

Optimising SpMV for FEM on FPGAs

Paul Grigoras, Pavel Burovskiy, Wayne Luk, Spencer Sherwin






Finite Element Methods - Solve PDEs over large, unstructured geometries

- PDEs: Incompressible Navier Stokes, Shallow Water etc.
- Applications: computational fluid dynamics, biomedicine, geoscience, etc.

Finite Element Methods


Source: www.nektar.info



**Mesh over
unstructured domain**

Finite Element Methods


Source: www.nektar.info



**Mesh over
unstructured domain**




Mesh elements



Finite Element Methods


Source: www.nektar.info



**Mesh over
unstructured domain**



Mesh elements




Assembly

Sparse Matrix

Finite Element Methods


Source: www.nektar.info



**Mesh over
unstructured domain**



Mesh elements




Assembly



Finite Element Methods


Source: www.nektar.info



**Mesh over
unstructured domain**




Source: www.nektar.info



CFD Simulation

Finite Element Methods


Source: www.nektar.info



**Mesh over
unstructured domain**



Mesh elements



Assembly


Sparse Matrix



PDE Solver

Finite Element Methods


Source: www.nektar.info



**Mesh over
unstructured domain**



Mesh elements



Assembly

Sparse Matrix




PDE Solver



Linear Solver

Finite Element Methods


Source: www.nektar.info



**Mesh over
unstructured domain**



Mesh elements



Assembly

Sparse Matrix




PDE Solver



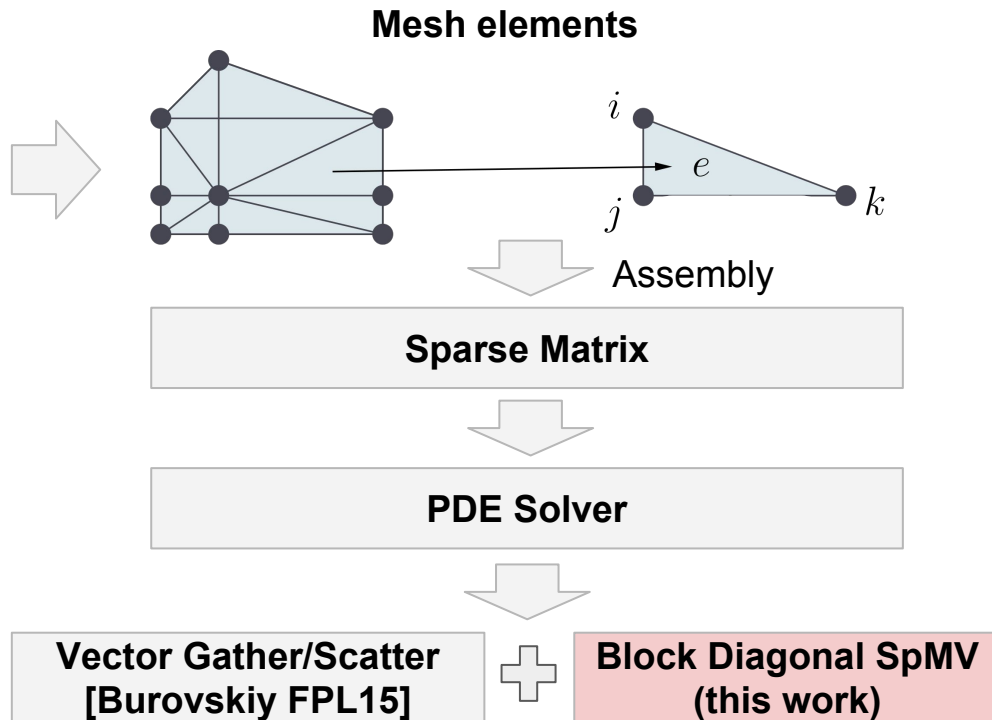
Iterative Linear Solver \Rightarrow SpMV

Finite Element Methods

Source: www.nektar.info






**Mesh over
unstructured domain**



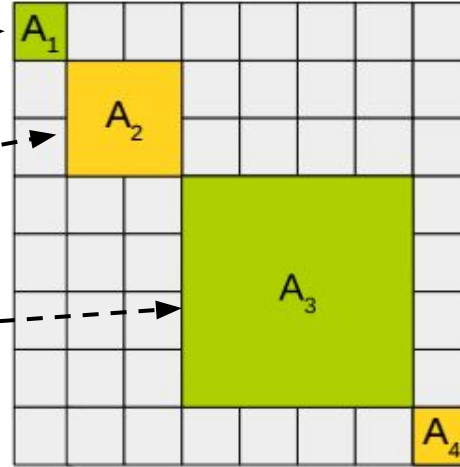
Overview

- *Point of departure*: focus on high order, spectral HP FEM, with local assembly
 - **block diagonal SpMV** (this work) vs **generic SpMV** (prior work)


Block SpMV



Block Diagonal Sparse Matrix



Dense Vectors



- Each dense block corresponds to one element
- **Larger dense blocks \Rightarrow More structured computation**

SpMV

A

x

X

=

b



Dense Matrix Blocks



Zero Elements

Overview

- *Point of departure*: focus on high order, spectral HP FEM, with local assembly
 - **block diagonal SpMV** (this work) vs **generic SpMV** (prior work)

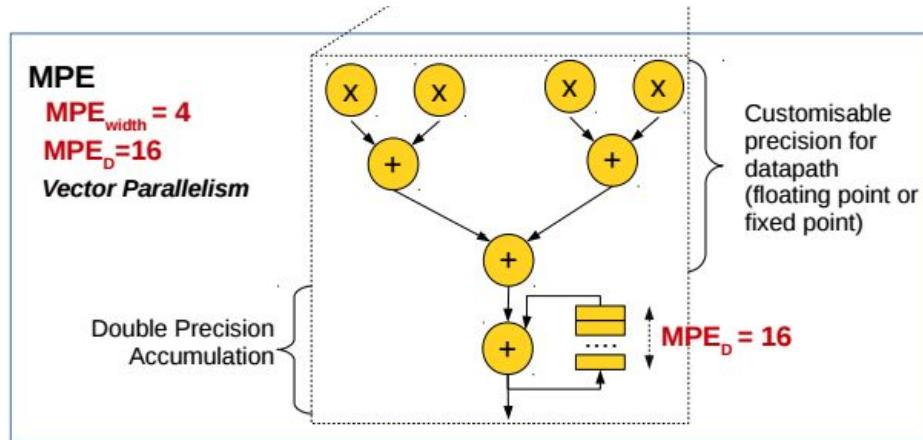
Overview

- *Point of departure*: focus on high order, spectral HP FEM, with local assembly
 - **block diagonal SpMV** (this work) vs **generic SpMV** (prior work)
- *Contributions*:
 - Optimised architecture and implementation for block diagonal SpMV
 - Resource constrained performance model for the proposed architecture
 - Automated method to customise the architecture based on mesh parameters

Overview



- *Point of departure*: focus on high order, spectral HP FEM, with local assembly
 - **block diagonal SpMV** (this work) vs **generic SpMV** (prior work)
- *Contributions*:
 - Optimised architecture and implementation for block diagonal SpMV
 - Resource constrained performance model for the proposed architecture
 - Automated method to customise the architecture based on mesh parameters
- *Result*: a custom, mesh-specific architecture generator
 - Maximise throughput/area \Rightarrow fit larger meshes & improve performance

Architecture



- Each MPE has
 - Independent memory channel
 - Customisable precision datapath
 - Variable depth FIFO - support block variations at runtime


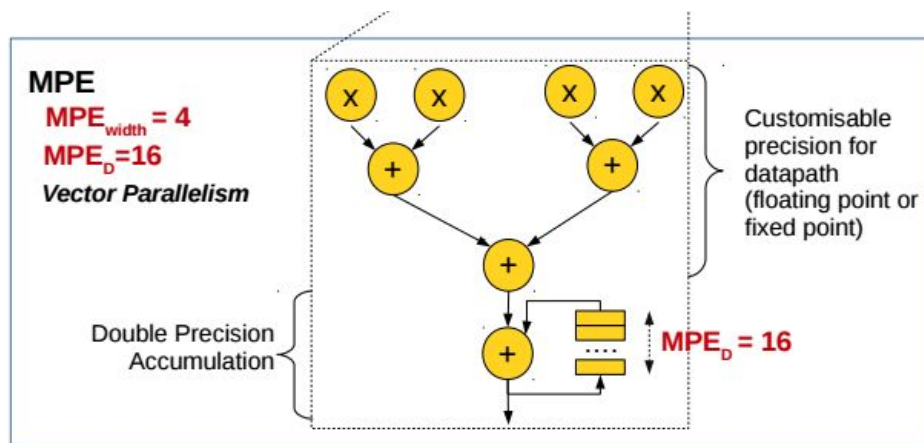
Architecture



- Each MPE has
 - Independent memory channel
 - Customisable precision datapath
 - Variable depth FIFO - support block variations at runtime

- Design:
 - Parametric: NMPEs, MPEwidth
 - Task vs Data Parallelism tradeoff
 - \Rightarrow Mesh specific optimal config.

Architecture



- Each MPE has
 - Independent memory channel
 - Customisable precision datapath
 - Variable depth FIFO - support block variations at runtime
- Design:
 - Parametric: NMPEs, MPEwidth
 - Task vs Data Parallelism tradeoff
 - \Rightarrow Mesh specific optimal config.
- Block SpMV advantages:
 - \Rightarrow Simplified control (format decoding)
 - \Rightarrow Reduced metadata
 - \Rightarrow Simplified reduction circuit

Parameter Extraction

- Assume matrix is block diagonal
- Extract mesh parameters: size & number of blocks for each element
- In DSE: find and synthesise optimal architectures
- At runtime: select the appropriate architecture

Algorithm 2 Extract BS_i, NB_i for the sparse matrix A

```
1: function EXTRACTBLOCKMATRIXPARAMS(A)
2:   SORTENTRIES(A)           ▷ Sort by row and column index
3:    $occ \leftarrow \{\}$        ▷ Dictionary of  $BS_i$  and  $NB_i$ 
4:   for  $i \in 0 \dots A.nrows$  do
5:      $(first, last) \leftarrow \text{ROWSPAN}(i)$ 
6:      $bs \leftarrow last - first$        ▷ Assume new block size
7:      $startPosition \leftarrow i$      ▷ Upper left corner of block
8:     while  $i - startPosition < bs$  do
9:       if  $last > startPosition + bs$  then
10:        ERROR! Not a block diagonal matrix
11:       else if  $first < startPosition$  then
12:        ERROR! Not a block diagonal matrix
13:       end if
14:        $i \leftarrow i + 1$ 
15:        $(first, last) \leftarrow \text{ROWSPAN}(i)$ 
16:     end while
17:      $occ[bs] \leftarrow occ[bs] + 1$    ▷ update block count
18:   end for
19:   return  $occ$            ▷ return the frequency map of block sizes
20: end function
```

Performance Model

- Mesh parameters \Rightarrow optimal architecture parameters

- Performance:
$$N_{Cycles} = \sum_i CyclesPerBlock(B_i) \times Tasks(B_i) \quad (3)$$

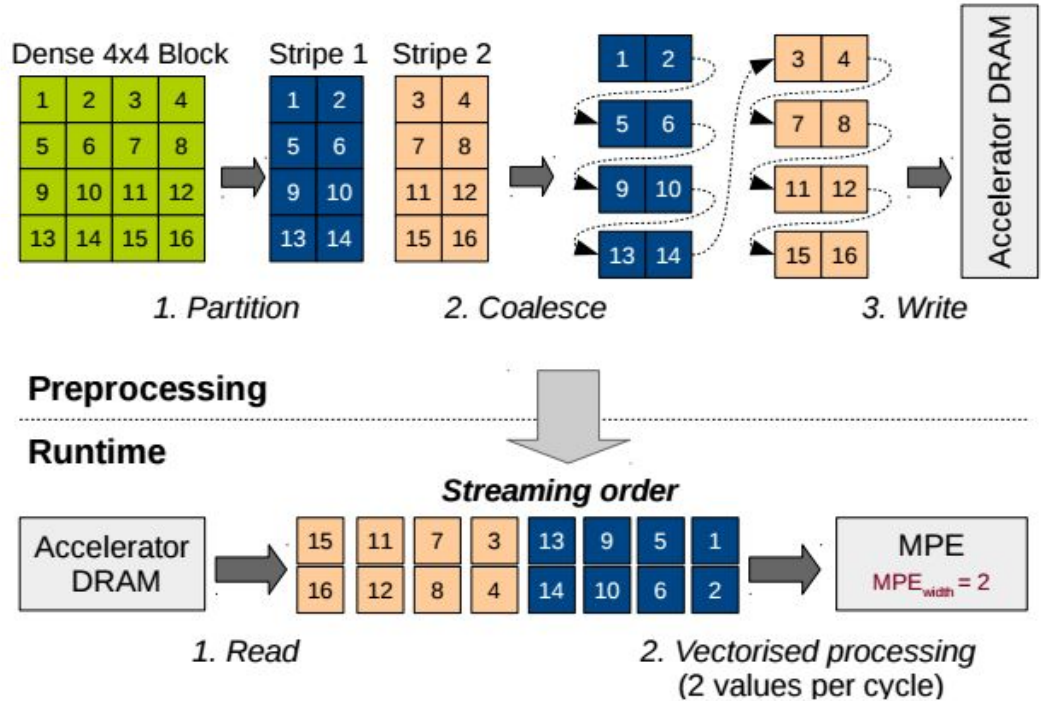
$$= \sum_i \left(\left\lceil \frac{BS_i}{MPE_{width}} \right\rceil \times BS_i \right) \times \left\lceil \frac{NB_i}{N_{MPE}} \right\rceil \quad (4)$$

- Resource usage:
$$R = R_{Static} + N_{MPE} \times (MPE_{width} \times (R_+ + R_X) + R_B + R_{DRAM}) \quad (5)$$

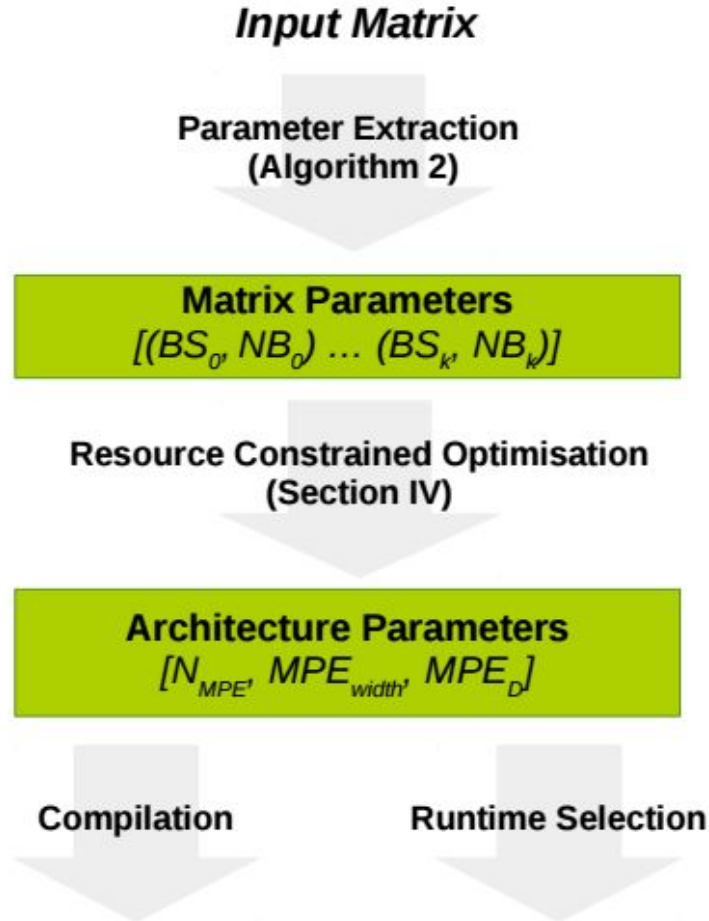
- Functional, hardware constraints \Rightarrow See paper for details

Runtime

- Software layer - can be integrated in existing FEM software packages
- Reorder to **enforce linear access pattern in DRAM**
 - Maximise throughput
 - Minimise control logic

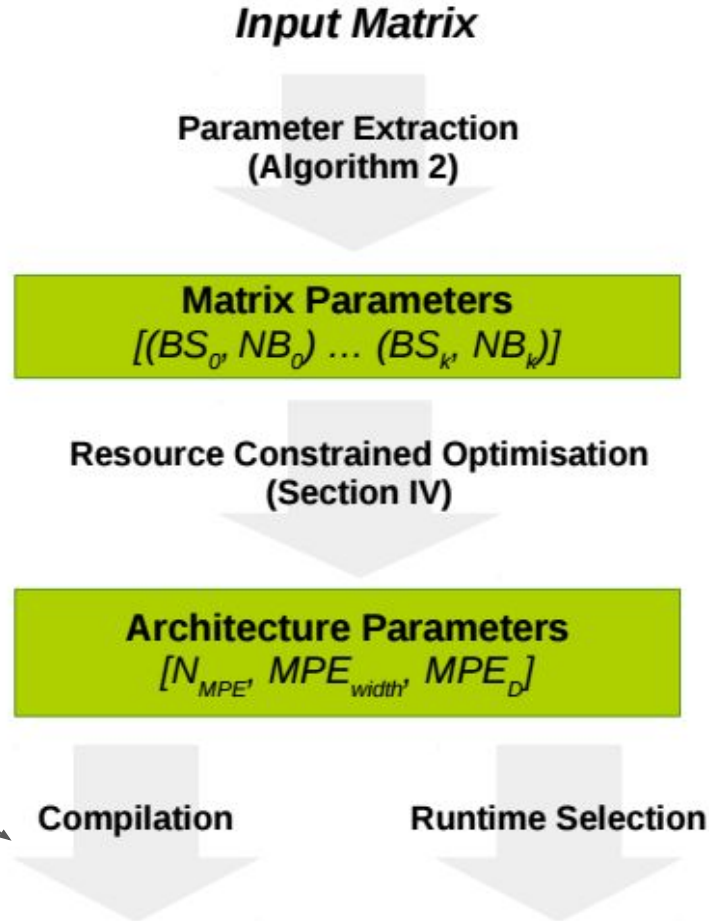


Putting it Together



Putting it Together

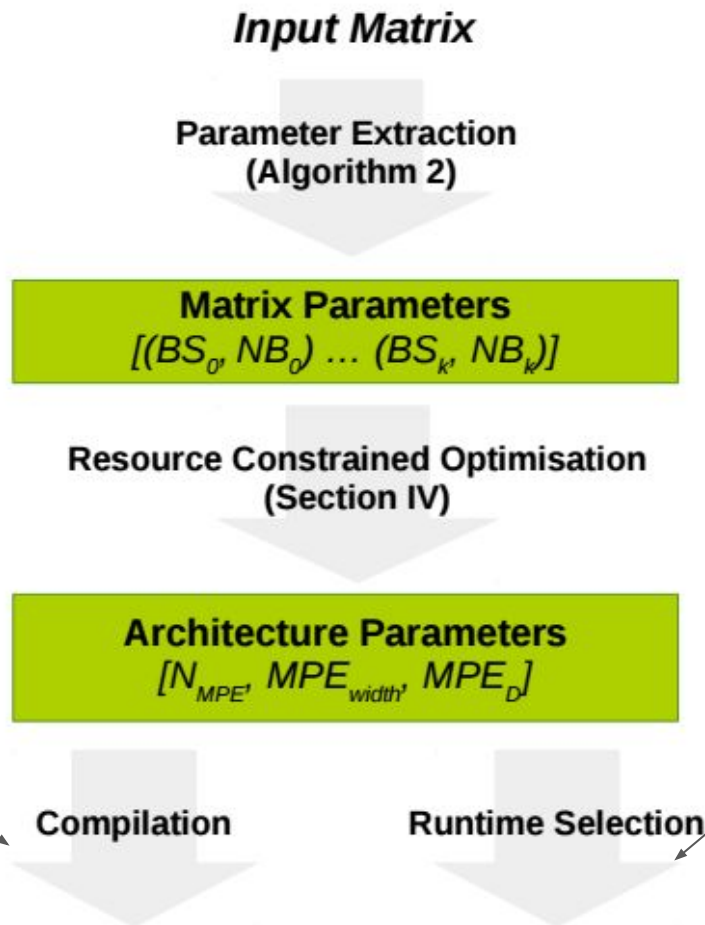
Offline tuning: build a repository of customised architectures from a set of mesh instances



Putting it Together

Offline tuning: build a repository of customised architectures from a set of mesh instances

Runtime: select the optimal architecture for an input mesh instance



Evaluation

Evaluation

- **Implementation**

- Design: MaxCompiler + MaxJ dataflow language
- FPGA Server: Maxeler Max 4 Maia (Stratix VSG, 48GB DRAM, per board)
- Software: C++14, G++ 5.2
- CPU Server: Dual Intel Xeon E5-2640, 64GB DRAM, Infiniband QSFP
- Place and route with Altera Quartus 14.1
- Available as extension to the CASK framework [Grigoras et al, FPGA 16]:
 - <http://caskorg.github.io/cask/>

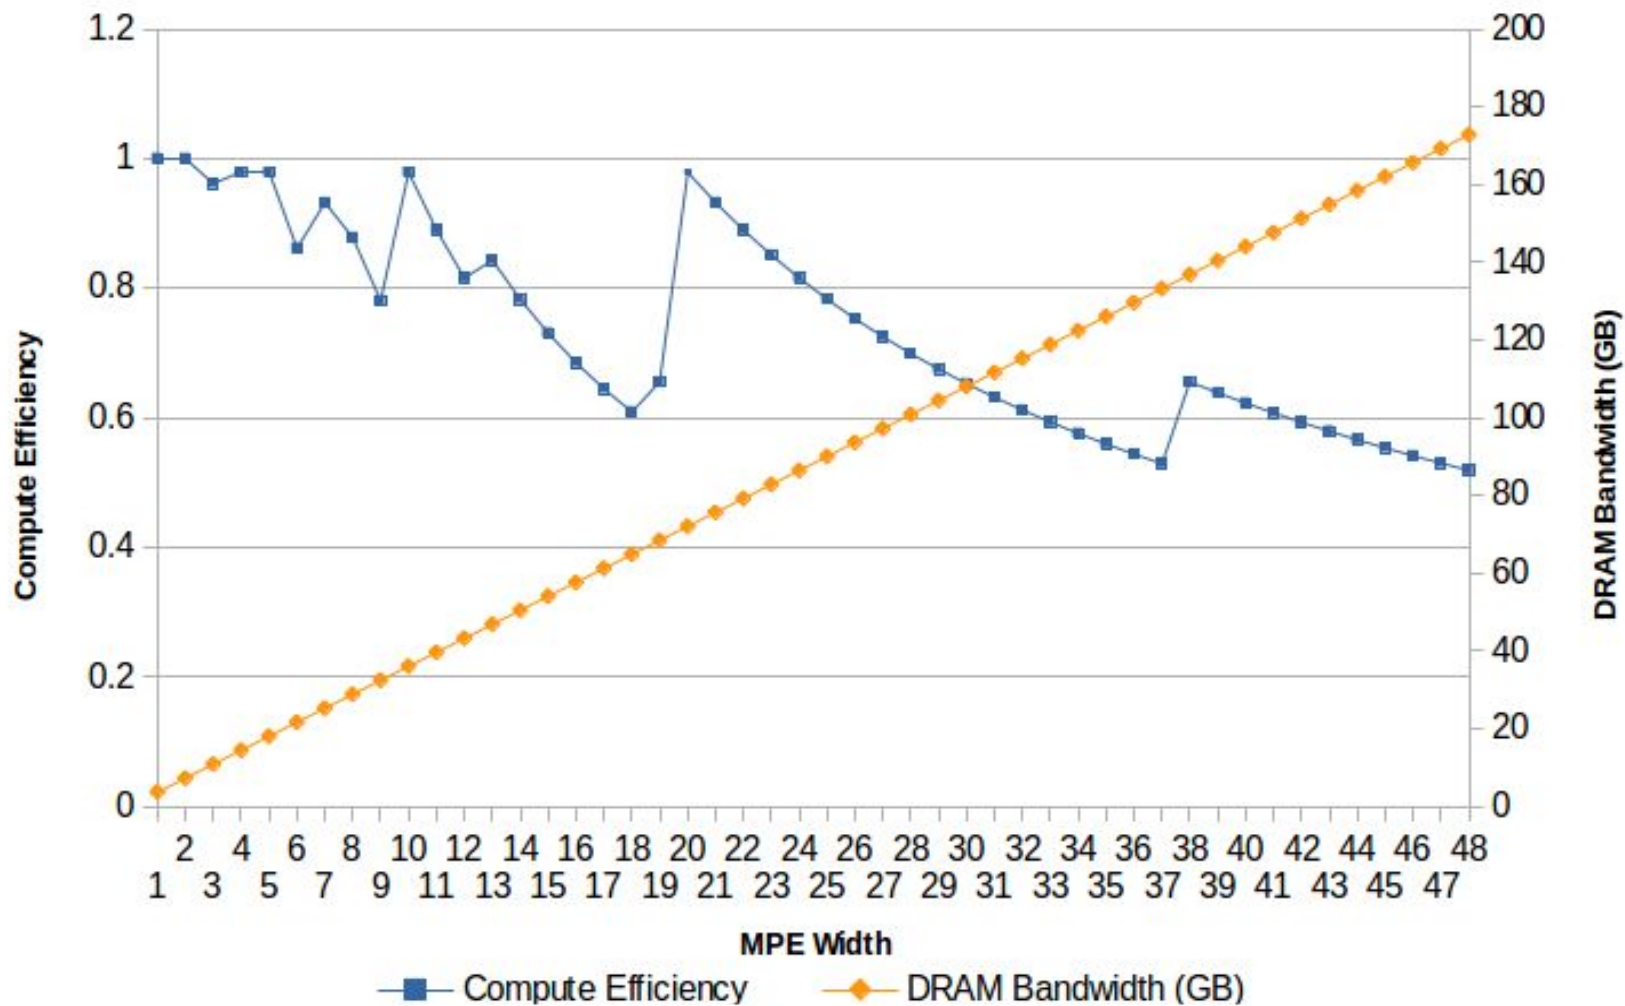
- **Reference software** - Nektar++ FEM Package, <http://www.nektar.info/>


- **Reference hardware**


- [Burovskiy et al, FPL 15], Nektar++ Accelerated FEM


Experiments

1. *What is the benefit of tuning architecture based on mesh properties?*
 - a. Fixed mesh (NACA 1L, [Burovskiy et al, FPL 2015]) - optimal architecture







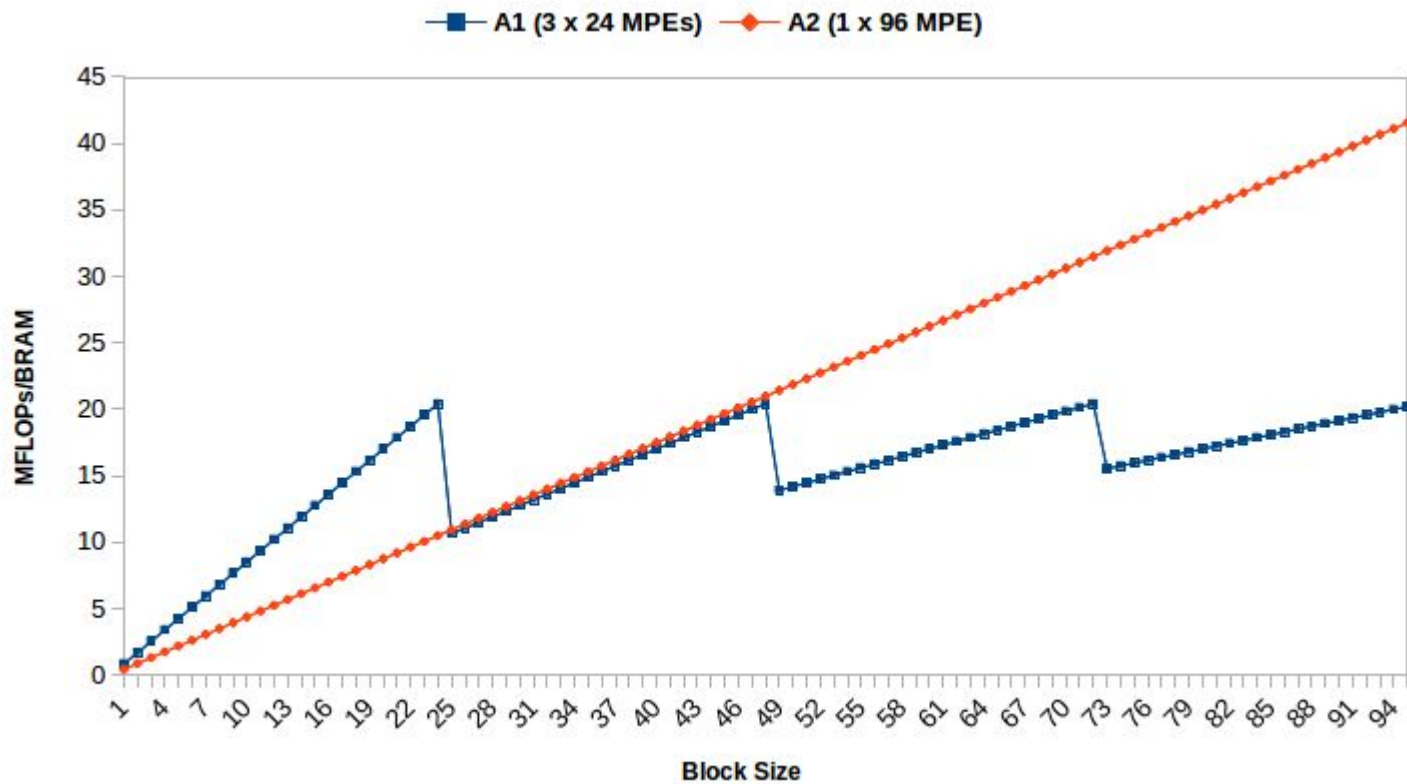


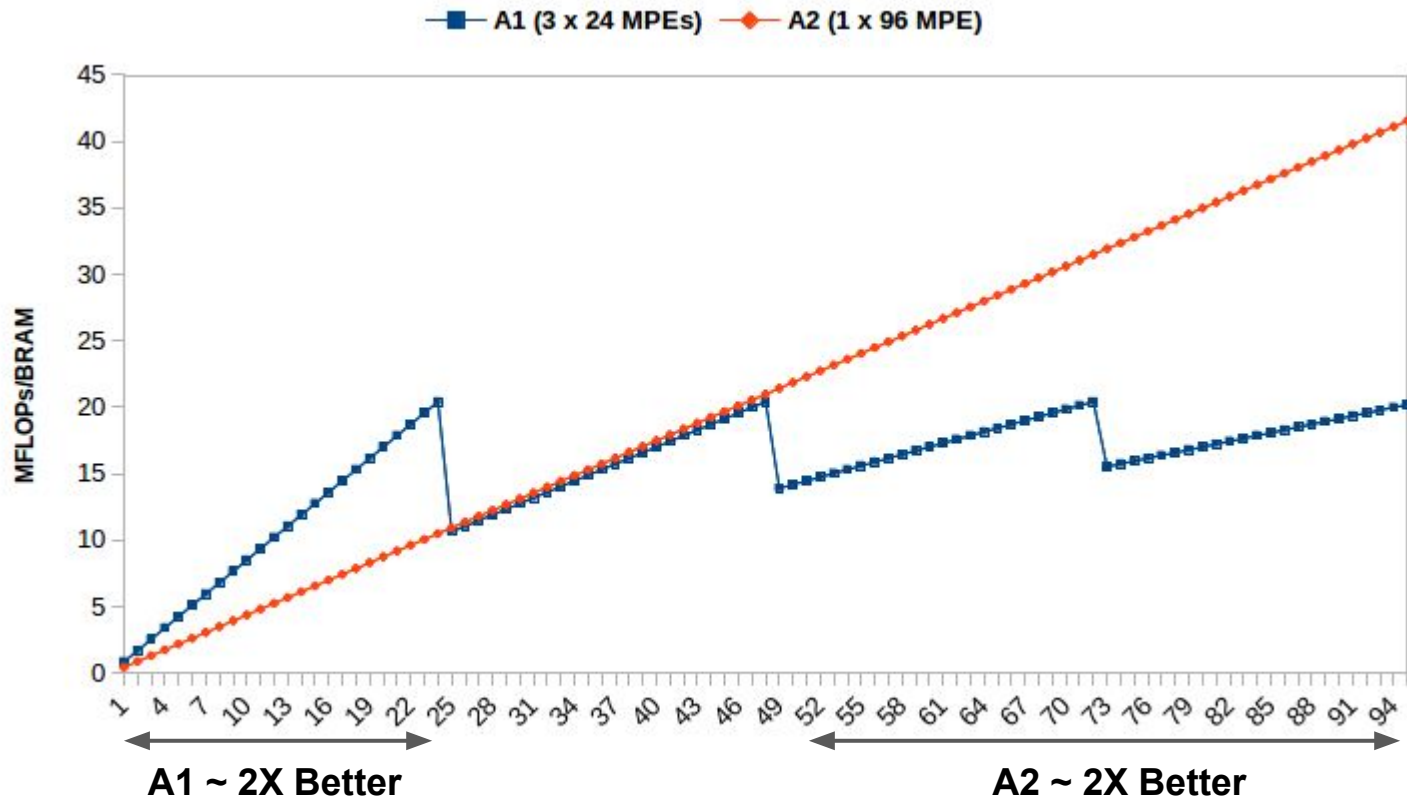
Experiments

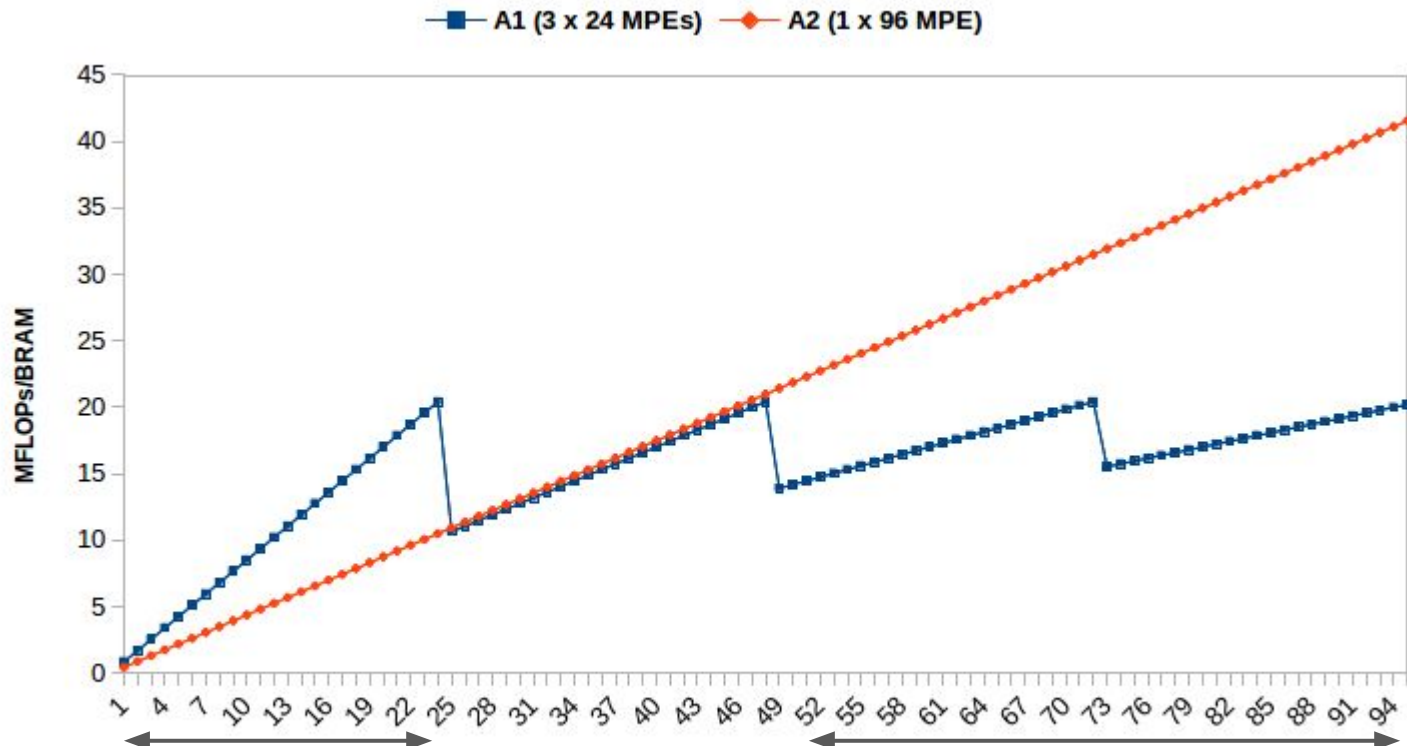
1. *What is the benefit of tuning architecture based on mesh properties?*
 - a. Fixed mesh (NACA 1L, [Burovskiy et al, FPL 2015]) - optimal architecture
 - ⇒ find architecture with good efficiency
 - ⇒ improve performance s.t. resource usage

Experiments

1. *What is the benefit of tuning architecture based on mesh properties?*
 - a. Fixed mesh (NACA 1L, [Burovskiy et al, FPL 2015]) - optimal architecture
 - ⇒ find architecture with good efficiency
 - ⇒ improve performance s.t. resource usage
 - b. Fixed architecture, variable mesh, data parallel vs task parallel








A1 ~ 2X Better

+Mem. Chan., -Vector lanes, 1058 BRAMs

A2 ~ 2X Better

- Mem. Chan., + Vector lanes, 686 BRAMs




A1 ~ 2X Better

+Mem. Chan., -Vector lanes, 1058 BRAMs
 ⇒ **Good for small blocks**

A2 ~ 2X Better

- Mem. Chan., + Vector lanes, 686 BRAMs
 ⇒ **Good for large blocks**



A1 ~ 2X Better

+Mem. Chan., -Vector lanes, 1058 BRAMs
 ⇒ **Good for small blocks**

A2 ~ 2X Better


- Mem. Chan., + Vector lanes, 686 BRAMs
 ⇒ **Good for large blocks**

Experiments

1. *What is the benefit of tuning architecture based on mesh properties?*
 - a. Fixed mesh (NACA 1L, [Burovskiy et al, FPL 2015]) - optimal architecture
 - ⇒ find architecture with good efficiency
 - ⇒ improve performance s.t. resource usage
 - b. Fixed architecture, variable mesh, data parallel vs task parallel
 - ⇒ select optimal MPE Width and N Mpe for given mesh
 - ⇒ improve performance, reduce resource usage

Experiments

1. *What is the benefit of tuning architecture based on mesh properties?*
 - a. Fixed mesh (NACA 1L, [Burovskiy et al, FPL 2015]) - optimal architecture
 - ⇒ find architecture with good efficiency
 - ⇒ improve performance s.t. resource usage
 - b. Fixed architecture, variable mesh, data parallel vs task parallel
 - ⇒ select optimal MPE Width and N Mpe for given mesh
 - ⇒ improve performance, reduce resource usage
2. *What is the expected benefit for a full FEM implementation?*
 - a. Baseline, Nektar++ implementation from [Burovskiy et al, FPL 2015]



Experiments

1. *What is the benefit of tuning architecture based on mesh properties?*
 - a. Fixed mesh (NACA 1L, [Burovskiy et al, FPL 2015]) - optimal architecture
 - ⇒ find architecture with good efficiency
 - ⇒ improve performance s.t. resource usage
 - b. Fixed architecture, variable mesh, data parallel vs task parallel
 - ⇒ select optimal MPE Width and N Mpe for given mesh
 - ⇒ improve performance, reduce resource usage
2. *What is the expected benefit for a full FEM implementation?*
 - a. Baseline, Nektar++ implementation from [Burovskiy et al, FPL 2015]
 - ⇒ enabling larger problem sizes, not supported by previous work.
 - ⇒ enable a good proportion of the projected speedup (3X over CPU)

Conclusion

1. *Proposed:*
 - a. FPGA architecture optimised for variable-size block diagonal SpMV
 - b. method to extract customisation parameters directly from mesh instance
 - c. software to integrate with existing FEM package, Nektar++

2. *Achieved:*
 - a. Fit larger FEM problems on a single FPGA
 - b. 3X speedup over optimised CPU

3. *Future:* exploration of additional trade-offs & parameters

That's it folks! Thank you!