

Accelerating high-order mesh optimisation with an architecture-independent programming model

Jan Eichstädt^{*1}, Mashy Green¹, Michael Turner¹, Joaquim Peiró¹ and David Moxey²

¹ Department of Aeronautics, Imperial College London, Exhibition Road, SW7 2AZ
London, jan.eichstaedt13@imperial.ac.uk

² College of Engineering, Mathematics and Physical Sciences, University of Exeter,
D.Moxey@exeter.ac.uk

Keywords: *high-order mesh generation, architecture-independent programming model, Kokkos, portability, parallel hardware, variational framework*

Heterogeneous manycore performance-portable programming models and libraries, such as *Kokkos* [1], have been developed to facilitate portability and maintainability of high-performance computing codes and enhance their resilience to architectural changes. Here we investigate the suitability of the *Kokkos* programming model for optimizing the performance of the high-order mesh generator *NekMesh* [2], which has been developed to efficiently generate meshes containing millions of elements for industrial problems involving complex geometries. We describe the variational approach for *a posteriori* high-order mesh optimisation employed within *NekMesh* and its parallel implementation. We discuss its implementation for modern manycore massively parallel shared-memory CPU and GPU platforms using *Kokkos* and demonstrate that we achieve increased performance on multicore CPUs and accelerators compared with a native *Pthreads* implementation. Further, we show that we achieve additional speedup and cost reduction by running on GPUs without any hardware-specific code optimisation.

REFERENCES

- [1] H. Carter Edwards, C. R. Trott, D. Sunderland, Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *Journal of Parallel and Distributed Computing*, Vol. **74** (12), pp. 3202–3216, 2014.
- [2] M. Turner, D. Moxey, S. J. Sherwin, J. Peiró, Automatic Generation of 3D Unstructured High-Order Curvilinear Meshes. *ECCOMAS Congress 2016 VII European Congress on Computational Methods in Applied Sciences and Engineering*, pp. 5–10, 2016.