Modelling Homogeneous Generative Meta-Programming*

Martin Berger[1], Laurence Tratt[2], and Christian Urban[3]

1   University of Sussex, Brighton, United Kingdom
2   King's College London, United Kingdom
3   King's College London, United Kingdom

Abstract

Homogeneous generative meta-programming (HGMP) enables the generation of program fragments at compile-time or run-time. We present a foundational calculus which can model both compile-time and run-time evaluated HGMP, allowing us to model, for the first time, languages such as Template Haskell. The calculus is designed such that it can be gradually enhanced with the features needed to model many of the advanced features of real languages. We demonstrate this by showing how a simple, staged type system as found in Template Haskell can be added to the calculus.

1   Introduction

Homogeneous generative meta-programming (HGMP) enables program fragments to be generated by a program as it is being either compiled or executed. Lisp was the first language to support HGMP, and for many years its only well known example. More recent languages such as MetaML [35, 36] and Template Haskell [33] support varying kinds of HGMP. Simplifying slightly, *homogeneous* systems are those where the program that creates the fragment is written in the same language as the fragment itself, in contrast to heterogeneous systems such as the C preprocessor where two different languages and/or systems are involved in the generation [32, 37]. Similarly, we use *generative* to distinguish HGMP from other forms of meta-programming such as reflection which focus on analysing (but, in general, not changing) a system.

Perhaps surprisingly, given its long history, HGMP's semantics have largely been defined by implementations [20]. Some aspects such as hygiene [1, 20] have been studied in detail. There has also been extensive work on compile-time type-checked, run-time evaluated HGMP, primarily in the context of MetaML and its descendants [5, 16, 22, 35, 36].

While we do not wish to advocate one style of HGMP over another, we are not aware of work which provides a natural formal basis for the style of HGMP found in languages such as Template Haskell and Converge [37] (broadly: compile-time evaluation of normal code with staged or dynamic type-checking). Our intention in this paper is to directly model, without encodings, a wider range of HGMP concepts than previously possible, in a simple

---

Language	Strings	ASTs	UpMLs	Compile-time HGMP	Run-time HGMP
Converge	●	●	●	●	●
JavaScript	●	○	○	○	●
Lisp	●	●	●	●	●
MetaML	○	○	●	○	●
Template Haskell	○	●	●	●	○
Scala (`scala.meta`)	○	●	●	●	●

Figure 1 A high-level characterisation of various HGMP languages.

yet expressive way, facilitating greater understanding of how these concepts relate to one another.

The system we construct is based on a simple untyped λ-calculus, to which we gradually add features and complexity, including a type system. This allows us to model, for the first time, HGMP which is evaluated at both compile-time (e.g. Template Haskell-ish) and run-time (e.g. MetaML-ish) HGMP. As a side benefit, this also gives clear pointers for how similar features can be added to 'real' programming languages.

To summarise, this paper's key contributions are:

- The first clear description of the design space of multiple languages and confusingly similar, yet distinct, meta-programming systems.
- The first calculus to be naturally to naturally model languages such as Template Haskell.
- The first calculus to be able to semi-systematically deal with syntactically rich languages.
- A demonstration of the calculus's simplicity by showing how it can be easily extended to model (monotyped) systems such as Template Haskell's.

2   HGMP design space

Although HGMP can seem an easy topic to discuss, in reality its various flavours and communities suffer greatly from incommensurability: important differences are ignored; and similarities are obscured by terminology or culture. Since we are not aware of previous work which tries to unify the various branches of the HGMP family, it is vital that we start by sketching the major points in the design space, so that we can be clear about both general concepts and the specific terminology we use.

Figure 1 summarises how some well known approaches sit within this classification. We use 'Lisp' as an over-arching term for a family of related languages (from Common Lisp to Scheme) and 'MetaML' to refer to MetaML and its descendants (e.g. MetaOCaml). Similarly, we use 'JavaScript' to represent what are, from this paper's perspective, similar languages (e.g. Python and Ruby).

2.1   The HGMP subset of meta-programming

The general area of meta-programming can be categorised in several different ways. In this paper we consider homogeneous, generative meta-programming.

We use Sheard's definition of homogeneous and heterogeneous systems: "homogeneous systems [are those] where the meta-language and the object language are the same, and heterogeneous systems [are those] where the meta-language is different from the object-language" [32]. The most well known example of a heterogeneous generative meta-programming system is C, where the preprocessor is both a separate system and language from C itself.

Heterogeneous systems are more flexible, but their power is difficult to tame and reason about [38].

In homogeneous systems in particular, we can then differentiate between generative and reflective. Reflection can introspect on run-time structures and behaviour (as in e.g. Smalltalk or Self [4]). In contrast, generative meta-programming explicitly constructs and executes program fragments.

2.2 Program fragment representation

An important, yet subtle, choice HGMP languages must make is how to represent the fragments that a program can generate. Three non-exclusive options are used in practice: strings; abstract syntax trees (ASTs); and upMLs (often called backquotes or quasi-quotes). We now define each of these, considering their suitability for our purposes.

2.2.1 Strings

In most cases, the simplest representation of a program is as a plain string. Bigger strings can be built from smaller strings and eventually evaluated. Evaluation typically occurs via a dedicated eval function at run-time; a handful of languages provide a compile-time equivalent, which allows arbitrary strings to be compiled into a source file. If such features are not available, then a string can simply be saved to a temporary file, compiled, and then run.

Representing program fragments as strings is trivial, terse, and can express any program valid in the language's concrete syntax.[1] However, strings can express nonsensical (e.g. syntactically invalid) programs and prevent certain properties (e.g. hygiene or certain notions of type-safety) from being enforced. Because of this, we believe that representing programs as strings is too fragile to serve as a sound basis for a foundational model.

2.2.2 ASTs

ASTs represent program fragments as a tree. For example, $2 + 3$ may be represented by the AST $\mathsf{ast_{add}}(\mathsf{ast_{int}}(2), \mathsf{ast_{int}}(3))$. ASTs are thus a simplification of a language's concrete syntax. Exactly how the concrete syntax should be simplified is influenced by an AST designer's personal tastes and preferences, and different languages – and, occasionally, different implementations of the same language – can take different approaches. In general, ASTs are designed to make post-parsing stages of a system easier to work with (e.g. type-checkers and code generators). HGMP languages which expose an AST datatype also enable users directly to instantiate new ASTs. Although ASTs generally allow semantically nonsensical programs to be created (e.g. referencing variables that are not defined), ASTs provide fewer opportunities for representing ill-formed programs than strings.

By definition, every valid piece of concrete syntax must have a valid AST representation (although not every valid AST may have a direct concrete syntax representation; this possibility is rarely exploited, but see e.g. [38] where it is used to help ensure lexical scoping for AST fragments generated in module M and inserted into M'). STs are therefore the most fundamental representation of programs and we use them as the basis of our calculus.

---

[1] For most languages this means that every possible program can be represented by strings. A few languages forbid concrete syntax representations of valid ASTs e.g. Converge prevents variables beginning with $ from being parsed, as part of its hygiene system.

2.3 UpMLs

UpML (Up MetaLevel) is the name for the concept traditionally called quasi-quote or back-quote[2], which allow AST (or AST-like) structures to be represented by quoted chunks of normal program syntax[3]. We have chosen the term 'upML' to highlight an important relationship with downMLs (see Section 3.2). UpMLs are often used because they enable a familiar means of representing code. They can also be used statically to guarantee various properties.

There are two distinct styles of upMLs in HGMP languages, which we now discuss. The most common style of upMLs is found in languages such as Lisp and Template Haskell, where they are used as a 'front-end' for creating ASTs. While ASTs in such languages are powerful, even small syntax fragments lead to deeply nested, unwieldy, ASTs. UpMLs allow AST fragments to be directly built from a concrete syntax fragment e.g. the upML expression

$$\uparrow\{2 + 3\}$$

evaluates to the AST

$$\mathsf{ast_{add}}(\mathsf{ast_{int}}(2), \mathsf{ast_{int}}(3)).$$

UpMLs can contain holes which are expressed using a downML $\downarrow\{\ldots\}$. The expression in the hole is expected to evaluate to an AST. For example, if the function f returns the AST equivalent of $2 + 3$ – in other words, f returns $\mathsf{ast_{add}}(\mathsf{ast_{int}}(2), \mathsf{ast_{int}}(3))$ – then

$$\uparrow\{\downarrow\{f\ ()\}\ * 4\}$$

will have an intermediate evaluation equivalent to $\uparrow\{(2 + 3) * 4\}$, leading to the eventual AST

$$\mathsf{ast_{mult}}(\mathsf{ast_{add}}(\mathsf{ast_{int}}(2), \mathsf{ast_{int}}(3)), \mathsf{ast_{int}}(4)).$$

In this model, ASTs are the fundamental construct and upMLs a convenience.

In their less common style – found only, to the best of our knowledge, in MetaML and its descendants – upMLs are a datatype in and of their own right, and do not represent ASTs. This has the shortcoming that it cannot represent some reasonable forms of meta-programming. For example, as discussed in [33], one can not use this form of upMLs to create projection functions such as:

$$(n, i) \quad \mapsto \quad \text{code of } \lambda(x_0, ..., x_{n-1}).x_i$$

In contrast, one can always use (perhaps laboriously) AST constructors to create such functions. A related problem relates to the tight coupling of upMLs and a language's concrete grammar which may not allow sub-constructs (e.g. an `else` clause) to be used in isolation and/or the location of holes can be ambiguous (if a hole follows an `if` construct, is it expected to be filled with an `else` clause, or a separate expression that follows, but is not part of, the `if`?).

---

[2] Quasi-quotes were developed by Quine for working with logics [40].
[3] Quotation is typically indicated by syntactic annotations such as brackets, but it is also possible to eschew explicit markers altogether and use types to distinguish between programs and code as data [23, 31].

Our experience is that languages without upMLs are prohibitively difficult to use, and find little traction with users. We therefore add upMLs as an optional extension to our calculus. We choose upMLs as a 'front-end' for creating ASTs, due to the greater expressivity and ubiquity of this approach.

2.4   Compile-time vs. run-time execution

In order to be useful, HGMP program fragments must at some point be run. Broadly speaking, execution happens at either compile-time (when the wider program is being compiled but not executed) or run-time (as part of normal program execution). Different languages allow evaluation at compile-time and/or run-time depending on the language: Template Haskell can evaluate ASTs only at compile-time; JavaScript can evaluate strings only at run-time; while Lisp can evaluate ASTs and strings at both compile-time and run-time.

Run-time evaluation is conceptually simple: a normal user-level function, conventionally called eval, takes in a program fragment (as a string or AST, depending on the language) and evaluates it. Every time the program is run, eval is called anew, as any other user-level function.

Compile-time evaluation is trickier and is represented in our approach with top-level downMLs $\downarrow\{\dots\}$ (i.e. a downML that is not nested inside an upML). When a top-level downML is encountered, the code inside it is evaluated before the surrounding program; that code must evaluate to an AST which then overwrites the downML before normal compilation resumes. In other words, all top-level downMLs are evaluated and 'eliminated' before run-time execution. Once the top-level downMLs are evaluated, a new program is constructed. No matter how many times the new program is evaluated, the top-level downMLs are not – cannot be! – reevaluated. We sometimes say that the effects of compile-time evaluation are 'frozen' in the resulting program.

The practical effects of run-time and compile-time evaluation are best seen by example. For example, the following program evaluates code at run-time using eval and prints 1 3 6 (where ';' is the sequencing operator):

```
print(1);
print(2 + eval(print(3); ASTInt(4)))
```

Replacing the eval with a downML leads to a program which prints 3 1 6:

```
print(1);
print(2 + ↓{print(3); ASTInt(4)})
```

These two possibilities have various implications. For example, compile-time evaluation allows ASTs to interact with early stages of the programming language's semantics and can introduce new variables into scope. In contrast, run-time evaluation can reference variables but not change those in scope in any way. There is also a significant performance difference: if a calculation can be moved from run-time to compile-time, it then has no run-time impact.

2.5   Implicit and explicit HGMP

Compile-time evaluation can also be subdivided into explicit and implicit flavours. In Lisp, 'macros' are special constructs explicitly identified at their point of definition by a user; in contrast, 'function calls' whose name references a macro are identified by the compiler and the macro evaluated at compile-time with the 'function call' arguments passed to it. Since one cannot tell by looking at a Lisp function call in isolation whether its arguments

will be evaluated at compile-time or run-time, it allows Lisp programmers to 'extend' the language transparently to the user. Although implicit evaluation has traditionally been seen as more problematic in syntactically rich languages, Honu [30] shows that a Lisp-like macro system can be embedded in such languages. However, languages such as Template Haskell take a different approach, explicitly identifying the locations where compile-time evaluation will happen using downMLs, allowing the user to call arbitrary user code. To an extent, the difference between implicit and explicit HGMP is cultural and without wishing to pick sides, in this paper we concentrate on explicit HGMP. This allows us to concentrate on the fundamentals of HGMP without the parsing considerations that are generally part of implicit HGMP.

2.6 HGMP vs. macro expansion

Systems such as Template Haskell share the same compile-time and run-time language (with the small exception that the run-time language does not feature upMLs): the compile-time evaluation of code uses the same evaluation rules as run-time evaluation. While some Lisp systems share this model, some do not. Most notably Scheme, and its descendent Racket, use a macro expander [13]. Rather than evaluate arbitrary Lisp code, Scheme and Racket macros share the same syntax as, but a different evaluation semantics to, their surrounding language. Thus, in our terminology, `define-syntax` is not a HGMP system. Modern Scheme and Racket systems have an additional macro system `syntax-case` which does allow HGMP.

3   A simple HGMP calculus

In this section, we define the minimal calculus which does interesting HGMP so that we can focus on the core features. In later sections we enrich this calculus with more advanced constructs. Our starting point is a standard call-by-value (CBV) λ-calculus whose grammar is as follows:

$$M \quad ::= \quad x \mid MN \mid \lambda x.M \mid c \mid M + N \mid \dots$$

Here, x ranges over variables and c over constants (e.g strings, integers). We include $+$ as an example of a wide class of common syntactic constructs.

3.1 ASTs

As discussed in Section 2.2.2, ASTs are the most fundamental form of representing programs in HGMP. In essence, every element of the calculus must have a representation as an AST. Of the syntactic constructs in the λ-calculus, constants, applications, and additions are most easily modelled; both variables and λ-abstractions require the representation of variables. We model variables as strings, which makes modelling later HGMP features easier and matches 'real' systems. We thus extend the λ-calculus as follows:

$$M \quad ::= \quad \dots \mid \mathsf{ast_t}(\tilde{M})$$

$$t \quad ::= \quad \mathsf{var} \mid \mathsf{app} \mid \mathsf{lam} \mid \mathsf{int} \mid \mathsf{string} \mid \mathsf{add} \mid \dots$$

We write $\tilde{M}$ for tuples $(M_1, ..., M_n)$ with $|\tilde{M}|$ denoting the length of the tuple. An AST constructor $\mathsf{ast_t}(\tilde{M})$ takes $|M| + 1$ arguments. The first argument t is a *tag* which specifies the specific AST datatype, and the rest of the arguments are then relative to that datatype. For example $\mathsf{ast_{var}}("x")$ is the AST representation of the variable x, $\mathsf{ast_{lam}}(\mathsf{ast_{string}}("x"), M)$ is the AST representation of $\lambda x.N$, and $\mathsf{ast_{int}}(3)$ is the AST representation of the constant 3.

$$\frac{}{x \Downarrow_{ct} x} \text{ Var ct} \qquad \frac{M \Downarrow_{ct} A \quad N \Downarrow_{ct} B}{MN \Downarrow_{ct} AB} \text{ App ct} \qquad \frac{M \Downarrow_{ct} A}{\lambda x.M \Downarrow_{ct} \lambda x.A} \text{ Lam ct} \qquad \frac{}{c \Downarrow_{ct} c} \text{ Const ct}$$

$$\frac{M \Downarrow_{ct} A \quad N \Downarrow_{ct} B}{M + N \Downarrow_{ct} A + B} \text{ Add ct} \qquad \frac{M_i \Downarrow_{ct} A_i}{\mathsf{ast_t}(\tilde{M}) \Downarrow_{ct} \mathsf{ast_t}(\tilde{A})} \text{ Ast}_c \text{ ct}$$

$$\frac{M \Downarrow_{ct} A \quad A \Downarrow_{\lambda} B \quad B \Downarrow_{dl} C}{\downarrow\{M\} \Downarrow_{ct} C} \text{ DownML ct}$$

$$\frac{}{\mathsf{ast_{var}}("x") \Downarrow_{dl} x} \text{ Var dl} \qquad \frac{M \Downarrow_{dl} M' \quad N \Downarrow_{dl} N'}{\mathsf{ast_{app}}(M, N) \Downarrow_{dl} M'N'} \text{ App dl}$$

$$\frac{M \Downarrow_{dl} "x" \quad N \Downarrow_{dl} N'}{\mathsf{ast_{lam}}(M, N) \Downarrow_{dl} \lambda x.N'} \text{ Lam dl} \qquad \frac{}{\mathsf{ast_{int}}(n) \Downarrow_{dl} n} \text{ Int dl}$$

$$\frac{}{\mathsf{ast_{string}}("x") \Downarrow_{dl} "x"} \text{ String dl} \qquad \frac{M \Downarrow_{dl} M' \quad N \Downarrow_{dl} N'}{\mathsf{ast_{add}}(M, N) \Downarrow_{dl} M' + N'} \text{ Add dl}$$

$$\frac{}{\lambda x.M \Downarrow_{\lambda} \lambda x.M} \text{ Lam} \qquad \frac{M \Downarrow_{\lambda} \lambda x.M' \quad N \Downarrow_{\lambda} N' \quad M'[N'/x] \Downarrow_{\lambda} L}{MN \Downarrow_{\lambda} L} \text{ App}$$

$$\frac{... \quad M_i \Downarrow_{\lambda} N_i \quad ...}{\mathsf{ast_t}(\tilde{M}) \Downarrow_{\lambda} \mathsf{ast_t}(\tilde{N})} \text{ Ast}_c$$

■ **Figure 2** Key big-step reduction rules for the CBV semantics of our simple calculus. Some standard rules (e.g. $\Downarrow_{\lambda}$ for addition) are omitted for brevity.

3.2 Compile-time HGMP

To model compile-time HGMP, we extend the calculus with a new construct downML, which provides a way of syntactically defining the points in a program where compile-time HGMP should occur (we do not need a tag for downMLs for reasons that will become clear later):

$$M \quad ::= \quad ... \mid \downarrow\{M\} \qquad t \quad ::= \quad ...$$

In essence, a downML $\downarrow\{M\}$ is an expression which must be evaluated at compile-time, i.e. before the rest of the program is executed. To model this, we find ourselves in the most complex and surprising part of our calculus: we have distinct but interacting reduction relations for the compile-time and run-time stages.

Figure 2 shows the reduction rules for our simple system. We use big-step semantics for brevity. There are three reduction relations:

$\Downarrow_{ct}$ models a compiler. It takes a program, possibly containing downMLs, as input and produces a program with no downMLs as output. It does this by recursively scanning through the input program looking for downMLs and evaluating them. Normal λ-calculus terms are copied from input to output unchanged. When a downML is encountered, the rule [DownML ct], explained below, evaluates the expression inside the downML. Assuming that expression returns an AST, the $\Downarrow_{dl}$ relation turns the AST into a normal program which then overwrites the downML. For example, $(\lambda z.z) \downarrow\{\mathsf{ast_{string}}((\lambda y.y)"x")\} \Downarrow_{ct} (\lambda z.z)"x"$.

$\Downarrow_{dl}$ models the conversion of ASTs into 'normal' programs. In our case, this means converting ASTs into programs (for example $\mathsf{ast}_{\mathsf{string}}(\text{"}x\text{"}) \Downarrow_{dl} \text{"}x\text{"}$). As this may suggest, $\Downarrow_{dl}$ is a simple relation which can be semi-mechanically created from the AST structure of the language.

$\Downarrow_{\lambda}$ models run-time execution. The rules are normal λ-calculus CBV reduction rules augmented with the minimum rules to evaluate ASTs.

Put another way, $\Downarrow_{ct}$ and $\Downarrow_{\lambda}$ are the key reduction relations, which allow us to accurately capture the reality that a program is compiled once but run many times:

$$\underbrace{M \Downarrow_{ct}}_{\text{compile-time}} A \underbrace{\Downarrow_{\lambda} V}_{\text{run-time}}$$

The key rule in $\Downarrow_{ct}$ is [DOWNML CT]. Its left-most premise $M \Downarrow_{ct} A$ first recursively scans for downMLs nested in M. The middle premise $A \Downarrow_{\lambda} B$ is the heart of the rule, evaluating the expression to produce an AST using normal λ-calculus evaluation. The simplicity of this premise belies its importance: the expression A can perform arbitrary computation. For the time being, we assume that the expression returns an AST (we defer consideration of erroneous computations to Section 6.2). That AST is then converted into a normal program which overwrites the downML. The program resulting from the $\Downarrow_{ct}$ relation can then be run as many times with the $\Downarrow_{\lambda}$ relation as desired. Figure 11 shows a fully worked-out example of a compile-time program and the $\Downarrow_{ct}$ relation.

3.2.1 Scoping

Our simple calculus intentionally allows variables to be captured dynamically. Although this naturally follows from the reduction rules, we now explicitly explain how this occurs and why. First, we note that the [APP] rule in the definition of $\Downarrow_{\lambda}$ uses traditional capture-avoiding substitution $M[N/x]$. Note that we do not need to extend the definition to downMLs, which will have been removed by $\Downarrow_{ct}$ before substitution is applied.

We can create AST representations of variables containing arbitrary variables (for example $\mathsf{ast}_{\mathsf{var}}(\text{"}x\text{"})$) which downMLs will then turn into normal programs. Consider the following two programs and their compilation:

- $\lambda x. \downarrow\{\mathsf{ast}_{\mathsf{var}}(\text{"}x\text{"})\} \Downarrow_{ct} \lambda x.x.$
- $\lambda y. \downarrow\{\mathsf{ast}_{\mathsf{var}}(\text{"}x\text{"})\} \Downarrow_{ct} \lambda y.x.$

As these examples suggest, our calculus is not hygienic, and thus allows variables to be captured. This is a deliberate design decision for two reasons. First, not all languages that we wish to model have an explicit notion of hygiene, instead providing a function which generates fresh (i.e. unique) names (conventionally called gensym). Second, there is not, as yet, a single foundational style of hygiene, and different languages take subtly different approaches. We talk about possible future directions for hygiene in Section 8.

3.3 Run-time HGMP

Having introduced compile-time HGMP, we now have all the basic tools needed to introduce run-time HGMP. We follow the Lisp tradition and use a function called eval. Unlike downMLs, evals are not eliminated at compile-time: they are, in essence, normal λ-calculus functions. We extend the calculus grammar (including an AST equivalent) as follows:

$$M ::= ... \mid \mathsf{eval}(M) \qquad\qquad t ::= ... \mid \mathsf{eval}$$

$$\frac{L \Downarrow_\lambda M \quad M \Downarrow_{dl} N \quad N \Downarrow_\lambda N'}{\mathsf{eval}(L) \Downarrow_\lambda N'} \text{ Eval rt} \qquad \frac{M \Downarrow_{ct} N}{\mathsf{eval}(M) \Downarrow_{ct} \mathsf{eval}(N)} \text{ Eval ct}$$

$$\frac{M \Downarrow_{dl} N}{\mathsf{ast_{eval}}(M) \Downarrow_{dl} \mathsf{eval}(N)} \text{ Eval dl}$$

Figure 3 Additional reduction rules for eval.

$$\frac{\ldots \quad M_i \Downarrow_{ct} A_i \quad \ldots}{\mathsf{ast_{promote}}(\mathsf{tag_t}, \tilde{M}) \Downarrow_{ct} \mathsf{ast_{promote}}(\mathsf{tag_t}, \tilde{A})} \text{ Promote ct}$$

$$\frac{L \Downarrow_{dl} \mathsf{tag_t} \quad t \neq \mathsf{promote} \quad \ldots \quad M_i \Downarrow_{dl} N_i \quad \ldots}{\mathsf{ast_{promote}}(L, \tilde{M}) \Downarrow_{dl} \mathsf{ast_t}(\tilde{N})} \text{ Promote dl 1}$$

$$\frac{L \Downarrow_{dl} \mathsf{tag_{promote}} \quad M \Downarrow_{dl} \mathsf{tag_t} \quad \ldots \quad N_i \Downarrow_{dl} N_i' \quad \ldots}{\mathsf{ast_{promote}}(L, M, \tilde{N}) \Downarrow_{dl} \mathsf{ast_{promote}}(\mathsf{tag_t}, \tilde{N'})} \text{ Promote dl 2}$$

$$\frac{\ldots \quad M_i \Downarrow_\lambda A_i \quad \ldots}{\mathsf{ast_{promote}}(\mathsf{tag_t}, \tilde{M}) \Downarrow_\lambda \mathsf{ast_{promote}}(\mathsf{tag_t}, \tilde{A})} \text{ Promote}$$

Figure 4 Additional rules defining $\Downarrow_{ct}$, $\Downarrow_{dl}$ and $\Downarrow_\lambda$ for AST promotion.

The additional reduction rules for eval are shown in Figure 3. [Eval rt] reduces M to a value, which must be an AST, and which is then turned into a normal λ term and executed. Note that unlike compile-time HGMP, eval cannot introduce new variables into a scope. A detailed example of running eval is given in Figure 11.

4 Enriching the calculus

The simple calculus of the previous section allows readers to concentrate on the core of our approach. However, it is too spartan to model important properties of real languages. In this section, we enrich the simple calculus with further features which add complexity but allow us to model real languages.

4.1 Higher-order ASTs

The simple system in Section 3.1 does not allow higher-order meta-programming (e.g. meta-meta-programming). While the simple ASTs we introduced in Section 3.1 are sufficient to represent normal λ-calculus terms as an AST, programs with ASTs cannot be represented as ASTs. While not all real languages (e.g. Template Haskell) allow higher-order meta-programming, many do (e.g. MetaML and Converge). We thus introduce higher-order ASTs now to make the presentation of later features consistent.

Higher-order ASTs need a means to represent programs that can create ASTs as ASTs themselves. There are many plausible ways that this could be done: the mechanism we settled upon allows extra syntactic elements to be easily added by further extensions. The basis of our approach is a new datatype $\mathsf{ast_{promote}}(M, \tilde{N})$ which allows an arbitrary AST with a tag M and parameters $\tilde{N}$ to be promoted up a meta-level. We thus need to introduce

$$\frac{M \Downarrow_{ul} A}{\uparrow\{M\} \Downarrow_{ct} A} \text{ UpML ct} \qquad \frac{M \Downarrow_{ct} A}{\downarrow\{M\} \Downarrow_{ul} A} \text{ DownML ul}$$

$$\frac{}{"x" \Downarrow_{ul} \mathsf{ast}_{\mathsf{string}}("x")} \text{ String ul} \qquad \frac{M \Downarrow_{ul} A \quad N \Downarrow_{ul} B}{MN \Downarrow_{ul} \mathsf{ast}_{\mathsf{app}}(A,B)} \text{ App ul}$$

$$\frac{M \Downarrow_{ul} A}{\mu g.\lambda x.M \Downarrow_{ul} \mathsf{ast}_{\mathsf{rec}}(\mathsf{ast}_{\mathsf{string}}("g"), \mathsf{ast}_{\mathsf{string}}("x"), A)} \text{ Rec ul}$$

$$\frac{M \Downarrow_{ul} A}{\lambda x.M \Downarrow_{ul} \mathsf{ast}_{\mathsf{lam}}(\mathsf{ast}_{\mathsf{string}}("x"), A)} \text{ Lam ul} \qquad \frac{}{\mathsf{tag_t} \Downarrow_{ul} \mathsf{tag_t}} \text{ Tag ul}$$

$$\frac{M \Downarrow_{ul} A}{\mathsf{eval}(M) \Downarrow_{ul} \mathsf{ast}_{\mathsf{eval}}(A)} \text{ Eval ul} \qquad \frac{M \Downarrow_{ul} A \quad A \Downarrow_{ul} B}{\uparrow\{M\} \Downarrow_{ul} B} \text{ UpML ul}$$

$$\frac{}{x \Downarrow_{ul} \mathsf{ast}_{\mathsf{var}}("x")} \text{ Var ul} \qquad \frac{... \quad M_i \Downarrow_{ul} A_i \quad ...}{\mathsf{ast_t}(\tilde{M}) \Downarrow_{ul} \mathsf{ast}_{\mathsf{promote}}(\mathsf{tag_t}, \tilde{A})} \text{ Ast ul}$$

■ **Figure 5** Additional rules for upMLs.

a way for programs to reference tags arbitrarily, and extend the syntax as follows:

$$M ::= \dots \mid \mathsf{tag_t} \qquad\qquad t ::= \dots \mid \mathsf{promote}$$

The corresponding reduction rules are in Figure 4. Note that tags are normal values in the calculus so that one can write programs which can create higher-order ASTs. Promoted ASTs can then be reduced one meta-level with the existing $\Downarrow_{dl}$ relation. For example, $\mathsf{ast}_{\mathsf{promote}}(\mathsf{string}, \mathsf{ast}_{\mathsf{string}}("x")) \Downarrow_{dl} \mathsf{ast}_{\mathsf{string}}("x")$.

4.2 UpMLs

Using AST constructors alone to perform HGMP is tiresome—while it gives complete flexibility, the sheer verbosity of such an approach quickly overwhelms even the most skilled and diligent programmer. UpMLs ameliorate this problem by allowing concrete syntax to be used to represent ASTs (see Section 2.3). Since, depending on a language's syntax, upMLs can be less expressive than ASTs, we model upMLs as a transparent compile-time expansion to the equivalent AST constructor calls e.g. $\uparrow\{2\} \Downarrow_{ct} \mathsf{ast}_{\mathsf{int}}(2)$.

To add UpMLs to our language, we first extend the grammar as follows:

$$M ::= \dots \mid \uparrow\{M\} \qquad\qquad t ::= \dots$$

Note that, like downMLs, upMLs have disappeared after the compile-time stage, so we have no need to make an AST equivalent of them.

Figure 5 shows the reduction rules needed for upMLs including the new $\Downarrow_{ul}$ reduction relation which handles the upML to AST conversion. When, during the recursive sweep of a program by the $\Downarrow_{ct}$ reduction relation, an upML is encountered, it is handed over to the $\Downarrow_{ul}$ reduction relation which translates a λ-term into its AST equivalent.

The major subtlety in the new rules relates to an important practical need. UpMLs on their own can only construct ASTs of a fixed 'shape' and are thus rather limited. Languages with upMLs (or their equivalents) therefore allow holes to be put into them where arbitrary ASTs can be inserted; in essence, the upML serves as a template with defined points of

variability. In some languages (e.g. Converge) holes inside upMLs are syntactically differentiated from holes outside, but we use downMLs to represent such 'inner' holes. In the same way as top-level downMLs, inner downMLs are expected to return an AST; unlike top-level downMLs, they are evaluated at run-time not compile-time. The [DOWNML UL] rule therefore simply runs the expression inside it through the $\Downarrow_{ct}$ reduction relation and uses the result as-is. This allows examples such as the following:

$$\uparrow\{2+\downarrow\{\uparrow\{3+4\}\}\} \quad \Downarrow_{ct} \quad \mathsf{ast_{add}}(\mathsf{ast_{int}}(2), \mathsf{ast_{add}}(\mathsf{ast_{int}}(3), \mathsf{ast_{int}}(4)))$$

In our model, upMLs are simple conveniences for AST construction, rather as they were in early Lisp implementations. More recent languages (e.g. Scheme, Template Haskell) use upMLs in addition as a means of ensuring referential transparency and hygiene [7]. Our formulation of upMLs is designed to open the door for such possibilities, but it is beyond the scope of this paper to tackle them.

4.2.1 The relationship between compile-time levels

Readers may wonder why we have chosen the names upML and downML for what are often called backquote / quasi-quote and macro call / splice respectively. We build upon an observation from MetaLua [14] that these two operators are more deeply connected than often considered, though our explanation is somewhat different. Our starting point is to note that, during compilation, there are three stages that a compiler can go through: normal compilation; converting upMLs to ASTs; and running user code in a downML. UpMLs / downMLs not only control which stage the compiler is in at any point during compilation, but have a fundamental relation to ASTs which we now investigate.

We call the normal compilation stage level 0. AST constructors in normal λ-terms are simply normal datatype constructors. When we encounter a top-level upML, we shift stage 'up' to level +1. In this level we take code and convert it into an AST which represents the code. When we encounter a top-level downML, we shift stage 'down' to level -1. In this level we take code and run it.

The basic insight is that the compiler level corresponds to the ASTs created or consumed: at level 0 we neither create or consume ASTs; at level 1 we create them (with upMLs); and at level -1, we consume ASTs (downMLs must evaluate to an AST, which is then converted to a normal λ-term).

Building upon this, we can see that this notion naturally handles downMLs nested within upMLs (and vice versa), bringing out the symmetry between the two operators, which can cancel each other out. Consider a program which nests a downML in an upML (i.e. $\uparrow\{\downarrow\{M\}\}$). How is the program M dealt with? Compilation starts at level 0; the upML shifts it to level 1; and the downML shifts it back to level 0. Thus we can see that M is handled at the normal compilation level and neither creates or consumes ASTs at compile-time (the fact that, at run-time, M ultimately needs to evaluate to an AST is irrelevant from a compile-time perspective). Similarly, consider an upML nested inside a downML (i.e. $\downarrow\{\uparrow\{M\}\}$). Since the downML shifts the compiler to level -1 and the upML shifts it back to level 0, we can see that the end effect is that M is dealt with as if it had always been at the normal compilation level.

In fact, the notion of these 3 levels (-1, 0, +1) is sufficient to explain arbitrarily nested downMLs and upMLs. For example, we can clearly see that two nested upMLs (i.e. $\uparrow\{\uparrow\{M\}\}$) create an AST representation of M (at level 2) which can be turned back into a normal λ-term by two nested downMLs (operating at level -2). As this suggests, unlike

$$\frac{M \Downarrow_{ct} A \quad A \Downarrow_\lambda B \quad N[B/x] \Downarrow_{ct} C}{\text{let}_\downarrow x = \ M \text{ in } N \Downarrow_{ct} C} \text{ {\scriptsize LET CT}}$$

$$
\begin{aligned}
\downarrow\{M\}[N/x] &= \downarrow\{M[N/x]\} \\
\uparrow\{M\}[N/x] &= \uparrow\{M[N/x]\} \\
(\text{let}_\downarrow x = \ M \text{ in } N)[L/y] &= \begin{cases} \text{let}_\downarrow x = \ M[L/y] \text{ in } N[L/y] & x \neq y \\ \text{let}_\downarrow x = \ M \text{ in } N & x = y \end{cases}
\end{aligned}
$$

Figure 6 The additional reduction rule, as well as the substitution rules, for letdownMLs.

systems such as MetaML, we do not need to label expressions as belonging to a certain level, nor do we need to do anything special to handle levels extending to $-\infty$ or $+\infty$.

4.3 Lifting

Most HGMP languages allow semi-arbitrary run-time values to be lifted up a meta-level (e.g. an integer 3 to be converted to an AST $\text{ast}_{int}(3)$). In some languages lifting is implicit (e.g. in Template Haskell, a variable inside an upML which references a definition outside the upML, and which is of a simple type such as integers, is implicitly lifted), while in others it is explicit (e.g. Converge forces all lifting to be explicit). All the languages we are aware of that use implicit lifting determine this statically, and can be trivially translated to explicit lifting. We thus choose to model explicit lifting. We extend the grammar as follows:

$$M \quad ::= \quad ... \ | \ \text{lift}(M) \qquad t \quad ::= \quad ... \ | \ \text{lift}$$

Figure 7 shows the additional reduction rules. The rules for the $\Downarrow_{ct}$, $\Downarrow_{dl}$, and $\Downarrow_{ul}$ relations are mechanical. Capture-avoiding substitution is given as $\text{lift}(M)[N/x] = \text{lift}(M[N/x])$. The rules for $\Downarrow_\lambda$ show that lift is a polymorphic function, turning values of type T into an AST ast_T e.g. $\text{lift}(2+3) \Downarrow_\lambda \text{ast}_{int}(5)$.

The relation between upMLs and lift can be seen from the following examples (where $\circ$ represents relational composition):

- $\uparrow\{2+3\} \Downarrow_{ct} \text{ast}_{add}(\text{ast}_{int}(2), \text{ast}_{int}(3))$
- $\uparrow\{2+3\} \ (\Downarrow_{ct} \circ \Downarrow_\lambda) \ \text{ast}_{add}(\text{ast}_{int}(2), \text{ast}_{int}(3))$
- $\text{lift}(2+3) \Downarrow_{ct} \text{lift}(2+3)$
- $\text{lift}(2+3) \ (\Downarrow_{ct} \circ \Downarrow_\lambda) \ \text{ast}_{int}(5)$

4.4 Cross-level variable scoping

The downMLs modelled in Section 3.2 run each expression in a fresh environment with no link to the outside world. While in theory this is sufficiently expressive, in practice it is restrictive: downMLs cannot share code, and so each downML must include within it a copy of every library function it wishes to use. Languages such as Template Haskell therefore allow variables defined outside downMLs (e.g. functions) to be referenced within a downML. Different languages have subtly different mechanisms to define which variables are available within a downML (e.g. Converge allows, with some restrictions, variables defined within a module M to be used in a downML within that module; Template Haskell only allows

variables imported from other modules to be used in a downML), and we do not wish to model the specifics of any one language's scheme.

We therefore provide a simple abstraction which can be used to model the scoping rules of different languages. The *letdownML* construct $\mathsf{let}_\downarrow x = M$ in N makes a program M available as x to N at compile-time (i.e. including inside downMLs). We extend the grammar as follows:

$$M \quad ::= \quad ... \mid \mathsf{let}_\downarrow x = M \text{ in } N \qquad t \quad ::= \quad ...$$

The additional reduction rule for letdownML is given in Figure 6. As this shows, letdownMLs are let bindings that are performed at compile-time rather than run-time. Figure 6 therefore also defines the additional substitution rules required.

4.5 Examples

The staged power function [8] has become a standard way of comparing HGMP approaches. The idea is to specialise the function $\lambda n x . x^n$ with respect to its first argument. This is more efficient than implementations with variable exponent, provided the cost of specialisation is amortised through repeated use at run-time. To model this in our calculus, we assume the existence of the standard recursion operator $\mu g . \lambda x . M$ that makes g available for recursive calls in M. The staged power function then becomes:

$$M = \mu p . \lambda n . \text{if } n = 1 \text{ then } \uparrow\{x\} \text{ else } \uparrow\{x \times \downarrow\{p\ (n-1)\}\}$$
$$power = \lambda n . \uparrow\{\lambda x . \downarrow\{M\ n\}\}$$

For example *power* 3 reduces to an AST equivalent to that generated by $\uparrow\{\lambda x . x \times x \times x\}$. The function *power* can be used to specialise code at compile-time:

$$\mathsf{let}\ cube\ = \downarrow\{power\ 3\} \text{ in } (cube\ 4) + (cube\ 5)$$

and at run-time:

$$\mathsf{let}\ cube\ = \mathsf{eval}(power\ 3) \text{ in } (cube\ 4) + (cube\ 5)$$

By stretching the example somewhat, we can also show higher-order HGMP in action. Assume we wish to produce a variant of *power* which takes one part of the exponent early on in a calculation, with the second part known only later (e.g. because we want to compute $\lambda n . x^{m+n}$ frequently for a small number of different n that become available after m is known). We can then use the following higher-order meta-program for this purpose:

$$power_{ho} = \lambda m . \uparrow\{\lambda n . \uparrow\{\lambda x . \downarrow\{M\ (m+n)\}\}\}$$

and use it in a number of different ways e.g.:

$$\mathsf{let}\ cube\ = \downarrow\{\downarrow\{power_{ho}\ 1\}\ 2\} \text{ in } cube\ 4$$

which specialises both arguments at compile-time, or

$$\mathsf{let}\ f\ = \downarrow\{power_{ho}\ 1\} \text{ in } \mathsf{let}\ cube\ = f\ 2 \text{ in } \mathsf{eval}(cube)\ 4$$

where the first argument is specialised at compile-time, and the second at run-time.

$$\frac{M \Downarrow_{ct} N}{\mathsf{lift}(M) \Downarrow_{ct} \mathsf{lift}(N)} \quad \frac{M \Downarrow_{dl} N}{\mathsf{ast}_{\mathsf{lift}}(M) \Downarrow_{dl} \mathsf{lift}(N)} \quad \frac{M \Downarrow_{ul} N}{\mathsf{lift}(M) \Downarrow_{ul} \mathsf{ast}_{\mathsf{lift}}(N)}$$

$$\frac{c \text{ is an integer}}{\mathsf{lift}(c) \Downarrow_\lambda \mathsf{ast}_{\mathsf{int}}(c)} \quad \frac{c \text{ is a string}}{\mathsf{lift}(c) \Downarrow_\lambda \mathsf{ast}_{\mathsf{string}}(c)}$$

Figure 7 Additional rules for lifting. For simplicity, we only define lifting for integers and strings, but one can define lifting for any type desired.

5   A recipe for creating HGMP calculi

Nothing in the presentation of our calculus has relied on the λ-calculus as starting point. In this section, we build upon an observation from Converge that HGMP can easily be detached from the 'base' language it has been added to [37]: we informally show how one can add HGMP features to a typical programming language.

Let us imagine that we have a language L which we wish to extend with HGMP features to create L_{mp}. We assume that L has its syntax given as an algebraic signature (with an indication of bindings), and the semantics as a rule system over the syntax. We require L to have a string-esque datatype to represent variables; such a datatype can be trivially added to L if not present. We can then create L_{mp} as follows:

- Mirror every syntactic element of L with an AST and a tag.
- Add eval, $\mathsf{ast}_{\mathsf{promote}}$ and their corresponding tags.
- Add upMLs, downMLs, and letdownMLs.
- Add appropriate reduction rules for ASTs, upMLs, downMLs, and letdownMLs.

Semi-formally, we define L_{mp}'s syntax as follows. Assuming that C is the set of L's program constructors, L_{mp}'s constructors and tags are defined as follows:

$$T = C \cup \{\mathsf{eval}, \mathsf{promote}\}$$

$$C_{mp} = C \;\cup\; \{\mathsf{eval}, \downarrow\!\{\_\}, \uparrow\!\{\_\}, \mathsf{let}_\downarrow\} \cup \{\mathsf{ast}_\mathsf{t} \mid t \in T\} \cup \{\mathsf{tag}_\mathsf{t} \mid t \in T\}$$

The arities and binders of the new syntax are as follows:
- If $c \in C$ then its arity and binders are unchanged in C_{mp}.
- ast_c has the same arity as $c \in C$ and no binders.
- $\mathsf{ast}_{\mathsf{promote}}$ has variable arity, or, equivalently has arity 2, with the second argument being of type list. There are no binders.
- $\mathsf{ast}_{\mathsf{eval}}$ has arity 1 and no binders.
- tag_t has arity 0 and no binders for $t \in T$.
- eval, $\downarrow\!\{\_\}$, and $\uparrow\!\{\_\}$ have arity 1 and no binders.
- $\mathsf{let}_\downarrow$ has arity 3 and its first argument is binding.

L_{mp} inherits all of L's reduction rules. L_{mp}'s $\Downarrow_\lambda$ reduction relation is then augmented with the following rules:

$$\frac{t \in T}{t \Downarrow_\lambda t} \qquad \frac{L \Downarrow_\lambda M \quad M \Downarrow_{dl} N \quad N \Downarrow_\lambda N'}{\mathsf{eval}(L) \Downarrow_\lambda N'} \qquad \frac{... \quad M_i \Downarrow_\lambda N_i \quad ... \quad t \in T}{\mathsf{ast}_\mathsf{t}(\tilde{M}) \Downarrow_\lambda \mathsf{ast}_\mathsf{t}(\tilde{N})}$$

A $\Downarrow_{dl}$ relation must then be added to L_{mp}. The definition of this relation follows the same pattern as that of $\Downarrow_{dl}$ in the calculus presented earlier in the paper: each L constructor c

must have a rule in $\Downarrow_{dl}$ to convert it from an AST to a normal calculus term. If a constructor c has no binders, the corresponding rule is simple:

$$\frac{...\quad M_i \Downarrow_{dl} N_i \quad ...}{\mathsf{ast_c}(\tilde{M}) \Downarrow_{dl} c(\tilde{N})}$$

Two examples of such rules are as follows:

$$\overline{\mathsf{ast_{var}}("x") \Downarrow_{dl} x} \qquad \overline{\mathsf{ast_{string}}("x") \Downarrow_{dl} "x"}$$

Constructors with binders are most easily explained by example. If c has arity 2, with the first argument being a binder, the following rule must be added:

$$\frac{M \Downarrow_{dl} "x" \quad N \Downarrow_{dl} N'}{\mathsf{ast_c}(M, N) \Downarrow_{dl} c(x, N')}$$

The following rules must be added for higher-order ASTs:

$$\frac{M \Downarrow_{dl} N}{\mathsf{ast_{eval}}(M) \Downarrow_{dl} \mathsf{eval}(N)} \qquad \frac{L \Downarrow_{dl} \mathsf{tag_t} \quad M_i \Downarrow_{dl} N_i \quad t \in T}{\mathsf{ast_{promote}}(L, \tilde{M}) \Downarrow_{dl} \mathsf{ast_c}(\tilde{N})}$$

$$\frac{L \Downarrow_{dl} \mathsf{tag_{promote}} \quad M \Downarrow_{dl} \mathsf{tag_t} \quad N_i \Downarrow_{dl} R_i}{\mathsf{ast_{promote}}(L, M, \tilde{N}) \Downarrow_{dl} \mathsf{ast_{promote}}(\mathsf{tag_t}, \tilde{R})} \qquad \frac{t \in T}{\mathsf{tag_t} \Downarrow_{dl} \mathsf{tag_t}}$$

Assuming we wish to enable compile-time HGMP, a $\Downarrow_{ct}$ relation must be added:

$$\frac{M \in \{x, "x"\} \cup \{\mathsf{tag_t} \mid t \in T\}}{M \Downarrow_{ct} M} \qquad \frac{M \Downarrow_{ct} N}{\mathsf{eval}(M) \Downarrow_{ct} \mathsf{eval}(N)}$$

$$\frac{M_i \Downarrow_{ct} N_i \quad c \in C}{c(\tilde{M}) \Downarrow_{ct} c(\tilde{N})} \qquad \frac{M_i \Downarrow_{ct} N_i \quad t \in T}{\mathsf{ast_t}(\tilde{M}) \Downarrow_{ct} \mathsf{ast_t}(\tilde{N})}$$

$$\frac{t \in T}{\mathsf{tag_t} \Downarrow_{ct} \mathsf{tag_t}} \qquad \frac{M \Downarrow_{ct} A \quad A \Downarrow_\lambda B \quad B \Downarrow_{dl} C}{\Downarrow\{M\} \Downarrow_{ct} C}$$

$$\frac{M \Downarrow_{ct} A \quad A \Downarrow_\lambda B \quad N[B/x] \Downarrow_{ct} C}{\mathsf{let_\downarrow} x = M \text{ in } N \Downarrow_{ct} C}$$

Note that the last two rules (for downMLs and letdownMLs) are added unchanged from earlier in the paper (we assume for letdownMLs that L has a suitable notion of capture-avoiding substitution, which is extended to L_{mp} as described in Section 3). The rules for upMLs are given by a new relation $\Downarrow_{ul}$ which is a trivial variation of that in Figure 5 and omitted for brevity.

While semi-mechanically creating L_{mp} from L easily results in a new language with HGMP, we cannot guarantee that L_{mp} will always respect the 'spirit' of L. For example, adding HGMP to the π-calculus in this fashion would lead to HGMP that executes sequentially (e.g. in the evaluation of downMLs) which may not be desirable (although the resulting HGMPified π-calculus would be a good starting point for developing message-passing based forms of HGMP). Nevertheless, for most sequential programming languages, we expect L_{mp} to be in the spirit of L. As this shows, the HGMP features of L_{mp} are easily considered separately. We suggest this helps explain how such systems have been retro-fitted on languages such as Haskell, and gives pointers for designers of other languages who wish to consider adding HGMP.

Language	Monotyped	Parameterised	Hybrid	Dynamic	Staged	Upfront
Converge	○	○	○	●	n/a	n/a
Lisp	○	○	○	●	n/a	n/a
MetaML	○	●	○	○	○	●
Template Haskell	●	○	○	○	●	○
`scala.meta`	●	○	○	○	●	○

Figure 8 How different HGMP languages approach typing.

6   Example: staged typing and HGMP

In conventional programming languages, static typing provides compile-time guarantees that certain classes of error cannot happen at run-time. However, HGMP blurs the lines between compile-time and run-time, causing complications in typing that have not yet been fully resolved. The purpose of this section is to demonstrate that our calculus can also be useful for studying static typing in an HGMP language. We do this by defining a type system which is conceptually close to Template Haskell's.

6.1   Design issues

The three major design questions for static typing in the face of HGMP are: what does type-safety mean in multi-staged languages? When should static types be enforced? And: what static types should ASTs have?

Type-safety normally means that programs cannot get 'stuck' at run-time, in the sense that the program gets to a point where no reduction rules can be applied to it, but it has not yet reached a value. The static typing system identifies such programs and prevents them from being run. Alas, concepts like "stuck" and even "value" are not straightforward in HGMP languages. This section will only outline some of the key issues. We leave a detailed investigation as further work.[4]

There are two main choices for when static types are to be checked in an HGMP language: upfront or in stages. Upfront typing as found in MetaML guarantees that any program which statically type-checks cannot get stuck in any later stage [35, 36]. This strong guarantee comes at a price: many seemingly reasonable meta-programs fail to type-check, at least for simple typing systems (e.g. admitting type inference). We therefore believe that – except, perhaps, for verification-focused languages – staged type-checking is the more practical of the two approaches. It guarantees only that, whenever an AST is $\Downarrow_{dl}$ converted to a normal λ-term as a result of a downML or eval, the program will not get stuck before the next such conversion. Thus, type-checking might need to be carried out more than once, and the guarantees at each stage are weaker than in upfront checking. In the rest of this paper, we only consider staged type-checking.

There are three main ways that code (be that ASTs or MetaML-esque quasi-quotes) can be statically typed. In a *monotyped* system, every program representing code has the same type Code. In a *parameterised* system, code of type $\mathsf{Code}\langle\alpha\rangle$ can be shifted a meta-level (at

---

[4] For example, the computation described by $\Downarrow_{ul}$ relation does not have a syntactic notion of value. Consider the term $\mathsf{ast_{int}}(3)$. Whether one should consider it as a value with respect to $\Downarrow_{ul}$ depends on whether $\Downarrow_{ul}$ was previously applied to 3 or not. This complicates defining a small-step semantics corresponding to $\Downarrow_{ul}$.

$$\frac{M \Downarrow_{ct} N}{\mathsf{eval}^{\alpha}(M) \Downarrow_{ct} \mathsf{eval}^{\alpha}(N)} \quad \frac{M \Downarrow_{dl} N}{\mathsf{ast}_{\mathsf{eval}(\alpha)}(M) \Downarrow_{dl} \mathsf{eval}^{\alpha}(N)}$$

$$\frac{M \Downarrow_{ul} A}{\mathsf{eval}^{\alpha}(M) \Downarrow_{ul} \mathsf{ast}_{\mathsf{eval}(\alpha)}(A)}$$

◼ **Figure 9** $\Downarrow_{ct}$, $\Downarrow_{dl}$, $\Downarrow_{ul}$ reduction relations for the altered eval.

compile-time or run-time) to a program of type α. Finally, it is possible to bridge these two extremes with a *hybrid* system which allows both parameterised and monotyped code types. MetaML uses parameterised code types. Template Haskell is currently monotyped (though there are proposals for it to move to a hybrid system [28]) and we thus use that as the basis of our typing system. Figure 8 surveys how different HGMP languages approach typing.

6.2 Staged typing for the foundational calculus

The key properties in the staged typing we define are as follows:

1. All code that is evaluated by the $\Downarrow_{\lambda}$ relation will have been previously type-checked and thus cannot get stuck. This is done by type-checking the expressions inside downMLs and evals, and type-checking the complete program after all downMLs have been removed.
2. All code that is converted by the $\Downarrow_{dl}$ relation will have been previously type-checked to ensure that the ASTs involved are properly formed and thus applying $\Downarrow_{dl}$ cannot get stuck.
3. Since the only possible places where $\Downarrow_{ct}$ could get stuck are where it references $\Downarrow_{\lambda}$ and $\Downarrow_{dl}$, (1, 2) guarantee that $\Downarrow_{ct}$ doesn't get stuck.
4. Since $\Downarrow_{ul}$ could only get stuck where it references $\Downarrow_{ct}$, (3) guarantees that $\Downarrow_{ul}$ doesn't get stuck.

To make this form of typing concrete, we create a type system for this paper's calculus (modifying eval for reasons that will soon become clear). The type system can be seen as an extension of Template Haskell's, augmented with higher-order HGMP and run-time HGMP.

We first need to extend the calculus grammar to introduce types as follows:

$$
\begin{aligned}
M \quad &::= \quad ... \mid \mu g.\lambda x.M \mid \mathsf{eval}^{\alpha}(M) \\
t \quad &::= \quad ... \mid \mathsf{rec} \mid \mathsf{eval}(\alpha) \\
\alpha \quad &::= \quad \mathsf{Int} \mid \mathsf{Bool} \mid \alpha \to \beta \mid \mathsf{String} \mid \mathsf{Code} \mid \mathsf{Tag}_t
\end{aligned}
$$

We assume readers are acquainted with static types for the basic λ-calculus. ASTs have type Code; each tag_t has a corresponding static type Tag_t. The only surprising change is the type annotation of $\mathsf{eval}^{\alpha}(M)$ and the corresponding tag $\mathsf{eval}(\alpha)$. The type annotation α is used for type-checking the program N, obtained from M by evaluation to an AST and subsequent $\Downarrow_{dl}$ conversion back to a normal λ-term: all we have to do is verify that N has type α. Without this type annotation, we would have to type-check N and ensure that N's type is compatible with its context. The additional reduction rules for all relations except $\Downarrow_{\lambda}$ are shown in Figure 9.

The core of the approach is to intersperse type-checking with reduction, making sure that we can never run code that has not been type-checked. We therefore first add a 'normal'

$$\frac{\Gamma, x : \alpha \vdash M : \beta}{\Gamma \vdash \lambda x M : \alpha \to \beta} \quad \frac{\Gamma, g : \alpha \to \beta, x : \alpha \vdash M : \beta}{\Gamma \vdash \mu g.\lambda x.M : \alpha \to \beta} \quad \overline{\Gamma, x : \alpha \vdash x : \alpha}$$

$$\frac{\Gamma \vdash M : \alpha \to \beta \quad \Gamma \vdash N : \alpha}{\Gamma \vdash MN : \beta} \quad \frac{\Gamma \vdash M : \mathsf{String} \quad \Gamma \vdash N : \mathsf{Code}}{\Gamma \vdash \mathsf{ast_{lam}}(\mathsf{ast_{string}}(M), N) : \mathsf{Code}}$$

$$\overline{\Gamma \vdash \mathsf{tag_t} : \mathsf{Tag}_k} \quad \frac{\Gamma \vdash L : \mathsf{String} \quad \Gamma \vdash M : \mathsf{String} \quad \Gamma \vdash N : \mathsf{Code}}{\Gamma \vdash \mathsf{ast_{rec}}(\mathsf{ast_{string}}(L), \mathsf{ast_{string}}(M), N) : \mathsf{Code}}$$

$$\frac{t \neq \mathsf{lam, rec, promote, string, int} \quad ... \quad \Gamma \vdash M_i : \mathsf{Code} \quad ...}{\Gamma \vdash \mathsf{ast_t}(\tilde{M}) : \mathsf{Code}}$$

$$\frac{\Gamma \vdash M : \mathsf{Int}}{\Gamma \vdash \mathsf{ast_{int}}(M) : \mathsf{Code}} \quad \frac{\Gamma \vdash M : \mathsf{String}}{\Gamma \vdash \mathsf{ast_{string}}(M) : \mathsf{Code}}$$

$$\frac{\Gamma \vdash M : \mathsf{Code}}{\Gamma \vdash \mathsf{eval}^\alpha(M) : \alpha} \quad \frac{\Gamma \vdash M : \mathsf{Tag}_t \quad ... \quad \Gamma \vdash N_i : \mathsf{Code} \quad ...}{\Gamma \vdash \mathsf{ast_{promote}}(M, \tilde{N}) : \mathsf{Code}}$$

Figure 10 Type-checking with type Code for programs not containing upMLs and downMLs. Some straightforward cases omitted.

type-checking phase between compile-time and run-time (i.e for programs which do not use downMLs or eval, this phase is the only type-check invoked):

$$\underbrace{M \Downarrow_{ct} A}_{\text{compile-time}} \quad \overbrace{\vdash A : \alpha}^{\text{type-checking}} \quad \underbrace{A \Downarrow_\lambda V}_{\text{run-time}}$$

Second, we must add a type-checking phase to ensure that code generated at compile-time and inserted into the program by downMLs is type-safe (i.e. the expression in a downML has a static type of Code). We therefore alter [DownML CT] as follows:

$$\frac{M \Downarrow_{ct} A \quad \vdash A : \mathsf{Code} \quad A \Downarrow_\lambda B \quad B \Downarrow_{dl} C}{\Downarrow\{M\} \Downarrow_{ct} C} \text{ \scriptsize DownML CT}$$

Finally, we alter [EVAL RT] to perform a type-check on the code it will evaluate at run-time:

$$\frac{L \Downarrow_\lambda M \quad M \Downarrow_{dl} N \quad \vdash N : \alpha \quad N \Downarrow_\lambda N'}{\mathsf{eval}^\alpha(L) \Downarrow_\lambda N'} \text{ \scriptsize EVAL RT}$$

We define the typing judgement $\Gamma \vdash M : \alpha$ as follows. M is a program that does not contain upMLs or downMLs (which will have been removed by $\Downarrow_{ct}$ before type-checking). Γ is an environment (i.e. a finite map) from variables to types such that all of M's free variables are in the domain of Γ. Note that type-checking needs to be applied only to programs without downMLs and upMLs, hence the free variables of a program not containing upMLs and downMLs can be defined as usual for a program, and we omit the details. We write $\Gamma, x : \alpha$ for the typing environment that extends Γ with a single entry, mapping x to α, assuming that x is not in Γ's domain. We write $\vdash M : \alpha$ to indicate that the environment is empty.

The rules defining $\Gamma \vdash M : \alpha$ are given in Figure 10. The rules for variables, function abstraction, recursion, and application are as in conventional λ-calculus. For ASTs $\mathsf{ast_t}(\tilde{M})$ where t is not one of lam, rec, promote, string, or int, if all the arguments have type Code then $\mathsf{ast_t}(\tilde{M})$ also has type Code. ASTs representing a binding construct (e.g. $\mathsf{ast_{lam}}(M, N)$)

have type Code if: the terms representing binders are of the form $\mathsf{ast_{string}}(L)$ with L having type String; and N has type Code.

6.3 Examples

With the typing system defined, a few examples can help understand how and when it operates. First we note that terms such as $\mathsf{ast_{lam}}((\lambda x.x), M)$ that would get stuck without the type system do not type-check in our system. Second we can see that some expressions pass one type-check and fail a later one. Consider the following program:

$$2 + \Downarrow\{ \ \mathsf{ast_{lam}}(\mathsf{ast_{string}}("x"), \mathsf{ast_{var}}("x")) \ \}$$

The downML $\Downarrow_{ct}$ reduces to $\mathsf{ast_{lam}}(\mathsf{ast_{string}}("x"), \mathsf{ast_{var}}("x"))$ which successfully type-checks as being of type code. The entire program then $\Downarrow_{ct}$-reduces to $2 + \lambda x.x$, which is neither a value nor has any $\Downarrow_\lambda$-reductions and fails to type-check.

The type system can also check more complex properties, such as AST constructors with the wrong number of arguments. Let M be the program $\mathsf{ast_{promote}}(\mathsf{tag_{int}}, \mathsf{ast_{int}}(1))$ in the following program:

$$\mathsf{ast_{promote}}(\mathsf{tag_{promote}}, \mathsf{tag_{int}}, M, M)$$

When run through a downML or eval for the first time it will yield:

$$\mathsf{ast_{promote}}(\mathsf{tag_{int}}, \mathsf{ast_{int}}(1), \mathsf{ast_{int}}(1))$$

which type-checks correctly. However, if this AST is run through a downML or eval it results in:

$$\mathsf{ast_{int}}(1, 1)$$

which fails to type-check.

7 Related work

Meta-programming is such a long-studied subject that a full related work section would be a paper in its own right. We have referenced many real-world systems in previous sections; in this section, we therefore concentrate on related work that has a foundational or formal bent, and which has not been previously mentioned.

Run-time HGMP has received more attention than compile-time HGMP, with MetaML and the $reFL^{ect}$ language [18] being amongst the well known examples. MetaML is the closest in spirit to this paper, though it has two major, and two minor, differences. The minor differences are that MetaML is typed and hygienic, whereas our system can model untyped and non-hygienic systems, enabling people to experiment with different notions of each. The first major difference is that MetaML does not model compile-time evaluation of arbitrary code (see below for a discussion of MacroML, which partly addresses this). The second major difference is upMLs and ASTs: MetaML has only the former, while our system has both, with ASTs the 'fundamental' construct and upMLs a convenience atop them. As discussed in Section 2.3, this restricts the programs – and hence programming languages – that can be expressed.

Run-time HGMP is also the primary object of study in *unstaging translations* (see e.g. [6, 21, 24]). These are semantics-preserving embeddings of an HGMP language into a language

without explicit constructs for representing code as data. Depending on the particular pairing of source and target language, the unstaging translation can be extremely complex, making it difficult to use as a mechanism for understanding the fundamental constructs. We are also not aware of unstaging translations that treat compile-time and run-time HGMP in a unified way. An open research question is whether our general approach in Section 5 can be unstaged in a generic way.

Compile-time HGMP research has mostly focused on Lisp macros (e.g. [3, 20]) or C++ templates (e.g. [17]). Perhaps the work most similar to ours is the formal model of a large subset of Racket's macro system [13]. However, this formalises Racket's `define-syntax` system which is dynamically typed, and not a HGMP system in our definition (see Section 2.6). The system we define is closer in spirit to a statically typed version of Racket's `syntax-case` system. MacroML [16] investigates Lisp-style macro systems by translation into MetaML. The key insight is that macros are special constructs which must be entirely expanded in a separate stage before any normal code is evaluated. Our approach instead models systems where normal code can be evaluated at both compile-time and run-time.

Research on types for HGMP and the relationship with modal logics via a Curry-Howard correspondence started with work by Davis and Pfenning [10, 9]. In recent years, more expressive typing systems along these lines have been investigated (see e.g. [27, 39]). The axiomatic semantics of HGMP is explored in [2]. Some original approaches towards the foundations of run-time HGMP are: M-LISP [26] which provides an operational semantics for a simplified Lisp variant with `eval` but without macros; Archon [34], which is based on the untyped λ-calculus but without an explicit representation of code; the two-level λ-calculus [15] which is based on nominal techniques; and the ρ-calculus [25] which combines ideas from Conway games and π-calculus.

Issues closely related to HGMP have been studied in the field of logic, often under the heading of reflection [19]. Little work seems to have been done towards unification of the multiple approaches to meta-programming. Farmer et al.'s concept of syntax frameworks [11, 12] may well have been the first foray in this direction but are not intended to be a full model of meta-programming, whether homogeneous or heterogeneous. In particular, they do not capture the distinction between compile-time HGMP and run-time HGMP.

8   Conclusions

In this paper we presented the first foundational calculus for modelling compile-time and run-time HGMP as found in languages such as Template Haskell. The calculus is designed to be considered in increments, and adjusted as needed to model real-world languages. We provided a type system for the calculus. We hope that the calculus provides a solid basis for further research into HGMP.

The most obvious simplification in the calculus is its treatment of names. The calculus deliberately allows capture and is not hygienic since there are different styles of hygiene, and various possible ways of formalising it. A system similar to Template Haskell's, for example, where names in upMLs are preemptively α-renamed to fresh names would be a simple addition, but other, sometimes more complex, notions are possible (e.g. determining which variables should be fresh and which should allow capture). We hypothesise that the formal definition of hygiene in [1], which is based on nominal techniques [29], can be adapted to our foundational calculus.

───── **References** ─────

1 Michael D. Adams. Towards the essence of hygiene. In *Proc. POPL*, 2015.

2 Martin Berger and Laurence Tratt. Program Logics for Homogeneous Metaprogramming. In *Proc. LPAR*, pages 64–81, 2010.

3 Ana Bove and Laura Arbilla. A confluent calculus of macro expansion and evaluation. In *Proc. LFP*, pages 278–287, 1992.

4 Gilad Bracha and David Ungar. Mirrors: design principles for meta-level facilities of object-oriented programming languages. In *Proc. OOPSLA*, pages 331–344, 2004.

5 Cristiano Calcagno, Walid Taha, Liwen Huang, and Xavier Leroy. Implementing Multi-stage Languages Using ASTs, Gensym, and Reflection. In *Proc. GPCE*, pages 57–76, 2003.

6 Wontae Choi, Baris Aktemur, Kwangkeun Yi, and Makoto Tatsuta. Static Analysis of Multi-staged Programs via Unstaging Translation. In *Proc. POPL*, pages 81–92, 2011.

7 William Clinger and Jonathan Rees. Macros that work. In *Proc. POPL*, pages 155–162, 1991.

8 Krzysztof Czarnecki, John O'Donnell, Jörg Striegnitz, and Walid Taha. DSL Implementation in MetaOCaml, Template Haskell, and C++. In *Proc. Dagstuhl Workshop on Domain-specific Program Generation*, volume 3016, pages 50–71, 2004.

9 Rowan Davies. A temporal-logic approach to binding-time analysis. In *Proc. LICS*, pages 184–195, 1996.

10 Rowan Davies and Frank Pfenning. A modal analysis of staged computation. *J. ACM*, 48(3):555–604, 2001.

11 William M. Farmer. The formalization of syntax-based mathematical algorithms using quotation and evaluation. In J. Carette, editor, *Intelligent Computer Mathematics*, pages 35–50, 2013.

12 William M. Farmer and Pouya Larjani. Frameworks for reasoning about syntax that utilize quotation and evaluation. McSCert Report 9, McMaster University, 2013. Available at `http://imps.mcmaster.ca/doc/syntax.pdf`.

13 Matthew Flatt, Ryan Culpepper, David Darais, and Robert Bruce Findler. Macros that work together: Compile-time bindings, partial expansion, and definition contexts. *JFP*, 22(2):181–216, March 2012.

14 Fabien Fleutot and Laurence Tratt. Contrasting compile-time meta-programming in Metalua and Converge. In *Workshop on Dynamic Languages and Applications*, July 2007.

15 Murdoch J. Gabbay and Dominic P. Mulligan. Two-level lambda-calculus. In *Proc. WFLP*, volume 246, pages 107–129, 2009.

16 Steven E. Ganz, Amr Sabry, and Walid Taha. Macros as multi-stage computations: Type-safe, generative, binding macros in MacroML. In *Proc. ICFP*, pages 74–85, 2001.

17 Ronald Garcia and Andrew Lumsdaine. Toward Foundations for Type-Reflective Metaprogramming. *SIGPLAN Not.*, 45(2):25–34, October 2009.

18 Jim Grundy, Tom Melham, and John O'Leary. A reflective functional language for hardware design and theorem proving. *JFP*, 16(2):157–196, 2006.

19 John Harrison. Metatheory and Reflection in Theorem Proving: A Survey and Critique. Technical Report CRC-053, SRI International, 1995.

20 David Herman and Mitchell Wand. A theory of hygienic macros. In *Proc. ESOP*, pages 48–62, March 2008.

21 Jun Inoue, Oleg Kiselyov, and Yukiyoshi Kameyama. Staging beyond terms: Prospects and challenges. In *Proc. PEPM*, PEPM '16, pages 103–108, 2016.

22 Jun Inoue and Walid Taha. Reasoning about multi-stage programs. In *Proc. ESOP*, pages 357–376, 2012.

23 Ulrik Jørring and William L. Scherlis. Compilers and staging transformations. In *Proc. POPL*, pages 86–96. ACM, 1986.

24 Yukiyoshi Kameyama, Oleg Kiselyov, and Chung-chieh Shan. Closing the stage: From staged code to typed closures. In *Proc. PEPM*, pages 147–157, 2008.

25 L. Gregory Meredith and Matthias Radestock. A Reflective Higher-order Calculus. *ENTCS*, 141(5):49–67, 2005.

26 Robert Muller. M-LISP: A representation-independent dialect of LISP with reduction semantics. *TOPLAS*, 14(4):589–616, October 1992.

27 Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. *Transactions on Computational Logic*, 2007.

28 Simon Peyton Jones. New directions for Template Haskell. `http://ghc.haskell.org/trac/ghc/blog/TemplateHaskellProposal`, October 2010.

29 Andrew M. Pitts. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge University Press, 2013.

30 Jon Rafkind and Matthew Flatt. Honu: Syntactic extension for algebraic notation through enforestation. In *GPCE*, pages 122–131, 2012.

31 Tiark Rompf and Martin Odersky. Lightweight modular staging: a pragmatic approach to runtime code generation and compiled dsls. *Commun. ACM*, 55(6):121–130, 2012.

32 Tim Sheard. Accomplishments and research challenges in meta-programming. *Proc. SAIG '01*, 2196:2–44, September 2003.

33 Tim Sheard and Simon Peyton Jones. Template meta-programming for Haskell. In *Proc. Haskell workshop*, pages 1–16, 2002.

34 Aaron Stump. Directly reflective meta-programming. *Higher Order Symbol. Comput.*, 22(2):115–144, June 2009.

35 Walid Taha. *Multi-Stage Programming: Its Theory and Applications*. PhD thesis, Oregon Graduate Institute of Science and Technology, 1993.

36 Walid Taha and Michael Florentin Nielsen. Environment classifiers. In *Proc. POPL*, pages 26–37, 2003.

37 Laurence Tratt. Compile-time meta-programming in a dynamically typed OO language. In *Proc. DLS*, pages 49–64, October 2005.

38 Laurence Tratt. Domain specific language implementation via compile-time meta-programming. *TOPLAS*, 30(6):1–40, 2008.

39 Takeshi Tsukada and Atsushi Igarashi. A Logical Foundation for Environment Classifiers. *Logical Methods in Computer Science*, 6(4:8):1–43, 2010.

40 Willard van Orman Quine. *Mathematical Logic*. Harvard Univ. Press, 2003.

$$\cfrac{\cfrac{\cfrac{x \Downarrow_{ct} x}{\lambda x.x \Downarrow_{ct} \lambda x.x} \quad \cfrac{7 \Downarrow_{ct} 7}{\mathsf{ast_{int}}(7) \Downarrow_{ct} \mathsf{ast_{int}}(7)}}{(\lambda x.x)\,\mathsf{ast_{int}}(7) \Downarrow_{ct} (\lambda x.x)\,\mathsf{ast_{int}}(7)} \quad \cfrac{\cfrac{\dots}{(\lambda x.x)\,\mathsf{ast_{int}}(7)\} \Downarrow_{ct} 7}}{}}{\cfrac{(\lambda x.x) \Downarrow \{(\lambda x.x)\,\mathsf{ast_{int}}(7)\} \Downarrow_{ct} (\lambda x.x)7}{}}$$

$$\cfrac{x \Downarrow_{ct} x}{\lambda x.x \Downarrow_{ct} \lambda x.x}$$

$$\cfrac{\cfrac{\lambda x.x \Downarrow_\lambda \lambda x.x \quad \cfrac{(\lambda x.x)\,\mathsf{ast_{int}}(7) \Downarrow_\lambda \mathsf{ast_{int}}(7) \quad \cfrac{\mathsf{ast_{int}}(7) \Downarrow_{dl} 7 \quad 7 \Downarrow_\lambda 7}{x[7/x] \Downarrow_\lambda 7}}{\mathsf{eval}((\lambda x.x)\,\mathsf{ast_{int}}(7)) \Downarrow_\lambda 7}}{(\lambda x.x)(\mathsf{eval}((\lambda x.x)\,\mathsf{ast_{int}}(7))) \Downarrow_\lambda 7}}{}$$

$$(\lambda x.x)\,\mathsf{ast_{int}}(7) \Downarrow_\lambda \mathsf{ast_{int}}(7) \qquad \mathsf{ast_{int}}(7) \Downarrow_{dl} 7$$

Figure 11 Examples of compile-time HGMP (top) and run-time HGMP (bottom).