

Parallel finite element analysis on the Intel Xeon Phi

Lee Margetts, David Arregui-Mena
Terry Hewitt, Luke Mason

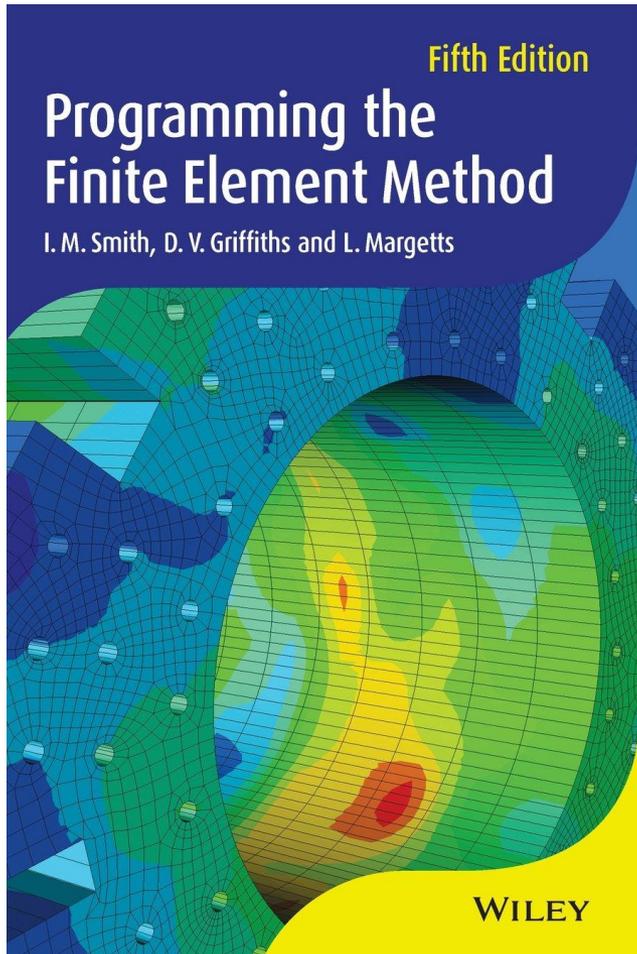
Emerging Technology Conference EMiT 2016, 2-3 June, Barcelona, Spain

Overview

- The ParaFEM project
- Basic characteristics of the Xeon Phi
- The hotspot in the finite element solver
- Preliminary performance results
- Path to improved performance
- Summary

ParaFEM

Open Source Parallel FEA



- ParaFEM
- Open source library + ~70 mini-apps
- ~64,000 cores
- >1 billion degrees of freedom
- Used for teaching and research
- 930+ registered on website
- ~1100 citations of text book

<http://parafem.org.uk>

<http://www.amazon.com/Programming-Finite-Element-Method-Smith/dp/1119973341>

Problem-driven Development



Research Question or Engineering Problem

- Is HPC required to solve the problem?
- Does the existing software need customization?



Implementation

- Test using analytical problem
- Compare with ISV software. If differences, fix or explain

**BETTER
SOFTWARE
BETTER
RESEARCH**

Sustainability

- Source code committed to repository
- Publish modifications

simpleware 

 **CCFE**
CULHAM CENTRE FOR
FUSION ENERGY

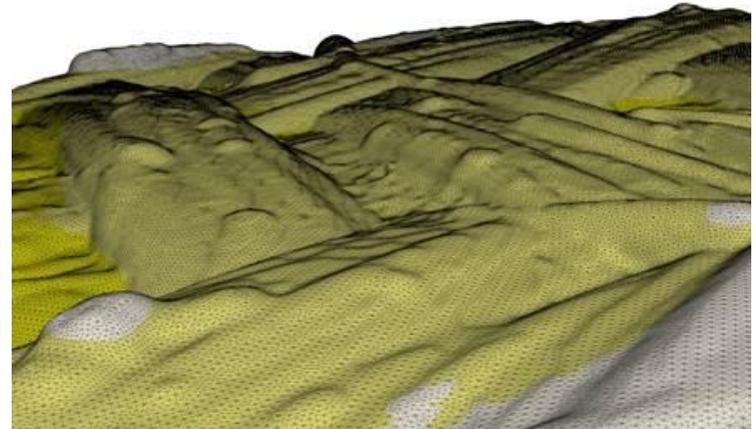
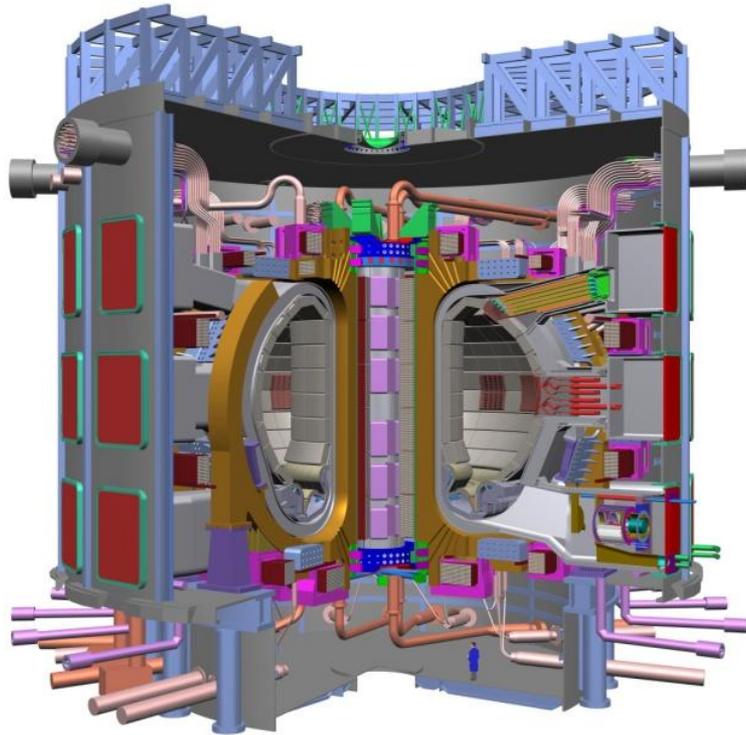
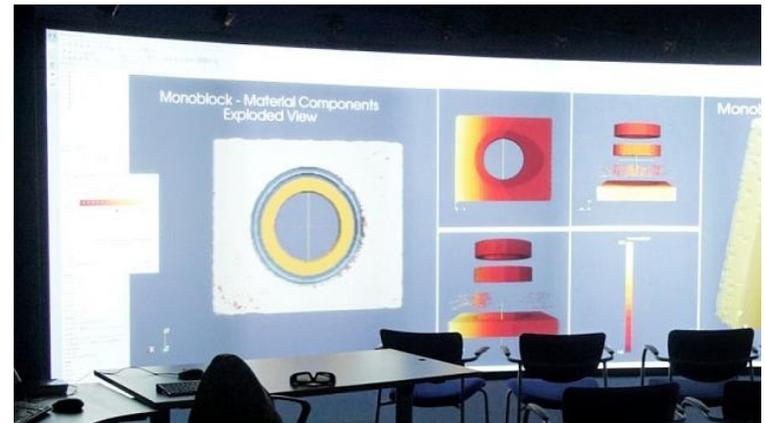


Image-based FEA



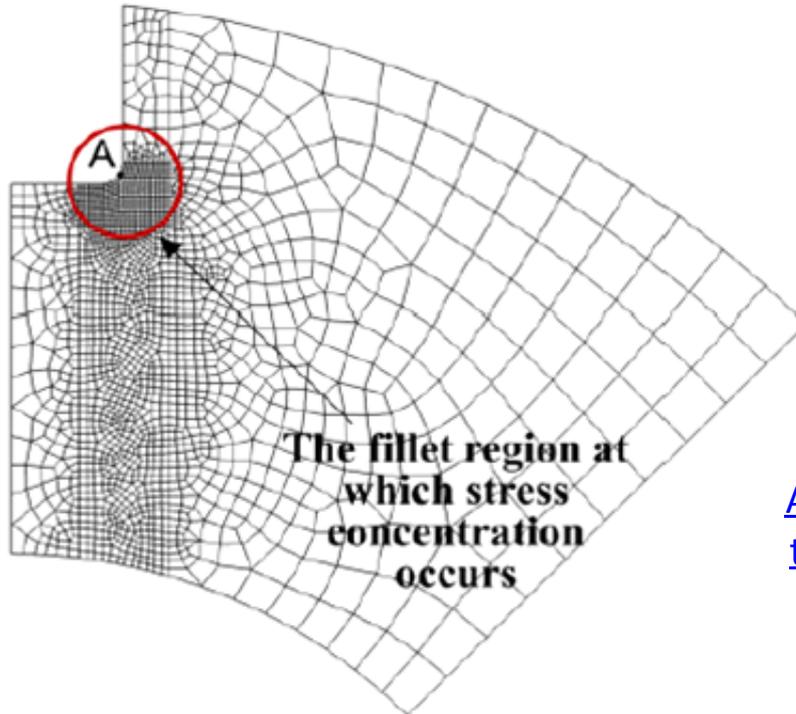
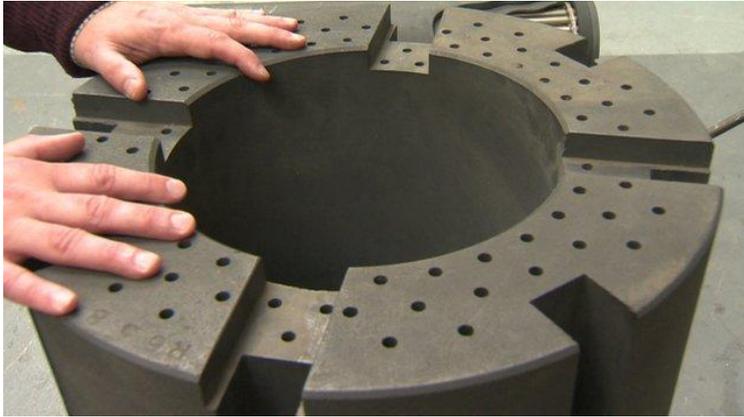
[Evans et al. "Transient thermal finite element analysis of CFC-Cu ITER monoblock using X-ray tomography data", Journal of Fusion Engineering Design, 2015](#)

Random Fields

NATIONAL NUCLEAR
LABORATORY

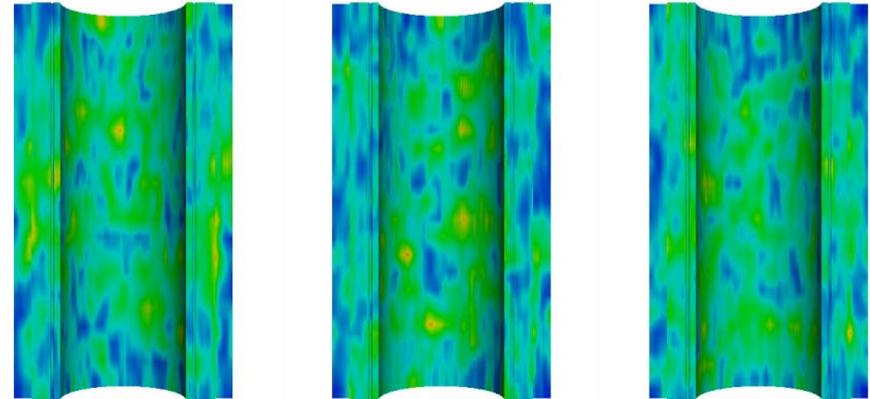


CONACYT

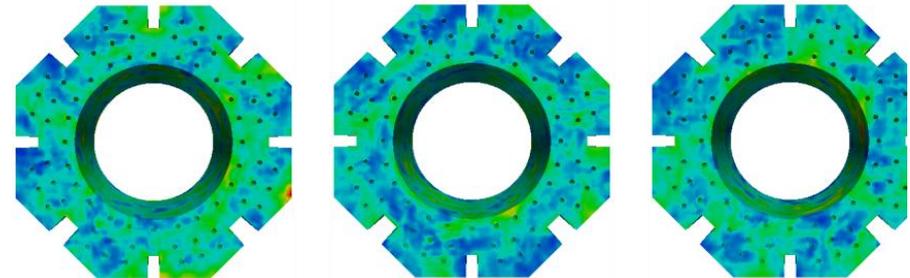


Vertical cross section

Stochastic



Horizontal cross section



[Arregui et al. "Random spatial variability in the coefficient of thermal expansion induces pre-service stresses in computer models of virgin gilsocarbon bricks", Journal of Nuclear Materials, 2015](#)



THE UNIVERSITY
of EDINBURGH

MANCHESTER
1824

The University of Manchester
Aerospace Research Institute



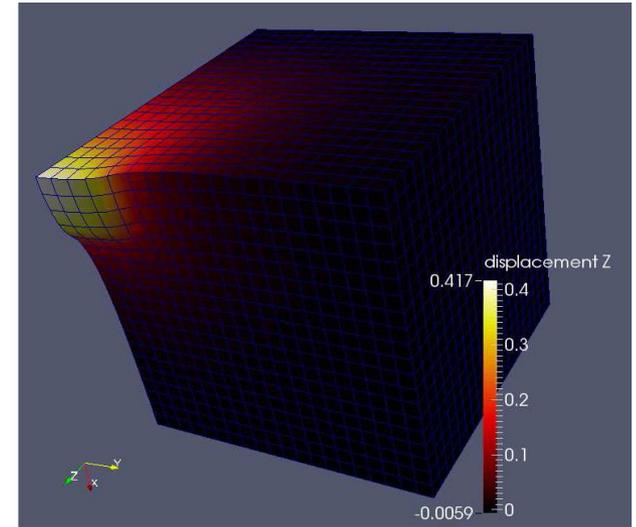
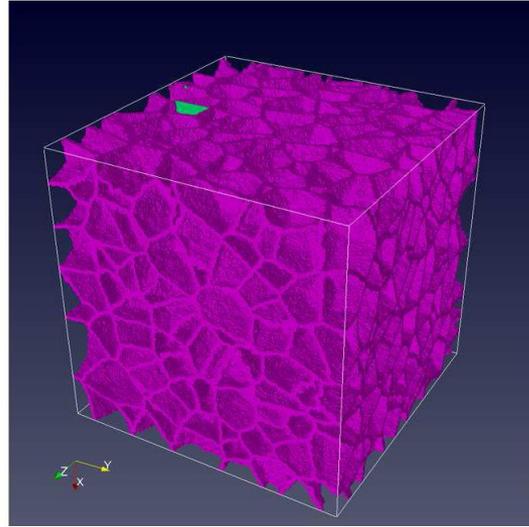
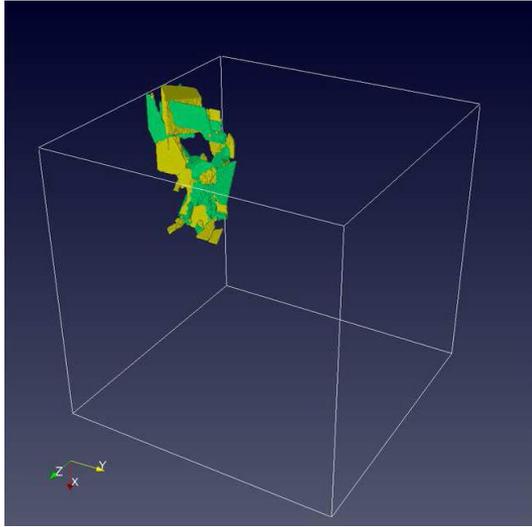
Nonlinear geometry with
plasticity as Abaqus UMATs

Cohesive fracture modelling
of CFC through Abaqus UELs

Parallel capability for Abaqus users

[Levrero et al. "Evaluating the macroscopic yield behaviour of trabecular bone using a nonlinear homogenisation approach", Journal of the Mechanical Behavior of Biomedical Materials, 2016.](#)

Cellular Automata + Finite Elements



Cellular automata at the mesoscale and
continuum mechanics at macroscale

[Shterenlikht A, Margetts L. Three-dimensional cellular automata modelling of cleavage propagation across crystal boundaries in polycrystalline microstructures. Royal Society of London. Proceedings A. Mathematical, Physical and Engineering Sciences. 2015](#)

ALSTOM

EPSRC CASE PhD Studentship

Couple OpenFOAM + ParaFEM

Fluid-structure interaction in wind farms





Random field generator

OpenFOAM FSI

ParaView Viz

Cellular Automata



ParaFEM Open Source Library

Solvers

Maths libraries

MPI

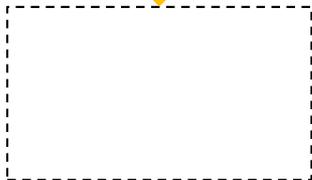
OpenMP

GPUs

Xeon Phi

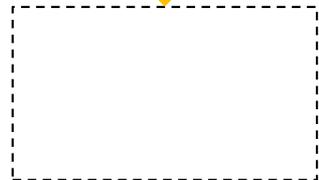
ARM

Cloud



ABAQUS UMAT/UEL

ENSIGHT



Xeon Phi

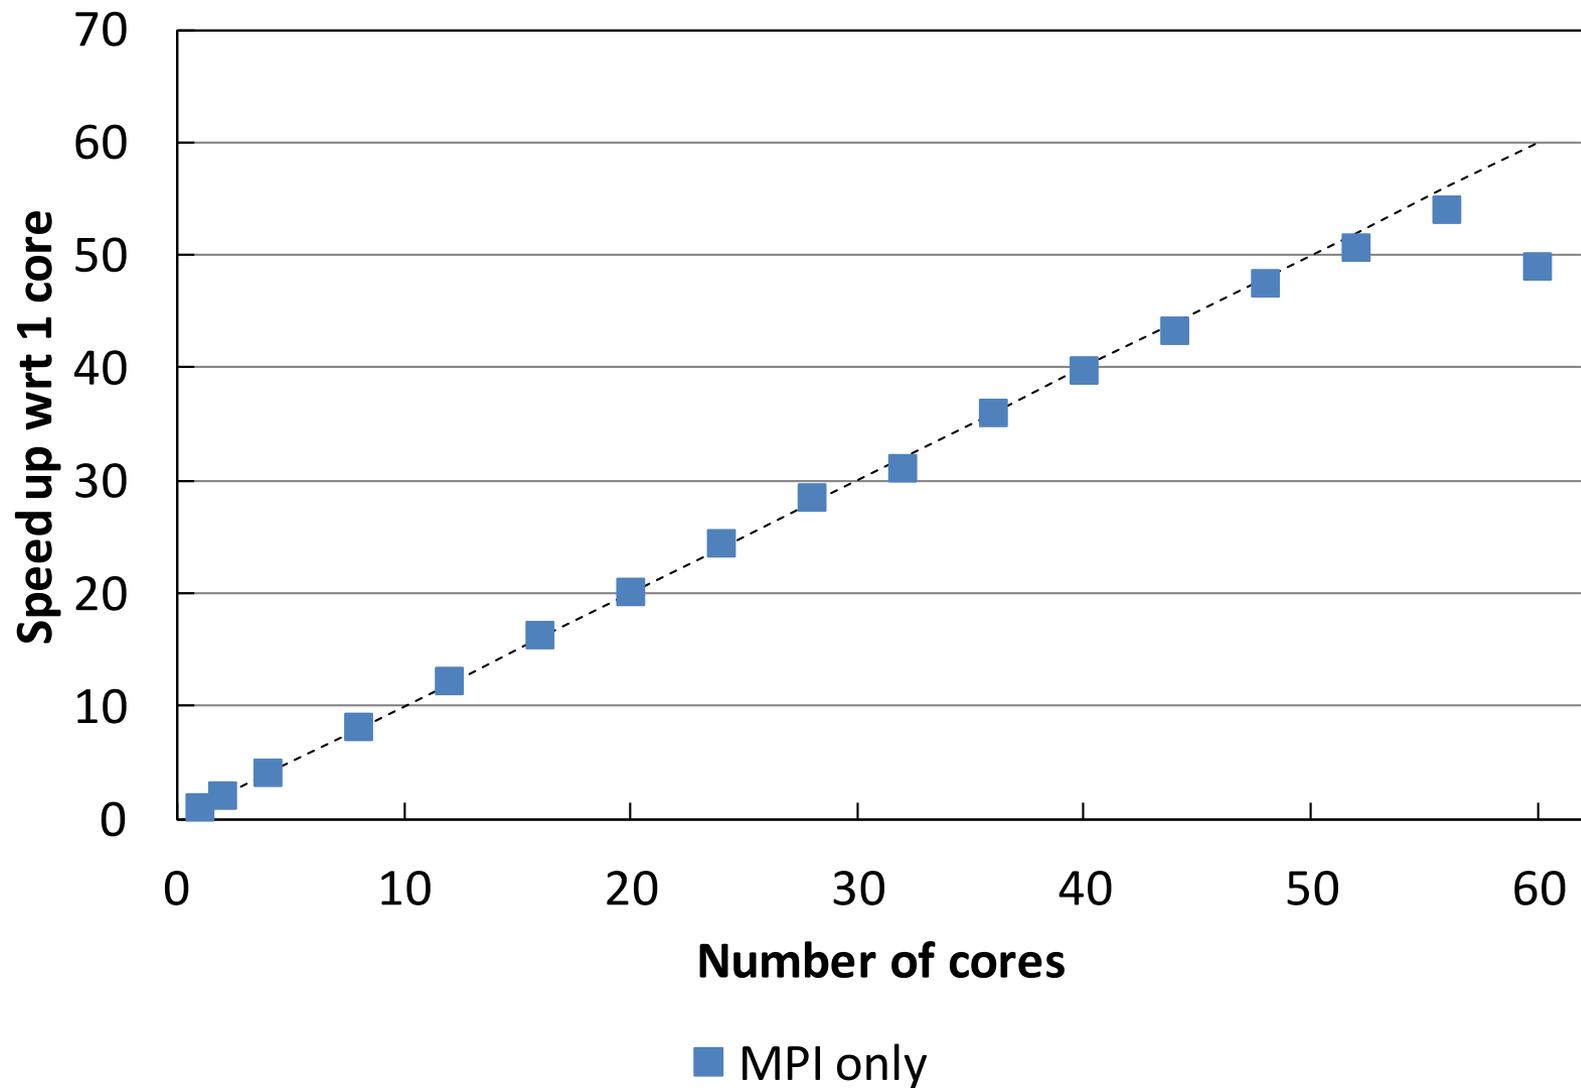
Xeon Phi Basics

- iDataplex at the Hartree Centre, UK
 - 84 nodes with 2 x 12 core Intel Xeon processors
 - Ivy Bridge E5-2697y2 2.7GHz
 - 42 nodes have Intel Xeon Phi 5110P accelerator
 - Each Phi has 60 cores
 - Each core can support 4 threads
 - Theoretical peak is 1 Tflop (double) or 2 Tflops (single)
- The whole code is compiled & run on the Phi
 - Offloading from host to accelerator is not considered here

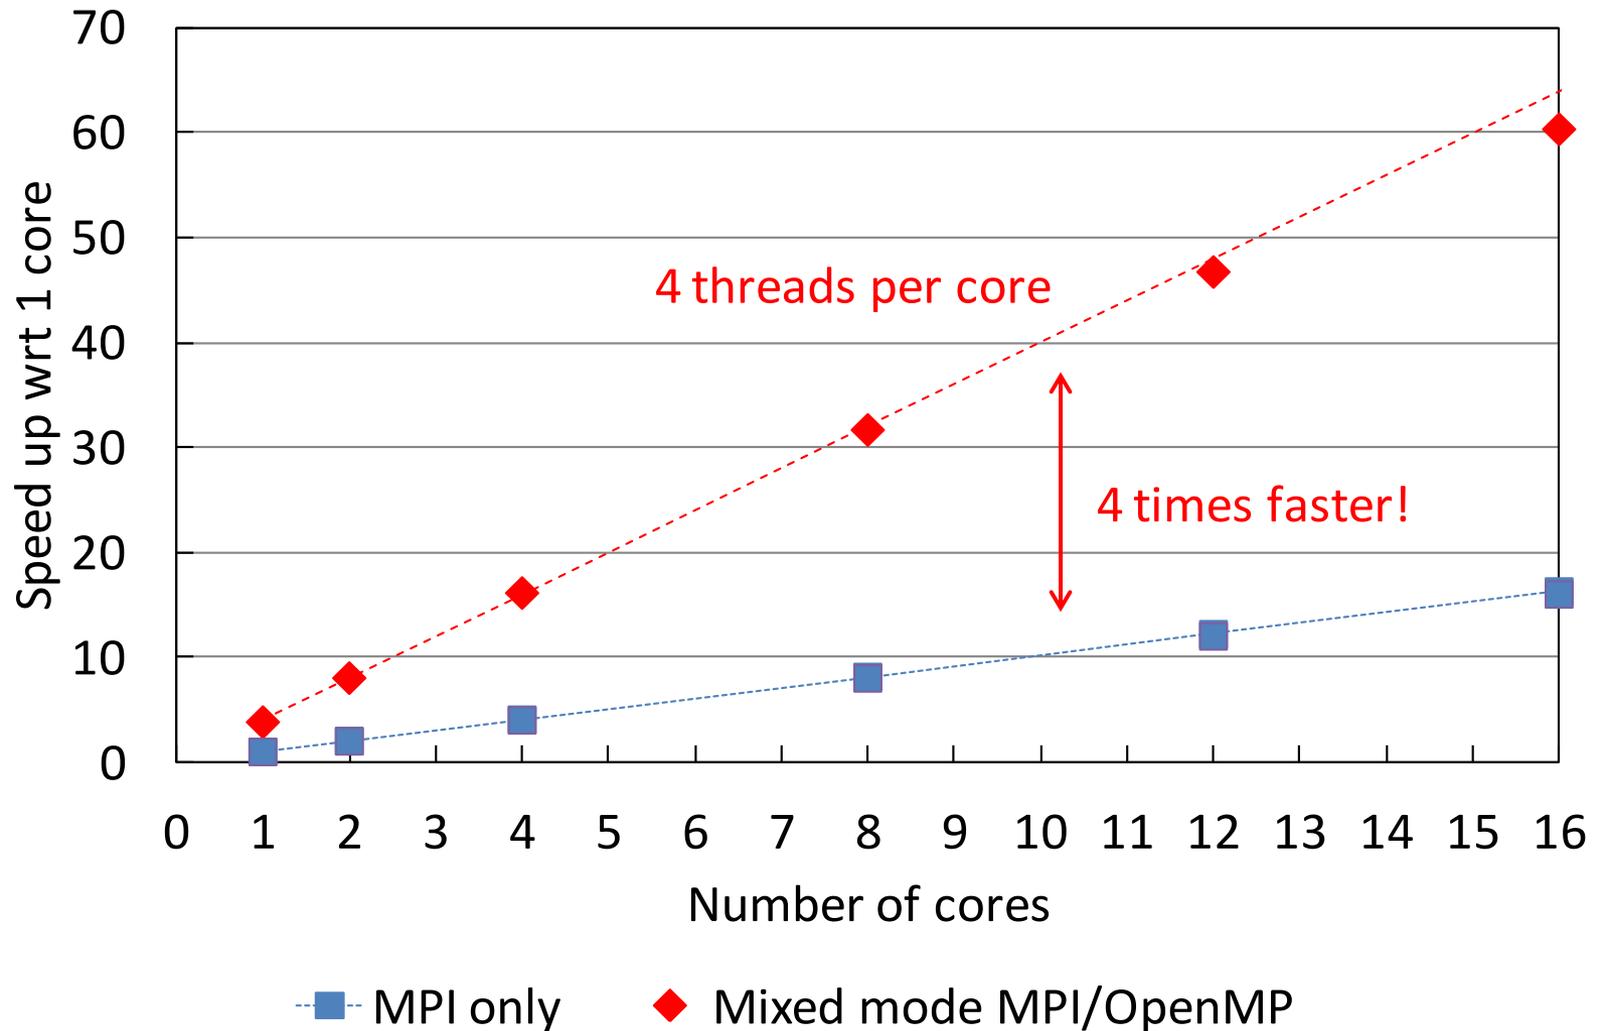
ParaFEM Hotspot

```
iterations: DO
  iters=iters+1; u_pp=zero; pmul_pp=zero; utemp_pp=zero
  CALL gather(p_pp,pmul_pp)
  elements_3: DO iel=1,nels_pp
    CALL dgemv('n',ntot,ntot,1.0_iwp,storkm_pp(:, :, iel), &
      ntot,pmul_pp(:, iel),1,0.0_iwp,utemp_pp(:, iel),1)
  END DO elements_3
  CALL scatter(u_pp,utemp_pp)
  up=DOT_PRODUCT_P(r_pp,d_pp)
  alpha=up/DOT_PRODUCT_P(p_pp,u_pp)
  xnew_pp=x_pp+p_pp*alpha
  r_pp=r_pp-u_pp*alpha
  d_pp=diag_precon_pp*r_pp
  beta=DOT_PRODUCT_P(r_pp,d_pp)/up
  p_pp=d_pp+p_pp*beta
  CALL checon_par(xnew_pp,tol,converged,x_pp)
  IF(converged.OR.iters==limit)EXIT
END DO iterations
```

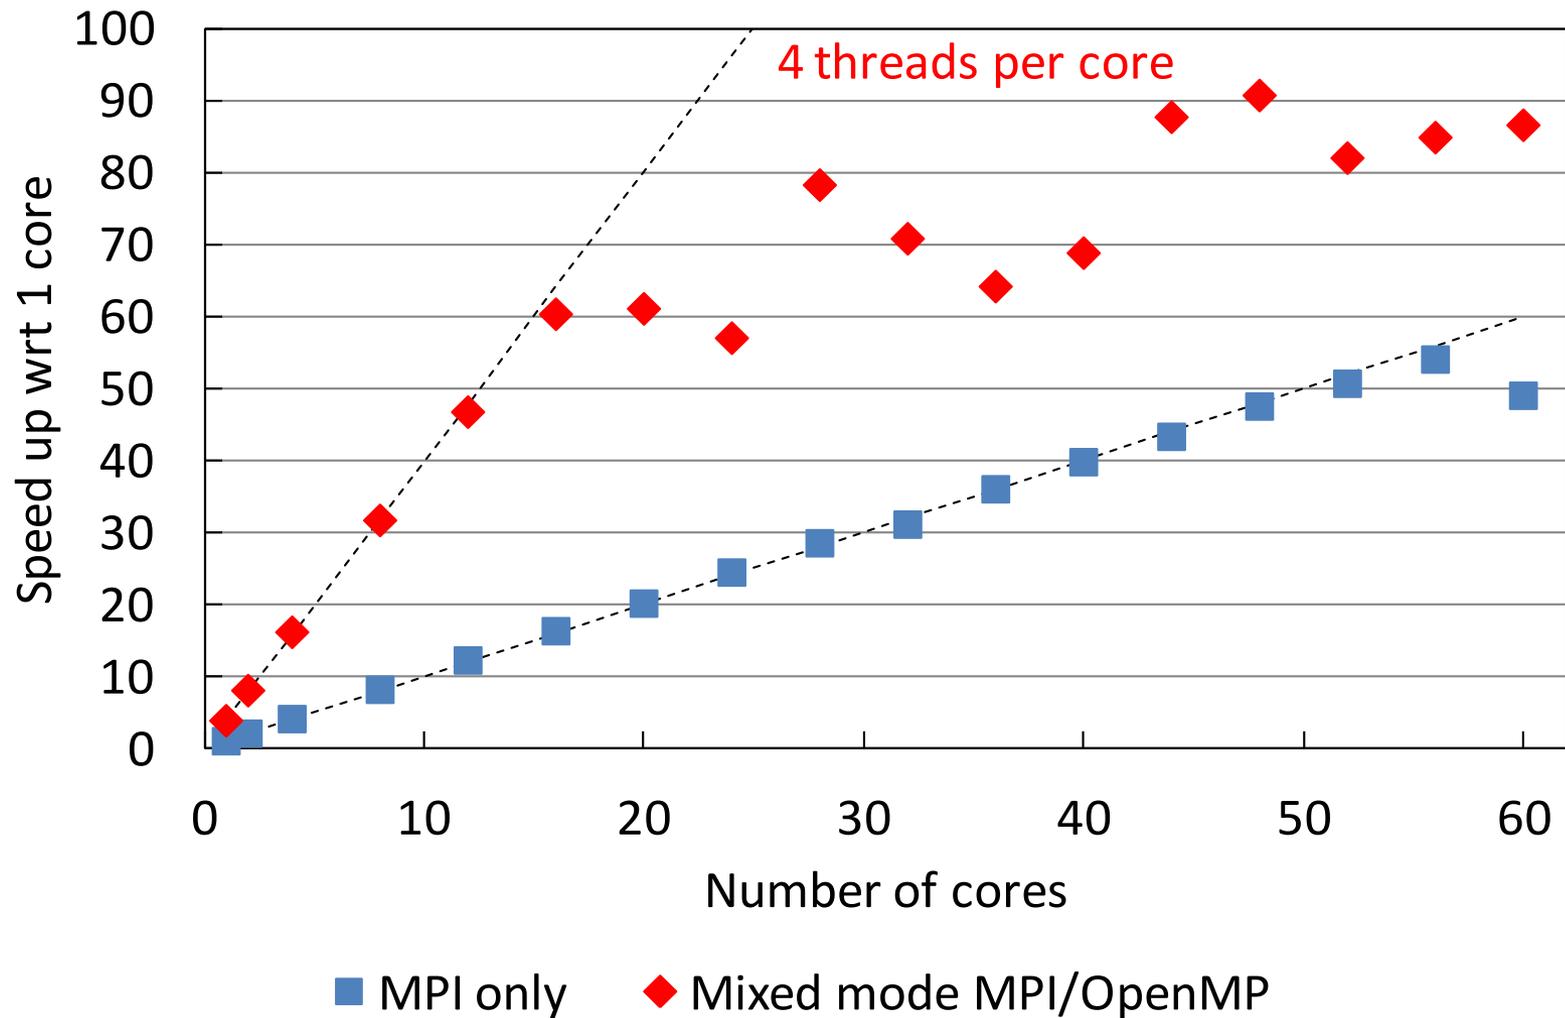
Original MPI Only



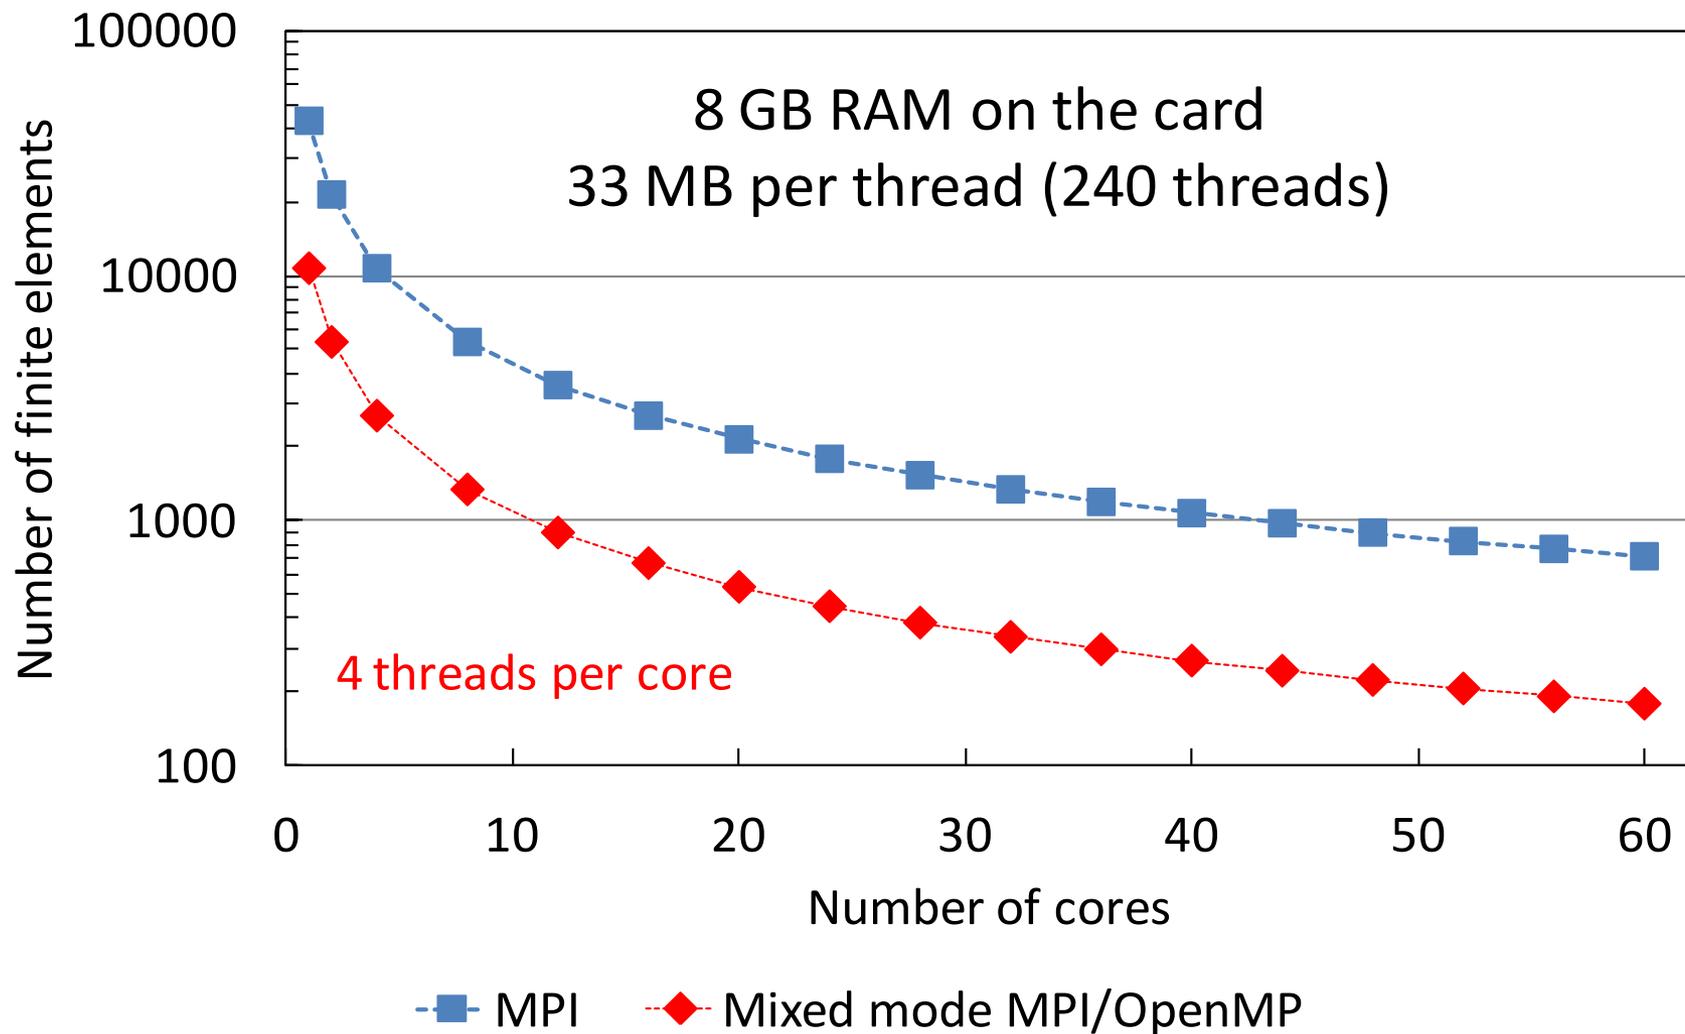
Mixed Mode MPI/OpenMP



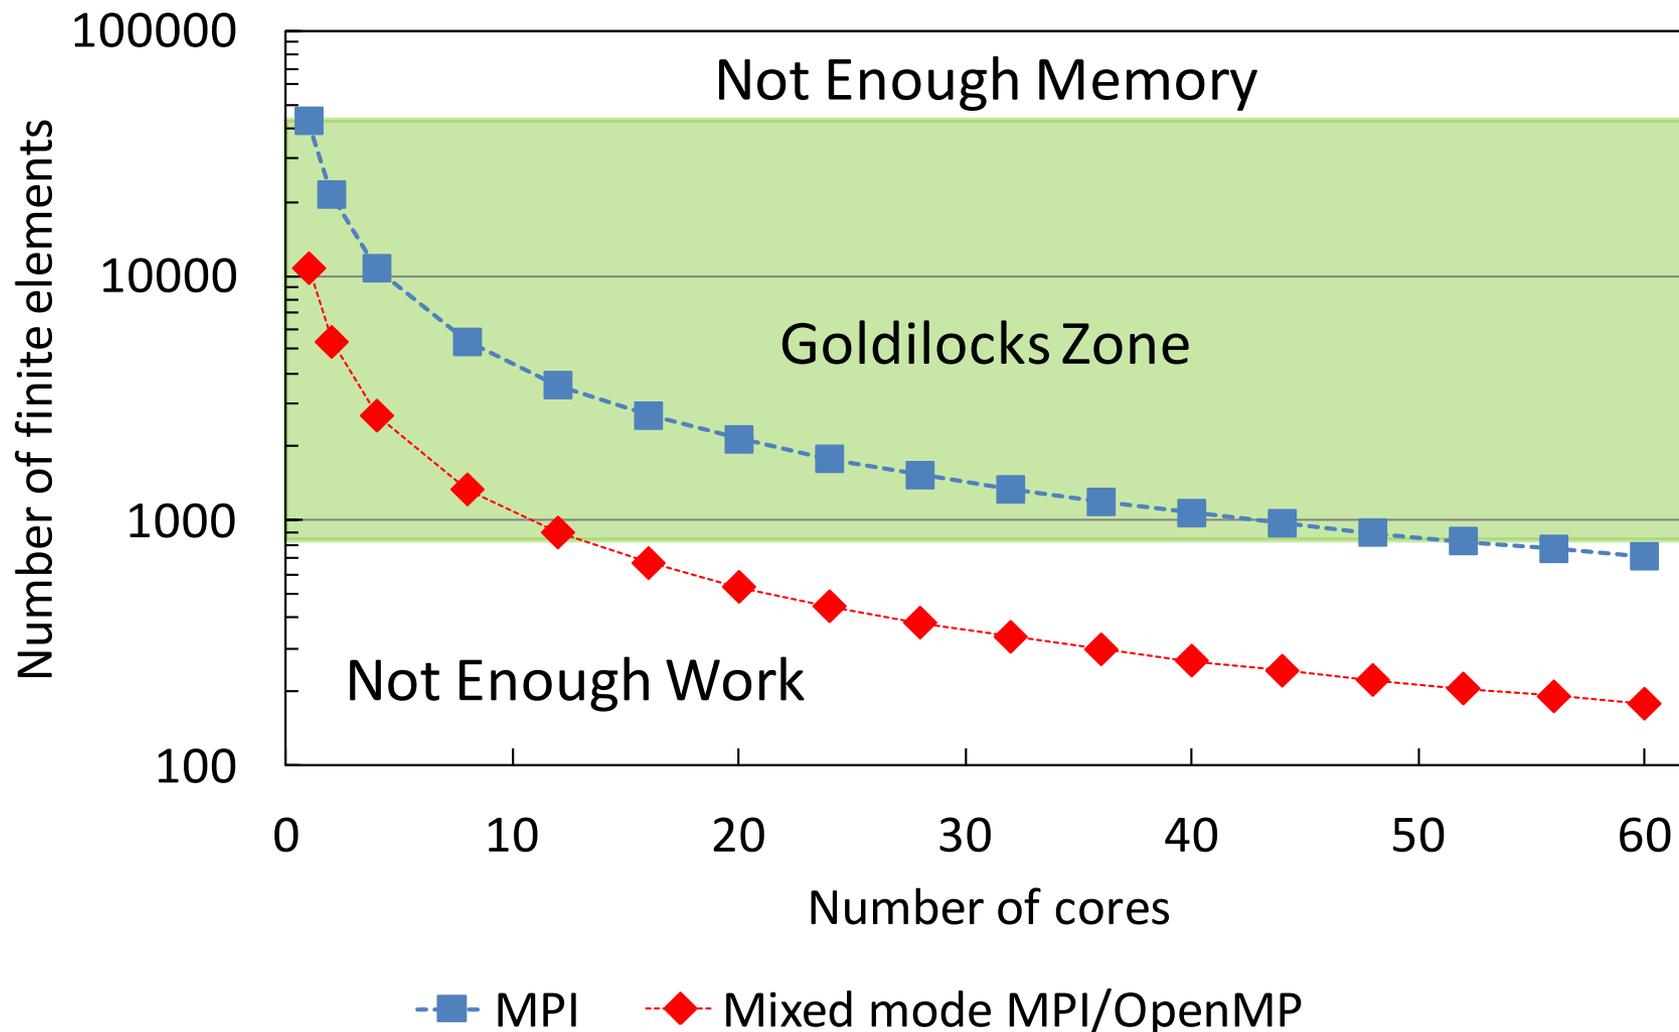
Mixed Mode MPI/OpenMP



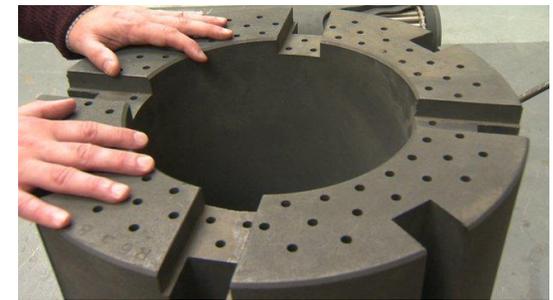
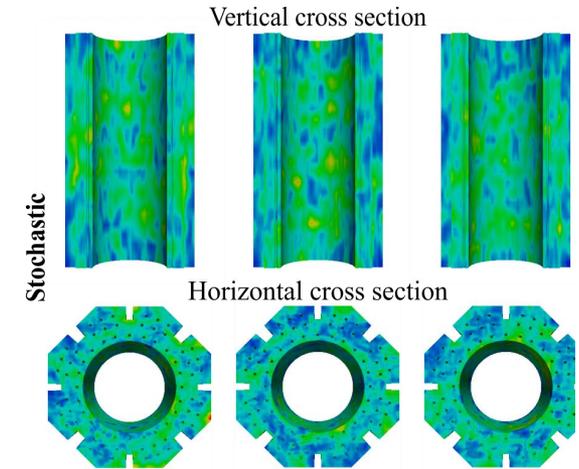
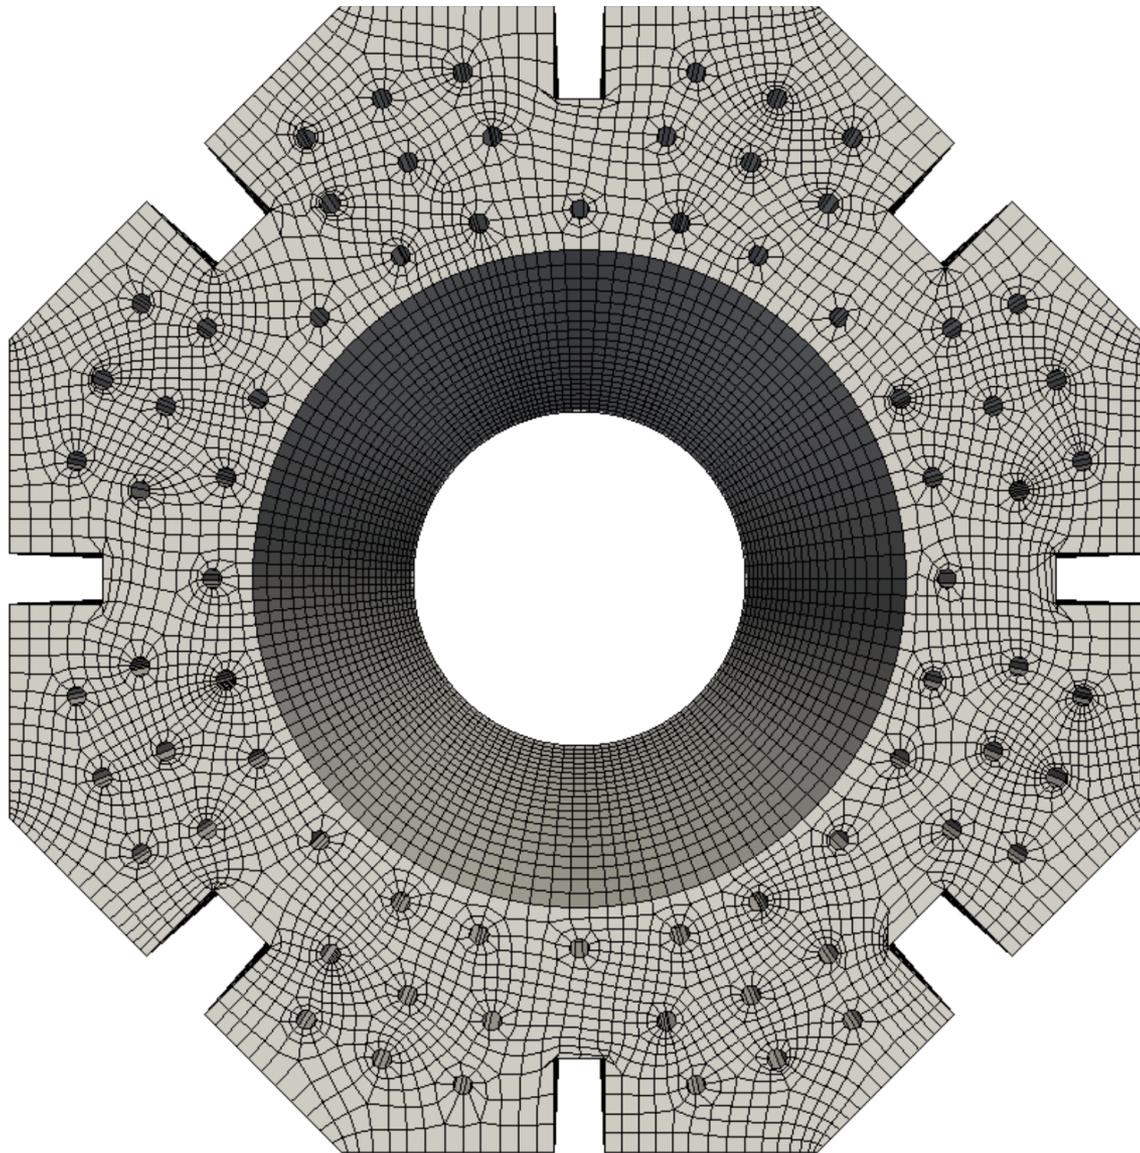
Maximum #Finite Elements Per Thread



Maximum #Finite Elements Per Thread

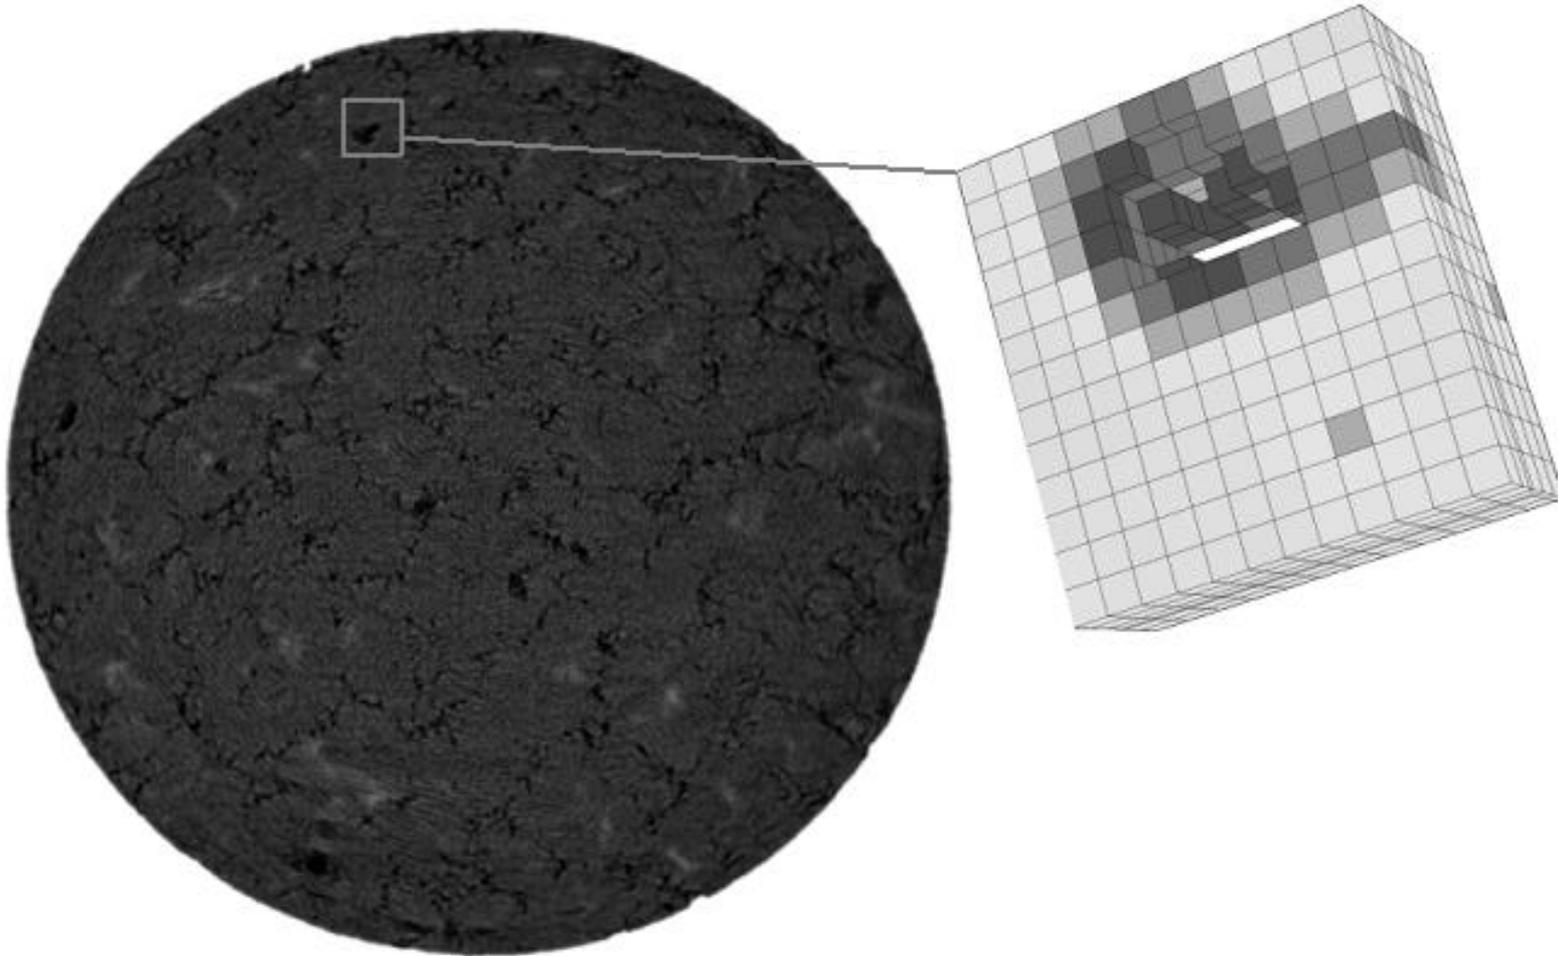


Reducing Storage Requirements



1 matrix for 1 set of
identical elements

Voxelised Meshes

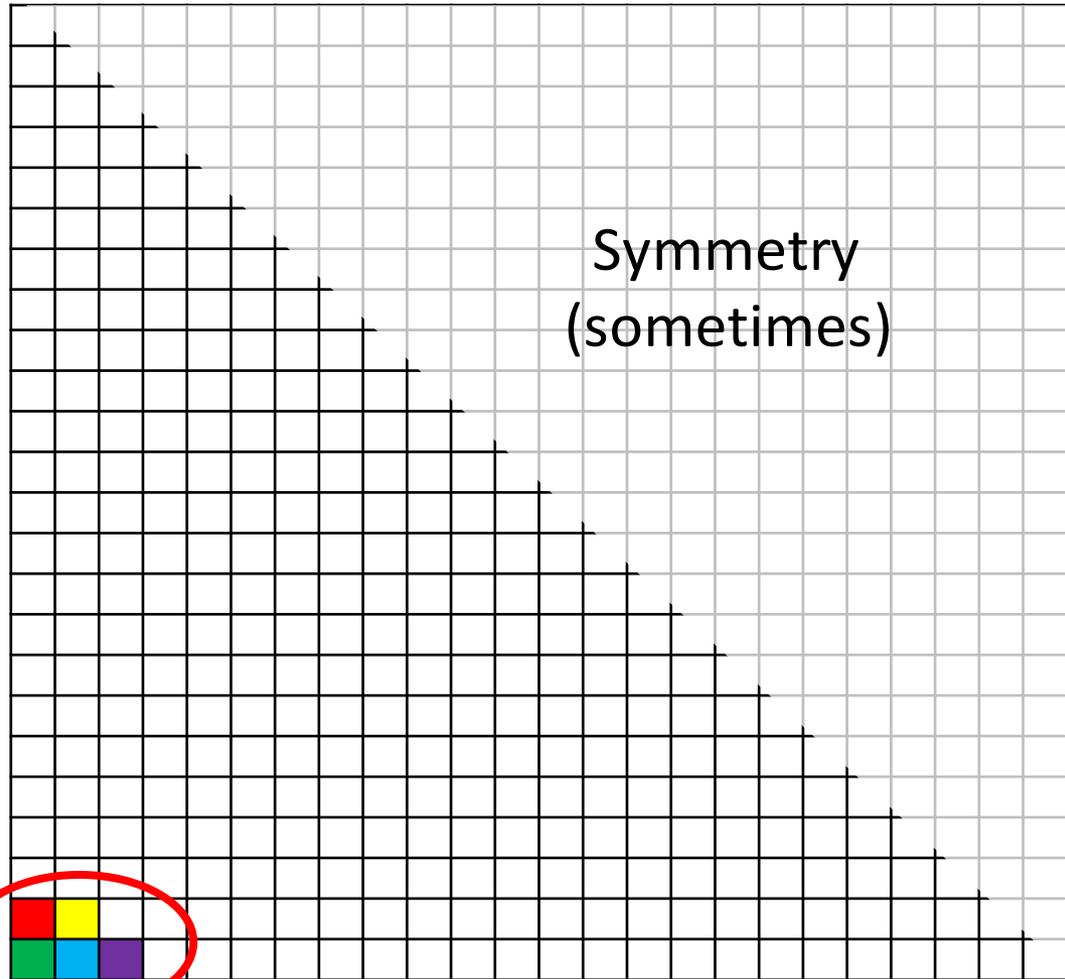


Voxels can be converted into “perfect” hexahedra.
The stiffness matrix for a **general** 8 node hexahedron is a 24 by 24 matrix .
A **voxel-based** 8 node hexahedron only requires **>1%** of this storage.

Element Storage for Voxelised Meshes

$j = 1, 24$

$i = 1, 24$



10 unique numbers
from 5 absolute
values (special case
for voxel meshes)

This much storage!

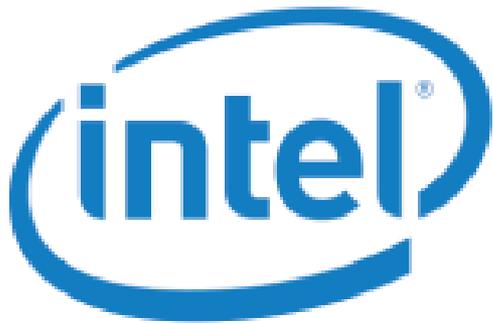
Summary

Summary

- Mixed mode MPI/OpenMP is promising!
- Xeon Phi is around 2 times faster than the 2 x 12 core Xeon host (with caveats).
- Small memory footprint per thread means performance hit when work per thread is low.
- Next steps are to evaluate new “low memory” finite element implementation, do a bit more tuning and then enjoy ...

Acknowledgement

This work was supported by the
Intel Parallel Computing Centre at the
IBM/STFC Hartree Centre in Warrington, UK



Intel Parallel Computing Centers

