

# Nested Data Parallelism in Haskell

Manuel M. T. Chakravarty  
University of New South Wales

Joint work with Gabriele Keller

Slide 1

- ① Nesl-style parallelism in Haskell
- ② Flattening Transformation

## PARALLEL ARRAYS

$$\text{Sparse matrix: } \begin{pmatrix} 5 & 0 & 8 \\ 0 & 0 & 0 \\ 9 & 0 & 0 \end{pmatrix} * \begin{pmatrix} -3 \\ 2 \\ 9 \end{pmatrix}$$

**type** *SparseRow* = `[(Int, Float)]`  
**type** *SparseMatrix* = `[:SparseRow:]`

E.g., `[::(0, 5), (2, 8):], [:], [:(0, 9):]`

— Sparse matrix vector multiplication

*smvm* :: *SparseMatrix* → `[:Float:]` → *Float*  
*smvm sm vec* =  

$$[:\text{sumP } [x * (\text{vec} \text{!} \text{ col}) \mid (\text{col}, x) \leftarrow \text{row}] \mid \text{row} \leftarrow \text{sm}:]$$
products of one row

## New data type:

- Parallel array with  $\alpha$  elements: `[: $\alpha$ :]`
- For example,

```
data RoseTree  $\alpha$  = Node  $\alpha$  [:RoseTree  $\alpha$ :]
-- e.g., useful to implement Barnes-Hut N-body algorithm
```

## List-like operations:

- Same special syntax, but with `[: . :]` brackets  
E.g., `[x + y | x ← xs | y ← ys:]`
- Prelude functions with suffix *P*  
E.g., `lengthP xs`, `mapP (+1) xs`

## Semantics:

- All elements are demanded simultaneously  
`lengthP [!1, !, 3:] = 3`, whereas `[!1, !, 3:] !: 0 = !`
- Finite length

## HOW SHALL WE IMPLEMENT PARALLEL ARRAYS?

- ① **Very fine-grained multi-threading:**  
`[foo x | x ← xs:]` generates `lengthP xs` threads
  - ✓ Conceptually simple
  - ✓ Provably efficient thread scheduling
  - ✗ Thread granularity
- ② **Flattening transformation:**  
`[foo x | x ← xs:]` becomes `foo† xs` where `foo†` lifted
  - ✓ Improves array performance already on uniprocessors
  - ✓ Portability (DM, GPUs & multicores with vector instructions)
  - ✗ Requires sophisticated compiler technology

➔ Let's look more closely at flattening...

## FLATTENING

What is flattening?

```
foo :: Int -> Int -> Int
foo x y = x * 2 + y
```

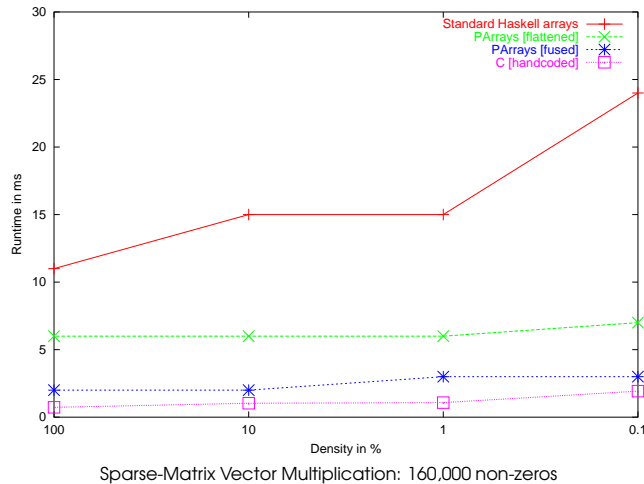
```
[foo x y | x <- xs | y <- ys]
=
foo↑ xs ys
```

```
foo↑ :: [:Int:] -> [:Int:] -> [:Int:]
foo↑ xs ys = [x * 2 + y | x <- xs | y <- ys:]
           = xs *↑ (replicateP (lengthP xs) 2) +↑ ys
```

Slide 5


→ In its full glory more tricky as it has to deal with recursion, higher-order functions, etc.

Sequential performance with flattening:



Slide 6

Slide 7



Portability of flattening...

... to distributed-memory machines:

- There surely will be clusters of multicores
- Flattening gives us a handle on controlling load balancing
- Flat arrays are easier to partition than nested arrays (and other irregular structures)

Slide 8

... to graphical processing units (GPUs):

- Recently became interesting for general-purpose computing
- Stream processing on GPUs generalises classic vector processing
- Flattening seems attractive to widen application domain

... to multicores with vector instructions:

- Use multiple cores, hyperthreads, and vector instructions simultaneously

## IMPLEMENTING FLATTENING FOR HASKELL

### Flattening itself:

- ① Flattening of data structures
- ② Vectorising functions
- ③ Rewriting of closures to support vanilla and vectorised versions of functions

Slide 9

### Supporting transformations:

- Type-indexed array primitives
- Partitioning into threads (guided by types)
- Equational array fusion (improves locality of access)

## CONCLUSIONS


### Nested data parallelism:

- Generalises regular data parallelism
- Convenient programming model for a wide range of applications
- Fits nicely into functional programming languages

Slide 11

### Flattening:

- Transforms nested into flat data parallelism
- Already improves sequential array performance in Haskell
- Promises portability
- Requires further transformations (array fusion)
- Requires significant implementation effort



Slide 10