



Nektar++: Enabling high-fidelity modelling on heterogeneous HPC platforms

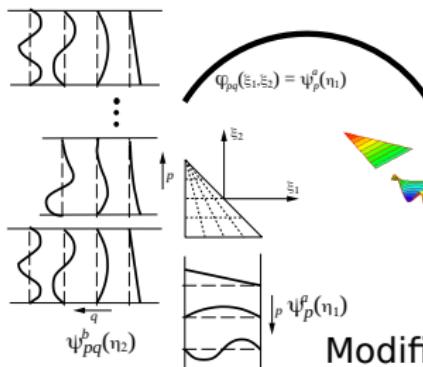
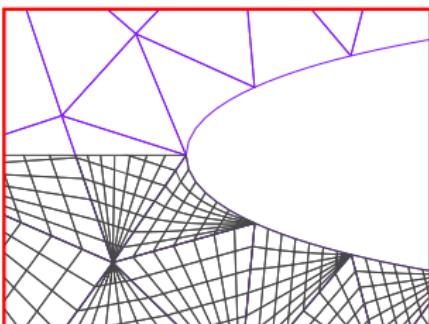
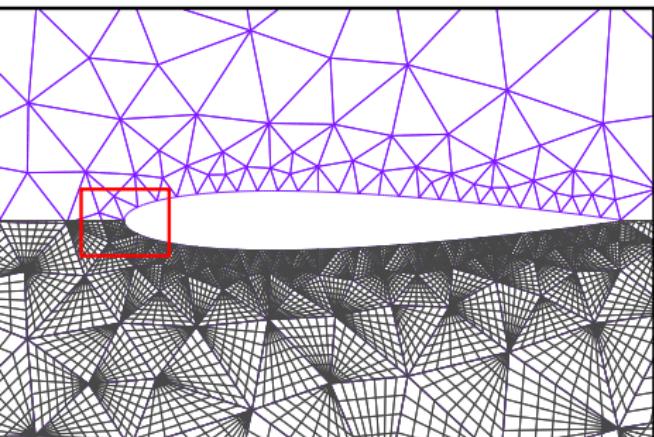
Chris Cantwell (ICL), David Moxey (KCL),
Jacques Xing (ICL), Boyang Xia (KCL), Alexandra Liosi (ICL),
Henrik Wustenberg (ICL), Chi Hin Chan (ICL),
Mike Kirby (Utah), Spencer Sherwin (ICL)

ARCHER2 Celebration of Science, Edinburgh
15th May 2025

Overview

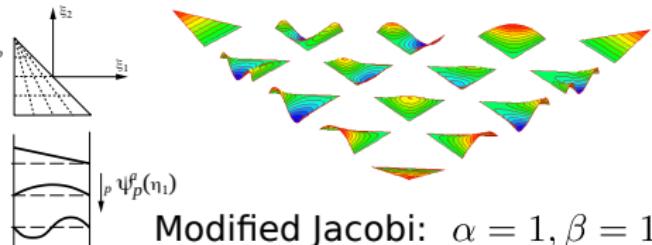
- Spectral/hp element methods and applications
- Nektar++ and its design philosophy
- Core code design
- Current performance results

Nektar++ and Spectral/hp element methods



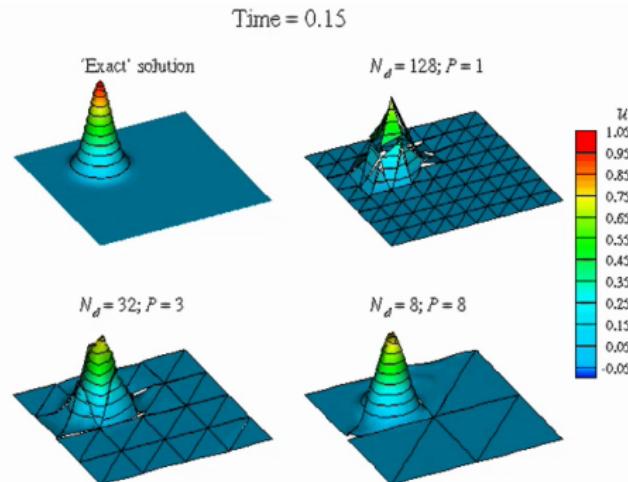
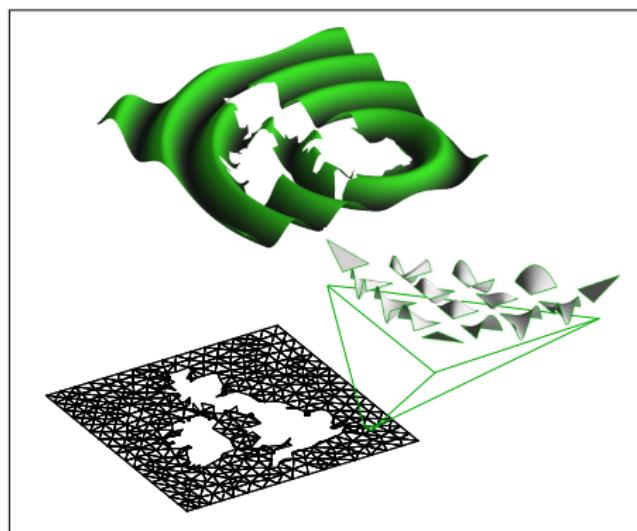
- High-order extension of linear finite element methods
- Enrich each element with high-order polynomial shape functions
- Coarser elements
- Nodal/Modal bases
- Tensor-product bases

$$P_n^{\alpha, \beta}(\xi)$$



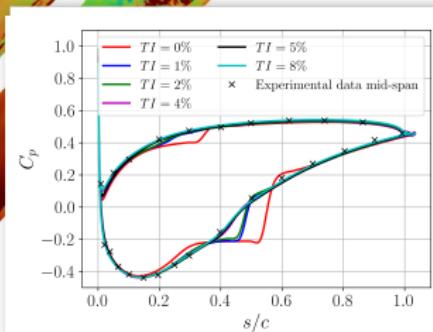
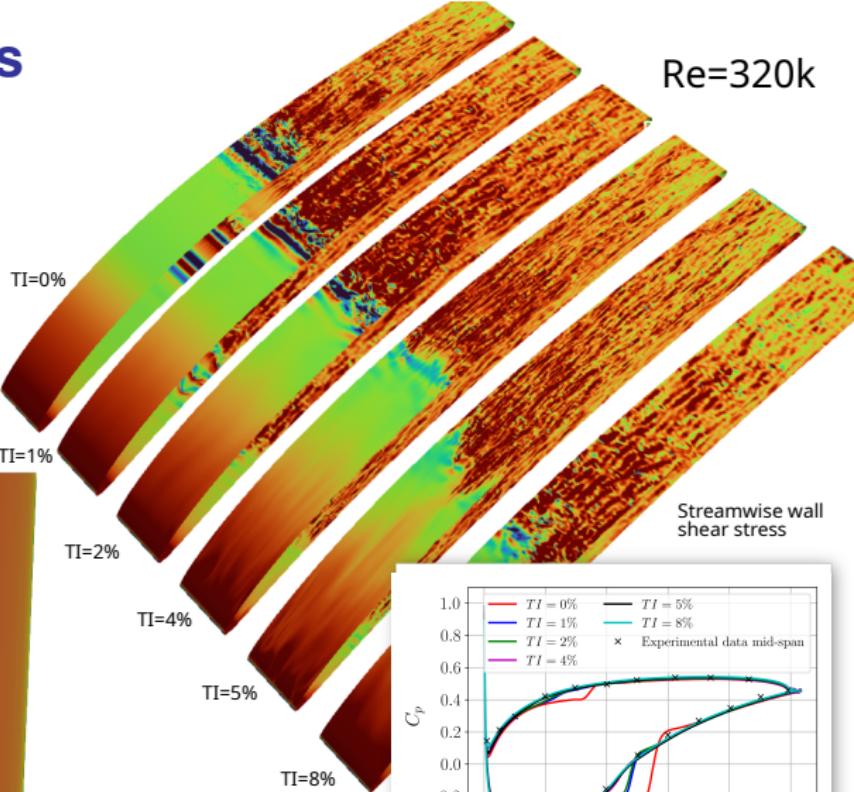
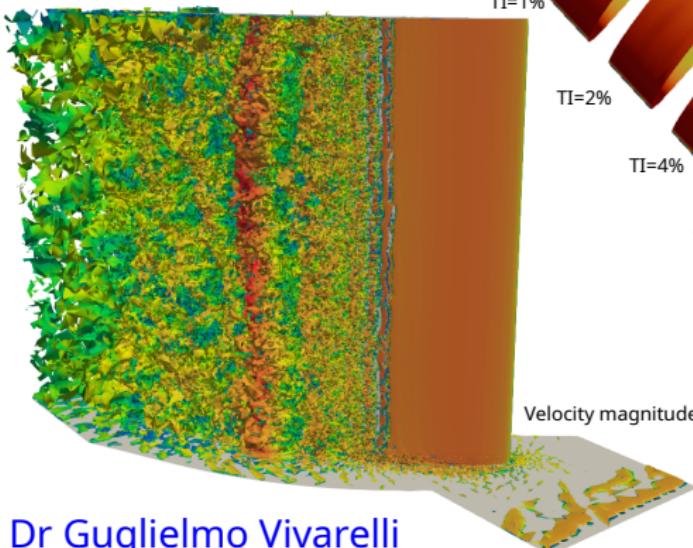
Nektar++ and spectral/hp element methods

Why are spectral/hp methods beneficial?



Enables improved capturing of multi-scale phenomena over longer time periods than low-order methods.

Compressor blades

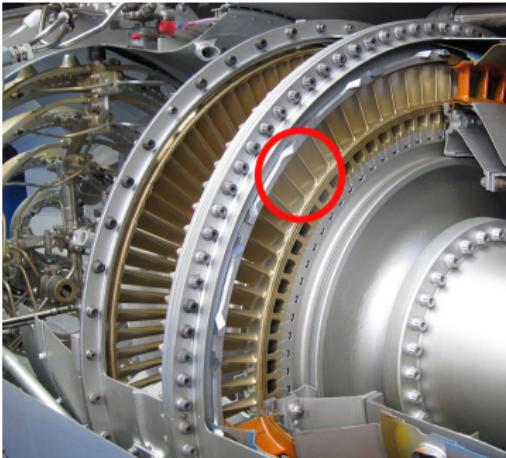
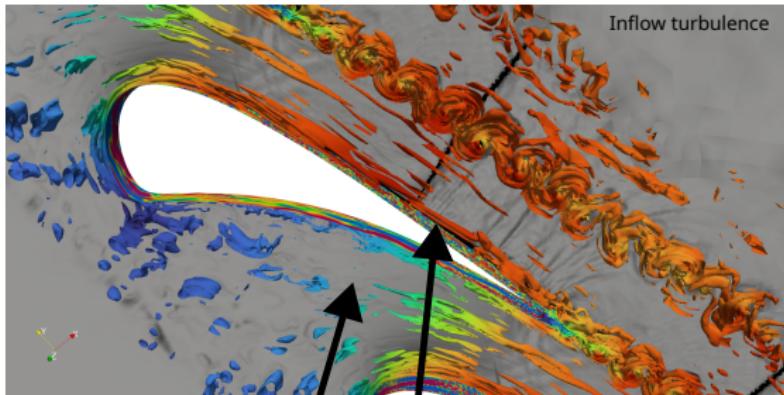


Dr Guglielmo Vivarelli



Rolls-Royce®

High-pressure Turbine blades



$Re=800k\sim 1M$

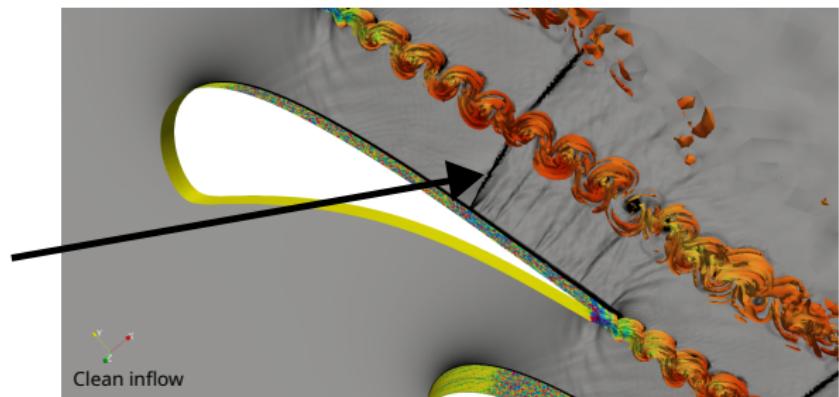
Ma 0.15 to 0.9-1.1
in 5-10cm

Boundary layer ~1mm

Shocks between blades

High resolution and
stabilisation required

Dr Guglielmo Vivarelli

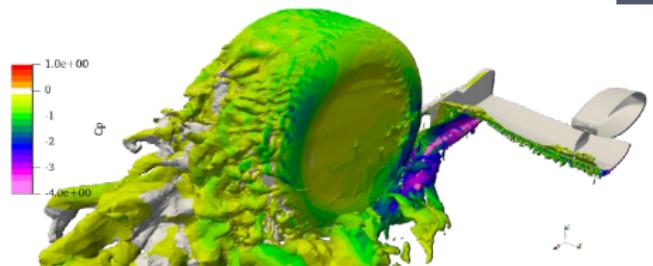
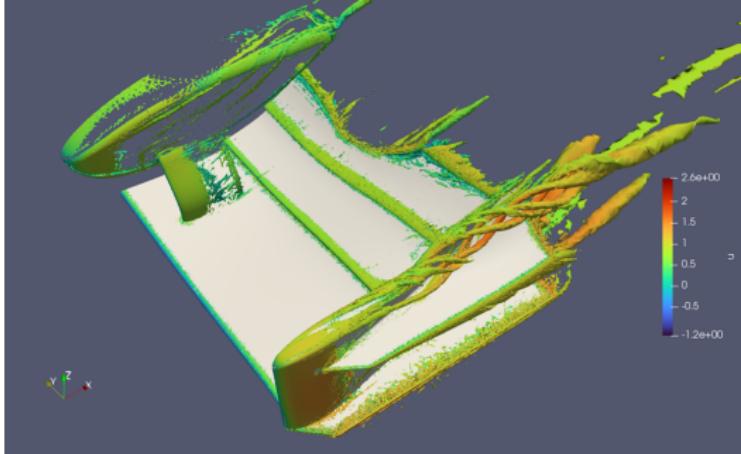
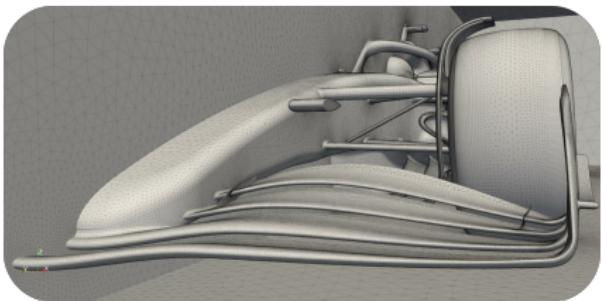


LS89 Test Case

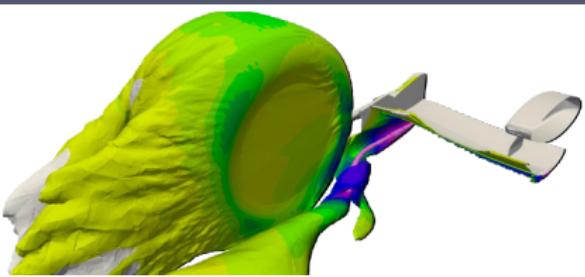


Rolls-Royce®

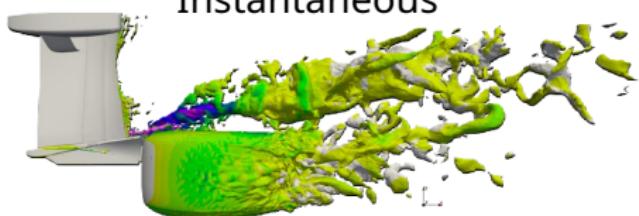
Formula 1



Instantaneous

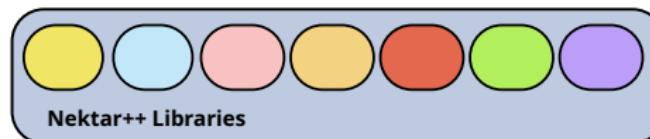
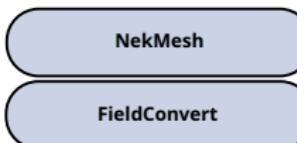
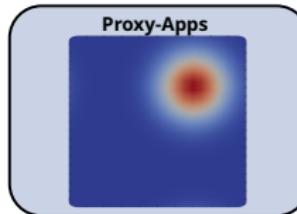
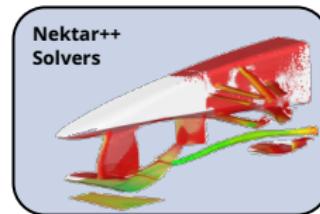


Time-averaged



Parv Khurana; Alexandra Liosi

Nektar++ framework



BLAS

LAPACK

Scotch

METIS

Boost

TinyXML

HDF5

gslib

MPI

Tetgen

FFTW

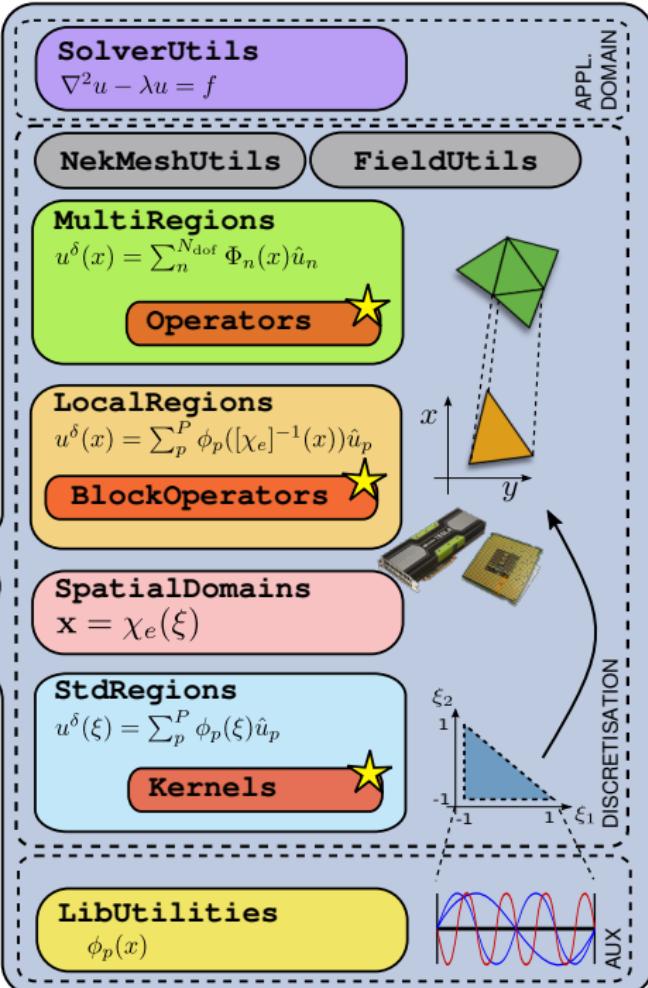
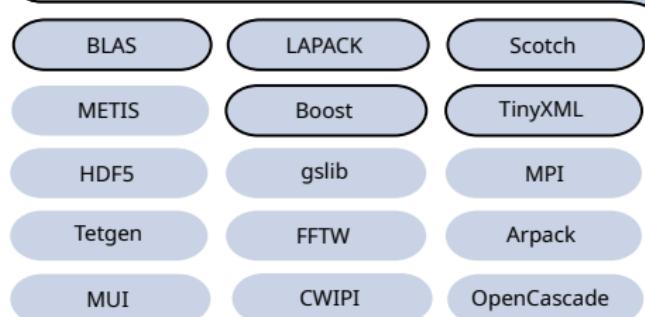
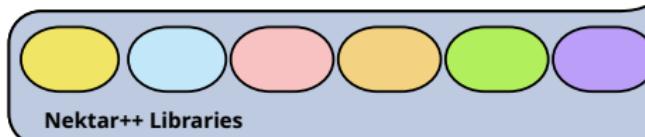
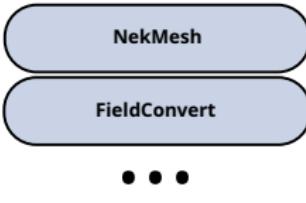
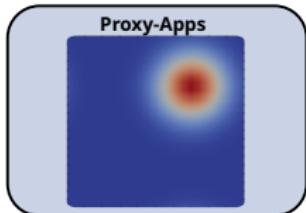
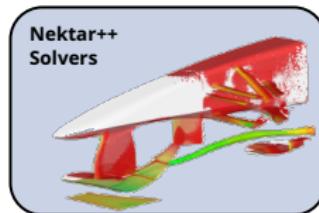
Arpack

MUI

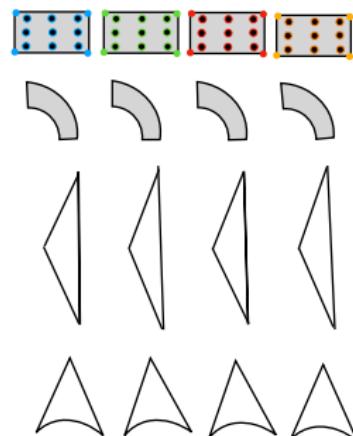
CWIPPI

OpenCascade

Nektar++ framework



Need for algorithmic flexibility



Polynomial order, Shape, Deformation

Interpolation/Backwards Transformation

$$\begin{matrix} \text{StdMatrix} \\ = \\ \begin{bmatrix} \text{StdMatrix} & \text{ColorMatrix} \end{bmatrix} \end{matrix}$$
$$\begin{matrix} \text{Sum-factorisation} \\ = \\ \begin{bmatrix} \text{SumMatrix} & \text{ColorMatrix} \end{bmatrix} \end{matrix}$$
$$\begin{matrix} \text{Sum-factorisation by Element} \\ (\text{vectorised}) \\ = \\ \begin{bmatrix} \text{SumMatrix} & \text{ColorMatrix} \end{bmatrix} \end{matrix}$$
$$\begin{matrix} \text{Sum-factorisation by Entry} \\ (\text{threaded}) \\ = \\ \begin{bmatrix} \text{SumMatrix} & \text{ColorMatrix} \end{bmatrix} \end{matrix}$$

- Different elements shapes / orders / deformation.
- Target architecture / CPU / GPU.
- Performance of algorithms a function of both.

Nektar++ on heterogeneous platforms

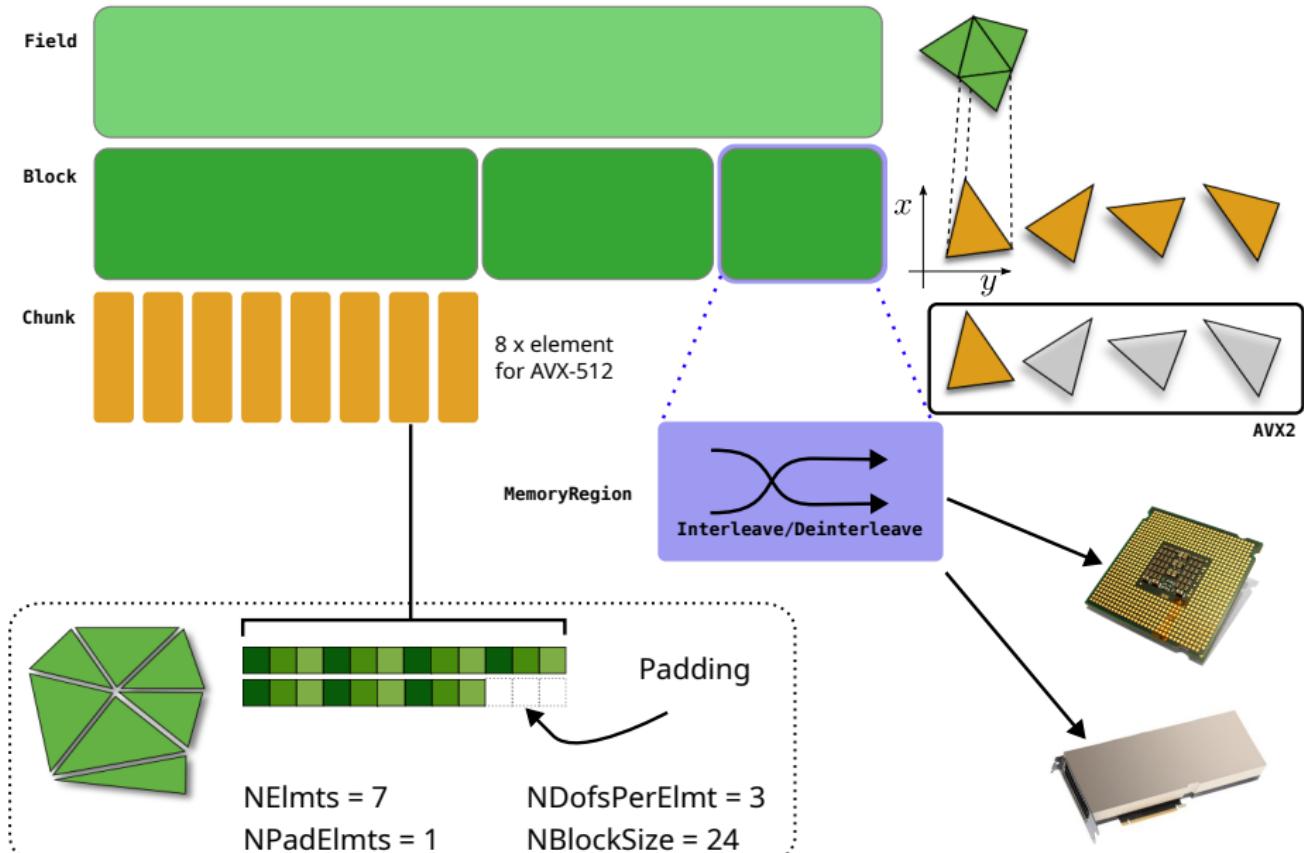
Design considerations

- Minimise data movement
- Align/interleave data appropriately to support SIMD / SIMT approaches
- Improve separation of data and algorithms
- Portability (hardware platforms / heterogeneity) – SYCL / Kokkos?
- Support platform-targeted optimisation (e.g. native kernels)
- Future-proofing: capability to support future hardware / programming models

Design decisions

- Abstraction layer to support multiple architectures / programming models and native kernels
- SYCL support to leverage both current and future platforms
- Templating of kernels to enable maximum compiler optimisation

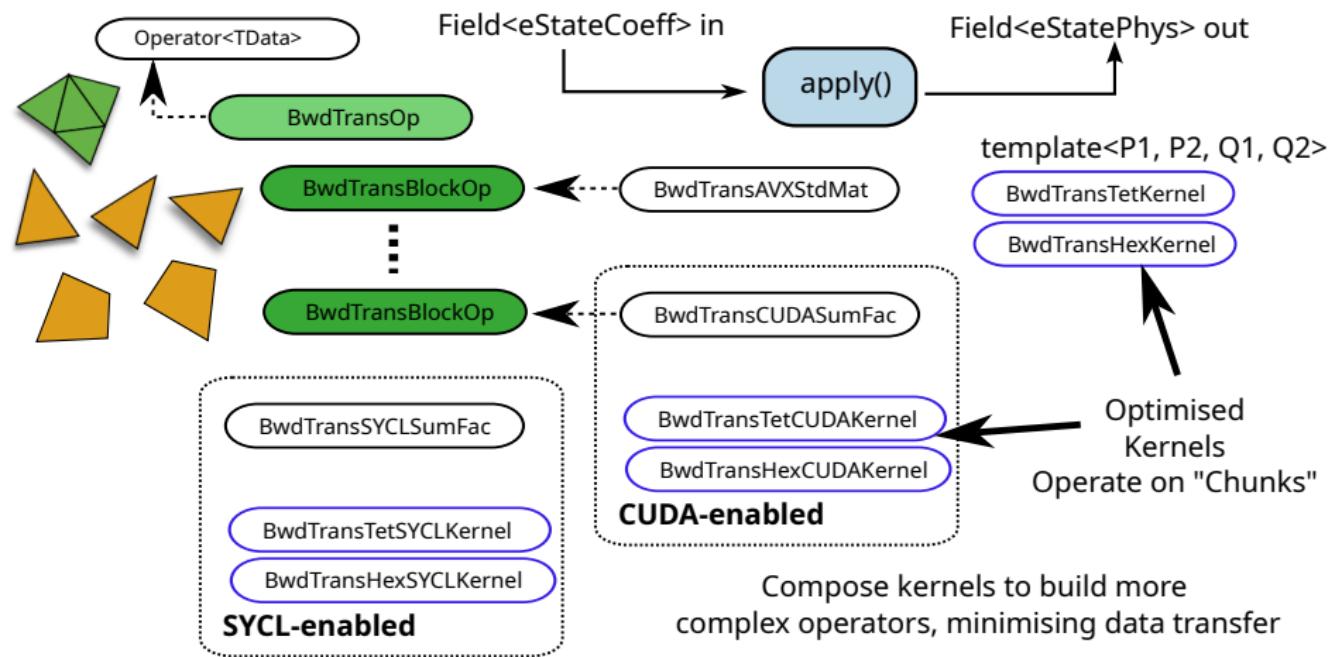
Code design: data management



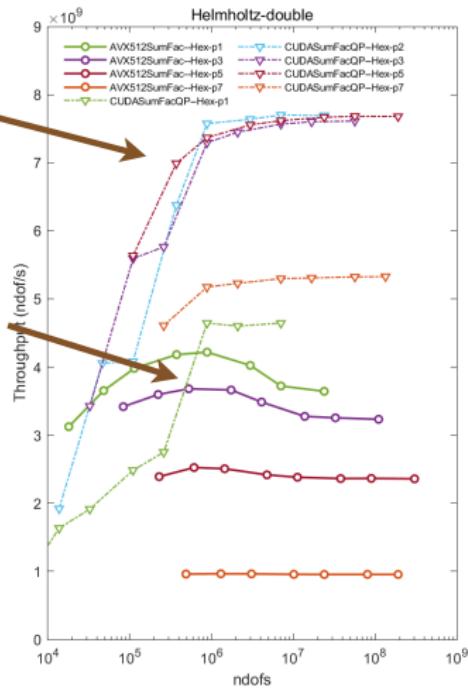
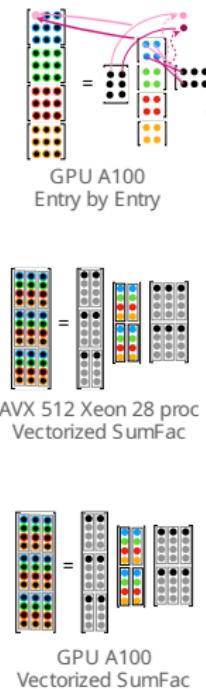
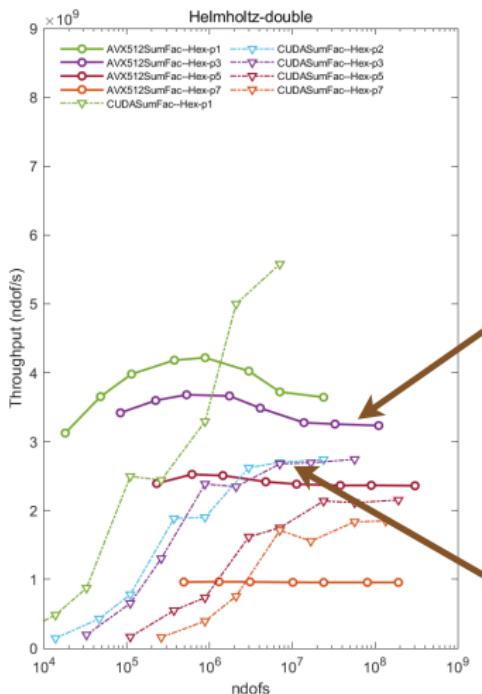
Code design: Operators and BlockOperators

Represent algorithmic steps in reusable, composable and efficient forms

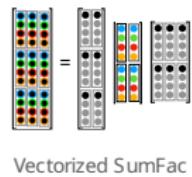
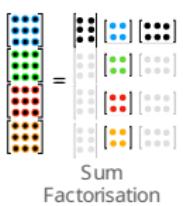
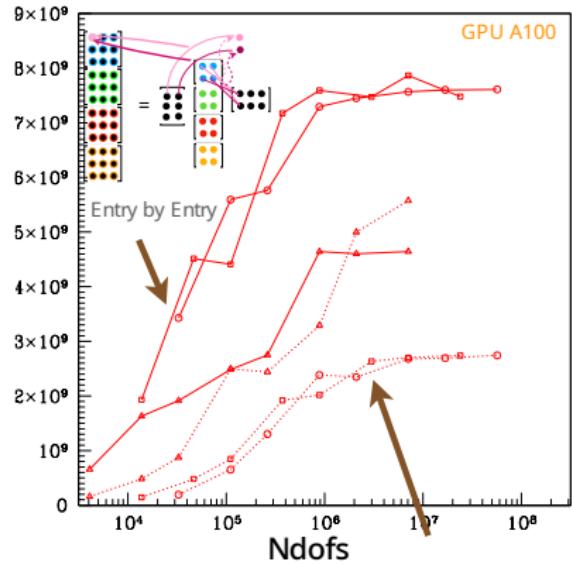
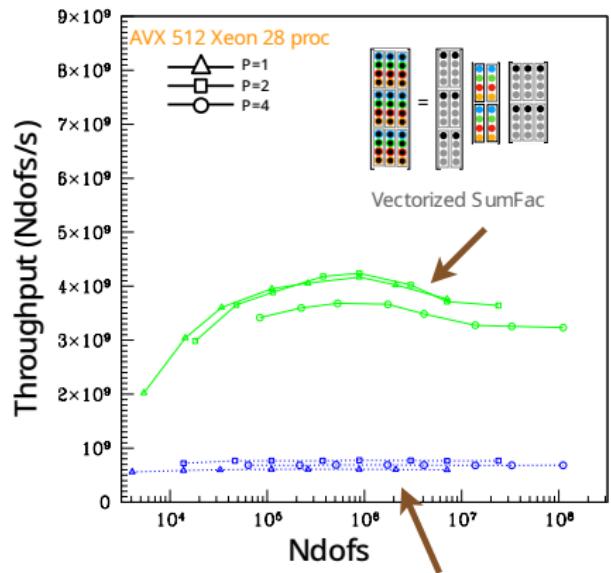
- **Abstraction:** solvers use base-class components - hiding specific implementation
- **Multiple implementations:** architecture-optimised, support programming models
- **High-performance kernels:** operations underpinned by compiler-optimised kernels



Performance



Performance



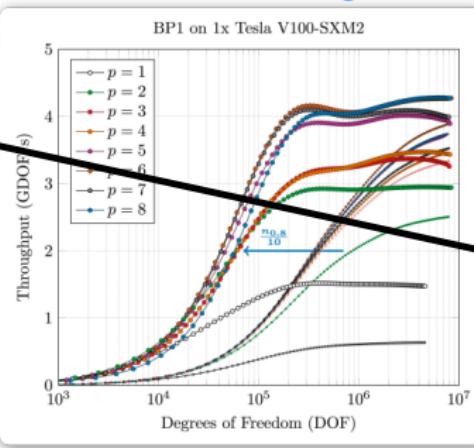
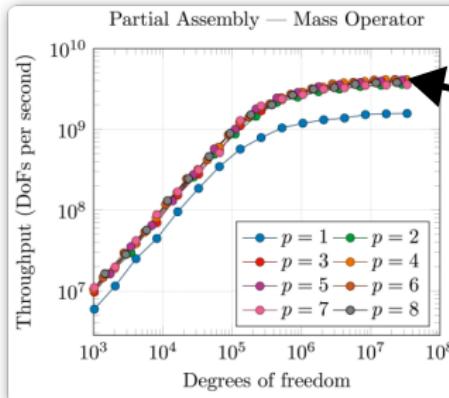
Comparable performance with other codes

Available access | Research article | First published online June 14, 2024

High-performance finite elements with MFEM

Julian Andrej, Nabil Atallah, Jan-Phillip Bäcker, Jean-Sylvain Camier, Dylan Copeland, Veselin Dobrev, Yohann Dudouit, Tobias Duswald, Brendan Keith, Dohyun Kim, Tzanio Kolev, Boyan Lazarov, Ketan Mittal, Will Pazner, Socratis Petrides, Syun'ichi Shiraiwa, Mark Stowell, and Vladimir Tomov  View all authors and affiliations

Volume 38, Issue 5 | <https://doi.org/10.1177/10943420241261981>



CUDA Kernels
Tetrahedral elements

Nektar++ Helmholtz ~ 7×10^9 DOFs/s

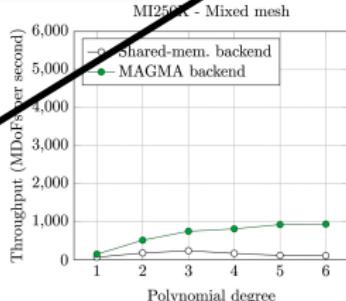
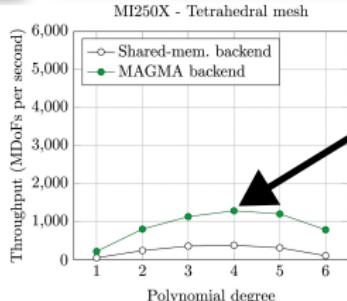
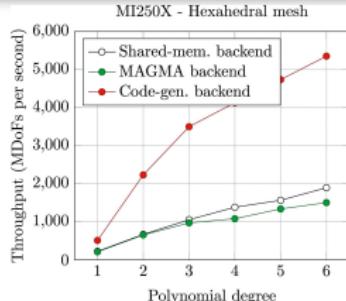
NVIDIA A100 GPU

MFEM Mass ~ 3×10^9 DOFs/s

NVIDIA V100 GPU

MFEM Mass ~ 1.5×10^9 DOFs/s

AMD MI250 GPU



Summary

- Nektar++ and spectral/hp element methods enable high-fidelity modelling in complex geometries
- Successfully used at scale on large CPU systems (such as ARCHER2)
- On-going work to redesign code to efficiently leverage GPU systems
- Requirement to switch between algorithms to maximise performance
- Current GPU performance competitive with other leading packages
- Performance benefits from tensor-product bases on simplex elements

Acknowledgements



ARCHER2 eCSE GPU project GPU-eCSE01-48

Other support:



Engineering and
Physical Sciences
Research Council



Rolls-Royce®



...and everyone who has contributed to the Nektar++ project!

Thank you for listening!