FUNCTIONAL PEARL

Typed quote/antiquote or: Compile-time parsing

RALF HINZE

*Department of Computer Science, University of Oxford, Wolfson Building, Parks Road,
Oxford OX1 3QD, England, UK
(e-mail: ralf.hinze@cs.ox.ac.uk)*

1 Introduction

Haskell (Peyton Jones, 2003) is often used as a host language for embedding other languages. Typically, the abstract syntax of the guest language is defined by a collection of datatype declarations; parsers and pretty-printers convert between the concrete syntax and its abstract representation. A quote/antiquote mechanism permits a tighter integration of the guest language into the host language by allowing one to use phrases in the guest language's *concrete syntax*.

For a simple example, assume that the abstract syntax of the guest language is given by the following datatype of binary trees.

data *Tree* = *Leaf* | *Fork Tree Tree*

To dispense with the need for parentheses, we choose prefix notation for the concrete syntax. The following interactive session illustrates the use of quotations.

```
Main⟩   « fork fork leaf leaf leaf »
Fork (Fork Leaf Leaf ) Leaf
Main⟩   size (« fork fork leaf leaf leaf ») + 1
4
```

A *quotation* is delimited by guillemets (« and ») and consists of a phrase in concrete syntax, in our case, a prefix expression. The concrete syntax is a sequence of terminal symbols, written in typewriter font. A quotation evaluates to abstract syntax and can be freely mixed with ordinary Haskell expressions. In our example, a quotation yields a value of type *Tree* and may therefore serve as an argument to *size*, which computes the size of a tree.

Perhaps surprisingly, our quote mechanism guarantees that the guest-language phrase is well formed: the malformed quotations « fork » and « leaf leaf » are both rejected by Haskell's type-checker. This is a big advantage over the use of strings, which often serve as an approximation to quotations.

The relationship between host and guest language also suggests a notion of antiquotation: the ability to splice a host-language expression into the middle of a guest-language phrase. Continuing the example above, here is a session that

demonstrates the use of antiquotations:

> *Main*⟩   « fork ` (*full* 2) leaf »
> *Fork* (*Fork* (*Fork Leaf Leaf*) (*Fork Leaf Leaf*)) *Leaf*
> *Main*⟩   **let** *foo t* = « fork ` *t* leaf »
> *Main*⟩   *foo* (« fork leaf fork leaf leaf »)
> *Fork* (*Fork Leaf* (*Fork Leaf Leaf*)) *Leaf*

An *antiquotation* is written as a back-quote (`) followed by an atomic Haskell expression, for instance, an identifier or a parenthesised expression. The Haskell expression typically generates a piece of abstract syntax, for instance, in the first expression above, a fully balanced binary tree of depth 2.

A quote/antiquote mechanism usually requires an extension of the host language. The purpose of this pearl is to show that one can program such a mechanism within Haskell itself. The technique is based on Okasaki's flattening combinators (Okasaki, 2002, 2003), which we shall review in the next section. To make the idea fly, I assume that we can use an arbitrary terminal symbol in typewriter font *as a Haskell identifier*. If you think that this assumption undermines the argument, then you should read the pearl as an exercise in compile-time parsing.

2 Background: the other continuation-passing style monad

To illustrate the basic idea consider a very simple example, which implements concrete syntax for the naturals.

> *Main*⟩   (« | | | », « | | | | | » + 7)
> (3, 12)

We have only one terminal symbol, the vertical bar, where a sequence of *n* bars represents the number *n*.

The succession of symbols « | | | » looks like a sequence of terminals enclosed in guillemets. But, of course, this is an illusion; the sequence is, in fact, a nested application of functions. If we take "«," "»," and "|" as aliases for *quote*, *endquote*, and *tick*, then « | | | » abbreviates the fully parenthesised expression (((*quote tick*) *tick*) *tick*) *endquote*. In what follows, we shall use "«" and *quote*, "»" and *endquote*, "`" and *antiquote* interchangeably.

Now, if Haskell used postfix function application, then we could simply define *quote* = 0, *tick* = *succ*, *endquote* = *id* and we would be done. For Haskell's prefix function application, we must additionally arrange that functions and arguments are swapped:

> *quote*   *f* = *f* 0
> *tick*   *n f* = *f* (*succ n*)
> *endquote n*   = *n*

The stepwise evaluation of « | | | » shows that *tick* increments the counter, initialised to 0 by *quote*, and then passes control to the next function, which is either

another *tick* or *endquote*.

$$\begin{aligned}
& \textit{quote tick tick tick endquote} \\
=\ & \textit{tick } 0 \textit{ tick tick endquote} \\
=\ & \textit{tick } 1 \textit{ tick endquote} \\
=\ & \textit{tick } 2 \textit{ endquote} \\
=\ & \textit{endquote } 3 \\
=\ & 3
\end{aligned}$$

The evaluation is solely driven by the terminal symbols, which is why we call them *active terminals.* This technique of passing control to a function argument is reminiscent of *continuation-passing style* (CPS). And indeed, if we call the *CPS* Monad to mind[1]

> **type** *CPS* $\alpha = \forall ans \,.\, (\alpha \rightarrow ans) \rightarrow ans$
>
> **instance** *Monad CPS* **where**
> $return\ a = \lambda\kappa \rightarrow \kappa\ a$
> $m \ggg k \;= \lambda\kappa \rightarrow m\ (\lambda a \rightarrow k\ a\ \kappa)$

we can identify *quote* as *return* 0 and *tick* as *lift succ* where *lift* turns a pure function into a monadic one:

> **type** $\alpha \rightarrowtail \beta = \alpha \rightarrow CPS\ \beta$
>
> $lift \qquad :: (\alpha \rightarrow \beta) \rightarrow (\alpha \rightarrowtail \beta)$
> $lift\ f\ a = return\ (f\ a)$

However, the bind of the monad, "$\ggg$," seems unrelated: in the *CPS* monad the continuation represents the "rest of the computation" whereas in our example the continuation only stands for the next parsing step.

    It may come as a surprise that the instance declaration above is not the only possibility for turning *CPS* into a monad. Here is a second instance introducing the monad of *partial continuations.*

> **instance** *Monad CPS* **where**
> $return\ a = \lambda\kappa \rightarrow \kappa\ a$
> $m \ggg k \;= m\ k$

The definition of *return* is unchanged; "$\ggg$" is now a type-restricted instance of function application. Actually, it is a combination of type application—the universally quantified variable *ans* in the type of *m* is instantiated to *CPS* β— and function application, but this is not visible in Haskell. Since "$\ggg$" is postfix application of "effectful" functions, this CPS monad implements postfix function application! Consequently, the quotation « | | | » can be seen as a monadic computation in disguise:

$$(((quote\ tick)\ tick)\ tick)\ endquote \quad = \quad run\ (quote \ggg tick \ggg tick \ggg tick)$$

---

[1] The instance declaration, which is not legal Haskell, serves only illustrative purposes. We shall only need *return* and only at this particular type.

where *run* encapsulates a *CPS* computation:

$$run \quad :: CPS \; \alpha \rightarrow \alpha$$
$$run \; m = m \; id$$

Generalising from the example, quotations are based on the identity

$$quote \; act_1 \; \dots \; act_n \; endquote \quad = \quad run \; (quote \ggg act_1 \ggg \dots \ggg act_n)$$

where $quote :: CPS \; \tau_1$, $act_i :: \tau_i \rightarrowtail \tau_{i+1}$, and $endquote = id$. It is useful to think of the τ_i as state types and the act_i as transitions: *quote* initialises the state; each active terminal act_i transforms the state. In our example, we had a single state type but this need not be the case in general. In fact, choosing precise state types is the key to "typed quotes/antiquotes."

Just in case you were wondering, none of the two *CPS* monads is a very exciting one in terms of expressiveness: they are both isomorphic to the *identity monad* with *return* and *run* converting between them. In other words, *CPS* offers no effects. Without loss of generality, we may therefore assume that *quote* and act_i are liftings: $quote = return \; a$ and $act_i = lift \; f_i$ for some a and suitable f_i. The following calculation summarises our findings:

$$quote \; act_1 \; \dots \; act_n \; endquote$$
$$= \quad \{ \text{ definition of '} \ggg \text{' and } run \text{, and } endquote = id \; \}$$
$$run \; (quote \ggg act_1 \ggg \dots \ggg act_n)$$
$$= \quad \{ \; CPS \text{ is a pure monad: } quote = return \; a \text{ and } act_i = lift \; f_i \; \}$$
$$run \; (return \; a \ggg lift \; f_1 \ggg \dots \ggg lift \; f_n)$$
$$= \quad \{ \text{ monad laws } \}$$
$$run \; (return \; (f_n \; (\dots \; (f_1 \; a) \; \dots)))$$
$$= \quad \{ \; run \cdot return = id \; \}$$
$$f_n \; (\dots \; (f_1 \; a) \; \dots)$$

In the toy example and in the formal development above, *endquote* was always the identity. This is, however, not quite adequate, as the desired value of a quotation is not necessarily identical to the last state. Fortunately, *endquote* can be any function since we can fuse a post-processing step with the final continuation: $post \; (run \; m) = m \; post$. This is an immediate consequence of the free theorem for the type $CPS \; \alpha$ (Wadler, 1989).

To summarise, a quotation of type τ is of the form

$$quote \; act_1 \; \dots \; act_n \; endquote$$

where $quote :: CPS \; \tau_1$, $act_i :: \tau_i \rightarrowtail \tau_{i+1}$ and $endquote :: \tau_{n+1} \rightarrow \tau$.

Since the evaluation of a quotation is driven by the terminal symbols, the implementation of a quote/antiquote mechanism for a particular guest language goes hand in hand with the development of a parser for the concrete syntax. The

following sections are ordered by the underlying parser's level of sophistication: Section 3 shows how to implement simple postfix and prefix parsers, Section 4 deals with predictive top-down parsers, and finally, Section 5 introduces quotations that are based on bottom-up parsers.

3 Parsing datatypes

Continuing the example from the introduction, we show how to parse elements of datatypes in postfix and in prefix notation. Section 3.1 is an excerpt of Okasaki's extensive treatment of postfix languages (see Okasaki, 2002).

3.1 Postfix notation

In postfix notation, also known as reverse Polish notation, functions follow their arguments. Postfix notation dispenses with the need for parentheses, if the arity of functions is statically known. This is generally not the case in higher order typed languages, but it is true of **data** constructors (ignoring the fact that they are curried in Haskell).

Evaluation of postfix expressions is naturally stack-based: a function pops its arguments from the stack and pushes the result back onto it. To parse datatypes in postfix notation, we introduce for each data constructor $C :: \tau_1 \to \cdots \to \tau_n \to \tau$ a *postfix constructor*:

$$
\begin{array}{l}
c :: (((st, \tau_1), \ldots), \tau_n) \to (st, \tau) \\
c \quad (((st, t_1), \ldots), t_n) = (st, C \ t_1 \ \ldots \ t_n)
\end{array}
$$

The stack, represented by a nested pair, grows from left to right. The modification of the stack is precisely captured in the type: c is only applicable if the stack contains at least n arguments and the topmost n have the correct types. For the *Tree* type, this specialises to

$$
\begin{array}{ll}
leaf & :: st \to (st, Tree) \\
leaf \ st & = (st, Leaf) \\
fork & :: ((st, Tree), Tree) \to (st, Tree) \\
fork \ ((st, l), r) = (st, Fork \ l \ r)
\end{array}
$$

Given these prerequisites, we can instantiate the framework of Section 2.

$$
\begin{array}{ll}
quote & :: CPS \ () \\
quote & = return \ () \\
\texttt{leaf} & :: st \twoheadrightarrow (st, Tree) \\
\texttt{leaf} & = lift \ leaf \\
\texttt{fork} & :: ((st, Tree), Tree) \twoheadrightarrow (st, Tree) \\
\texttt{fork} & = lift \ fork \\
endquote & :: ((), Tree) \to Tree \\
endquote \ ((), t) = t
\end{array}
$$

The function *quote* initialises the state to the empty stack; *endquote* extracts the quoted tree from a singleton stack.

It is instructive to step through the static and dynamic elaboration of a quotation. Type checking statically guarantees that a quotation constitutes a well-formed postfix expression.

$$
\begin{aligned}
&\textit{quote} &&:: CPS\ () \\
&\textit{quote}\ \texttt{leaf} &&:: CPS\ ((),\textit{Tree}) \\
&\textit{quote}\ \texttt{leaf leaf} &&:: CPS\ (((),\textit{Tree}),\textit{Tree}) \\
&\textit{quote}\ \texttt{leaf leaf fork} &&:: CPS\ ((),\textit{Tree}) \\
&\textit{quote}\ \texttt{leaf leaf fork leaf} &&:: CPS\ (((),\textit{Tree}),\textit{Tree}) \\
&\textit{quote}\ \texttt{leaf leaf fork leaf fork} &&:: CPS\ ((),\textit{Tree}) \\
&\textit{quote}\ \texttt{leaf leaf fork leaf fork}\ \textit{endquote} &&:: \textit{Tree}
\end{aligned}
$$

In each step, the state type precisely mirrors the stack layout. Consequently, pushing too few or too many or the wrong types of arguments results in a static type-error. The dynamic evaluation shows how the state evolves.

$$
\begin{aligned}
&\quad \textit{quote}\ \texttt{leaf leaf fork leaf fork}\ \textit{endquote} \\
&= \texttt{leaf}\ ()\ \texttt{leaf fork leaf fork}\ \textit{endquote} \\
&= \texttt{leaf}\ ((),\textit{Leaf}\,)\ \texttt{fork leaf fork}\ \textit{endquote} \\
&= \texttt{fork}\ (((),\textit{Leaf}\,),\textit{Leaf}\,)\ \texttt{leaf fork}\ \textit{endquote} \\
&= \texttt{leaf}\ ((),\textit{Fork}\ \textit{Leaf}\ \textit{Leaf}\,)\ \texttt{fork}\ \textit{endquote} \\
&= \texttt{fork}\ (((),\textit{Fork}\ \textit{Leaf}\ \textit{Leaf}\,),\textit{Leaf}\,)\ \textit{endquote} \\
&= \textit{endquote}\ ((),\textit{Fork}\ (\textit{Fork}\ \textit{Leaf}\ \textit{Leaf}\,)\ \textit{Leaf}\,) \\
&= \textit{Fork}\ (\textit{Fork}\ \textit{Leaf}\ \textit{Leaf}\,)\ \textit{Leaf}
\end{aligned}
$$

The state is always passed as the first argument. This is something to bear in mind when implementing additional functionality, such as an antiquote mechanism.

$$
\begin{aligned}
&\textit{antiquote} &&:: st \to \textit{Tree} \to (st,\textit{Tree}) \\
&\textit{antiquote}\ st\ t = \textit{return}\ (st,t)
\end{aligned}
$$

The tree is spliced into the current position simply by pushing it onto the stack.

3.2 Prefix notation

Postfix notation was easy; its dual, prefix notation, is slightly harder. Prefix notation was invented in 1920 by Jan Łukasiewicz, a Polish logician, mathematician, and philosopher. Because of its origin, prefix notation is also known as Polish notation.

In postfix notation, a function follows its arguments; so a stack of arguments is a natural choice for the state. In prefix notation, a function precedes its arguments. Consequently, the state becomes a stack of *pending* arguments. For each data constructor $C :: \tau_1 \to \cdots \to \tau_n \to \tau$, we introduce a *prefix constructor*:

$$
\begin{aligned}
&c^\circ :: ((\tau,st) \to \alpha) \to ((\tau_1,(\ldots,(\tau_n,st)))) \to \alpha) \\
&c^\circ \quad ctx \qquad = \lambda(t_1,(\ldots,(t_n,st))) \to ctx\ (C\ t_1\ \ldots\ t_n,st)
\end{aligned}
$$

The stack, again represented by a nested pair, now grows from right to left. The first argument of c° can be seen as a request for a value of type τ (a request is also known as a context or as an expression with a hole). The prefix constructor can satisfy this request but, in turn, generates requests for its arguments.[2] For the *Tree* type, we obtain

$$
\begin{aligned}
&leaf^\circ &&:: ((Tree, st) \to \alpha) \to (st \to \alpha) \\
&leaf^\circ\ ctx = \lambda st \to ctx\ (Leaf, st) \\
&fork^\circ &&:: ((Tree, st) \to \alpha) \to ((Tree, (Tree, st)) \to \alpha) \\
&fork^\circ\ ctx = \lambda(t, (u, st)) \to ctx\ (Fork\ t\ u, st)
\end{aligned}
$$

The implementation of quotations and antiquotations is again a straightforward application of the general framework:

$$
\begin{aligned}
&quote &&:: CPS\ ((Tree, ()) \to Tree) \\
&quote &&= return\ (\lambda(t, ()) \to t) \\
&\texttt{leaf} &&:: ((Tree, st) \to \alpha) \rightarrowtail (st \to \alpha) \\
&\texttt{leaf} &&= lift\ leaf^\circ \\
&\texttt{fork} &&:: ((Tree, st) \to \alpha) \rightarrowtail ((Tree, (Tree, st)) \to \alpha) \\
&\texttt{fork} &&= lift\ fork^\circ \\
&endquote &&:: (() \to Tree) \to Tree \\
&endquote\ ctx = ctx\ () \\
&antiquote &&:: ((Tree, st) \to \alpha) \to Tree \rightarrowtail (st \to \alpha) \\
&antiquote\ ctx\ t = return\ (\lambda st \to ctx\ (t, st))
\end{aligned}
$$

The stack is initialised to one pending argument; we are done if there are no pending arguments left. Let us again step through an example.

$$
\begin{aligned}
&quote &&:: CPS\ (((Tree, ()) \to Tree) \\
&quote\ \texttt{fork} &&:: CPS\ ((Tree, (Tree, ())) \to Tree) \\
&quote\ \texttt{fork}\ \texttt{fork} &&:: CPS\ ((Tree, (Tree, (Tree, ()))) \to Tree) \\
&quote\ \texttt{fork}\ \texttt{fork}\ \texttt{leaf} &&:: CPS\ ((Tree, (Tree, ())) \to Tree) \\
&quote\ \texttt{fork}\ \texttt{fork}\ \texttt{leaf}\ \texttt{leaf} &&:: CPS\ (((Tree, ()) \to Tree) \\
&quote\ \texttt{fork}\ \texttt{fork}\ \texttt{leaf}\ \texttt{leaf}\ \texttt{leaf} &&:: CPS\ (() \to Tree) \\
&quote\ \texttt{fork}\ \texttt{fork}\ \texttt{leaf}\ \texttt{leaf}\ \texttt{leaf}\ endquote &&:: CPS\ Tree
\end{aligned}
$$

The types show how the stack of pending arguments grows and shrinks. For instance, when the first two `fork`s have been processed, three further subtrees are required: the left and the right subtree of the second `fork` and the right subtree of the first `fork`. The stepwise evaluation makes this explicit:

---

[2] The type variable α that appears in the type signature of c° corresponds to the type of the entire quotation and can be safely instantiated to *Tree*. The polymorphic type is only vital if c° is used in quotations of different types.

$\quad$ *quote* `fork fork leaf leaf leaf` *endquote*
$= $ `fork` $(\lambda(t,()) \to t)$ `fork leaf leaf leaf` *endquote*
$= $ `fork` $(\lambda(t,(u,())) \to Fork\ t\ u)$ `leaf leaf leaf` *endquote*
$= $ `leaf` $(\lambda(t',(u',(u,()))) \to Fork\ (Fork\ t'\ u')\ u)$ `leaf leaf` *endquote*
$= $ `leaf` $(\lambda(u',(u,())) \to Fork\ (Fork\ Leaf\ u')\ u)$ `leaf` *endquote*
$= $ `leaf` $(\lambda(u,()) \to Fork\ (Fork\ Leaf\ Leaf)\ u)$ *endquote*
$= $ *endquote* $(\lambda() \to Fork\ (Fork\ Leaf\ Leaf)\ Leaf)$
$= Fork\ (Fork\ Leaf\ Leaf)\ Leaf$

Again, if we pass too few or too many or the wrong types of arguments, then we get a static type-error.

Remark 1
The deeply nested pairs can be avoided if we curry the prefix constructors:

$$c^\circ :: (\tau \to \alpha) \to (\tau_1 \to \cdots \to \tau_n \to \alpha)$$
$$c^\circ \quad ctx \quad = \quad \lambda t_1 \to \cdots \to \lambda t_n \to ctx\ (C\ t_1\ \ldots\ t_n)$$

Additionally, we have generalised the result type of requests from $st \to \alpha$ to α. The adaptation of the remaining code is left as an exercise to the reader. $\qquad\square$

4 Top-down parsing

The main reason for treating prefix parsers is that they pave the way for the more expressive class of LL(1) parsers. The basic setup remains unchanged; we need only one additional programming technique. To keep the learning curve smooth, however, we shall go through one intermediate step and treat grammars in Greibach normal form (GNF) first.

4.1 Greibach normal form

A context-free grammar is in GNF iff all productions are of the form $A \to a\omega$, where a is a terminal symbol and ω is a possibly empty sequence of nonterminal symbols. A grammar in GNF is (syntactically) *unambiguous* iff each pair of productions $A_1 \to a\omega_1$ and $A_2 \to a\omega_2$ satisfies $A_1 = A_2 \implies \omega_1 = \omega_2$. Unambiguous grammars in GNF generalise datatype declarations, as a terminal (data constructor) may appear in different productions (datatype declarations).

$\quad$ Here is an example grammar for a simple imperative language and an equivalent grammar in GNF.

$$
\begin{array}{rcl}
S & \to & \texttt{id} := E \\
  & | & \texttt{if}\ E\ S\ S \\
  & | & \texttt{while}\ E\ S \\
  & | & \texttt{begin}\ B\ \texttt{end} \\
E & \to & \texttt{id} \\
B & \to & S\ |\ S\ ;\ B
\end{array}
\qquad \Longrightarrow \qquad
\begin{array}{rcl}
S & \to & \texttt{id}\ C\ E \\
  & | & \texttt{if}\ E\ S\ S \\
  & | & \texttt{while}\ E\ S \\
  & | & \texttt{begin}\ S\ R \\
C & \to & := \\
E & \to & \texttt{id} \\
R & \to & \texttt{end}\ |\ ;\ S\ R
\end{array}
$$

Remark 2

If we add the production $S \rightarrow$ if E S, then the grammar becomes ambiguous, an instance of the (in-) famous *dangling-else problem*. $\qquad\square$

The abstract syntax of the imperative language is given by

type *Var* $=$ *String*
data *Stat* $=$ *Set Var Var* | *If Var Stat Stat* | *While Var Stat* | *Begin* [*Stat*]

As an example, the quotation

```
« begin
  x := y ;
  if x
     y := z
     z := y
  end »
```

evaluates to *Begin* [*Set* "x" "y", *If* "x" (*Set* "y" "z") (*Set* "z" "y")]. The definition of the variables x, y, and z poses a minor problem, which we shall discuss later.

A parser for a grammar in GNF is very similar to a prefix parser: the state is a stack of *pending* nonterminal symbols. An active terminal selects a production by looking at the topmost symbol on the stack. If the grammar is unambiguous, then there is at most one suitable production. The nonterminal is then replaced by the right-hand side of the production (omitting the leading terminal).

As before, we want to guarantee statically that a quotation is well formed, so that its parse does not fail. To this end, we represent nonterminals by types:

newtype S α $=$ S (*Stat* $\rightarrow \alpha$)
newtype C α $=$ C (α)
newtype E α $=$ E (*Var* $\rightarrow \alpha$)
newtype R α $=$ R ([*Stat*] $\rightarrow \alpha$)

The declarations also list the types of the semantic values that are attached to the nonterminals. Using these type-level nonterminals, we can program the type-checker to parse quotations: each production $A \rightarrow a B_1 \ldots B_n$ is mapped to a function a, the active terminal, of type A $\alpha \twoheadrightarrow B_1$ ($\cdots$ (B_n α) $\cdots$), that implements the expansion of A. There is one hitch, however: the terminal a may appear in different productions, so it cannot possibly translate to a single function. Rather, an active terminal stands for a family of functions, represented in Haskell by a multiple-parameter type-class.

We introduce one class for each terminal symbol that appears more than once. In our case, the only such terminal is id, so we need just one class.

class *Id lhs rhs* | *lhs* $\rightarrow$ *rhs* **where**
    id :: *String* $\rightarrow$ (*lhs* $\twoheadrightarrow$ *rhs*)

The functional dependency *lhs* $\rightarrow$ *rhs* avoids ambiguities during type-inference making use of the fact that the underlying grammar in GNF is unambiguous.

Each production is translated into an equation.

instance *Id* (*S* α) (*C* (*E* α)) **where**
   id *l* = *lift* (λ(*S ctx*) → *C* (*E* (λr → *ctx* (*Set l r*))))
if   = *lift* (λ(*S ctx*) → *E* (λc → *S* (λt → *S* (λe → *ctx* (*If c t e*)))))
while = *lift* (λ(*S ctx*) → *E* (λc → *S* (λs → *ctx* (*While c s*))))
begin = *lift* (λ(*S ctx*) → *S* (λs → *R* (λr → *ctx* (*Begin* (*s* : *r*)))))
:=   = *lift* (λ(*C ctx*) → *ctx*)
instance *Id* (*E* α) α **where**
   id *i* = *lift* (λ(*E ctx*) → *ctx i*)
end   = *lift* (λ(*R ctx*) → *ctx* [])
;   = *lift* (λ(*R ctx*) → *S* (λs → *R* (λr → *ctx* (*s* : *r*))))
quote   = *return* (*S* (λs → *s*))
endquote s = *s*

The *quote* function pushes the start symbol on the stack; *endquote* simply returns the final value of type *Stat*.

The terminal symbol id is a bit unusual in that it takes an additional argument, the string representation of the identifier. This has the unfortunate consequence that identifiers must be enclosed in parentheses as in « (id "x") := (id "y") ». We can mitigate the unwelcome effect by introducing shortcuts

x, y :: (*Id lhs rhs*) ⇒ *lhs* → *rhs*
x = id "x"
y = id "y"

so that the quotation becomes « x := y ». Alternatively, we may swap the two arguments of id. In this case, the parentheses can, in fact, must be dropped, so that the quotation is written « id "x" := id "y" ».

It is instructive to walk through a derivation.

    « while x y := z »
= while (*S* (λs → *s*)) x y := z »
= x (*E* (λc → *S* (λs → *While c s*))) y := z »
= y (*S* (λs → *While* "x" *s*)) := z »
= := (*C* (*E* (λr → *While* "x" (*Set* "y" *r*)))) z »
= z (*E* (λr → *While* "x" (*Set* "y" *r*))) »
= » (*While* "x" (*Set* "y" "z"))
= *While* "x" (*Set* "y" "z")

To summarise, for each nonterminal symbol *A*, we define a type: **newtype** *A* α = *A* (*Val* → α), where *Val* is the type of the semantic values attached to *A*. For each terminal symbol *a*, we introduce a class: **class** *a lhs rhs* | *lhs* → *rhs* **where** *a* :: *lhs* → *rhs*. Finally, each production *A* → *aB_1 . . . B_n* gives rise to an instance declaration:

instance *a* (*A* α) (*B_1* (· · · (*B_n* α) · · ·)) **where**
   a = *lift* (λ(*A ctx*) → *B_1* (λv_1 → · · · → *B_n* (λv_n → *ctx* (*f v_1 . . . v_n*))))

where *f* is the semantic action associated with the production. Of course, if a terminal appears only once, then there is no need for overloading.

```
newtype I  α = I  (Expr              → α)   -- id
newtype A  α = A  (                  α)   -- +
newtype M  α = M  (                  α)   -- *
newtype O  α = O  (                  α)   -- (
newtype C  α = C  (                  α)   -- )
newtype E  α = E  (Expr              → α)
newtype E' α = E' ((Expr → Expr) → α)
newtype T  α = T  (Expr              → α)
newtype T' α = T' ((Expr → Expr) → α)
newtype F  α = F  (Expr              → α)
class Id    old new | old → new where id :: String → old ↠ new
class Add   old new | old → new where + ::        old ↠ new
class Mul   old new | old → new where * ::        old ↠ new
class Open  old new | old → new where ( ::        old ↠ new
class Close old new | old → new where ) ::        old ↠ new
class Endquote old              where » ::        old → Expr
```

Fig. 1. The LL(1) parser for the expression grammar, part 1.

4.2 LL(1) *parsing*

We are well prepared by now to tackle the first major challenge: implementing quotations whose syntax is given by an LL(1) grammar. As before, we shall work through a manageable example. This time, we implement arithmetic expressions given by the grammar on the left below (see Aho *et al.* 2006, p. 193).

$$
\begin{array}{ccc}
\begin{aligned}
E &\rightarrow E + T \mid T \\
T &\rightarrow T * F \mid F \\
F &\rightarrow (\ E\ ) \mid id
\end{aligned}
& \Longrightarrow &
\begin{aligned}
E &\rightarrow T\ E' \\
E' &\rightarrow +\ T\ E' \mid \epsilon \\
T &\rightarrow F\ T' \\
T' &\rightarrow *\ F\ T' \mid \epsilon \\
F &\rightarrow (\ E\ ) \mid id
\end{aligned}
\end{array}
$$

The expression grammar is *not* LL(1) due to the left recursion; eliminating the left recursion yields the equivalent LL(1) grammar on the right.

The abstract syntax of arithmetic expressions is given by

data *Expr* = *Id String* | *Add Expr Expr* | *Mul Expr Expr*

The semantic actions that construct values of type *Expr* are straightforward to define for the original expression grammar. They are slightly more involved for the LL(1) grammar: E' and T' yield expressions with a hole, where the hole stands for the missing left argument of the operator (see Figures 1 and 2).

The main ingredient of a predictive top-down parser is the *parsing table*. Here is the table for the LL(1) grammar above (see Aho *et al.* 2006, p. 225).

	id	+	*	(	)	»
E	$E \to T\ E'$			$E \to T\ E'$		
E'		$E' \to +\ T\ E'$			$E' \to \epsilon$	$E' \to \epsilon$
T	$T \to F\ T'$			$T \to F\ T'$		
T'		$T' \to \epsilon$	$T' \to *\ F\ T'$		$T' \to \epsilon$	$T' \to \epsilon$
F	$F \to id$			$F \to (\ E\ )$		

instance Id $(E$ $\alpha)$ $(T'$ $(E'$ $\alpha))$ **where**   -- expand: $E \to T\ E'$
  id s $(E$ $ctx)$ = id s $(T$ $(\lambda t \to E'$ $(\lambda e' \to ctx$ $(e'\ t))))$
instance Id $(T$ $\alpha)$ $(T'$ $\alpha)$ **where**   -- expand: $T \to F\ T'$
  id s $(T$ $ctx)$ = id s $(F$ $(\lambda f \to T'$ $(\lambda t' \to ctx$ $(t'\ f))))$
instance Id $(F$ $\alpha)$ α **where**   -- expand: $F \to id$
  id s $(F$ $ctx)$ = id s $(I$ $(\lambda v \to ctx$ $v))$
instance Id $(I$ $\alpha)$ α **where**   -- pop
  id s $(I$ $ctx)$ = $return$ $(ctx$ $(Id$ $s))$

instance Add $(E'$ $\alpha)$ $(T$ $(E'$ $\alpha))$ **where**   -- expand: $E' \to + T\ E'$
  + $(E'$ $ctx)$ = + $(A$ $(T$ $(\lambda t \to E'$ $(\lambda e' \to ctx$ $(\lambda l \to e'$ $(Add\ l\ t))))))$
instance $(Add$ α $\alpha')$ $\Rightarrow$ Add $(T'$ $\alpha)$ α' **where**   -- expand: $T' \to \epsilon$
  + $(T'$ $ctx)$ = + $(ctx$ $(\lambda e \to e))$
instance Add $(A$ $\alpha)$ α **where**   -- pop
  + $(A$ $ctx)$ = $return$ ctx

instance Mul $(T'$ $\alpha)$ $(F$ $(T'$ $\alpha))$ **where**   -- expand: $T' \to * F\ T'$
  * $(T'$ $ctx)$ = * $(M$ $(F$ $(\lambda f \to T'$ $(\lambda t' \to ctx$ $(\lambda l \to t'$ $(Mul\ l\ f))))))$
instance Mul $(M$ $\alpha)$ α **where**   -- pop
  * $(M$ $ctx)$ = $return$ ctx

instance $Open$ $(E$ $\alpha)$ $(E$ $(C$ $(T'$ $(E'$ $\alpha))))$ **where**   -- expand: $E \to T\ E'$
  ($(E$ $ctx)$ = ($(T$ $(\lambda t \to E'$ $(\lambda e' \to ctx$ $(e'\ t))))$
instance $Open$ $(T$ $\alpha)$ $(E$ $(C$ $(T'$ $\alpha)))$ **where**   -- expand: $T \to F\ T'$
  ($(T$ $ctx)$ = ($(F$ $(\lambda f \to T'$ $(\lambda t' \to ctx$ $(t'\ f))))$
instance $Open$ $(F$ $\alpha)$ $(E$ $(C$ $\alpha))$ **where**   -- expand: $F \to ($ E $)$
  ($(F$ $ctx)$ = ($(O$ $(E$ $(\lambda e \to C$ $(ctx$ $e))))$
instance $Open$ $(O$ $\alpha)$ α **where**   -- pop
  ($(O$ $ctx)$ = $return$ ctx

instance $(Close$ α $\alpha')$ $\Rightarrow$ $Close$ $(E'$ $\alpha)$ α' **where**   -- expand: $E' \to \epsilon$
) $(E'$ $ctx)$ =) $(ctx$ $(\lambda e \to e))$
instance $(Close$ α $\alpha')$ $\Rightarrow$ $Close$ $(T'$ $\alpha)$ α' **where**   -- expand: $T' \to \epsilon$
) $(T'$ $ctx)$ =) $(ctx$ $(\lambda e \to e))$
instance $Close$ $(C$ $\alpha)$ α **where**   -- pop
) $(C$ $ctx)$ = $return$ ctx

instance $(Endquote$ $\alpha)$ $\Rightarrow$ $Endquote$ $(E'$ $\alpha)$ **where**   -- expand: $E' \to \epsilon$
  » $(E'$ $ctx)$ = » $(ctx$ $(\lambda e \to e))$
instance $(Endquote$ $\alpha)$ $\Rightarrow$ $Endquote$ $(T'$ $\alpha)$ **where**   -- expand: $T' \to \epsilon$
  » $(T'$ $ctx)$ = » $(ctx$ $(\lambda e \to e))$
instance $Endquote$ $Expr$ **where**   -- pop
  » e = e

Fig. 2. The LL(1) parser for the expression grammar, part 2.

The table also includes a column for "»," which serves as an end marker.

   The state is now a stack of pending symbols, both terminal and nonterminal. The symbol X on top of the stack determines the action of the current active terminal a. If $X = a$, then the terminal pops X from the stack and passes control to the next active terminal (pop action). If X is a nonterminal, then the terminal looks up the production indexed by X and a in the parsing table, replaces X by the right-hand side of the production, and remains active (expand action). Again, we need not consider error conditions as the type-checker will statically guarantee that parsing does not fail.

Since the state is a stack of symbols, we must introduce types both for terminal and nonterminal symbols (we only show some representative examples here, the complete code is listed in Figures 1 and 2):

> **newtype** E $\alpha = E$ $(Expr \to \alpha)$ -- E
>
> **newtype** I $\alpha = I$ $(Expr \to \alpha)$ -- id

Like nonterminals, terminal symbols may carry semantic information: the terminal id s, for instance, returns the semantic value Id s of type $Expr$. For each terminal symbol including "»," we introduce a class and an instance that implements the pop action:

> **class** *Add old new* | *old* $\to$ *new* **where**
> $+ :: old \rightarrowtail new$
>
> **instance** *Add* $(A$ $\alpha)$ α **where**
> $+ (A$ $ctx) = return$ ctx

Finally, each entry of the parsing table gives rise to an instance declaration that implements the expand action. Here are the instances for "+":

> **instance** *Add* $(E'$ $\alpha)$ $(T$ $(E'$ $\alpha))$ **where**
> $+ (E'$ $ctx) = + (A$ $(T$ $(\lambda t \to E'$ $(\lambda e' \to ctx$ $(\lambda l \to e'$ $(Add$ l $t))))))$
>
> **instance** $(Add$ α $\alpha') \Rightarrow Add$ $(T'$ $\alpha)$ α' **where**
> $+ (T'$ $ctx) = + (ctx$ $(\lambda e \to e))$

Since the look-ahead token is unchanged, the second instance requires an additional context, *Add* α α', which accounts for the "recursive" call to "+." The first instance also contains a call to "+" but this occurrence can be statically resolved: it refers to the "pop instance." In general, an instance for the production $A \to \omega$ requires a context iff $\epsilon \in \mathscr{L}(\omega)$, as in this case the stack shrinks.[3] The instance head always reflects the parsing state *after* the final pop action. Consider the *id* instance

> **instance** *Id* $(E$ $\alpha)$ $(T'$ $(E'$ $\alpha))$ **where**
> id s $(E$ $ctx) = id$ s $(T$ $(\lambda t \to E'$ $(\lambda e' \to ctx$ $(e'$ $t))))$

The expansion phase proceeds $E \to T$ $E' \to F$ T' $E' \to id$ T' E'. Consequently, the instance head records that the state changes from E to T' E'.

It remains to implement *quote* and *antiquote*.

> *quote* $= return$ $(E$ $(\lambda e \to e))$
>
> *antiquote* $:: E$ $st \to Expr \rightarrowtail st$
> *antiquote* $(E$ $st)$ $e = return$ $(st$ $e)$

The type of *antiquote* dictates that we can only splice an expression into a position where the nonterminal E is expected. This can be achieved by enclosing the antiquotation in "(" and ")."

---

[3] The standard construction of parsing tables using *First* and *Follow* sets already provides the necessary information: the parsing table contains the production $A \to \omega$ for look-ahead a if either $a \in First(\omega)$, or $\epsilon \in \mathscr{L}(\omega)$ and $a \in Follow(A)$. Only in the second case is a context required.

$Main\rangle$ « x + (` ($foldr1$ Add [Id ($show$ i) | $i \leftarrow$ [1..3]])) + y »
Add (Add (Id "x") (Add (Id "1") (Add (Id "2") (Id "3")))) (Id "y")

Remark 3
We have implemented the parsing actions in a rather ad-hoc way. Alternatively, they
can be written using *lift* and monadic composition:

instance Id (E α) (T' (E' α)) **where**
    id s = id s $\circ$ $lift$ (λ(E ctx) $\to$ T ($\lambda t \to E'$ ($\lambda e' \to ctx$ (e' t)))))

($\circ$)   :: ($b \to CPS$ c) $\to$ ($a \to CPS$ b) $\to$ ($a \to CPS$ c)
$q \circ p$ = $\lambda a \to p$ $a \ggg q$

The rewrite nicely separates the expansion step from the "recursive call."   □

5 Bottom-up parsing: LR(0) parsing

Let us move on to our final challenge: quotations whose concrete syntax is based on
an LR(0) grammar–LR(1) grammars are also doable but we resist the temptation
to spell out the details. Unlike LL parsing, the LR method is generally not suitable
for constructing parsers by hand. For that reason, we shall base the treatment on a
very simple example, the language of balanced parentheses.

$$P \quad \to \quad \epsilon \mid P \ (\ P \ )$$

The abstract syntax is given by the *Tree* datatype: $P \to \epsilon$ constructs a *Leaf*,
$P \to P$ (P) a *Fork*.

An LR parser is similar to a postfix parser in that the state is a stack of symbols
the parser has already seen. In contrast, the state of an LL parser is a stack of
symbols it expects to see.

For efficiency reasons, an LR parser maintains additional information that
summarises the stack configuration in each step. This is accomplished by a finite-state
machine, the so-called LR(0) *automaton*. Here is the automaton for the grammar
above (S is a new start symbol, "»" serves as an end marker).

The automaton has six states; the production(s) contained in the states illustrate
the progress of the parse with the dots marking the borderline between what we
have seen and what we expect to see. If the dot appears before a terminal symbol,

we have a *shift state* (colored in yellow/light gray). If the dot has reached the end in one of the productions, we have a *reduce state* (colored in blue/dark gray). In a shift state, the parser consumes an input token and pushes in onto the stack. In a reduce state, the right-hand side of a production resides on top of the stack, which is then replaced by the left-hand side.

In our example, the parser first reduces $P \to \epsilon$ moving from the start state 1 to state 2. Then, it shifts either "»" or "(." Each transition is recorded on the stack. This information is used during a reduction to determine the next state. Consider state 6; two sequences of moves end in this state: $1 \xrightarrow{P} 2 \xrightarrow{(} 5 \xrightarrow{P} 4 \xrightarrow{)} 6$ and $5 \xrightarrow{P} 4 \xrightarrow{(} 5 \xrightarrow{P} 4 \xrightarrow{)} 6$. Removing $P(P)$ from the stack means returning to either state 1 or state 5. Pushing P onto the stack, we move forward to either state 2 or state 4. In short, reducing $P \to P(P)$ is accomplished by replacing the above transitions by either $1 \xrightarrow{P} 2$ or $5 \xrightarrow{P} 4$. The point is that, in general, there are several transitions for a single production.

Turning to the implementation, the parser's state is a stack of LR(0) states. Each LR(0) state carries the semantic value of the unique symbol that annotates the ingoing edges.

```
data S₁   = S₁           -- S
data S₂ st = S₂ Tree st   -- P
data S₃ st = S₃      st   -- »
data S₄ st = S₄ Tree st   -- P
data S₅ st = S₅      st   -- (
data S₆ st = S₆      st   -- )
```

Each state of the automaton is translated into a function that performs the corresponding action. A shift state simply delegates the control to the next active terminal. A reduce state pops the transitions corresponding to the right-hand side from the stack and pushes a transition corresponding to the left-hand side. If there are several possible transitions, then a reduce action is given by a family of functions represented as a type class.

```
quote                = state₁ S₁                         -- start
state₁ st            = state₂ (S₂ Leaf st)               -- reduce
state₂ st            = return st                         -- shift
state₃ (S₃ (S₂ t S₁)) = t                               -- accept
state₄ st            = return st                         -- shift
state₅ st            = state₄ (S₄ Leaf st)               -- reduce
class State₆ old new | old → new where
   state₆            :: old ↠ new
instance State₆ (S₆ (S₄ (S₅ (S₂ S₁)))) (S₂ S₁) where     -- reduce
   state₆ (S₆ (S₄ u (S₅ (S₂ t S₁))))     = state₂ (S₂ (Fork t u) S₁)
instance State₆ (S₆ (S₄ (S₅ (S₄ (S₅ st))))) (S₄ (S₅ st)) where  -- reduce
   state₆ (S₆ (S₄ u (S₅ (S₄ t (S₅ st))))) = state₄ (S₄ (Fork t u) (S₅ st))
```

The pattern $S_6 (S_4\ u\ (S_5 (S_4\ t\ (S_5\ st))))$ nicely shows the interleaving of states and semantic values. Since the stack is nested to the right, u is the topmost semantic value and consequently becomes the right subtree in *Fork t u*.

The active terminals implement the shift actions.

```
class Open old new | old → new where
  (                       :: old ⇸ new
instance Open (S₂ st) (S₄ (S₅ (S₂ st))) where
  ( st @(S₂ _ _)        = state₅ (S₅ st)
instance Open (S₄ st) (S₄ (S₅ (S₄ st))) where
  ( st @(S₄ _ _)        = state₅ (S₅ st)
  ) st @(S₄ _ _)        = state₆ (S₆ st)
  endquote st @(S₂ _ _) = state₃ (S₃ st)
```

We need a class if a terminal annotates more than one edge. Again, the instance types are not entirely straightforward as they reflect the stack modifications up to the next shift: for instance, "(" moves from S_2 to S_5 and then to S_4, which is again a shift state.

The implementation technique also works for LR(1) grammars. In this case, the active terminals implement both shift and reduce actions.

6 Conclusion

Quotations provide a new, amusing perspective on parsing: terminal symbols turn active and become the driving force of the parsing process. It is quite remarkable that all major syntax analysis techniques can be adapted to this technique.

Typed quotations provide static guarantees: using type-level representations of symbols Haskell's type-checker is instrumented to scrutinise whether a quotation is syntactically correct. Of course, this means that syntax errors become type errors, which are possibly difficult to decipher. Adding proper error handling is left as the obligatory "instructive exercise to the reader."

References

Aho, A. V., Lam, M. S., Sethi, R. & Ullman, J. D. (2006) *Compilers: Principles, Techniques, and Tools.* 2nd ed. Addison-Wesley.

Okasaki, C. (2002) Techniques for embedding postfix languages in Haskell. In *Proceedings of the 2002 ACM SIGPLAN Haskell Workshop*, Chakravarty, M. (ed), ACM, pp. 105–113.

Okasaki, C. (2003) Theoretical Pearls: Flattening combinators: Surviving without parentheses. *J. Funct. Program.,* **13**(4), 815–822.

Peyton Jones, S. (2003) *Haskell 98 Language and Libraries.* Cambridge University Press.

Wadler, P. (1989) Theorems for free! In *the Fourth International Conference on Functional Programming Languages and Computer Architecture (FPCA'89)*. London, UK: Addison-Wesley, pp. 347–359.