


Notes on “The Derivative of a Regular Type is its Type of One-Hole Contexts” by Conor McBride

Felix Filozov

Monday 20th October, 2014

Suppose we have the following binary tree

Timestamp:
20141020161700



and we would like to make *local modifications*, efficiently. Meaning, we’re only interested in modifying nodes within a subtree. For this particular example, we’re going to focus on modifying ⑤, we’ll call the *target node*.

Using a language like C, which allows mutation, modifying ⑤ can be done very efficiently. We simply create a pointer to ⑤, and modify the node through the pointer in $O(1)$.

This seems to be a problem for languages that do not allow mutation, pure languages, or so it seems. Trying to modify ⑤ the naive way, would mean recursing down to ⑤, changing it, and then recursing back up. This would have the effect of copying every node on the path from ① to ⑤. Here’s an example of the naive approach in Haskell:

```
data Tree = Empty | Node Int Tree Tree deriving (Show)

-- Search for 5 in a Tree and replace it with a new value.
naiveModify :: Tree -> Int -> Tree
naiveModify (Node oldVal t1 t2) newVal =
  if oldVal == 5
  then (Node newVal t1 t2)
  else (Node oldVal (naiveModify t1 newVal)
                  (naiveModify t2 newVal))
naiveModify Empty newVal = Empty
```

Once *naiveModify* finds the target node, and replaces the value, it will create new nodes as it recurses back up to the root of the tree.


The goal is to come up with a way so that local modifications to the above binary tree can be done efficiently in a pure language.

An elegant solution to this problem is to use a data structure called the *zipper* [1] [3]. Using the zipper, we can traverse the data structure, and make modifications to it. To make local modifications to any “regular” data structure efficiently, we simply have to create a zipper for it. This means other data structures can have their own zippers. For example, lists can also have zippers. However, all zippers work on the same principles. For the purpose of this talk, I will describe how a zipper works for binary trees, so from now on when I mention “zipper”, I mean a zipper for binary trees, unless stated otherwise.

Before I tell you how a zipper actually looks like, let me first tell you what operations a zipper supports. A zipper supports the following operations: *left()*, *right()*, *up()*, and *modify(v)*. We could add more, but these are sufficient.

For example, if our location is ②, and we call *up()*, then the zipper moves to ①. If we call *left()*, then the zipper moves to ④. If we call *right()*, then the zipper moves to ⑤. If we call *modify(v)*, then the value of ② will be replaced with *v*. It should be clear that using these operations we can walk around the tree, and have the ability to modify any node.

Here’s how a zipper looks like:




What we see above is the initial state of the zipper. The first element is the empty list. We’ll see later that this list is used by *left()*, *right()*, and *up()*. The fourth element is the *focus node*, that is, the node where we’re currently located. The second element is the left branch of the focus node, and the third element is the right branch of the focus node.

At this point you might be thinking that I’ve sold you snake oil! Is this really what a zipper is? “But it’s just the tree broken apart and put into a tuple!”, you say. Yes, but the magic actually happens in the operations. So let’s look at how they work.


As the zipper moves down the tree using *left()* and *right()*, it will be recording all the branches that it did not move to. So if it moves left, it will record the right branch into the previously mentioned list. If it moves right, it will record the left branch. For example, if the starting position is ①, then calling *left()* would move the zipper to ②, and in addition it would also record the right branch, ③, into the list.

Starting from the same position, if we call *right()*, the zipper would move to ③, and record the left branch




So as the zipper moves down the tree, it doesn't just discard the branches it passes by, instead it saves them. The metaphor for this is to imagine the zipper unzipping the tree. As it moves down, it “unzips” the tree, that is, it splits the tree along a path.

Let's look concretely at how `left()` works. For convenience, the initial state of the zipper is repeated:



After a single call to `left()`, the zipper looks like this:




Let's focus on the entry that was added to the list. The *L* is a marker which means that the zipper moved along the left branch. The second element is the branch the zipper did not move to as it moved left. It did not move to the right branch of $\textcircled{1}$, that is $\begin{array}{c} \textcircled{3} \\ \textcircled{6} \textcircled{7} \end{array}$, therefore the right branch was saved. The

third element is the element from which the zipper moved away, which was $\textcircled{1}$.

Every time the zipper moves left or right, this list will be augmented with additional entries.

Also, notice that the rest of the tuple also changed. The focus node has become $\textcircled{2}$, and the left and right branch of the focus node has been recorded in the tuple as well.


If we then call `right()`, the zipper changes to:



A new entry was added to the list. The *R* is a marker to indicate that the zipper moved right. Since the zipper moved right, the left branch of $\textcircled{2}$, that is $\begin{array}{c} \textcircled{4} \\ \textcircled{8} \textcircled{9} \end{array}$, was recorded. Since it moved away from $\textcircled{2}$, it recorded that node as well.

The rest of the tuple also changed. The new focus node is $\textcircled{5}$, and since it has no children, the zipper recorded *NIL*.

We've reached the node we wanted to modify. Let's modify this node by calling `modify(10)`. The zipper looks like this:



Let's modify it again by calling `modify(20)`. The zipper looks like this:

$$\left(\left[\left(R, \begin{array}{c} \textcircled{4} \\ \textcircled{8} \text{ --- } \textcircled{9} \end{array}, \textcircled{2} \right), \left(L, \begin{array}{c} \textcircled{3} \\ \textcircled{6} \text{ --- } \textcircled{7} \end{array}, \textcircled{1} \right) \right], \text{NIL}, \text{NIL}, \textcircled{20} \right)$$

We can continue modifying this particular node, and this modification will be done in $O(1)$, because we're only changing one entry in the tuple, the focus node. We're no longer paying the heavy price by having to copy all the nodes along the path to our target node. Hurray! We've achieved our goal.


After we've done all the necessary modifications to $\textcircled{5}$ (it became $\textcircled{20}$), suppose we decide to go back up to the root of the tree. We simply call `up()` twice. Because the zipper saved all the branches it did not move to, as it moved to the target node, the zipper will be able to reconstruct the tree as it moves up.

If we call `up()`, the zipper looks like this:

$$\left(\left[\left(L, \begin{array}{c} \textcircled{3} \\ \textcircled{6} \text{ --- } \textcircled{7} \end{array}, \textcircled{1} \right) \right], \begin{array}{c} \textcircled{4} \\ \textcircled{8} \text{ --- } \textcircled{9} \end{array}, \textcircled{20}, \textcircled{2} \right)$$

The focus node was $\textcircled{20}$, and by calling `up()` the zipper moved to the parent node $\textcircled{2}$. By moving to $\textcircled{2}$, the zipper reconstructed the entire subtree rooted at $\textcircled{2}$. It was able to do that by using the head entry in the list, $\left(R, \begin{array}{c} \textcircled{4} \\ \textcircled{8} \text{ --- } \textcircled{9} \end{array}, \textcircled{2} \right)$.

The entry tells the zipper that the parent node of $\textcircled{20}$ is $\textcircled{2}$. The left branch of $\textcircled{2}$ is $\begin{array}{c} \textcircled{4} \\ \textcircled{8} \text{ --- } \textcircled{9} \end{array}$. The *R* marker indicates that the zipper moved to the right from $\textcircled{2}$ to get to $\textcircled{20}$, which means $\textcircled{20}$ is the right branch of $\textcircled{2}$. Putting all this together the zipper reconstructed the subtree rooted at $\textcircled{2}$ as:



The zipper recorded $\textcircled{2}$ as the new focus node and its children into the rest of the tuple.


Let's call `up()` again. This time the result is:

$$\left(\left[\right], \begin{array}{c} \textcircled{2} \\ \textcircled{4} \text{ --- } \textcircled{20} \end{array}, \begin{array}{c} \textcircled{3} \\ \textcircled{6} \text{ --- } \textcircled{7} \end{array}, \textcircled{1} \right)$$

The zipper moved to the parent node $\textcircled{1}$ of the focus node $\textcircled{2}$. By doing so, it reconstructed the entire subtree rooted at $\textcircled{1}$, which happens to be the tree itself. Once again, it used the head entry in the list, $\left(L, \begin{array}{c} \textcircled{3} \\ \textcircled{6} \text{ --- } \textcircled{7} \end{array}, \textcircled{1} \right)$.

The marker L indicates that the zipper moved left to get to $\textcircled{2}$. The entry also tells us the zipper moved away from $\textcircled{1}$. This means the subtree rooted at $\textcircled{2}$ is the left branch of $\textcircled{1}$. In addition, the entry also has the right branch of $\textcircled{1}$, that is $\textcircled{6} \textcircled{3} \textcircled{7}$. Putting all this together, the reconstructed subtree

rooted at $\textcircled{1}$ is



The zipper then put the new focus


node and its children into the rest of the tuple.

Going up the tree is called zipping up the tree. It's a metaphor. The list of entries in the zipper contains separated branches. As the zipper moves up the tree, it puts those branches back together until we get the tree back.

Let's introduce some terminology.

$$\left(\underbrace{\left[\left(L, \textcircled{6} \textcircled{3} \textcircled{7}, \textcircled{1} \right) \right]}_{\text{one-hole context}}, \textcircled{8} \textcircled{4} \textcircled{9}, \textcircled{5}, \underbrace{\textcircled{2}}_{\text{focus node}} \right)$$

Looking at a zipper above, we've already given the fourth element in the tuple a special name, the focus node. Everything else in the tuple is called the *one-hole context*. The one-hole context represents everything around the focus node. That is, it's the tree itself without the focus node, the hole. Here's a way of showing the one-hole context:



This is how the zipper is able to reconstruct the tree as it moves up. It uses the one-hole context, and the focus node to rebuild the tree.

Since this paper is about types, let's write down the types of the data structures we've seen so far:

$$\begin{aligned} BTree &= 1 + Int * BTree * BTree \\ BTreeZipper &= [(2 * BTree * Int)] * BTree * BTree * Int \end{aligned}$$

Conor's paper [2] tells us that the type of $BTree$ and the type of the one-hole context are related by the rules of differentiation. In particular, we can derive

the type of the one-hole context, $[(2 * BTree * BTree)] * BTree * BTree$, from the type of $BTree$. This means that we can derive the type of the zipper of $BTree$ as well. In fact, we can derive the type of the zipper of any “regular” type.

I should mention that he represents recursive types differently than what I’ve shown. He uses the μ operator instead. He would write the type of the tree like this:

$$BTree = \mu y.(1 + Int * y * y)$$

From the little I know about recursive type theory, if we treat μ like a lambda abstraction and repeatedly replace y with $\mu y.(1 + Int * y * y)$, then we’ll get an infinite tree type. For example, if we replace once, we get:


$$BTree = (1 + Int * \mu y.(1 + Int * y * y) * \mu y.(1 + Int * y * y))$$

If we continue replacing, we’ll get:

$$BTree = (1 + Int * (1 + Int * (1 + Int * \dots * \dots) * (1 + Int * \dots * \dots)) * (1 + Int * (1 + Int * \dots * \dots) * (1 + Int * \dots * \dots)))$$

where “...” are further expansions.

Another way of representing the infinite tree type:



This infinite tree type is equivalent to the fixed-point of $\mu y.(1 + Int * y * y)$. That’s as far as I know about recursive types, if you know more, please find me and tell me more about them. From this point on, I’m simply going to use it as notation.

Conor defines the differentiation operator as follows:

$$\begin{aligned}
\partial_x &\mapsto 1 \\
\partial_x(y \setminus x) &\mapsto 0 \\
\partial_x 0 &\mapsto 0 \\
\partial_x(S + T) &\mapsto \partial_x S + \partial_x T \\
\partial_x 1 &\mapsto 0 \\
\partial_x(S \times T) &\mapsto \partial_x S \times T + S \times \partial_x T \\
\partial_x(\mu y.F) &\mapsto \mu z. \underline{\partial_x F|y = \mu y.F}_z + \underline{\partial_y F|y = \mu y.F}_z \times z \\
\partial_x(F|y = S) &\mapsto \partial_x F|y = S + \partial_y F|y = S \times \partial_x S \\
\partial_x \underline{T}_x &\mapsto 0 \\
\partial_x \underline{T}_{y \setminus x} &\mapsto \underline{\partial_x T}_y
\end{aligned}$$

We can use our intuition to make sense of some of these. For example, if our data structure is $S + T$, then the one-hole context of such a data structure will clearly be the one-hole context of S , or T . It wouldn't make sense for the one-hole context of the data structure to just be the one-hole context of S , because what if we have an instance of type T ?

With regards to the product rule, if we put a hole in S , then we'll have to record T . If we put a hole in T , then we'll have to record S . This is similar to what we did with our tree. The zipper moved left, that is it put a hole in the left branch, therefore it saved the right one.

Let's apply the rules to $BTree$:

$$\begin{aligned}
\partial_{Int}(BTree) &= \partial_{Int} \mu y. (1 + Int * y * y) \\
&= \mu z. \partial_{Int} (1 + Int * y * y) | y = BTree + \partial_y (1 + Int * y * y) | y = BTree \times z \\
&= \mu z. (\partial_{Int} 1) + \partial_{Int} (Int * y * y) | y = BTree + (\partial_y 1) + \partial_y (Int * y * y) | y = BTree \times z \\
&= \mu z. \partial_{Int} (Int * y * y) | y = BTree + \partial_y (Int * y * y) | y = BTree \times z \\
&= \mu z. (y * y) | y = BTree + (\partial_y (Int) \times y * y + \partial_y (y * y) * Int) | y = BTree \times z \\
&= \mu z. (BTree * BTree) + (\partial_y (y * y) * Int) | y = BTree \times z \\
&= \mu z. (BTree * BTree) + ((\partial_y (y) * y + y * \partial_y (y)) * Int) | y = BTree \times z \\
&= \mu z. (BTree * BTree) + (y * Int + y * Int) | y = BTree \times z \\
&= \mu z. (BTree * BTree) + (2 * y * Int) | y = BTree \times z \\
&= \mu z. (BTree * BTree) + (2 * BTree * Int) \times z \\
&= [(2 * BTree * Int)] * (BTree * BTree) \\
&= [(2 * BTree * Int)] * BTree * BTree
\end{aligned}$$

Observe that $\partial_{Int}(BTree)$ equals to the type of the one-hole context. From here, deriving the type of the zipper is simple.

In the paper we can also find the derivation of the one-hole context for the list data structure.

References

- [1] Huet Gérard, *The Zipper*, J. Funct. Program., September 1997, Vol. 7 Num. 5, 1997
- [2] McBride Conor, *The Derivative of a Regular Type is its Type of One-Hole Contexts (Extended Abstract)*, 2001.
- [3] Miran Lipovaca, *Learn You a Haskell for Great Good!: A Beginner's Guide*, <http://learnyouahaskell.com/zippers>, 2011.