

Pure Reasoning

& Raw Performance



Functional Binary Search Trees

- Great for **reasoning**:
 - Typical exercise in verification courses
 - Can be automatically synthesized from a specification
 - Amortized time complexity can be proven automatically
- But **performance** is lacking..
- Let's consider *splay trees* as an example

```
fun splay(pivot, t)
  match t
  Node(a, x, b) -> if x > pivot then match a // search for pivot in left subtree
  Node(a1, y, a2) -> if y < pivot then // search for pivot in right subtree
  val (small, big) = splay(pivot, a2) // recursively partition the subtree
  (Node(a1, y, small), Node(big, x, b)) // add subtrees to the appropriate pile
  ...
```

- New keys are inserted at the root with two children: the existing keys smaller and bigger than the new key, as computed by *splay* (above).
- But this allocates $\mathcal{O}(\log n)$ new nodes and uses $\mathcal{O}(\log n)$ stack. **Can we do better?**

In-Place Reuse

- One can safely update an immutable data structure in-place given unique ownership
- We use reference-counting to detect at runtime which blocks are unique
- Our compiler then detects reuse opportunities and removes allocations:


For unique data this function is fully in-place

```
fip fun reverse( xs, acc )
  match xs
  Cons(x,xx) -> reuse opportunity
  reverse( xx, Cons(x, acc))
  Nil -> Nil
```

```
fip fun reverse( xs, acc )
  match xs
  Cons(x,xx) -> val loc = if is-unique(xs) // refcount==1?
  then &xs // then reuse the memory, else allocate:
  else { dup(x); dup(xx); decref(xs); alloc(2) }
  reverse( xx, Cons@loc(x, acc))
  Nil -> Nil
```

Bottom-Up Algorithms

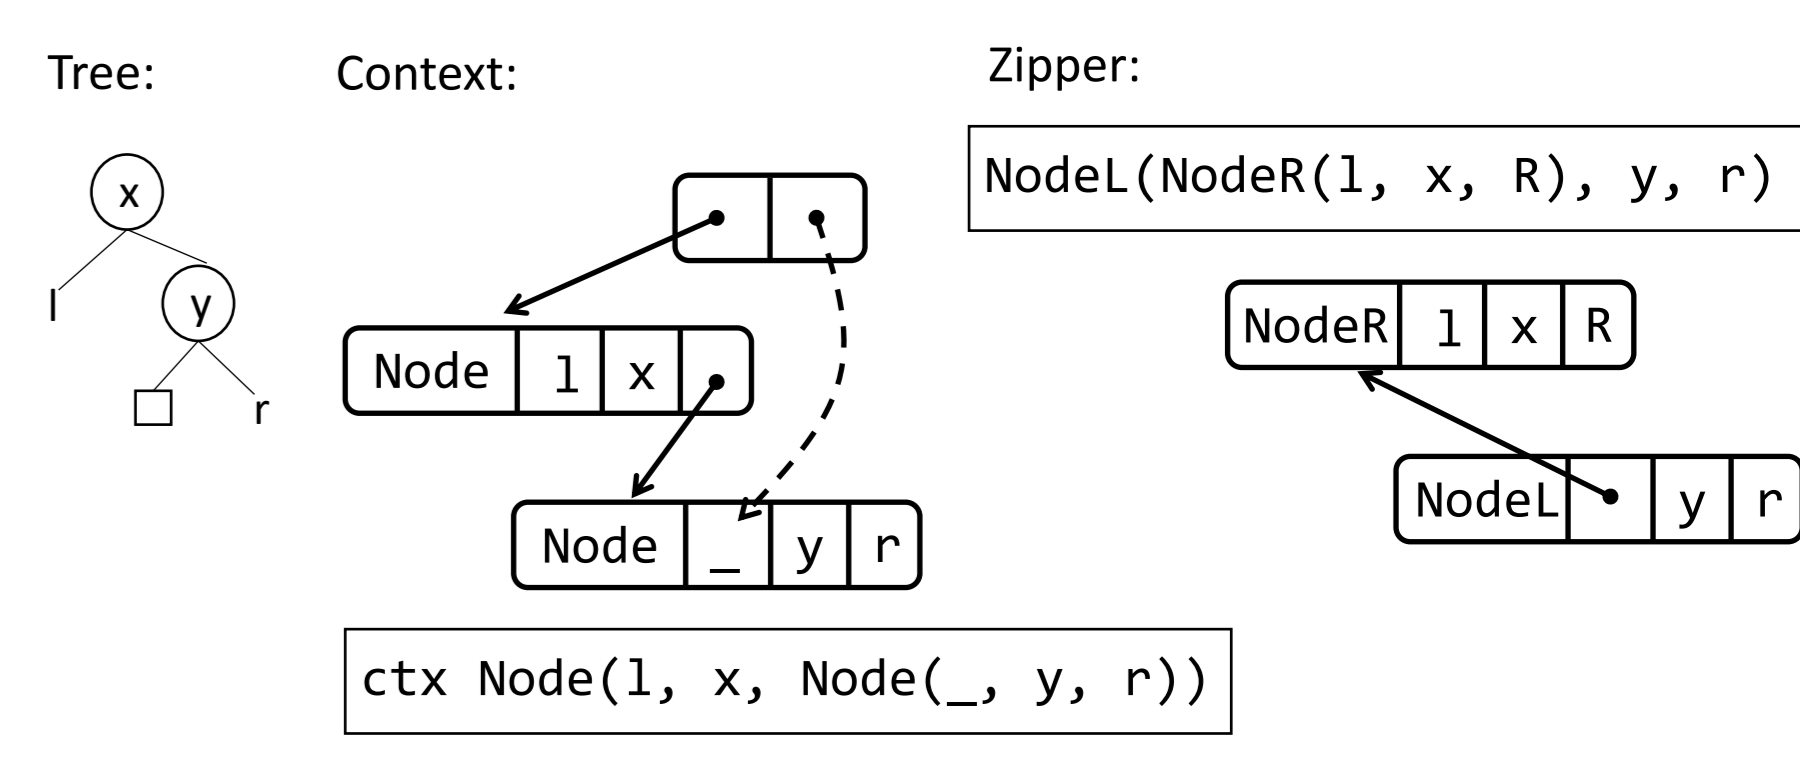
- A bottom-up insertion algorithm traverses down the tree searching for the new key and then *splays on the way back up*
- We represent the path upwards through the tree using a *zipper*
- Our functional algorithm uses $\mathcal{O}(1)$ new nodes and $\mathcal{O}(1)$ stack: It corresponds exactly to the description by Sleator & Tarjan (picture above)



```
fip fun splay( zipper, b, x, c )
  match zipper
  NodeR(a,y,NodeL(up,z,d)) -> reuse
  up.splay( Node(a,y,b), x, Node(c,z,d) )
  ...
```


Constructor Contexts & Zippers

- Both describe data structures with constant time access to a single hole
- Constructor contexts store the path from the root to the hole and an extra pointer to the hole
- But zippers invert pointers: they store the path from the hole to the root




```
(++) : cctx<a,b> -> b -> a
(++): cctx<a,b> -> cctx<b,c> -> cctx<a,c>
```

- Both are semantically immutable: If contexts are shared, our runtime copies the path from root to hole and creates a new hole-pointer



Top-Down Algorithms

- A top-down algorithm traverses down the tree searching for the new key and *splays while going down*
- We represent the already-traversed tree using two *constructor contexts* L and R
- Our functional algorithm uses $\mathcal{O}(1)$ new nodes and $\mathcal{O}(1)$ stack: It corresponds exactly to the top-down description by Sleator & Tarjan (Fig. 11 of their paper, reproduced above)



```
fip fun splay( t, k, l_ctx, r_ctx )
  match t
  Node(ayzb,x,c) -> if x > k then match ayzb
  Node(a,y,zb) -> if y < k then
  splay(zb,k, l_ctx ++ ctx Node(a,y,_), r_ctx ++ ctx Node(_,x,c))
  ...
```

Original, imperative algorithms

```
Definition heap_ptr_insert_td : val :=
fun ( name, root ) =>
var: left_dummy := #0 in
var: right_dummy := #0 in
var: node := root in
var: left_hook := &left_dummy in
var: right_hook := &right_dummy in
while ( true ) {
  if ( node == #0 ) {
    if ( !node->value == name ) {
      *left_hook = node->left;
      *right_hook = node->right;
      root := node;
      break;
    }
  }
  else {
    if ( !node->value > name ) {
      *right_hook = node;
      *right_hook = &(node->right);
      node := node->left;
    }
    else {
      *left_hook = node;
      *left_hook = &(node->left);
      node := node->right;
    }
  }
}
*left_hook = #0;
*right_hook = #0;
root := &root #0 #0;
root->value = name;
break;
};
root->left = left_dummy;
root->right = right_dummy;
ret: root
;
```

- We formalize precisely the imperative programs from the original papers in Iris HeapLang
- Using only loop invariants extracted from the functional programs, we can show functional correctness
- This shows that the functional programs capture the essence of the original algorithms

Benchmarks

- Our new algorithms in Koka perform on-par with the original C implementations of *move-to-root*, *splay*, *zip* and *red-black* trees

