

A Specification for Typed Template Haskell

ANONYMOUS AUTHOR(S)

Multi-stage programming is a proven technique that provides predictable performance characteristics by controlling code generation. We propose a core semantics for Typed Template Haskell, an extension of Haskell that supports multi staged programming that interacts well with polymorphism and qualified types. Our semantics relates a declarative source language with qualified types to a core language based on the the polymorphic lambda calculus augmented with multi-stage constructs.

Additional Key Words and Phrases: staging, polymorphism, constraints

1 INTRODUCTION

Producing optimal code is a difficult task that is greatly assisted by *staging annotations*, a technique which has been extensively studied and implemented in a variety of languages [Kiselyov 2014; Rompf and Odersky 2010; Taha and Sheard 2000]. These annotations give programmers fine control by instructing the compiler to generate code in one stage of compilation that can be used in another.

The classic example of staging is the *power* function, where the value n^k can be efficiently computed for a fixed k by generating code where the required multiplications have been unrolled and inlined. The incarnation of staged programming in Typed Template Haskell¹ benefits from type classes, one of the distinguishing features of Haskell [Hall et al. 1996; Peyton Jones et al. 1997], allowing a definition that can be reused for any type that is qualified to be numeric:

$$\begin{aligned} \text{power} &:: \text{Num } a \Rightarrow \text{Int} \rightarrow \text{Code } a \rightarrow \text{Code } a \\ \text{power } 0 \text{ } cn &= \llbracket 1 \rrbracket \\ \text{power } k \text{ } cn &= \llbracket \$(cn) * \$(\text{power } (k - 1) \text{ } cn) \rrbracket \end{aligned}$$

Any value $n::\text{Int}$ can be quoted to create $\llbracket n \rrbracket::\text{Code } \text{Int}$, then spliced in the expression $\$(\text{power } 5 \llbracket n \rrbracket)$ to generate $n * (n * (n * (n * (n * 1))))$. Thanks to type class polymorphism, this works when n has any fixed type that satisfies the *Num* interface, such as *Integer*, *Double* and countless other types.

It is somewhat surprising, then, that the following function fails to compile in the latest implementation of Typed Template Haskell in GHC 8.10.1:

$$\begin{aligned} \text{power5} &:: \text{Num } a \Rightarrow a \rightarrow a \\ \text{power5 } n &= \$(\text{power } 5 \llbracket n \rrbracket) \end{aligned}$$

Currently, GHC complains that there is no instance for *Num a* available, which is strange because the type signature explicitly states that *Num a* may be assumed. But this is not the only problem with this simple example: in the definition of *power*, the constraint is used inside a quotation but is bound outside. As we will see, this is an ad-hoc decision that then leads to subtle inconsistencies.

This paper sets out to formally answer the question of how a language with polymorphism and qualified types should interact with a multi-stage programming language. As well as the fact that all code generation happens at compile time, these are features that distinguish Typed Template Haskell, which until now has had no formalism even though it is an extension already fully integrated into GHC. We extend previous work that looked at cross-stage persistence in this

¹Typed Template Haskell is a variation of the Template Haskell [Sheard and Jones 2002] that adds types in the style of MetaML [Taha and Sheard 2000]. The abstract datatype *Code a* in this paper is a newtype wrapper around $Q (TExp a)$. In GHC, typed quotes are implemented by $\llbracket \mid \mid \rrbracket$ and typed splices by $\$(\$ (\))$, rather than $\llbracket \cdot \rrbracket$ and $\$(\cdot)$. It was implemented in 2013 by Geoffrey Mainland (who also used *power* as a motivating example) under a proposal by Simon Peyton Jones.

50 setting [Pickering et al. 2019a] by more closely modelling the reality of the type system implemented
 51 in GHC. In particular the implications of compile-time code generation are considered, as are the
 52 impredicativity restriction and the interaction with instance definitions.

53 Providing a formalism for the interaction of polymorphism, qualified types, and multi-stage
 54 programming has important consequences for Typed Template Haskell and other future imple-
 55 mentations that combine these features. This specification resolves the uncertainty about how it
 56 should interact with new language features. It has also been unclear how to fix a large number of
 57 bugs in the implementation because there is no precise semantics it is supposed to operate under.
 58 Presently these deficiencies hampered the adoption of meta-programming in Haskell as a trusted
 59 means of producing code with predictable performance characteristics. This untapped potential
 60 has been demonstrated to be beneficial in other staged systems in MetaML [Taha and Sheard 2000],
 61 Scala’s LMS [Rompf and Odersky 2010], and MetaOCaml [Kiselyov 2014].

62 At first glance it may appear that constructing programs using only quotations and splices
 63 is restrictive. However, multi-stage programming is widely applicable to any domain in which
 64 statically known abstractions should be eliminated. Some particular examples include removing
 65 the overhead of using parser combinators [Jonnalagedda et al. 2014; Krishnaswami and Yallop 2019;
 66 Willis et al. 2020], generic programming [Yallop 2017] and effect handlers [Schuster et al. 2020].
 67 Furthermore, there are certain interesting techniques and tricks which are useful when constructing
 68 multi-stage programs that can vastly improve performance [Kiselyov 2018; Taha 2004].

69 Following a brief introduction to staging with Typed Template Haskell (Section 2), this paper
 70 makes the following contributions:

- 71 • We demonstrate the problems that arise from the subtle interaction between qualified con-
 72 straints and staging, which highlights the need for a clear specification (Section 3).
- 73 • We introduce a type system for a source language which models Typed Template Haskell
 74 and introduces a new constraint form (Section 4).
- 75 • We introduce an explicitly typed core language which is System F extended with multi-stage
 76 features (Section 5).
- 77 • We give the elaboration semantics from the source language into the core language (Section 6).

78 Future directions including type inference, the latest developments in supporting impredicativity
 79 and notes about an implementation are then discussed (Section 7). This work is put into the context
 80 of related work (Section 8), before finally concluding (Section 9).

82 2 BACKGROUND: MULTI-STAGED PROGRAMMING

83 This section describes the fundamental concepts of Typed Template Haskell as currently im-
 84 plemented in GHC. In these aspects, the current implementation coincides with our proposed
 85 specification of Typed Template Haskell, as we are excluding polymorphism and type classes. This
 86 serves as preparation for more interesting examples that are discussed later (Section 3).

87 Typed Template Haskell is an extension to Haskell which implements the two standard staging
 88 annotations of multi-stage programming: *quotes* and *splices*. An expression $e :: a$ can be quoted
 89 to generate the expression $\llbracket e \rrbracket :: \text{Code } a$. Conversely, an expression $c :: \text{Code } a$ can be spliced to
 90 extract the expression $\$(c)$. An expression of type $\text{Code } a$ is a representation of an expression of
 91 type a . Given these definitions, it may seem that quotes and splices can be used freely so long as
 92 the types align: well-typed problems don’t go wrong, as the old adage says; but things are not so
 93 simple for staged programs. As well as being well-typed, a program must be well-staged, whereby
 94 the level of variables is considered.

95 The concept of a level is of fundamental importance in multi-stage programming. The *level* of
 96 an expression is an integer given by the number of quotes that surround it, minus the number of
 97

99 splices: quotation increases the level and splicing decreases the level. Negative levels are evaluated
 100 at compile time, level 0 is evaluated at runtime and positive levels are future unevaluated stages.²
 101 The goal of designing a staged calculus is to ensure that the levels of the program can be evaluated
 102 in order so that an expression at a particular level can only be evaluated when all expressions it
 103 depends on at previous levels have been first evaluated.

104 In the simplest setting, a program is *well-staged* if each variable it mentions is used only at the
 105 level in which it is bound. Doing so in any other stage may simply be impossible, or at least require
 106 special attention. The next three example programs, *timely*, *hasty*, and *tardy*, are all well-typed,
 107 but only the first is well-staged.

108 Using a variable that was defined in the *same* stage is permitted. A variable can be introduced
 109 locally using a lambda abstraction inside a quotation and then used freely within that context:

```
110     timely :: Code (Int → Int)
111     timely = [ λx → x ]
```

113 Of course, the variable is still subject to the usual scoping rules of abstraction.

114 Using a variable at a level *before* it is bound is problematic because at the point we wish to
 115 evaluate the prior stage, we will not yet know the value of the future stage variable and so the
 116 evaluation will get stuck:

```
117     hasty :: Code Int → Int
118     hasty c = $(c)
```

120 Here, we cannot splice *c* without knowing the concrete representation that *c* will be instantiated to.
 121 There is no recovery from this situation without violating the fact that lower levels must be fully
 122 evaluated before higher ones.

123 Using a variable at a stage *after* it is bound is problematic because the variable will not generally
 124 be in the environment. It may, for instance, be bound to a certain known value at compile time but
 125 no longer present at runtime.

```
127     tardy :: Int → Code Int
128     tardy x = [ x ]
```

129 In contrast to the *hasty* example, the situation is not hopeless here: there are ways to support
 130 referencing previous-stage variables, which is called *cross-stage persistence* [Taha and Sheard 1997].

131 One option is to interpret a variable from a previous stage using a *lifting* construct which copies
 132 the current value of the variable into a future-stage representation. As the process of lifting is akin
 133 to serialisation, this can be achieved quite easily for base types such as strings and integers, but
 134 more complex types such as functions are problematic.

135 Another option is to use *path-based persistence*: for example, a top-level identifier defined in
 136 another module can be persisted, because we can assume that the other module has been separately
 137 compiled, so the top-level identifier is still available at the same location in future stages.

138 GHC currently implements the restrictions described above. For cross-stage persistence, it
 139 employs both lifting and path-based persistence. Top-level variables defined in another module
 140 can be used at any level. Lifting is restricted to types that are instances of a type class *Lift* and
 141 witnessed by a method

```
143     lift :: Lift a ⇒ a → Code a
```

145 ²Unlike MetaML [Taha and Sheard 1997], code generation in Typed Template Haskell is only supported at compile time,
 146 and therefore there is no run operation.

148 which notably excludes function types and other abstract types such as *IO*. An example such as
 149 *tardy* can then be viewed as being implicitly rewritten to

```
150     tardy' :: Int → Code Int
151     tardy' x = [ $(lift x) ]
```

153 in which the reference of *x* occurs at the same level as it is bound. For this reason, our formal
 154 languages introduced in Section 4 and Section 5 will consider path-based persistence, but not
 155 implicit lifting, as this is easy to add separately in the same way as GHC currently implements it.

156 Let us consider the example from the introduction once again:

```
157     power 0 cn = [ 1 ]
158     power k cn = [ $(cn) * $(power (k - 1) cn) ]
```

160 This example makes use of quotes and splices. The references to the variables *power*, *k* and *cn* are
 161 all at the same level as their binding occurrence, as they occur within one splice and one quote.
 162 The static information is the exponent *k* and the run-time information is the base *cn*. Therefore
 163 by using the staged function the static information can be eliminated by partially evaluating the
 164 function at compile-time by using a top-level splice. The generated code does not mention the
 165 static information.

166 When using Typed Template Haskell, it is common to just deal with two stages, a compile-time
 167 stage and a run-time stage. The compile-time stage contains all the statically known information
 168 about the domain, the typing rules ensure that information from the run-time stage is not needed
 169 to evaluate the compile-time stage and then the program is partially evaluated at compile-time to
 170 evaluate the compile-time fragment to a value. It is a common misconception that Typed Template
 171 Haskell only supports two stages, but there is no such restriction.³

172 3 STAGING WITH TYPE CLASSES AND POLYMORPHISM

174 The examples in the previous section were simple demonstrations of the importance of considering
 175 levels in a well-staged program. This section discusses more complicated cases which involve class
 176 constraints, top-level splices, instance definitions and type variables. It is here where our proposed
 177 version of Typed Template Haskell deviates from GHC's current implementation. Therefore, for
 178 each example we will report how the program *should* behave and contrast it with how the current
 179 implementation in GHC behaves.

180 3.1 Constraints

182 Constraints introduced by type classes have the potential to cause staging errors: type classes
 183 are implemented by passing dictionaries as evidence, and the implicit use of these dictionaries
 184 must adhere to the level restrictions discussed in Section 2. This requires a careful treatment of the
 185 interaction between constraints and staging, which GHC currently does not handle correctly.

186 *Example C1.* Consider the use of a type class method inside a quotation, similarly to how *power*
 187 is used in *power5* in the introduction:

```
188     c1 :: Show a ⇒ Code (a → String)
189     c1 = [ show ]
```

191 Thinking carefully about the levels involved, the signature indicates that the evidence for *Show a*,
 192 which is available at stage 0, can be used to satisfy the evidence needed by *Show a* at stage 1.

194 ³There is, however, an artificial restriction about nesting quotation brackets which makes writing programs with more than
 195 two stages difficult. It is ongoing work to lift this restriction.

In the normal dictionary passing implementation of type classes, type class constraints are elaborated to a function which accepts a dictionary which acts as evidence for the constraint. Therefore we can assume that the elaborated version of c_1 looks similar to the following:

$$d_1 :: ShowDict a \rightarrow Code (a \rightarrow String)$$

$$d_1 dShow = \llbracket show dShow \rrbracket$$

Now this reveals a subtle problem: naively elaborating without considering the *levels of constraints* has introduced a cross-stage reference where the dictionary variable $dShow$ is introduced at level 0 but used at level 1. As we have learned in Section 2, one remedy of this situation is to try to make $dShow$ cross-stage persistent by lifting. However, in general, lifting of dictionaries is not straightforward to implement. Recall that lifting in GHC is restricted to instances of the *Lift* class which excludes functions – but type class dictionaries are nearly always a record of functions, so automatic lifting given the current implementation is not possible.

GHC nevertheless accepts this program, with the underlying problem only being revealed when subsequently trying to splice the program. We instead argue that the program c_1 is ill-typed and should therefore be rejected at compilation time. Our solution is to introduce a new constraint form, *CodeC C*, which indicates that constraint C is available to be used in the next stage. Using this, the corrected type signature for example C1 is as follows:

$$c'_1 :: CodeC (Show a) \Rightarrow Code (a \rightarrow String)$$

$$c'_1 = \llbracket show \rrbracket$$

The *CodeC (Show a)* constraint is introduced at level 0 but indicates that the *Show a* constraint will be available to be used at level 1. Therefore the *Show a* constraint can be used to satisfy the *show* method used inside the quotation. The corresponding elaborated version is similar to the following:

$$d'_1 :: Code (ShowDict a) \rightarrow Code (a \rightarrow String)$$

$$d'_1 cdShow = \llbracket show \$(cdShow) \rrbracket$$

As $cdShow$ is now the representation of a dictionary, we can splice the representation inside the quote. The reference to $cdShow$ is at the correct level and the program is well-staged.

Example C2. The example C1 uses a locally provided constraint which causes us some difficulty. We have a different situation if a constraint can be solved by a concrete global type class instance:

$$c_2 :: Code (Int \rightarrow String)$$

$$c_2 = \llbracket show \rrbracket$$

In c_2 , the global *Show Int* instance is used to satisfy the *show* constraint inside the quotation. Since this elaborates to a reference to a top-level instance dictionary, the reference can be persisted using path-based persistence. Therefore it is permitted to use top-level instances to satisfy constraints inside a quotation, in the same way that it is permitted to refer to top-level variables.

GHC currently correctly accepts this program. However, it does not actually perform the dictionary translation until the program is spliced, which is harmless in this case, because the same *Show Int* instance is in scope at both points.

3.2 Top-level splices

A splice that appears in a definition at the top-level scope of a module⁴ introduces new scoping challenges because it can potentially require class constraints to be used at levels prior to the ones where they are introduced.

⁴Do not confuse this use of “top-level” with the staging level.

Since a top-level splice is evaluated at compile time, it should be clear that no run-time information must be used in a top-level splice definition. In particular, the definition should not be permitted to access local variables or local constraints defined at a higher level. This is because local information which is available only at runtime cannot be used to influence the execution of the expression during compilation.

Example TS1. In the following example there are two modules. Module *A* defines the *Lift* class and contains the definition of a global instance *Lift Int*. The *lift* function from this instance is used in module *B* in the definition of ts_1 , which is a top-level definition. The reference to *lift* occurs inside a top-level splice, thus at level -1 , and so the *Lift Int* instance is needed a compile time.

<pre> 256 module A where 257 class Lift a where 258 lift :: a → Code a 259 instance Lift Int where 260 ... 261 ... 262 ... </pre>	<pre> 256 module B where 257 import A 258 ts₁ :: Int 259 ts₁ = \$(lift 5) </pre>
---	--

This program is currently accepted by GHC, and it should be, because the evidence for a top-level instance is defined in a top-level variable and top-level definitions defined in other modules are permitted to appear in top-level splices.

The situation is subtly different to top-level quotations as in example C2, because the instance must be defined in another module which mirrors the restriction implemented in GHC that top-level definitions can only be used in top-level splices if they are defined in other modules.

Example TS2. On the other hand, constraints introduced locally by a type signature for a top-level definition must not be allowed. In the following example, the *Lift A* constraint is introduced at level 0 by the type signature of ts_2 but used at level -1 in the body:

```

272     data A = A
273     ts2 :: Lift A ⇒ A
274     ts2 = $(lift A)

```

We assume *lift* is imported from another module as in Example TS1 and therefore cross-stage persistent. However, the problem is that the dictionary that will be used to implement the *Lift A* constraint will not be known until runtime. Therefore the definition of *Lift A* is not available to be used at compile-time in the top-level splice, and GHC correctly rejects this program. It is evident why this program should be rejected considering the dictionary translation, in a similar vein to C1. The elaboration would amount to a future-stage reference inside the splice.

Example TS3. So far we have considered the interaction between constraints and quotations separately to the interaction between constraints and top-level splices. The combination of the two reveals further subtle issues. The following function ts_3 is at the top-level and uses a constrained type. This example is both well-typed and well-staged.

```

287     ts3 :: Ord a ⇒ a → a → Ordering
288     ts3 = $([ compare ])

```

In ts_3 , the body of the top-level splice is a simple quotation of the *compare* method. This method requires an *Ord* constraint which is provided by the context on ts_3 . The constraint is introduced at level 0 and also used at level 0, as the top-level splice and the quotation cancel each other out. It is therefore perfectly fine to use the dictionary passed to ts_3 to satisfy the requirements of *compare*.

Unfortunately, the current implementation in GHC rejects this program. It generally excludes local constraints from the scope inside top-level splices, in order to reject programs like Example TS2. Our specification accepts the example by tracking the levels of local constraints.

3.3 The *power* function revisited

Having discussed constraints and also their interaction with top-level splices, we can now fully explain why *power5* function discussed in Section 1 works in current GHC when spliced at a concrete type, but fails when spliced with an overloaded type, whereas it should work for both:

```
power5 :: Int → Int      -- Option M
power5 :: Num a ⇒ a → a -- Option P
power5 n = $(power 5 [ n ])
```

The two problems are that GHC accepts *power* at the incorrect type

```
power :: Num a ⇒ Int → Code a → Code a
```

(as in Example C1) and that it does not actually perform a dictionary translation for this until this is spliced. Option M works in GHC, because it finds the *Num Int* global instance when splicing, similarly to Example TS1, and everything is (accidentally) fine. However, Option P fails because the local constraint is not made available in the splice as in Example TS2, and even if it was, it would be a reference at the wrong level.

As with Example C1, we argue that the function *power* should instead only be accepted at type

```
power :: CodeC (Num a) ⇒ Int → Code a → Code a
```

Then, Option M is fine due to the cross-stage persistence of the global *Num Int* instance declaration; and Option P works as well, as the program will elaborate to code that is similar to:

```
power5' :: NumDict a → a → a
power5' dNum n = $(power [ dNum ] 5 [ n ])
```

By quoting *dNum*, the argument to *power* is a representation of a dictionary as required, and reference is at the correct level.

3.4 Instance Definitions

The final challenge to do with constraints is dealing with instance definitions which use top-level splices. This situation is of particular interest as there are already special typing rules in GHC for instance methods which bring into scope the instance currently being defined in the body of the instance definition. This commonly happens in recursive datatypes where the instance declaration must be recursively defined.

Example I1. In the same manner as a top-level splice, the body of an instance method cannot use a local constraint, in particular the instance currently being defined or any of the instance head in order to influence the code generation in a top-level splice. On the other hand, the generated code should certainly be permitted to use the instances from the instance context and the currently defined instance. In fact, this is necessary in order to generate the majority of instances for recursive datatypes! Here is an instance definition that is defined recursively.

```
data Stream = Cons { hd :: Int, tl :: Stream }
instance Eq Stream where
  s1 == s2 = hd s1 == hd s2 && tl s1 == tl s2
```

344 The instance for *Eq Stream* compares the tails of the streams using the equality instance for
 345 *Eq Stream* that is currently being defined. This code works well, as it should, so there is reason to
 346 believe that a staged version should also be definable.

347 *Example I2.* The following is just a small variation of Example I1 where the body of the method
 348 is wrapped in a splice and a quote:
 349

350 **instance** *Eq Stream* **where**

351 $(==) = \$(\llbracket \lambda s_1 s_2 \rightarrow hd\ s_1 == hd\ s_2 \ \&\&\ tl\ s_1 == tl\ s_2 \rrbracket)$

352 This program should be accepted and equivalent to I1, because in general, $\$(\llbracket e \rrbracket) = e$. However, it
 353 is presently rejected by GHC, with the claim that the recursive use of $(==)$ for comparing the tails is
 354 not stage-correct, when in principle it could and should be made cross-stage persistent.
 355

356 *Example I3.* Yet another variant of Example I2 is to try to move the recursion out of the instance
 357 declaration, as follows:

358 $eqStream :: CodeC\ (Eq\ Stream) \Rightarrow Code\ (Stream \rightarrow Stream \rightarrow Bool)$

359 $eqStream = \llbracket \lambda s_1 s_2 \rightarrow hd\ s_1 == hd\ s_2 \ \&\&\ tl\ s_1 == tl\ s_2 \rrbracket$

360 Then, in a different module, we should be able to say:
 361

362 **instance** *Eq Stream* **where**

363 $(==) = \$(eqStream)$

364 It is important that the definition of *eqStream* uses the new constraint form *CodeC (Eq Stream)*
 365 so that the definition of *eqStream* is well-staged. As *Eq Stream* is the instance which is currently
 366 being defined, the *Eq Stream* constraint is introduced at level 0 when typing the instance definition.
 367 Therefore, in a top-level splice the local constraint can only be used by functions which require
 368 a *CodeC (Eq Stream)* constraint. This example is similar to *power* in Section 3.3 but the local
 369 constraint is introduced implicitly by the instance definition.
 370

371 It is somewhat ironic that while one of the major use cases of untyped Template Haskell is gen-
 372 erating instance definitions such as this one, it seems impossible to use the current implementation
 373 of Typed Template Haskell for the same purpose.

374 3.5 Type Variables

375 The previous examples showed that type constraints require special attention in order to ensure
 376 correct staging. It is therefore natural to consider type variables as well, and indeed previous work
 377 by Pickering et al. [2019a] ensured that type variables were level-aware in order to ensure a
 378 sound translation. However, our calculus will show that this is a conservative position in a language
 379 based on the polymorphic lambda calculus which enforces a phase distinction [Cardelli 1988] where
 380 typechecking happens entirely prior to running any stages of a program.
 381

382 *Example TV1.* Consider the following example that turns a list of quoted values into a quoted
 383 list of those values:

384 $list :: \forall a. [Code\ a] \rightarrow Code\ [a]$

385 $list [] = \llbracket [] @a \rrbracket$

386 $list (x : xs) = \llbracket (: @a \$ (x) \$ (list @a xs) \rrbracket$

387
 388 The type variable *a* is bound at level 0 and used both at level 1 (in its application to the type
 389 constructors), and at level 0 (in its recursive application to the *list* function). If we took the lead
 390 from values then this program would be rejected by the typechecker. However, because of the
 391 in-built phase distinction, evaluation cannot get stuck on an unknown type variable as all the types
 392

will be fixed in the program before execution starts. We just need to make sure that the substitution operation can substitute a type inside a quotation.

As it happens, this definition is correctly accepted by GHC, which is promising. However, it cannot actually be used since GHC produces compile error when given an expression such as $\$(list [\ [\ True \], [\ [\ False \]])$, since it complains that the type variable a is not in scope during type checking, and eventually generates an internal error in GHC itself.

3.6 Intermediate Summary

This section has uncovered the fact that the treatment of constraints is rather arbitrary in the current implementation of GHC. The examples we discussed have been motivated by three different ad-hoc restrictions that the current implementation exhibits. Firstly, prior stage constraint references (Section 3.1) are permitted without any checks. Secondly, top-level splices (Section 3.2) are typechecked in an environment isolated from the local constraint environment which is an overly conservative restriction. Thirdly, in an instance definition (Section 3.4), any reference to the instance currently being defined inside a top-level splice is rejected even if it is level-correct. Additionally we showed that although constraints have a role to play in understanding staging, type variables need no special treatment beyond the correct use of type application (Section 3.5).

Do the problems uncovered threaten the soundness of the language? In the case of prior-stage references, they do. The current representation form of quotations in GHC is untyped and therefore any evidence is erased from the internal representation of a quotation. The instance is therefore “persisted by inference”, which operates under the assumption that enough information is present in an untyped representation to re-infer all the contextual type information erased after typechecking. This assumption leads to unsoundness as generated programs will fail to typecheck, as discussed already by Pickering et al. [2019a]. Future-stage references are forbidden by conservative checks, which we will aim to make more precise in our calculus in order to accept more programs.

The problems observed so far are the result of interactions between class constraints and staging. Our goal now will be to develop a formal calculus to model and resolve these problems. The approach will be to introduce a source language that includes type constraints in addition to quotes and splices (Section 4), and a core language that is a variant of the explicit polymorphic lambda calculus with multi-stage constructs (Section 5). We then show how the source language can be translated into a core language where constraints have been elaborated away into a representation form for a quotation that is typed and where all contextual type information is saved (Section 6).

4 SOURCE LANGUAGE

The source language we introduce has been designed to incorporate the essential features of a language with metaprogramming and qualified types that is able to correctly handle the situations discussed in Section 3. We try to stay faithful to GHC’s current implementation of Typed Template Haskell where possible, with the addition of the quoted constraint form *CodeC* for the reasons discussed in Section 3. The key features of this language are:

- (1) Values *and constraints* are always indexed by the level at which they are introduced to ensure that they are well-staged.
- (2) The constraint form *CodeC* indicates that a constraint is available at the next stage.
- (3) Types (including type variables) are not level-indexed and can be used at any level.
- (4) There is no **run** operation in the language. All evaluation of code is performed at compile-time by mean of top-level splices which imply the existence of negative levels.
- (5) Path-based cross-stage persistence is supported.

```

442   $pgm ::= e \mid def; pgm \mid cls; pgm \mid inst; pgm$ 
443   $def ::= \mathbf{def} \ k = e$ 
444   $cls ::= \mathbf{class} \ TC \ a \ \mathbf{where} \ \{k :: \sigma\}$ 
445   $inst ::= \mathbf{instance} \ \overline{C} \Rightarrow TC \ \tau \ \mathbf{where} \ \{k = e\}$ 
446   $e ::=$  expressions
447   $\quad x \mid k$  variables / globals
448   $\quad \mid \lambda x : \tau. e \mid e \ e$  abstraction / application
449   $\quad \mid \llbracket e \rrbracket \mid \$(e)$  quotation / splice
450
451   $\tau ::=$  types
452   $\quad a \mid H$  variables / constants
453   $\quad \mid \tau \rightarrow \tau$  functions
454   $\quad \mid Code \ \tau$  representation
455   $\rho ::= \tau \mid C \Rightarrow \rho$  qualified types
456   $\sigma ::= \rho \mid \forall a. \sigma$  quantification
457   $C ::= TC \ \tau \mid CodeC \ C$  constraint / representation
458   $\Gamma ::= \bullet \mid \Gamma, x : (\tau, n) \mid \Gamma, a \mid \Gamma, (C, n)$  type environment
459   $P ::= \bullet \mid P, (\forall a. \overline{C} \Rightarrow C) \mid P, k : \sigma$  program environment
460

```

Fig. 1. Source Syntax

These features have been carefully chosen to allow a specification of Typed Template Haskell that addresses all the shortcomings identified in Section 3 while staying close to the general flavour of meta-programming that is enabled in its current implementation in GHC.

4.1 Syntax

The syntax of the source language (Figure 1) models programs with type classes and metaprogramming features. This syntax closely follows the models of languages with type classes of Bottu et al. [2017] and Chakravarty et al. [2005], but with the addition of multi-stage features and without equality or quantified constraints.

A program pgm is a sequence of value definitions def , class definitions cls , and instance definitions $inst$ followed by a top-level expression e .

Top-level definitions def are added to the calculus to model separate compilation and path-based cross-stage persistence in a way similar to what GHC currently implements: only variables previously defined in a top-level definition can be referenced at arbitrary levels.

Both type classes and instances are in the language, but simplified as far as possible: type class definitions cls have precisely one method and no superclasses; instance definitions $inst$ are permitted to have an instance context.

The expression language e is a standard λ -calculus with the addition of the two multi-stage constructs, quotation $\llbracket e \rrbracket$ and splicing $\$(e)$ which can be used to separate a program into stages.

Just as in Haskell, our language is *predicative*: type variables that appear as class parameters and in quantifiers range only over monotypes. This is modelled by the use of a Damas-Milner style type system which distinguishes between monotypes τ , qualified types ρ and polytypes σ . The type argument to the representation constructor $Code \ \tau$ must be a monotype.

A constraint C is either a type class constraint $TC \ \tau$, or a representation of a constraint $CodeC \ C$.

The environment Γ is used for locally introduced information, including value variables $x : (\tau, n)$, type variables a , and local type class axioms (C, n) . The environment keeps track of the (integer)

level n that value and constraint variables are introduced at; the typing rules ensure that the variables are only used at the current level.

The program theory P is an environment of the type class axioms introduced by instance declarations and of type information for names introduced by top-level definitions. The axioms are used to dictate whether the usage of a type class method is allowed or not. The names indicate which variables can be used in a cross-stage persistent fashion.

4.2 Typing Rules

The typing rules proceed in a predicable way for anyone familiar with qualified type systems. The difference is that the source expression typing judgement (Figure 2) and constraint entailment judgement (Figure 3) are now indexed by a level. Furthermore, since these rules are the basis for elaboration into the core language, they have been combined with the elaboration rules, as highlighted by a grey box whose contents can be safely ignored until the discussion of elaboration (Section 6).

The typing rules also refer to type and constraint formation rules (Figures 8 and 9), which check just that type variables are well-scoped (our language does not include an advanced kind system).

4.2.1 Levels. A large part of the rules are indexed by their (integer) level. It is standard to index the expression judgement at a specific level but less standard to index the constraint entailment judgement. The purpose of the level index is to ensure that expression variables and constraints can only be used at the level they are introduced. The result is a program which is separated into stages in such a way that when imbued with operational semantics, the fragments at stage n will be evaluated before the fragments at stage $n + 1$.

The full program is checked at level 0. Quotation (`E_QUOTE`) increases the level by one and splicing (`E_SPLICE`) decreases the level. It is permitted to reach negative levels in this source language. The negative levels indicate a stage which should evaluate at compile-time.

In more detail, the consequence of the level-indexing is that:

- A variable can be introduced at any specific level n by a λ -abstraction (`E_ABS`) and used at precisely that level (`E_VAR`).
- Variables introduced by top-level definitions can be used at any level (`E_VAR_TOPLEVEL`).
- Type variables can be introduced (`E_TABS`) and used at any level (`E_TAPP`).
- Constraints can be introduced at any level (`E_C_ABS`), but in practice, due to the system being predicative, this level will always be 0 (see Section 4.2.2). Constraints are also introduced into the program logic by instance declarations. Constraints are appropriately eliminated by application (`E_C_APP`). The entailment relation (Figure 3) ensures that a constraint is available at the current stage.

Note that we do not allow implicit lifting as discussed in Section 2 and implemented in GHC. Explicit lifting via a *Lift* class can easily be expressed in this language, and implicit lifting, if desired, can be orthogonally added as an additional elaboration step.

The rules furthermore implement the top-level splice restriction present in GHC. A top-level splice will require its body to be at level -1 and thereby rule out the use of any variables defined outside of the splice, with the exception of top-level variables.

4.2.2 Level of constraints. Inspecting the `E_C_ABS` rule independently it would appear that a local constraint can be introduced at any level and then be used at that level. In fact, the `E_C_ABS` rule can only ever be applied at level 0 because the `E_QUOTE` rule requires that the enclosed expression has a τ type. This forbids the ρ type as produced by `E_C_ABS` and more closely models the restriction to predicative types as present in GHC.

$$\begin{array}{c}
\boxed{P; \Gamma \vdash^n e : \sigma \rightsquigarrow t \mid TSP} \\
\frac{x : (\tau, n) \in \Gamma}{P; \Gamma \vdash^n x : \tau \rightsquigarrow x \mid \bullet} \text{E_VAR} \qquad \frac{k : \sigma \in P}{P; \Gamma \vdash^n k : \sigma \rightsquigarrow k \mid \bullet} \text{E_VAR_TOPLEVEL} \\
\frac{P; \Gamma, x : (\tau_1, n) \vdash^n e : \tau_2 \rightsquigarrow t \mid TSP \quad \Gamma \vdash_{\text{ty}} \tau_1 \rightsquigarrow \tau'_1}{P; \Gamma \vdash^n \lambda x : \tau_1. e : \tau_1 \rightarrow \tau_2 \rightsquigarrow \lambda x : \tau'_1. t \mid TSP} \text{E_ABS} \\
\frac{P; \Gamma \vdash^n e_1 : \tau_1 \rightarrow \tau_2 \rightsquigarrow t_1 \mid TSP_1 \quad P; \Gamma \vdash^n e_2 : \tau_1 \rightsquigarrow t_2 \mid TSP_2}{P; \Gamma \vdash^n e_1 e_2 : \tau_2 \rightsquigarrow t_1 t_2 \mid TSP_1 \cup TSP_2} \text{E_APP} \\
\frac{a \notin \Gamma \quad P; \Gamma, a \vdash^n e : \sigma \rightsquigarrow t \mid TSP}{P; \Gamma \vdash^n e : \forall a. \sigma \rightsquigarrow \Lambda a. t \mid TSP} \text{E_TABS} \qquad \frac{P; \Gamma \vdash^n e : \forall a. \sigma \rightsquigarrow t \mid TSP \quad \Gamma \vdash_{\text{ty}} \tau \rightsquigarrow \tau'}{P; \Gamma \vdash^n e : \sigma[\tau/a] \rightsquigarrow e \langle \tau' \rangle \mid TSP} \text{E_TAPP} \\
\frac{P; \Gamma \vdash^{n+1} e : \tau \rightsquigarrow t \mid TSP}{P; \Gamma \vdash^n \llbracket e \rrbracket : \text{Code } \tau \rightsquigarrow \llbracket t \rrbracket_{TSP_n} \mid \llbracket TSP \rrbracket^n} \text{E_QUOTE} \\
\frac{P; \Gamma \vdash^{n-1} e : \text{Code } \tau \rightsquigarrow t \mid TSP \quad \text{fresh } sp \quad \Gamma \vdash_{\text{ty}} \tau \rightsquigarrow \tau' \quad \Gamma \rightsquigarrow \Delta}{P; \Gamma \vdash^n \$(e) : \tau \rightsquigarrow sp \mid \{\Delta \vdash sp : \tau' = t\} \cup_{n-1} TSP} \text{E_SPLICE} \\
\frac{\Gamma \vdash_{\text{ct}} C \rightsquigarrow \tau \quad P; \Gamma, \text{ev} : (C, n) \vdash^n e : \rho \rightsquigarrow t \mid TSP \quad \text{fresh ev}}{P; \Gamma \vdash^n e : C \Rightarrow \rho \rightsquigarrow \lambda \text{ev} : \tau. t \mid TSP} \text{E_C_ABS} \\
\frac{P; \Gamma \vdash^n e : C \Rightarrow \rho \rightsquigarrow t_1 \mid TSP_1 \quad P; \Gamma \vdash^n C \rightsquigarrow t_2 \mid TSP_2}{P; \Gamma \vdash^n e : \rho \rightsquigarrow t_1 t_2 \mid TSP_1 \cup TSP_2} \text{E_C_APP}
\end{array}$$

Fig. 2. Source Expression Typing with Elaboration Semantics

The new constraint form *CodeC* C can be used to allow local constraints to be used at positive levels in E_C_APP . The rules C_INCR and C_DECR in constraint entailment can convert from *CodeC* C and C and vice versa. During elaboration, these rules will insert splices and quotes around the dictionary arguments as needed.

Constraints satisfied by instance declarations can be used at any level by means of C_GLOBAL .

Let us illustrate this by revisiting Example C1 and considering the expression $\llbracket \text{show} \rrbracket$. In this case, we assume the program environment to contain the type of *show*:

$$P = \bullet, \text{show} : \forall a. \text{Show } a \Rightarrow a \rightarrow \text{String}$$

From this we can use E_VAR_TOPLEVEL to conclude that

$$P; \Gamma \vdash^1 \text{show} : \forall a. \text{Show } a \Rightarrow a \rightarrow \text{String}$$

Because $\llbracket \text{show} \rrbracket$ must have type *Code* τ for some monotype τ , we cannot apply E_QUOTE yet, but must first apply E_TAPP and E_C_APP . We can conclude

$$\begin{array}{c}
\boxed{P; \Gamma \vDash^n C \rightsquigarrow t \mid TSP} \\
\hline
\frac{\text{ev}:(\forall a. \overline{C}_i \Rightarrow C) \in P \quad \Gamma \vdash_{\text{ty}} \tau \rightsquigarrow \tau' \quad \overline{P; \Gamma \vDash^n C_i[\tau/a] \rightsquigarrow t_i \mid TSP_i}}{P; \Gamma \vDash^n C[\tau/a] \rightsquigarrow \text{ev} \langle \tau' \rangle \overline{t}_i \mid \bigcup_i TSP_i} \text{C_GLOBAL} \\
\hline
\frac{\text{ev}:(C, n) \in \Gamma}{P; \Gamma \vDash^n C \rightsquigarrow \text{ev} \mid \bullet} \text{C_LOCAL} \quad \frac{P; \Gamma \vDash^{n+1} C \rightsquigarrow t \mid TSP}{P; \Gamma \vDash^n \text{CodeC } C \rightsquigarrow \llbracket t \rrbracket_{TSP_n} \mid \llbracket TSP \rrbracket^n} \text{C_DECR} \\
\hline
\frac{P; \Gamma \vDash^{n-1} \text{CodeC } C \rightsquigarrow t \mid TSP \quad \Gamma \vdash_{\text{ct}} C \rightsquigarrow \tau \quad \text{fresh } sp \quad \Gamma \rightsquigarrow \Delta}{P; \Gamma \vDash^n C \rightsquigarrow sp \mid \{\Delta \vdash sp: \tau = t\} \cup_{n-1} TSP} \text{C_INCR}
\end{array}$$

Fig. 3. Source Constraints with Elaboration

$P; \Gamma \vdash^1 \text{show} : a \rightarrow \text{String}$

if the entailment

$P; \Gamma \vDash^1 \text{Show } a$

holds. As we do not assume any instances for *Show* in this example, the constraint can only be justified via the local environment Γ . Because, as discussed above, any constraint can only have been introduced at level 0, the only way to move a local constraint from level 0 to level 1 is rule *C_INCR*, which requires us to show

$P; \Gamma \vDash^0 \text{CodeC } (\text{Show } a)$

This in turn follows from *C_LOCAL* if we assume that Γ contains *CodeC* (*Show* *a*):

$\Gamma = \bullet, (\text{CodeC } (\text{Show } a), 0)$

At this point, we know that

$P; \Gamma \vdash^0 \llbracket \text{show} \rrbracket : \text{Code } (a \rightarrow \text{String})$

and by applying *E_C_ABS* and *E_TABS*, we obtain

$P; \bullet \vdash^0 \llbracket \text{show} \rrbracket : \forall a. \text{CodeC } (\text{Show } a) \Rightarrow \text{Code } (a \rightarrow \text{String})$

as desired.

4.2.3 Program Typing. A program (Figure 7) is a sequence of value, class and instances declarations followed by an expression. The declaration forms extend the program theory which is used to typecheck subsequent definitions. Value definitions extend the list of top-level definitions available at all stages. The judgement makes it clear that the top-level of the program is level 0 and that each expression is checked in an empty local environment.

Class definitions extend the program theory with the qualified class method. Instance definitions extend the environment with an axiom for the specific instance which is being defined. The rule checks that the class method is of the type specified in the class definition.

Notice that before the type class method is checked, the local environment Γ is extended with an axiom for the instance we are currently checking. This is important to allow recursive definitions (Example 11) to be defined in a natural fashion. The constraint is introduced at level 0,

$$\begin{array}{c}
\boxed{pgm \mid P \vdash_{\text{def}} \text{def} : P \rightsquigarrow pgm} \\
P; \bullet \vdash^0 e : \sigma \rightsquigarrow t \mid TSP \quad \bullet \vdash_{\text{ty}} \sigma \rightsquigarrow \tau \quad d = \text{def } k : \tau = t \quad ds = \text{collapse}(-1, TSP, d; pgm) \\
\hline
pgm \mid P \vdash_{\text{def}} \text{def } k = e : P, k : \sigma \rightsquigarrow ds
\end{array}
\text{DEF}$$

Fig. 4. Source Definition Typing with Elaboration

$$\begin{array}{c}
\boxed{pgm \mid P \vdash_{\text{cls}} \text{cls} : P \rightsquigarrow pgm} \\
\bullet, a \vdash_{\text{ty}} \sigma \\
\hline
pgm \mid P \vdash_{\text{cls}} \text{class } TC \ a \ \text{where } \{k :: \sigma\} : P, k : \forall a. TC \ a \Rightarrow \sigma \rightsquigarrow pgm
\end{array}
\text{CLS}$$

Fig. 5. Source Class Typing with Elaboration

$$\begin{array}{c}
\boxed{pgm \mid P \vdash_{\text{inst}} \text{inst} : P \rightsquigarrow pgm} \\
\text{class } TC \ a \ \text{where } \{k :: \sigma\} \quad \overline{b_j} = \text{fv}(\tau) \quad \bullet \vdash_{\text{ty}} \sigma[\tau/a] \rightsquigarrow \tau' \quad \bullet, \overline{b_j} \vdash_{\text{ct}} C_i \rightsquigarrow \tau_i \\
P; \bullet, \overline{b_j}, \text{ev} : (TC \ \tau, 0), (\overline{C_i}, 0) \vdash^0 e : \sigma[\tau/a] \rightsquigarrow t \mid TSP \\
dl = \text{def } \text{ev} : (\forall \overline{b_j}. \overline{\tau_i} \rightarrow \tau') = t \quad ds = \text{collapse}(-1, TSP, dl; pgm) \quad \text{fresh ev} \\
\hline
pgm \mid P \vdash_{\text{inst}} \text{instance } \overline{C_i} \Rightarrow TC \ \tau \ \text{where } \{k = e\} : P, \text{ev} : (\forall \overline{b_j}. \overline{C_i} \Rightarrow TC \ \tau) \rightsquigarrow ds
\end{array}
\text{INST}$$

Fig. 6. Source Instance Typing with Elaboration

the top-level of the program. It is important to introduce the constraint to the local environment rather than program theory because this constraint should not be available at negative levels. If the constraint was introduced to the program theory, then it could be used at negative levels, which would amount to attempting to use the instance in order to affect the definition of said instance. This nuance in the typing rules explains why Example I2 is accepted and the necessity of the *CodeC* constraint in Example I3.

5 THE CORE LANGUAGE

In this section we describe an explicitly typed core language which is suitable as a compilation target for the declarative source language we described in the Section 4.

5.1 Syntax

The syntax for the core language is presented in Figure 10. It is a variant of the explicit polymorphic lambda calculus with multi-stage constructs, top-level definitions and top-level splice definitions.

Quotes and splices are represented by the syntactic form $\llbracket e \rrbracket_{SP}$, which is a quotation with an associated splice environment which binds *splice variables* for each splice point within the quoted expression. A splice point is where the result of evaluating a splice will be inserted. Splice variables to represent the splice points are bound by splice environments and top-level splice definitions. The expression syntax contains no splices, splices are modelled using the splice environments which are attached to quotations and top-level splice definitions.

687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735

$$\begin{array}{c}
 \boxed{P \vdash_{\text{pgm}} \text{pgm} : \sigma \rightsquigarrow \text{pgm}} \\
 \frac{P; \bullet \vdash^0 e : \sigma \rightsquigarrow t \mid TSP \quad \bullet \vdash_{\text{ty}} \sigma \rightsquigarrow \tau \quad p = t : \tau \quad ds = \text{collapse}(-1, TSP, p)}{P \vdash_{\text{pgm}} e : \sigma \rightsquigarrow ds} \text{P_EXPR} \\
 \\
 \frac{P_2 \vdash_{\text{pgm}} \text{pgm} : \sigma \rightsquigarrow \text{pgm} \quad \text{pgm} \mid P_1 \vdash_{\text{def}} \text{def} : P_2 \rightsquigarrow \text{pgm}_1}{P_1 \vdash_{\text{pgm}} \text{def}; \text{pgm} : \sigma \rightsquigarrow \text{pgm}_1} \text{P_DEF} \\
 \\
 \frac{P_2 \vdash_{\text{pgm}} \text{pgm} : \sigma \rightsquigarrow \text{pgm} \quad \text{pgm} \mid P_1 \vdash_{\text{cls}} \text{cls} : P_2 \rightsquigarrow \text{pgm}_2}{P_1 \vdash_{\text{pgm}} \text{cls}; \text{pgm} : \sigma \rightsquigarrow \text{pgm}_1} \text{P_CLS} \\
 \\
 \frac{P_2 \vdash_{\text{pgm}} \text{pgm} : \sigma \rightsquigarrow \text{pgm} \quad \text{pgm} \mid P_1 \vdash_{\text{inst}} \text{inst} : P_2 \rightsquigarrow \text{pgm}_1}{P_1 \vdash_{\text{pgm}} \text{inst}; \text{pgm} : \sigma \rightsquigarrow \text{pgm}_1} \text{P_INST}
 \end{array}$$

Fig. 7. Source Program Typing with Elaboration

$$\begin{array}{c}
 \boxed{\Gamma \vdash_{\text{ty}} \sigma \rightsquigarrow \tau} \\
 \\
 \frac{a \in \Gamma}{\Gamma \vdash_{\text{ty}} a \rightsquigarrow a} \text{T_VAR} \quad \frac{}{\Gamma \vdash_{\text{ty}} H \rightsquigarrow H} \text{T_CONST} \quad \frac{\Gamma \vdash_{\text{ty}} \tau_1 \rightsquigarrow \tau'_1 \quad \Gamma \vdash_{\text{ty}} \tau_2 \rightsquigarrow \tau'_2}{\Gamma \vdash_{\text{ty}} \tau_1 \rightarrow \tau_2 \rightsquigarrow \tau'_1 \rightarrow \tau'_2} \text{T_ARROW} \\
 \\
 \frac{\Gamma \vdash_{\text{ct}} C \rightsquigarrow \tau_1 \quad \Gamma \vdash_{\text{ty}} \rho \rightsquigarrow \tau_2}{\Gamma \vdash_{\text{ty}} C \Rightarrow \rho \rightsquigarrow \tau_1 \rightarrow \tau_2} \text{T_CARROW} \quad \frac{\Gamma, a \vdash_{\text{ty}} \sigma \rightsquigarrow \tau}{\Gamma \vdash_{\text{ty}} \forall a. \sigma \rightsquigarrow \forall a. \tau} \text{T_FOR_ALL} \\
 \\
 \frac{\Gamma \vdash_{\text{ty}} \tau \rightsquigarrow \tau'}{\Gamma \vdash_{\text{ty}} \text{Code } \tau \rightsquigarrow \text{Code } \tau'} \text{T_CODE}
 \end{array}$$

Fig. 8. Source Type Formation with Elaboration

$$\begin{array}{c}
 \boxed{\Gamma \vdash_{\text{ct}} C \rightsquigarrow \tau} \\
 \\
 \frac{\text{class } TC \text{ a where } \{k :: \sigma\} \quad \Gamma \vdash_{\text{ty}} \sigma[\tau/a] \rightsquigarrow \tau'}{\Gamma \vdash_{\text{ct}} TC \tau \rightsquigarrow \tau'} \text{C_TC} \quad \frac{\Gamma \vdash_{\text{ct}} C \rightsquigarrow \tau}{\Gamma \vdash_{\text{ct}} \text{Code } C \rightsquigarrow \text{Code } \tau} \text{C_CODEC}
 \end{array}$$

Fig. 9. Source Constraint Formation with Elaboration

The splice environment maps a splice point sp to a local type environment Δ , a type τ and an expression e which we write as $\Delta \vdash sp : \tau = e$. The typing rules will ensure that the expression e has type $\text{Code } \tau$. The purpose of the environment Δ is to support open code representations which lose their lexical scoping when lifted from the quotation.

A top-level splice definition is used to support elaborating from negative levels in the source language. These top level declarations are of the form $\text{spdef } \Delta \vdash^n sp : \tau = e$ which indicates that

the expression e will have type *Code* τ at level n in environment Δ . Top-level splice definitions also explicitly record the level of the original splice so that the declaration can be typechecked at the correct level. The level index is not necessary for the quotation splice environments because the level of the whole environment is fixed by the level at which the quotation appears.

5.2 Typing Rules

The expression typing rules for the core language are for the most part the same as those in the source language. The rules that differ are shown in Figure 11.

The splice environment typing rules are given in Figure 12. A splice environment is well-typed if each of its definitions is well-typed. We check that the body of each definition has type *Code* τ in an environment extended by Δ . In the E_QUOTE rule, the body of the quotation is checked in an environment Γ extended with the contents of the splice environment. A splice variable sp is associated with an environment Δ , a type τ and a level n in Γ . When a splice variable is used, the E_SPlice_VAR rule ensures that the claimed local environment aligns with the actual environment, the type of the variable is correct and the level matches the surrounding context.

Top-level splice definitions are typed in a similar manner in Figure 13. The body of the definition is checked to have type *Code* τ at level n in environment Γ , the program theory is then extended with a splice variable $sp : (\Gamma, \tau)$ which can be used in the remainder of the program.

5.3 Dynamic Semantics

The introduction of splice environments makes the evaluation order of the core calculus evident and hence a suitable target for compilation.

Usually in order to ensure a well-staged evaluation order, the reduction relation must be level-indexed to evaluate splices inside quotations. In our calculus, there is no need to do this because the splices have already been lifted outside of the quotation during the elaboration process. This style is less convenient to program with, but easy to reason about and implement.

In realistic implementations, the quotations are compiled to a representation form for which implementing substitution can be difficult. By lifting the splices outside of the representation, the representation does not need to be inspected or traversed before it is spliced back into the program. The representation can be treated in an opaque manner which gives us more implementation freedom about its form. In our calculus this is evidenced by the fact there is no reduction rule which reduces inside a quotation.

The program evaluation semantics evaluate each declaration in turn from top to bottom. Top-level definitions are evaluated to values and substituted into the rest of the program. Top-level splice definitions are evaluated to a value of type *Code* τ which has the form $\llbracket e \rrbracket_{sp}$. The splice variable is bound to the value e with the substitution SP applied and then substituted into the remainder of the program.

Using splice environments and top-level splice definitions is reminiscent of the approach taken in logically inspired languages by [Nanevski \[2002\]](#) and [Davies and Pfenning \[2001\]](#).

5.4 Module Restriction

In GHC, the module restriction dictates that only identifiers bound in other modules can be used inside top-level splices. This restriction is modelled in our calculus by the restriction that only identifiers previously bound in top-level definitions can be used inside a top-level splice. The intention is therefore to consider each top-level definition to be defined in its own “module” which is completely evaluated before moving onto the next definition.

This reflects the situation in a compiler such as GHC which supports separate compilation. When compiling a program which uses multiple modules, one module may contain a top-level splice

785	$pgm ::= e:\tau \mid def; pgm \mid spdef; pgm$	
786	$def ::= \mathbf{def} \ k:\tau = e$	
787	$spdef ::= \mathbf{spdef} \ \Delta \vdash^n \ sp:\tau = e$	
788		
789	$e, t ::=$	elaborated expressions
790	$x \mid sp \mid k$	variables / splice variables / globals
791	$\mid \lambda x:\tau. e \mid e e$	abstraction / application
792	$\mid \Lambda a. e \mid e \langle \tau \rangle$	type abstraction / application
793	$\mid \llbracket e \rrbracket_{sp}$	quotation
794	$SP ::= \bullet \mid SP, \Delta \vdash \ sp:\tau = e$	splice environment
795	$\tau ::=$	core types
796	$a \mid H$	variables / constants
797	$\mid \tau \rightarrow \tau$	functions
798	$\mid \mathit{Code} \ \tau$	representation
799	$\mid \forall a. \tau$	quantification
800	$\Gamma, \Delta ::= \bullet \mid \Gamma, x:(\tau, n) \mid \Gamma, sp:(\Delta, \tau, n) \mid \Gamma, a$	type environment
801	$P ::= \bullet \mid P, k:\tau \mid P, sp:(\Delta, \tau)$	program environment
802		

Fig. 10. Core Language Syntax

which is then used inside another top-level splice in a different module. Although syntactically both top-level splices occur at level -1 , they effectively occur at different stages. In the same way as different modules occur at different stages due to separate compilation, in our calculus, each top-level definition can be considered to be evaluated at a new stage.

At this point it is worthwhile to consider what exactly we mean by compile-time and run-time. So far we have stated that the intention is for splices at negative levels to represent compile-time evaluation so we should state how we intend this statement to be understood in our formalism. The elaboration procedure will elaborate each top-level definition to zero or more splice definitions (one for each top-level splice it contains) followed by a normal value definition. Then, during the evaluation of the core program the splice definitions will be evaluated prior to the top-level definition that originally contained them.

The meaning of “compile-time” is therefore that the evaluation of the top-level splice happens before the top-level definition is evaluated. It is also possible to imagine a semantics which partially evaluates a program to a residual by computing and removing as many splice definitions as possible.

If we were to more precisely model a module as a collection of definitions then the typing rules could be modified to only allow definitions from previously defined modules to be used at the top-level and all splice definitions could be grouped together at the start of a module definition before any of the value definitions. Then it would be clearly possible to evaluate all of the splice definitions for a module before commencing to evaluate the module definitions.

6 ELABORATION

In this section we describe the process of elaboration from the source language to the core language. Elaboration starts from a well-typed term in the source language and hence is defined by induction over the typing derivation tree (Figure 2).

6.1 Elaboration procedure

There are three key aspects of the elaboration procedure:

834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882

$$\begin{array}{c}
 \boxed{P; \Gamma \vdash^n e : \tau} \\
 \frac{\Delta \vdash sp : (\tau, n) \in \Gamma \quad \Delta \subseteq \Gamma}{P; \Gamma \vdash^n sp : \tau} \text{E_SPLICE_VAR} \quad \frac{sp : (\Delta, \tau) \in P \quad \Delta \subseteq \Gamma}{P; \Gamma \vdash^n sp : \tau} \text{E_TOP_SPLICE_VAR} \\
 \frac{P; \Gamma, a \vdash^n e : \tau}{P; \Gamma \vdash^n \Lambda a. e : \forall a. \tau} \text{E_TABS} \quad \frac{P; \Gamma \vdash^n e : \forall a. \tau_2}{P; \Gamma \vdash^n e \langle \tau_1 \rangle : \tau_2[\tau_1/a]} \text{E_TAPP} \\
 \frac{P; \Gamma, sp_i : (\Delta_i, \tau_i, n+1) \vdash^{n+1} e : \tau \quad P; \Gamma \vdash_{SP}^n SP \quad SP = \Delta_i \vdash sp_i : \tau_i = e_i}{P; \Gamma \vdash^n \llbracket e \rrbracket_{SP} : \text{Code } \tau} \text{E_QUOTE}
 \end{array}$$

Fig. 11. Core Expression Typing

$$\begin{array}{c}
 \boxed{P; \Gamma \vdash_{SP}^n SP} \\
 \frac{}{P; \Gamma \vdash_{SP}^n \bullet} \text{SP_EMPTY} \quad \frac{P; \Gamma \vdash_{SP}^n SP \quad P; \Gamma, \Delta \vdash^n e : \text{Code } \tau}{P; \Gamma \vdash_{SP}^n SP, \Delta \vdash sp : \tau = e} \text{SP_CONS}
 \end{array}$$

Fig. 12. Splice Environment Typing

$$\begin{array}{c}
 \boxed{P \vdash_{\text{def}} \text{def} : P} \quad \boxed{P \vdash_{\text{spdef}} \text{spdef} : P} \\
 \frac{P; \bullet \vdash^0 e : \tau}{P \vdash_{\text{def}} \text{def } k : \tau = e : P, k : \tau} \text{C_DEF} \quad \frac{P; \Gamma \vdash^n e : \text{Code } \tau}{P \vdash_{\text{spdef}} \text{spdef } \Gamma \vdash^n sp : \tau = e : P, sp : (\Gamma, \tau)} \text{C_SPDEF}
 \end{array}$$

Fig. 13. Core Definition Typing

$$\begin{array}{c}
 \boxed{P \vdash_{\text{pgm}} \text{pgm}} \\
 \frac{P; \bullet \vdash^0 e : \tau}{P \vdash_{\text{pgm}} e : \tau} \text{CP_MAIN} \quad \frac{P_1 \vdash_{\text{def}} \text{def} : P_2 \quad P_2 \vdash_{\text{pgm}} \text{pgm}}{P_1 \vdash_{\text{pgm}} \text{def}; \text{pgm}} \text{CP_DEF} \\
 \frac{P_1 \vdash_{\text{spdef}} \text{spdef} : P_2 \quad P_2 \vdash_{\text{pgm}} \text{pgm}}{P_1 \vdash_{\text{pgm}} \text{spdef}; \text{pgm}} \text{CP_SPDEF}
 \end{array}$$

Fig. 14. Core Program Typing

$$\begin{array}{c}
 \boxed{e \rightsquigarrow e} \\
 \frac{e_1 \rightsquigarrow e'_1}{e_1 e_2 \rightsquigarrow e'_1 e_2} \text{D_E_APP_L} \quad \frac{e_2 \rightsquigarrow e'_2}{e_1 e_2 \rightsquigarrow e_1 e'_2} \text{D_E_APP_R} \quad \frac{}{(\lambda x : \tau. e) v \rightsquigarrow e[v/x]} \text{D_E_BETA} \\
 \frac{e \rightsquigarrow e'}{e \langle \tau \rangle \rightsquigarrow e' \langle \tau \rangle} \text{D_E_TAPP} \quad \frac{}{(\Lambda a. e) \langle \tau \rangle \rightsquigarrow e[\tau/a]} \text{D_E_TBETA} \quad \frac{SP \rightsquigarrow SP'}{\llbracket e \rrbracket_{SP} \rightsquigarrow \llbracket e \rrbracket_{SP'}} \text{D_E_QUOTE}
 \end{array}$$

Fig. 15. Dynamic Semantics

883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931

$$\frac{\boxed{SP \rightsquigarrow SP'}}{SP, \Delta \vdash sp:\tau = e \rightsquigarrow SP', \Delta \vdash sp:\tau = e} \text{D_SP_BODY} \quad \frac{e \rightsquigarrow e'}{SP, \Delta \vdash sp:\tau = e \rightsquigarrow SP, \Delta \vdash sp:\tau = e'} \text{D_SP_CONS}$$

Fig. 16. Dynamic Semantics for Splice Environment

$$\frac{\boxed{def \rightsquigarrow def'}}{def\ k:\tau = e \rightsquigarrow def\ k:\tau = e'} \text{D_DEF} \quad \frac{\boxed{spdef \rightsquigarrow spdef'}}{spdef\ \Gamma \vdash^n sp:\tau = e \rightsquigarrow spdef\ \Gamma \vdash^n sp:\tau = e'} \text{D_SPDEF}$$

Fig. 17. Dynamic Semantics for Definitions

$$\frac{\boxed{pgm \rightsquigarrow pgm}}{def;\ pgm \rightsquigarrow def';\ pgm} \text{D_P_DEF} \quad \frac{\boxed{pgm \rightsquigarrow pgm}}{spdef;\ pgm \rightsquigarrow spdef';\ pgm} \text{D_P_SPDEF} \quad \frac{e \rightsquigarrow e'}{e:\tau \rightsquigarrow e':\tau} \text{D_P_MAIN}$$

$$\frac{}{def\ k:\tau = v;\ pgm \rightsquigarrow pgm[v/k]} \text{D_P_DEF_BETA}$$

$$\frac{}{spdef\ \Delta \vdash^n sp:\tau = \llbracket e \rrbracket_{SP};\ pgm \rightsquigarrow pgm[e[SP]/sp]} \text{D_P_SPDEF_BETA}$$

Fig. 18. Dynamic Semantics for Program

- (1) Splices at positive levels are removed in favour of a splice environment. The elaboration process returns a splice environment which is attached to the quotation form.
- (2) Splices at non-positive levels are elaborated to top-level splice definitions, which are bound prior to the expression within which they were originally contained.
- (3) Type class constraints are converted to explicit dictionary passing. We describe how to understand the new constraint form *CodeC* C in terms of quotation.

In order to support elaboration, all implicit evidence from the program theory and local environment are annotated with variables. The idea is that in the elaborated program, the evidence for each particular construct will be bound to a variable of that name. Top-level evidence by top-level variables and local evidence by λ -bound variables. The modifications to the typing rules for the source language syntax necessary in order to explain elaboration are highlighted by a grey box.

$$\Gamma ::= \bullet \mid \Gamma, x:(\tau, n) \mid \Gamma, a \mid \Gamma, \text{ev}:(C, n) \quad \text{type environment}$$

$$P ::= \bullet \mid P, \text{ev}:(\forall a. \overline{C} \Rightarrow C) \mid P, k:\sigma \quad \text{program environment}$$

The judgement $P; \Gamma \vdash^n e:\tau \rightsquigarrow t \mid TSP$ states that in program theory P and environment Γ , the expression e has type τ at level n and elaborates to term t whilst producing splices TSP .

6.2 Splice Elaboration

The TSP is a function from a level to a splice environment SP . In many rules, we perform level-pointwise union of produced splices, written \cup .

During elaboration, all splices are initially added to the TSP at the level of their contents, so a splice that occurs at level n is added at level $n-1$ by means of the operation \cup_{n-1} as in rule E_SPlice .

What happens with the splices contained in the TSP depends on their level. If they occur at a positive level, they will be bound by a surrounding quotation in rule E_QUOTE . The notation TSP_n denotes the projection of the splices contained in TSP at level n . They become part of the splice environment associated with the quotation. Via $\lfloor TSP \rfloor^n$, we then truncate TSP so that it is empty at level n and above.

If a splice occurs at non-positive level, it is a top-level splice and will become a top-level splice definition in rules DEF or $INST$, in such a way that the splice definitions are made prior to the value definition which yielded them. The splice definitions are created by the *collapse* judgement (Figure 19), which takes a splice environment returned by the expression elaboration judgement and creates top-level splice declarations for each negative splice which appeared inside a term. To guarantee a stage-correct execution, the splices are inserted in order of their level.

We maintain the invariant on the TSP where it only contains splices of levels prior to the current level of the judgement. Therefore, when the judgement level decreases as in E_QUOTE , the splices for that level are removed from the splice environment and bound at the quotation. As the top-level of the program is at level 0, the splice environment returned by an elaboration judgement at level 0 will only contain splices at negative levels, which is why the appeal to the *collapse* function starts from level -1 .

6.3 Constraint Elaboration

The second point of interest are the constraint elaboration rules given in Figure 3. Since in our language type classes have just a single method, we use the function corresponding to the method itself as evidence for a class instance in rule $INST$.

Constraints of the new constraint form $CodeC$ C are elaborated into values of type $Code$ τ . Therefore inspecting the entailment elaboration form C_DECR and C_INCR must be understood in terms of quotation. The C_DECR entailment rule is implemented by a simple quotation and thus similar to E_QUOTE . The C_INCR rule is conceptually implemented using a splice, but as the core language does not contain splices it is understood by adding a new definition to the splice environment, which mirrors E_SPlice . These rules explain the necessity of level-indexing constraints in the source language; the elaboration would not be well-staged if the stage discipline was not enforced.

The remainder of the elaboration semantics which elaborate the simple terms and constraint forms are fairly routine.

6.4 Example

As an example of elaboration, let us consider the expression $\$(c'_1)$ where $c'_1 = \llbracket show \rrbracket$ from Example C1. There are two points of interest here: there is a top-level splice which will be floated to a top-level splice definition, and the $CodeC$ $(Show\ a)$ constraint of c'_1 must be elaborated into quoted evidence using rule C_DECR .

We assume that the program environment contains c'_1 :

$$P = \bullet, c'_1 : \forall a. CodeC\ (Show\ a) \Rightarrow Code\ (a \rightarrow String)$$

As our program comprises a single main expression, we have to use rule P_EXPR . This rule requires us to first elaborate the expression itself at level 0 in an empty type environment. We obtain

$$P; \bullet \vdash^0 \$(c'_1) : \forall a. Show\ a \Rightarrow a \rightarrow String$$

$$\rightsquigarrow \Lambda a. \lambda ev. a \rightarrow String. sp \mid \{ \Delta \vdash sp : a \rightarrow String = c'_1 \langle a \rangle \llbracket ev \rrbracket \bullet \} \cup_{-1} \bullet$$

$$\begin{array}{c}
\boxed{\text{collapse}(n, TSP, \text{pgm}) = ds} \\
\frac{}{\text{collapse}(n, \bullet, \text{pgm}) = \text{pgm}} \text{C_EMPTY} \\
\frac{ds = \text{collapse}(n-1, [TSP]^n, \overline{\text{spdef}_i; \text{pgm}}) \quad \overline{\text{spdef}_i = \text{spdef } \Delta_i \vdash^n \text{sp}_i : \tau_i = e_i} \\
\quad \overline{TSP_n = \Delta_i \vdash \text{sp}_i : \tau_i = e_i}}{\text{collapse}(n, TSP, \text{pgm}) = ds} \text{C_STRIP}
\end{array}$$

Fig. 19. Definition of *collapse*

$$\begin{array}{c}
\boxed{\Gamma \rightsquigarrow \Delta} \\
\frac{}{\bullet \rightsquigarrow \bullet} \text{TE_EMPTY} \quad \frac{\Gamma \rightsquigarrow \Delta \quad \Gamma \vdash_{\text{ty}} \tau \rightsquigarrow \tau'}{\Gamma, x : (\tau, n) \rightsquigarrow \Delta, x : (\tau', n)} \text{TE_VAR} \quad \frac{\Gamma \rightsquigarrow \Delta}{\Gamma, a \rightsquigarrow \Delta, a} \text{TE_TYVAR} \\
\frac{\Gamma \rightsquigarrow \Delta \quad \Gamma \vdash_{\text{ct}} C \rightsquigarrow \tau}{\Gamma, \text{ev} : (C, n) \rightsquigarrow \Delta, \text{ev} : (\tau, n)} \text{TE_CTX}
\end{array}$$

Fig. 20. Elaboration of Type Environments

where

$$\Delta = \bullet, a, \text{ev} : (a \rightarrow \text{String}, 0)$$

The splice point *sp* has been introduced for the splice via rule E_SPLICE. Evidence for the *Show* constraint is introduced into the type environment at level 0 via rule E_C_ABS. It is captured in Δ for use in the splice. Because the use of the evidence occurs at level -1 and the required constraint is *CodeC Show*, rule C_DECR is used to quote the evidence. The splice environment attached to the quote is empty, because there are no further splices.

Back to rule P_EXPR, we furthermore obtain that the resulting main expression is of the form

$$p = \Lambda a. \lambda \text{ev} : a \rightarrow \text{String}. \text{sp} : \forall a. (a \rightarrow \text{String}) \rightarrow a \rightarrow \text{String}$$

The type of *p* results from elaborating the original *Show a* constraint to the type of its method $a \rightarrow \text{String}$ via the rule C_TC. We can now look at *collapse* and observe that it will extract the one splice at level -1 passed to it into a top-level splice definition that ends up before *d*, the result being

$$\text{spdef } \Delta \vdash^{-1} \text{sp} : a \rightarrow \text{String} = c'_1 \langle a \rangle \llbracket \text{ev} \rrbracket \bullet ; p$$

The operational semantics for the language evaluates the definition of *sp* first to a quotation before the quoted expression is substituted into the remainder of the program so that evaluation can continue.

7 PRAGMATIC CONSIDERATIONS

Now that the formal developments are complete and the relationship between the source language and core language has been established by the elaboration semantics, it is time to consider how our formalism interacts with other language extensions, and to discuss implementation issues.

7.1 Integration into GHC

The implementation and integration of this specification into GHC's sophisticated architecture naturally requires certain design decisions to be made, which we now discuss.

1030 7.1.1 *Type Inference*. Type inference for our new constraint form should be straightforward to
1031 integrate into the constraint solving algorithm used in GHC [Vytiniotis et al. 2011]. The key
1032 modification is to keep track of the level of constraints and only solve goals with evidence at the
1033 right level. If there is no evidence available for a constraint at the correct level, then either the
1034 C_INCR or C_DECR rule can be invoked in order to correct the level of necessary evidence.

1035
1036 7.1.2 *Substituting Types*. Substituting inside quotations poses some implementation challenges
1037 depending on the quotation representation. In previous implementations which used low-level
1038 representations to represent quotations [Pickering et al. 2019a; Roubinchtein 2015] the solution
1039 was to maintain a separate environment for free variables which can be substituted into without
1040 having to implement substitution in terms of the low-level representation. The idea is then that by
1041 the time the low-level representation is evaluated all the variables bound in the environment are
1042 already added to the environment and so computation can proceed as normal when it encounters
1043 what was previously a free variable.

1044 Therefore we should treat every free type variable in a quotation as a splice point, create an
1045 environment to attach to the quotation which maps the splice point to the variable of the name
1046 and then when the quotation is interpreted substitute the type into the quotation.

1047 7.1.3 *Erasing Types*. The operational semantics of System F do not depend on the type information
1048 and so the types can be erased before evaluation. It is this observation that leads us to accept
1049 Example TV1. However, there is another point in the design space – we could have elaborated to
1050 an erased version of System F where types were replaced by placeholders. This would save us the
1051 complications of having to substitute types inside quotations. However, the option of maintaining
1052 the type information is important for practical purposes. GHC has an optional internal typechecking
1053 phase called core lint which verifies the correctness of code generation and optimisation, it would
1054 be a shame to lose this pass in any program which used metaprogramming. In a language where
1055 the type information dwarfs the runtime content of terms, it would be desirable to also explore the
1056 option to store erased or partially erased terms [Brady et al. 2003; Jay and Peyton Jones 2008].
1057

1058 7.1.4 *Cross-Stage Persistence*. Earlier work on cross-stage persistence [Pickering et al. 2019a] has
1059 suggested that implementing cross-stage persistence for instance dictionaries should be possible
1060 because a dictionary was ultimately a collection of top-level functions so some special logic could
1061 be implemented in the compiler to lift a dictionary.

1062 Additional complexity arises when functions with local constraints are passed an arbitrary
1063 dictionary which could have been constructed from other dictionaries, which in turn come from
1064 dictionaries. At the point the dictionary is passed, the required information about its structure has
1065 been lost so it is impossible to interpret into a future stage. A solution to this could be to use a
1066 different evidence form which passes the derivation tree for a constraint to a function as evidence
1067 before the function constructs the required dictionary at the required stage. This would increase
1068 the runtime overhead of using type classes and would not be practical to implement.
1069

1070 7.2 Interaction with Existing Features

1071 So far we have considered how metaprogramming interacts with qualified types, but of course
1072 there are other features that are specific to GHC that need to be considered.
1073

1074 7.2.1 *GADTs*. Local constraints can be introduced by pattern matching on a GADT. For simplicity
1075 our calculus did not include GADTs or local constraints but they require similar treatment to other
1076 constraints introduced locally. The constraint solver needs to keep track of the level that a GADT
1077 pattern match introduces a constraint and ensure that the constraint is only used at that level.
1078

Note that this notion of a “level” is the stage of program execution where the constraint is introduced and not the same idea of a level the constraint solver uses to prevent existentially quantified type variables escaping their scope. Each nested implication constraint increases the level so type variables introduced in an inner scope are forbidden from unifying with type variables which are introduced at a previous level.

7.2.2 Quantified Constraints. The quantified constraints extension [Bottu et al. 2017] relaxes the form of the constraint schemes allowed in method contexts to also allow the quantified and implication forms which in our calculus are restricted to top-level axioms. Under this restriction there are some questions about how the *CodeC* constraint form should interact especially with implication constraints. In particular, whether constraint entailment should deduce that *CodeC* $(C_1 \Rightarrow C_2)$ entails *CodeC* $C_1 \Rightarrow \text{CodeC } C_2$ or the inverse and what consequences this has for type inference involving these more complicated constraint forms.

7.3 Interaction with Future Features

GHC is constantly being improved and extended to encompass new and ambitious features, and so we consider how metaprogramming should interact with features that are currently in the pipeline.

7.3.1 Impredicativity. For a number of the examples that we have discussed in this paper, an alternative would be to use impredicative instantiation to more precisely express the binding position of a constraint. For instance, the function $c_1 :: \text{Show } a \Rightarrow \text{Code } (a \rightarrow \text{String})$ from Example C1 might instead have been expressed in the following manner:

$$\begin{aligned} c'_1 &:: \text{Code } (\forall a. \text{Show } a \Rightarrow a \rightarrow \text{String}) \\ c'_1 &= \llbracket \text{show} \rrbracket \end{aligned}$$

The type of c'_1 now binds and uses the *Show a* constraint at level 1 without the use of *CodeC*.

However, despite many attempts [Peyton Jones et al. 2007], GHC has never properly supported impredicative instantiation due to complications with type inference. Recent work [Serrano et al. 2020, 2018] has proposed the inclusion of restricted impredicative instantiations, and these would accept c'_1 . In any case, impredicativity is not a silver bullet, and there is still a need for the *CodeC* constraint form. Without the *CodeC* constraint form there is no way to manipulate “open” constraints. That is, constraints which elaborate to quotations containing free variables.

Experience has taught us that writing code generators which manipulate open terms is a lot more convenient than working with only closed terms. We predict that the same will be true of working with the delayed constraint form as well. In particular, features such as super classes, instance contexts and type families can be used naturally with *CodeC* constraints. So whilst relaxing the impredicativity restriction will have some positive consequences to the users of Typed Template Haskell, it does not supersede our design but rather acts as a supplement.

7.3.2 Dependent Haskell. Our treatment of type variables is inspired by the in-built phase distinction of System F. As Haskell barrels at an impressive rate to a dependently typed language [Eisenberg 2016; Gundry 2013], the guarantees of the phase distinction will be lost in some cases. At this point it will be necessary to revise the specification in order to account for the richer phase structure.

The specification for Dependent Haskell [Weirich et al. 2017] introduces the so-called relevance quantifier in order to distinguish between relevant and irrelevant variables. The irrelevant quantifier is intended to model a form of parametric polymorphism like the \forall in System F, the relevant quantifier is written as Π after the dependent quantifier from dependent type theory.

Perhaps it is sound to modify their system in order to enforce the stage discipline for relevant type variables not irrelevant ones. It may be that the concept of relevance should be framed in

1128 terms of stages, where the irrelevant stage proceeds all relevant and computation stages – in which
1129 case it might be desirable to separate the irrelevant stage itself into multiple stages which can be
1130 evaluated in turn. The exact nature of this interaction is left as a question for future work.

1131 8 RELATED WORK

1133 Multi-stage languages with explicit staging annotations were first suggested by [Taha and Sheard](#)
1134 [1997, 2000]. Since then there has been a reasonable amount of theoretical interest in the topic
1135 which has renewed in recent years. There are several practical implementations of multi-stage
1136 constructs in mainstream programming languages such as BER MetaOCaml [[Kiselyov 2014](#)], Typed
1137 Template Haskell and Doty [[Stucki et al. 2018](#)].

1138 At a first glance there is surprisingly little work which attempts to reconcile multi-stage program-
1139 ming with language features which include polymorphism. Most presented multi-staged calculi are
1140 simply typed despite the fact that all the languages which practically implement these features
1141 support polymorphism. The closest formalism is by [Kokaji and Kameyama \[2011\]](#) who consider
1142 a language with polymorphism and control effects. Their calculus is presented without explicit
1143 type abstraction and application. There is no discussion about cross-stage type variable references
1144 or qualified types and their primary concern, similar to [Kiselyov \[2017\]](#) is the interaction of the
1145 value restriction and staging. Our calculus in contrast, as it models Haskell, does not contain any
1146 effects so we have concentrated on qualified types. [Calcagno et al. \[2003\]](#) present a similar ML-like
1147 language with let-generalisation.

1148 Combining together dependent types and multi-stage features is a more common combination.
1149 The phase distinction is lost in most dependently typed languages as the typechecking phase
1150 involves evaluating expressions. Therefore in order to ensure a staged evaluation, type variables
1151 must also obey the same stage discipline as value variables. This is the approach taken by [Kawata](#)
1152 [and Igarashi \[2019\]](#). [Pašalic \[2004\]](#) defines the dependently-typed multi-stage language Meta-D but
1153 doesn't consider constraints or parametric polymorphism. [Concoction \[Fogarty et al. 2007\]](#) is an
1154 extension to MetaOCaml where Coq terms appear in types. The language is based on λ_{HO} [[Pašalic](#)
1155 [et al. 2002](#)] which includes dependent types but is motivated by removing tags in the generated
1156 program. [Brady and Hammond \[2006\]](#) observe similarly that it is worthwhile to combine together
1157 dependent types and multi-stage programming to turn a well-typed interpreter into a verified
1158 compiler. The language presented does not consider parametric polymorphism nor constraints.

1159 We are not aware of any prior work which considers the implications of relevant implicit
1160 arguments formally, although there is an informal characterization by [Pickering et al. \[2019a\]](#).

1162 *Formalising Template Haskell.* With regards to formalisms of Template Haskell, a brief description
1163 of Untyped Template Haskell is given by [Sheard and Jones \[2002\]](#). The language is simply-typed
1164 and does not account for multiple levels. The language has also diverged since their formalism as
1165 untyped quotations are no longer typechecked before being converted into their representation.
1166 Their formalism does account for the Q monad which provides operations that allow a programmer
1167 to “reify” types, declarations and so on in order to inspect the internal structure. These are typically
1168 used in untyped Template Haskell in order to generically define instances or other operations.

1169 It is less common to use the reification functions in Typed Template Haskell programs and so we
1170 have avoided their inclusion in our formalism for the sake of simplicity. In our calculus we wanted
1171 to precisely understand the basic interaction between constraints and quotations, and the existence
1172 of the Q monad is orthogonal to this.

1173 Code generators are typically effectful in order to support operations such as let insertion or
1174 report errors so it is an important question to define a calculus with effects. From GHC 8.12, the
1175 type of quotations will be generalised [[Pickering 2019](#)] from $Q (TExp\ a)$ to a minimal interface
1176

1177 $\forall m. \text{Quote } m \Rightarrow m \text{ (TExp } a)$ so a user will have more control over which effects they are allowed to
1178 use in their code generators. We leave formalising this extension open to future work.

1179 *Metaprogramming In GHC.* GHC implements many different forms of metaprogramming from the
1180 principled to the ad-hoc. At the principled end of the spectrum in a similar vein as Typed Template
1181 Haskell there is Cloud Haskell [Epstein et al. 2011], which implements a modality for distributed
1182 computing. Another popular principled style is generic programming [Magalhães et al. 2010;
1183 Rodriguez et al. 2008] which allows the representation of datatypes to be inspected and interpreted
1184 at runtime. Untyped Template Haskell is used for untyped code generation in the combinator
1185 style. As well as generating expressions it can be used to generate patterns, declarations and types.
1186 Programs are typechecked after being generated rather than typechecking the generators as in
1187 Typed Template Haskell. Untyped Template Haskell has a limited interface into the typechecker but
1188 Source Plugins [Pickering et al. 2019b] allow unfettered access to the internal state and operations
1189 of the typechecker and other compiler phases.

1191 *Modal Type Systems.* Type systems motivated by modal logics have more commonly contemplated
1192 the interaction of modal operators and polymorphism. In particular attention has turned recently
1193 to investigating dependent modal type theories and the complex interaction of modal operators
1194 in such theories [Gratzer et al. 2020]. It seems probable that ideas from this line of research can
1195 give a formal account of the interaction of the code modality [Davies and Pfenning 2001] and the
1196 parametric quantification from System F which can also be regarded as a modality [Nuyts and
1197 Devriese 2018; Pfenning 2001].

1198 In recent times, Fitch-Style Modal calculi [Clouston 2018; Gratzer et al. 2019] have become a
1199 popular way of specifying a modal type system due to their good computational properties. It
1200 would be interesting future work to attempt to modify our core calculus to a Fitch-Style system
1201 which was not level indexed. In particular the calculus for Simple RaTT [Bahr et al. 2019] looks
1202 like a good starting point.

1203 In short, we are sure there is a lot to learn from the vast amount of literature on modal type
1204 systems but we are not experts in this field and the literature does not deal with our practical
1205 concerns regarding implementing and writing programs in staged programming languages. The
1206 goal of this paper was not to uncover a logically inspired programming language but to give a
1207 practical and understandable practical specification which can be understood by people without
1208 extensive background in modal type theories.

1209 9 CONCLUSION

1211 Now that we have presented the first formalism of Typed Template Haskell, the way is clear
1212 for future researchers to understand and extend the basic system. We envisage that the system
1213 will be useful for two different communities. Firstly, researchers into extensions such as Linear
1214 Haskell [Bernardy et al. 2017] and Dependent Haskell [Weirich et al. 2017] now have the possibility
1215 to consider how their features interact with stages so that the language remains sound with respect
1216 to staged evaluation. Secondly, users interested in multi-stage programming now have a firmer
1217 foundation to base further extensions to the multi-stage features in GHC.

1218 In the long-term our hope is to make multi-stage programming a more popular and accessible
1219 paradigm for functional programmers. Even with the present implementation of Typed Template
1220 Haskell, Yallop et al. [2018] have shown how different features implemented in GHC can be used
1221 together in order to express elegant code generators. We anticipate that a firm foundation for Typed
1222 Template Haskell that supports finer control over the type of qualified constraints will enable more
1223 practitioners to explore this unexplored territory, and begin to reach for staging as a useful tool in
1224 their toolbox of techniques to enhance the predictability and performance of their code.

1225

REFERENCES

- 1226
1227 Patrick Bahr, Christian Uldal Graulund, and Rasmus Ejlers Møgelberg. 2019. Simply RaTT: A Fitch-Style Modal Calculus
1228 for Reactive Programming without Space Leaks. *Proc. ACM Program. Lang.* 3, ICFP, Article 109 (July 2019), 27 pages.
1229 <https://doi.org/10.1145/3341713>
- 1230 Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. 2017. Linear
1231 Haskell: Practical Linearity in a Higher-Order Polymorphic Language. *Proc. ACM Program. Lang.* 2, POPL, Article 5 (Dec.
1232 2017), 29 pages. <https://doi.org/10.1145/3158093>
- 1233 Gert-Jan Bottu, Georgios Karachalias, Tom Schrijvers, Bruno C. d. S. Oliveira, and Philip Wadler. 2017. Quantified Class
1234 Constraints. In *Proceedings of the 10th ACM SIGPLAN International Symposium on Haskell* (Oxford, UK) (*Haskell 2017*).
1235 Association for Computing Machinery, New York, NY, USA, 148–161. <https://doi.org/10.1145/3122955.3122967>
- 1236 Edwin Brady and Kevin Hammond. 2006. A Verified Staged Interpreter is a Verified Compiler. In *Proceedings of the 5th
1237 International Conference on Generative Programming and Component Engineering* (Portland, Oregon, USA) (*GPCE '06*).
1238 Association for Computing Machinery, New York, NY, USA, 111–120. <https://doi.org/10.1145/1173706.1173724>
- 1239 Edwin Brady, Conor McBride, and James McKinna. 2003. Inductive families need not store their indices. In *International
1240 Workshop on Types for Proofs and Programs*. Springer, 115–129. https://doi.org/10.1007/978-3-540-24849-1_8
- 1241 Cristiano Calcagno, Walid Taha, Liwen Huang, and Xavier Leroy. 2003. Implementing multi-stage languages using ASTs,
1242 Gensym, and reflection. In *Proceedings of the 2nd international conference on Generative programming and component
1243 engineering* (Erfurt Germany) (*GPCE03*). Association for Computing Machinery, 57–76. [https://doi.org/10.5555/954186.
1244 954190](https://doi.org/10.5555/954186.954190)
- 1245 Luca Cardelli. 1988. Phase distinctions in type theory. Unpublish Manuscript.
- 1246 Manuel M. T. Chakravarty, Gabriele Keller, and Simon Peyton Jones. 2005. Associated Type Synonyms. *SIGPLAN Not.* 40, 9
1247 (Sept. 2005), 241–253. <https://doi.org/10.1145/1090189.1086397>
- 1248 Ranald Clouston. 2018. Fitch-style modal lambda calculi. In *International Conference on Foundations of Software Science and
1249 Computation Structures*. Springer, 258–275. https://doi.org/10.1007/978-3-319-89366-2_14
- 1250 Rowan Davies and Frank Pfenning. 2001. A Modal Analysis of Staged Computation. *J. ACM* 48, 3 (May 2001), 555–604.
1251 <https://doi.org/10.1145/382780.382785>
- 1252 Richard A Eisenberg. 2016. *Dependent types in Haskell: Theory and practice*. University of Pennsylvania.
- 1253 Jeff Epstein, Andrew P. Black, and Simon Peyton Jones. 2011. Towards Haskell in the Cloud. In *Proceedings of the 4th ACM
1254 Symposium on Haskell* (Tokyo, Japan) (*Haskell '11*). ACM, New York, NY, USA, 118–129. [https://doi.org/10.1145/2034675.
1255 2034690](https://doi.org/10.1145/2034675.2034690)
- 1256 Seth Fogarty, Emir Pasalic, Jeremy Siek, and Walid Taha. 2007. Concoqion: Indexed Types Now!. In *Proceedings of the 2007
1257 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation* (Nice, France) (*PEPM '07*).
1258 Association for Computing Machinery, New York, NY, USA, 112–121. <https://doi.org/10.1145/1244381.1244400>
- 1259 Daniel Gratzer, GA Kavvos, Andreas Nuyts, and Lars Birkedal. 2020. Multimodal Dependent Type Theory. (2020). In
1260 submission.
- 1261 Daniel Gratzer, Jonathan Sterling, and Lars Birkedal. 2019. Implementing a Modal Dependent Type Theory. *Proc. ACM
1262 Program. Lang.* 3, ICFP, Article 107 (July 2019), 29 pages. <https://doi.org/10.1145/3341711>
- 1263 Adam Gundry. 2013. *Type inference, Haskell and dependent types*. Ph.D. Dissertation. University of Strathclyde.
- 1264 Cordelia V. Hall, Kevin Hammond, Simon L. Peyton Jones, and Philip Wadler. 1996. Type Classes in Haskell. *ACM Trans.
1265 Program. Lang. Syst.* 18, 2 (1996), 109–138. <https://doi.org/10.1145/227699.227700>
- 1266 Barry Jay and Simon Peyton Jones. 2008. Scrap your type applications. In *International Conference on Mathematics of
1267 Program Construction*. Springer, 2–27. https://doi.org/10.1007/978-3-540-70594-9_2
- 1268 Manohar Jonnalagedda, Thierry Coppey, Sandro Stucki, Tiark Rompf, and Martin Odersky. 2014. Staged Parser Combinators
1269 for Efficient Data Processing. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming
1270 Systems Languages & Applications* (Portland, Oregon, USA) (*OOPSLA '14*). Association for Computing Machinery, New
1271 York, NY, USA, 637–653. <https://doi.org/10.1145/2660193.2660241>
- 1272 Akira Kawata and Atsushi Igarashi. 2019. A Dependently Typed Multi-stage Calculus. In *Asian Symposium on Programming
1273 Languages and Systems*. Springer, 53–72. https://doi.org/10.1007/978-3-030-34175-6_4
- 1274 Oleg Kiselyov. 2014. The Design and Implementation of BER MetaOCaml. In *Functional and Logic Programming*, Michael
1275 Codish and Eijiro Sumii (Eds.). Springer International Publishing, Cham, 86–102. [https://doi.org/10.1007/978-3-319-
1276 07151-0_6](https://doi.org/10.1007/978-3-319-07151-0_6)
- 1277 Oleg Kiselyov. 2017. Generating Code with Polymorphic let: A Ballad of Value Restriction, Copying and Sharing. *Electronic
1278 Proceedings in Theoretical Computer Science* 241 (Feb 2017), 1–22. <https://doi.org/10.4204/eptcs.241.1>
- 1279 Oleg Kiselyov. 2018. Reconciling Abstraction with High Performance: A MetaOCaml approach. *Foundations and Trends in
1280 Programming Languages* 5, 1 (2018), 1–101. <https://doi.org/10.1561/25000000038>
- 1281 Yuichiro Kokaji and Yuki Yoshi Kameyama. 2011. Polymorphic multi-stage language with control effects. In *Asian Symposium
1282 on Programming Languages and Systems*. Springer, 105–120. https://doi.org/10.1007/978-3-642-25318-8_11

- 1275 Neelakantan R. Krishnaswami and Jeremy Yallop. 2019. A Typed, Algebraic Approach to Parsing. In *Proceedings of the*
1276 *40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) (*PLDI 2019*).
1277 Association for Computing Machinery, New York, NY, USA, 379–393. <https://doi.org/10.1145/3314221.3314625>
- 1278 José Pedro Magalhães, Atze Dijkstra, Johan Jeuring, and Andres Löh. 2010. A Generic Deriving Mechanism for Haskell. In
1279 *Proceedings of the Third ACM Haskell Symposium on Haskell* (Baltimore, Maryland, USA) (*Haskell '10*). Association for
1280 Computing Machinery, New York, NY, USA, 37–48. <https://doi.org/10.1145/1863523.1863529>
- 1281 Aleksandar Nanevski. 2002. Meta-Programming with Names and Necessity. (2002), 206–217. <https://doi.org/10.1145/581478.581498>
- 1282 Andreas Nuyts and Dominique Devriese. 2018. Degrees of Relatedness: A Unified Framework for Parametricity, Irrelevance,
1283 Ad Hoc Polymorphism, Intersections, Unions and Algebra in Dependent Type Theory. In *Proceedings of the 33rd Annual*
1284 *ACM/IEEE Symposium on Logic in Computer Science* (Oxford, United Kingdom) (*LICS '18*). Association for Computing
1285 Machinery, New York, NY, USA, 779–788. <https://doi.org/10.1145/3209108.3209119>
- 1286 Emir Pašalic. 2004. *The role of type equality in meta-programming*. Ph.D. Dissertation. OGI School of Science & Engineering
1287 at OHSU.
- 1288 Emir Pašalic, Walid Taha, and Tim Sheard. 2002. Tagless Staged Interpreters for Typed Languages. In *Proceedings of the*
1289 *Seventh ACM SIGPLAN International Conference on Functional Programming* (Pittsburgh, PA, USA) (*ICFP '02*). Association
1290 for Computing Machinery, New York, NY, USA, 218–229. <https://doi.org/10.1145/581478.581499>
- 1291 Simon Peyton Jones, Mark Jones, and Erik Meijer. 1997. Type classes: an exploration of the design space. In *Haskell Workshop*.
1292 Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Mark Shields. 2007. Practical Type Inference for Arbitrary-
1293 Rank Types. *J. Funct. Program.* 17, 1 (Jan. 2007), 1–82. <https://doi.org/10.1017/S0956796806006034>
- 1294 Frank Pfenning. 2001. Intensionality, extensionality, and proof irrelevance in modal type theory. In *Proceedings 16th Annual*
1295 *IEEE Symposium on Logic in Computer Science*. IEEE, 221–230.
- 1296 Matthew Pickering. 2019. Overloaded Quotations. GHC proposal. [https://github.com/ghc-proposals/ghc-proposals/blob/
1297 master/proposals/0246-overloaded-bracket.rst](https://github.com/ghc-proposals/ghc-proposals/blob/master/proposals/0246-overloaded-bracket.rst)
- 1298 Matthew Pickering, Nicolas Wu, and Csongor Kiss. 2019a. Multi-Stage Programs in Context. In *Proceedings of the 12th ACM*
1299 *SIGPLAN International Symposium on Haskell* (Berlin, Germany) (*Haskell 2019*). Association for Computing Machinery,
1300 New York, NY, USA, 71–84. <https://doi.org/10.1145/3331545.3342597>
- 1301 Matthew Pickering, Nicolas Wu, and Boldizsár Németh. 2019b. Working with Source Plugins. In *Proceedings of the 2019*
1302 *ACM SIGPLAN Symposium on Haskell* (Berlin, Germany) (*Haskell '19*). ACM, New York, NY, USA. [https://doi.org/10.
1303 1145/3331545.3342599](https://doi.org/10.1145/3331545.3342599)
- 1304 Alexey Rodriguez, Johan Jeuring, Patrik Jansson, Alex Gerdes, Oleg Kiselyov, and Bruno C. d. S. Oliveira. 2008. Comparing
1305 Libraries for Generic Programming in Haskell. In *Proceedings of the First ACM SIGPLAN Symposium on Haskell* (Victoria,
1306 BC, Canada) (*Haskell '08*). Association for Computing Machinery, New York, NY, USA, 111–122. [https://doi.org/10.1145/
1307 1411286.1411301](https://doi.org/10.1145/1411286.1411301)
- 1308 Tiark Rompf and Martin Odersky. 2010. Lightweight Modular Staging: A Pragmatic Approach to Runtime Code Generation
1309 and Compiled DSLs. In *Proceedings of the Ninth International Conference on Generative Programming and Component*
1310 *Engineering* (Eindhoven, The Netherlands) (*GPCE '10*). ACM, New York, NY, USA, 127–136. [https://doi.org/10.1145/
1311 1868294.1868314](https://doi.org/10.1145/1868294.1868314)
- 1312 Evgeny Roubinchtein. 2015. *IR-MetaOCaml: (re)implementing MetaOCaml*. Master’s thesis. University of British Columbia.
1313 <https://doi.org/10.14288/1.0166800>
- 1314 Philipp Schuster, Jonathan Immanuel Brachthäuser, and Klaus Ostermann. 2020. Zero-cost Effect Handlers by Staging.
1315 (2020). In submission.
- 1316 Alejandro Serrano, Jurriaan Hage, Simon Peyton Jones, and Dimitrios Vytiniotis. 2020. A quick look at impredicativity.
1317 (January 2020). <https://www.microsoft.com/en-us/research/publication/a-quick-look-at-impredicativity/> In submission.
- 1318 Alejandro Serrano, Jurriaan Hage, Dimitrios Vytiniotis, and Simon Peyton Jones. 2018. Guarded Impredicative Polymorphism.
1319 In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Philadelphia,
1320 PA, USA) (*PLDI 2018*). ACM, New York, NY, USA, 783–796. <https://doi.org/10.1145/3192366.3192389>
- 1321 Tim Sheard and Simon Peyton Jones. 2002. Template Meta-programming for Haskell. In *Proceedings of the 2002 ACM*
1322 *SIGPLAN Workshop on Haskell* (Pittsburgh, Pennsylvania) (*Haskell '02*). ACM, New York, NY, USA, 1–16. [https://doi.org/
1323 10.1145/581690.581691](https://doi.org/10.1145/581690.581691)
- 1324 Nicolas Stucki, Aggelos Biboudis, and Martin Odersky. 2018. A Practical Unification of Multi-Stage Programming and
1325 Macros. In *Proceedings of the 17th ACM SIGPLAN International Conference on Generative Programming: Concepts and*
1326 *Experiences* (Boston, MA, USA) (*GPCE 2018*). Association for Computing Machinery, New York, NY, USA, 14–27. <https://doi.org/10.1145/3278122.3278139>
- 1327 Walid Taha. 2004. *A Gentle Introduction to Multi-stage Programming*. Springer Berlin Heidelberg, Berlin, Heidelberg, 30–50.
1328 https://doi.org/10.1007/978-3-540-25935-0_3

- 1324 Walid Taha and Tim Sheard. 1997. Multi-stage Programming with Explicit Annotations. In *Proceedings of the 1997 ACM*
1325 *SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation* (Amsterdam, The Netherlands)
1326 *(PEPM '97)*. ACM, New York, NY, USA, 203–217. <https://doi.org/10.1145/258993.259019>
- 1327 Walid Taha and Tim Sheard. 2000. MetaML and multi-stage programming with explicit annotations. *Theor. Comput. Sci.* 248,
1328 1-2 (2000), 211–242. [https://doi.org/10.1016/S0304-3975\(00\)00053-0](https://doi.org/10.1016/S0304-3975(00)00053-0)
- 1329 Dimitrios Vytiniotis, Simon Peyton Jones, Tom Schrijvers, and Martin Sulzmann. 2011. OutsideIn(X) Modular Type Inference
1330 with Local Assumptions. *J. Funct. Program.* 21, 4-5 (Sept. 2011), 333–412. <https://doi.org/10.1017/S0956796811000098>
- 1331 Stephanie Weirich, Antoine Voizard, Pedro Henrique Azevedo de Amorim, and Richard A. Eisenberg. 2017. A Specification
1332 for Dependent Types in Haskell. *Proc. ACM Program. Lang.* 1, ICFP, Article 31 (Aug. 2017), 29 pages. <https://doi.org/10.1145/3110275>
- 1333 Jamie Willis, Nicolas Wu, and Matthew Pickering. 2020. Staged Selective Parser Combinators. (2020). In submission.
- 1334 Jeremy Yallop. 2017. Staged Generic Programming. *Proc. ACM Program. Lang.* 1, ICFP, Article 29 (Aug. 2017), 29 pages.
1335 <https://doi.org/10.1145/3110273>
- 1336 Jeremy Yallop, Tamara von Glehn, and Ohad Kammar. 2018. Partially-Static Data as Free Extension of Algebras. *Proc. ACM*
1337 *Program. Lang.* 2, ICFP, Article 100 (July 2018), 30 pages. <https://doi.org/10.1145/3236795>
- 1338
- 1339
- 1340
- 1341
- 1342
- 1343
- 1344
- 1345
- 1346
- 1347
- 1348
- 1349
- 1350
- 1351
- 1352
- 1353
- 1354
- 1355
- 1356
- 1357
- 1358
- 1359
- 1360
- 1361
- 1362
- 1363
- 1364
- 1365
- 1366
- 1367
- 1368
- 1369
- 1370
- 1371
- 1372