

Functional programming and parallelism

Conal Elliott

May 2016

What makes a language good for parallelism?

...

What makes a language *bad* for parallelism?

- Sequential bias
 - Primitive: assignment (state change)
 - Composition: *sequential* execution
 - “Von Neumann” languages (Fortran, C, Java, Python, ...)
- *Over-linearizes* algorithms.
- Hard to isolate accidental sequentiality.

Can we fix sequential languages?

- Throw in parallel composition.



- Oops:
 - Nondeterminism
 - Deadlock
 - Intractable reasoning



Can we *un-break* sequential languages?

*Perfection is achieved not when there is nothing left to add,
but when there is nothing left to take away.*

Antoine de Saint-Exupéry

Applications perform zillions of simple computations.

- Compute all at once?
- Oops — dependencies.
- Minimize dependencies!

Dependencies

- Three sources:
 - ① Problem
 - ② Algorithm
 - ③ Language
- Goals: eliminate #3, and reduce #2.

Dependency in sequential languages

- Built into sequencing: $A; B$
- Semantics: B begins where A ends.
- Why sequence?

Idea: remove *all* state

- And, with it,
 - mutation (assignment),
 - sequencing,
 - statements.
- Expression dependencies are specific & explicit.
- Remainder can be parallel.
- Contrast: “ $A; B$ ” vs “ $A + B$ ” vs “ $(A + B) \times C$ ”.

Programming without state

- Programming is calculation/math:
 - Precise & tractable reasoning (algebra),
 - ... including optimization/transformation.
- No loss of expressiveness!
- “Functional programming” (value-oriented)
- Like arithmetic on big values

Sequential sum

C:

```
int sum(int arr[], int n) {  
    int acc = 0;  
    for (int i=0; i<n; i++)  
        acc += arr[i];  
    return acc;  
}
```

Haskell:

```
sum = sumAcc 0  
where  
    sumAcc acc []      = acc  
    sumAcc acc (a : as) = sumAcc (acc + a) as
```

$sum = foldl (+) 0$

where


$$\begin{aligned} foldl\ op\ acc\ [] &= acc \\ foldl\ op\ acc\ (a : as) &= foldl\ op\ (acc\ 'op'\ a)\ as \end{aligned}$$

$$sum = foldr (+) 0$$

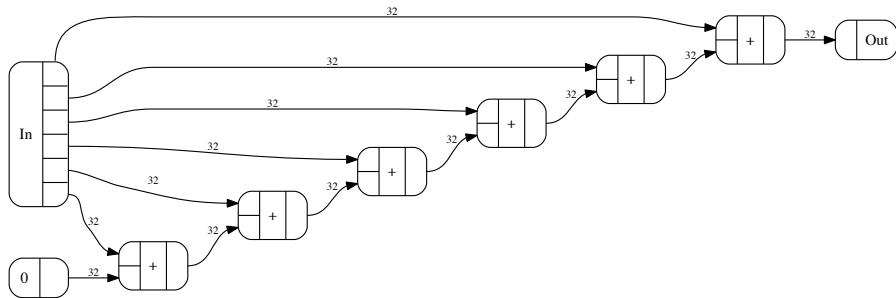
where

$$\begin{aligned} foldr\ op\ e\ [] &= e \\ foldr\ op\ e\ (a : as) &= a\ 'op'\ foldr\ op\ e\ as \end{aligned}$$

Sequential sum — left



Sequential sum — right



Parallel sum — how?

Left-associated sum:

$$\textit{sum} [a, b, \dots, z] \equiv (\dots((0 + a) + b)\dots) + z$$

How to parallelize?

Divide and conquer?

Balanced data

data $Tree\ a = L\ a \mid B\ (Tree\ a)\ (Tree\ a)$

Sequential:

$sum = sumAcc\ 0$

where

$sumAcc\ acc\ (L\ a) = acc + a$

$sumAcc\ acc\ (B\ s\ t) = sumAcc\ (sumAcc\ acc\ s)\ t$

Again, $sum = foldl\ (+)\ 0$.


Parallel:

$sum\ (L\ a) = a$

$sum\ (B\ s\ t) = sum\ s + sum\ t$

Equivalent? Why?

Balanced tree sum — depth 4



Balanced computation

- Generalize beyond $+$, 0 .
- When valid?

Associative folds

Monoid: type with associative operator & identity.

$$\textit{fold} :: \textit{Monoid } a \Rightarrow [a] \rightarrow a$$

Not just lists:

$$\textit{fold} :: (\textit{Foldable } f, \textit{Monoid } a) \Rightarrow f\ a \rightarrow a$$

Balanced data structures lead to balanced parallelism.

Two associative folds

$$\text{fold} :: \text{Monoid } a \Rightarrow [a] \rightarrow a$$

$$\text{fold } [] = \emptyset$$

$$\text{fold } (a : as) = a \oplus \text{fold } as$$

$$\text{fold} :: \text{Monoid } a \Rightarrow \text{Tree } a \rightarrow a$$

$$\text{fold } (L a) = a$$

$$\text{fold } (B s t) = \text{fold } s \oplus \text{fold } t$$

Derivable automatically from types.

Trickier algorithm: prefix sums


C:

```
int prefixSums(int arr[], int n) {  
    int sum = 0;  
    for (int i=0; i<n; i++) {  
        int next = arr[i];  
        arr[i] = sum;  
        sum += next;  
    }  
    return sum;  
}
```

Haskell:

$$\text{prefixSums} = \text{scanl } (+) 0$$

Sequence prefix sum



Sequential prefix sums on trees

$prefixSums = scanl (+) 0$

$scanl\ op\ acc\ (L\ a) = (L\ acc, acc\ 'op'\ a)$


$scanl\ op\ acc\ (B\ u\ v) = (B\ u'\ v', vTot)$

where

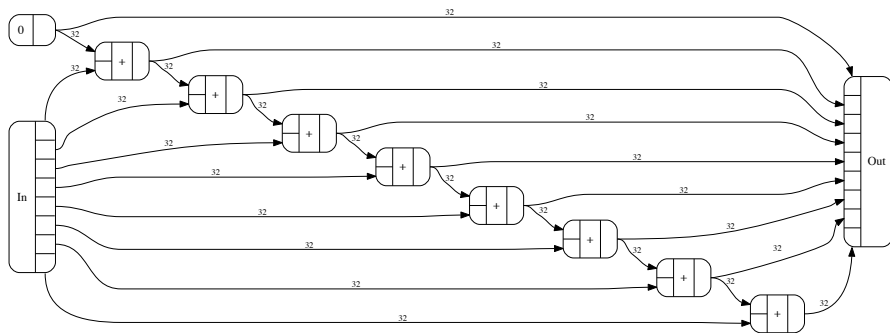
$(u', uTot) = scanl\ op\ acc\ u$

$(v', vTot) = scanl\ op\ uTot\ v$

Sequential prefix sums on trees — depth 2



Sequential prefix sums on trees — depth 3



Sequential prefix sums on trees

$$\text{prefixSums} = \text{scanl } (+) 0$$

$$\text{scanl } op \text{ acc } (L \ a) = (L \ acc, acc \ 'op' \ a)$$

$$\text{scanl } op \text{ acc } (B \ u \ v) = (B \ u' \ v', vTot)$$

where

$$(u', uTot) = \text{scanl } op \text{ acc } \ u$$

$$(v', vTot) = \text{scanl } op \ uTot \ v$$

- Still very sequential.
- Does associativity help as with *fold*?


Parallel prefix sums on trees

On trees:


$$\begin{aligned} \text{scan } (L \ a) &= (L \ \emptyset, a) \\ \text{scan } (B \ u \ v) &= (B \ u' \ (\text{fmap } \text{adjust } v'), \text{adjust } vTot) \\ \text{where} \\ (u', uTot) &= \text{scan } u \\ (v', vTot) &= \text{scan } v \\ \text{adjust } x &= uTot \oplus x \end{aligned}$$

- If balanced, dependency depth $O(\log n)$, work $O(n \log n)$.
- Can reduce work to $O(n)$. (*Understanding efficient parallel scan*).
- Generalizes from trees.
- *Automatic from type.*

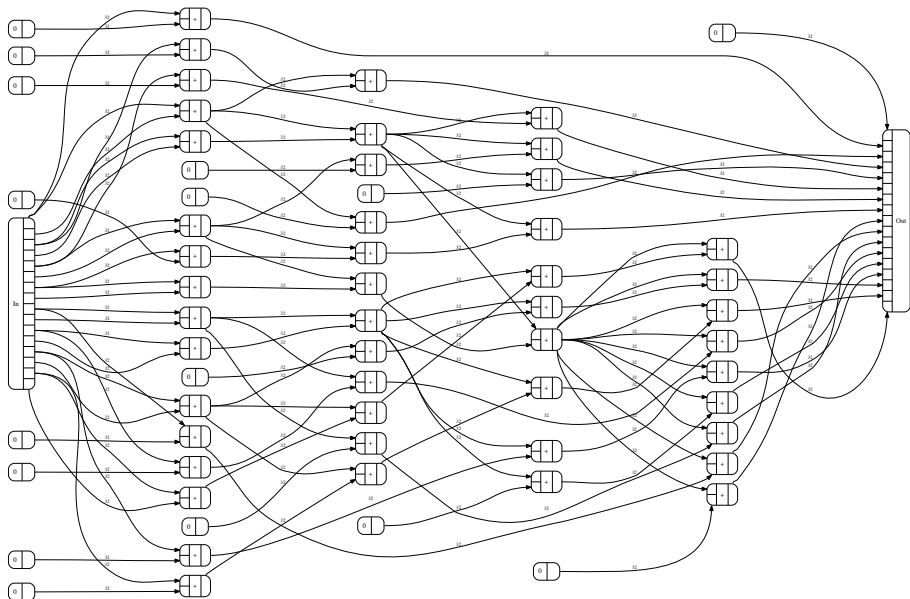
Balanced parallel prefix sums — depth 2




Balanced parallel prefix sums — depth 3




Balanced parallel prefix sums — depth 4




Balanced parallel prefix sums — depth 2, optimized



Balanced parallel prefix sums — depth 3, optimized



Balanced parallel prefix sums — depth 4, optimized



Why functional programming?

- Parallelism
- Correctness
- Productivity

R&D agenda: elegant, massively parallel FP

- Algorithm design:
 - Functional & richly typed
 - Parallel-friendly
 - Easily composable
- Compiling for highly parallel execution:
 - Convert to algebraic vocabulary (CCC).
 - Interpret vocabulary as “circuits” (FPGA, silicon, GPU).
 - Other interpretations.

Composable data structures

- Data structure tinker toys:

data *Empty* $a = \text{Empty}$

data *Id* $a = \text{Id } a$

data $(f + g) \ a = L \ (f \ a) \mid R \ (g \ a)$

data $(f \times g) \ a = \text{Prod } (f \ a) \ (g \ a)$

data $(g \circ f) \ a = O \ (g \ (f \ a))$

- Specify algorithm version for each.
- Automatic, type-directed composition.



$$\overbrace{Id \times \cdots \times Id}^{n \text{ times}}$$

Right-associated:

type family $RVec\ n$ **where**
 $RVec\ Z = Empty$
 $RVec\ (S\ n) = Id \times RVec\ n$

Left-associated:

type family $LVec\ n$ **where**
 $LVec\ Z = Empty$
 $LVec\ (S\ n) = LVec\ n \times Id$

Perfect binary leaf trees

$$\overbrace{Pair \circ \dots \circ Pair}^{n \text{ times}}$$

Right-associated:

type family $RBin\ n$ **where**

$$RBin\ Z = Id$$

$$RBin\ (S\ n) = Pair \circ RBin\ n$$

Left-associated:

type family $LBin\ n$ **where**

$$LBin\ Z = Id$$

$$LBin\ (S\ n) = LBin\ n \circ Pair$$

Uniform pairs:

$$\mathbf{type}\ Pair = Id \times Id$$

Generalized trees

$$\overbrace{h \circ \dots \circ h}^{n \text{ times}}$$

Right-associated:

type family *RPow* *h n* **where**
 RPow h Z = *Id*
 RPow h (S n) = *h* \circ *RPow h n*

Left-associated:

type family *LPow* *h n* **where**
 LPow h Z = *Id*
 LPow h (S n) = *LPow h n* \circ *h*

Binary:

type *RBin n* = *RPow Pair n*
type *LBin n* = *LPow Pair n*

Composing scans

class *LScan* *f* **where**

lscan :: *Monoid* *a* \Rightarrow *f* *a* \rightarrow (*f* \times *Id*) *a*

pattern *And1* *fa* *a* = *Prod* *fa* (*Id* *a*)

instance *LScan* *Empty* **where**

lscan *fa* = *And1* *fa* \emptyset

instance *LScan* *Id* **where**

lscan (*Id* *a*) = *And1* (*Id* \emptyset) *a*

instance (*LScan* *f*, *LScan* *g*) \Rightarrow *LScan* (*f* \times *g*) **where**

lscan (*Prod* *fa* *ga*) = *And1* (*Prod* *fa'* *ga'*) *gx*

where

And1 *fa'* *fx* = *lscan* *fa*


And1 *ga'* *gx* = *adjust* *fx* (*lscan* *ga*)

Composing scans


instance (*LScan g, LScan f, Zip g*) \Rightarrow *LScan (g \circ f)* **where**
 lscan (O gfa) = And1 (O (zipWith adjust tots' gfa')) tot
 where
 (*gfa', tots*) = *unzipAnd1 (fmap lscan gfa)*
 And1 tots' tot = lscan tots

adjust :: (*Monoid a, Functor t*) \Rightarrow *a* \rightarrow *t a* \rightarrow *t a*
adjust a t = fmap (a \oplus) t


Scan — *R*Pow Pair N5




Scan — *L*Pow Pair N5



Scan — $R\text{Pow} (L\text{Vec } N3) N3$



Scan — $RPow (LPow Pair N2) N3$



Polynomial evaluation


$$a_0 \cdot x^0 + \dots + a_n \cdot x^n$$

evalPoly coeffs x = coeffs · powers x


powers = lproducts ∘ pure

lproducts = underF Product lscan

Powers — $RBin\ N_4$



Polynomial evaluation — *RBin* N_4



Fast Fourier transform

DFT:

$$X_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N}nk}$$

FFT for $N = N_1 \cdot N_2$ (Gauss / Cooley-Tukey):

$$X_k = \sum_{n_1=0}^{N_1-1} \left[e^{-\frac{2\pi i}{N}n_1k_2} \right] \left(\sum_{n_2=0}^{N_2-1} x_{N_1n_2+n_1} e^{-\frac{2\pi i}{N_2}n_2k_2} \right) e^{-\frac{2\pi i}{N_1}n_1k_1}$$

Fast Fourier transform

```
class FFT f where  
  type FFO f :: * → *  
  fft :: RealFloat a ⇒ f (Complex a) → FFO f (Complex a)
```

```
instance FFT Id where  
  type FFO Id = Id  
  fft = id
```

```
instance FFT Pair where  
  type FFO Pair = Pair  
  fft (a :# b) = (a + b) :# (a - b)
```

FFT — composition (Gauss / Cooley-Tukey)

instance... \Rightarrow *FFT* ($g \circ f$) **where**

type *FFO* ($g \circ f$) = *FFO* $f \circ$ *FFO* g

fft = $O \circ \text{traverse } \text{fft} \circ \text{twiddle} \circ \text{traverse } \text{fft} \circ \text{transpose} \circ \text{unO}$

twiddle :: ... $\Rightarrow g (f (\text{Complex } a)) \rightarrow g (f (\text{Complex } a))$

twiddle = ($\text{zipWith} \circ \text{zipWith}$) (*) *twiddles*

where

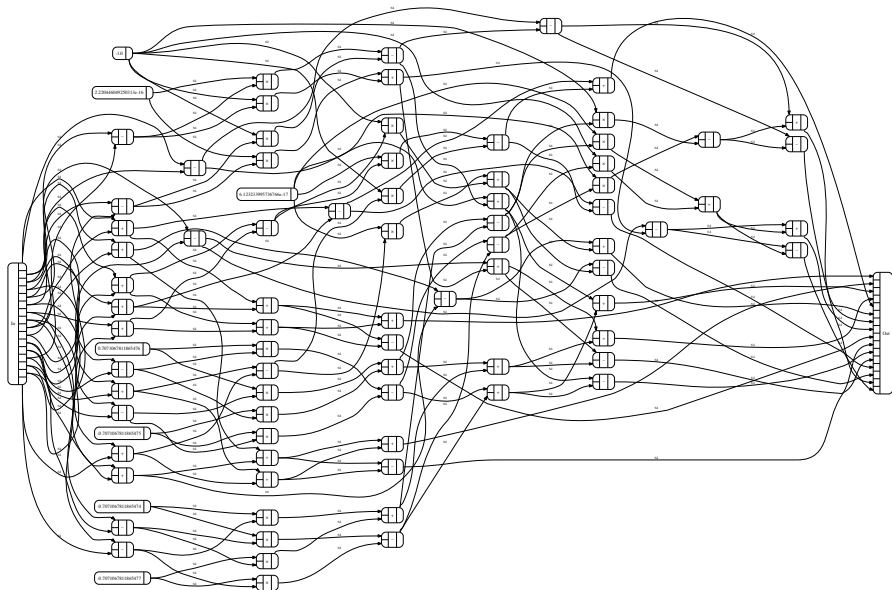
n = $\text{size}@(g \circ f)$

twiddles = $\text{fmap powers } (\text{powers } \omega)$

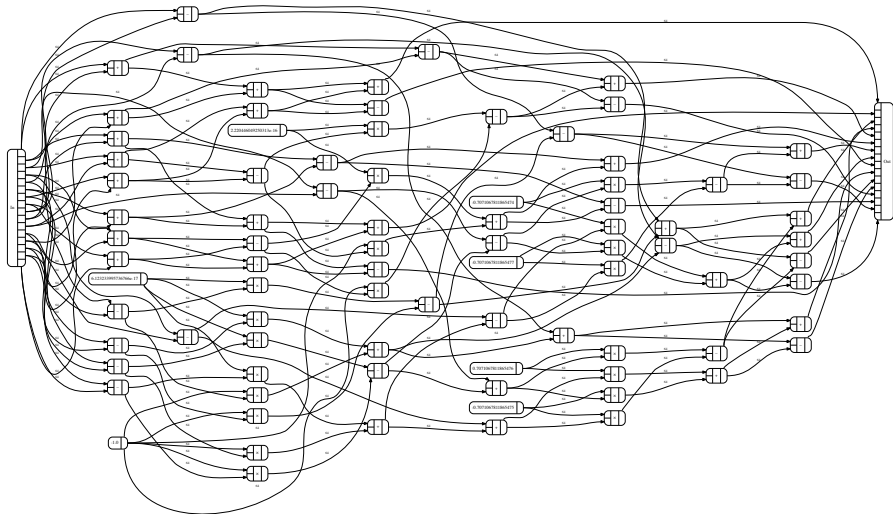
ω = $\text{cis } (-2 * \pi / \text{fromIntegral } n)$

$\text{cis } a$ = $\cos a :+ \sin a$

FFT — *RBin N3* (“Decimation in time”)




FFT — *LBin N3* (“Decimation in frequency”)




Bitonic sort


Bitonic sort — depth 1




Bitonic sort — depth 2



Bitonic sort — depth 3




Bitonic sort — depth 4



Manual vs automatic placement

ENIAC, 1946:



Manual vs automatic placement

- Programmers used to explicitly place computations in space.
- Mainstream programming *still* manually places in time.
- Sequential composition: crude placement tool.
- Threads: notationally clumsy & hard to manage correctly.
- If we relinquish control, automation can do better.