Accelerate High Performance C

High Performance Computing in Haskell Gabriele Keller Trevor McDonell, Ivo Gabe de Wolff, David van Balen, Josh Meredith





many others: Robbert van der Helm, Tom Smeding, Bart Wijgers, Nara Prasetya, Rick van Hoef, Hugo Peters, Tijn Janssen, Jason van den Hurk



What is Accelerate?

Where we want to take it

? What's currently happening

- with Manuel Chakravarty, Sean Lee, Trevor McDonell
- Now most team members based in Utrecht
- DPH:
 - efficiently to concrete architectures via GHC

Started as a spin-off project of DPH (Data Parallel Haskell) at UNSW, Sydney

based on a generalisation of NESL's irregular, nested model, trying to map it

- Accelerate supports a fairly simple model (nested regular data parallelism)
- Idea is to generalise it for certain forms of irregular parallelism
- Initial aim:
- Don't want to depend on GHC for the performance critical part
 - deeply embedded in Haskell

give Haskell programmers a *low overhead way* to exploit the available hardware (GPU, multiple cores on the their desktop/laptop) to achieve performance

Haskell/Accelerate program

GHC

Haskell/Accelerate executable

Reify and optimise Accelerate program



 Data-parallel computations are expressed via operations on multi-dimensional arrays





Continuing Adam's Syntax benchmark:

NumPy matmul = lambda x, y: np.einsum ('ik,kj->ij', x,y)
SaC { [i,j] -> sum ({ [k] -> A[i,k] * B[k,j] }) }
Dex for i j. sum for k. x.i.k * y.k.j
Futhark def matmul i32 [n] [p] (A: [n][m] i32)

A)

(B: [m][p] i32) =
map (\A_row ->
map (\B_col -> reduce (+) 0 (map2 (*) A_row B_col))
 (transpose B))



Matrix multiplication in Accelerate

mmx1 :: Num a => Acc (Matrix a) -> Acc (Matrix a) -> Acc (Matrix a) mmx1 arr brr = sum (generate (I3 cA rB rA) where (I2 CA rA) = shape arr $(I2 \ CB \ rB) = shape \ brr$

mmx2 arr brr = sum (zipWith (*) arrRepl brrRepl) where

arrRepl = replicate (S3 All cB All) arr brrRepl = replicate (S3 rA All All) trr trr = transpose brr (12 cA rA) = shape arr (I2 cB rB) = shape brr

mmx3 arr brr = arr <> brr

(\(I3 i j k) -> arr!(I2 i k) * brr!(I2 k j)))

Index-based vs point-free

Accelerate offers a set of second order array operations

maps, zipWiths

- stencil (convolution matrix) operations
- scans, folds
- permutations and backpermutations/scatter and gathers

conditionals, while loops and such on the array and scalar level

Index-based vs point-free

Many of the specialised functions in Accelerate can be expressed via generate (zipWith, replicate, transpose)

generate shapeOfResult indexToValueFn

if possible, specialised functions should be used Accelerate doesn't try to detect access patterns in generate

memory access patterns are more explicit

mmx2 arr brr = sum (zipWith (*) arrRepl brrRepl) where trr = transpose brr arrRepl = replicate (S3 All cB All) arr brrRepl = replicate (S3 rA All All) trr (I2 CA rA) = shape arr $(I2 \ CB \ rB) = shape \ brr$

as long as the data of an array is only used once work is never copied as a result of fusion

intermediate structures are 'fused away'

A closer look at the types

mmx2 arr brr = sum (zipWith (*) arrRepl brrRepl)

sum = fold (+) 0

fold :: (Shape sh) Elt a) => (Exp a -> Exp a -> Exp a) -> Exp a -> Acc (Array (sh :. Int) a) -> Acc (Array(sh a)









Shape (Rank) polymorphism

The Shape type describes shapes of and indices to multi-dimensional arrays: data $\mathbf{Z} = \mathbf{Z}$

> infix1 : data tail : head = ta

Members of type class Shape have the form

type	DIMO		=	Z
type	DIM1		=	DIMO
type	Scalar	a		Array
type	Vector	a		Array

In the matrix multiplication program:

0	0.00	70707	000000	2202020	50,000	000000	00000	20000	202020	202020	00000	00000	50,000	02020	00000	202020	202020	00000	<u>59595</u>
	267	02020		Coco Col		00000000	000000	000000	50000	000000		000000		25252	00000	000000	0000000	00000	
	NON	oxoxo					o xoxox	o xo xo x	00000	O XO XO X	o cococ	o xo xo x	DXoXoXo	Kokoka	No No No	o xo xo x	o so so so	Noxox	
25	SOIL	0,000					00000	00000	20000	00000	00000	000000	DXOXOXO	Koxoxo)XOXOX	00000	00000000	De la como	
6	000	00000	00000			0000000	00000	00000		000000	000000		200000	596969		000000	000000	202020	
0		20202	000000		2020202	0000000	202020	202020	202020	202020	202020	202020	1020202	0000	10,0,0	202020	202020	10,00,0	
Q	000	000	000000		-0-0-0-	0000000	00000	00000	507070	00000	00000	00000	5050505	00000		00000			
R	202	02020		2000000		0000000	00000	000000	52626	000000	000000	000000		20202	525252	00000		525252	
XC	Xox	0%0%0	Kokoko	KOKOKOK	oxoxoxo	No So	00000	020202	o xo xo x	020202	00000	000000		Koxox	Noxox	00000	0202020	No Koka	
)ĞÇ	No S	0,0,0	00000			00000	00000	00000	00000	00000	00000		2202020	505050	2000	00000	0000000	250505	
ŏ:	59 5	20202	OCOL E	000000		0000000	000000	202020	202020	000000		202020	2020202	000000	O XOXO	00000	No No No		
0						0000000	207200	00000	XOXOXC	KOX KC		202020	0.0000	OXOXO			20 0 00	KO KOKO	

	5
000000000000000	
	3
	3
	3
	5
	3
XOXOXOXOXOXOXOX)>
	2
	2
	2

	69 6				OX.	
DXO-	000	Sos	0.0	XON	6X	2
520-	1000	202	0×0	NON	δX	
DXO-	KOX0	ROX	0×0	NOX	ōΥ.	0
			R •		02	0×
					OX.	07
	(03 b				0	07
2XO-	1070	NON	0,0	NO7	07	0ř
2202	070	SON	070	XOX	07	27
2202	000	NON	000	XOY	25	07
			R ⁰		Qň	Qř
				70-	OX.	



Slices

Similar to Shape, but components include Any, All

slice :: (Slice slix, Elt e) => Acc (Array (FullShape slix) e) -> Exp slix -> Acc (Array (SliceShape slix) e)

replicate :: (Slice slix, Elt e) => Exp slix -> Acc (Array (SliceShape slix) e) -> Acc (Array (FullShape slix) e)



The Elt type class



Ine Elt type class

- All the usual base types are members of this class
 - In floating points, ints of various sizes, char, boolean, and so on
 - shapes and indices
 - n-tuples
 - sum-types, like the Maybe type:

data Maybe a = Nothing



Just a

Elt a => Elt (Maybe a)

The Elt type class

User-defined data types:

Accelerate compiler needs to know

how to represent these types efficiently as matrix elements

class Elt a where type EltR a fromElt :: a -> EltR a toElt



:: EltR a. -> a

The Elt type class

- Default implementations for non-recursive data types are provided
 - product types (tuples, records)
 - arrays of tuples represented as tuples of arrays
 - sum types (alternatives)
 - tuples of arrays of data, flags



Ine Elt type class

- Arrays are not in the Elt type class
- Ragged/irregular matrices are not directly supported
- Segmented operations can be used to express manually flattened computations

foldSeg :: forall sh e i. (Shape sh, Elt e, IsIntegral i) => (Exp e -> Exp e -> Exp e) -> Exp e -> Acc (Array (sh :. Int) e) -> Acc (Segments i) -> Acc (Array (sh :. Int) e)



Ine Elt type class

smvm :: Num a => smvm smat vec = let (T2 segd (T2 inds vals)) = smat vecVals = backpermute inds vec products = zipWith (*) vecVals vals in foldSeg (+) 0 products segd



sum = fold (+) 0

fold :: (Shape sh, Elt a)
 => (Exp) a -> Exp a -> Exp a)
 -> Exp a
 -> (Acc) (Array (sh :. Int) a)
 -> Acc (Array sh a)

Exp a sequential calculation resulting in value of scalar type a
 Acc arr parallel calculation resulting in value of array type arr
 both just represent abstract syntax trees of the computations, not actual values

both just represent abstract syntax to of that type



In generate code and execute it on the CPUs or GPU, we need to call a version of run

a run :: Acc a -> a \mathbf{z} rung :: (Acc a -> Acc b) -> a -> b runN :: (Acc a -> Acc b) -> a -> b



Lifting into the embedded language can be done automatically for overloaded values:

> 0 :: Num a => a (+) :: <u>Num</u> a => a -> a -> a instance Num a => Num (Exp a) where (+) expr1 expr2 = ...



Not so easy for non-overloaded values

not :: Exp Bool -> Exp Bool not False = Trug not Prue = False

lift :: a -> Exp a

not Falk	
	000

not :: Exp Bool -> Exp Bool not b = b ? (lift False, lift True)

(simplified)

xp Bool -> Exp Bool = lift True
= lift False

(Exp t, Exp t) -> Exp t

How can we pattern match on embedded values?

Pattern matching on embedded expressions

Unlifting is not generally possible without evaluation

Exp a -> a

only possible it topmost constructor of a type is unique
 Exp (a, b) -> (Exp a, Exp b)

swap:: Exp (a, b) -> Exp (b, a)
swap xy =
 let
 (x, y) = unlift xy
 in lift (y, x)

Pattern matching on embedded expressions

Haskell pattern synonyms and view patterns

swap :: Exp (a, b) -> Exp (b, a) swap $(T2 \times y) = T2 \times y$

sum (generate (13 cA rB rA)

but this does not for types with multiple constructors - e.g. Bool

((13 i j k) -> arr!(12 i k) * brr!(12 k j)))

Pattern matching on embedded expressions

data Either a b = Left a | Right b
 deriving (Generic, Elt)

mkPattern ''Either

swapE :: Exp (Either a b) ->
swapE (Left_ a) = Right_ a
swapE (Right_ b) = Left_ b

match swapE

swapE :: Exp (Either a b) -> Exp (Either b a)

Current work



Rewrite of the compiler pipeline Ivo Gabe de Wolff & David van Balen

Improved fusion Destructive updates Specification of different schedules on the user level



Improved Fusion

- Currently
 - intermediate arrays whose data is used only once are reliably fused
- In the new pipeline:
 - horizontal, vertical and diagonal fusion
 - takes estimated cost of operations into account
 - using ILP solver to find more efficient schedules

work is never duplicated (even if that might be the right thing to do)

Destructive updates

- When iterating over the whole array, Accelerate is pretty good at handling memory efficiently
- However, destructive updates are important for some classes of algorithms

permute

- :: forall sh sh' a. (Shape sh, Shape sh', Elt a) => (Exp a -> Exp a -> Exp a) -- combination function -> Acc (Array sh' a) -> (Exp sh -> Exp (Maybe sh')) -> Acc (Array sh a) -> Acc (Array sh' a)
- updated destructively



-- default values -- index permutation function -- source values to be permuted

In the new pipeline, we can identify situations where arrays can and should be

Specifying schedules on the user-level

- Currently, user has to rely on Accelerate to pick the schedule
- No explicit representation of the schedule in the AST
- New pipeline will provide high-level constructs for the user to specify schedules and enable optimisations







More sophisticated stencil computations

- We're working with researchers from Geoscience at UU and NIOZ on simulations of coastal developments
- Essentially, stencil computations:

stencil:: (Stencil sh a stencil, Elt b) => (stencil -> Exp b) -> Boundary (Array sh a) -> Acc (Array sh a) -> Acc (Array sh b) stencilFn ((_,t,_), (l,c,r), $(_,b,_)) = t + 1 + c + r + b$

- More efficient scheduling on GPUs
- Language support for dynamic adaption of simulation speed
- Better interface for easy visualisation of data

Creative Commons Attribution-Share Alike 2.0 Generic, John Tushin







Automatic Differentiation

- Tom Smeding's work on reverse automatic differentiation for Accelerate identified some performance pain points
 - generated code nothing like hand-written code Accelerate optimises for
 - more sophisticated fusion required memory management issues need to be addressed
- Size inference?



Thank you!

acceleratehs.org

https://github.com/AccelerateHS/

