# The Kansas University rewrite engine

## A Haskell-Embedded Strategic Programming Language with Custom Closed Universes

NEIL SCULTHORPE, NICOLAS FRISBY and ANDY GILL*

*Information and Telecommunication Technology Center, The University of Kansas, USA*

(*e-mail:* {`neil,nfrisby,andygill`}`@ittc.ku.edu`)

## Abstract

When writing transformation systems, a significant amount of engineering effort goes into setting up the infrastructure needed to direct individual transformations to specific targets in the data being transformed. *Strategic programming languages* provide general-purpose infrastructure for this task, which the author of a transformation system can use for any algebraic data structure. The Kansas University Rewrite Engine (KURE) is a typed strategic programming language, implemented as a Haskell-embedded domain-specific language. KURE is designed to support *typed* transformations over *typed* data, and the main challenge is how to make such transformations compatible with generic traversal strategies that should operate over *any* type. Strategic programming in a typed setting has much in common with *datatype-generic programming*. Compared to other approaches to datatype-generic programming, the distinguishing feature of KURE's solution is that the user can configure the behaviour of traversals based on the *location* of each datum in the tree, beyond their behaviour being determined by the *type* of each datum. This article describes KURE's approach to assigning types to generic traversals, and the implementation of that approach. We also compare KURE, its design choices, and their consequences, with other approaches to strategic and datatype-generic programming.

## 1 Introduction

This article describes an approach for building rewrite engines over strongly typed data structures. Rewriting rules are often expressed as local correctness-preserving transformations, sometimes with preconditions and contextual requirements, such as lexical scope. A rewrite engine takes these local rewriting rules and applies them in systematic ways to achieve a global effect. Strategic programming (Visser, 2005), a style of programming that explicitly supports *programmable strategies* for composing rewriting rules, is one approach to structuring rewrite engines. The principal design decision in any *typed* strategic-programming implementation is how to specialise generic traversal strategies to operate over a typed syntax. This design decision becomes more challenging when we wish to support complex traversal strategies, such as potentially failing traversals, selective traversals, and traversals over mutually recursive data types.

The Kansas University Rewrite Engine (KURE) is a domain-specific language for typed strategic programming, and is the principal subject of this article. The KURE implementation began as a component of the HERA system (Gill, 2006), but was later abstracted out to a standalone Haskell library. Since then KURE has gone through several stages of development, with various experimental designs implemented. Two prior publications discuss KURE at earlier development stages: Gill (2009) described the first standalone version of KURE; Farmer *et al.* (2012) used an interim version of KURE as part of a larger system. The current implementation of KURE (Sculthorpe & Gill, 2014) has similarities with existing approaches to *datatype-generic programming*, but has different engineering compromises regarding data-structure traversal. KURE also builds on many years of research into strategic programming, especially the work on Stratego (Visser *et al.*, 1998; Visser, 2004) and Strafunski (Lämmel & Visser, 2002).

Specifically, the contributions of this article are:

- We present an approach of using *closed universes* to type generic traversals over mutually recursive data types (Section 3).
- We describe a Haskell implementation of strategic programming that uses these universes to index generic traversals (Section 4).
- We demonstrate how these universes can be used to define *statically selective traversals* (Section 4.3.2), and traversals whose behaviour can depend not just on the type of the values being traversed, but also on their *location* in the data structure (Section 4.3.3).
- We augment the implementation with a user-defined *context*, and show how it can be maintained automatically within a generic traversal (Section 4.4).
- We demonstrate the viability of using KURE for realistic applications by presenting a case study of KURE's usage in the HERMIT project (Section 5).
- We compare the KURE approach to generic traversals with that of Stratego, and with the SYB and Uniplate approaches to datatype-generic programming (Section 6).
- We compare KURE's performance with SYB and Uniplate, and explain the causes of the performance variations (Section 7).

## 2 Strategic and generic programming

Traditional *term rewriting* involves defining a set of rewriting rules over some object language, and applying those rules exhaustively to a term in that language. *Strategic programming* (Visser, 2005) builds on term rewriting by adding *programmable strategies* that allow the user to control when and where in a term rewriting rules are applied (Bravenboer *et al.*, 2008). Typically the object language is a programming or document-markup language, but the ideas generalise to any tree-structured data.

The term *generic programming* has multiple meanings (Gibbons, 2005), but in this article we always mean *datatype-generic programming* (also known as *polytypic programming*) (Hinze & Löh, 2007; Rodriguez Yakushev *et al.*, 2008; Hinze & Löh, 2009). This notion of generic programming involves defining functions that operate over typed data, but that operate based on the *shape* of the data rather than its *type*. Typically, *generic traversals* are used to navigate to a particular set of locations in a data type, and then a type-specific function is applied at those locations.

KURE is implemented as a Haskell library, providing rewriting strategies as Haskell combinators, and operating on Haskell data types. Hence, the KURE implementation can be viewed as either an embedded domain-specific language for strategic programming, or as a generic programming library.

In this section we will introduce the strategic programming paradigm by presenting Stratego (Visser *et al.*, 1998; Visser, 2004; Bravenboer *et al.*, 2008), the most widely used strategic rewriting system, and then overview other approaches to strategic and generic programming. First however, we will introduce our notation and terminology.

### 2.1 Terminology and notation

In Section 2.2 we will describe the Stratego language, and use Stratego syntax to do so. In Section 3 we will discuss assigning types to rewrites and strategies at a language-independent level, but use a Haskell-like notation for expressing types. In Sections 4 to 7 we use Haskell code, including several (commonly used) language extensions that are not part of the Haskell 2010 standard, but are provided by the Glasgow Haskell Compiler. In the Haskell code we also take the liberty of allowing type variables in top-level type signatures to scope over local type signatures, so that we can provide such type signatures as an aid to the reader (this could be achieved with a language extension and some additional syntactic clutter, but we believe that inhibits readability). In all sections, we typeset the program code `->` and `=>` as $\rightarrow$ and $\Rightarrow$, respectively, and in the Haskell code we typeset `\` as $\lambda$.

Throughout this paper we will be working with various tree-structured object languages. These object languages will be represented as algebraic data types, either in Stratego or in Haskell. In both cases, we typeset data-type constructors in sans-serif font.

As the distinction between the semantics of different strategic/generic languages can be quite subtle, we need to introduce some precise terminology for discussing algebraic data types. We recommend that the reader does not try to internalise all these definitions immediately, but refers back to this section as needed.

- The *proper components* of a value $x$ are the arguments to the constructor of $x$.
- The *components* of a value $x$ are its proper components, as well as $x$ itself.
- The *proper substructures* of $x$ are the proper components of $x$, and those components' proper substructures.
- The *substructures* of $x$ are its proper substructures, as well $x$ itself.
- Two substructures are *independent* if neither is a substructure of the other.
- The *similarly typed* proper components/substructures of $x$ are those that have the same type as $x$.
- The *maximal* elements of a set of substructures are those that are not a proper substructure of any other element.
- A *node* is a substructure that can be targeted by a rewriting rule.
- *Child nodes* are the maximal independent proper substructures of a node that are themselves also nodes.

Which substructures can be targets for a rewriting rule, and hence are considered to be nodes, varies between individual strategic/generic languages, so is not defined in general.

## 2.2 Stratego

Stratego is a strategic programming language designed to operate on an arbitrary object language. The Stratego language is part of the larger Stratego/XT toolkit (Visser, 2004; Bravenboer *et al.*, 2008), which provides additional tools such as parsers and pretty printers to convert between the object language and its representation as an algebraic data type within Stratego. Recently, Stratego has been integrated with the Eclipse integrated development environment, as part of the Spoofax Language Workbench (Kats & Visser, 2010).

As an example (adapted from Bravenboer *et al.*, 2008), a simple object language of integer addition could be defined in Stratego as follows:

**constructors**
$$\text{Add} : Exp * Exp \rightarrow Exp$$
$$\text{Int} \; : \textbf{INT} \qquad \rightarrow Exp$$

That is, an expression is either the addition of two expressions, or an integer. Integers, floats and strings are built-in types in Stratego. Although the constructors are assigned type signatures, Stratego is not a typed language — the Stratego compiler only checks that the arities of these constructor definitions match their usage (Bravenboer *et al.*, 2008).

Rewriting rules over this object language can be defined in the form $name : pat_1 \rightarrow pat_2$, where *name* is the name of the rule, $pat_1$ is a constructor pattern defining which nodes the rule can target, and $pat_2$ is the result of the rule. For example, the commutativity and left-unit laws of addition can be expressed as rewriting rules as follows:

**rules**
$$commutativity : \text{Add} \, (e_1, e_2) \rightarrow \text{Add} \, (e_2, e_1)$$
$$leftUnit \qquad : \text{Add} \, (\text{Int} \, (0), e) \rightarrow e$$

Note that $e$, $e_1$ and $e_2$ are meta-language variables that can match any node.

Rewriting rules can either succeed or fail. The typical cause of failure is when $pat_1$ does not match the node that the rule is being applied to. Stratego also provides two built-in rewriting rules:

- *id* is the identity rewriting rule that always succeeds;
- *fail* is the always-failing rewriting rule.

As well as rewriting rules, Stratego also provides some built-in *strategies*, which provide a means of combining rewriting rules. Some representative strategies are as follows:

- $r_1 \; ; \; r_2$ is *sequential composition*: apply $r_1$, then apply $r_2$, requiring both to succeed;
- $r_1 <+ r_2$ is *deterministic choice* (or "catch"): apply $r_1$, but if it fails then apply $r_2$;
- $all(r)$ is a *shallow traversal*: apply $r$ to all child nodes, requiring all to succeed.

Note that strategies (and rewriting rules) have a notion of a "current node", and that the $all \, (r)$ strategy shifts the current node to its children when applying $r$. In Stratego, a node's children are its proper components, excluding the built-in string, integer and float types.

Using these basic strategies, more complex strategies can be defined. For example, a strategy that catches a failed rewrite with the identity rewrite (and thus always succeeds) can be defined as follows:

**strategies**
$$try \, (r) = r <+ id$$

More interestingly, we can define strategies that traverse the entire tree. For example, the following strategies apply a rewriting rule to all nodes in the tree, as either a top-down (pre-order) or bottom-up (post-order) traversal:

**strategies**
  $alltd\,(r) = r\,;\,all\,(alltd\,(r))$
  $allbu\,(r) = all\,(allbu\,(r))\,;\,r$

We call these *deep* traversal strategies as they recursively descend through the nodes of the tree. In contrast, we call *all* a *shallow* traversal strategy because it descends only to a node's children, and no further.

Strategies can be invoked from within a rule by enclosing them in angled brackets. For example, the right-unit law of addition can be defined by sequencing *commutativity* and *leftUnit*, and applying them to a node:

**rules**
  $rightUnit : e \;\rightarrow\; \langle\, commutativity\,;\, leftUnit\,\rangle\; e$

Using this programming idiom, several rewriting systems have been implemented on top of Stratego and related tools, including Stratego itself, optimisers, static analysers, pretty printers, and frameworks for syntactic language extensions (Erdweg *et al.*, 2011, 2012; Erdweg & Rieger, 2013). We have only given a cursory overview of the essence of Stratego here, and the interested reader is referred to the Stratego literature (Visser *et al.*, 1998; Visser, 2004; Bravenboer *et al.*, 2008).

### 2.3 Selective traversals

Stratego's shallow traversal *all* descends into every child node, and the deep traversals *alltd* and *allbu* descend into every node in the tree. While this is often the required behaviour, it is sometimes desirable to exclude some nodes from a traversal. This could be for semantic reasons: for example, when performing single-variable substitution on a representation of the lambda calculus, a traversal should not descend past a rebinding of the variable being substituted (shadowing). Alternatively, it may be for performance reasons: excluding nodes that are known not to contain any substructures that will be modified by the traversal. This can significantly improve the efficiency of traversals, so much so that entire libraries (e.g. Alloy (Brown & Sampson, 2009); Uniplate (Mitchell & Runciman, 2007)) have been designed around enabling it.

A traversal that descends into some nodes but not others is called a *selective traversal*. Selectivity comes in two forms: *static selectivity*, where it is known at compile-time that certain nodes are never descended into, and *dynamic selectivity*, where whether to descend into a particular node is determined by a run-time predicate. Static selectivity is more efficient than dynamic selectivity, as it allows compiler optimisations to be applied, and avoids the need for the run-time checks.

### 2.4 Approaches to strategic and generic programming

Strategic programming has been implemented in several mainstream languages, including Standard ML (Visser *et al.*, 1998), Java (Balland *et al.*, 2008) and Haskell

(Lämmel & Visser, 2002). The Haskell implementation *Strafunski* was the original inspiration for KURE. Strafunski follows in the tradition of Stratego, but adds types to rewriting rules and strategies, and uses Scrap-Your-Boilerplate (SYB) (Lämmel & Peyton Jones, 2003) generic programming to implement generic traversals over arbitrary data types. We will not discuss the specific features and limitations of Strafunski in this article; instead we will discuss more generally the SYB approach to generic traversals.

Strategic programming in a typed setting resembles datatype-generic programming, to which there are a rich variety of approaches. To date, the most popular language for implementing approaches to generic programming has been Haskell, though there have also been implementations in Scala (Moors *et al.*, 2006; Oliveira & Gibbons, 2010) and Agda (Löh & Magalhães, 2011). In this article we compare only with the approaches to generic programming that are most similar to strategic programming: those that focus on traversing data rather than on reflecting the structure of types.

In Section 6 we extensively compare KURE to the Uniplate (Mitchell & Runciman, 2007) and conventional SYB (Lämmel & Peyton Jones, 2003) approaches. These are well-documented and have the best-maintained implementations of the traversal-centric approaches. Other similar approaches include Smash (Kiselyov, 2006), Alloy (Brown & Sampson 2009), Compos (Bringert & Ranta, 2008), and Multiplate (O'Connor, 2011). We already mentioned Alloy for its emphasis on selective traversal. Multiplate has the most similar representation of traversal to KURE: its central value is a record of rewrites, which is nearly isomorphic to a KURE rewrite over a universe (see Section 3.2). The distinguishing features of Smash and Compos are not relevant to this article, so we will not discuss them here.

Our focus is on generic traversals, but there are also alternative approaches that support functions beyond traversals. The primary line of research explores alternative ways to reflect the structure of types and thereby define extensible and generic functions. We refer to the interested reader to the surveys (Hinze & Löh, 2007; Rodriguez Yakushev *et al.*, 2008) for discussions of most such libraries before 2009. The subsequent work culminates in GHC Generics (Magalhães *et al.*, 2010), which situates the type reflection and extensibility in two corresponding language features: type families and type classes, respectively.

Finally, Van Noort *et al.* (2010) have applied the modern generic foundations specifically to rewriting. Their key innovation compared to other embedded implementations of generic programming is a generic intensional representation of rewriting rules as a *data type*, rather than as a *function*, in the host language. This allows rules to be expressed more concisely, particularly those that contain meta-variables. The data-type representation also makes rewriting rules more amenable to manipulation and analysis than an (opaque) host-language function, albeit at some efficiency cost.

## 2.5 *Related work*

There have been many systems for supporting the translation of internal representations of programs using strategic programming concepts. A number of such systems improve the ability to provide core language constructions, and support for user-visible extensibility. SugarJ (Erdweg *et al.*, 2011), which is built on Stratego, allows libraries to augment Java with syntactical extensions, and includes support for using Stratego rules directly.

SugarHaskell (Erdweg *et al.*, 2012) applies the same ideas but with Haskell as the base language, while Sugar* (Erdweg & Rieger, 2013) generalises the approach to abstract over the base language. With Fortified Macros (Culpepper & Felleisen, 2010; Culpepper, 2012), new syntax can be supported with improved support for error checking and better error messages, using pattern-match–style rewriting rules. KURE as a library has a fundamentally orthogonal concern: KURE is designed for rewriting *internal* syntax that is represented as an algebraic data structure.

Transformation systems have also been used as a basis for impressive language endeavours in the space of mechanised meta theory. For example, the K framework (Lazar *et al.*, 2012) has been used to provide a best-in-class operational semantics for C (Ellison & Roşu, 2012), and PLT Redex (Felleisen *et al.*, 2009) uses transformations as a basis for authoring operational semantics (Klein *et al.*, 2012). KURE has a significantly more modest goal: supporting the rewriting of tree-structured data using techniques from generic programming and design patterns from strategic programming.

# 3 Strongly typed strategic programming

In untyped strategic programming languages such as Stratego, a common source of programming errors is rewriting rules that can produce ill-formed terms (Erdweg *et al.*, 2014). A wide class of such errors can be prevented by using *typed* algebraic data types to represent nodes in the tree, thereby constraining the possible children that each node can have. However, this poses the challenge of how to assign types to rewriting rules and strategies. The main difficulty is assigning types to *generic* traversal strategies that should operate over nodes of *any* type. In this section we consider this challenge, and describe how KURE addresses it with a mixture of static and dynamic typing.

## 3.1 Typing transformations and rewrites

First, we make a distinction between rewriting rules that preserve the type of a node, and rewriting rules that can change the type of a node. We will refer to the former as *rewrites*, and the latter as *transformations*. In a strongly typed setting, a rewrite can be used to modify a node in the tree, whereas a transformation cannot because the resulting tree may not be type correct. However, a transformation can be used to project information from a node, or multiple transformations can be composed to form a rewrite.

Let the type $\mathbf{T}\ \sigma\ \tau$ denote a transformation from a node of type $\sigma$ to a node of type $\tau$. Then let $\mathbf{R}\ \tau$ denote a rewrite over a node of type $\tau$, and let us define $\mathbf{R}$ to be a synonym for the special case of $\mathbf{T}$ where both parameters are the same:

$$\mathbf{R}\ \tau = \mathbf{T}\ \tau\ \tau$$

Using $\mathbf{R}$ and $\mathbf{T}$, it is straightforward to assign types to the following Stratego strategies:

```
id    :: R τ                          -- identity rewrite
fail  :: T σ τ                        -- failing transformation
(;)   :: T ρ σ → T σ τ → T ρ τ        -- sequential composition
(<+ ) :: T σ τ → T σ τ → T σ τ        -- deterministic choice
```

### *3.2 Typing generic traversals*

Now consider assigning a type to Stratego's *all* strategy. If all nodes in the object language have the same type $\tau$, then this is straightforward:

$$all :: \mathbf{R}\ \tau \to \mathbf{R}\ \tau$$

However, if the object language is represented by multiple algebraic data types, then this is inadequate. Consider the situation where the types of some (or all) of the children of a node differ from their parent. In that case, a rewrite of the above type could only apply its argument rewrite to children of the same type as the parent — a significant shortcoming.

For example, consider the following algebraic data types (representing a small language of declarations and expressions):

$$Decl = Id \times Expr$$
$$Expr = Id + Expr \times BinOp \times Expr + Decl \times Expr$$

Given a rewrite $r :: \mathbf{R}\ Expr$, then we could apply *all r* to a node of type *Expr*, and *r* would be applied to all *Expr* children. However, applying *all r* to a *Decl* node would be prohibited as ill-typed, even though we should be able to apply *r* to the *Expr* child.

A possible alternative would be to give *all* a more general type:

$$all :: \mathbf{R}\ \sigma \to \mathbf{R}\ \tau$$

KURE does not use this type however, as it does not relate $\sigma$ and $\tau$ at all, and so *all* cannot use the argument rewrite in any meaningful way — *without run-time type comparisons*. The run-time comparison approach has been taken by Dolstra & Visser (2001); Dolstra (2001), who invent a language that includes a run-time type-case inside traversal combinators. This is also essentially the approach taken by Strafunski using SYB (Lämmel & Visser, 2002), as we will discuss in Section 6.

Ideally, we would like generic traversal strategies to accept a *set* of distinctly typed rewrites as an argument, one for each child. However, this approach would require us to introduce sets of rewrites as a new first-class notion, and then define additional strategies to operate on these sets, and to combine them with our existing **R**s and **T**s. This is possible, indeed it is the approach taken by the Multiplate library (O'Connor, 2011), but it is not the approach that KURE takes. Instead, we reuse the **R** type as the single argument to *all*. The challenge is finding a type $\upsilon$ such that $\mathbf{R}\ \upsilon$ encodes a set of rewrites.

One such type is an algebraic (disjoint) *sum type*. If a node of type $\tau$ has $n$ possible children with types $\tau_1, \tau_2, \ldots, \tau_n$, then we can define the sum type:

$$\upsilon = \tau_1 + \tau_2 + \cdots + \tau_n$$

and the type of *all* could then be:

$$all :: \mathbf{R}\ \upsilon \to \mathbf{R}\ \tau$$

For example, if $\tau$ is *Decl*, then we would have:

$$\upsilon = Id + Expr$$
$$all :: \mathbf{R}\ \upsilon \to \mathbf{R}\ Decl$$

A disadvantage of using a sum type in this way is that it admits some non-type-preserving rewrites. This can be seen if we expand the **R** synonym:

$$\mathbf{R}\ \upsilon = \mathbf{T}\ \upsilon\ \upsilon = \mathbf{T}\ (\tau_1 + \tau_2 + \cdots + \tau_n)\ (\tau_1 + \tau_2 + \cdots + \tau_n)$$

That is, a rewrite could transform a child of one type into a child of another type. To retain type safety with this approach, KURE performs dynamic type checks during generic traversals. Any such type errors are handled by having the rewrite fail. By comparison, Multiplate's set approach has more type precision, but leads to less composable transformations.

That said, the above type for *all* is not yet very composable — for example, the Stratego definitions of *alltd* and *allbu* (Section 2.2) would not type check, as $\upsilon$ and $\tau$ are not the same type. To address this, we generalise from a sum of child types to a sum of the types of all nodes in the tree. We call this generalised sum type a *universe type*. We then modify the type of *all* such that its argument and return types are rewrites over the same universe. For example:

$$\upsilon = Decl + Expr + Id + BinOp$$
$$all :: \mathbf{R} \, \upsilon \to \mathbf{R} \, \upsilon$$

This type is composable, and leads to deep traversals of the same type. For example:

$$alltd :: \mathbf{R} \, \upsilon \to \mathbf{R} \, \upsilon$$
$$alltd \, (r) = r \, ; \, all \, (alltd \, (r))$$

A universe need not contain all node types in the tree; it could contain only a subset of the node types. In this case, *all r* could only traverse nodes whose types inhabit that universe, and could only target children whose types inhabit that universe. That is, the universe acts as an index defining a *statically selective traversal* (Section 2.3). Consequently, *all* is an *ad-hoc polymorphic* (Strachey, 1967) function: it exhibits non-uniform behaviour when instantiated to different universe types.

## 4 KURE implementation

In this section we describe how the ideas in Section 3 are implemented as a Haskell-embedded domain-specific language. We begin by describing the main aspects of KURE that are required by most use cases: side-effecting rewrites (Section 4.1); strategies (Section 4.2); and generic traversals using a universe (Section 4.3.1). We then discuss more advanced aspects of KURE that may only be needed for some use cases: traversals over multiple universes (Section 4.3.2); traversals that distinguish nodes based on their location rather than their type (Section 4.3.3); support for maintaining a context during generic traversals (Section 4.4); and strategies and traversals involving change detection (Section 4.5).

### 4.1 Transformations, rewrites and side-effects

The central implementation decision is how to represent the **T** type. Haskell is a purely functional language, so simply defining **T** as a function,

**type** $T \, a \, b = a \to b$

would be inadequate, as it does not allow a transformation to perform any side effects. Side effects are an inherent part of strategic programming, most notably the side effects of

```
class Monad m where                              pure :: Monad m ⇒ a → m a
  return :: a → m a                              pure = return
  (>>=) :: m a → (a → m b) → m b                 (<$>) :: Monad m ⇒ (a → b) → m a → m b
  fail   :: String → m a                         f <$> ma = ma >>= λa → return (f a)
                                                 (<*>) :: Monad m ⇒ m (a → b) → m a → m b
                                                 mf <*> ma = mf >>= λf → (f <$> ma)
```

Fig. 1. Monads in Haskell, and some derived functions.

*failure*, and *catching* failure. Additionally, specific use cases may need to support arbitrary side-effecting operations, such as fresh name generation.

In Haskell, the standard way of encoding side effects is to use a *monadic* (Wadler, 1992) structure. Furthermore, the KURE library needs to be applicable to different use cases, with arbitrary side effects, so the library design should not commit to any specific set of effects. Therefore, rather than using a concrete monad, transformations are parameterised over an arbitrary monad. This leads to the following implementation of **T**, where *m* is the monad parameter[1]:

**type** $T\ m\ a\ b = a → m\ b$

As before, we will define rewrites as a synonym for a special case of transformations:

**type** $R\ m\ a = T\ m\ a\ a$

Note that in Haskell, the constraint that the parameter *m* must be a monad is specified on the operations over *T*, rather than on the *T* type itself.

However, an arbitrary monad does not support catching failure, and indeed it is somewhat of a historical accident that Haskell's *Monad* type class (Figure 1) even supports throwing failures. As catching failure is necessary to encode the deterministic-choice strategy, KURE provides a subclass of *Monad* that provides a catching operation:

**class** $Monad\ m ⇒ MonadCatch\ m$ **where**
  $catchM :: m\ a → (String → m\ a) → m\ a$

The purpose of the *String* argument in *fail* and *catchM* is to pass error messages; this allows KURE strategies to provide informative error messages for failing transformations. As error messages are not the topic of this article, we will not discuss them further.

Using the function space of the host language is not the only way to represent transformations; an alternative encoding is as pairs of expressions, augmented with some additional infrastructure to represent meta-variables. Most embedded implementations of strategic/generic programming take the former approach, but Van Noort *et al.* (2010) have shown that the latter approach can be taken without any additional boilerplate burden on the user.

Finally, we note that while an implementation of *T* as a (partial) function may appear to restrict us to deterministic transformations, the monadic structure allows non-deterministic

---

[1] Actually, *T* needs to be a **newtype**, rather than a type synonym, because type synonyms cannot be partially applied and to avoid ambiguous type class instances. However, we elide that detail in this article to avoid the syntactic clutter introduced by the data constructor.

**instance** $Monad\ m \Rightarrow Monad\ (T\ m\ a)$ **where**
$\quad return :: b \rightarrow T\ m\ a\ b$
$\quad return\ b = \lambda\_ \rightarrow return\ b$
$\quad fail :: String \rightarrow T\ m\ a\ b$
$\quad fail\ s = \lambda\_ \rightarrow fail\ s$
$\quad (>\!\!>\!\!=) :: T\ m\ a\ b \rightarrow (b \rightarrow T\ m\ a\ d) \rightarrow T\ m\ a\ d$
$\quad t >\!\!>\!\!= f = \lambda a \rightarrow (t\ a >\!\!>\!\!= \lambda b \rightarrow (f\ b)\ a)$
**instance** $MonadCatch\ m \Rightarrow MonadCatch\ (T\ m\ a)$ **where**
$\quad catchM :: T\ m\ a\ b \rightarrow (String \rightarrow T\ m\ a\ b) \rightarrow T\ m\ a\ b$
$\quad catchM\ t_1\ t_2 = \lambda a \rightarrow catchM\ (t_1\ a)\ (\lambda s \rightarrow (t_2\ s)\ a)$

Fig. 2. The *Monad* and *MonadCatch* type class instances for *T*.

transformations to be encoded, either by using a monad that supports pseudo-random choice, or a list-like monad that computes multiple results.

### 4.2 Strategies

We now need to implement strategies over *T*. The identity and sequential-composition strategies, which in KURE we denote as *idR* and $>\!\!>\!\!>$, can be defined straightforwardly as the identity and composition operations in the Kleisli category:

$idR :: Monad\ m \Rightarrow R\ m\ a$
$idR = return$
$(>\!\!>\!\!>) :: Monad\ m \Rightarrow T\ m\ a\ b \rightarrow T\ m\ b\ d \rightarrow T\ m\ a\ d$
$t_1 >\!\!>\!\!> t_2 = \lambda a \rightarrow (t_1\ a >\!\!>\!\!= t_2)$

The *T* type itself forms a monad, as defined in Figure 2. The definition is fairly standard: *T m a* forms a monad that corresponds to applying the Reader monad transformer (Liang *et al.*, 1995) to the underlying monad *m*, where the environment is the argument node *a*. This monadic structure corresponds to a strategy that applies multiple transformations to the same node, with the choice of subsequent transformations depending on the results of prior transformations. We will give some examples of this in Section 5.4.

Failure and catching infrastructure are provided by the *fail* and *catchM* methods of the *Monad* and *MonadCatch* classes. Exact analogues of Stratego's *fail* and $<\!\!+$ strategies (which do not pass error messages) can then be easily defined:

$failT :: Monad\ m \Rightarrow T\ m\ a\ b$
$failT = fail\ \texttt{"failT"}$
$(<\!\!+) :: MonadCatch\ m \Rightarrow T\ m\ a\ b \rightarrow T\ m\ a\ b \rightarrow T\ m\ a\ b$
$t_1 <\!\!+ t_2 = catchM\ t_1\ (\lambda\_ \rightarrow t_2)$

New strategies can then be constructed in the same manner as in Stratego, for example:

$tryR :: MonadCatch\ m \Rightarrow R\ m\ a \rightarrow R\ m\ a$
$tryR\ r = r <\!\!+ idR$

Note that the type signatures can be entirely inferred by the Glasgow Haskell Compiler, and so can be omitted by a KURE user if desired.

### 4.3 Universe-indexed generic traversals

We now turn to the main implementation challenge: implementing generic traversals using universe types. This subsection is divided into three parts:

- Section 4.3.1 demonstrates the main idea, and gives an example of standard KURE usage using a single universe that treats all nodes of the same type equally.
- Section 4.3.2 explains how to support statically selective traversals by using multiple universes.
- Section 4.3.3 describes how to define traversals that distinguish between nodes of the same type, based on their location in the tree.

#### 4.3.1 Walking the tree

Recall that the key idea from Section 3.2 is for *all* to be ad-hoc polymorphic in its universe type. KURE implements this using the following type class, where $u$ is the universe type:

```
class Walker u where
    allR :: MonadCatch m ⇒ R m u → R m u
```

Notice that the monad $m$ is not a class parameter — instead the *allR* method is *parametrically polymorphic* (Strachey, 1967) in any *MonadCatch m*. This choice brings expressiveness benefits, as we will explain in Section 4.5.

The KURE library provides a suite of traversal strategies defined using *allR*. For example, Stratego's *alltd* strategy is defined as follows:

```
alltdR :: (Walker u, MonadCatch m) ⇒ R m u → R m u
alltdR r = r >>> allR (alltdR r)
```

Providing the *Walker* instance for each universe is the responsibility of the KURE user, and this is how the user customises the traversal behaviour for each universe. To demonstrate how a KURE user would define a universe and *Walker* instance, we will use the following representation of an object language as a running example. To avoid confusion between KURE implementation code and user code, we enclose the latter in boxes.

```
data Prog = PrgCons Decl Prog | PrgNil
data Expr = Lit Literal | Var Name | Append Expr Expr | Let Decl Expr
data Decl = Decl Name Expr
type Name   = String
type Literal = String
```

This models a small expression language with non-recursive let-bindings for building strings. We deliberately designed this example to use strings in two semantically different ways (as variable names and as literals), so that we can later demonstrate how KURE can distinguish between the two during traversals.

The first task for the KURE user is to define a universe. Let us assume she wants to traverse programs, declarations and expressions, but not variable names or string literals. That is, she wants to define a statically selective traversal that omits strings. (In Haskell,

$$injectT :: (Monad\ m, Injection\ a\ u) \Rightarrow T\ m\ a\ u$$
$$injectT = \lambda a \rightarrow return\ (inject\ a)$$
$$projectT :: (Monad\ m, Injection\ a\ u) \Rightarrow T\ m\ u\ a$$
$$projectT = \lambda u \rightarrow \textbf{case}\ project\ u\ \textbf{of}$$
$$\qquad\qquad\qquad\qquad \mathsf{Just}\ a\ \ \rightarrow return\ a$$
$$\qquad\qquad\qquad\qquad \mathsf{Nothing} \rightarrow fail\ \texttt{"projectT"}$$
$$promoteR :: (Monad\ m, Injection\ a\ u) \Rightarrow R\ m\ a \rightarrow R\ m\ u$$
$$promoteR\ r = projectT >>> r >>> injectT$$
$$extractR :: (Monad\ m, Injection\ a\ u) \Rightarrow R\ m\ u \rightarrow R\ m\ a$$
$$extractR\ r = injectT >>> r >>> projectT$$

Fig. 3. Promotion and extraction combinators.

*String* is implemented as a synonym for a singly linked list of characters.) She begins by defining a universe consisting of the three node types she wishes to traverse:

> **data** $U = \mathsf{UP}\ Prog\ |\ \mathsf{UD}\ Decl\ |\ \mathsf{UE}\ Expr$

A common need when working with universe types is to inject nodes into a universe, or project them out of a universe. To aid this, KURE provides a type class *Injection*,

> **class** *Injection a u* **where**
> $\quad inject\ \ :: a \rightarrow u$
> $\quad project :: u \rightarrow Maybe\ a$

along with several conversion functions for converting transformations and rewrites to and from universe types (Figure 3). The important point about these conversion functions is that a failed projection will result in a failing transformation. Consequently, a rewrite over a universe type that is applied to a node will fail if it does not preserve the node type (this provides the dynamic type checking mentioned in Section 3.2). Injection instances are straightforward to define, as exemplified by the following instance for *Expr*:

> **instance** *Injection Expr U* **where**
> $\quad inject = \mathsf{UE}$
> $\quad project\ (\mathsf{UE}\ expr) = \mathsf{Just}\ expr$
> $\quad project\ \_\qquad\quad = \mathsf{Nothing}$

Having declared *Injection* instances for a universe, the user can then declare a *Walker* instance. This requires the user to decide which substructures should be considered child nodes, but the definition of *allR* is otherwise systematic. The essence of the definition is as follows. First, deconstruct the node by pattern matching. Then, for each child node, extract a specialised rewrite from the argument rewrite (which operates over the universe type), and apply that rewrite to child. Finally, reconstruct the node by reapplying the constructor(s). As an example, we present the *Walker* instance for the *U* universe in Figure 4.

Observe that the traversal does not descend below variable or literal nodes. These are the non-systematic parts of the definition that had to be decided by the user. These user decisions prevent this definition of *allR* from being generated automatically. Additionally, we will later (Section 4.4) augment the *Walker* class to support maintaining a context, which will also interfere with automatic generation of *allR* definitions.

```
instance Walker U where
    allR :: MonadCatch m ⇒ R m U → R m U
    allR r u = case u of
                UP p → UP    <$> allProg p
                UD d → UD    <$> allDecl d
                UE e → UE    <$> allExpr e
        where
        allProg :: R m Prog
        allProg PrgNil          = pure PrgNil
        allProg (PrgCons d p)   = PrgCons <$> (extractR r) d <*> (extractR r) p
        allDecl :: R m Decl
        allDecl (Decl n e)      = Decl <$> pure n <*> (extractR r) e
        allExpr :: R m Expr
        allExpr (Var n)         = Var    <$> pure n
        allExpr (Lit l)         = Lit    <$> pure l
        allExpr (Append e₀ e₁)  = Append <$> (extractR r) e₀ <*> (extractR r) e₁
        allExpr (Let d e)        = Let    <$> (extractR r) d  <*> (extractR r) e
```

Fig. 4. An example of a typical *Walker* instance.

However, it is possible to generate a *Walker* instance for a particular formulaic traversal. Indeed, an earlier version of KURE (Gill, 2009) used Template Haskell (Sheard & Peyton Jones, 2002) to generate both universe types and corresponding *Walker* instances. Those instances treated every proper component of a node as a child, and performed no context updates. For our example string language, that would mean treating the *String* representation of variables and literals as the sole child of a variable or literal node, and the head and tail of every *String* would be children of that *String*. Thus, for example, deep traversals would visit every character of every variable name and string literal.

Ideally, the user would be able to take a hybrid approach of mechanically deriving the systematic cases, while manually overwriting the non-systematic cases. Investigating support for this sort of partial derivation is ongoing work (Frisby *et al.*, 2012).

### 4.3.2 Alternative universes

We now consider using multiple universes to provide multiple statically selective traversals over the same object language. As an example, consider the situation where a user wants a traversal that descends into strings, without losing the existing capability to perform traversals that do not descend into strings. She begins by defining a new universe,

$$\textbf{data } U_1 = \mathsf{UP}_1 \; Prog \mid \mathsf{UD}_1 \; Decl \mid \mathsf{UE}_1 \; Expr \mid \mathsf{US}_1 \; String$$

and declaring corresponding *Injection* instances. She then defines a *Walker* instance for $U_1$. This is mostly the same as the instance for $U$, but with three additions to *allR*. First, there is an additional local function *allString*:

```
allString :: R m String
allString []        = pure []
allString (x : xs) = (:) <$> pure x <*> (extractR r) xs
```

Second, the Var and Lit clauses of *allExpr* are modified such that the traversal descends into the string:

> *allExpr* (Var *n*) = Var <$> (*extractR r*) *n*
> *allExpr* (Lit *l*) = Lit <$> (*extractR r*) *l*

Third, there is an extra case alternative when pattern matching on the universe:

> $US_1$ *s* → $US_1$ <$> *allString s*

This may seem like a lot of code duplication, but in practice *allR* definitions operating over the same object language can be built out of a set of reusable *congruence combinators* (Visser, 2004), which factor out the duplication. We will demonstrate this in Section 5.2.

When performing a traversal, the user now has the option of which universe to use; that is, she can choose between statically selective traversals. We will give some examples of this in Sections 4.4 and 5.4.

### 4.3.3 Locations not types

Thus far we have treated all nodes of the same type in the same manner. However, a KURE traversal can distinguish between nodes based on their location, even if their types are the same. For example, imagine a user desired the $U_1$ traversal to descend into string literals but not into variables names, despite them both having type *String*. This could be achieved by changing the Var clause of *allExpr* to the following:

> *allExpr* (Var *n*) = Var <$> *pure n*

Alternatively, if a user wished to support both behaviours, she could define a new universe with a *Walker* instance defined in this way. She could also define a universe for traversals that descend into variables but not literals.

It is also possible to define a traversal that descends into multiple nodes of the same type, yet treats them distinctly based on their location. As this is not quite so straightforward, we defer the details to the Appendix.

### 4.4 Maintaining a context

It is often useful for transformations to have access to a *context*, and for this context to be updated automatically during generic traversals. The motivating example is storing the definitions of variable bindings in scope (when working with an object language that has variable bindings), to support local *fold/unfold* (Burstall & Darlington, 1977) transformations. One way to encode a context would be to constrain *m* to be a Reader monad (Liang *et al.*, 1995), but, for subtle reasons that we will discuss in Section 4.5, it is beneficial to keep the context separate from the monad. Instead, we make the context type (*c*) an additional parameter of *T*, and make the context available as an argument to each transformation:

**type** *T c m a b* = *c* → *a* → *m b*

The strategy definitions from Section 4.2 need to be updated accordingly, but the changes are fairly minor. The changes correspond to applying (another) Reader monad transformation (Liang *et al.*, 1995), where the additional environment is the context *c*.

As an example, we will consider adding a context to our string expression language. Let us assume the desired context is a list of all declarations currently in scope:

**type** *Context* = [*Decl*]

Having these declarations available allows, for example, variables to be inlined locally. A rewrite to perform such inlining can be defined as follows[2]

```
inline :: Monad m ⇒ R Context m Expr
inline = λc e → case e of
                  Var n → lookupName n c
                  _     → fail "not a Var"
      where
        lookupName :: Monad m ⇒ Name → Context → m Expr
        lookupName _ []            = fail "Name not found"
        lookupName n (Decl n' e : c) = if n == n'
                                       then return e
                                       else  lookupName n c
```

We now need to consider how to correctly maintain the context during generic traversals. The key idea is to define all generic traversals in terms of *allR*. If we can ensure that *allR* performs all desired context updates, then all generic traversals will update the context. Furthermore, the definition of *allR* will be the *only* place that the context is modified — all other strategy definitions will only read the context.

To enable this, we add the context to the *Walker* class as an additional class parameter:

```
class Walker c u where
   allR :: MonadCatch m ⇒ R c m u → R c m u
```

It is important that, unlike the monad *m*, the context be a class parameter rather than parameter of the *allR* method. This is because we need to instantiate the context to a concrete data type so that we can modify it during the generic traversal.

As an example, consider the following representative fragment of the modified *Walker* instance for the *U* universe (from Section 4.3.1):

```
instance Walker Context U where
   allR :: MonadCatch m ⇒ R Context m U → R Context m U
   ...
      allExpr c (Append e_0 e_1) = Append <$> (extractR r) c e_0 <*> (extractR r) c e_1
      allExpr c (Let d e)        = Let     <$> (extractR r) c d  <*> (extractR r) (d : c) e
   ...
```

Observe that the context is passed as an argument to the rewrites being applied to the children, and that in the case of the let-body, the context is updated by adding the

---

[2] We assume no variable shadowing for simplicity.

let-bound declaration. We reiterate that because this choice of context update is made by the user, it is not possible to mechanically derive this *Walker* instance.

It is now possible to define deep generic traversals that make use of the context. For example, the following traversal will inline all variable occurrences in an expression:

$$inlineAll = alltdR\ (tryR\ (promoteR\ inline))$$

(The use of *promoteR* is necessary because *inline* is a rewrite over *Expr*, whereas *alltdR* expects a rewrite over a universe as its argument.)

As *inlineAll* has not been ascribed a type signature, its type is inferred to be polymorphic in its universe type. The traversal can thus be instantiated to any universe containing *Expr*; that is, the user can choose between statically selective traversals. Variables only occur as expressions, not within literals or names, so it would be more efficient to instantiate *inlineAll* to the $U$ universe than the $U_1$ universe.

### 4.5 Change detection and monad transformers

When defining a strategy, it can be useful to be able to determine not just whether a rewrite succeeded, but whether it actually modified the node it was applied to. To enable this, KURE supports an (optional) convention whereby rewrites that do not modify the node should fail, and that repeated application of a rewrite should fail after a finite number of iterations. When following this convention, a rewrite that can result in an identity rewrite (such as *tryR*) may be used, but only as a sub-component of a rewrite that is guaranteed to either make a change or fail. For example, the following *repeatR* strategy recursively applies a rewrite until it fails, returning the result before the failure:

$repeatR :: MonadCatch\ m \Rightarrow R\ c\ m\ a \rightarrow R\ c\ m\ a$
$repeatR\ r = r >>> tryR\ (repeatR\ r)$

To further support this convention, KURE provides a monad transformer that can be used to convert a rewrite defined for an arbitrary monad into a rewrite that succeeds if *at least one* sub-rewrite succeeds, converting any failures into identity rewrites. We omit the implementation details, but the main idea is to use a writer monad transformer over the monoid of Boolean disjunction, with the Boolean value representing whether at least one sub-rewrite has succeeded. This allows us to derive from *allR* traversals such as

$anytdR :: (Walker\ c\ u, MonadCatch\ m) \Rightarrow R\ c\ m\ u \rightarrow R\ c\ m\ u$

which is a variant of *alltdR* that succeeds if the argument rewrite succeeds at *any* node. Note that this differs from *alltdR* (*tryR r*), which will always succeed.

For this monad-transformer technique to work in combination with generic traversal strategies, it is essential that the *allR* class method be parametrically polymorphic in its monad parameter. If the monad was a class parameter, then the *Walker* instance would be specialised to that particular monad, and could not be used with a transformed monad. This is the principal reason why the monad is not a class parameter, and, consequently, why the monad and context are kept separate.

This monad-transformer technique is taken from the SYB library, and KURE uses several transformers to define a suite of generic traversal combinators. These derived traversals include transformations as well as rewrites. For example,

$$collectT :: (Walker\ c\ u, MonadCatch\ m) \Rightarrow T\ c\ m\ u\ b \rightarrow T\ c\ m\ u\ [b]$$

applies its argument transformation to every node in the tree, ignoring failures and collecting the results of successes in a list. Because *collectT* in defined in terms of *allR*, each application of the argument transformation to a sub-node will occur in an appropriately updated context. Without transforming the monad, it would not be possible to define *collectT* in terms of *allR*, and thus the user would be required to define an additional generic traversal from which *collectT could* be defined. Such an additional traversal would also need to include context updates in its definition, as it would be independent of *allR*.

## 5 Case study: HERMIT

The principal use-case of KURE to date has been in the HERMIT system (Farmer *et al.*, 2012; Sculthorpe *et al.*, 2013). Indeed, much of the development of KURE has been driven by the needs of HERMIT. As this is a "realistic" use of KURE, we will now present HERMIT as a case study of KURE usage. During the case study we will also introduce *congruence combinators*: a technique to allow code reuse when defining generic traversals over multiple universes.

HERMIT is an experimental plugin for the Glasgow Haskell Compiler (GHC) that provides a toolkit to support user-guided transformations during compilation. Haskell programs are represented as abstract syntax trees inside GHC, and HERMIT makes extensive use of KURE to rewrite these abstract syntax trees.

HERMIT and GHC are systems undergoing active development, and the details of their implementation, and the use of KURE, continue to fluctuate. Consequently, we will not attempt to present the exact HERMIT source code in this section. Instead, we will give a simplified presentation that demonstrates the key points of KURE usage in this project. As previously, we will enclose in boxes any code that is not provided by the KURE library.

### 5.1 GHC Core universes

HERMIT operates on GHC's internal intermediate language, GHC Core. GHC Core is an implementation of System $F_C$, which is System F (Girard, 1972; Reynolds, 1974) extended with let-binding, constructors, and first-class type equalities (Sulzmann *et al.*, 2007). KURE is used for rewriting nodes in the GHC Core abstract syntax tree, and for projecting information from nodes using transformations.

A representative subset of GHC Core is shown in Figure 5. The majority of HERMIT's transformations and rewrites act on *Program*, *Bind*, *Def*, *Expr* and *Alt* nodes, while only a few act on *Coercion* and *Type* nodes. Observe that *Bind*, *Def*, *Expr* and *Alt* nodes are mutually recursive, but that neither they nor *Program* nodes appear as substructures of *Coercion* or *Type* nodes. Consequently, a statically selective traversal that omits *Coercion* and *Type* nodes will not change the semantics of a traversal that only acts on the former set of nodes, but will yield significant efficiency gains by avoiding unnecessarily traversing *Coercion*

```
type Program = [Bind]
data Bind     = NonRec Var Expr | Rec [Def]
data Def      = Def Var Expr
data Expr     = Var Var | Lit Literal
              | App Expr Expr | Lam Var Expr
              | Let Bind Expr | Case Expr Var Type [Alt]
              | Cast Expr Coercion | Type Type
type Alt      = (AltCon, [Var], Expr)
data Coercion = CoVar Var | Refl Type | Sym Coercion | Trans Coercion Coercion | ...
data Type     = TyVar Var | Fun Type Type | ForAll Var Type | ...
```

Fig. 5.  A representative subset of GHC Core.

and *Type* nodes. To support this, HERMIT defines two universes, one that includes types and coercions (*CoreTC*), and one that does not (*Core*):

```
data Core   = CProg Program | CBind Bind | CDef Def | CExpr Expr | CAlt Alt
data CoreTC = CCoercion Coercion | CType Type | CCore Core
```

### 5.2  Congruence combinators

Rather than define *Walker* instances for our two universes directly, we shall first define *congruence combinators* (Visser, 2004) for the nodes in our tree. Congruence combinators provide an abstract way of traversing a node that factors out the commonality of different statically selective traversals. We will then define the *Walker* instances in terms of the congruence combinators, thereby avoiding code duplication.

Each congruence combinator is specialised to a single node constructor, and takes as arguments transformations to apply to each of the node's children, and a function to combine the results. For example, consider the following congruence combinator for App nodes:

```
appT :: Monad m ⇒ T c m Expr a₁ → T c m Expr a₂ → (a₁ → a₂ → b) → T c m Expr b
appT t₁ t₂ f = λc expr → case expr of
                         App e₁ e₂ → f <$> t₁ c e₁ <*> t₂ c e₂
                         _         → fail "not an App"
```

That is, an App node should be traversed by applying the argument transformations $t_1$ and $t_2$ to the child nodes $e_1$ and $e_2$, and then combining the results by mapping the function $f$ over them. If the node is not an App, then the transformation should fail.

A useful specialisation of a congruence combinator is when the argument transformations are rewrites, and the function to combine the results is the node constructor:

```
appR :: Monad m ⇒ R c m Expr → R c m Expr → R c m Expr
appR r₁ r₂ = appT r₁ r₂ App
```

This is essentially a version of *allR* specialised to App nodes, a point that will be important later. Note that because the number and type of its child nodes are known, *appR* can take

one argument rewrite for each child, with each rewrite matching the type of the child, rather than taking a single rewrite over the universe type as is done by *allR*.

The *appT* congruence combinator does not modify the context, and hence is polymorphic in its context parameter *c*. However, in general, if traversing the node should cause a context update, then the congruence combinator for that node should update the context accordingly. For example, the HERMIT context (*HermitC*) stores the set of variable bindings in scope (among other things). Thus, a congruence combinator for Let nodes needs to add the bindings to the context when traversing into the let-body:

$$letT :: Monad\ m \Rightarrow T\ HermitC\ m\ Bind\ a_1 \rightarrow T\ HermitC\ m\ Expr\ a_2 \rightarrow (a_1 \rightarrow a_2 \rightarrow b)$$
$$\rightarrow T\ HermitC\ m\ Expr\ b$$
$$letT\ t_1\ t_2\ f = \lambda c\ expr \rightarrow \textbf{case}\ expr\ \textbf{of}$$
$$\qquad\qquad Let\ bds\ e \rightarrow f <\$>\ t_1\ c\ bds <\!*\!>\ t_2\ (addBindings\ bds\ c)\ e$$
$$\qquad\qquad \_ \qquad\qquad \rightarrow fail\ \texttt{"not a Let"}$$
$$addBindings :: Bind \rightarrow HermitC \rightarrow HermitC$$
$$\text{-- definition not given}$$

While congruence combinator definitions are fairly systematic, they do require some semantic knowledge of the object language. For example, consider the top-level *Program* type, which is a list of bindings. Each *Bind* in the list could be considered a child of a *Program* node. However, another interpretation could be that a (non-empty) list has exactly two children: the binding group at the head of the list, and another program making up the tail of the list. The latter interpretation is the one HERMIT uses, as that matches the scoping behaviour of binding groups in GHC Core.

Defining congruence combinators for all nodes that we wish to traverse may seem like a lot of work, but our experience with HERMIT has been that there is a significant pay-off in the simplicity with which subsequent transformations can be defined. Furthermore, this allows the context updates to be localised to one place. Any *Walker* instances that require the same context updates can be defined using these congruence combinators (see Section 5.3), and any node-specific transformations can also use the congruence combinators to apply transformations to descendant nodes — in both cases the context will be correctly updated while doing so. This localisation makes it less likely that any contextual updates will be accidentally omitted by the programmer, and also aids code maintenance.

Finally, we note that the notion of congruence combinators is not specific to HERMIT or KURE, and can be used in any strategic programming language to express shallow traversals over nodes in any object language.

### 5.3 Walker instances

Once a complete set of congruence combinators and *Injection* instances have been defined for the nodes of the object language, then defining *Walker* instances is systematic. The instances for the *Core* and *CoreTC* universes are presented in Figure 6. The key point is that each congruence combinator is given *extractR r* as an argument for each child node that should be traversed, and *idR* as an argument for each child node that should be omitted. The main difference between the two instances is in *allRexpr*, as several of the *Expr* nodes have *Type* or *Coercion* children (you may find it helpful to refer back to Figure 5). Additionally,

```
instance Walker HermitC Core where
  allR :: MonadCatch m ⇒ R HermitC m Core → R HermitC m Core
  allR r = promoteR allRprog <+ promoteR allRbind <+ promoteR allRdef
        <+ promoteR allRalt   <+ promoteR allRexpr
    where
      allRprog :: R HermitC m Program
      allRprog = progNilR <+ progConsR (extractR r) (extractR r)
      allRbind :: R HermitC m Bind
      allRbind = nonRecR (extractR r) <+ recR (extractR r)
      allRdef  :: R HermitC m Def
      allRdef  = defR idR (extractR r)
      allRalt  :: R HermitC m Alt
      allRalt  = altR idR idR (extractR r)
      allRexpr :: R HermitC m Expr
      allRexpr = varR idR <+ litR idR <+ appR (extractR r) (extractR r)
              <+ lamR idR (extractR r) <+ letR (extractR r) (extractR r)
              <+ caseR (extractR r) idR idR (extractR r)
              <+ castR (extractR r) idR <+ typeR idR

instance Walker HermitC CoreTC where
  allR :: MonadCatch m ⇒ R HermitC m CoreTC → R HermitC m CoreTC
  allR r = promoteR allRprog       <+ promoteR allRbind <+ promoteR allRdef
        <+ promoteR allRalt        <+ promoteR allRexpr
        <+ promoteR allRcoercion <+ promoteR allRtype
    where
      allRprog     :: R HermitC m Program
      allRprog     = progNilR <+ progConsR (extractR r) (extractR r)
      allRbind     :: R HermitC m Bind
      allRbind     = nonRecR (extractR r) <+ recR (extractR r)
      allRdef      :: R HermitC m Def
      allRdef      = defR idR (extractR r)
      allRalt      :: R HermitC m Alt
      allRalt      = altR idR idR (extractR r)
      allRexpr     :: R HermitC m Expr
      allRexpr     = varR idR <+ litR idR <+ appR (extractR r) (extractR r)
                  <+ lamR idR (extractR r) <+ letR (extractR r) (extractR r)
                  <+ caseR (extractR r) idR (extractR r) (extractR r)
                  <+ castR (extractR r) (extractR r) <+ typeR (extractR r)
      allRcoercion :: R HermitC m Coercion
      allRcoercion = coVarR idR <+ reflR (extractR r) <+ symR (extractR r)
                  <+ transR (extractR r) (extractR r)
      allRtype     :: R HermitC m Type
      allRtype     = tyVarR idR <+ FunR (extractR r) (extractR r)
                  <+ forAllR idR (extractR r)
```

Fig. 6. *Walker* instances for the *Core* and *CoreTC* universes.

the *CoreTC* instance has functions for traversing *Type* and *Coercion* nodes, which were not needed in the *Core* case. In both cases, *Var*, *Literal* and *AltCon* nodes are not traversed.

  Notice that having defined the context updates within the congruence combinators, we do not need to update the context within each *Walker* instance. Furthermore, if HERMIT

later adds more universes, it will be possible to define the corresponding *Walker* instances using the existing congruence combinators.

### 5.4 Example rewrites and transformations

We will now give some examples of HERMIT rewrites and transformations. We begin with a rewrite that converts the application of a lambda to a non-recursive let binding (a form of $\beta$-reduction that preserves sharing). This can be defined succinctly as follows:

```
appLamToLet :: Monad m ⇒ R c m Expr
appLamToLet = do App (Lam v e₂) e₁ ← idR
                 return (Let (NonRec v e₁) e₂)
```

This definition exploits Haskell's monadic **do**-notation to implicitly handle failure: if the argument node does not match the pattern App (Lam $v$ $e_2$ $e_1$) then this rewrite will fail, invoking the *fail* method of the underlying monad $m$.

To apply this rewrite throughout the tree, we promote it to operate on the *Core* universe, and apply the *anytdR* traversal strategy:

```
appLamsToLets :: (Walker c Core, MonadCatch m) ⇒ R c m Core
appLamsToLets = anytdR (promoteR appLamToLet)
```

This rewrite will convert any applications of lambdas in the tree to let expressions, and will succeed if there is at least one such conversion. We use the *Core* universe rather than the *CoreTC* universe for efficiency, as there are no lambdas within *Type* or *Coercion* nodes.

Thus far, the only monadic side effects used in this article have been failure, and catching failure. As an example of an application-specific effect, some HERMIT rewrites need to be able to generate globally fresh variables. HERMIT defines its own concrete monad (*HermitM*), which provides an operation *newVar* :: *HermitM Var*, among others. This is useful when, for example, defining a rewrite that introduces a let binding:

```
letIntro :: R c HermitM Expr
letIntro = λ _ e → do v ← newVar
                      return (Let (NonRec v e) (Var v))
```

As a final example, consider the task of collecting all variable occurrences in the tree. This can be achieved using the library traversal *collectT* (from Section 4.5), in combination with a user-defined local transformation (*varOcc*) that succeeds only if the current node is a variable, returning that variable.

```
-- definition not given
varOcc :: T HermitC m CoreTC Var
collectVarOccs :: MonadCatch m ⇒ T HermitC m CoreTC [Var]
collectVarOccs = collectT varOcc
```

We choose the *CoreTC* universe because GHC Core contains type variables and coercion variables (beyond just value variables).

## 5.5 *User experience*

Many of the features of KURE were motivated by the needs of HERMIT, including support for: statically selective traversals; rewriting nodes of different types during the same traversal; automatic maintenance of a context during generic traversals; and the ability to detect when a successful rewrite actually modified its target. This combination of features is not provided by any other library for strategic or generic programming.

KURE transformations and strategies permeate the HERMIT code base, with the majority of user-facing commands being implemented by an underlying KURE transformation or rewrite. Overall, our experience (and that of the other HERMIT developers) of using KURE has been extremely positive. The library strategies saved a substantial amount of implementation effort, and the structured handling of context and monadic effects (especially failure) has helped avoid errors and made code refactoring easier.

Of particular note is that KURE's universe types have proved especially amenable to design changes. The initial HERMIT implementation (Farmer *et al.*, 2012) only traversed the *Core* universe, which was a deliberate design decision to avoid the inefficiency of traversing types and coercions. However, as HERMIT developed we began to need to apply transformations within types and coercions. The ability to define a new universe (*CoreTC*) after much infrastructure was already in place, and to do so using the existing congruence combinators, was immensely convenient.

Regarding programming style, we found that congruence combinators tend to be more useful than the *shallow* generic traversal strategies. There were no cases where we wanted to perform a shallow traversal without knowing the identity of the node constructor, and hence we were always able to use a (more precise) congruence combinator. However, HERMIT does make extensive use of *deep* generic traversal strategies.

Looking beyond HERMIT, the largest example of KURE use is the `html-kure` package (Gill, 2013), which layers the KURE interface on top of the HXT HTML parser and pretty printer (Schmidt *et al.*, 2012). Rewriting HTML is a task for which KURE was not specifically designed, so that this was straightforward provides evidence that KURE is more generally useful (indeed, we use this package to preprocess our research group's website). However, KURE needs to be used for many more applications before we are in a position to fairly assess its ease of use compared to other strategic rewriting systems.

## 6 Comparison of KURE to SYB, Uniplate and Stratego

In this section we compare and contrast KURE's approach to generic traversals with three other approaches: Stratego, Scrap Your Boilerplate (SYB) (Lämmel & Peyton Jones, 2003) and Uniplate (Mitchell & Runciman, 2007). The focus of our comparison is the differing choices that each approach makes as to which substructures of a node are considered "children", and how, given those choices, the generic traversal combinators of each approach are typed. We also consider the consequences of these choices regarding modularity and statically selective traversals.

### 6.1 Identifying children

In Stratego, the children of a node are its proper components, excluding the built-in string, integer and float types. SYB similarly takes the children of a node to be its proper components, but it does not exclude any types. Instead, primitive types are considered to be nodes with zero children.

In KURE, the user specifies which substructures of a node should be considered children. Typically, these children are zero or more independent proper substructures of the node (but may also be abstract *retractions* of the substructures, which allows different substructures of the same type to be treated distinctly).

In Uniplate, the children of a node are determined by type. Specifically, the children of a node of type $\tau$ are the maximal proper substructures of that same type $\tau$. Uniplate also provides an extension called Biplate, in which an additional type index $\sigma$ is used. The children of a Biplate node of type $\tau$ are the maximal substructures of that type $\sigma$. This use of an ad-hoc polymorphic type to index a family of traversals has similarities with KURE's use of a universe type to index a family of traversals.

Thus, in Stratego and SYB children are determined by *structure*, in Uniplate children are determined by *type*, and in KURE children are determined by *location*, as specified by the user.

### 6.2 Typing a generic traversal

We will now compare the types of the generic traversal strategies in KURE, SYB and Uniplate. We use Haskell for the comparison, but we present only the essence of each approach, not the actual Haskell library implementations. Strategies in Stratego are untyped, so we mostly omit Stratego from the discussion.

The focus of our comparison will be the shallow traversal that requires success at every child node: KURE's *allR* strategy. The analogue in SYB is *gmapM*, while Uniplate has two analogues: *descendM* and *descendBiM*. These can be defined as follows:

**class** *Data* $\tau$ **where**
    $gmapM :: Monad\ m \Rightarrow (\textbf{forall}\ \sigma.\ Data\ \sigma \Rightarrow \sigma \rightarrow m\ \sigma) \rightarrow (\tau \rightarrow m\ \tau)$

**class** *Uniplate* $\tau$ **where**
    $descendM :: Monad\ m \Rightarrow (\tau \rightarrow m\ \tau) \rightarrow (\tau \rightarrow m\ \tau)$
**class** *Uniplate* $\sigma \Rightarrow$ *Biplate* $\tau\ \sigma$ **where**
    $descendBiM :: Monad\ m \Rightarrow (\sigma \rightarrow m\ \sigma) \rightarrow (\tau \rightarrow m\ \tau)$

These classes are analogous to KURE's *Walker* class, with the type $\tau \rightarrow m\ \tau$ corresponding to a rewrite.

The type of each shallow traversal is a consequence of which substructures are considered children. In SYB the children are all proper components of the node being traversed, and so the argument rewrite to *gmapM* must be applicable at almost any type. The Haskell implementation of SYB achieves this via *rank-2 polymorphism* (Peyton Jones *et al.*, 2007).

The Uniplate shallow traversals have simpler types. As *descendM* treats only similarly typed proper substructures as children, its rewrite argument need only be applicable at that one type. The type of *descendBiM* is more general, allowing the type of the children to differ from their parent, but still requiring all children to have the same type.

Note that for both the SYB and Uniplate libraries, an *applicative functor* (McBride & Paterson, 2008) rather than a monad would suffice, whereas a *Monad is* required for KURE's *allR* traversal. This is because *allR* takes a rewrite that operates on a universe as an argument, and KURE needs the ability to fail in the (error) case that the argument rewrite changes the child to a different summand. This is a disadvantage of the rewrite-over-a-sum-type approach.

Finally, we note that while Stratego is conventionally an untyped language, there has been recent work on adding *dynamic* type checks to Stratego, in the form of *typesmart constructors* (Erdweg *et al.*, 2014). These typesmart constructors dynamically check that any node constructed by a transformation is well-typed, and cause the transformation to fail if not. While implemented rather differently, in essence this is similar to KURE's dynamic check that a rewrite operating on a universe does not change the summand type: in both cases the strategic programming notion of *failure* is invoked when a term is found to be ill-typed by a dynamic check.

### 6.3 *Defining a generic traversal*

We will now present some example definitions of the generic shallow traversal in SYB and Uniplate, using the same object language of string expressions that we introduced in Section 4.3. To keep the examples concise, we only define traversals for *Expr* and *Decl* nodes.

The intended SYB semantics entirely determine the definition of SYB's shallow traversal: it must target every proper component. A major advantage of this inflexibility is that the definition can be derived mechanically. The *Data* instances for *Decl* and *Expr* are:

```
instance Data Expr where
    gmapM :: Monad m ⇒ (forall σ. Data σ ⇒ σ → m σ) → (Expr → m Expr)
    gmapM r (Var n)       = Var    <$> r n
    gmapM r (Lit s)       = Lit    <$> r s
    gmapM r (Append e₀ e₁) = Append <$> r e₀ <*> r e₁
    gmapM r (Let d e)     = Let    <$> r d  <*> r e
instance Data Decl where
    gmapM :: Monad m ⇒ (forall σ. Data σ ⇒ σ → m σ) → (Decl → m Decl)
    gmapM r (Decl n e) = Decl <$> r n <*> r e
```

Because all proper components are targeted, the types of all proper components require *Data* instances to be declared for them, even if the user would prefer a statically selective traversal that does not descend into them. In this case, these instances rely on existing *Data* instances for lists and characters.

The intended Uniplate semantics entirely determine the definitions of Uniplate's shallow traversals: the targets are the maximal substructures of a specific type. These targets can be below several constructors, which leads to more complex instances. To gain some code reuse, the instances are usually defined mutually recursively; but this still leads to a quadratic number of instances. For example, the instances for traversing *Expr* and *Decl* nodes are as follows:

```
instance Uniplate Expr where
    descendM :: Monad m ⇒ (Expr → m Expr) → (Expr → m Expr)
    descendM _ (Var n)        = Var    <$> pure n
    descendM _ (Lit l)        = Lit    <$> pure l
    descendM r (Append e₀ e₁) = Append <$> r e₀              <*> r e₁
    descendM r (Let d e)      = Let    <$> descendBiM r d <*> r e
instance Uniplate Decl where
    descendM :: Monad m ⇒ (Decl → m Decl) → (Decl → m Decl)
    descendM r (Decl n e) = Decl <$> pure n <*> descendBiM r e
instance Biplate Decl Expr where
    descendBiM :: Monad m ⇒ (Expr → m Expr) → (Decl → m Decl)
    descendBiM r (Decl n e) = Decl <$> pure n <*> r e
instance Biplate Expr Decl where
    descendBiM :: Monad m ⇒ (Decl → m Decl) → (Expr → m Expr)
    descendBiM _ (Var n)        = Var    <$> pure n
    descendBiM _ (Lit l)        = Lit    <$> pure l
    descendBiM r (Append e₀ e₁) = Append <$> descendBiM r e₀ <*> descendBiM r e₁
    descendBiM r (Let d e)      = Let    <$> r d              <*> descendBiM r e
```

To avoid the quadratic code explosion, the Uniplate library offers alternative means of defining instances by making them polymorphic, using SYB to various degrees.

### 6.4 Statically selective traversals

As discussed in Section 2.3 there are two forms of selective traversal: static and dynamic. Static selectivity is more efficient than dynamic, but static selectivity can also limit the expressiveness of a traversal.

The SYB semantics mandate exactly one way to traverse a node, and that is to descend into all proper components of the node. That is, the user is not intended to define statically selective traversals. This is analogous to being limited to a single KURE universe that contains all substructures of the tree as summands. Consequently, dynamic selectivity is the only way to prevent a deep traversal from descending into all substructures of the tree. Stratego is essentially the same as SYB in this regard, in that it provides exactly one way to traverse a node.

Like SYB, the Uniplate semantics mandate how to traverse a node. Unlike SYB, Uniplate is fundamentally characterised by static selectivity: each traversal only ever has one target type, and it only descends into a substructure if it could contain that type. Essentially, Uniplate provides a separate statically selective traversal for each type of node, each of which corresponds to a KURE traversal indexed by a singleton universe.

However, the expressiveness cost of Uniplate's static selectivity is that if the user wishes to modify nodes of different types, then she must perform multiple traversals, one for each node type that she wishes to target. Decomposing a traversal into multiple sequential traversals in this way is inefficient, and could be problematic to express if the traversal uses monadic effects in such a way that they cannot be reordered.

SYB and Uniplate are the two endpoints of a spectrum regarding static selectivity, and a KURE traversal can lie anywhere in between. The universe type determines which locations are to be visited by a traversal, thereby allowing the user to customise the static selectivity. A large universe tends towards the SYB end of the spectrum; in which case

dynamic selectivity may offer significant efficiency gains. Note that, in contrast to Uniplate, the static selectivity of KURE is location-directed, not type-directed.

### 6.5 Data abstraction

While the intended SYB *semantics* mandate exactly one way to traverse a node, the Haskell *implementation* of SYB does not prevent the user from declaring semantically invalid *Data* instances that omit some proper components of a node during a traversal. However, this may cause third-party code to behave unexpectedly when used with such instances. Furthermore, as there can only be one *Data* instance in scope, this would not allow other statically selective traversals to be defined: in SYB the *Data* instance defines the one sole way to traverse all nodes of a given type.

There is a workaround for these problems though: the user can create new abstract types as wrappers around the node types to be traversed, with a manually defined *Data* instance for each abstract type that defines a statically selective traversal. This overcomes the limitation of there only being one way to traverse a node, as each abstract type has its own traversal behaviour. The disadvantages of this approach are that it violates the intended semantics of SYB, the code can no longer be mechanically generated, and it requires a set of abstract types to be defined for each statically selective traversal.

The Uniplate semantics specify how to traverse a node based on the type of the targeted children. However, as with SYB, it is possible to define semantically invalid instances that omit some children of the targeted type, or to introduce abstract-type wrappers to allow several statically selective traversals to be defined.

Ultimately, if the user is willing to disregard the intended semantics of SYB and Uniplate, then extensive use of abstract types and manual instance declarations allows either library to simulate the traversal behaviour of the other, or even of KURE. Meanwhile, KURE is designed to be configurable to different statically selective traversals, so simulating the behaviour of SYB or Uniplate, or behaving somewhere in between, is within its intended semantics, and does not require inventive uses of abstract types. Furthermore, as KURE uses the universe type to index its traversals, abstract-type wrappers are only required in the (rare) case that a traversal needs to traverse nodes of the same type in different locations, while treating them distinctly (see Appendix).

### 6.6 Modularity

We now consider the modularity of each library, by contrasting how easy it is to add new (statically selective) traversals, or to modify existing traversals. Stratego is the simplest case: there is exactly one way to traverse a node, and this is built-in to the Stratego language. Thus it is perfectly modular as it can traverse any object language with nothing required from the user — but that traversal cannot be modified in any way.

SYB is the second-most modular because it defines a single statically selective traversal using an open universe. The user can add new node types without modifying any existing code, extending the SYB traversal to operate over types that were not previously traversable.

Uniplate is founded on the premise that each (statically selective) traversal targets children of exactly one type. Thus, for each traversal, the set of child types is a closed singleton set. The Biplate extension allows for parent nodes of multiple types to be traversed, sharing code, but the child nodes being targeted must still be of that singleton type. Thus, the universe of traversable node types is open, and there is no need to modify existing code when adding new parent node types (though doing so may allow for more code reuse). That is, a Uniplate traversal is modular in the sense that it can be extended to operate over new traversable nodes, but it is not modular in the sense that rewrites targeting children of different types cannot be combined into a single traversal. To target children of a different type, an entirely new traversal must be defined.

In KURE, each universe is a closed sum type, and is used to define exactly one statically selective traversal. As the universe is closed, it is not possible to add new nodes without modifying existing code. However, new traversals can be defined by adding new universes (and their corresponding instances). Because KURE traversals are more configurable than those of Stratego, SYB and Uniplate, it is more likely that several similar traversals will be defined (e.g. over the *Core* and *CoreTC* universes of HERMIT). For these cases, some code reuse can be enabled by defining instances in terms of congruence combinators (Section 5.2).

## 6.7 Summary

If used as intended, SYB, Uniplate and KURE provide fairly distinct approaches to traversal, with SYB corresponding closely with Stratego. SYB provides a single traversal based on an open set of node types; KURE allows the user to specify statically selective traversals, each with a closed set of node types; and Uniplate allows multiple statically selective traversals, each with an open set of traversable node types but only a single target node type. In SYB everything is a target; in Uniplate all substructures of a specific type are targets, and in KURE the user specifies which locations are to be targets. They each take different approaches to typing: SYB uses rank-2 polymorphism; KURE uses a universe type; and Uniplate is able to give the type directly because it only targets one type. However, by using abstract types, it is possible for the implementation of each library to emulate the behaviour of the others.

Assuming the intended semantics of each approach, KURE can be considered more expressive because the user specifies the children to be targeted, whereas the targets are fixed by the semantics of SYB and Uniplate. However, this expressiveness comes at the cost of modularity: the universe type is closed and thus existing code has to be recompiled whenever a universe is extended.

An advantage of KURE and Stratego over SYB and Uniplate is that they provide an extensive library of strategic traversals, with domain-specific error messages. KURE also provides support for allowing generic traversals to automatically update a context. These are pragmatic benefits: it would be possible to implement a strategic programming language that contained similar infrastructure but used the traversal approach of SYB or Uniplate.

Finally, we note that the traversal semantics of Stratego, SYB and Uniplate are specific enough that the traversals can be mechanically derived. This is not possible in KURE, because children are user-specified locations.

## 7 Performance

This section compares the performance of KURE, SYB and Uniplate using three simple strategic programs. Our focus is on the generic shallow traversal operator of each library, so we have selected small benchmarks that are just sufficient to demonstrate the differences in traversal performance. We try to identify the causes of the different performance results, and analyse the impact of selective traversals and the most relevant GHC optimisations. This analysis is quite in depth, and is aimed at a reader interested in the efficient implementation of typed strategic or generic programming languages.

We do not provide a performance comparison with Stratego, because recent Stratego versions are implemented as an Eclipse plugin, with Stratego code being compiled to Java. The priority of this implementation effort is portability rather than efficiency, and the slow loading time of Java classes is intended to be amortised over a lengthy Eclipse session (Visser, 2013). Conversely, KURE, SYB and Uniplate are all implemented as Haskell libraries, and are significantly optimised by GHC.

Although GHC offers many optimisation parameters, we only used the standard level-O2 optimisation for our benchmarks, as we are interested in the performance of common usage. We measure using the Criterion (O'Sullivan, 2012) benchmarking library on an otherwise quiescent workstation; our measurements have more than 95% confidence. We used the most recently released version of each tool and library at the time of making the measurements, which were: KURE 2.16.1, SYB 0.4.1, Uniplate 1.6.12, Criterion 0.5.1.1 and GHC 7.8.2. By using the same compiler, compiler options and benchmarking tools, our results are a meaningful comparison of the relative performance of the differing traversal techniques. The benchmark code is available as supplementary material on the Journal of Functional Programming website.

### 7.1 Type class dictionaries and the static argument transformation

In GHC, type class constraints are implemented as functions that take a *class dictionary* as an argument (Wadler & Blott, 1989). A class dictionary is essentially a record containing a definition for each method of the type class, specialised to a concrete class parameter (or parameters). Thus each class instance generates a single dictionary. The key to performance in our two benchmarks is to eliminate higher-order functions, of which a major source is the passing of these class dictionaries. We will explain how various aspects of each library enable or inhibit this objective.

Inlining directly eliminates higher-order functions. In GHC, inlining also enables many other optimisations, including type specialisation. In particular, the *static argument transformation* (SAT) (Santos, 1995) makes it possible to specialise recursive definitions. The SAT pulls outside of the recursion any arguments — including dictionaries — that never change in recursive calls, thereby creating an outer definition that can be inlined in order to specialise the inner recursive definition. However, in a mutually recursive set of definitions, GHC heuristically chooses one definition to be a loop-breaker, which prohibits it from being inlined. Thus, none of the arguments bound by the loop-breaker can be specialised via inlining. GHC's implementation of the SAT is quite conservative, because the transformation by itself is often detrimental unless it enables further optimisations. However, as in

our context the transformation is beneficial, we write our definitions with explicitly static arguments whenever possible.

The SYB architecture precludes the *gmapM* method of the *Data* class from having explicitly static arguments when the class parameter is a recursive data type. In the *gmapM* method, any application of its rewrite argument must be supplied a corresponding *Data* dictionary, which creates mutual recursion between the *Data* dictionary and the traversal. By preventing inlining and hence subsequent optimisations, this is a major contributor to the poor performance of SYB.

Mutual recursion between dictionaries also arises in Uniplate, but the consequences are less severe as the rewrite and dictionary arguments of *descendBiM* are explicitly static arguments. The KURE universe type enables the SAT for the same reason that it is detrimental to modularity: all of the traversals are defined in a single *Walker* instance. Thus the rewrite and the dictionary are explicitly static arguments, even in the presence of mutually recursive node types.

### 7.2 Benchmark: Fibonacci

This benchmark evaluates the Fibonacci function by rewriting the terms of a simple singly recursive expression language. This exercises the bottom-up deep-traversal strategy. The rewrite rules are a pedantic transcription of the Fibonacci recurrence relation:

```
type R m a = a → m a
data Fib = Lit Int | Plus Fib Fib | Fib Fib
plusRule :: Monad m ⇒ R m Fib
plusRule (Plus (Lit x) (Lit y)) = return (Lit (x + y))
plusRule _                      = fail "plusRule"
fibBaseRule :: Monad m ⇒ R m Fib
fibBaseRule (Fib (Lit 0)) = return (Lit 0)
fibBaseRule (Fib (Lit 1)) = return (Lit 1)
fibBaseRule _             = fail "fibBaseRule"
fibStepRule :: Monad m ⇒ R m Fib
fibStepRule (Fib n) = return (Plus (Fib (Plus n (Lit (−2)))) (Fib (Plus n (Lit (−1)))))
fibStepRule _       = fail "fibStepRule"
```

For each variant, we define a function to reduce a *Fib* term to an integer. The following hand-written statically selective traversal (Hand) is our baseline variant:

```
reduce :: MonadCatch m ⇒ R m Fib
reduce = eval
   where
     eval    = allbu (try (plusRule <+ evalFib))
     evalFib = fibBaseRule <+ (fibStepRule >>> eval)
allbu :: Monad m ⇒ R m Fib → R m Fib
allbu r = go
   where
     go = all go >>> r
all :: Monad m ⇒ R m Fib → R m Fib
all r (Plus a b) = Plus <$> r a <*> r b
all r (Fib a)    = Fib  <$> r a
all _ x          = pure x
```

Table 1. *Fibonacci benchmark, slowdown with respect to* `Hand`

| code variant | geometric mean | geometric standard deviation | goodness of fit ($r^2$) |
|---|---|---|---|
| SYB-gmapM | 4.65 | 1.09 | 0.99 |
| SYB-sat | 2.84 | 1.08 | 0.99 |
| SYB-sat-static | 20.88 | 1.13 | 0.99 |
| KURE | 1.04 | 1.12 | 0.99 |
| Uni | 0.84 | 1.06 | 0.99 |

We omit the definitions of $>>>$, $<+$, and *try*, which are defined in the usual manner.

The performance results are shown in Table 1. For each variant, we measure the time to compute *reduce* (Fib (Lit $n$)), where $n$ ranges from 20 to 34. For the SYB variants we stop at 31 because they become slow. We list the geometric mean and geometric standard deviation of its slowdown with respect to the baseline. We also list the goodness of fit of each variant's measurements' logarithmic regression, since the expected complexity of the Fibonacci computation is exponential. The worst-case error in any given ratio of execution times or slowdowns was 2.2%.

We used the following variants in our experiment:

- `SYB-gmapM` uses the deep traversal *everywhereM* along with a hand-written *gmapM* definition, which outperforms the (omitted) GHC-derived *gmapM* definition by approximately 15%.
- `SYB-sat` uses the same *gmapM* definition from `SYB-gmapM`, but uses a SAT of *everywhereM*. This allows the traversal to be inlined and hence specialised to a concrete *Monad* dictionary.
- `SYB-sat-static` uses an abstract type (as discussed in Section 6.5) to define a statically selective traversal that omits *Int* nodes; it also uses the SAT of *everywhereM*.
- `KURE` simply uses *Fib* as a singleton universe.
- `Uni` defines the conventional *Uniplate* instance for *Fib*.

For this dense bottom-up traversal of a singly recursive type, it is evident that Uniplate is well-tuned for its intended purpose, outperforming even the hand-written shallow traversal by 15%. We suspect this is due to *constructor specialisation* (Peyton Jones, 2007), because the resulting GHC Core is relatively large with many similar loops. KURE nearly matches the hand-written baseline. The SYB variants are at least three times slower than `Hand`. For this particular traversal, the statically selective SYB variant (`SYB-sat-static`) is actually much slower, as any speedup due to static selectivity is overwhelmed by the abstract type interfering with other optimisations. To achieve these performance results, the user need only use level-`O2` optimisation and ensure that inlining and specialisation are possible by defining functions with explicitly static arguments.

### 7.3 Benchmark: Paradise

Our second benchmark is the *increase* function from the Paradise benchmark (Lämmel & Peyton Jones, 2003). This differs from the Fibonacci benchmark by traversing several data

types, some of which are mutually recursive. The performance results vary primarily because of the large potential for static selectivity — only 33% of the non-trivial substructures are the actual targets. The data types are as follows:

**newtype** *Company* = C [*Dept*]
**data** *Dept*　　　 = D *Name Employee* [*Unit*]
**data** *Unit*　　　 = PU *Employee* | DU *Dept*
**data** *Employee*　 = E *Person Salary*
**data** *Person*　　 = P *Name Address*
**newtype** *Salary*　 = S *Integer*
**type** *Name*　 = *String*
**type** *Address* = *String*

The objective is to increase all *Salary* values in a *Company* by a given increment. The following hand-written statically selective traversal (`Hand`) is the conventional definition of the traversal, and we use it as our baseline variant:

*increase* :: *Monad m* ⇒ *Integer* → *R m Company*
*increase k* = *allbuC* (*inc k*)
*inc* :: *Monad m* ⇒ *Integer* → *R m Salary*
*inc k* = $\lambda$ (S *x*) → **if** *x* < 0 **then** *fail* `"inc"` **else** *return* (S (*x* + *k*))

*allbuC f* (C *ds*)　　 = C　<$> *mapM* (*allbuD f*) *ds*
*allbuD f* (D *n m us*) = D *n* <$> *allbuE f m* <*> *mapM* (*allbuU f*) *us*
*allbuE f* (E *p s*)　　 = E *p* <$> *f s*
*allbuU f* (PU *e*)　　 = PU <$> *allbuE f e*
*allbuU f* (DU *d*)　　 = DU <$> *allbuD f d*

The performance results are shown in Table 2. Each traversal was applied to the same 100 arbitrary test inputs, whose sizes are evenly spread out along the line from 1000 to 8500 nodes. We list the goodness of fit for a linear model, since the traversal complexity is approximately linear with respect to the number of nodes. The worst-case error in any given ratio of execution times or slowdowns was 6.5%.

We measured the following variants of the traversal:

- `Hand-sat` is a manual application of the SAT to `Hand`; it enables specialisation with respect to both the monad and the rewrite argument.
- `SYB-dyn` uses a SAT of the deep traversal *everywhereBut*, which uses dynamic selectivity to descend only into nodes that could contain a *Salary*. The *Data* instance is derived by SYB.
- `SYB-static` uses a SAT of *everywhere* with an abstract type with a manually defined *Data* instance to encode a statically selective traversal that only visits *Salary* substructures. As the *Data* dictionaries for *Dept* and *Unit* are mutually recursive via the *gmapM* definitions, one of those definitions will be a loop-breaker.
- `SYB-static-sat` is a variant of `SYB-static` that eliminates the mutual recursion by including a distinct copy of the *gmapM* definition for *Unit* within the *gmapM* definition for *Dept*. The copy of *gmapM* has the SAT applied. This variant enables specialisation with respect to both the monad and the rewrite argument.
- `KURE` uses a conventional universe for the data types, and uses the *allbuR* strategy to reach all of the *Salary* substructures.

N. Sculthorpe et al.

Table 2. *Paradise benchmark, slowdown with respect to* `Hand`, *100 arbitrary inputs*

| code variant | geometric mean | geometric standard deviation | goodness of fit ($r^2$) |
|---|---|---|---|
| `Hand-sat` | 0.54 | 1.13 | 0.96 |
| `SYB-dyn` | 3.68 | 1.17 | 0.97 |
| `SYB-static` | 1.60 | 1.15 | 0.96 |
| `SYB-static-sat` | 0.54 | 1.14 | 0.86 |
| `KURE` | 0.68 | 1.14 | 0.88 |
| `KURE-sel` | 0.54 | 1.14 | 0.95 |
| `Uni` | 1.01 | 1.16 | 0.95 |

- `KURE-sel` uses an alternative universe that only descends into *Salary* substructures.
- `Uni` defines the conventional *Uniplate* and *Biplate* instances for the *Salary* target type. As the *Biplate* dictionaries for *Dept* and *Unit* are mutually recursive via the *descendBiM* definitions, one of the those definitions will be a loop-breaker.

The factor of two speedup of `Hand-sat` over `Hand` demonstrates the effectiveness of the SAT. SYB is 350% slower than `Hand`, even with dynamic selectivity. The static selectivity of `SYB-static` reduces the overhead to 60%, demonstrating the advantage of static over dynamic selectivity in SYB for a function like this benchmark. The `SYB-static-sat` variant further enables the SAT by breaking the mutual recursion among dictionaries, at the cost of some minor code duplication, and the result performs as well as the `Hand-sat` variant. The systemic SYB issue that precluded static arguments is avoided in the two statically selective variants: the target type *Salary* is not recursive, so the traversal does not visit any substructures that result in mutual recursion with the *Data* dictionary. And in the `SYB-static-sat` variant there is no recursion among dictionaries, so inlining can specialise the traversals as much as possible.

The Uniplate variant is about 100% slower than the best KURE and SYB variants. This is an unexpected result, and Neil Mitchell has communicated that it is a regression. We anticipate performance at least as fast as KURE.

The conventional KURE definition, visiting all substructures inhabiting one of the data types, achieves an efficiency just 20% slower than `Hand-sat`. The `KURE-sel` variant increases the static selectivity by only visiting *Salary* substructures, matching `Hand-sat`.

This benchmark emphasises the importance of explicitly static arguments for GHC optimisation and the potential speedup of static selectivity. Indeed, a traversal defined with SYB — the conventionally worst performing library — that was designed specifically for both of those concerns gives performance matching hand-written definitions.

### 7.4 Benchmark: Tseitin transformation

The Tseitin Transformation (Tseitin, 1968) is an algorithm for converting an arbitrary Boolean expression to Conjunctive Normal Form. The traversal differs from the previous benchmarks in two ways. First, it uses a state monad transformer to maintain a name supply and a writer monad transformer to collect the list of conjunctions. Second, it has aspects that are awkward to express as a strategic program. As with the Fibonacci benchmark, the

Table 3. *Tseitin benchmark, slowdown with respect to* `Hand`*, 100 arbitrary inputs*

| code variant | geometric mean | geometric standard deviation | goodness of fit ($r^2$) |
|---|---|---|---|
| KURE | 1.00 | 1.05 | 0.87 |
| SYB-gmapM | 3.32 | 1.15 | 0.82 |
| SYB-sat | 3.29 | 1.15 | 0.77 |
| SYB-sat-dyn | 1.86 | 1.16 | 0.73 |
| SYB-sat-static | 3.24 | 1.11 | 0.88 |
| Uni | 1.05 | 1.10 | 0.78 |

data type is singly recursive. The traversal uses the *allbuR* combinator and must visit all compound expressions, so static selectivity is relatively unimportant.

The data type we use for Boolean expressions is *BExp*:

**data** *BExp* = Lit *BLit* | And *BExp BExp* | Or *BExp BExp* | Not *BExp*
**data** *BLit* = Var *String* | TT | FF

The algorithm replaces every operator application with a freshly generated variable name and emits conjuncts coupling that new variable to its arguments' respective variable names. However, without a term representation based on pattern/shape functors (Swierstra, 2008), there is no way to express the type inhabited by the intermediate result after applying the algorithm to a node's sub-expressions before applying it to the node itself. We settle for wrapping and unwrapping the generated variable names with the Lit constructor.

We give the definition of the rule for normalising a conjunction; the omitted rules for disjunction and negation are similar. Each rule uses the state-effect to generate a unique variable name, uses the writer-effect to emit the encoding conjuncts as a *difference list* (Hughes, 1986), and returns the generated variable so that ancestor nodes can proceed in the same way. The *MT* monad transformer provides the state- and writer-effects, and the *allbu* combinator extends the rules to entire expressions.

```
tseitin :: MonadCatch m ⇒ R (MT m) BExp
tseitin = allbu (try (andRule <+ orRule <+ notRule))
andRule :: MonadCatch m ⇒ R (MT m) BExp
andRule (And (Lit al) (Lit bl)) = do v ← get
                                      put (v + 1)
                                      let a = Lit al
                                          b = Lit bl
                                          c = Lit (Var ("$x" ++ show v))
                                          d₁ = [Not a, Not b, c]
                                          d₂ = [a, Not c]
                                          d₃ = [b, Not c]
                                      tell ((d₁:) ∘ (d₂:) ∘ (d₃:))
                                      return c
andRule _ _ = fail "andRule"
```

The performance results are shown in Table 3. Each traversal was applied to the same 100 arbitrary test inputs, whose sizes are evenly spread out along the line from 3000 to 9500 nodes. On average, 13% of the nodes were Nots and 44% were Lits. We list the goodness of fit for a linear model, since the traversal complexity is approximately linear

with respect to the number of nodes. The worst-case error in any given ratio of execution times or slowdowns was 5.2%.

We measured the following variants of the traversal:

- `SYB-gmapM` uses the conventional *everywhereM* strategy with no selectivity along with a hand-written *gmapM* definition, as in the Fibonacci benchmark.
- `SYB-sat` uses a SAT of *everywhereM*.
- `SYB-sat-dyn` uses a SAT of the deep traversal strategy *everywhereMBut*, which dynamically avoids descending into literals.
- `SYB-sat-static` uses a SAT of *everywhereM* and uses an abstract type to statically avoid descending into Lit nodes.
- `KURE` simply uses *BExp* as a singleton universe.
- `Uni` defines the conventional *Uniplate* instance for *BExp*.

The results exhibit the same themes as discussed for the previous benchmarks: KURE and Uniplate are comparable to hand-written code, and `SYB-gmapM` has significant overhead. Thus the relative performances were not altered by the presence of monad transformers, nor the use of an algorithm that does not fit as cleanly into a strongly typed strategic programming paradigm. In particular, as with the Fibonacci benchmark, the benefit of static selectivity is outweighed by the interference of the abstract type with other optimisations.

### 7.5 Summary

A key optimisation for generic traversals is selectivity, and static selectivity is more efficient than dynamic selectivity. The performance gain of selectivity depends on the proportion of the data structure that needs to be traversed. In the Paradise benchmark, where much of the structure could be avoided, the gain was significant. Conversely, in the Fibonacci and Tseitin benchmarks, where most of the structure must be traversed, the effective overhead of encoding the selectivity outweighed the benefit.

Performing the SAT can bring significant benefits to the performance of generic traversals in Haskell, as it allows inlining and thence specialisation. The design of SYB interferes with the SAT, which contributes to SYB's relatively poor performance. Only by manually duplicating code were we able to apply the SAT for the SYB variant and achieve performance comparable with the other libraries. Note that the variants with hyphenated names (`-gmapM`, `-sat`, `-static`, etc.) required manual implementation of the corresponding optimisation; the unhyphenated variants represent a straightforward use of the libraries.

The relatively poor performance of SYB is well known (Rodriguez Yakushev *et al.*, 2008). Recent work by Adams *et al.* (2014) has attempted to address this by applying an (SYB-specific) mechanical transformation to user SYB code at compile-time, implemented using the HERMIT system described in Section 5. By using aggressive inlining and symbolic evaluation, they were able to improve the performance of SYB code to match hand-written code.

In summary, on all three benchmarks KURE was significantly better than SYB. Uniplate performed better than KURE for Fibonacci, roughly the same as KURE for Tseitin, and noticeably worse than KURE for Paradise. However, the latter result appears to be a regression in Uniplate, and we expect KURE-like performance once that is resolved.

## 8 Conclusion

In this article we have described our approach of using *universe types* to assign types to generic traversals, and to support statically selective traversals and traversals that can distinguish between data based on its *location*, not just its *type*. We then described the implementation of this idea as part of a Haskell-embedded strategic programming language. To demonstrate the viability of this approach, we presented its usage in the HERMIT system.

The main novelty of KURE compared to other strategic and generic rewriting systems is the way it uses universe types to index generic traversals. We gave a detailed comparison of KURE generic traversals, and their types, with those of SYB and Uniplate, two other approaches to typed generic programming. The crucial distinction is that the semantics of KURE traversals are not determined by the structure or types of the data structure: instead the KURE user can customise how traversals should treat each *location* in the data type being traversed. Consequently KURE is relatively configurable, and can be used to simulate either SYB or Uniplate semantics, or, more usually, can be given a semantics somewhere in between. However, this flexibility does impose a cost in modularity, with modifications to a traversal requiring more recompilation of existing code than in either of the other two libraries (though the necessary changes are fairly small and localised).

Finally, we compared the performance of the KURE implementation with the SYB and Uniplate libraries, examining the main optimisations that each library supports or inhibits. For the small benchmarks we used, we found KURE's performance to be roughly comparable with Uniplate, and superior to SYB.

In closing, we think that the KURE approach is a useful addition to the existing approaches to typed strategic and generic programming. The emphasis on customisable strategic traversals has allowed KURE to serve as the rewrite engine of HERMIT, and we expect to use KURE as a foundation for building higher-level rewriting languages in the future.

## Supplementary materials

For supplementary material for this article, please visit dx.doi.org/10.1017/10.1017/S0956796814000185.

## References

Adams, M. D., Farmer, A. & Magalhães, J. P. (2014) Optimizing SYB is easy! In *Workshop on Partial Evaluation and Program Manipulation*. New York, NY: ACM, pp. 71–82.

Balland, E., Moreau, P.-E. & Reilles, A. (2008) Rewriting strategies in Java. In *Eighth International Workshop on Rule Based Programming*, Electronic Notes in Theoretical Computer Science, vol. 219. Amsterdam: Elsevier, pp. 97–111.

Bravenboer, M., Kalleberg, K. T., Vermaas, R. & Visser, E. (2008) Stratego/XT 0.17. A language and toolset for program transformation. *Sci. Comput. Program.* **72**(1–2), 52–70.

Bringert, B. & Ranta, A. (2008) A pattern for almost compositional functions. *J. Funct. Program.* **18**(5–6), 567–598.

Brown, N. C. C. & Sampson, A. T. (2009) Alloy: Fast generic transformations for Haskell. In *Haskell Symposium*. New York, NY: ACM, pp. 105–116.

Burstall, R. M. & Darlington, J. (1977) A transformation system for developing recursive programs. *J. ACM* **24**(1), 44–67.

Culpepper, R. (2012) Fortifying macros. *J. Funct. Program.* **22**(4–5), 439–476.

Culpepper, R. & Felleisen, M. (2010) Fortifying macros. In *International Conference on Functional Programming*. New York, NY: ACM, pp. 235–246.

Dolstra, E. (2001) *First Class Rules and Generic Traversals for Program Transformation Languages*, Master's thesis, Utrecht University.

Dolstra, E. & Visser, E. (2001) *First-Class Rules and Generic Traversal*. Tech. rept. UU-CS-2001-38. Utrecht University, Utrecht.

Ellison, C. & Roşu, G. (2012) An executable formal semantics of C with applications. In *Symposium on Principles of Programming Languages*. New York, NY: ACM, pp. 533–544.

Erdweg, S., Rendel, T., Kästner, C. & Ostermann, K. (2011) SugarJ: Library-based syntactic language extensibility. In *International Conference on Object-Oriented Programming, Systems, Languages, and Applications*. New York, NY: ACM, pp. 391–406.

Erdweg, S. & Rieger, F. (2013) A framework for extensible languages. In *International Conference on Generative Programming: Concepts & Experiences*. New York, NY: ACM, pp. 3–12.

Erdweg, S., Rieger, F., Rendel, T. & Ostermann, K. (2012) Layout-sensitive language extensibility with SugarHaskell. In *Haskell Symposium*. New York, NY: ACM, pp. 149–160.

Erdweg, S., Vergu, V., Mezini, M. & Visser, E. (2014) Modular specification and dynamic enforcement of syntactic language constraints when generating code. In *International Conference on Modularity*. New York, NY: ACM, pp. 241–252.

Farmer, A., Gill, A., Komp, E. & Sculthorpe, N. (2012) The HERMIT in the machine: A plugin for the interactive transformation of GHC core language programs. In *Haskell Symposium*. New York, NY: ACM, pp. 1–12.

Felleisen, M., Findler, R. B. & Flatt, M. (2009) *Semantics Engineering with PLT Redex*. Cambridge, MA: MIT Press.

Frisby, N., Gill, A. & Alexander, P. (2012) A pattern for almost homomorphic functions. In *Workshop on Generic Programming*. New York, NY: ACM, pp. 1–12.

Gibbons, J. (2005) Patterns in datatype-generic programming. In *Workshop on Declarative Programming in the Context of Object-Oriented Languages 2003*, Striegnitz, J. & Davis, K. (eds), NIC Series, vol. 27. Jülich: NIC, pp. 277–289.

Gill, A. (2006) Introducing the Haskell equational reasoning assistant. In *Haskell Workshop*. New York, NY: ACM, pp. 108–109.

Gill, A. (2009) A Haskell hosted DSL for writing transformation systems. In *Working Conference on Domain-Specific Languages*, Taha, W. (ed), Lecture Notes in Computer Science, vol. 5658. Berlin: Springer, pp. 285–309.

Gill, A. (2013) `http://hackage.haskell.org/package/html-kure`.

Girard, J.-Y.. (1972) *Interprétation fonctionelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Ph.D. thesis, Université Paris Diderot, Paris.

Hinze, R. & Löh, A. (2007) Generic programming, now! *International spring school on datatype-generic programming 2006*, Backhouse, R., Gibbons, J., Hinze, R. & Jeuring, J. (eds), Lecture Notes in Computer Science, vol. 4719. Berlin: Springer, pp. 150–208.

Hinze, R. & Löh, A. (2009) Generic programming in 3D. *Sci. Comput. Program.* **74**(8), 590–628.

Hughes, R. J. M. (1986) A novel representation of lists and its application to the function "reverse". *Inf. Process. Lett.* **22**(3), 141–144.

Kats, L. C. L. & Visser, E. (2010) The Spoofax language workbench. Rules for declarative specification of languages and IDEs. In *International Conference on Object-Oriented Programming, Systems, Languages, and Applications*. New York, NY: ACM, pp. 444–463.

Kiselyov, O. (2006) *Smash your boiler-plate without class and typeable*.
http://article.gmane.org/gmane.comp.lang.haskell.general/14086.

Klein, C., Clements, J., Dimoulas, C., Eastlund, C., Felleisen, M., Flatt, M., McCarthy, J. A., Rafkind, J., Tobin-Hochstadt, S. & Findler, R. B. (2012) Run your research: On the effectiveness of lightweight mechanization. In *Symposium on Principles of Programming Languages*. New York, NY: ACM, pp. 285–296.

Lämmel, R. & Peyton Jones, S. (2003) Scrap your boilerplate: A practical design pattern for generic programming. In *Types in Languages Design and Implementation*. New York, NY: ACM, pp. 26–37.

Lämmel, R. & Visser, J. (2002) Typed combinators for generic traversal. *International Symposium on Practical Aspects of Declarative Programming*, Krishnamurthi, S. & Ramakrishnan, C. R. (eds), Lecture Notes in Computer Science, vol. 2257. Berlin: Springer, pp. 137–154.

Lazar, D., Arusoaie, A., Serbanuta, T. F., Ellison, C., Mereuta, R., Lucanu, D. & Roşu, G. (2012) Executing formal semantics with the K tool. In *International symposium on formal methods*, Giannakopoulou, D. & Méry, D. (eds), Lecture Notes in Computer Science, vol. 7436. Berlin: Springer, pp. 267–271.

Liang, S., Hudak, P. & Jones, M. (1995) Monad transformers and modular interpreters. In *Symposium on Principles of Programming Languages*. New York, NY: ACM, pp. 333–343.

Löh, A. & Magalhães, J. P. (2011) Generic programming with indexed functors. In *Workshop on Generic Programming*. New York, NY: ACM, pp. 1–12.

Magalhães, J. P., Dijkstra, A., Jeuring, J. & Löh, A. (2010) A generic deriving mechanism for Haskell. In *Haskell Symposium*. New York, NY: ACM, pp. 37–48.

McBride, C. & Paterson, R. (2008) Applicative programming with effects. *J. Funct. Program.* **18**(1), 1–13.

Mitchell, N. & Runciman, C. (2007) Uniform boilerplate and list processing. In *Haskell Workshop*. New York, NY: ACM, pp. 49–60.

Moors, A., Piessens, F. & Joosen, W. (2006) An object-oriented approach to datatype-generic programming. In *Workshop on Generic Programming*. New York, NY: ACM, pp. 96–106.

van Noort, T., Rodriguez Yakushev, A., Holdermans, S., Jeuring, J., Heeren, B. & Magalhães, J. P. (2010) A lightweight approach to datatype-generic rewriting. *J. Funct. Program.* **20**(3–4), 375–413.

O'Connor, R. (2011) Functor is to Lens as Applicative is to Biplate: Introducing Multiplate. *Workshop on generic programming*. Available at http://arxiv.org/abs/1103.2841.

Oliveira, B. C. D. S. & Gibbons, J. (2010) Scala for generic programmers. *J. Funct. Program.* **20**(3–4), 303–352.

O'Sullivan, B. (2012). http://hackage.haskell.org/package/criterion.

Peyton Jones, S. (2007) Call-pattern specialisation for Haskell programs. In *International Conference on Functional Programming*. New York, NY: ACM, pp. 327–337.

Peyton Jones, S., Vytiniotis, D., Weirich, S. & Shields, M. (2007) Practical type inference for arbitrary-rank types. *J. Funct. Program.* **17**(1), 1–82.

Reynolds, J. C. (1974) Towards a theory of type structure. In *Colloque sur la programmation*, Robinet, B. (ed), Lecture Notes in Computer Science, vol. 19. Berlin: Springer, pp. 408–423.

Rodriguez Yakushev, A., Jeuring, J., Jansson, P., Gerdes, A., Kiselyov, O. & Oliveira, B. (2008) Comparing libraries for generic programming in Haskell. In *Haskell Symposium*. New York, NY: ACM, pp. 111–122.

Santos, A. (1995) *Compilation by Transformation in Non-Strict Functional Languages*. Ph.D. thesis, University of Glasgow, Glasgow.

Schmidt, U., Schmidt, M. & Kuseler, T. (2012) http://hackage.haskell.org/package/hxt.

Sculthorpe, N., Farmer, A. & Gill, A. (2013) The HERMIT in the tree: Mechanizing program transformations in the GHC core language. In *International Symposium on Implementation and Application of Functional Languages 2012*, Hinze, R. (ed), Lecture Notes in Computer Science, vol. 8241. Berlin: Springer, pp. 86–103.

Sculthorpe, N. & Gill, A. (2014) http://hackage.haskell.org/package/kure.

Sheard, T. & Peyton Jones, S. (2002) Template metaprogramming for Haskell. In *Haskell Workshop*. New York, NY: ACM, pp. 1–16.

Strachey, C. (1967) *Fundamental Concepts in Programming Languages*. Lecture Notes, International Summer School in Computer Programming, Copenhagen. Reprinted in *Higher-Order and Symbolic Computation*, **13**(1–2), 11–49, 2000.

Sulzmann, M., Chakravarty, M. M. T., Peyton Jones, S. & Donnelly, K. (2007) System F with type equality coercions. In *International Workshop on Types in Language Design and Implementation*. New York, NY: ACM, pp. 53–66.

Swierstra, W. (2008) Data types à la carte. *J. Funct. Program*, **18**(4), 423–436.

Tseitin, G. S. (1968). On the complexity of derivations in the propositional calculus. In *Structures in Constructive Mathematics and Mathematical Logic, Part II*, Slisenko, A.O. (ed), Seminars in Mathematics, vol. 8. Moscow: Steklov Mathematical Institute, pp. 115–125. Reprinted in Pages 466–483 of: *Automation of reasoning*, Springer, 1983.

Visser, E. (2004) Program transformation with Stratego/XT: Rules, strategies, tools, and systems in StrategoXT-0.9. In *International Seminar on Domain-Specific Program Generation 2003*, Lengauer, C., Batory, D., Consel, C. & Odersky, M. (eds), Lecture Notes in Computer Science, vol. 3016. Berlin: Springer, pp. 216–238.

Visser, E. (2005) A survey of strategies in rule-based program transformation systems. *J. Symb. Comput.* **40**(1), 831–873.

Visser, E. (2013). Personal communication.

Visser, E., Benaissa, Zine-el-Abidine, & Tolmach, A. (1998) Building program optimizers with rewriting strategies. In *International Conference on Functional Programming*. New York, NY: ACM, pp. 13–26.

Wadler, P. (1992) The essence of functional programming. *Symposium on Principles of Programming Languages*. New York, NY: ACM, pp. 1–14.

Wadler, P. & Blott, S. (1989) How to make ad-hoc polymorphism less ad hoc. In *Symposium on Principles of Programming Languages*. New York, NY: ACM, pp. 60–76.

## Appendix. Traversing similarly typed nodes distinctly

Sometimes it is desirable to define a traversal that descends into multiple nodes of the same type, yet treats them distinctly based on their location. This is a generalisation of the technique in Section 4.3.3, which only made a binary choice of *whether* to descend into a node or not, based on its location. This appendix will demonstrate how such a traversal can be defined in KURE, using the String language from Section 4.3 as an example.

Consider the (somewhat contrived) situation where a user wants to apply a rewrite that operates on substrings of literals, and another that operates on substrings of variable names, during the same traversal. The first step is to define abstract types to represent names and literals, thereby providing a type distinction between them:

```
newtype StringLit   = StringLit Literal
newtype StringName = StringName Name
```

Any rewrites that should succeed for only one of *Literal* or *Name* should then be defined over these abstract types, rather than over *String*. The user then defines a universe, with separate summands for names and literals:

```
data U_2 = UP_2 Prog | UD_2 Decl | UE_2 Expr | UL_2 Literal | UN_2 Name
```

The *Injection* instances for $U_2$ use the *StringLit* and *StringName* types, rather than *String*. This is necessary to avoid ambiguity, as otherwise there would be two overlapping *Injection* instances for *String*.

However, we now cannot use *extractR* as directly as before in the definition of *allR*. Instead we define two variants of *extractR*, one for names and one for literals, which additionally add and remove the **newtype** wrappers. The variant for literals is as follows:

```
extractRlit :: Monad m ⇒ R m U₂ → R m Literal
extractRlit r = λl → unLit <$> (extractR r) (StringLit l)
   where
      unLit (StringLit l) = l
```

We then use these functions to define the Lit and Var clauses of *allExpr*, as well as two new local functions, *allName* (not shown here) and *allLit*:

```
allExpr (Var n) = Var <$> (extractRname r) n
allExpr (Lit l)  = Lit  <$> (extractRlit r) l
allLit :: R m Literal
allLit []       = pure []
allLit (x : xs) = (:) <$> pure x <*> (extractRlit r) xs
```

A traversal using the $U_2$ universe will now descend into both names and literals, but, for example, will only apply a rewrite of type *R m StringLit* to substrings of literals, not to substrings of names.