

Speed and Simplicity for Incremental Sequence Computation

How do we make Incremental Computation easier to use?

How will it compete with the speed of native code?

PROGRES

Incremental Code

memo table

The diagram shows a Scala function `max` with annotations. A vertical line separates the function body into two regions: a light blue region on the left labeled "Explicit memo" and a light grey region on the right labeled "Table handling".

- Explicit memo:** Contains the function signature `fn max(l: List) -> Num` and the initial `let` binding `let (a,b) = match split(l) { ... }`.
- Table handling:** Contains the `fn memo(l) -> Num` function, the `let ma = match ttable.get(a) { ... }` block, and the `let ma = max(l); ttable.put(a,ma); ma` block.

Red arrows indicate the flow of control and data:

- An arrow from the `let` binding in the "Explicit memo" region to the `memo` function in the "Table handling" region.
- An arrow from the `memo` function back to the `let` binding.
- An arrow from the `let ma = match ttable.get(a) { ... }` block to the `let ma = max(l); ttable.put(a,ma); ma` block.
- An arrow from the `let ma = max(l); ttable.put(a,ma); ma` block back to the `let ma = match ttable.get(a) { ... }` block.
- An arrow from the `let ma = max(l); ttable.put(a,ma); ma` block to the `let (a,b) = match split(l) { ... }` block.
- An arrow from the `let (a,b) = match split(l) { ... }` block back to the `let ma = match ttable.get(a) { ... }` block.

```
table = new Memo;
fn max(l: List) -> Num {
  fn memo(l) -> Num {
    let ma = match ttable.get(a){
      Some(a) => a,
      None => {
        let ma = max(l);
        ttable.put(a,ma);
        ma
      }
    }
  }
  let (a,b) = match split(l) {
    None => 1.pop(),
    Some((a,b)) => {
      bin_max(memo(a),memo(b))
    }
  }
}
```

library

The diagram illustrates the flow of data and dependencies between two code snippets. The left snippet is a Haskell-like function:

```
fn max(m1: MemoList) -> Num {
  let l = read(m1);
  let (a,b) = match split(l) {
    None => l.pop(),
    Some((a,b)) => {
      bin_max(
        memo(max(write(a))),
        memo(max(write(b)))
      )
    }
  }
}
```

The right snippet is a simplified version of the same function:

```
fn max(m1: MemoList) -> Num {
  let l = read(m1);
  let (a,b) = match split(l) {
    None => l.pop(),
    Some((a,b)) => {
      bin_max(
        max(a),
        max(b)
      )
    }
  }
}
```

Arrows indicate the relationship between the two snippets:

- A yellow arrow points from the `read(m1)` line in the left snippet to the `read(m1)` line in the right snippet, labeled "read/write data for".
- A red arrow points from the `split(l)` line in the left snippet to the `split(l)` line in the right snippet, labeled "dependency tracking".
- A red arrow points from the `bin_max` line in the left snippet to the `bin_max` line in the right snippet, labeled "Explicit subsequence".
- A yellow arrow points from the `memo(max(write(a)))` line in the left snippet to the `max(a)` line in the right snippet, labeled "automatic change propagation".
- A yellow arrow points from the `memo(max(write(b)))` line in the left snippet to the `max(b)` line in the right snippet, labeled "automatic change propagation".

collections

The diagram illustrates the transformation of code into action labels and basic code. The code is as follows:


```
fn max(l: Giraz) -> Num {
  namespace("calc_max", || {
    l.fold_lr(|a,b| {
      bin_max(a,b)
    })
  })
}
```

The annotations are:

- action labels**: Points to the `namespace("calc_max", || {` line.
- basic code**: Points to the `l.fold_lr(|a,b| {` line.
- automatic change propagation**: Points to the `bin_max(a,b)` line.

Incremental computation works best with **large data sets** or long-running computations. Incremental libraries are well-suited for long-running computations, but not for simple computations on large datasets. The collections strategy makes up for this deficiency.

Incremental Max



Dependency Graphs

This thunk computes max

Transitively

“dirty” graph

Mutate

4 2 5 1 4 9 3

Dependency graphs are important for general purpose incremental computation. Nodes in the graph store code. When changing an input node, all dependents are marked as needing re-computation. In this way, dependency graphs can abstract some of the incremental code management away from the user.

Prior incremental code, like the **memo-table and library versions** to the left, required users to think directly about the incremental strategies used. This becomes limiting when trying to divide up work into subproblems. The user must consider rearranging code that spans multiple recursive calls.

The Giraz is based on the RAZ data structure. It has an **edit mode and a compute mode**. When in edit mode, the user can change data as if they had a pointer between two linked lists. In compute mode, the user calls one of the methods with some non-code they wrote. All the work is done by the Giraz.

To test the Giraz we compare with native code. We run a similar computation over the Giraz and a single array. We insert a new item in a random location and re-run the computation. The Giraz completes the re-computation much faster than the array!

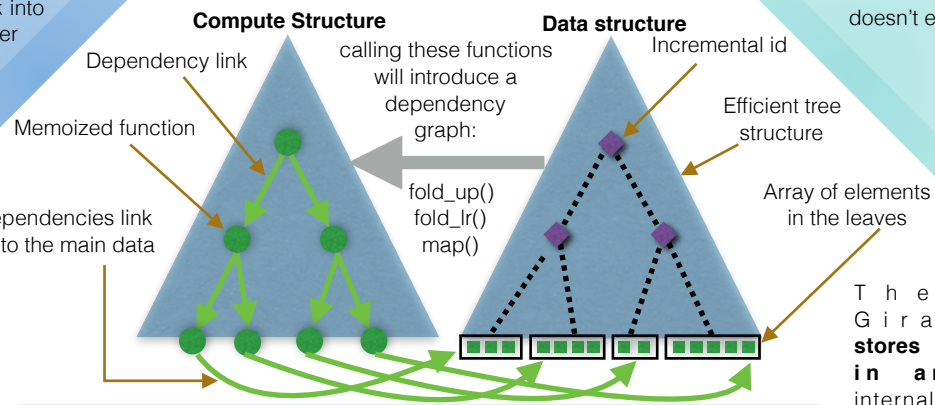
max	find the largest number in the sequence	58x Speedup
quicksort	find the subset of points that surround the others	4x Speedup
calc	parse numbers and symbols and perform the calculation	24x Speedup
to_string	change all the numbers to text	450x Speedup
reverse	change the order of the elements	22x Speedup

There were 10^6 elements and most tests updated in under 1ms.

Introducing:

Giraz

A collection for incremental sequences



The Giraz **stores data in arrays** internally, which may or may not be exposed to the user, depending on the incremental method call. The array boundaries are defined when the data is added to the collection. The user calls **“archive” functions** instead of inserting a data element. These calls insert markers used internally to structure the whole collection.

Systems language



Incremental computation requires a lot of memory management. References and garbage collectors in some languages can get in the way. The best choice is a systems language like rust that does not have a garbage collector.

"Rust is a systems programming language that runs blazingly fast, prevents segfaults, and guarantees thread safety."


Garbage collectors are usually great for languages that use references and linked data goes out of scope during execution.

With incremental computation, large dependency graphs are stored and their scope doesn't necessarily match the program structure.


The garbage collector ends up traversing these graphs, spending a lot of time searching for stale data that doesn't exist.

Data structure

Overhead Reduction




Many incremental computation libraries use one memoized thunk per user function. This requires a lot of overhead.



Grouping up functions and running them together when the thunk requires recompilation can save a lot of time looking up parameters.


Cache Coherency



Arrays require far fewer memory accesses than linked lists. Using arrays prevents pauses in computation. In incremental computation, reads have to check whether the data has changed or if it's in the memo table, requiring more time.

The Adaption Incremental Computation Engine uses “names” as memo table keys. They allow more **complex usage of memorization primitives**. Making good use of names in complex situations is a difficult research problem, but we can encode some working strategies into our collections abstractions. As more research comes out, we can improve the library without asking users to change their code.

Incremental Quickhull



by Kyle Headley
More info at kyleheadley.github.io



CUPLV
Programming Languages and Verification
University of Colorado Boulder