, `BINDPA` (Figure 6), and `EMPTY` (Figure 5).*

These properties show that our proposal aligns well with the existing behaviour of type inference. This is not just theoretically important, but also shows that our proposal is natural and unsurprising for existing users.

7 Implementation and extensions

We have developed an implementation of our proposal in the de facto standard Haskell compiler GHC. GHC’s inferencer is based on the `OUTSIDEIN(X)` type inference system. As a result, our proposal fits relatively nicely into the compiler’s inference infrastructure. Nevertheless, GHC’s actual inferencer is (unavoidably) more complex than Vytiniotis et al.’s elegant theory, notably when it comes to the inference and generalisation of mutually recursive blocks and higher-rank types. Hence, our prototype currently implements only part of our theoretical development. More specifically, it correctly unifies wildcards with closed types, but does not yet support unifying with open types, generalisation and extra-constraints wildcards. The prototype code is available for download at <http://github.com/mrBliss/ghc>. We still intend to check and ensure compatibility with the `SCOPEDTYPEVARIABLES` [14] and `CONSTRAINTKINDS` [15,16] extensions, but we expect no major problems there.

8 Related work

Vytiniotis et al. provide a comprehensive overview of work on constraint-based type systems and type inference for advanced type system features that we do not repeat here [6], except to discuss aspects related to *partial* type signatures. Vytiniotis claim that their presentation is the first one that deals with local assumptions introduced by type signatures and data constructors, and where those local assumptions may include type equalities.

The idea of partial type signatures is not new. The topic regularly comes up on the Haskell community mailing lists. In two 2006 tickets on the Haskell Prime wiki (where the Haskell community proposes and tracks future language changes), Malcolm Wallace proposes a form of partial type signatures [10,9]. His

proposal seems similar to ours, but it does not contain a lot of detail. A GHC feature request has also been logged to request a form of constraint wildcards [5].

The programming language *Alice*, an extended version of Standard ML focused on concurrent and constraint programming, also features type wildcards. From the short paragraph in the online manual [8] on type wildcards, they appear similar to our partial type signatures (including w.r.t. the generalisation) but there are no constraint wildcards, as *Alice* does not have constraints. We have not found a formal description of *Alice*'s type wildcards.

The *Agda* programming language [7] has a dependent type system, that does not make a strict distinction between types and values. The type system allows a powerful form of type-level computation, so that type inferencing becomes harder. On the other hand, terms may appear in types, so that the inferencer can sometimes infer terms as well. In *Agda*, any value or type can be replaced by an underscore, in which case *Agda* will try to infer the value from the available type information using a conservative unification-based inference approach that performs well in practice. *Agda*'s type inference does not perform any form of generalisation: if the type checker cannot infer the value of such a meta-variable, it just reports an error.

Our work was inspired by the partial signatures in Dijkstra and Swierstra's *Explicit Haskell* [11][17, Chapter 10]. They also use wildcards and allow predicate wildcards very similar to our extra-constraints wildcards. However, where we follow Vytiniotis et al. in using a rather standard form of HM style type generalisation, Dijkstra and Swierstra use *quantifier location inference* rules that differ significantly, both for normal and partial type signatures. They argue that depending on the structure of the type in which a type variable appears, it should either be existentially or universally quantified to align with user expectations. For example, the type $a \rightarrow a$ is interpreted as $\forall a.a \rightarrow a$ but $(a \rightarrow a) \rightarrow Int$ is interpreted as $(\forall a.a \rightarrow a) \rightarrow Int$, unlike Haskell. In a product type, the variables are quantified existentially instead of universally, e.g. (a, a) is interpreted as $\exists a.(a, a)$ and $(a, a) \rightarrow Int$ as $(\exists a.(a, a)) \rightarrow Int$. Dijkstra and Swierstra formalise Explicit Haskell, but do not prove results like our theorems 1 and 2.

For the *Chameleon* programming language, Sulzmann and Wazny describe a form of *existential type signatures*, supported in addition to standard *universal signatures* [12,13]. Type variables in a universal signature $f :: a \rightarrow a$ are interpreted in the same way as Haskell, i.e. as $f :: \forall a.a \rightarrow a$. However, in an existential type signature $f :: a \rightarrow a$ (note: three colons) the variables are interpreted in more or less like our named wildcards, so that the signature becomes equivalent to our $f :: _a \rightarrow _a$. A mixture of existential and universal annotations is not supported, but can be encoded by nesting an existential annotation in a universal one. Analogously to the *scoped type variables* extension, the variables in these signatures share the same scope.

Both *F_{ML}* [18] and *HMF* [19] combine the expressiveness of System F, including first-class polymorphism, with the convenience of Hindley-Milner type inference, while remaining a conservative extension of ML and HM respectively. Both solutions employ *partial type annotations* to avoid the *guessing* of poly-

morphic types during type inference. These partial type annotations are similar to the ones described in our discussions of the *scoped type variables* Haskell extension, but without the scoping aspect. Furthermore, they support partial type annotations of the following form: $e :: \exists \bar{\alpha}. \sigma$, where the free variables $\bar{\alpha}$ in σ are locally bound. The annotation should be read as “for some (monomorphic) types $\bar{\alpha}$, the expression e has type σ .” and the $\bar{\alpha}$ correspond to our named wildcards. The authors formalised these partial type annotations, including generalisation, in the context of a HM-based type system. Neither Rémy nor Leijen consider GADTs or local type assumptions.

9 Conclusion

Partial type signatures are a useful feature that has often been requested and proposed for Haskell. They bridge the gap between complete type annotations and none at all. Our proposal pins down the precise behaviour and we formally prove its well-behavedness. The result fits naturally in both the existing formal description of GHC’s type inferencer (OUTSIDEIN(X)) and the implementation. The idea of partial type signatures is not novel, but we believe our proposal is the first that supports all the features necessary for present-day GHC Haskell, esp. local constraint assumptions.

References

1. Pierce, B.: Types and programming languages. MIT Press (2002)
2. Peyton Jones, S., Vytiniotis, D., Weirich, S., Shields, M.: Practical type inference for arbitrary-rank types. *J. Funct. Program.* **17**(1) (2007) 1–82
3. Kiselyov, O.: Partial signatures (August 2004) Visited on 31/12/2012.
4. Claessen, K., Axelsson, E.: The patch-combinators package. Hackage (2012)
5. Various authors: Infer type context in a type signature. GHC Ticket (2011)
6. Vytiniotis, D., Peyton Jones, S., Schrijvers, T., Sulzmann, M.: OutsideIn(X): Modular type inference with local assumptions. *J. Funct. Program.* **21**(4-5) (2011) 333–412
7. Norell, U.: Towards a practical programming language based on dependent type theory. PhD thesis, Chalmers University and Göteborg University (2007)
8. Alice development team: Alice manual - extensions to the SML type language (March 2007) Visited on 30/12/2012.
9. Wallace, M.: Partial type signatures/annotations. Haskell Prime Wiki (Feb. 2006)
10. Wallace, M., et al.: Partial type signatures. Haskell Prime Wiki (Jan. 2006)
11. Dijkstra, A., Swierstra, D.S.: Making implicit parameters explicit. Technical Report UU-CS-2005-032, Universiteit Utrecht (2005)
12. Sulzmann, M., Wazny, J.: Lexically scoped type annotations. manuscript (2005)
13. Wazny, J.: Type inference and type error diagnosis for Hindley/Milner with extensions. PhD thesis, University of Melbourne (2006)
14. Peyton Jones, S., Shields, M.: Lexically-scoped type variables. (2004)
15. Bolingbroke, M.: Constraint kinds for GHC. online (2011)
16. Yorgey, B.A., Weirich, S., Cretin, J., Peyton Jones, S., Vytiniotis, D., Magalhães, J.P.: Giving haskell a promotion. In: TLDI. (2012) 53–66

17. Dijkstra, A.: Stepping Through Haskell. PhD thesis, Universiteit Utrecht (2005)
18. Rémy, D.: Simple, partial type-inference for System F based on type-containment. In: ICFP, ACM (2005) 130–143
19. Leijen, D.: HMF: simple type inference for first-class polymorphism. In: ICFP, ACM (2008) 283–294

A Modified constraint solver lemma

Because of the changed invariants of the constraint solver judgment $\mathcal{Q}; Q_{given}; \bar{a}_{tch} \vdash^{solv} C_{wanted} \rightsquigarrow Q_{residual}; \theta$ that we discussed in section 4.3, we need to modify one of the lemmas employed in `OUTSIDEIN(X)`.

Lemma 1. *Assume that $\Gamma \vdash e : \tau \rightsquigarrow C$. Then, for all C_{ext} , if $\mathcal{Q}; Q_g; \bar{\beta} \vdash^{solv} C \wedge C_{ext} \rightsquigarrow Q_r; \theta$ then there exists Q such that $Q; \theta \Gamma \vdash e : \theta\tau$ and $\mathcal{Q} \wedge \theta Q_g \wedge Q_r \vdash Q$.*

This lemma replaces Vytiniotis et al.’s lemma 5.1. The only modification is that in the conclusion, we apply the produced substitution θ to the given constraints Q_g as well, since given constraints can now also contain unification variables, produced from wildcards. This change follows from the changes made in section 4.3 and more background can be found there.

B Proof for Theorem 1

Proof. To prove the soundness of the `OUTSIDEIN(X)` algorithm, we need to prove that when a program is well-typed according to the algorithmic typing rules, it will also be well-typed according to the natural typing rules. The authors of `OUTSIDEIN(X)` deemed this proof straightforward [6, Theorem 5.1, p. 44], but we shall explicitly formulate the proof for the adapted `BINDPA` rule.

We need to prove from:

$$\Gamma \vdash e : v \rightsquigarrow C \quad (1)$$

$$\underline{Q}; \underline{\tau} \Rightarrow_a Q; \tau; extra \quad (2)$$

$$\mathcal{Q}; Q; fuv(v) \cup fuv(C) \cup fuv(\tau, Q) \vdash^{solv} C \wedge v \sim \tau \rightsquigarrow Q_{res}; \theta \quad (3)$$

$$extra \vee (Q_{res} = \epsilon) \quad (4)$$

$$\bar{b} \text{ fresh} \quad (5)$$

$$\bar{\beta} = fuv(\theta\tau) \cup fuv(\theta Q \wedge Q_{res}) \quad (6)$$

$$\mathcal{Q}; \Gamma, (f : \forall \bar{a} \bar{b}. [\bar{\beta} \mapsto \bar{b}] (\theta Q \wedge Q_{res} \Rightarrow \theta\tau)) \vdash prog \quad (7)$$

$$\mathcal{Q}; \Gamma \vdash f :: \forall \bar{a}. \underline{Q} \Rightarrow \underline{\tau} = e, prog \quad (8)$$

the following statements:

$$\underline{Q}; \underline{\tau} \Rightarrow Q; \tau \quad (9)$$

$$Q_1; \Gamma \vdash e : \tau \quad (10)$$

$$\bar{a} \cup \bar{b} = \text{ftv}(Q) \cup \text{ftv}(\tau) \quad (11)$$

$$\underline{Q} \wedge Q \Vdash Q_1 \quad (12)$$

$$\underline{Q}; \Gamma, (f : \forall \bar{a} \bar{b}. Q \Rightarrow \tau) \vdash \text{prog} \quad (13)$$

$$\underline{Q}; \Gamma \vdash f :: \forall \bar{a}. \underline{Q} \Rightarrow \underline{\tau} = e, \text{prog} \quad (14)$$

First, we appeal to Lemma 2 with $[\bar{\beta} \mapsto \bar{b}] \theta$ and Q_{res} , from (3) and (6) to attain (15) (note that $\theta Q_{\text{res}} = Q_{\text{res}}$, as $\text{dom}(\theta) \# \text{ftv}(Q_{\text{res}})$ [6, p. 20]). Lemma 2 imposes some conditions on the given Q_{res} and θ , namely $\text{dom}(\theta) \supseteq \text{ftv}(\tau) \cup \text{ftv}(Q)$ and $\text{extra} \vee (Q_{\text{res}} = \epsilon)$. The latter is satisfied by (4) and the former is satisfied because the solver will unify free unification variables from τ and Q , but not necessarily all free unification variables. The remaining free unification variables in $\theta\tau$ and θQ are handled by the additional substitution $[\bar{\beta} \mapsto \bar{b}]$.

$$\underline{Q}; \underline{\tau} \Rightarrow [\bar{\beta} \mapsto \bar{b}] \theta (Q \wedge Q_{\text{res}}); [\bar{\beta} \mapsto \bar{b}] \theta \tau \quad (15)$$

We then appeal to Lemma 1 as follows: $\Gamma \vdash e : v \rightsquigarrow C$ (1), then choose $C_{\text{ext}} = v \sim \tau$. Given (3), there exists Q_1 such that:

$$Q_1; \theta \Gamma \vdash e : \theta v \quad (16)$$

$$\underline{Q} \wedge \theta Q_g \wedge Q_\tau \vdash Q_1 \quad (17)$$

In the statements (9) to (14), choose for Q and τ respectively $[\bar{\beta} \mapsto \bar{b}] (\theta Q \wedge Q_{\text{res}})$ and $[\bar{\beta} \mapsto \bar{b}] \theta \tau$ from the statements (1) to (8). Most statements from (9) to (14) are then directly proved. We now go over the remaining statements. (10) is proved because $\theta \Gamma = \Gamma$ as Γ is a closed top-level environment. For statement (13) we rely on induction.

Lemma 2. *Assume that $\underline{Q}; \underline{\tau} \Rightarrow_a Q; \tau; \text{extra}$. Then, for all θ and Q_{res} , if $\text{dom}(\theta) \supseteq \text{ftv}(\tau) \cup \text{ftv}(Q)$, and $\text{extra} \vee (Q_{\text{res}} = \epsilon)$, then $\underline{Q}; \underline{\tau} \Rightarrow \theta Q \wedge Q_{\text{res}}; \theta \tau$.*

Proof. By induction on the derivation of $\underline{Q}; \underline{\tau} \Rightarrow_a Q; \tau; \text{extra}$. We consider cases corresponding to which rule was used.

- Case ANAMEDWC. We apply rule NAMEDWC. For proving $[\bar{a} \mapsto \bar{v}] \underline{\tau} \Rightarrow^t \theta \tau$ from $[\bar{a} \mapsto \bar{w}] \underline{\tau} \Rightarrow_a^t \tau$, we appeal to Lemma 3 as follows. We decorate the variable names corresponding with those from the lemma with a prime, the other variables were defined in this case. Let:

- $\underline{\tau} = [\bar{a} \mapsto \bar{w}] \underline{\tau}$.
- $\tau' = \tau$
- $\varphi' = [\bar{w} \mapsto \bar{v}]$
- $\theta' = \theta \setminus \text{dom}(\varphi')$.

Lemma 3 then gives us $\varphi' \underline{\tau}' \Rightarrow^t \theta' \varphi' \tau'$. $\varphi' \underline{\tau}' = [\bar{w} \mapsto \bar{v}] [\bar{a} \mapsto \bar{w}] \underline{\tau} = [\bar{a} \mapsto \bar{v}] \underline{\tau}$ and $\theta' \varphi' \tau' = \theta' [\bar{w} \mapsto \bar{v}] \tau = \theta \tau$, as desired. Analogously for $[\bar{a} \mapsto \bar{v}] \underline{Q} \Rightarrow^c \theta Q$, appealing to Lemma 4, with $Q_{\text{res}} = \epsilon$.

- Case AEXTRAWC. Apply rule EXTRAWC and use the induction hypothesis. Q_{res} need not be ϵ as $extra = true$.

Lemma 3. *Assume that $\underline{\tau} \Rightarrow_a^t \tau$. Then, for all θ and φ , if $dom(\theta) \supseteq fuv(\tau) \setminus fuv(\underline{\tau})$, $dom(\varphi) \supseteq fuv(\underline{\tau})$, and $dom(\theta) \cap dom(\varphi) = \emptyset$, then $\varphi \underline{\tau} \Rightarrow^t \theta \varphi \tau$.*

Proof. By induction on the derivation of $\underline{\tau} \Rightarrow_a^t \tau$. We consider cases corresponding to which rule was used. The substitution φ has as domain the unification variables that take the place of named wildcards and substitution θ has as domain the unification variables that take the place of non-named wildcards.

- Case ATYWC. We have $\underline{\tau} = _$, $\tau = \omega$, and want $\varphi \underline{\tau} \Rightarrow^t \theta \varphi \tau$. $\varphi \underline{\tau} = \underline{\tau}$, as $fuv(\underline{\tau}) = \emptyset$, and $\theta \varphi \tau = \theta \varphi \omega$ results in the monotype τ , because $fuv(\tau) = \{\omega\} \subseteq dom(\theta)$. Thus, we can apply rule TYWC: $\underline{\tau} \Rightarrow^t \tau$.
- Case ATYNOWC. We have that $\tau \Rightarrow_a^t \tau$. As τ cannot contain any (free) unification variables, $\theta \varphi \tau = \varphi \tau$, in which case we can apply rule TYNOWC: $\varphi \tau \Rightarrow^t \varphi \tau$.
- Case ATYAPP. We have that $T \underline{\tau}_i \Rightarrow_a^t T \bar{\tau}_i$. As substitution is distributive over the application of the type constructor T , we can use the induction hypothesis and apply rule TYAPP.

Lemma 4. *Assume that $Q^w \Rightarrow_a^c Q$. Then, for all θ and φ , if $dom(\theta) \supseteq fuv(Q) \setminus fuv(Q^w)$, $dom(\varphi) \supseteq fuv(Q^w)$, and $dom(\theta) \cap dom(\varphi) = \emptyset$, then $\varphi Q^w \Rightarrow^c \theta \varphi Q$.*

Proof. By induction on the derivation of $Q^w \Rightarrow_a^c Q$. We consider cases corresponding to which rule was used. The substitution φ has as domain the unification variables that take the place of named wildcards and substitution θ has as domain the unification variables that take the place of non-named wildcards.

- Case ACONNOWC. We have that $Q \Rightarrow_a^c Q$. As Q cannot contain any (free) unification variables, $\theta \varphi Q = \varphi Q$, in which case we can apply rule CONNOWC: $\varphi Q \Rightarrow^c \varphi Q$.
- Case ACONCONJ. We can use the induction hypothesis and apply rule CONCONJ.
- Case ACONTc. Substitution is distributive over the application of the type-class constraint D . We can apply rule CONTc using Lemma 3.
- Case ACONEQ. Substitution is distributive over an equality constraint. We can apply rule CONEQ using Lemma 3.

C Proof for Theorem 2

Proof. Induction on the size of $prog$, using Lemma 5 for definitions with type signatures without wildcards, and Lemma 6 for definitions without a type signature.

Lemma 5. *For a definition f with a type signature without wildcards, BINDPA (Figure 6) is equivalent with BINDA (Figure 5).*

Proof. We shall prove that when f has a type signature without wildcards, the premise as well as the conclusion of BINDPA are equivalent with those of BINDA. This implies that the rules are equivalent in the case of a type signature without wildcards.

The statements in the premise of BINDA are:

$$\Gamma \vdash e : v \rightsquigarrow C \quad (1)$$

$$\mathcal{Q}; Q; fuv(v) \cup fuv(C) \stackrel{solv}{\vdash} C \wedge v \sim \tau \rightsquigarrow \epsilon; \theta \quad (2)$$

$$\mathcal{Q}; \Gamma, (f : \forall \bar{a}. Q \Rightarrow \tau) \vdash prog \quad (3)$$

From these statements, the conclusion $\mathcal{Q}; \Gamma \vdash f :: \forall \bar{a}. Q \Rightarrow \tau = e, prog$ can be proved.

We will begin by proving the equivalence of BINDA's premise with BINDPA's premise.

BINDA and BINDPA will both infer the same v and C in the expression $\Gamma \vdash e : v \rightsquigarrow C$, as the judgment is applied with identical input parameters. Next, the wildcard desugaring judgment is applied:

$$Q; \tau \Rightarrow_a Q; \tau; \text{false} \quad (4)$$

Given that Q nor τ contain wildcards, the output parameters of the judgment will be the same as the input parameters, and *extra* will be false.

It is impossible for Q and τ to have free unification variables, as they do not contain wildcards. Thus:

$$fuv(\tau) \cup fuv(Q) = \emptyset \quad (5)$$

The call to the constraint solver happens with the same input parameters as the call in (2), resulting in the same output, except for the residual constraints, which are captured in Q_{res} instead of forced to be ϵ . However, as *extra* = false (4), $Q_{res} = \epsilon$ must be true, which makes the output parameters of both calls to the solver identical after all.

As there are no free unification variables in τ and Q (5), the substitution θ will affect Q (6) nor τ , (7).

$$\theta Q = Q \quad (6)$$

$$\theta \tau = \tau \quad (7)$$

$$\bar{\beta} = fuv(\theta \tau) \cup fuv(\theta Q \wedge Q_{res}) = fuv(\tau) \cup fuv(Q \wedge \epsilon) = \emptyset \quad (8)$$

Both $\bar{\beta}$ and \bar{b} will be empty, combined with $Q_{res} = \epsilon$, (6), and (7) results in the following type for f :

$$f : \forall \bar{a} \bar{b}. [\bar{\beta} \mapsto \bar{b}] (\theta Q \wedge Q_{res} \Rightarrow \theta \tau) \quad (9)$$

$$f : \forall \bar{a}. Q \Rightarrow \tau \quad (10)$$

This is the same type as the one in (3), thus, we have now proved the equivalence of the premises.

The conclusions of both rules are also equivalent as the difference between the rules lies in the fact that BINDPA's conclusion allows for wildcards, of which we have said that there are none. Thus, the \underline{Q} and $\underline{\tau}$ in BINDPA's conclusion

can simply be replaced with Q and τ , which makes it equivalent with BINDA's conclusion.

Lemma 6. *For a definition f without a type signature, ALTBIND (Figure 6) is equivalent with BIND (Figure 5).*

Proof. When f has no type signature, the rule ALTBIND applies, which requires a proof of its premise $\mathcal{Q}; \Gamma \vdash f :: _ \rightarrow _ = e, prog$. To prove this statement, we apply the BINDPA rule, of which the premise consists of the following statements:

$$\Gamma \vdash e : v \rightsquigarrow C \quad (1)$$

$$\mathcal{Q}; \underline{\tau} \Rightarrow_a \mathcal{Q}; \tau; extra \quad (2)$$

$$\mathcal{Q}; \mathcal{Q}; fuv(v) \cup fuv(C) \cup fuv(\tau) \cup fuv(Q) \vdash^{solv} C \wedge v \sim \tau \rightsquigarrow Q_{res}; \theta \quad (3)$$

$$extra \vee Q_{res} = \epsilon \quad (4)$$

$$\bar{b} \text{ fresh} \quad (5)$$

$$\bar{\beta} = fuv(\theta\tau) \cup fuv(\theta Q \wedge Q_{res}) \quad (6)$$

$$\mathcal{Q}; \Gamma, (f : \forall \bar{a}\bar{b}. [\bar{\beta} \mapsto \bar{b}] (\theta Q \wedge Q_{res} \Rightarrow \theta\tau)) \vdash prog \quad (7)$$

We shall now prove that for $\underline{Q} = _$ and $\underline{\tau} = _$ the premise of BIND is equivalent with the premise of BINDPA. By consequence, BIND will be equivalent with ALTBIND for the same values of \underline{Q} and $\underline{\tau}$.

The premise of BIND consists of the following statements:

$$\Gamma \vdash e : \tau \rightsquigarrow C \quad (8)$$

$$\mathcal{Q}; \epsilon; fuv(\tau) \cup fuv(C) \vdash^{solv} C \rightsquigarrow Q; \theta \quad (9)$$

$$\bar{a} \text{ fresh} \quad (10)$$

$$\bar{\alpha} = fuv(\theta\tau) \cup fuv(Q) \quad (11)$$

$$\mathcal{Q}; \Gamma, (f : \forall \bar{a}. [\bar{\alpha} \mapsto \bar{a}] (Q \Rightarrow \theta\tau)) \vdash prog \quad (12)$$

It is clear that (8) and (1) are equivalent. The rule (2) will be applied with $\underline{Q} = _$ and $\underline{\tau} = _$. The following rules will be applied by the wildcard desugaring judgment: AEXTRAWC ($\underline{Q} = _ = \epsilon \wedge _$), which results in $extra = true$, ANAMEDWC, which only applies ATYWC and ACONNOWC, as there are no named wildcards present in \underline{Q}^w and $\underline{\tau}$. ATYWC will replace the type wildcard with a fresh unification variable ω . ACONNOWC is applied because $\underline{Q}^w = \epsilon$. This results in the following expression:

$$\underline{Q}; \underline{\tau} \Rightarrow_a \epsilon; \omega; true \quad (13)$$

In the next step, the constraint solver is called (3). We shall prove that the call (3) happens with input parameters equivalent with those from the call in (9), which will result in equivalent output parameters. The first parameter, \mathcal{Q} , will be identical for both calls. The second parameter in (9) is ϵ , just like in (3), see (13). The third parameter differs, the $fuv(\tau) \cup fuv(C)$ from (9) will be identical to the $fuv(v) \cup fuv(C)$ from (3), but to the latter $fuv(\tau) \cup fuv(Q)$ will be added, namely $\{\omega\}$. The fourth parameters also differ slightly; in (3) there is an extra constraint: $v \sim \tau (= \omega)$. These two differing input parameters will only result

in an extra substitution, namely $\omega \mapsto v$. Thus, the substitution θ from (9) is equivalent to the one from (3). Rule (4) will not require $Q_{res} = \epsilon$ given that $extra = true$ (13).

We shall now prove that the $\bar{\alpha}$ (11) are identical to the $\bar{\beta}$ from (6). To avoid confusion between variables from (11) and (6), we shall suffix variables from (11) with an α -subscript. Given that $\tau = \omega$, $\theta = [\omega \mapsto v, \theta_\alpha]$, and $v = \tau_\alpha$, (14) is true. (15) is true because both Q_{res} and Q_α are output parameters of equivalent invocation of the constraints solver. From (14) and (15) follows (16).

$$\theta\tau = \theta\omega = [\omega \mapsto v, \theta_\alpha]\omega = \theta_\alpha v = \theta_\alpha \tau_\alpha \quad (14)$$

$$Q_{res} = Q_\alpha \quad (15)$$

$$\bar{\alpha} = fuv(\theta_\alpha \tau_\alpha) \cup fuv(Q_\alpha) = fuv(\theta\tau) \cup fuv(\epsilon \wedge Q_{res}) = \bar{\beta} \quad (16)$$

The final type in (7) will be the following:

$$f : \forall \bar{a} \bar{b}. [\bar{\beta} \mapsto \bar{b}] (\theta Q \wedge Q_{res} \Rightarrow \theta\tau) \quad (17)$$

$$f : \forall \bar{b}. [\bar{\beta} \mapsto \bar{b}] (\theta Q \wedge Q_{res} \Rightarrow \theta\tau) \quad (18)$$

$$f : \forall \bar{b}. [\bar{\beta} \mapsto \bar{b}] (Q_{res} \Rightarrow \theta\tau) \quad (19)$$

$$f : \forall \bar{b}. [\bar{\beta} \mapsto \bar{b}] (Q_\alpha \Rightarrow \theta\tau) \quad (20)$$

$$f : \forall \bar{b}. [\bar{\beta} \mapsto \bar{b}] (Q_\alpha \Rightarrow \theta_\alpha \tau_\alpha) \quad (21)$$

$$f : \forall \bar{a}. [\bar{\alpha} \mapsto \bar{a}] (Q_\alpha \Rightarrow \theta_\alpha \tau_\alpha) \quad (22)$$

In (18), the \bar{a} disappear because the annotated type $(_ \Rightarrow _)$ did not contain any type variables. In the next line (19), θQ disappears because it is empty (13). In (20) Q_{res} is replaced by Q_α because of (15). After this step, $\theta\tau$ is replaced by $\theta_\alpha \tau_\alpha$ because of (14). As $\bar{\alpha} = \bar{\beta}$ (16), we can replace $\bar{\beta}$ in (22) by $\bar{\alpha}$, where \bar{a} and \bar{b} are fresh type variables. This is the same type as in (12).

We have now proved that for a definition f without a type signature, ALT-BIND, which uses BINDPA, is equivalent with BIND.