

Can Tensor Programming Be Liberated from the Fortran Data Paradigm?

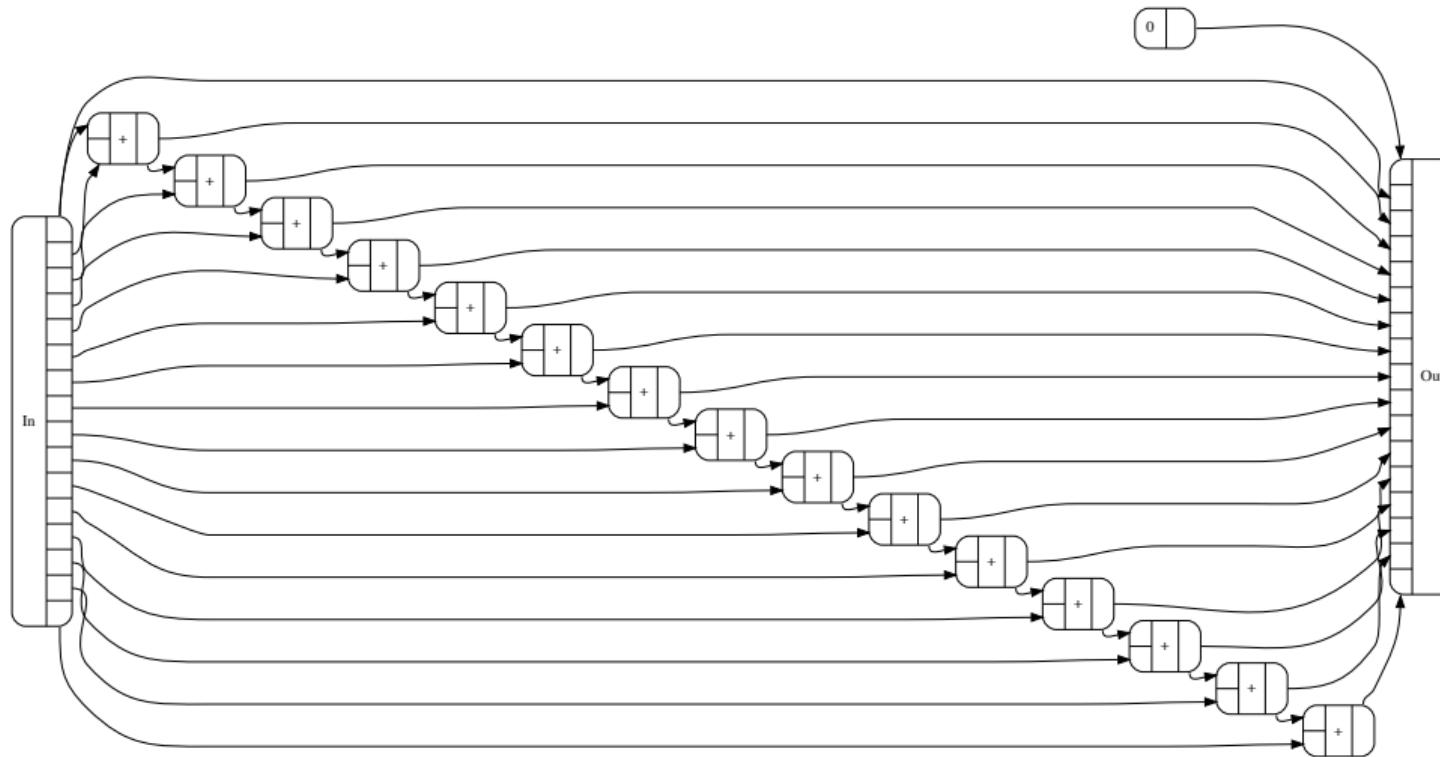
Conal Elliott

October 2021

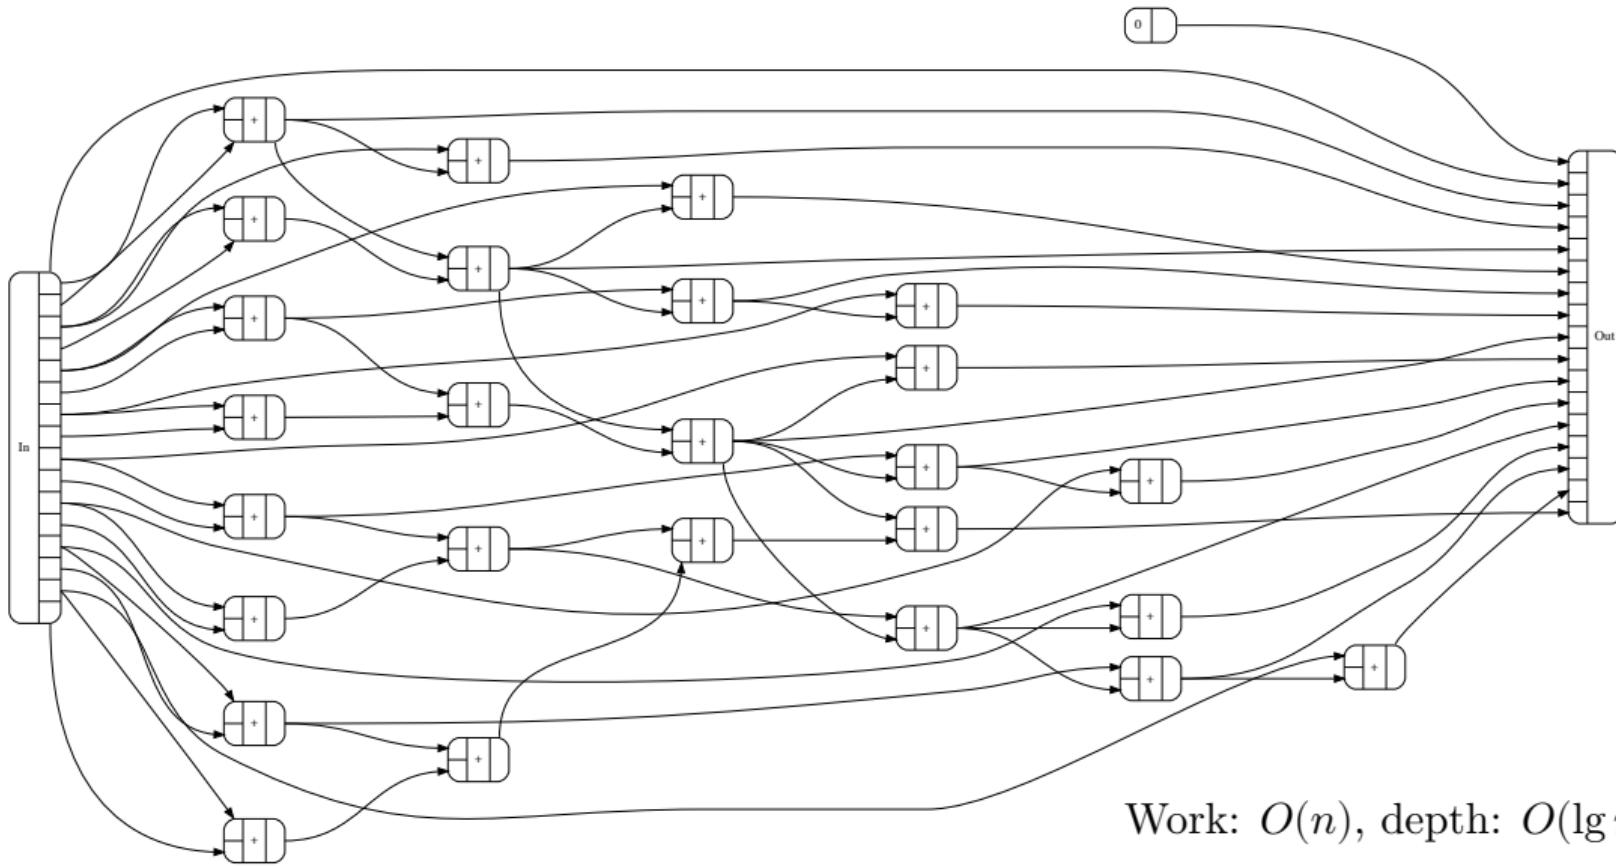
Prefix sum (left scan)

$$b_k = \sum_{i < k} a_i$$

Efficient sequential scan



Efficient parallel scan



An efficient array program (CUDA C)

```
--global__ void prescan(float *g_odata, float *g_idata, int n) {
    extern __shared__ float temp[]; // allocated on invocation
    int thid = threadIdx.x;
    int offset = 1;
    // load input into shared memory
    temp[2*thid] = g_idata[2*thid];
    temp[2*thid+1] = g_idata[2*thid+1];
    // build sum in place up the tree
    for (int d = n>>1; d > 0; d >>= 1) {
        __syncthreads();
        if (thid < d) {
            int ai = offset*(2*thid+1)-1;
            int bi = offset*(2*thid+2)-1;
            temp[bi] += temp[ai];
        }
        offset *= 2;
    }
    // clear the last element
    if (thid == 0) { temp[n - 1] = 0; }
    // traverse down tree & build scan
    for (int d = 1; d < n; d *= 2) {
        offset >>= 1;
        __syncthreads();
        if (thid < d) {
            int ai = offset*(2*thid+1)-1;
            int bi = offset*(2*thid+2)-1;
            float t = temp[ai];
            temp[ai] = temp[bi];
            temp[bi] += t;
        }
        __syncthreads();
    }
    // write results to device memory
    g_odata[2*thid] = temp[2*thid];
    g_odata[2*thid+1] = temp[2*thid+1];
}
```

Source: Harris, Sengupta, and Owens in *GPU Gems 3*, Chapter 39



WAT

```
function scan(a) =
  if #a == 1 then [0]
  else
    let es = even_elts(a);
        os = odd_elts(a);
        ss = scan({e+o: e in es; o in os})
    in interleave(ss,{s+e: s in ss; e in es})
```

Source: Guy Blelloch in *Programming parallel algorithms*, 1990

Still, why does it work, and how does it (correctly) generalize?

It's not naturally an array algorithm.

What is it?

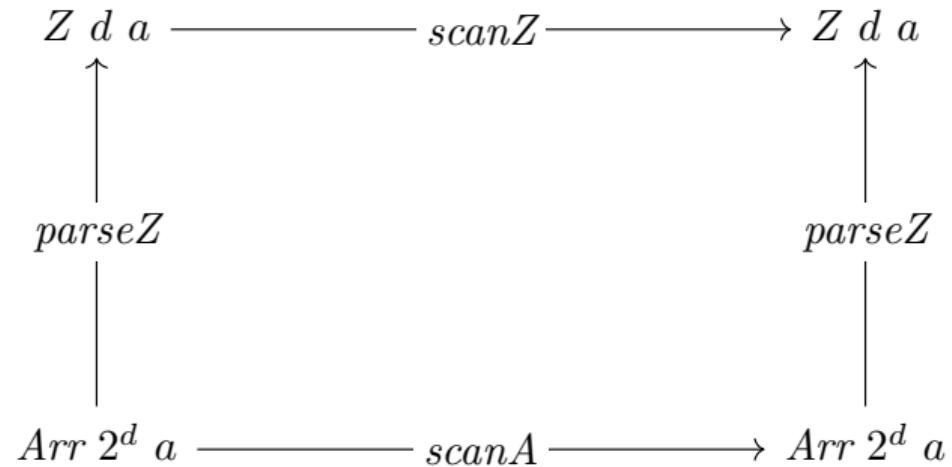
Clarifying the question

$$Arr\ 2^d\ a \xrightarrow{\quad} scanA \xrightarrow{\quad} Arr\ 2^d\ a$$

Clarifying the question

$$Z \ d \ a \xrightarrow{\text{scanZ}} Z \ d \ a$$
$$\text{Arr } 2^d \ a \xrightarrow{\text{scanA}} \text{Arr } 2^d \ a$$

Clarifying the question



Where scanZ is simple to state, prove, and generalize; and parseZ is formulaic in Z .

A compositional refinement

$$\begin{array}{ccc} Z \ d \ a & \xrightarrow{\quad scanZ \quad} & Z \ d \ a \times a \\ \uparrow \text{parseZ} & & \uparrow \text{parseZ} \otimes \text{id} \\ Arr \ 2^d \ a & \xrightarrow{\quad scanA \quad} & Arr \ 2^d \ a \times a \end{array}$$

Where $scanZ$ is simple to state, prove, and generalize; and $parseZ$ is formulaic in Z .

data $T_{\downarrow} :: * \rightarrow *$ **where**

$L :: a \rightarrow T_{\downarrow} a$

$B :: T_{\downarrow} a \rightarrow T_{\downarrow} a \rightarrow T_{\downarrow} a$

deriving instance Functor

$scanT_{\downarrow} :: Monoid a \Rightarrow T_{\downarrow} a \rightarrow T_{\downarrow} a \times a$

$scanT_{\downarrow}(L x) = (L \varepsilon, x)$

$scanT_{\downarrow}(B u v) = (B u' (fmap (u_{tot} \oplus) v'), u_{tot} \oplus v_{tot})$

where

$(u', u_{tot}) = scanT_{\downarrow} u$

$(v', v_{tot}) = scanT_{\downarrow} v$

Work: $O(n \lg n)$, depth: $O(\lg n)$.

Refined wrong guess: top-down, binary, *perfect* trees

```
data  $T_{\downarrow} :: Nat \rightarrow * \rightarrow *$  where  
   $L :: a \rightarrow T_{\downarrow} 0 a$   
   $B :: T_{\downarrow} d a \rightarrow T_{\downarrow} d a \rightarrow T_{\downarrow} (d + 1) a$ 
```

deriving instance $\text{Functor}(T_{\downarrow} d)$

$scanT_{\downarrow} :: \text{Monoid } a \Rightarrow T_{\downarrow} d a \rightarrow T_{\downarrow} d a \times a$

$scanT_{\downarrow}(L x) = (L \varepsilon, x)$

$scanT_{\downarrow}(B u v) = (B u' (\text{fmap } (u_{tot} \oplus) v'), u_{tot} \oplus v_{tot})$

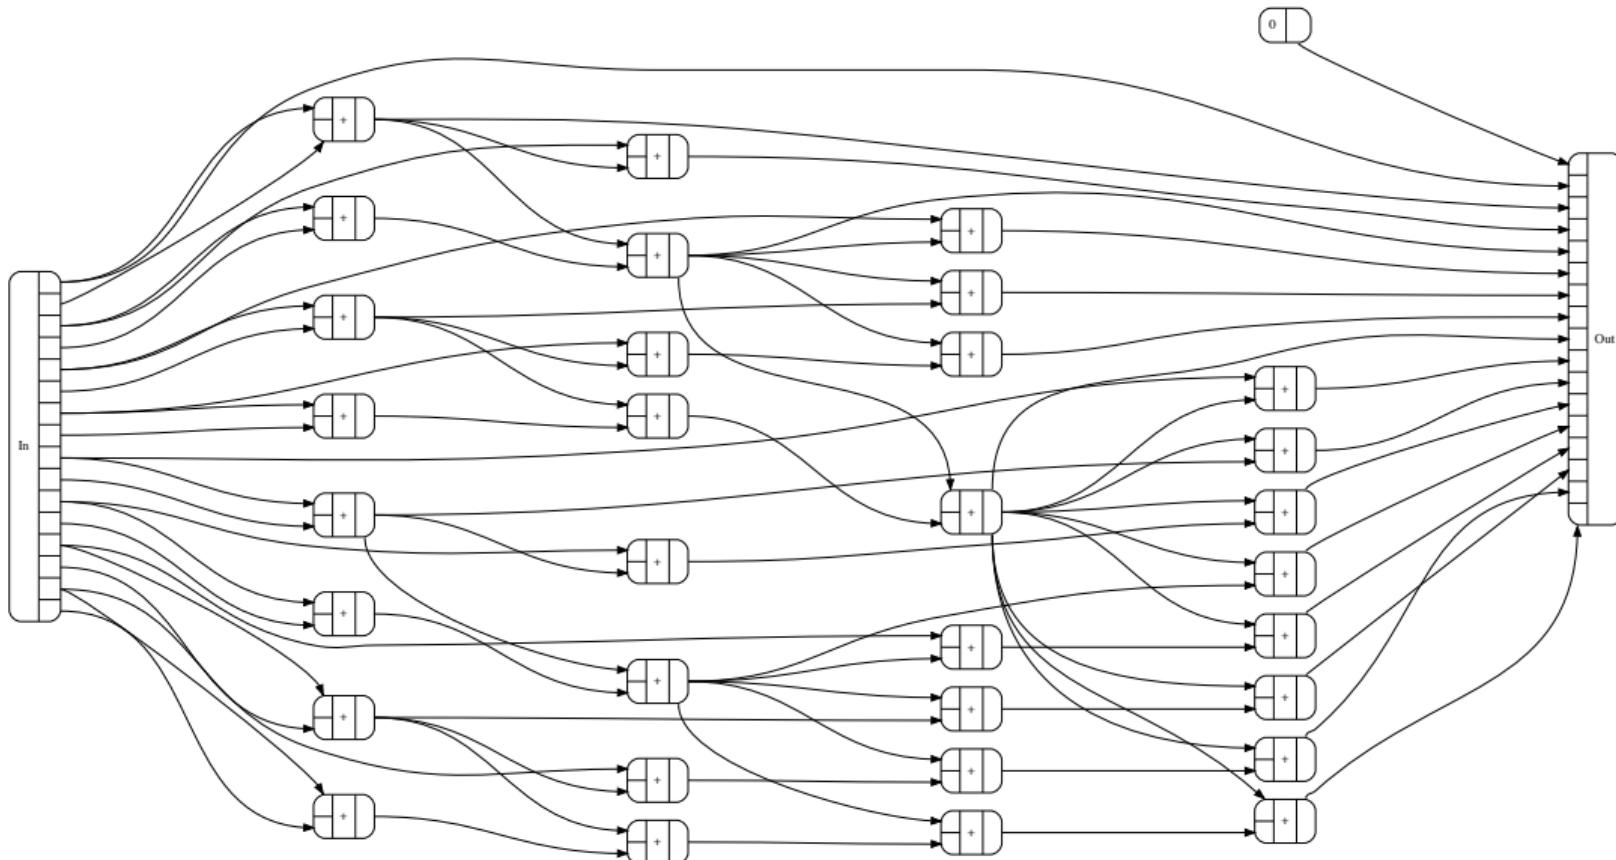
where

$(u', u_{tot}) = scanT_{\downarrow} u$

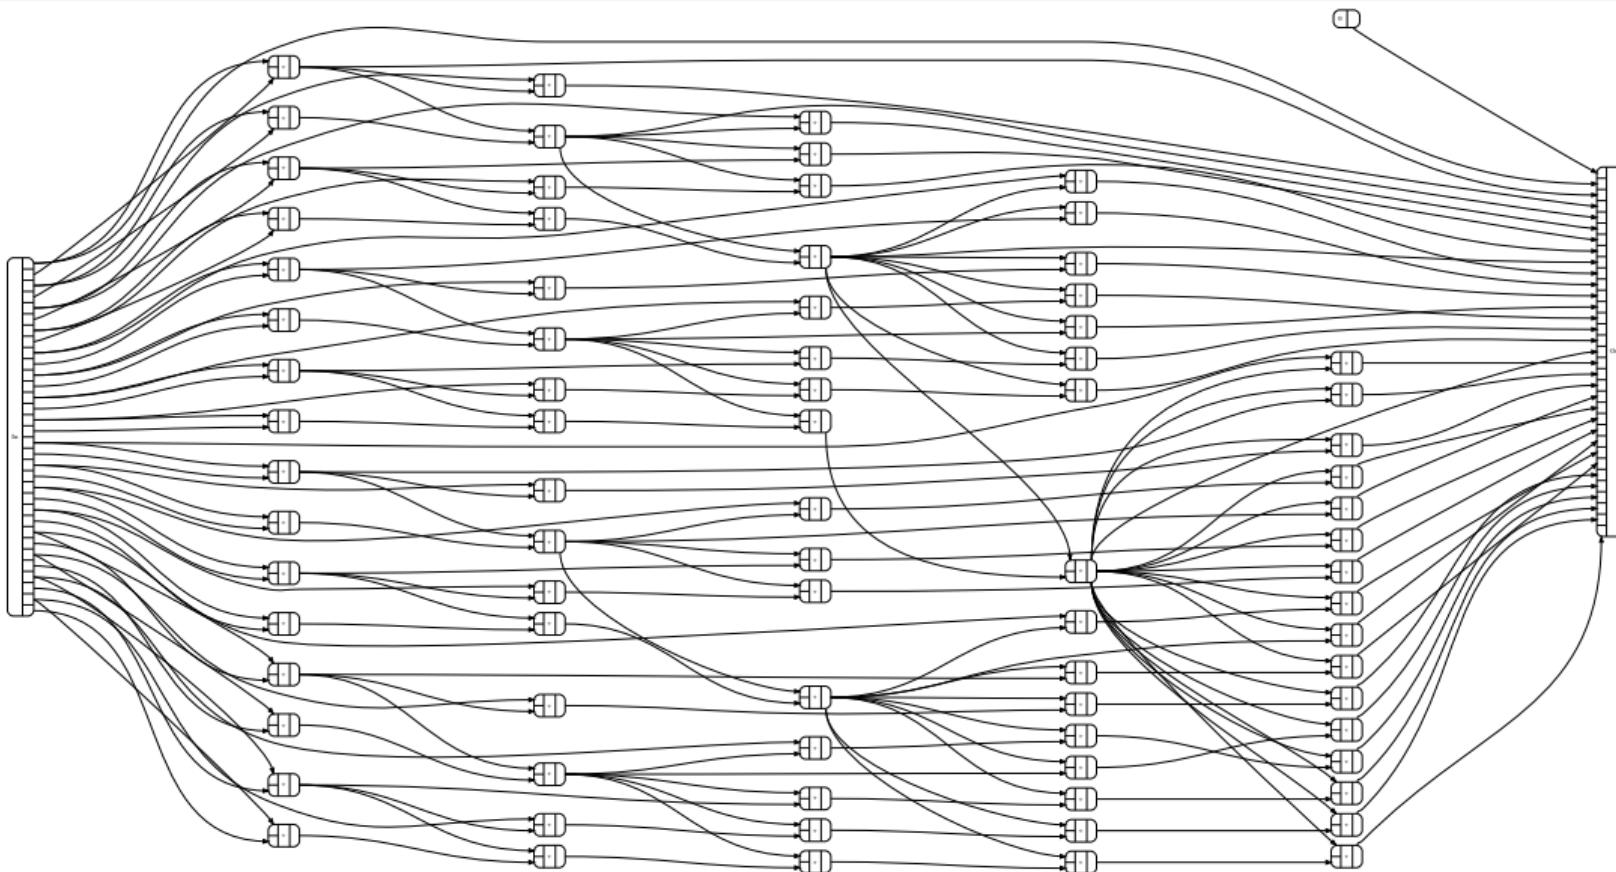
$(v', v_{tot}) = scanT_{\downarrow} v$

Work: $O(n \lg n)$, depth: $O(\lg n)$.

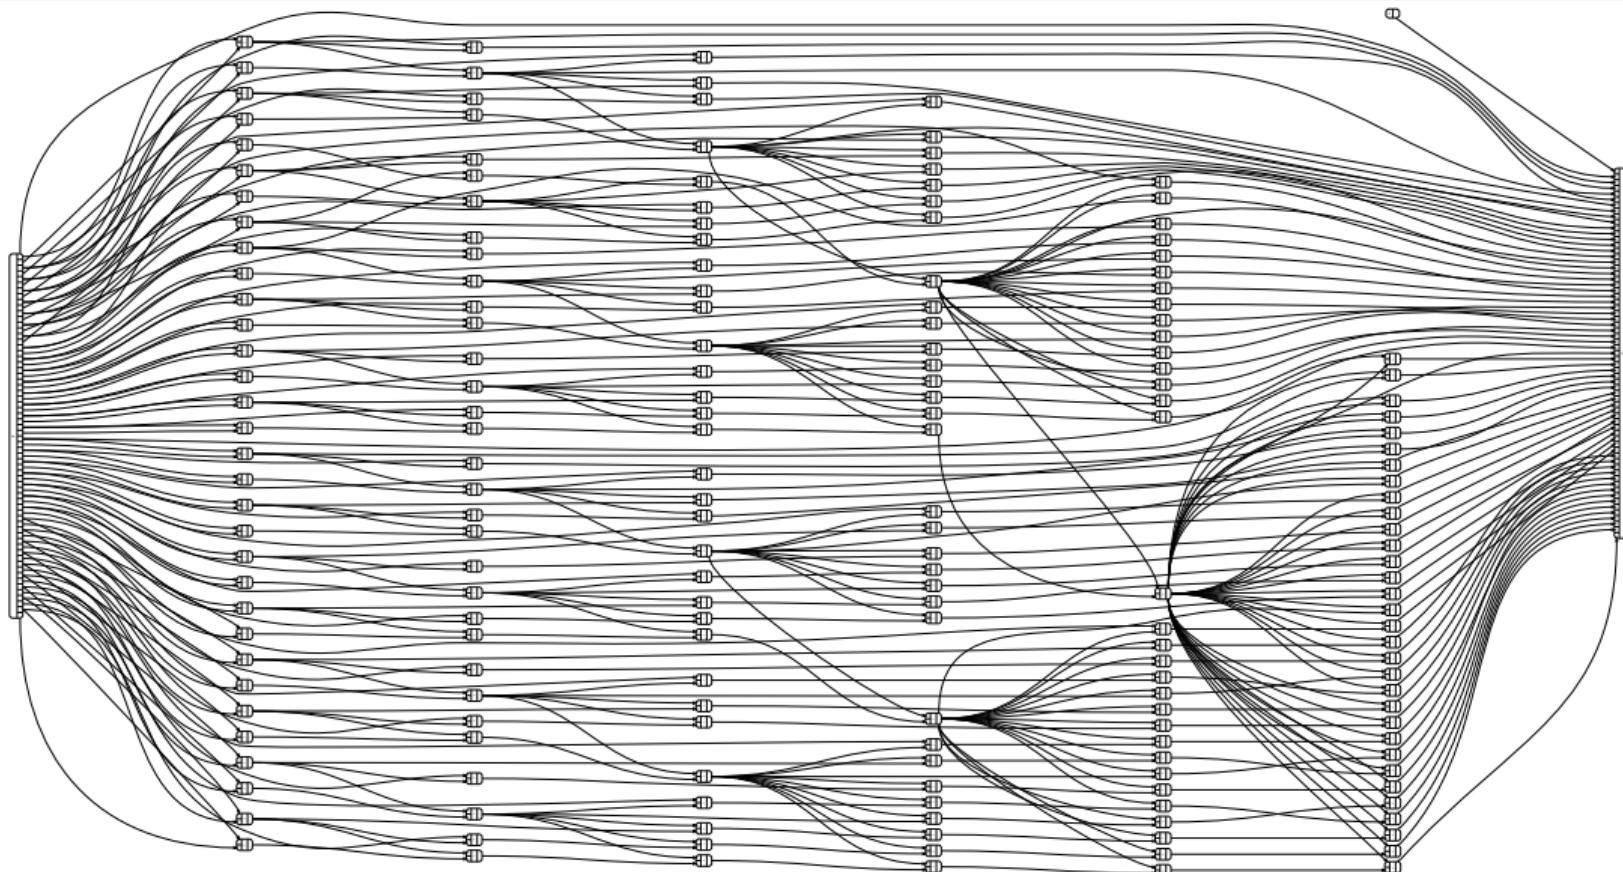
Top-down tree scan



Top-down tree scan



Top-down tree scan



But right answer

$$\begin{array}{ccc} T_{\downarrow} d \ a & \xrightarrow{\quad scanT_{\downarrow} \quad} & T_{\downarrow} d \ a \times a \\ \uparrow & & \uparrow \\ parseT_{\downarrow} & & parseT_{\downarrow} \otimes id \\ \downarrow & & \downarrow \\ Arr 2^d \ a & \xrightarrow{\quad scanA \quad} & Arr 2^d \ a \times a \end{array}$$

Refined wrong guess: top-down, binary, *perfect* trees

```
data  $T_{\downarrow} :: Nat \rightarrow * \rightarrow *$  where
   $L :: a \rightarrow T_{\downarrow} 0 a$ 
   $B :: T_{\downarrow} d a \rightarrow T_{\downarrow} d a \rightarrow T_{\downarrow} (d + 1) a$ 
deriving instance Functor ( $T_{\downarrow} d$ )
```

$scanT_{\downarrow} :: Monoid a \Rightarrow T_{\downarrow} d a \rightarrow T_{\downarrow} d a \times a$
 $scanT_{\downarrow} (L x) = (L \varepsilon, x)$
 $scanT_{\downarrow} (B u v) = (B u' (fmap (u_{tot} \oplus) v'), u_{tot} \oplus v_{tot})$

where

$$(u', u_{tot}) = scanT_{\downarrow} u$$
$$(v', v_{tot}) = scanT_{\downarrow} v$$

Wrong guess refactored

```
data P a = a :# a deriving Functor
```

```
data T↓ :: Nat → * → * where
```

```
    L :: a → T↓ 0 a
```

```
    B :: P (T↓ d a) → T↓ (d + 1) a
```

```
deriving instance Functor (T↓ d)
```

$scanT_{\downarrow} :: Monoid a \Rightarrow T_{\downarrow} d a \rightarrow T_{\downarrow} d a \times a$

$scanT_{\downarrow} (L x) = (L \varepsilon, x)$

$scanT_{\downarrow} (B (u :# v)) = (B (u' :# fmap (u_{tot} \oplus) v'), u_{tot} \oplus v_{tot})$

where

$(u', u_{tot}) = scanT_{\downarrow} u$

$(v', v_{tot}) = scanT_{\downarrow} v$

Wrong guess: top-down, binary, perfect trees

```
data P a = a :# a deriving Functor
```

```
data T↓ :: Nat → * → * where
```

```
    L :: a → T↓ 0 a
```

```
    B :: P (T↓ d a) → T↓ (d + 1) a
```

```
deriving instance Functor (T↓ d)
```

```
scanP :: Monoid a ⇒ P a → P a × a
```

```
scanP (x :# y) = (ε :# x, x ⊕ y)
```

```
scanT↓ :: Monoid a ⇒ T↓ d a → T↓ d a × a
```

```
scanT↓ (L x) = (L ε, x)
```

```
scanT↓ (B ts) = (B (zipWithP tweak tots' ts'), tot)
```

where

```
(ts', tots) = unzipP (fmap scanT↓ ts)
```

```
(tots', tot) = scanP tots
```

```
tweak x = fmap (x ⊕)
```

Work: $O(n \lg n)$, depth: $O(\lg n)$.

Right guess: *bottom-up*, perfect, binary trees

```
data P a = a :# a deriving Functor
```

```
data T↑ :: Nat → * → * where
```

```
  L :: a → T↑ 0 a
```

```
  B :: T↑ d (P a) → T↑ (d + 1) a
```

```
deriving instance Functor (T↑ d)
```

```
scanP :: Monoid a ⇒ P a → P a × a
```

```
scanP (x :# y) = (ε :# x, x ⊕ y)
```

```
scanT↑ :: Monoid a ⇒ T↑ d a → T↑ d a × a
```

```
scanT↑ (L x) = (L ε, x)
```

```
scanT↑ (B ps) = (B (zipWithT↑ tweak tots' ps'), tot)
```

where

```
(ps', tots) = unzipT↑ (fmap scanP ps)
```

```
(tots', tot) = scanT↑ tots
```

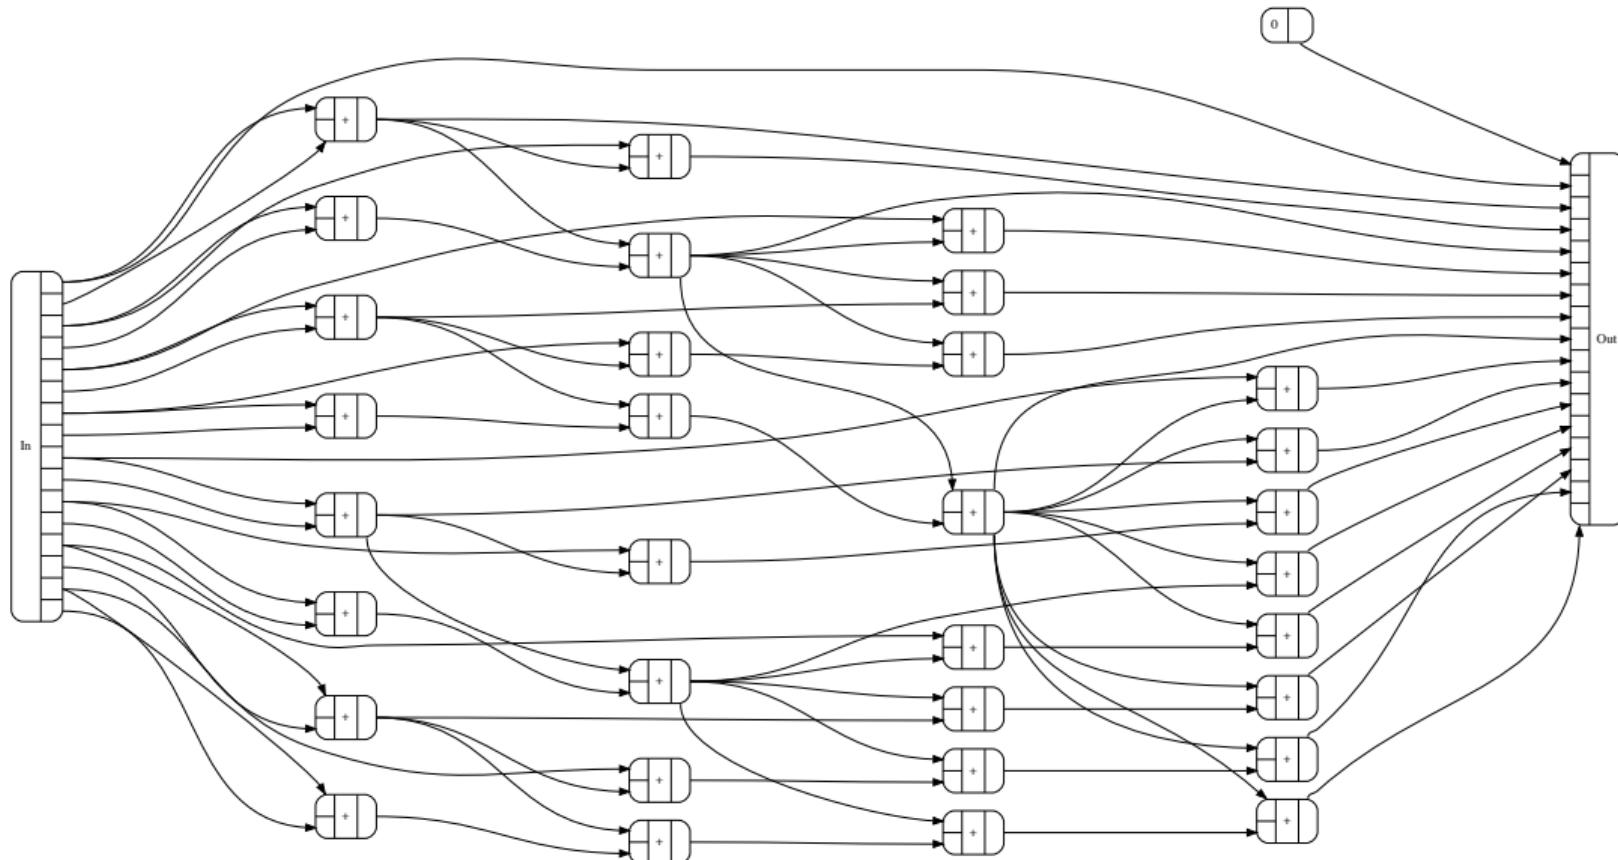
```
tweak x = fmap (x ⊕)
```

Work: $O(n)$, depth: $O(\lg n)$.

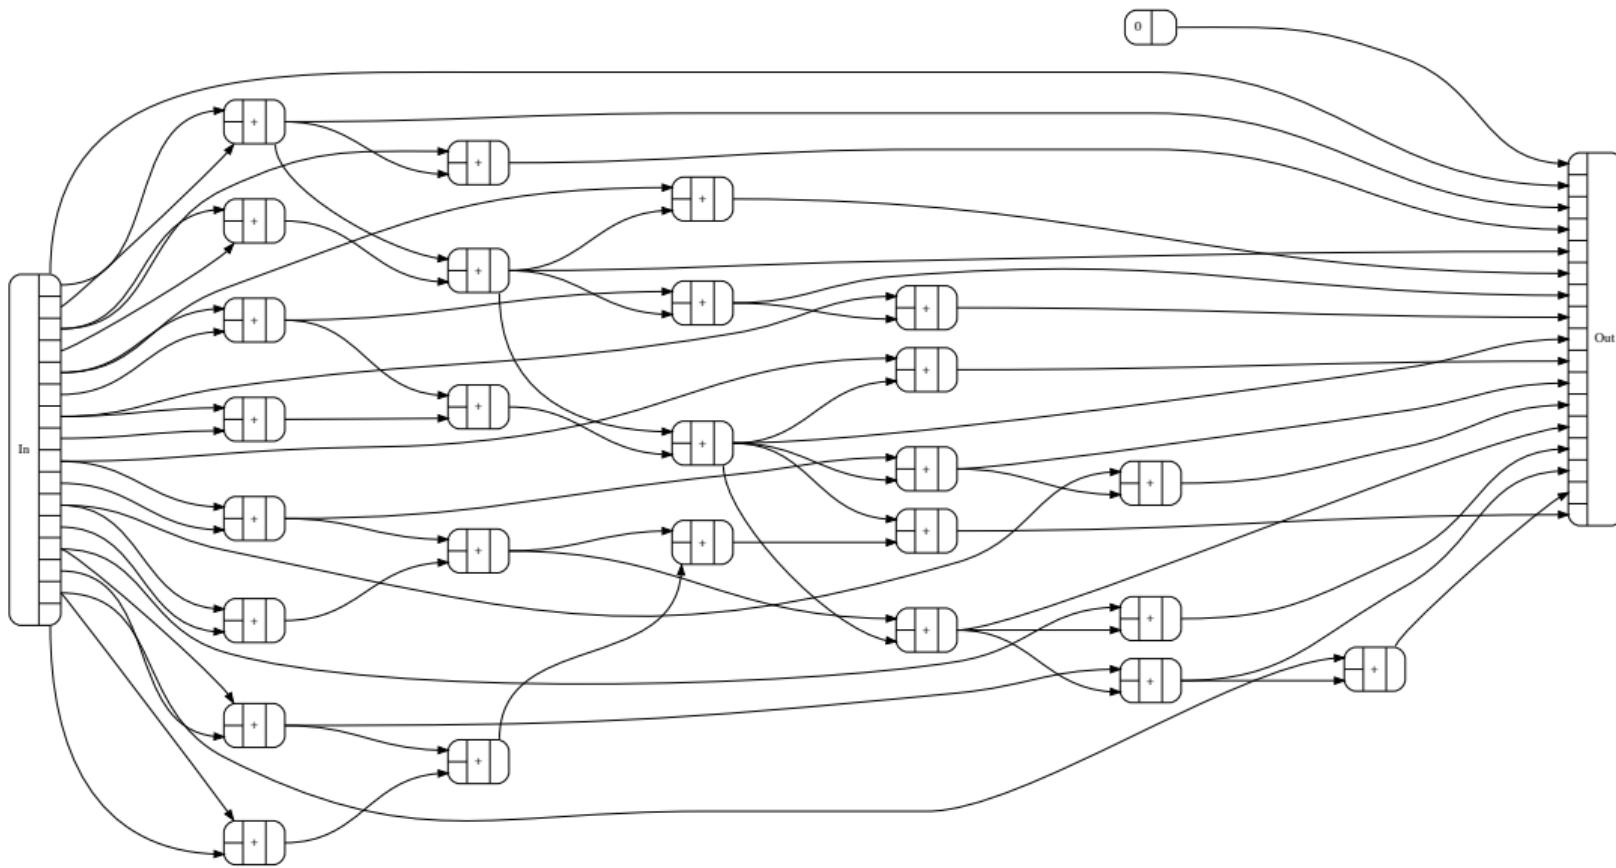
And right answer

$$\begin{array}{ccc} T_{\uparrow} d \ a & \xrightarrow{\quad scanT_{\uparrow} \quad} & T_{\uparrow} d \ a \times a \\ \uparrow & & \uparrow \\ parseT_{\uparrow} & & parseT_{\uparrow} \otimes id \\ \downarrow & & \downarrow \\ Arr 2^d \ a & \xrightarrow{\quad scanA \quad} & Arr 2^d \ a \times a \end{array}$$

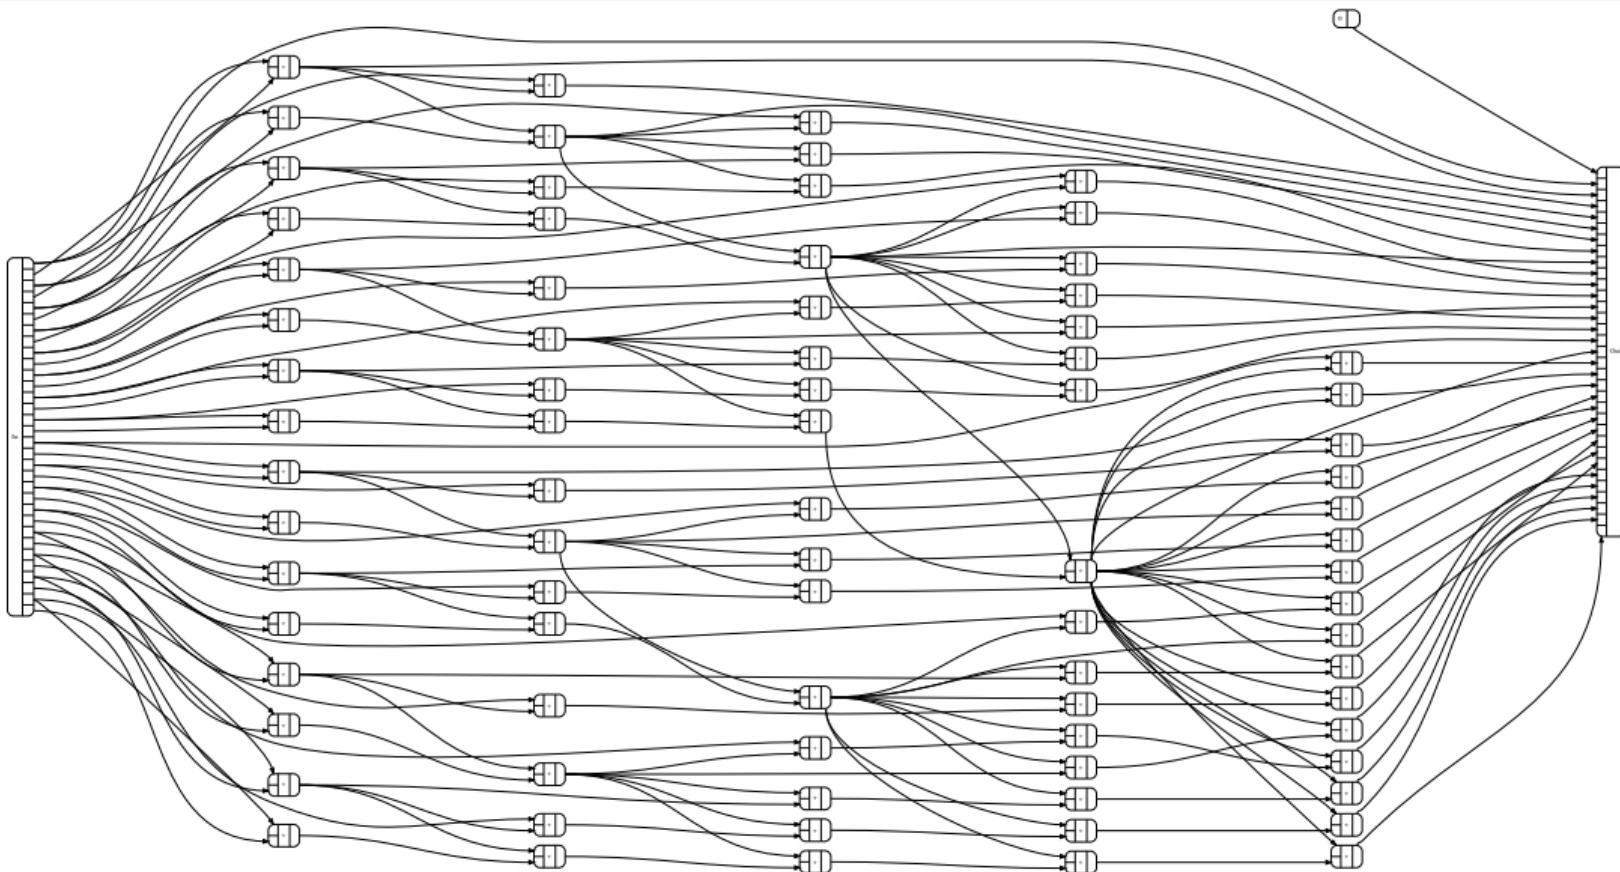
Top-down tree scan



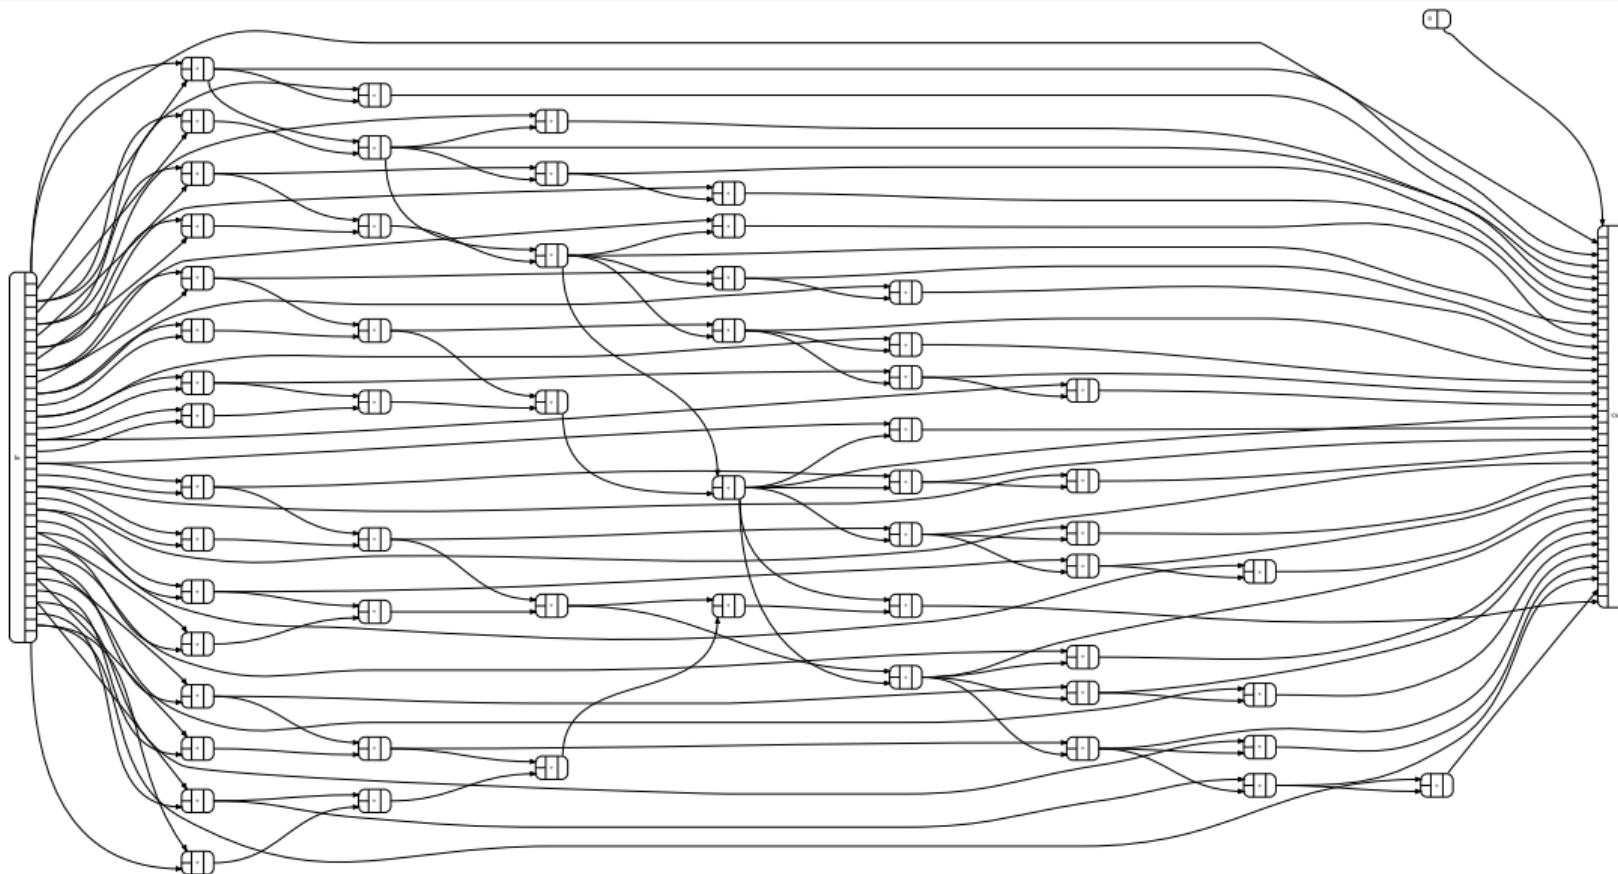
Bottom-up tree scan



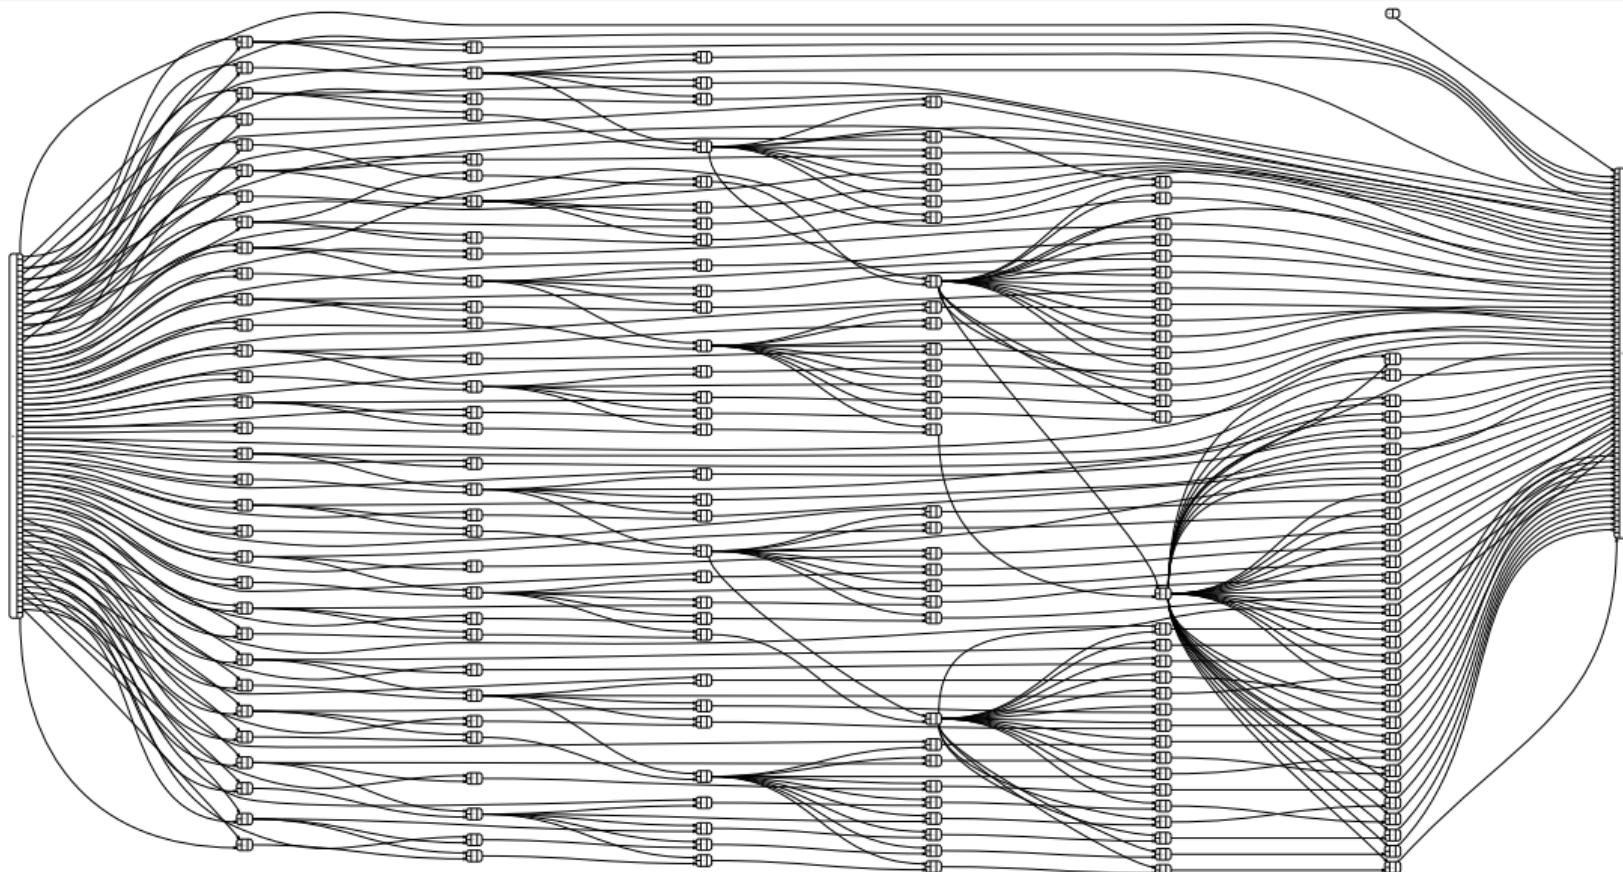
Top-down tree scan



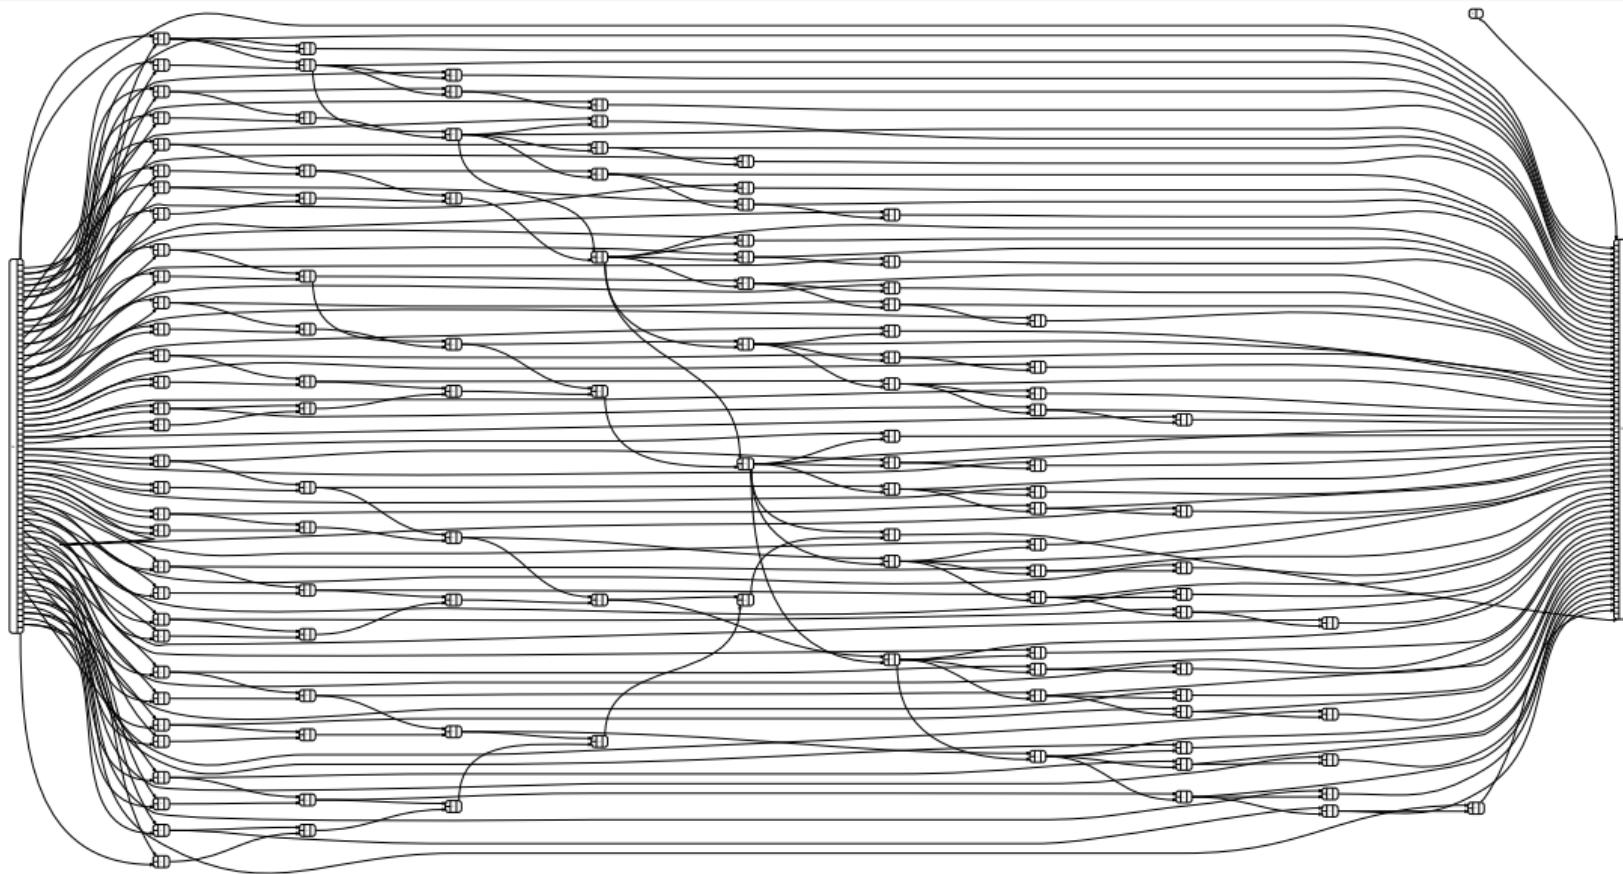
Bottom-up tree scan



Top-down tree scan



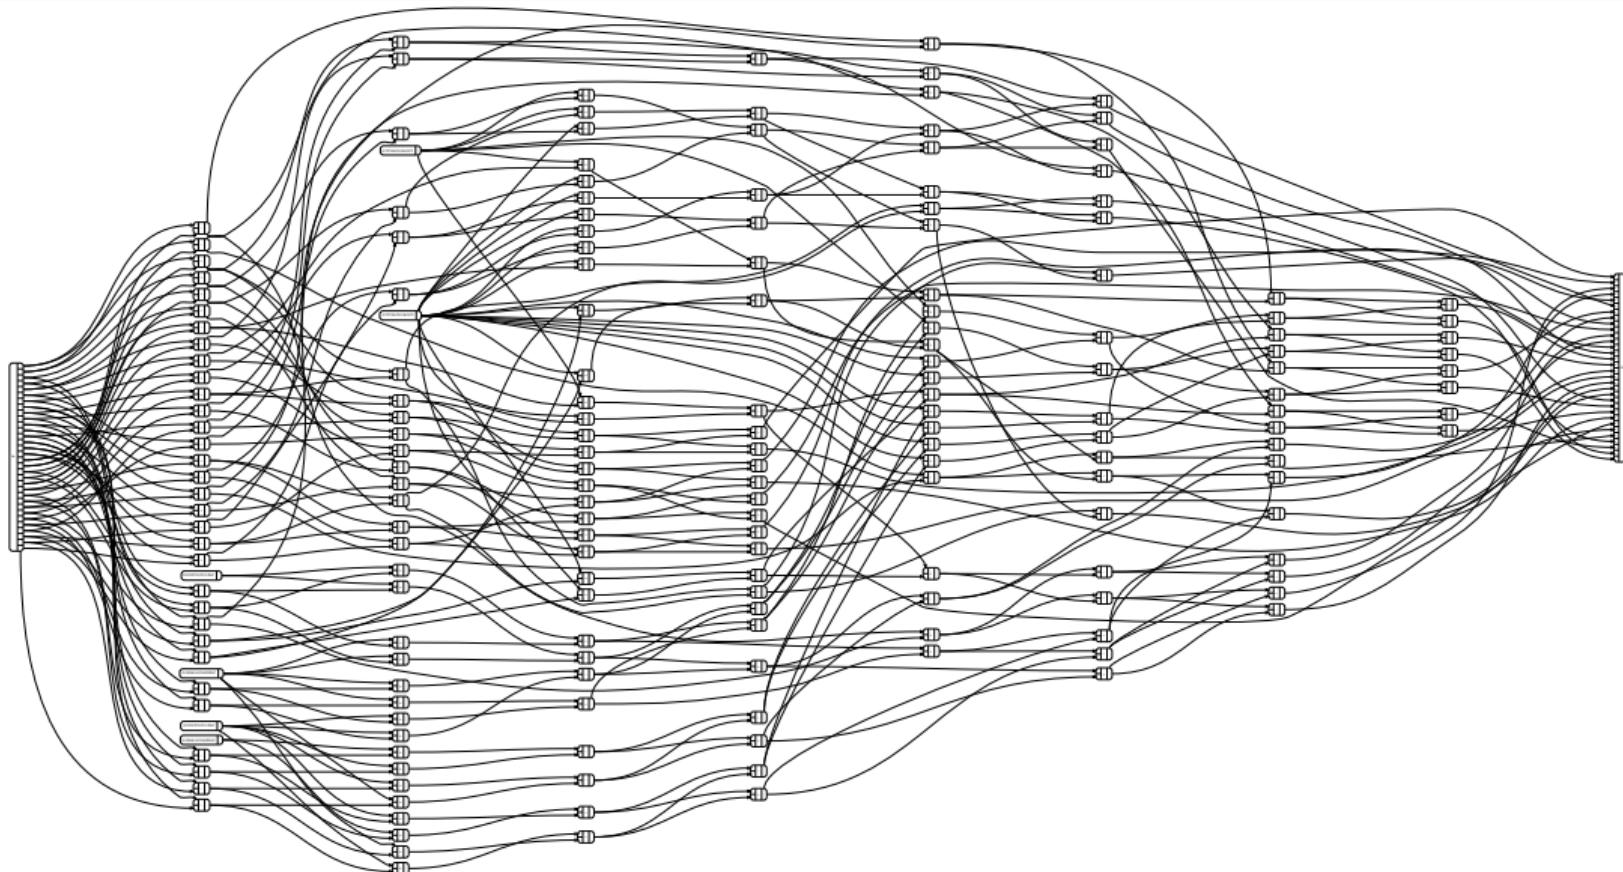
Bottom-up tree scan



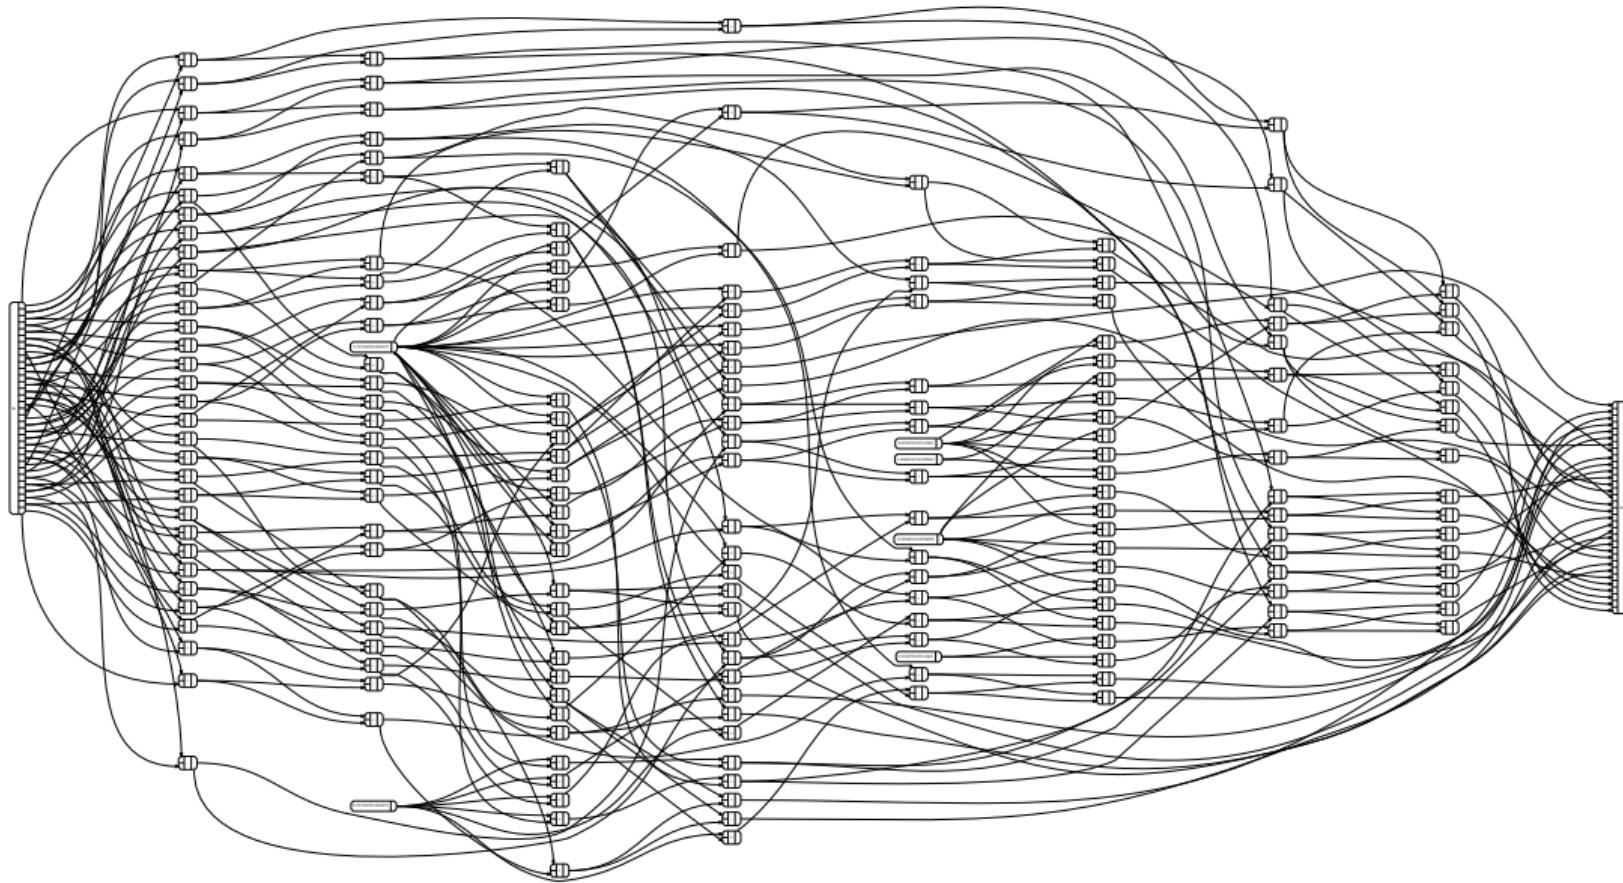
FFT decomposes similarly, yielding classic DIT & DIF algorithms.

See *Generic functional parallel algorithms: Scan and FFT* (ICFP 2017).

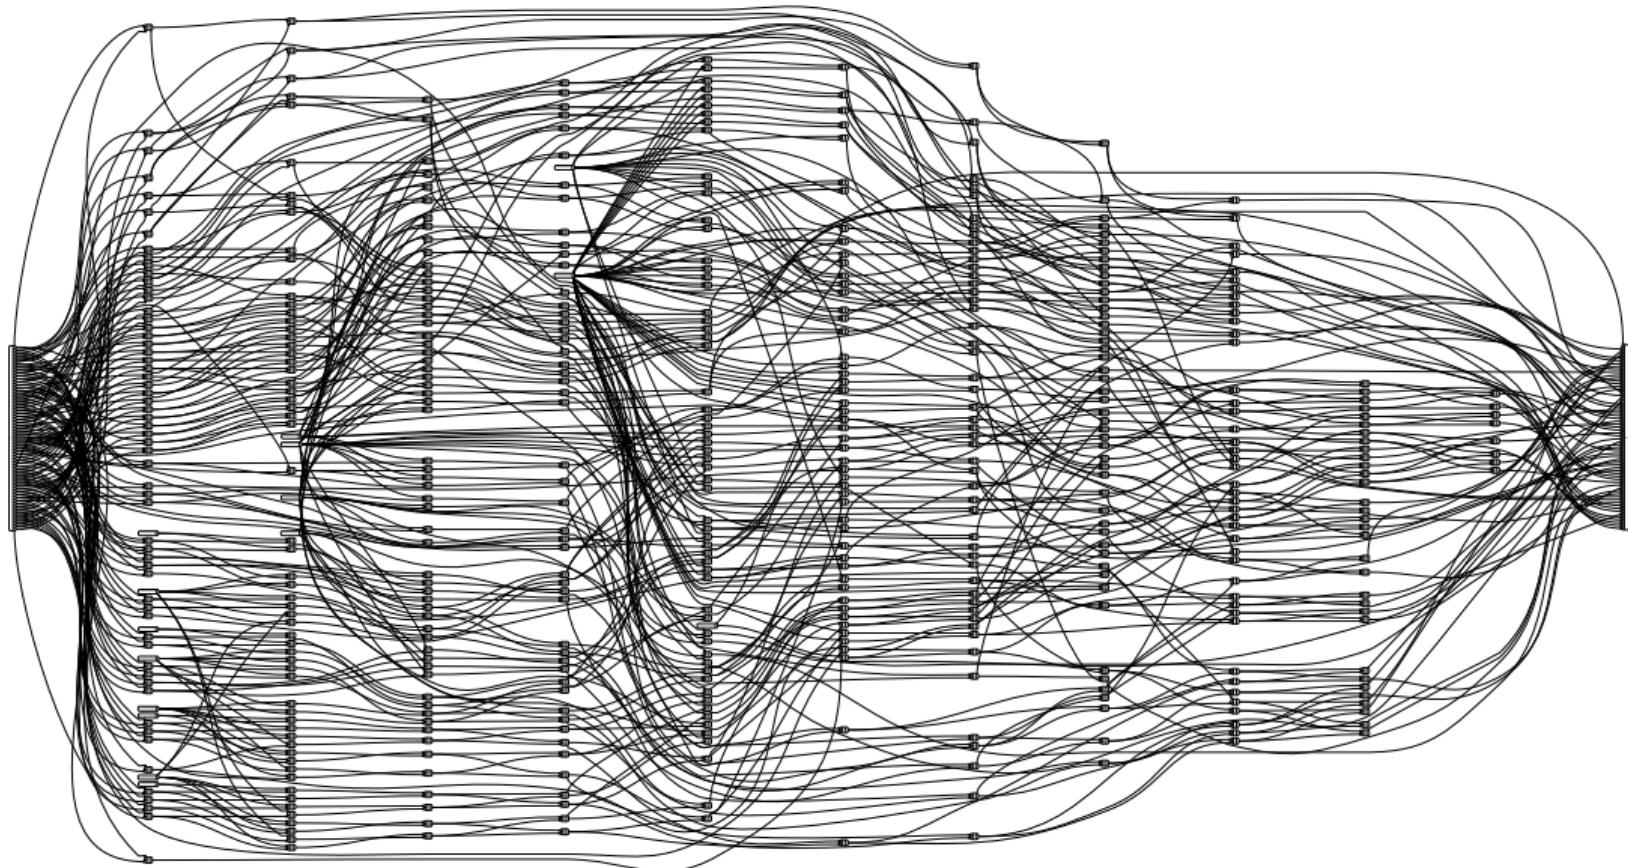
Top-down tree FFT (DIT)



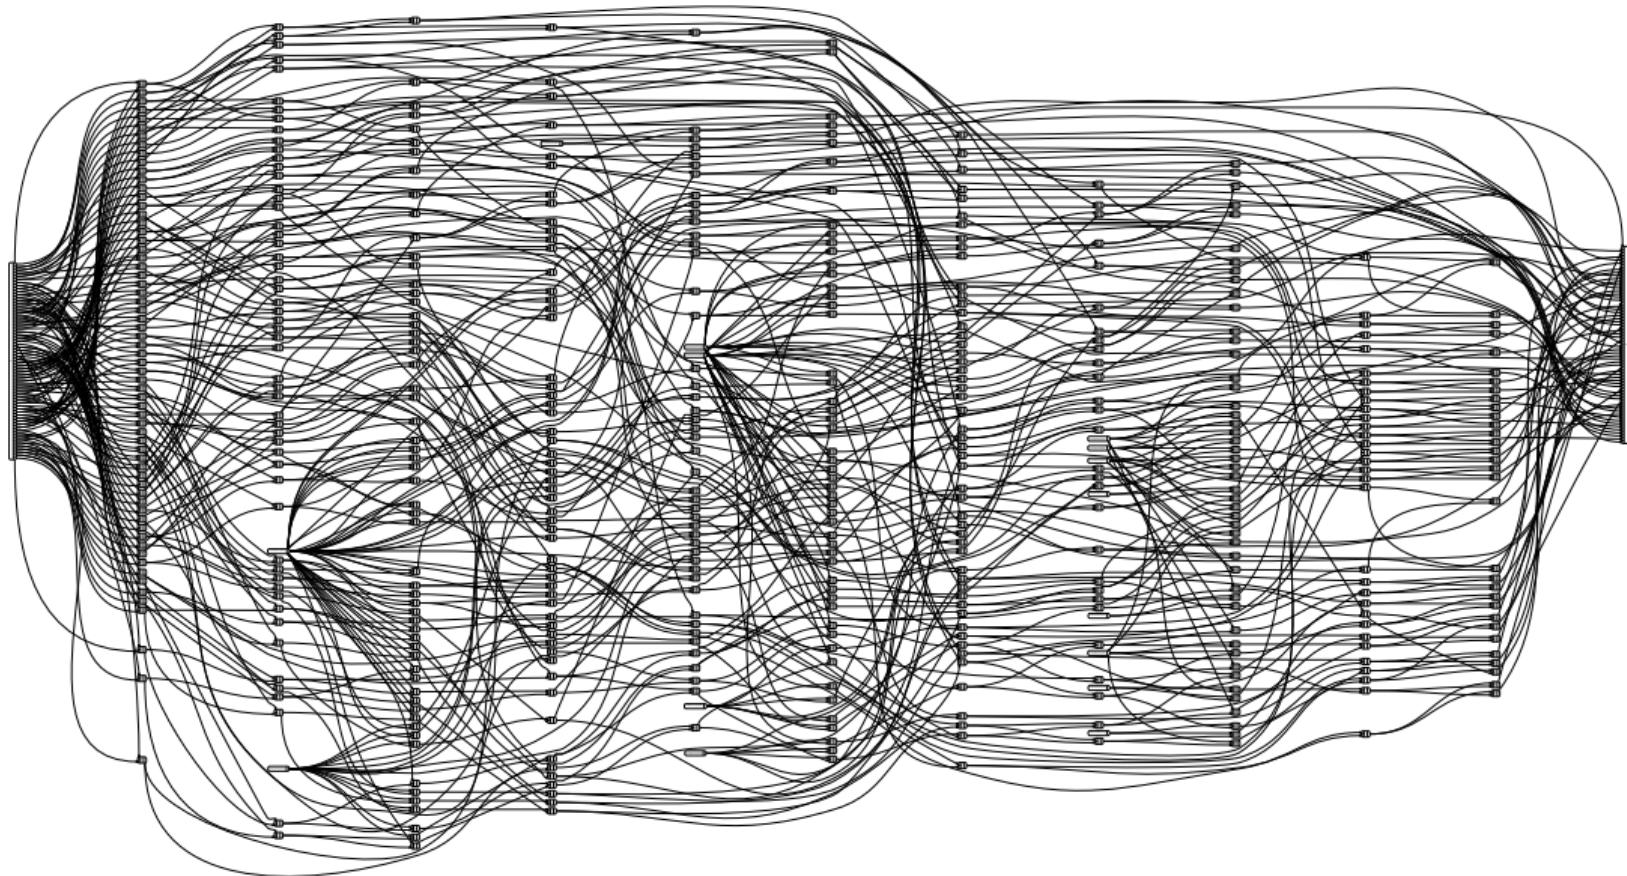
Bottom-up tree FFT (DIF)



Top-down tree FFT (DIT)



Bottom-up tree FFT (DIF)



Generalizing

Re-express parallel algorithm via singletons, products, and *compositions*.

Recomposing yields infinite family of *correct* parallel algorithms on tries.

All such tries are isomorphic to arrays (“parsing/unparsing”).

Arrays are numeric exponentials

$$a^0 = 1$$

$$a^1 = a$$

$$a^{m+n} = a^m \times a^n$$

$$a^{m \times n} = (a^m)^n$$

Tries are general exponentials

data U $a = U$ $\text{Arr } 0$ $\cong U$

data I $a = I a$ $\text{Arr } 1$ $\cong I$

data $(f \times g)$ $a = X(f a) (g a)$ $\text{Arr } (m + n) \cong \text{Arr } m \times \text{Arr } n$

data $(g \circ f)$ $a = O(g(f a))$ $\text{Arr } (m \times n) \cong \text{Arr } n \circ \text{Arr } m$

Tries are general exponentials

```
data U      a = U          ˜0   :: Arr 0 ≅ U
data I      a = I a        ˜1   :: Arr 1 ≅ I
data (f × g) a = X (f a) (g a)  (˜+) :: Arr m ≅ f → Arr n ≅ g → Arr (m + n) ≅ f × g
data (g ° f) a = O (g (f a))  (˜*) :: Arr m ≅ f → Arr n ≅ g → Arr (m * n) ≅ g ° f
```

-- “Parse/unparse”

```
data (≈) :: (* → *) → (* → *) → * where
  Iso :: (forall a. f a → g a) → (forall a. g a → f a) → f ≈ g
  -- Plus isomorphism proof
```

Vectors

$$\bar{n} = \overbrace{I \times \cdots \times I}^{n \text{ times}}$$

Right-associated (“cons”):

type family \vec{d} **where**

$$\vec{0} = U$$

$$\overrightarrow{d+1} = I \times \vec{d}$$

Left-associated (“snoc”):

type family \overleftarrow{d} **where**

$$\overleftarrow{Z} = U$$

$$\overleftarrow{d+1} = \overleftarrow{d} \times I$$

Perfect trees

$$h^d = \overbrace{h \circ \cdots \circ h}^{d \text{ times}}$$

Right-associated (“top-down”):

```
type family h↓d where
  h↓0 = I
  h↓d+1 = h  $\circ$  h↓d
```

Left-associated (“bottom-up”):

```
type family h↑d where
  h↑0 = I
  h↑d+1 = h↑d  $\circ$  h
```

```
type family  $2^{2^d}$  where
```

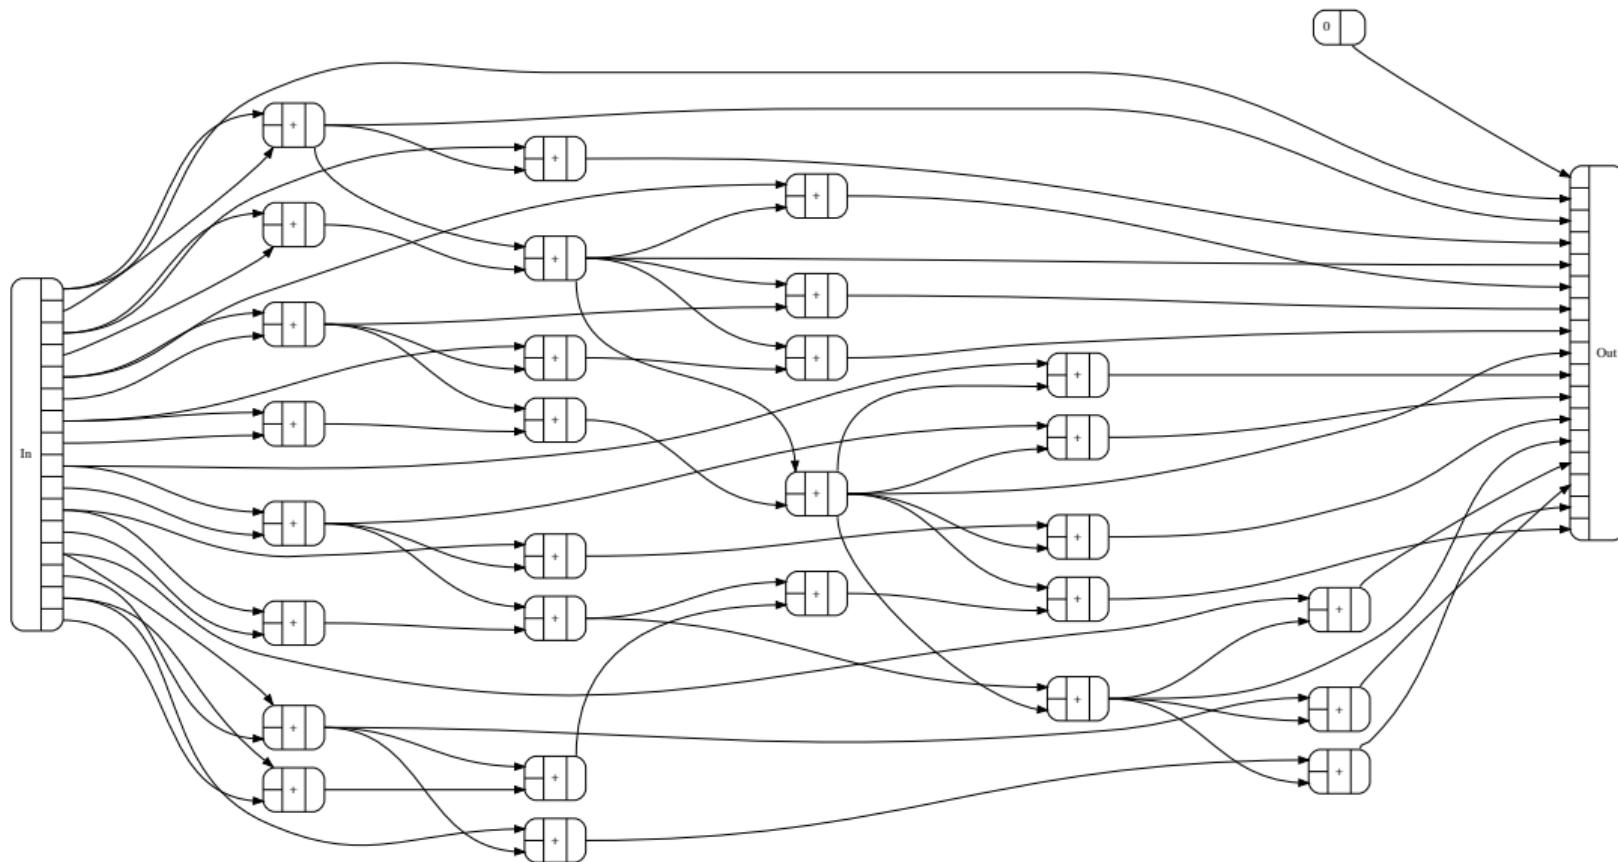
$$2^{2^0} = 2$$

$$2^{2^{d+1}} = 2^{2^d} \circ 2^{2^d}$$

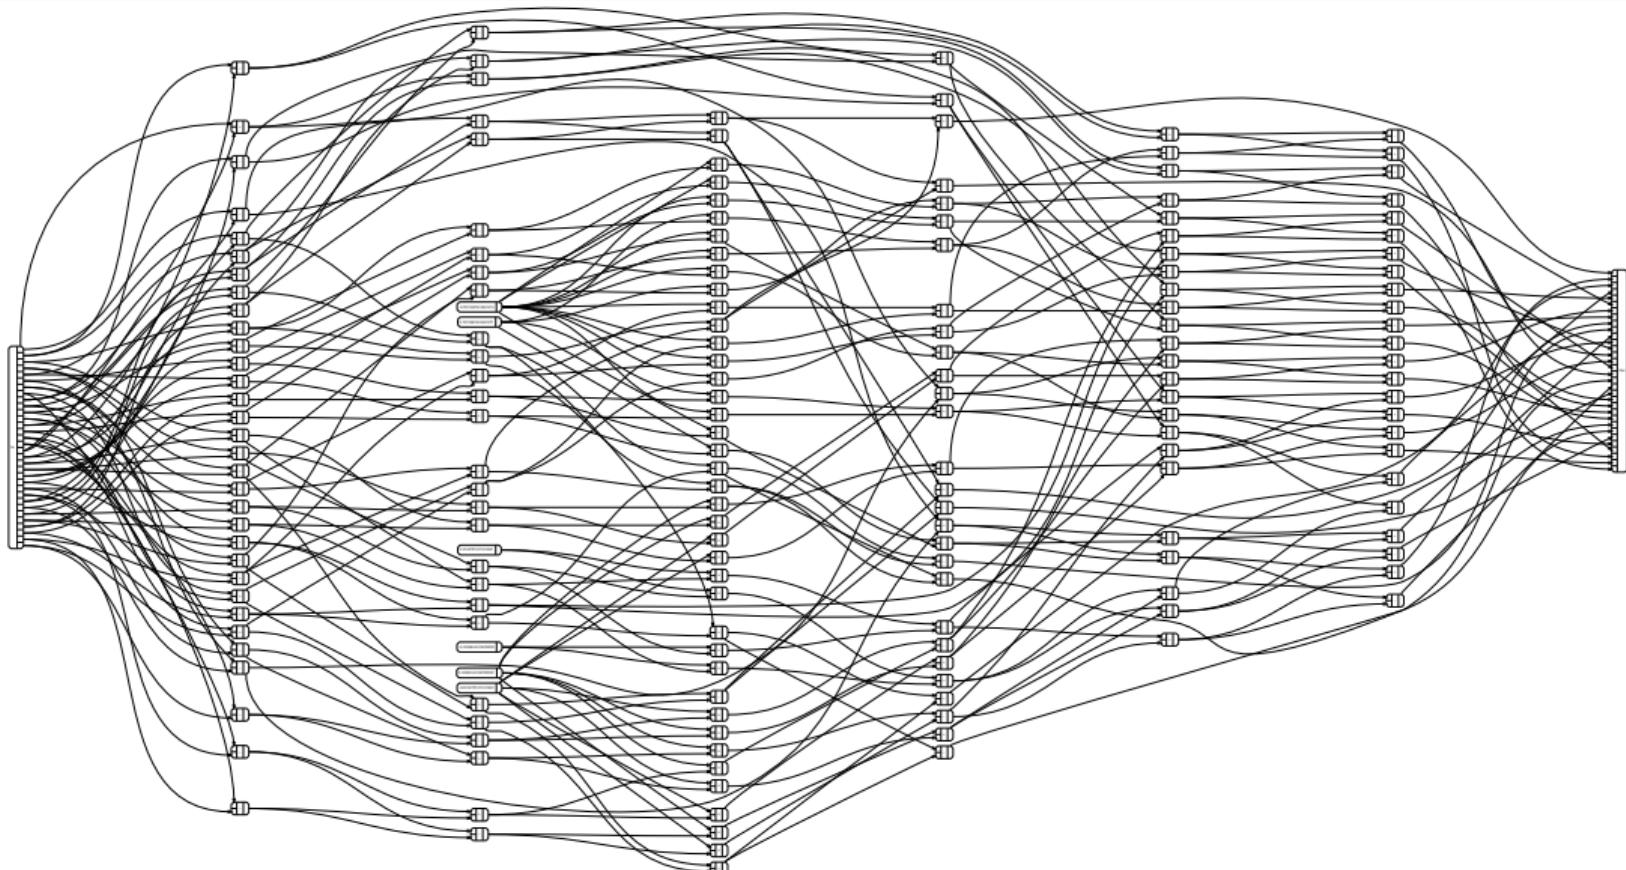
Notes:

- Variation of *Bush* type in [*Nested Datatypes*](#) by Bird & Meertens.
- Size 2^{2^n} , i.e., 2, 4, 16, 256, 65536,
- Composition-balanced counterpart to $2^{\uparrow 2^d}$ and $2^{\downarrow 2^d}$.
- Easily generalizes beyond pairing and squaring.

Bush scan



Bush FFT



Scan comparison

Size 16:

Type	work	depth
$2^{\downarrow 4}$	32	4
$2^{\uparrow 4}$	26	6
2^{2^2}	29	5

Size 256:

Type	work	depth
$2^{\downarrow 8}$	1024	8
$2^{\uparrow 8}$	502	14
2^{2^3}	718	10

FFT comparison

Size 16:

Type	+	\times	-	work	depth
$2^{\downarrow 4}$	74	40	74	188	8
$2^{\uparrow 4}$	74	40	74	188	8
2^{2^2}	72	32	72	176	6

Size 256:

Type	+	\times	-	work	depth
$2^{\downarrow 8}$	2690	2582	2690	7692	20
$2^{\uparrow 8}$	2690	2582	2690	7692	20
2^{2^3}	2528	1922	2528	6978	14

“But computer memory is arrays!”

- Systematically convert natural algorithms to accommodate.
- It's really $2^{\uparrow d}$ or $2^{\downarrow d}$.

Conclusions

- Alternative to array programming:
 - Reveals algorithm essence, connections, and generalizations.
 - Free of index computations (safe and uncluttered).
 - Translates to array program safely and systematically.
- Four well-known parallel algorithms: $h^{\downarrow d}$, $h^{\uparrow d}$.
- Two new and possibly useful algorithms: 2^{2^d} .
- Other examples: arithmetic, linear algebra, polynomials, bitonic sort.
- *Optimization matters but harms clarity and composability, so do it late.*