

Generics in Data Parallel Haskell

Roman Leshchinskiy

Programming Languages and Systems
University of New South Wales

Joint work with
Manuel Chakravarty
Gabriele Keller
Simon Peyton Jones

Programming in DPH

Multiply sparse vector with dense vector

```
svvm :: [:(Int, Float):] → [:(Float):] → Float  
svvm v w = sumP [x * (w !: i) | (i, x) ← v:]
```

Programming in DPH

Multiply sparse vector with dense vector

```
svvm :: [:(Int, Float):] → [:(Float):] → Float  
svvm v w = sumP [x * (w !: i) | (i, x) ← v:]
```

Multiply sparse matrix with dense vector

```
smvm :: [:[:(Int, Float):]:] → [:(Float):] → Float  
smvm m w = [svvm v w | v ← m:]
```

Programming in DPH

Quicksort

```
qsort xs | lengthP xs ≤ 1 = xs  
         | otherwise      = yss !: 0 +:+ eqs +:+ yss !: 1
```

where

```
m = xs !: 0
```

```
lts = [:y | y ← xs, y < m:]
```

```
eqs = [:y | y ← xs, y == m:]
```

```
gts = [:y | y ← xs, y > m:]
```

```
yss = [: qsort ys | ys ← [:lts, gts:]:]
```

Nested data parallelism

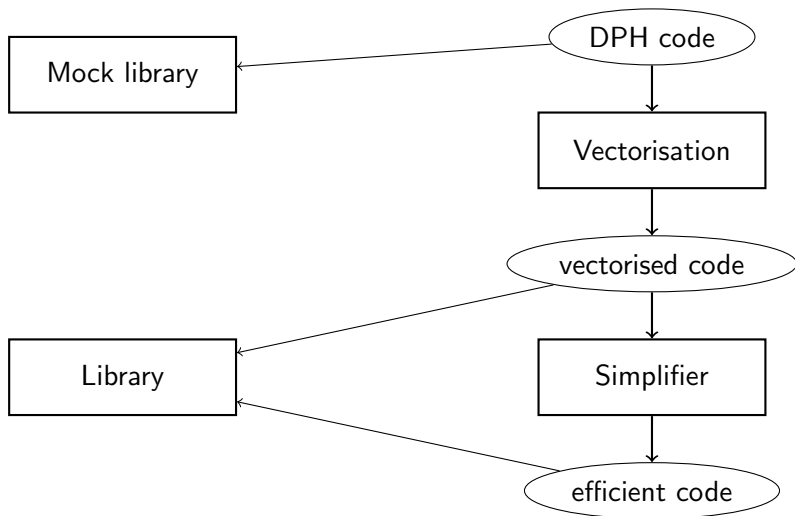
Great for programmers

- single control flow
- implicit synchronisation and communication
- high-level specification of parallelism
- transparent integration into Haskell

Bad for implementors

- extensive program transformations
- needs a very good optimiser

Architecture



Library

Compiler



Library

Compiler

Type families

Library

Compiler

Type families
Rewrite rules

Library

Compiler

Simplifier
Type families
Rewrite rules

Library

Compiler



Simplifier
Type families
Rewrite rules
Parallel RTS

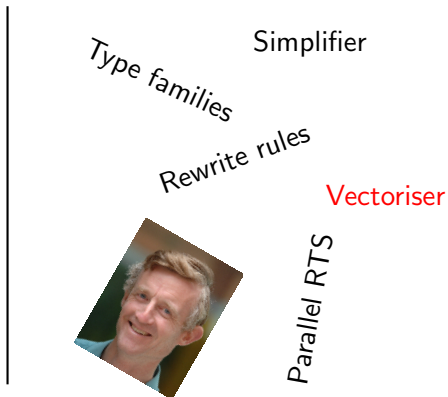
Library

Compiler



Library


Compiler



Library

Fusion

Compiler



Library

Fusion

Gang parallelism

Compiler

Type families

Rewrite rules

Simplifier

Vectoriser




Parallel RTS

Library

Fusion
Gang parallelism
Type transformation

Compiler

Simplifier
Type families
Rewrite rules
Vectoriser
Parallel RTS



Library

Generic operations

Fusion

Gang parallelism

Type transformation

Compiler

Simplifier

Type families

Rewrite rules

Vectoriser




Parallel RTS

Vectorisation – Basic idea

Flattening transformation (Blelloch 1995)

Convert a nested data parallel program into a flat data parallel program

- store nested arrays as flat arrays + segmenting information
- modify computations accordingly




Vectorisation – Basic idea

Flattening transformation (Blelloch 1995)

Convert a nested data parallel program into a flat data parallel program

- store nested arrays as flat arrays + segmenting information
- modify computations accordingly




Vectorisation – Basic idea

Flattening transformation (Blelloch 1995)

Convert a nested data parallel program into a flat data parallel program

- store nested arrays as flat arrays + segmenting information
- modify computations accordingly



svvm[↑]

Representing arrays

Non-parametric representation

data family $[a:]$

data instance $[Double:] = PDouble ByteArray$

data instance $[(a, b):] = PPair [a:] [b:]$

data instance $[[a]:] = PNest Segd [a:]$

- arrays of primitive types \implies memory blocks, no boxing!
- arrays of tuples \implies tuples of arrays
- nested arrays \implies flat data array + segmenting information
- arrays of trees \implies trees of arrays

Sparse matrices

$[[:(0, 1.3), (2, 0.5):], [(1, 3.4):], [:], [(0, 4.3), (1, 2.2):]]$
 $\implies PNest \langle 2, 1, 0, 2 \rangle (PPair [0, 2, 1, 0, 1:] [1.3, 0.5, 3.4, 4.3, 2.2:])$

Operations on parallel arrays

How do we implement $\text{indexP} :: [a] \rightarrow \text{Int} \rightarrow a$?

Operations on parallel arrays

How do we implement $indexP :: [a] \rightarrow Int \rightarrow a$?

```
class PR a where
```

```
  indexPR :: [a]  $\rightarrow$  Int  $\rightarrow$  a
```

```
instance PR Double where
```

```
  indexPR (PDouble bytes) i = indexU bytes i
```

```
instance (PR a, PR b)  $\Rightarrow$  PR (a, b) where
```

```
  indexPR (PPair xs ys) i = (indexPR xs i, indexPR ys i)
```

```
indexPV :: PR a  $\Rightarrow$  [a]  $\rightarrow$  Int  $\rightarrow$  a
```

```
indexPV = indexPR
```

Operations on parallel arrays

How do we implement $indexP :: [a] \rightarrow Int \rightarrow a$?

```
data PR a = PR { indexPR :: [a]  $\rightarrow$  Int  $\rightarrow$  a }
```

```
dPR_Double :: PR Double
```

```
dPR_Double = PR { indexPR =  $\lambda$ (PDouble bs) i  $\rightarrow$  indexU bs i }
```

```
dPR_Pair :: PR a  $\rightarrow$  PR b  $\rightarrow$  PR (a, b)
```

```
dPR_Pair pra prb = PR { indexPR =  $\lambda$ (PPair xs ys) i  $\rightarrow$   
  (indexPR pra xs i, indexPR prb ys i) }
```

```
indexPV :: PR a  $\rightarrow$  [a]  $\rightarrow$  Int  $\rightarrow$  a
```

```
indexPV = indexPR
```


Handling user-defined types

```
data Complex a = Complex a a
```

```
data instance [:Complex a:] = PComplex [:a:] [:a:]
```

Handling user-defined types

```
data Complex a = Complex a a  
data instance [:Complex a:] = PComplex [:a:] [:a:]
```

Easy?

```
dPR_Complex :: PR a → PR (Complex a)  
dPR_Complex pa = PR { indexPR = λ(PComplex xs ys) i →  
    Complex (indexPR pa xs i) (indexPR pa ys i) }
```

Handling user-defined types

```
data Complex a = Complex a a  
data instance [:Complex a:] = PComplex [:a:] [:a:]
```

Easy?

```
dPR_Complex :: PR a → PR (Complex a)  
dPR_Complex pa = PR { indexPR = λ(PComplex xs ys) i →  
    Complex (indexPR pa xs i) (indexPR pa ys i) }
```

Problem

- *PR* has 14 operations at the moment
- there will be more in the future
- compiler has to know how to generate them

Handling user-defined types

```
data Complex a = Complex a a  
data instance [:Complex a:] = PComplex [:a:] [:a:]
```

Easy with generic programming!

- associate user-defined types with a generic product-sum representation
- fixed set of representation types
- vectoriser generates conversion functions
- “real” operations are implemented in the library

Generic representations

```
type family PRepr a  
type instance PRepr (a, b) = (a, b)  
type instance PRepr (Complex a a) = (a, a)  
type instance PRepr (Maybe a) = Sum2 a ()
```

Generic representations

```
type family PRepr a  
type instance PRepr (a, b) = (a, b)  
type instance PRepr (Complex a a) = (a, a)  
type instance PRepr (Maybe a) = Sum2 a ()
```

```
data PA a = PA {  
  toPRepr      :: a → PRepr a  
  fromPRepr    :: PRepr a → a  
  toArrPRepr   :: [a] → [PRepr a]  
  fromArrPRepr :: [PRepr a] → [a]  
  dictPRepr    :: PR (PRepr a)}
```

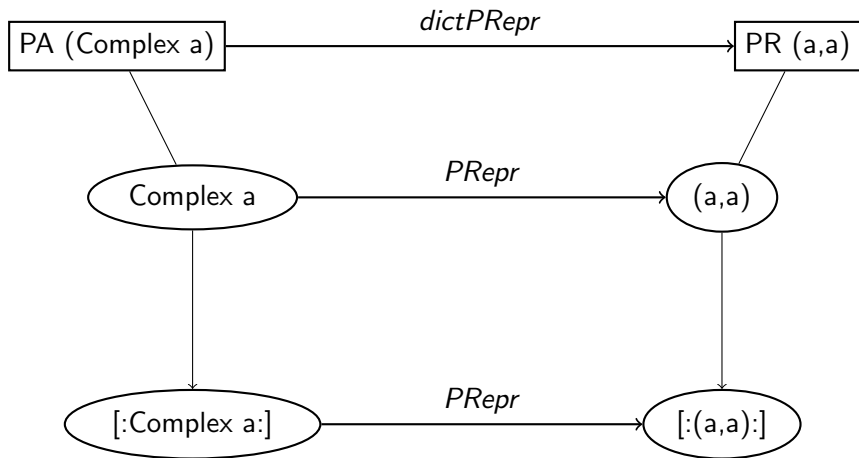
Generic representations

```
type family PRepr a  
type instance PRepr (a, b) = (a, b)  
type instance PRepr (Complex a a) = (a, a)  
type instance PRepr (Maybe a) = Sum2 a ()
```

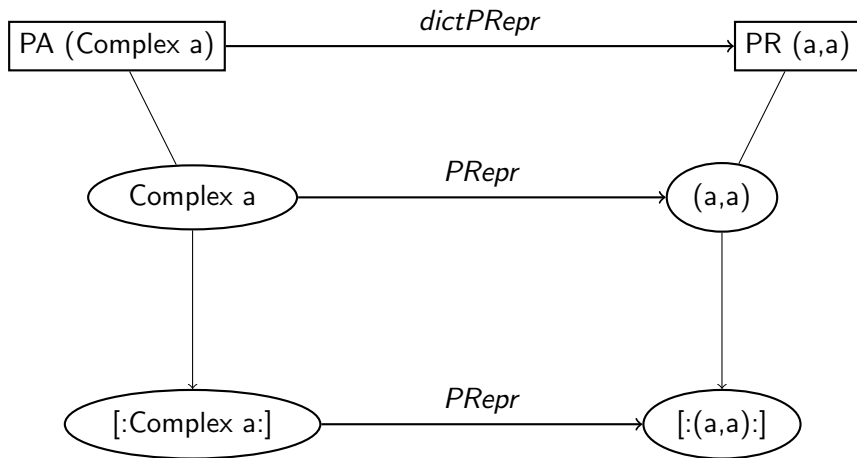
```
data PA a = PA {  
  toPRepr      :: a → PRepr a  
  fromPRepr    :: PRepr a → a  
  toArrPRepr   :: [a] → [PRepr a]  
  fromArrPRepr :: [PRepr a] → [a]  
  dictPRepr    :: PR (PRepr a)}
```

```
dPA_Complex :: PA a → PA (Complex a)  
dPA_Complex pa = PA { toPRepr = λ(Complex x y) → (x, y), ... }
```

Generic operations



Generic operations



type instance *PRepr* [:a:] = [:*PRepr* a:]

Implementing generic operations

Library

$$\begin{aligned} dPR_Pair &:: PR\ a \rightarrow PR\ b \rightarrow PR\ (a, b) \\ dPR_Pair\ pra\ prb &= PR\ \{ indexPR = \lambda(PPair\ xs\ ys)\ i \rightarrow \\ &\quad (indexPR\ pra\ xs\ i, indexPR\ prb\ ys\ i) \} \end{aligned}$$
$$\begin{aligned} indexP_V &:: PA\ a \rightarrow [a] \rightarrow Int \rightarrow a \\ indexP_V\ pa\ xs\ i &= fromPRepr\ pa \\ &\quad (indexPR\ (dictPRepr\ pa)\ (toArrPRepr\ pa\ xs)\ i) \end{aligned}$$

Implementing generic operations

Library

```
dPR_Pair :: PR a → PR b → PR (a, b)
dPR_Pair pra prb = PR { indexPR = λ(PPair xs ys) i →
  (indexPR pra xs i, indexPR prb ys i) }
```

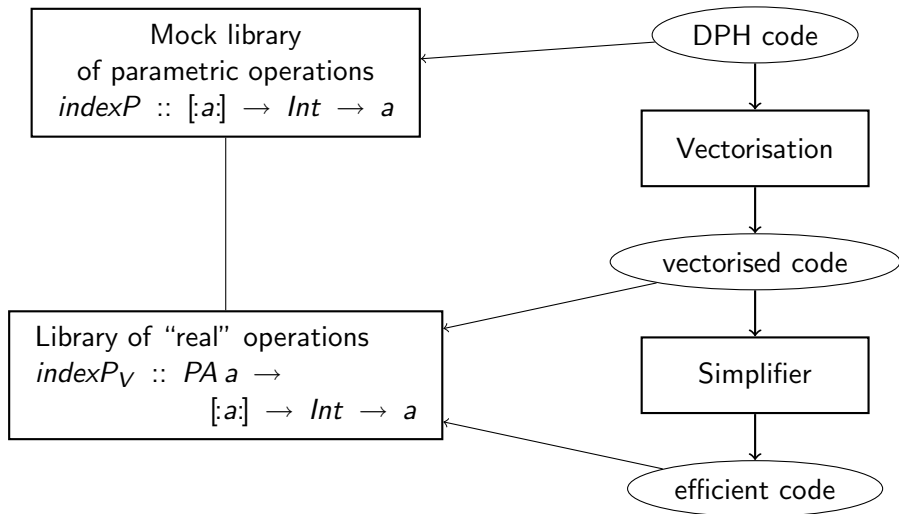
```
indexPV :: PA a → [:a:] → Int → a
indexPV pa xs i = fromPRepr pa
  (indexPR (dictPRepr pa) (toArrPRepr pa xs) i)
```

Vectoriser

```
data instance [:Complex a:] = Complex [:a:] [:a:]
type instance PRepr (Complex a) = (a, a)
```

```
dPA_Complex :: PA a → PA (Complex a)
dPA_Complex pa = PA { toPRepr = λ(Complex x y) → (x, y) ... }
```

Architecture – now with one more line!



Type families + generics are good ...

Type families + generics are good ...

... for DPH

- minimise knowledge built into the vectoriser
- keep library code in the library
- very efficient

Type families + generics are good ...

... for DPH

- minimise knowledge built into the vectoriser
- keep library code in the library
- very efficient
- we can give talks about it!

Type families + generics are good ...

... for DPH

- minimise knowledge built into the vectoriser
- keep library code in the library
- very efficient
- we can give talks about it!

... for everything else

- new approach to implementing generics
- design pattern directly supported by the language
- no performance penalty
- we can write papers about it!