

# A comparison of the shared-memory parallel programming models *OpenMP*, *OpenACC* and *Kokkos* in the context of implicit solvers for high-order FEM

Jan Eichstädt<sup>1</sup>, Martin Vymazal<sup>1</sup>, David Moxey<sup>2</sup>, Joaquim Peiró<sup>1</sup>

---

## Abstract

We consider the application of three performance-portable programming models in the context of a high-order spectral element, implicit time-stepping solver for the Navier-Stokes equations. We aim to evaluate whether the use of these models allows code developers to deliver high-performance solvers for computational fluid dynamics simulations that are capable of effectively utilising both many-core CPU and GPU architectures. Using the core elliptic solver for the Navier-Stokes equations as a benchmarking guide, we evaluate the performance of these models on a range of unstructured meshes and give guidelines for the translation of existing codebases and their data structures to these models.

*Keywords:* Shared-memory parallel programming models, *OpenMP*, *OpenACC*, *Kokkos*, Helmholtz equation, FEM

---

## 1. Introduction

The use of computational fluid dynamics (CFD) in engineering design is now well-established practice. Such usage requires quick turnover of simulations in order to provide meaningful input into the design process. Simpler models such as Reynolds-Averaged Navier-Stokes (RANS) and Detached Eddy Simulation (DES) provide a route to computationally less expensive simulations, but at the cost of substantial reduction in accuracy, particularly for geometries admitting highly unsteady and transitional flows. Although higher fidelity models and, in particular, large eddy simulations (LES)

---

<sup>1</sup>Department of Aeronautics, Imperial College London

<sup>2</sup>College of Engineering, Mathematics and Physical Sciences, University of Exeter

provide a route to increased accuracy for these flows, they come at substantial increases in computational cost. Major efforts within CFD are therefore targeted at developing efficient LES codes and give the capability for overnight production runs of industrial scale LES.

One of the main issues that has been identified in producing such codes is the choice of numerical method that underpins the simulation. Low-order finite volume methods have long been the staple choice of industrial CFD software. However, such methods come with relatively low levels of arithmetic intensity: the balance between floating point operations (FLOP) and memory bandwidth. At the same time, current architectural trends in high performance computing (HPC) systems steer algorithm design towards higher arithmetic intensity, since memory latency and bandwidth continuously undergo a slower improvement compared with FLOP performance increases. This naturally limits the performance gains that can be realised on modern hardware using lower-order methods. To date therefore, the combination of low-order methods with LES models has typically not shown a sufficient increase in fidelity in the resulting solution fields to justify the additional computational cost.

This has driven CFD practitioners to contemplate alternative numerical discretisations with more favourable performance characteristics. In particular, higher order finite element methods and their variants, such as discontinuous Galerkin and flux reconstruction methods, offer an attractive alternative to classical low-order methods: numerically, low diffusion and dispersion error enable flow structures to be more accurately propagated across long convective lengths, and their higher arithmetic intensity allows for a more effective balance of floating point operations against memory bandwidth, as demonstrated in for example [1].

Although the theoretical performance of modern hardware is continuously improving, the hardware itself is becoming increasingly heterogeneous. Software is thus increasingly challenging to code and maintain for different architectures, particularly those that rely on a mixture of CPU and GPU hardware. This has led to interest in how different programming models can be constructed in order to better separate algorithmic from implementational concerns, providing performance portable and architecture independent models.

The aim of this work is to assess three such programming models, namely *Kokkos* [2], *OpenMP* [3], and *OpenACC* [4], in the context of high-order methods for CFD. Aspects to be compared are the implementational effort

and code maintainability, as well as the performance and its portability across architectures. In particular, we focus on implicit time-stepping Navier-Stokes solvers, as turbulent flow physics drives the algorithmic choices towards these schemes which avoid severe CFL number restrictions associated with explicit schemes.

This work provides novelty from a number of perspectives. Firstly, although performance portability aspects of individual programming models are regularly assessed, the corresponding results are still not meaningful enough for our purpose. Common benchmark problems are most often based on kernels for BLAS operations, or so-called *mini-applications* or *mini-apps* for a variety of solver types. Examples of these can be found in a number of parallel benchmark suites, e.g. [5, 6, 7, 8, 9]. In this approach, stripped-down standalone versions of the core algorithms that form complete solvers are implemented. This allows for a simpler development environment, whilst still ensuring generally representative performance characteristics. Such mini-apps then allow for easier performance benchmarking and experimentations with programming models and other aspects, such as utilised hardware or compiler versions. While each mini-app solver can consist either of bespoke or of standard BLAS-type kernels, the relative weights of these kernels and their interactions will still be an important factor in the overall performance, so that isolated evaluations of individual kernels alone will not yield meaningful enough results. The mini-application that comes the closest to our purpose would be *miniFE* of the *Mantevo* suite [6]: however, this is only for a linear finite element solver, whereas we are interested in high-order finite element solvers.

In this article we will follow this approach to develop a mini-app to model the performance characteristics of a complete Navier-Stokes flow solver, as implemented within the spectral/*hp* element framework *Nektar++* [10]. While we make use of the *Nektar++* framework for pre- and post-processing, the time-stepping loop that we monitor for performance has been cleared of typical *Nektar++* features and flattened to generic *C++* code, before applying the parallel programming models. Further to this, we note that the complete Navier-Stokes implementation requires particularly heavy implementation effort that does not necessarily affect computational performance (e.g. from specialised boundary conditions). We therefore select a simpler, more lightweight solver, that still contains the relevant features of the complete solver; in our case this will be an implicit solver for the 2D Helmholtz equation. We note that this forms the essential implicit solve for both pressure

Poisson and velocity correction steps within the commonly used fractional time-stepping scheme for Navier-Stokes equations [11]. The matrices arising from the discretisation are positive definite and, unless the mesh is highly distorted, well-conditioned and thus we employ the Conjugate Gradient method with a Jacobi preconditioner. The algorithm is executed in a global matrix-free fashion, that avoids assembly and storing of the large global system, whilst considering only local elemental matrices. As spatial discretisation we employ triangular elements with a modal continuous Galerkin (CG) formulation from [12]. The modal basis is formed with a tensor product, so that we can use the efficient sum-factorisation technique within elemental operations.

A second novel aspect of this work is the focus on implicit versus explicit time-stepping routines. Certainly the combination of using explicit time-stepping for high-order methods alongside performance portable and architecture independent methods running on both CPUs and GPUs has been demonstrated in solvers such as *PyFR* [13]. Additionally, work by other finite element groups such as *deal.II* [14] and *Dune* [15] on mostly tensor product quadrilateral and hexahedral elements has been conducted, but typically from the CPU perspective. In this work, we aim to demonstrate how architecture-independent programming can be applied from the context of implicit time-stepping.

The paper is structured as follows. In Section 2 we give a brief overview of current portable-performance programming models that we will consider in this work, alongside the architectures under which they will be applied. Section 3 outlines the mathematical model under consideration for the Helmholtz model, and its implementation is discussed in Section 4. We conduct a thorough performance evaluation in Section 5 on both many-core CPU and GPU hardware, before giving conclusions and a summary in Section 6.

## 2. Parallel hardware and portable programming models

Current HPC architectures are characterised by a high degree of parallelism. It does not only come in the more traditional form of distributed memory parallelism, as would be the case for clusters of single-core CPUs, but also in the form of intra-node or shared memory parallelism over multiple compute cores. It is this shared memory parallel hardware and the programming model to make use of it, that we address in this work.

### 2.1. Multi-core and many-core architectures

The parallelism of multi-core CPUs comes on two different levels, firstly over the multiple compute cores, and secondly over the vector lanes within each core. These are addressed by the compiler using AVX instructions. The length of the vector lane of the CPU that we use for our performance benchmarks has 256 bits or 4 doubles.

Nvidia GPUs, an instance of many-core architectures, also provide their parallelism in two levels. The first level consists of a number of streaming multiprocessors (abbreviated as SM or SMX), with the second being the individual streaming processors (SP) or *CUDA cores* that lie on each SMX. Groups of 32 *CUDA* cores have to execute as an individual SIMT unit, making these comparable to vector lanes on CPU architectures.

At the same time, memory bandwidth of both these architectures is limited, particularly on GPUs. Algorithms therefore need a high arithmetic intensity — the ratio of compute operations to memory accesses — to benefit from the available computational resources. For example, the Nvidia P100 GPU has a double-precision FLOP-to-byte ratio of 9.66 FLOP/byte, based on a theoretical peak of 5,304 GFLOP/s and 549GB/s memory transfer. This level of arithmetic intensity therefore requires level-3 BLAS or similar operations to fully utilise all *CUDA* cores.

To address the parallelism of the CPU, the programming models (for example *OpenMP* or *Pthreads*) schedule one threads per core (or two in case of hyper-threading). Vector level parallelism is much harder to be achieved; it can be done explicitly by using assembler-level intrinsic instructions or with a number of implicit options like compiler auto-vectorisation, by using Intel *Cilk Plus C/C++ Extensions for Array Notations* or the Intel *IPP* or *MKL* libraries [16].

Nvidia GPUs are addressed using the *CUDA* [17] programming model, which allows operations to be explicitly assigned to all parallelism levels. The use of object-orientation programming must, however, be done with care: while seamless integration of *CUDA* kernels into *C++* code is possible, there are restrictions on what can be done within a kernel in terms of data management, which could otherwise have been conveniently hidden with object-oriented features. Most notably in our work are the limitations to offload instances of classes (for example individual mesh elements of different types) to the GPU and iterate over all instances and implicitly calling their member functions and data, especially if they stem from class inheritance and virtual functions.

All these aspects combined explain the difficulty to achieve good shared-memory parallel performance, compared with distributed memory performance, enabled by domain decomposition. A good *MPI* distributed memory implementation should achieve perfect weak and strong scaling, as long as the problem size per *MPI*-rank is sufficient and communication costs can be hidden. However, to further reduce the execution time per problem size, shared memory parallelism need to be additionally employed. Yet, the additional fine-grained level of shared memory parallelism on the hardware needs to be exploited by the application, both through parallel algorithms, as well as programming-models and compilers that can map the algorithmic parallelism to the underlying hardware. It is this aspect that often result in less ideal strong-scaling within a compute node.

## 2.2. Kokkos *versus* OpenMP and OpenACC

The lifecycles of scientific applications are rather large, often spanning more than a decade. In comparison, hardware developments are deemed to proceed at an accelerated pace. HPC hardware has witnessed the advent and disappearance of vector processing machines over the 1970s and 1980s, and intermediate dominance of single-core multi-purpose CPUs until the mid 2000s. At this time the processor clock-speed has reached its energetic limit, resulting in the development of multi-core processors and subsequently longer vector-lanes. Additionally, through the introduction of *OpenCL* [18] and *CUDA* it became possible to program massively parallel but light-weight graphics processing units (GPUs). The recent strong demand for machine learning hardware, that requires massively-parallel high-throughput and low-precision architectures will further shift the available hardware for scientific computing. Many current scientific computing frameworks origin from the single-core CPU era and hence require updated capabilities to benefit from the increased parallelism. The evolutionary approach will allow to adjust the existing frameworks step-by-step with a compatible parallel programming model. This avoids extensive rewriting of code, that often is in the order of 100,000 lines. It is this approach that we adopt in this work.

Alternatively, the code-base could be rewritten in a different low-level language like *CUDA*. However, this only allows to address the currently popular GPU hardware, and in view of the rapidly changing HPC architectures does not result in a sustainable solution. The low-level language *OpenCL* is portable between CPUs and almost any device architectures including Nvidia GPUs, AMD GPUs and FPGAs. *OpenCL* could therefore under this premise

be a viable option for many users. However, it requires either substantial tuning efforts to yield good performance – which limits portability in practice, or the use of specialised libraries – which would either limit performance or complexity of our algorithms.

Another alternative that has been adopted by multiple groups is to design a new framework that comprises enough flexibility to address a broad range of current and future architectures. Examples of this include *PyFR* [13] for flux reconstruction methods, *Firedrake* [19] for FEM methods, and *OpenSBLI* [20] for finite difference methods. These frameworks consist of a high-level interface to describe the system that is to be solved, with an abstraction layer underneath this which translates high-level syntax to various low-level languages that can target different backends. Using a set of optimised BLAS libraries and/or bespoke kernels for each type of architecture, alongside good auto-tuning algorithms, can yield very competitive performance.

For this work we have identified three programming models that can be added to our existing *Nektar++* framework. As a requirement they needed to support both CPU and GPU architectures to expose their parallel capabilities, have open-access compiler support and a syntax that allows seamless implementation in an existing *C++* codebase.

*Kokkos* [2] is a library consisting of a performance portable abstraction layer in *C++*, with *OpenMP* [3] and *Pthreads* [21] backend for multi-core CPUs and a *CUDA* [17] backend for Nvidia GPUs. At the time of conducting our study it only supported data- and no task-parallelism, which prevents load sharing between host and device. The compiler support is very wide, spanning the *gcc*, *clang-llvm*, and *pgi* open-access compilers. The *Kokkos* syntax is based on *Kokkos* data arrays, that map the data between CPU and device memory. Its unique quality is the polymorphic layout of the 2D arrays, that can automatically switch between row- and column-major layout, to achieve coalesced memory access on different architectures. Up to three levels of hierarchical parallelism can be specified and data arrays can be explicitly placed on different cache levels or texture memories. Other similar models based on libraries with multiple backends are *RAJA* [22], which has been found in a less mature state than *Kokkos* and *OCCA* [23].

*OpenMP* [3] is a programming model based on compiler directives or so called *pragmas*, that can be added to legacy codebases and specify parallel constructs. Initially it only supported parallel executions on multi-core CPUs, but from the *OpenMP4.0* specifications onwards offloading to target devices is supported. The specification supports both task- and data paral-

lelism, which would allow for a real heterogeneous execution and load sharing between both host and device. The compiler support for the target offloading, however, is very sparse. The *gcc-7* compiler supports only the most basic *parallel-for* construct. The most mature open-access compiler has been the experimental *clang-ykt* compiler branch. We will show, that whilst it was possible to implement, compile and execute our benchmark problem using this compiler, it is still lacking crucial features that prevent its enrolment within a complete framework like *Nektar++*.

*OpenACC* is a similar programming model to *OpenMP* that is based on compiler directives, but primarily targeted at offloading to devices. It supports both task and data parallelism. The *gcc-6* compiler support very simple offloading directives and is hence not complete enough for our purpose. The *pgi* compiler, however, that supports Nvidia GPUs has a very mature ecosystem and will be employed in this work. More recently not only a host fallback is supported by the *pgi*-compiler, but also proper multi-core CPUs. Without using a combined programming model of *OpenACC* for the device and *OpenMP* for the host parallelism, no heterogeneous execution on host and device is possible.

A rough equivalent of *OpenACC*, but with an *OpenCL* and not a *CUDA* backend is *SYCL* [24]. However, our preliminary assessment of the supporting ecosystem and first performance results were not encouraging and we did not explore it further.

### 3. Method

We consider the solution of the inhomogeneous Helmholtz equation for a scalar field variable  $u$ , defined in a domain  $\Omega$  with boundary  $\Gamma = \Gamma_D \cup \Gamma_N$ , represented by the system

$$\nabla^2 u + \lambda u = f \quad \text{on } \Omega \quad (1)$$

$$u = u^D \quad \text{on } \Gamma_D \quad (2)$$

$$\frac{\partial u}{\partial n} = g \quad \text{on } \Gamma_N, \quad (3)$$

where  $\lambda$  is a real coefficient,  $f$  is a forcing function,  $\Gamma_D$  and  $\Gamma_N$  denote the parts of the boundary where we impose Dirichlet and Neumann boundary conditions, respectively.



Defining the inner products

$$(a, b)_\Omega = \int_\Omega ab \, d\Omega \quad \langle a, b \rangle_\Gamma = \int_\Gamma ab \, d\Gamma, \quad (4)$$

a weak Galerkin approximation to the system is obtained through the inner product with a shape function  $v$ , i.e.

$$(v, \nabla^2 u)_\Omega + \lambda(v, u)_\Omega = (v, f)_\Omega \quad (5)$$

We can now apply the divergence theorem to reduce continuity requirements and get

$$(\nabla v, \nabla u)_\Omega + \lambda(v, u)_\Omega = (v, f)_\Omega + \left\langle v, \frac{\partial u}{\partial n} \right\rangle_\Gamma \quad (6)$$

where the boundary integral allows to directly specify the Neumann boundary conditions. The Dirichlet boundary conditions are lifted from the solution,  $u = u^H + u^D$ , so that we retain the unknown homogeneous term,  $u^H$ , in the left-hand-side of the equation and all the known terms are moved to the right-hand side, namely

$$(\nabla v, \nabla u^H)_\Omega + \lambda(v, u^H)_\Omega = (v, f)_\Omega + \langle v, g \rangle_{\Gamma_N} - (\nabla v, \nabla u^D)_\Omega - \lambda(v, u^D)_\Omega \quad (7)$$

We can now discretize this expression to obtain a linear system of the form

$$H \hat{u}^H = \hat{b} \quad (8)$$

where  $H$  is the matrix of coefficients of the discrete Helmholtz operator,  $\hat{u}^H$  is a vector containing the discrete unknowns and  $\hat{b}$  is the vector that results from the evaluation of the right-hand side using the known values of the forcing function and boundary conditions.

### 3.1. Implicit solver

The linear Helmholtz equation system is to be solved iteratively using the conjugate gradient method. This method is suitable, as the Helmholtz operator is symmetric and positive definite.

The method is initialised with an arbitrary solution  $\hat{u}_0^H$ :

$$r_0 = \hat{b} - H^H \hat{u}_0^H \quad (9)$$

$$w_0 = Cr_0 \quad (10)$$

$$s_0 = H^H w_0 \quad (11)$$

$$\beta = 0 \quad (12)$$

$$\alpha = \frac{r_0^T w_0}{r_0^T w_0}, \quad (13)$$

where  $r$  is the residual vector,  $C$  a suitable preconditioner,  $w$  and  $s$  are the conjugate gradient vectors, and  $\beta$  and  $\alpha$  are step sizes. The method then loops over the following steps  $k$ , until the residual is smaller than a set tolerance  $\epsilon > \epsilon_{tol}$ .

$$p_{k+1} = \beta p_k + w_k \quad (14)$$

$$q_{k+1} = \beta q_k + s_k \quad (15)$$

$$\hat{u}_{k+1}^H = \alpha p_{k+1} + \hat{u}_k^H \quad (16)$$

$$r_{k+1} = r_k - \alpha q_{k+1} \quad (17)$$

$$w_{k+1} = Cr_{k+1} \quad (18)$$

$$s_{k+1} = H^H w_{k+1} \quad (19)$$

$$\beta = \frac{r_{k+1}^T w_{k+1}}{r_k^T w_k} \quad (20)$$

$$\alpha = \frac{r_{k+1}^T w_{k+1}}{s_{k+1}^T w_{k+1} - r_{k+1}^T k + 1 w_{k+1} \cdot \beta_{k+1} / \alpha_k} \quad (21)$$

$$\epsilon = r_{k+1}^T r_{k+1} \quad (22)$$

Here  $p$  and  $q$  are search direction vectors and  $\hat{u}_{k+1}^H$  is the updated solution to the linear system that is to be solved. As the preconditioner  $C$ , we use a simple diagonal Jacobi preconditioner.

### 3.2. Discretisation of the Helmholtz equation

We now give a brief overview of the discretisation of the Helmholtz equation, which is described further in [12]. The previous algorithm is applicable to any spatial discretisation, but the size, structure and numerical properties of the matrices will depend on our choice of discretization. Here we will be using a continuous Galerkin projection over triangular elements with a

hierarchical modal basis. We note the central aspect of using a matrix-free scheme, that avoids assembling the complete global matrix of the Helmholtz operator. Instead we are breaking down the global operator into as a sum of elemental operators. This implies that our global vector of solution coefficients  $\hat{u}^H$  needs to be mapped first to the element local solution coefficients in a scatter-type operation:

$$\hat{u}_l = \mathbf{A}\hat{u}_g \quad (23)$$

Since the global-to-local matrix  $\mathbf{A}$  is very sparse it is implemented using a mapping vector  $v_{map}$  and a sign vector  $v_{sign}$  of length  $n_{local}$  as

$$\hat{u}_l[i] = v_{sign}[i] \hat{u}_g[v_{map}[i]] \quad \text{for } i \in [0 : n_{local}] \quad (24)$$

On each element  $e$  the elemental Helmholtz operator can then be applied as

$$\hat{s}_l^e = \mathbf{H}^e \hat{u}_l^e \quad \text{for } e \in [0 : n_{el}] \quad (25)$$

Finally, the local solution coefficients need to be mapped back to the global solution coefficients in a gather-type operation:

$$\hat{s}_g = \mathbf{A}^T \hat{s}_l \quad (26)$$

Equivalently to the global-to-local operation this is implemented as

$$\hat{s}_g[v_{map}[i]] = \hat{s}_g[v_{map}[i]] + v_{sign}[i] \hat{s}_l[i] \quad \text{for } i \in [0 : n_{local}] \quad (27)$$

Here it can be seen that boundary modes of neighbouring elements need to be added. In a parallel implementation this basic implementation can lead to a race-condition.

The elemental Helmholtz operation for a 2D element is the combination of the Laplacian operator and the mass operator and is written as

$$\hat{s}_l^e = \mathbf{H}^e \hat{u}_l^e \quad (28)$$

$$\hat{s}_l^e = [\mathbf{L}^e + \lambda \mathbf{M}^e] \hat{u}_l^e \quad (29)$$

$$\hat{s}_l^e = [(\mathbf{D}_{x_1} \mathbf{B})^T \mathbf{W} \mathbf{D}_{x_1} \mathbf{B} + (\mathbf{D}_{x_2} \mathbf{B})^T \mathbf{W} \mathbf{D}_{x_2} \mathbf{B} + \lambda \mathbf{B}^T \mathbf{W} \mathbf{B}] \hat{u}_l^e \quad (30)$$

The basis matrix  $\mathbf{B}$  with  $B_{i,j} = \phi_j(\xi_i)$  denotes the backward-transform from coefficient space to physical space, its transform  $\mathbf{B}^T$  is the forward operator from physical space to coefficient space. The diagonal weight matrix  $\mathbf{W}$  with  $W_{i,j} = J_{i,j} w_i w_j$  includes the Jacobian  $J$ , that is the mapping between

the standard element defined on  $\xi \in [-1, 1]$  and the curvilinear element on coordinates  $x$ , as well as the quadrature weightings  $w_i$  and  $w_j$ . Following [12] we select a Gauss-Lobatto-Legendre quadrature for  $w_i$  and a Gauss-Radau quadrature for  $w_j$  to mitigate the effects of the singularity in the collapsed coordinate Duffy mapping. The derivative matrices  $\mathbf{D}_{x_1}$  and  $\mathbf{D}_{x_2}$  are

$$\mathbf{D}_{x_1} = \mathbf{\Lambda} \left( \frac{\partial \xi_1}{\partial x_1} \right) \mathbf{D}_{\xi_1} + \mathbf{\Lambda} \left( \frac{\partial \xi_2}{\partial x_1} \right) \mathbf{D}_{\xi_2} \quad (31)$$

$$\mathbf{D}_{x_2} = \mathbf{\Lambda} \left( \frac{\partial \xi_1}{\partial x_2} \right) \mathbf{D}_{\xi_1} + \mathbf{\Lambda} \left( \frac{\partial \xi_2}{\partial x_2} \right) \mathbf{D}_{\xi_2}, \quad (32)$$

with the diagonal coefficient matrices  $\mathbf{\Lambda}$ . The block-diagonal derivative matrices  $\mathbf{D}_{\xi_1}$  and  $\mathbf{D}_{\xi_2}$  are

$$\mathbf{D}_{\xi_1} = \frac{\partial}{\partial \xi_1} \quad (33)$$

$$\mathbf{D}_{\xi_2} = \frac{\partial}{\partial \xi_2} \quad (34)$$

To minimise the operator count for an efficient CPU implementation, the elemental matrices  $\mathbf{W}$  for the mass operator have been precomputed and stored. For the elemental Laplacian operators, the following matrices have been precomputed and stored:

$$\mathbf{L}_{11} = \mathbf{D}_{x_1}^T \mathbf{W} \mathbf{\Lambda} \left( \frac{\partial \xi_1}{\partial x_1} \right)$$

$$\mathbf{L}_{12} = \mathbf{D}_{x_1}^T \mathbf{W} \mathbf{\Lambda} \left( \frac{\partial \xi_2}{\partial x_1} \right)$$

$$\mathbf{L}_{21} = \mathbf{D}_{x_2}^T \mathbf{W} \mathbf{\Lambda} \left( \frac{\partial \xi_1}{\partial x_2} \right)$$

$$\mathbf{L}_{22} = \mathbf{D}_{x_2}^T \mathbf{W} \mathbf{\Lambda} \left( \frac{\partial \xi_2}{\partial x_2} \right)$$

Where possible, we additionally exploit the tensor-product construction of the  $C^0$  basis to apply the sum-factorisation technique [25]. This is a widely used strategy for attaining substantial reductions in operator count for tensor-product operations, including backwards transforms and derivatives. For example, the matrix-vector product  $\mathbf{B}\hat{\mathbf{u}}$  would naively be calculated as:

$$u(\vec{\xi}) = \sum_p \sum_q \hat{u}_{pq} \phi_p(\xi_1, \xi_2)$$

which has around  $P^4$  floating-point operations. Using sum-factorisation however, and noting that  $\phi_{pq}(\xi_1, \xi_2) = \psi_p^1(\xi_1)\psi_q^2(\xi_2)$ , this matrix vector product can be expressed as:

$$u(\vec{\xi}) = \sum_p \psi_p(\xi_1) \left[ \sum_q \hat{u}_{pq} \psi_q(\xi_2) \right]$$

where the brackets denote temporary storage, at a cost of order  $P^3$ . Note that, in the above, these smaller one-dimensional products can be represented as matrix-matrix products. For example, taking  $\mathbf{B}_1$  and  $\mathbf{B}_2$  to denote the basis matrices of the one-dimensional shape functions  $\phi^1$  and  $\phi^2$ , we have that

$$\mathbf{B}\hat{\mathbf{u}} = [\mathbf{B}_1 \hat{\mathbf{u}}_{[Q_1]}]_{[Q_2]} \mathbf{B}_2^\top,$$

where  $[Q_1]$  denotes the reinterpretation of a vector as a  $Q_1 \times Q_2$  matrix with  $Q_i$  denoting the order of the quadrature in each direction. Existing implementations therefore focus on the use of optimised matrix-matrix routines such as the BLAS level-3 `dgemm`, as in [26], to perform these operations efficiently, or through the use of optimised libraries for small-matrix multiplications such as `libxsmm` [27]. In our implementation, we therefore consider an application of the Helmholtz operator in a matrix-free fashion that utilises these sum-factorisation kernels for each of the components of the Helmholtz equation above, rather than storing the elemental matrices  $\mathbf{H}^e$ . More details of this technique applied to the spectral element method on triangles and other higher dimensional shape types can be found in references [26, 12, 28].

#### 4. Implementation and parallelisation strategy

In the following, we discuss the individual components of the method, as introduced in Section 3. For each one we develop an individual parallel kernel, that ideally performs well on both CPU and GPU architectures. As will be seen, this is achieved apart from the *gather*-type reduction kernels, see Section 4.2.

The first kernel implements the AXPY operations specified by Equation 14 to Equation 18. All three programming models, *Kokkos*, *OpenMP*, and *OpenACC* have a suitable syntax and implementation for these *parallel-for-loop*-type of equations. The same applies to the *scatter*-type operation, specified by Equation 24.

The *gather*-type operation specified by Equation 27 can not readily be parallelised without causing a race-condition, as multiple local modes are added up to one global mode and hence need to write to the same memory space. How we still achieve reasonable parallel performance is discussed in Section 4.2.

The calculation of the conjugate gradient step sizes  $\alpha$  and  $\beta$  and the residual  $\epsilon$ , as given by Equation 21, Equation 20, and Equation 22 are classical reduction operations. These can be parallelised using the *parallel-reduce*-type syntax. Even though all three programming models support this operation, the *llvm-ykt* compiler branch for *OpenMP4.0* did not support it at the time of conducting this work. Instead, we wrote our own kernel for parallel-reductions. This kernel is based on a two level-implementation as follows. A first level that spawns multiple blocks to utilise the individual streaming multiprocessors of the GPU, so that each reduces a chunk of the overall array down to one scalar. The second level then takes all this intermediate results and performs another reduction on a single SMX to yield the final reduced scalar. We use zero-padding of the arrays on both levels to improve the performance. Although this custom kernel does not achieve the optimised performance of an optimal BLAS library implementation, it does not compromise the overall performance of the method. This will be shown in the results section (see Section 5).

All of the components above have costs that roughly scale with the degrees of freedom in the system. The computationally most demanding part, though, is the evaluation of the elemental Helmholtz operator, as given by Equation 30. A careful evaluation of the parallel algorithm is hence performed in Section 4.1.

#### 4.1. Interleaved data structures and kernels

Each of the elemental Helmholtz operations (see Equation 30) can be performed independently and thus in parallel. On a CPU without vectorisation, we would assign one thread per element or set of data, but on the GPU, there are multiple options. These are primarily explained by the constraint of GPUs and the *CUDA* programming model, which forces all 32 *CUDA* threads in one warp to perform the same mathematical or logical operation in any one cycle. Otherwise the warp will diverge so that the different operations are executed sequentially over multiple cycles, decreasing the performance. Three different work distributions can be distinguished here:

- (a) Each *CUDA* kernel or GPU processes one set of data.

- (b) Each *CUDA* block or streaming multiprocessor (SMX) processes one set of data.
- (c) Each *CUDA* thread or core processes one set of data.

Option (a) is only suitable for a single, very large operation and hence can be discarded for our purpose here. On the Nvidia P100 GPU for example, a matrix size of  $128 \times 16384$  on a `sgemm` kernel is necessary to get more than 50% of the peak performance of *CuBLAS* `sgemm` operations. Option (b) is mostly suited for operations on rather few, but large sets of data. The individual sets might vary in size and might undergo different operations, without inducing performance breakdowns. Option (c) is only suited for operations on many, small sets of data. For good performance, however, the individual sets should all have the same size and undergo the exact same operations.

In option (b), vectors or matrices need to be sufficiently large so that the overall operation can be efficiently split into multiple warps, in which all threads perform the same elemental mathematical operation. This is because in order to be fully occupied, each SMX of a Nvidia GPU needs to schedule 32 or 64 warps at the same time. Operations over multiple *CUDA* blocks are not synchronised, so that different operations can be executed without performance penalty, as long as a good load balancing between the different *CUDA* blocks is ensured. In terms of data structures, option (b) can operate efficiently on data that is arranged one set after the other. For matrix-matrix multiplication, techniques like *tiling* are further employed to split the overall operation efficiently into elemental operations on which each thread in a warp operates on coalesced data. This option is commonly implemented as *batched* or *strided* BLAS operations. Both these operations are collections or batches of BLAS operations that are performed concurrently, so that the batches will be automatically distributed over the individual SMX. As a special case, strided operations assume the same lengths of data arrays across all batches. The *MAGMA* library [29], the Intel *MKL* library, and the *CuBLAS* library [30] all support batched `gemm` operations, the latter also strided batched operations. Again on the Nvidia P100, 128 batched single precision matrix multiplications with matrix size of  $128 \times 128$  will result in about 50% peak performance. This is already an improvement from individual `sgemm` kernels, but lower polynomial order elemental matrix sizes are still far smaller than  $128 \times 128$ . E.g. a `gemm` operation occurring in the sum-factorisation kernel of a triangular element with polynomial order 10 has matrix sizes of  $11 \cdot 12$ .

In option (c), all data sets need to have the same size and undergo the same overall operation, so that 32 threads can always be efficiently combined into a warp. Here a sufficiently large number of the small data sets is required to fully occupy the number of concurrent warps that need to be scheduled. For this approach however, the naive memory layout of having data sets consecutively in memory would violate any coalesced data access. This is because in this case, the specific data entries of each set are separated by the whole length of each data set. Yet, all these specific data entries have to be loaded for the warp operations at the same time. For an efficient memory access, it follows that all specific data entries of a group of sets need to be coalesced. The group of sets is chosen as a multiple of the warp-size of 32 threads. As a result, the data sets of one such group become interleaved. The distance or stride between following entries of one data set is equal to the number of data sets that are combined in one group. Interpreting the simple data layout in which one data set is stored after each other, corresponds to a stride of one. For CPUs, this concept of explicit vectorisation over multiple elements using interleaved memory layouts has also been implemented in the *DUNE* library [15], and the *deal.II* library [14, 31] to benefit from the wider vector lanes.

For level-1 and level-2 BLAS operations, stride can be specified as a variable, so that no major code adjustment is necessary. Standard level-3 BLAS operations like `dgemm` which are used for our sum-factorisation kernels, however, do not possess a stride variable, so that we decided to implement a custom strided `dgemm` function, that can operate on the interleaved memory layout. Following the function call syntax of other strided batched `gemm`, we introduced two additional variables for each of the three data arrays representing the 3 sets of matrices. The two variables are the stride between the following entries, which is equal to the size of the group, and a variable denoting the number of the group a specific thread operates on. The function call syntax and the operations the algorithms of our custom `dgemm` kernel is given in Listing 1.

Listing 1: Custom interleaved `dgemm` syntax and algorithm

```
int interleavedDgemm(char transa, char transb,
    int M, int N, int K,
    const double alpha,
    const double *A, int lda, int strideA, int instA,
    const double *B, int ldb, int strideB, int instB,
    const double beta,
```



```

        double *C, int ldc, int strideC, int instC)
{
    // one of four cases shown
    if (transa == 'T',transb == 'N') {
        for (int m = 0; m < M; ++m) {
            for (int n = 0; n < N; ++n) {
                double c_mnp = 0;
                for (int k = 0; k < K, ++k)
                    c_mnp += A[(m + k*ldA) * strideA + instA]
                        * B[(k + n*ldB) * strideB + instB];
                C[(m + n*ldC) * strideC + instC] =
                    alpha * c_mnp +
                    beta * C[(m + n*ldC) * strideC + instC];
            }
        }
    }
}

```

Having developed an efficient and working implementation for option (c), and because the performance of option (b) would be severely limited due to the small matrix sizes considered here, we proceed with option (c) throughout this work.

There is another limitation of using a batched function as in option (b), because a new kernel needs to be scheduled for each of these BLAS functions. This is not ideal in conjunction with our framework, in which the code is structured in a way that describes the operations to be executed for each element or data set. Ideally, the data will be loaded into the cache at the beginning of such extended kernels and then operated on by a sequence of BLAS functions and, at the end, stored back into global memory. Using *CuBLAS* functions that start a kernel for each BLAS function will require multiple memory copies between global memory and cache, compromising efficiency. Further, it would require to re-factor the code so that the extended kernel is broken up into chunks corresponding to one BLAS function each.

We further found, that explicitly using the interleaved memory layout and strided kernel has only been strictly necessary for kernels in which dynamic memory allocation within the kernel is specified. This is the natural way to write with the *Kokkos* syntax and the natural extension from our initial CPU code. The variable length of arrays depending on the element polynomial order can more readily be accounted for using dynamic memory allocation. However, the *OpenMP*-target compiler did not cover this functionality and the *OpenACC* implementation was poorly performing. Here we had to write

---

**Algorithm 1** Initial local-to-global reduction algorithm

---

```
1: procedure LOCALTOGLOBALINITIAL( $\hat{s}_g, \hat{s}_g, v_{sign}, v_{map}$ )
2:   for all global coefficients  $j$  do in parallel
3:      $\hat{s}_g[j] \leftarrow 0$ 
4:   end for
5:   for all elements  $e$  do in serial
6:     for all coefficients  $c$  do in parallel
7:        $i \leftarrow e * n_{coeffs} + c$ 
8:        $\hat{s}_g[v_{map}[i]] \leftarrow \hat{s}_g[v_{map}[i]] + v_{sign}[i] * \hat{s}_l[i]$ 
9:     end for
10:  end for
11: end procedure
```

---

our kernels with static memory allocation and use templating to account for the varying array lengths of varying polynomial orders. We found, that in this case, the compiler optimisation can implicitly deal with this situation and create an interleaved memory layout.

#### 4.2. Parallel element colouring and reduction operations

The local-to-global mapping, as given in Equation 27 is an operation that is not trivial to parallelise. This *gather*-type or reduction operation can be expressed as a double for-loop, as in algorithm 1. The inner loop over all local coefficients or modes can be implemented in parallel, the outer loop over all elements needs to be in serial to avoid a race-condition.

Although we found this algorithm to perform very well on the CPU, on the GPU performance was far more limited, since even at reasonable element polynomial orders, not even a single SMX can be fully utilised. To enable more parallelism we introduced an initial elemental colouring of the mesh. This ensures that no neighbouring elements are in the same colourgroup, and hence no write conflict of multiple local modes to global modes can occur. With this colouring in place, it is now safe to parallelise the for-loop over all elements within one colourgroup and introduce a third outermost serial loop over all colourgroups, as in algorithm 2. Our colouring algorithm uses a re-balancing step so we yield colourgroups that are typically within 25% of the average group size.

The algorithm requires a global synchronisation between the different colourgroups, which can be realised using two options on the GPU. This is

---

**Algorithm 2** Parallel local-to-global reduction algorithm using element colouring

---

```

1: procedure LOCALTOGLOBALPARALLEL( $\hat{s}_g, \hat{s}_g, v_{sign}, v_{map}, C$ )
2:   for all global coefficients  $j$  do in parallel
3:      $\hat{s}_g[j] \leftarrow 0$ 
4:   end for
5:   for all coloursets  $c \in C$  do in serial
6:     for all elements  $e \in c$  do in parallel
7:       for all coefficients  $c$  do in parallel
8:          $i \leftarrow e * n_{coeffs} + c$ 
9:          $\hat{s}_g[v_{map}[i]] \leftarrow \hat{s}_g[v_{map}[i]] + v_{sign}[i] * \hat{s}_l[i]$ 
10:      end for
11:    end for
12:  end for
13: end procedure

```

---

either launching a new kernel per colourgroup, or launching a single kernel for all colourgroups, but only utilise one SMX, since it is not possible to synchronise between SMX in a single kernel. The first option comes with a greater overhead, whereas the second utilises only a fraction of the GPU. For reasonably large mesh sizes, each colourgroup kernel receives enough work so that the first option will be more performant. Another challenge occurs as the inner parallel-loop is relatively skinny and the outer parallel loop is relatively fat, which does not map well to the GPU hardware. However, since the two loops are perfectly nested, a good compiler implementation will be able to collapse the two loops, resulting in better performance.

For all CPU executions we utilise the initial algorithm 1, for all GPU executions we utilise the mesh-colouring approach and the parallel algorithm 2.

#### 4.3. Ease of implementation and maintainability of the programming models

All three tested programming models, *Kokkos*, *OpenACC*, and *OpenACC* advertise their capability to incrementally parallelise and port legacy applications written in *C++* to multicore CPUs and GPUs. While this is certainly the case, we noted a few obstacles that indicate this process is not as seamless. In fact, care has to be taken when the initial code-base features many object-oriented programming structures. While they can be parallelised over

the cores of a CPU, they do not map to the *CUDA* programming model, that underlies all GPU backends of the considered programming models.

For example, with *CUDA* it is not possible to parallelise over an array of element objects, in which each object holds its data as member-variables and calls its overloaded member-functions, depending on which type of element it is. In this case, the functions and hierarchies of objects need to be de-tangled and implemented in a more straightforward *C*-syntax, in which the corresponding variables and functions for each element are taken care of explicitly. This can result in a complete re-writing of the involved compute-heavy kernels.

Similarly, encapsulated data-structures can not be mapped directly to the GPU memory. Instead the raw pointers would need to be passed, which in turn requires all functions further down the call-tree to be adapted for the plain data-type. *Kokkos* uses its own encapsulated data-arrays, that map between different memory spaces. It follows that, if in the case of *Nektar++*, encapsulated data-arrays are employed, a large initial implementation effort is necessary for all three programming models. If the legacy application already uses plain *C* arrays, *OpenMP* and *OpenACC* would be favourable.

The limitation of using dynamic memory allocation within the kernels also requires the introduction of templating for variable array lengths. This is a further step away from initial OOP-oriented *C++* code-bases.

When it comes to the actual syntax of specifying parallel loops, the three programming models are comparable, but the *Kokkos* syntax is slightly more complex than that of the *OpenMP* and *OpenACC* compiler directives. *Kokkos* allows for the possibility to explicitly specifying an array to be placed into shared memory, which can be rather cumbersome to optimise when dealing with different element polynomial orders and attempting to align this placement with the shared memory size restriction. *OpenACC* can place arrays into shared memory implicitly, whereas the tested *OpenMP-clang-ykt* implementation does not cover this functionality. In addition, the *OpenMP* and *OpenACC* syntaxes are so similar that a change between them should be straightforward.

## 5. Performance evaluation

Our test case involves the solution of the Helmholtz equation

$$\nabla^2 u + \lambda u = -(\lambda + 2\pi^2) \cos(\pi x) \cos(\pi y) \quad (35)$$

with  $\lambda = 2.5$ , on a square domain  $\Omega = \{(x, y) : -1 \leq x \leq 1, -1 \leq y \leq 1\}$ . On the sides  $x = -1$  and  $x = 1$  we apply the Dirichlet boundary condition

$$u = \cos(\pi x) \cos(\pi y) \tag{36}$$

and the Neumann boundary condition

$$\frac{\partial u}{\partial n} = \pi \cos(\pi x) \sin(\pi y) \tag{37}$$

is imposed on the sides  $y = -1$  and  $y = 1$ .

We compare the same algorithms as presented in Section 3 that solve this equation implemented with three different programming models: *Kokkos*, *OpenMP* and *OpenACC*. We use the same codebase for both CPU and GPU applications and only at compile time specify the relevant compile flags for each architecture. The only algorithmic difference between architectures is the parallelism in the local-to-global reduction operation discussed in Section 4.2.

The individual meshes for the spatial discretisation consist of triangular straight-sided elements, but different element polynomial orders between 2nd and 11th order. The number of degrees of freedom in each mesh is fixed at 4 million, so that the meshes consist of between 66k to 1975k elements from highest to lowest order respectively. We iterate the conjugate gradient loop until the solution is converged to a tolerance of  $10^{-9}$ .

In terms of compiler toolchains, for our *Kokkos* implementation we use the GNU compiler *gcc-5.4.0* with optimisation flag `-O3`. For the *OpenACC* implementation we use the PGI compiler version *pgc++-18.4* with optimisation flag `-O3`. The *OpenMP* implementation is compiled with the experimental *llvm-clang* compiler branch *clang-ykt* and optimisation flag `-O3`.

At the time of conducting this work the *clang-ykt* branch has not been feature complete. Even though the *OpenMP4.0* specifications support *parallel-reduce* operations, it has not been implemented, yet. As a workaround, we wrote an own kernel for parallel reductions, that is based on a two level-implementation.

The *llvm-ykt* compiler branch also did not support linking to external libraries at the time of conducting this work. Thus, our *OpenMP*-GPU implementation had to be completely decoupled from the *Nektar++* framework.

For a better performance on the CPU we set the environmental variable `OMP_PROC_BIND=true`. For the GPU versions we always use *CUDA8.0*.

### 5.1. Comparison of Kokkos, OpenMP and OpenACC on multi-core architectures

For the multi-core system, we use a Intel E5-2690 v0 CPU with 16 cores and a base frequency of 2.9GHz and 64GB RAM with 1600Hz. The peak FP64 performance considering FMA and AVX is 371.2GFLOPs, the maximum memory bandwidth is 51.2GB/s, the thermal design point (TDP) is 135W.

As a benchmark, we also present the results obtained with our basic *MPI* implementation within *Nektar++*. It uses the same algorithms, but realises its parallel execution using domain decomposition.

#### 5.1.1. Strong scaling exercise

As is customary for multi-core CPU applications, we conduct a strong scaling exercise between 1 and 16 threads on a mesh consisting of 7th-order triangles with 116k elements. This allows us to evaluate the performance of each model and identify potential performance overheads. The results are given in Figure 1. Here we see that the *OpenMP* and *OpenACC* implementations perform better than the *Kokkos* implementation; however their relative performance varies with polynomial order and will be clearer to evaluate in this context in Section 5.1.2. More important is the observation that all three implementations scale very similarly with increasing numbers of threads and do not reach the almost ideal scaling of our *MPI* implementation. This is probably due to the reduction operations that do not allow us to fully utilise all threads continuously.

#### 5.1.2. Efficiency of different polynomial orders

For this test series we use meshes with elements of varying polynomial orders where the degree of freedom count is fixed across the polynomial orders. We utilise all 16 threads on our multi-core CPU. The runtimes for one iteration of the conjugate gradient solve of all four applications scaled by the exact number of DOFs are given in the upper left subplot of Figure 2. The runtimes improve with increasing element order up to about 6th order and then plateau. The low element orders have a lower FLOP count per DOF, but also a lower FLOP-per-byte ratio or arithmetic intensity. As our applications are memory-bound, they do not benefit from a lower FLOP count *per se*. Overall, the *MPI* implementation is the fastest across the range of polynomial orders, but closely matched by the *OpenMP* and the *OpenACC*

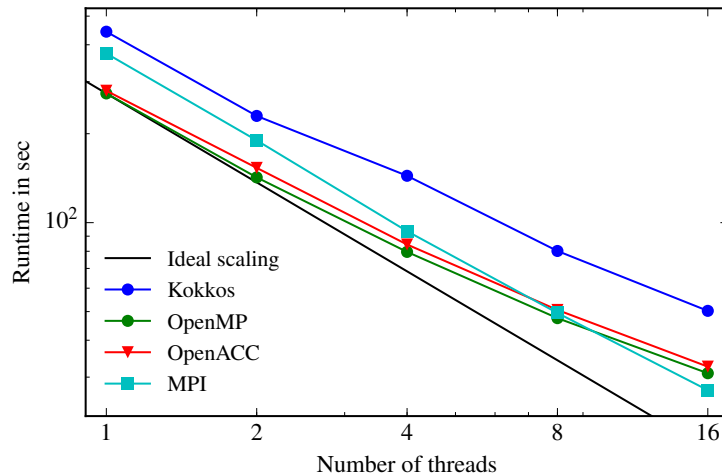


Figure 1: Strong scaling on 16-core CPU

versions. The *Kokkos* version gives the worst performance generally, with around twice the runtime.

The remaining three subplots of Figure 2 present the runtime-splits between the five different kernels of the application, as discussed in Section 4. The elemental Helmholtz operation consumes about 50% of the overall runtimes. The local-to-global reduction also shows bad performance for the lower elemental orders, as the involved arrays are too short in this case. Overall, the splitting of runtime between the *Kokkos*, *OpenMP*, and *OpenACC* is comparable.

### 5.1.3. Detailed performance evaluation

In order to evaluate how well the underlying CPU hardware is utilised we consider two metrics, the achieved double precision (FP64) FLOPs and the DRAM memory bandwidth. The FLOPs that the different kernels have achieved are given in Figure 6. The gather and scatter operations, as well as the conjugate gradient reductions and for-loops perform very poorly, as they are strongly memory bound. Only very few operations are performed on each data element, which is read and written back to DRAM memory. Only for the parallel reduction kernel data can be re-used from higher cache levels, leading to faster loading to registers. The gather and scatter operations perform especially poor, as the data is usually not read and written in continuous cache blocks. The Helmholtz operation is the critical kernel and achieves

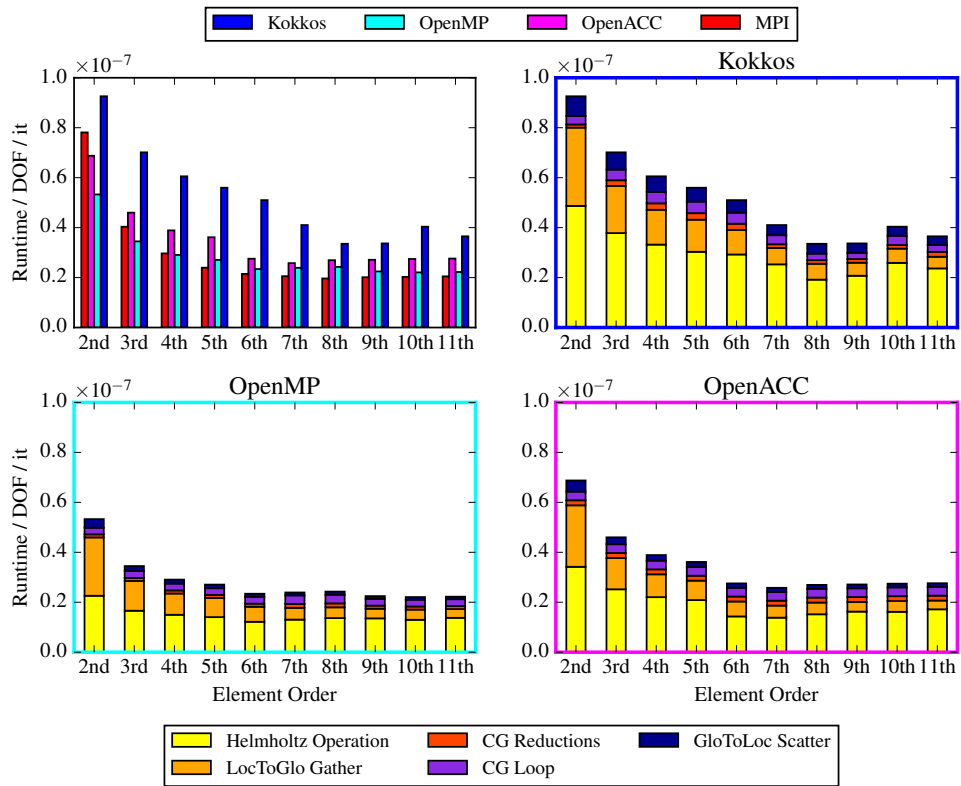


Figure 2: Runtime comparison of different CPU implementations over a range of elemental orders



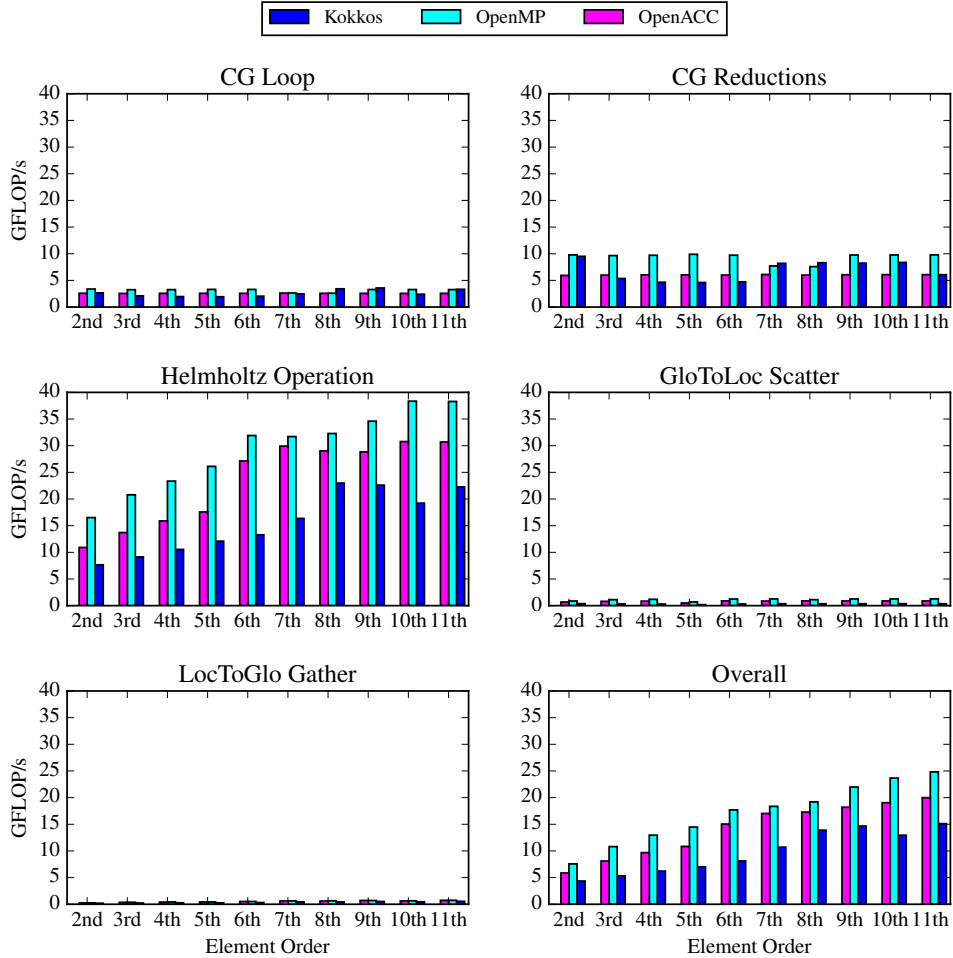


Figure 3: Achieved FP64 FLOPs for the main kernels on 16-core CPU

much better performance of up to 38GFLOP/s for the *OpenMP* version, 31GFLOP/s for the *OpenACC* version, but only 23GFLOP/s for the *Kokkos* version. For the *OpenMP* version, this translates to utilising more than 10% of peak FLOPs. The trend for higher FLOPs with higher element orders is very clear, as it requires more compute operations per loaded byte, i.e. a higher arithmetic intensity. The overall FLOPs performance is dominated by the Helmholtz operation kernel, but pulled down by about one third by the less performant kernels.

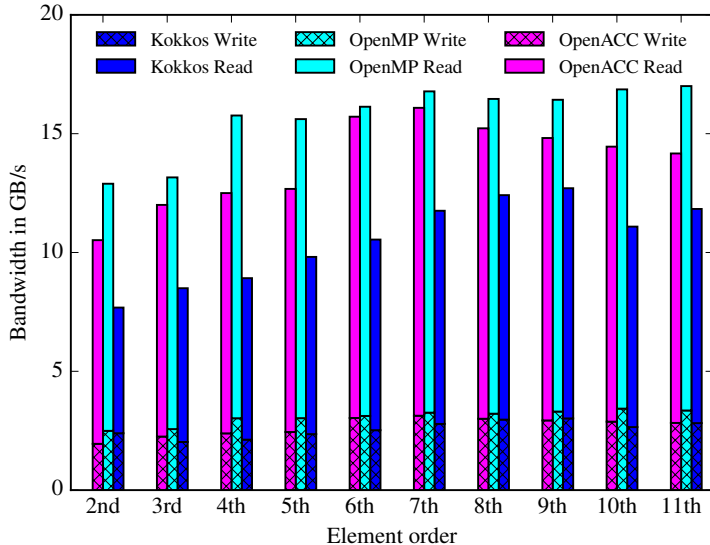


Figure 4: DRAM memory bandwidth on 16-core CPU

The DRAM memory bandwidth is given in Figure 4. Generally it can be seen from the data that especially the *OpenMP* and to a slightly lower degree also the *OpenACC* implementation utilise the memory more efficiently than the *Kokkos* implementation. As a tendency, higher polynomial orders are processed more efficiently, as can be expected due to the larger arrays. The *OpenMP* implementation achieves a combined *read* and *write* memory bandwidth of 17GB/s, which amounts to about 33% of the theoretical bandwidth.

## 5.2. Comparison of Kokkos, OpenMP and OpenACC on GPU architectures

As the GPU system we use a Nvidia Tesla P100 with 60 streaming multi-processors (SMX), a peak FP64 performance of 5304GFLOPS, a maximum bandwidth of 549GB/s, and a TDP of 300W.

### 5.2.1. Efficiency of different polynomial orders

We employ the same test methodology as in the multi-core CPU case in Section 5.1.2. The runtimes for one iteration scaled by the DOFs is given in Figure 5. Across the whole range of elemental orders, the *OpenACC* implementation is the most performant, while the variation between elemental

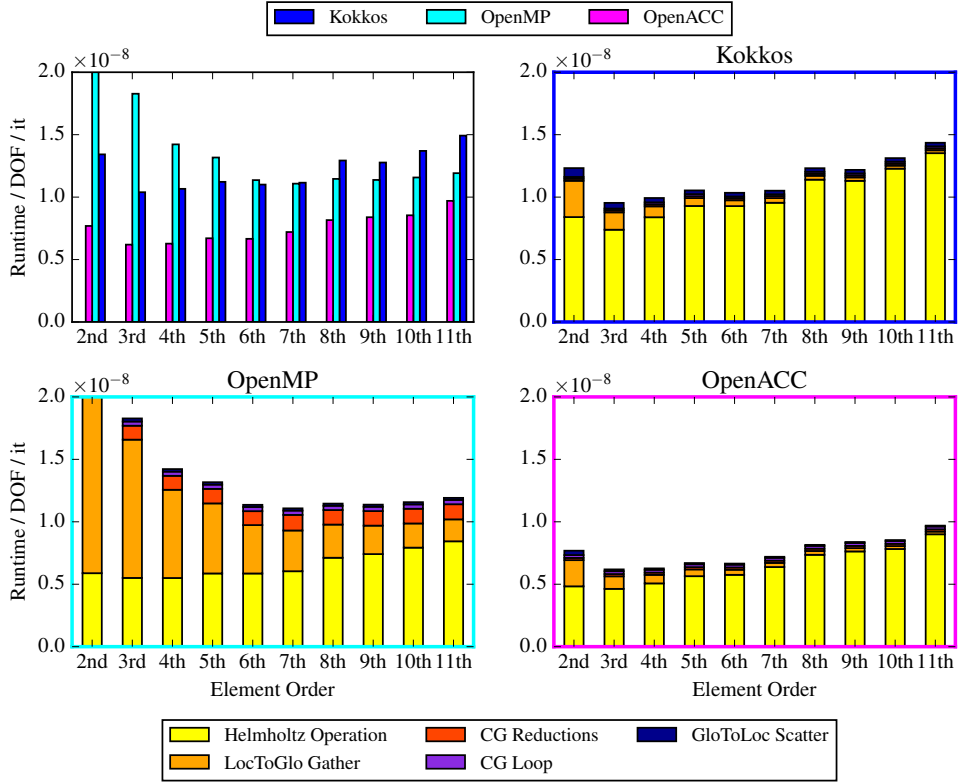


Figure 5: Runtime comparison of different GPU implementations over a range of elemental orders on P100

orders is low. The runtime splits show a much larger proportion of time spent in the elemental Helmholtz operation, that varies between about 65% for 2nd order and 90% for 11th order. The performance for the *OpenMP* implementation is only about 25% worse for high polynomial orders, but suffers greatly for low polynomial orders. It becomes clear that primarily the bad performance of the local-to-global reduction is responsible and only secondarily the lower performance of our custom-written CG-reduction kernel. The embarrassingly parallel elemental Helmholtz kernel is *on par* with the *OpenACC* version. The performance of the *Kokkos* implementation is about 40% worse than the *OpenACC* implementation across the range of elemental orders. Examining the runtime splits of the *Kokkos* version, no individual kernel can be pointed out that would be responsible, indicating a general performance issue with this version.

### 5.2.2. Detailed performance evaluation

Runtimes alone do not inform about how well the implementation is utilising the underlying hardware. To this end we first provide the achieved double precision (FP64) FLOP/s in Figure 6. Comparing the three different programming models, it can be seen again that *Kokkos* does not achieve the same performance as *OpenACC* and *OpenMP* for the Helmholtz operation kernel and hence for the overall application. Furthermore, the poor performance of the *OpenMP* reduction/gather kernels becomes obvious. In general, the achieved FLOP/s translate to utilising actually less than 2% of the theoretical FP64 performance, which indicates that the memory bandwidth is actually the performance bottleneck.

To prove this, we provide the DRAM memory bandwidth in Figure 7. Here it can be seen, that our implementations operate very close to the theoretical maximum across all elemental orders for the CG-Loop, the CG-reductions and the global-to-local scatter kernels. The most important kernel, the Helmholtz operation, though, is not fully utilising the DRAM bandwidth indicating that either the bandwidth of a higher level memory or cache level is limiting the overall performance or the memory latency results in a stall. The three programming models perform very similar, just the *OpenMP* application is falling off slightly.

Comparing the utilisation or bandwidth of the global (Figure 8) against the local and shared memory (Figure 9) for the Helmholtz operation kernel, we can see that the *Kokkos* kernel is utilising more of the slower global memory (residing on DRAM, L2 and L1 cache) than the faster local and shared memory that resides only on the L1 cache. This could explain the performance deficit of the *Kokkos* implementation. It has to be acknowledged, though, that a better performance could be achieved if tuning the Helmholtz operation kernel for each elemental order by explicitly specifying the utilisation of L1 cache for certain arrays. This tuning approach however defeats the idea of a performance portable programming model.

It can be concluded, that our applications are heavily memory bound. It shows, that just porting a legacy CPU-optimised algorithm can not fully utilise modern GPUs. More effort to reduce the memory dependency of the algorithms and hence increase the arithmetic intensity are necessary.

### 5.3. Cost comparison between different architectures

We use system acquisition costs to identify the most efficient architectures in terms of run-time and operational costs.

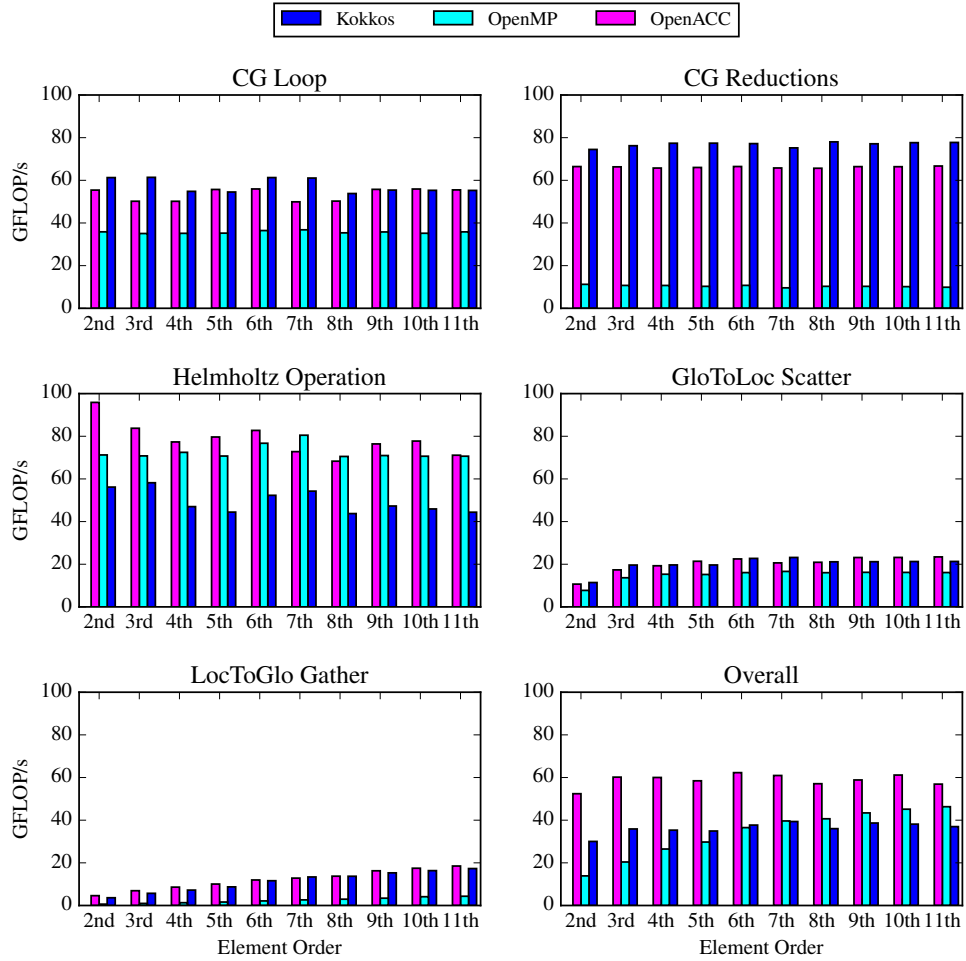


Figure 6: Achieved FP64 flops for the main kernels

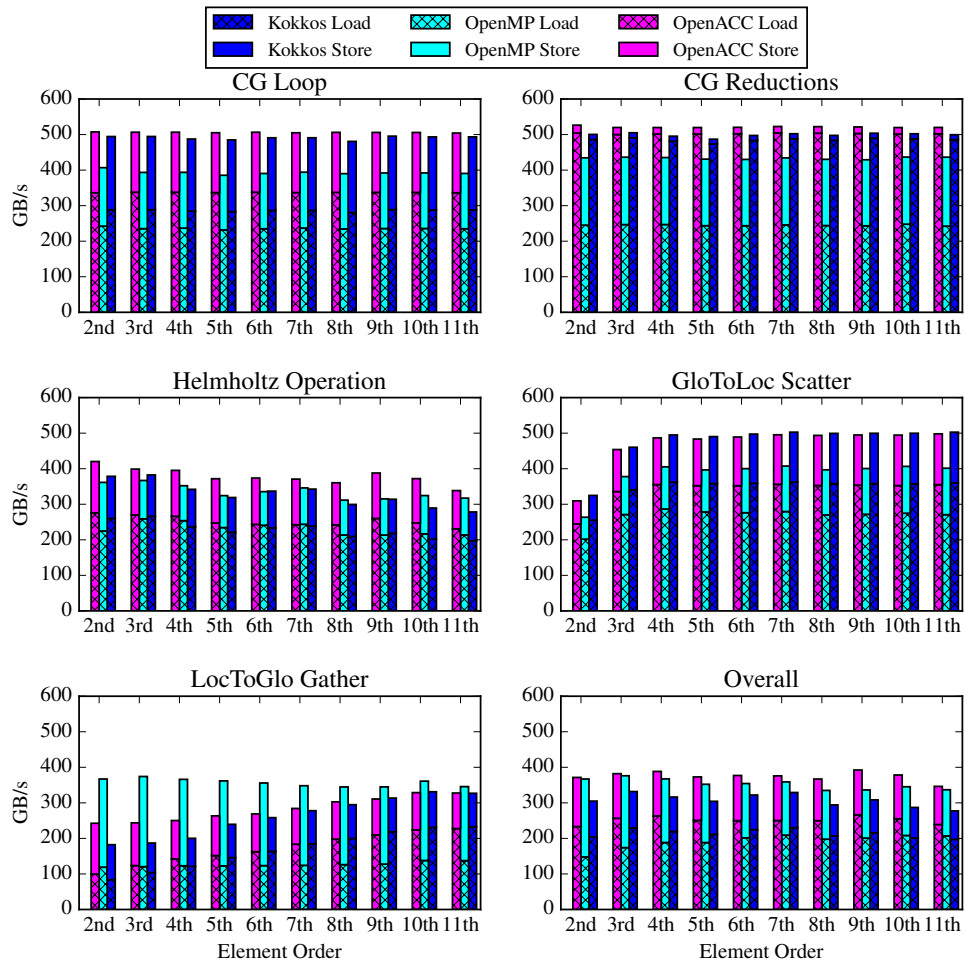


Figure 7: DRAM memory bandwidth for the main kernels

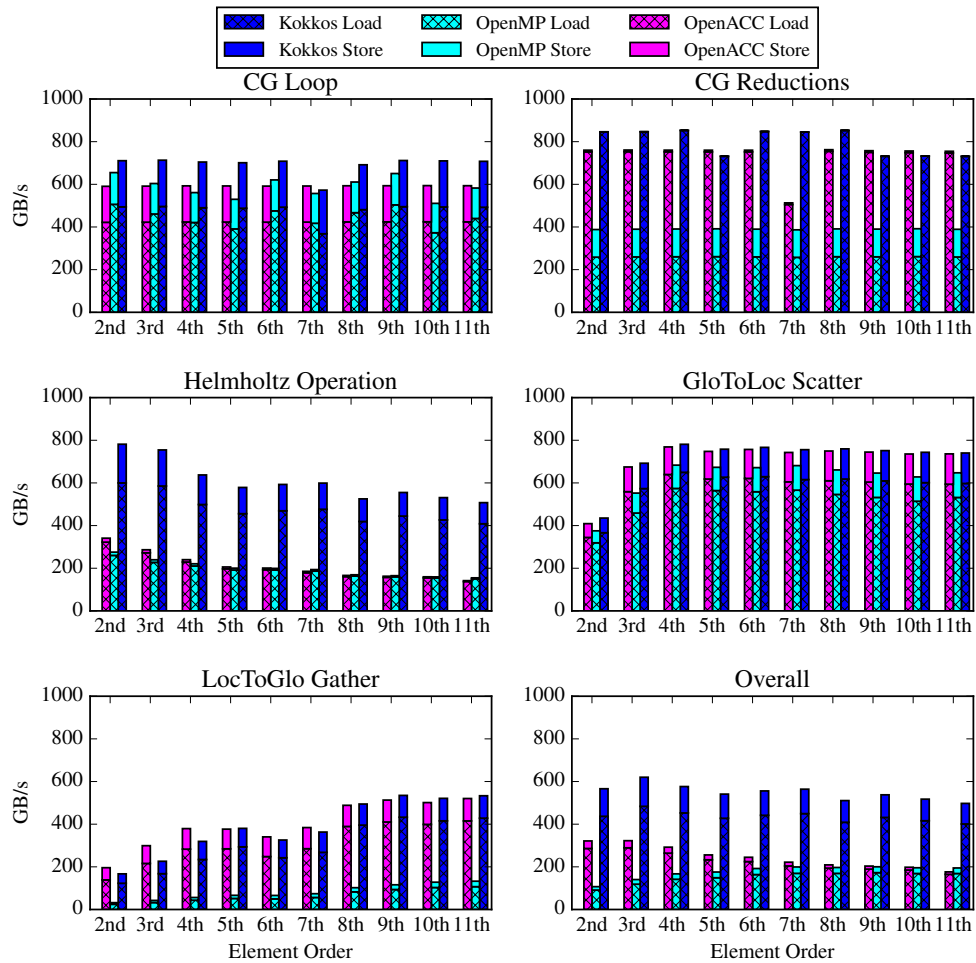


Figure 8: Global memory bandwidth for the main kernels

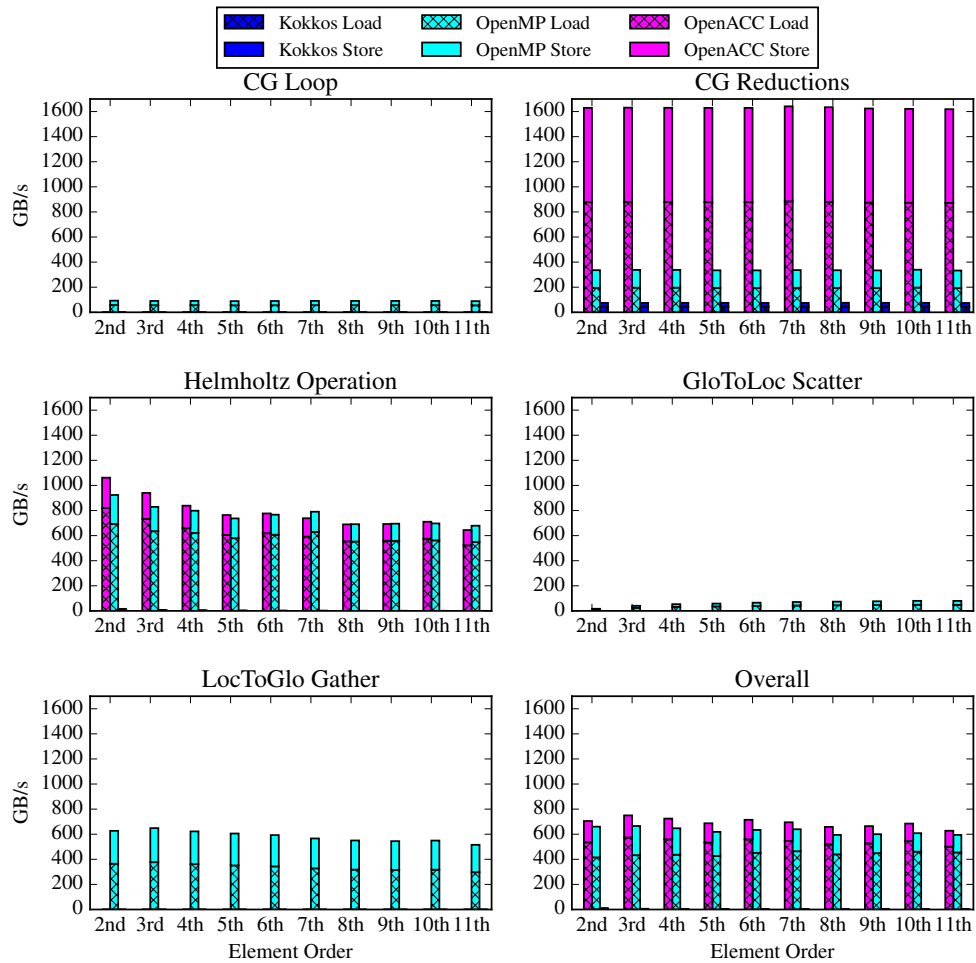


Figure 9: Local and shared memory bandwidth for the main kernels



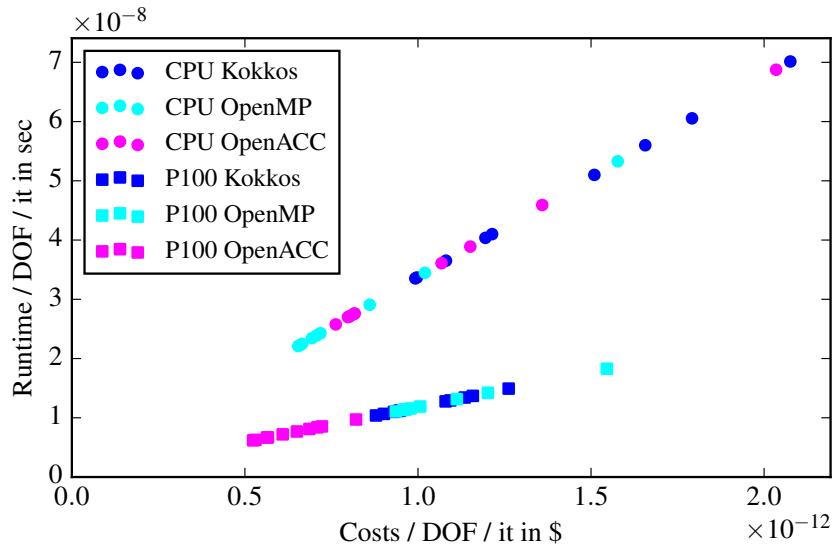


Figure 10: Cost vs time comparison of the CPU and GPU system

The acquisition costs are the sum of the CPU or GPU by itself plus a hosting system. We estimate the current price of the Nvidia Tesla P100 GPU as \$6000, the Intel E5-2690 v0 CPU with 2 sockets or 16 cores as \$800, and the hosting system as \$2000. We use a time frame of 3 years for complete linear depreciation, to obtain hourly prices for both the complete CPU and GPU system. The costs of code execution is then the product of run-time with hourly system costs. This follows the approach presented in [32]. The result is given in Figure 10. It illustrates very well the benefit of employing GPU systems for the presented type of algorithms, as both the runtime and costs are lower on the P100, the costs by 20%, the runtimes by 72%, considering the most performant cases. Considering only the two *OpenMP* implementations, though, the costs of utilising the P100 GPU is higher than for the CPU, although it is faster. In terms of both cost and time performance metrics, the *OpenACC* version for the GPU is the most beneficial. This result would only become more pronounced, when optimising the algorithms further and replace pre-computations and loading from memory with on-the-fly calculations to reduce the required memory bandwidth that so far slows down the GPU implementations to a higher degree than the CPU implementations.

## 6. Conclusions

There are two main conclusions to be drawn from this study.

Firstly, when comparing three different performance-portable programming models, namely *Kokkos*, *OpenACC* and *OpenMP* using open-access compilers in a high-order FEM setting, we conclude that the choice of programming model is still a complex question.

All three programming models require very similar implementational efforts. There are syntactical differences, which are minor between *OpenMP* and *OpenACC* (and would enable a seamless switch between the two) and slightly deeper compared with *Kokkos*, mostly due to the requirement of using *Kokkos*-specific arrays. *Kokkos* further allows addressing many aspects of performance tuning which are connected with native *CUDA* applications, allowing more individual trade-offs between portability and performance. The main current drawback of the *OpenMP* model is the lack of robust open-access compiler implementations. The best compiler currently available, which is the *llvm-ykt* branch, still lacks a parallel reduction functionality, has no linking support for external libraries, and showed bad performance of the gather-type operation, probably due to the implementation of the *collapsing*-clause and the implicit compiler choice for block-sizes. The *OpenACC* model has shown the best overall performance results, yielding the fastest runtimes on Nvidia GPUs, and only slightly slower runtimes on Intel CPUs than *OpenMP*.

Secondly, this study has helped us to identify the steps to be undertaken in order to port a legacy CPU code to modern multi-core CPUs and many-core GPUs and achieve good performance. The mapping of the algorithmic parallelisation to the hardware parallelisation, together with the right data layout, plays a fundamental role in achieving reasonable performance on modern hardware. Here we have employed a very fine grained parallelism in which we process a single mesh element on each thread. This requires the processing of elements to be vectorised either on the AVX vector lanes or over the *CUDA* threads of a warp. This vectorisation in turn necessitates the use of interleaved data structures and custom BLAS functions to operate on them. Reduction operations can become a major bottleneck in massively parallel applications. We employ mesh colouring to counteract this aspect and regain more parallelism.

The performance of our three applications is very reasonable on multi-core CPUs, but lacks the ability to fully leverage the arithmetic potential of GPUs.

To do so, it is not enough to translate the traditional algorithms, that often just minimise FLOP count, as we did in this study. Instead of pre-computing data and loading it when required, more data should be recomputed on the fly, as it can often be faster than loading a large array from DRAM memory, when latencies are of the order of  $10^4$  cycles. For example, the framework *OpenSBLI* [20] offers the option to variably execute either or a mix of these two approaches.

These observations show that performance portability might not be realisable in practice just by employing a performance portable programming model. The algorithmic choices between current CPUs and GPUs that have to be made now are fundamentally different. To deal with this increasing code complexity, auto-tuning frameworks and just-in-time compilation, that can make algorithmic choices based on the underlying architecture at runtime, should be considered.

### *Acknowledgements*

JE acknowledges the constructive discussions with Niki Loppi that led to the development of the strided kernels. JE also acknowledges the support through the President’s Scholarship of Imperial College London. MV acknowledges support from the EU Horizon 2020 project ExaFLOW (Grant no. 671571). DM and JP acknowledge support from the EPSRC under grant EP/R029423/1. The Quadro P5000 GPU used for this work has been kindly donated by the NVIDIA corporation.

### **References**

- [1] J.-E. W. Lombard, D. Moxey, S. J. Sherwin, J. F. A. Hoessler, S. Dhandapani, M. J. Taylor, Implicit large-eddy simulation of a wingtip vortex, *AIAA Journal* 54 (2) (2016) 506–518. doi:10.2514/1.J054181.
- [2] H. Carter Edwards, C. R. Trott, D. Sunderland, Kokkos: Enabling manycore performance portability through polymorphic memory access patterns, *Journal of Parallel and Distributed Computing* 74 (12) (2014) 3202–3216. doi:10.1016/j.jpdc.2014.07.003.
- [3] OpenMP 4.5 Specifications (2015).  
URL <http://www.openmp.org/wp-content/uploads/openmp-4.5.pdf>

- [4] OpenACC Programming and Best Practices Guide (2015).  
URL [http://www.openacc.org/sites/default/files/OpenACC\\_Programming\\_Guide\\_0.pdf](http://www.openacc.org/sites/default/files/OpenACC_Programming_Guide_0.pdf)
- [5] J. J. Dongarra, P. Luszczek, A. Petite, The LINPACK benchmark: Past, present and future, *Concurrency Computation Practice and Experience* 15 (9) (2003) 803–820. doi:10.1002/cpe.728.
- [6] M. A. Heroux, D. W. Doerfler, P. S. Crozier, J. M. Willenbring, H. C. Edwards, A. Williams, M. Rajan, E. R. Keiter, H. K. Thornquist, R. W. Numrich, Improving performance via Mini-applications, Tech. Rep. September, Sandia National Labs (2009).  
URL <https://mantevo.github.io/pdfs/MantevoOverview.pdf>
- [7] S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, S. H. Lee, K. Skadron, Rodinia: A benchmark suite for heterogeneous computing, in: *Proceedings of the 2009 IEEE International Symposium on Workload Characterization, IISWC 2009*, IEEE, 2009, pp. 44–54. doi:10.1109/IISWC.2009.5306797.
- [8] A. Danalis, G. Marin, C. McCurdy, J. S. Meredith, P. C. Roth, K. Spafford, V. Tipparaju, J. S. Vetter, The Scalable Heterogeneous Computing (SHOC) Benchmark Suite Categories and Subject Descriptors, in: *Proceedings of the 3rd Workshop on General-Purpose Computation on Graphics Processing Units, 2010*, pp. 63–74. arXiv:1107.1714v2, doi:10.1145/1735688.1735702.
- [9] J. A. Stratton, C. Rodrigues, I.-J. Sung, N. Obeid, L.-W. Chang, N. Ansari, G. D. Liu, W.-m. W. Hwu, Parboil: A Revised Benchmark Suite for Scientific and Commercial Throughput Computing, Tech. rep., University of Illinois at Urbana-Champaign (2012).  
URL <http://impact.crhc.illinois.edu/parboil.php>
- [10] C. D. Cantwell, D. Moxey, A. Comerford, A. Bolis, G. Rocco, G. Mengaldo, D. De Grazia, S. Yakovlev, J. E. Lombard, D. Ekelschot, B. Jordi, H. Xu, Y. Mohamied, C. Eskilsson, B. Nelson, P. Vos, C. Biotto, R. M. Kirby, S. J. Sherwin, Nektar++: An open-source spectral/hp element framework, *Computer Physics Communications* 192 (2015) 205–219. doi:10.1016/j.cpc.2015.02.008.

- [11] G. E. Karniadakis, M. Israeli, S. A. Orszag, High-Order Splitting Methods for the Incompressible Navier-Stokes Equations, *Journal of Computational Physics* 97 (2) (1991) 414–443. doi:10.1016/0021-9991(91)90007-8.
- [12] G. Karniadakis, S. Sherwin, *Spectral/hp Element Methods for Computational Fluid Dynamics*, 2nd Edition, Oxford University Press, 2005. doi:10.1093/acprof:oso/9780198528692.001.0001.
- [13] F. D. Witherden, A. M. Farrington, P. E. Vincent, PyFR: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach, *Computer Physics Communications* 185 (11) (2014) 3028–3040. doi:10.1016/j.cpc.2014.07.011.
- [14] M. Kronbichler, K. Kormann, Fast Matrix-Free Evaluation of Discontinuous Galerkin Finite Element Operators, *ACM Transactions on Mathematical Software* 45 (3) (2019) 1–40. doi:10.1145/3325864.
- [15] P. Bastian, C. Engwer, J. Fahlke, M. Geveler, D. GÖddeke, O. Iliev, O. Ippisch, R. Milk, J. Mohring, S. Müthing, M. Ohlberger, D. Ribbrock, S. Turek, Hardware-Based Efficiency Advances in the EXA-DUNE Project, in: H.-J. Bungartz, P. Neumann, W. E. Nagel (Eds.), *Software for Exascale Computing - SPPEXA 2013-2015*, Springer International Publishing, Cham, 2016, pp. 3–23. doi:10.1007/978-3-319-40528-5\_1.
- [16] Intel Corporation, *Using AVX Without Writing AVX Code* (2012).  
URL <https://software.intel.com/en-us/articles/using-avx-without-writing-avx-code>
- [17] Nvidia, *CUDA8.0 Release Note* (2016).  
URL [http://docs.nvidia.com/cuda/pdf/CUDA\\_Toolkit\\_Release\\_Notes.pdf](http://docs.nvidia.com/cuda/pdf/CUDA_Toolkit_Release_Notes.pdf)
- [18] Khronos Group, *OpenCL Overview* (2019).  
URL <https://www.khronos.org/opencv/>
- [19] F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. T. McRae, G.-T. Bercea, G. R. Markall, P. H. J. Kelly, Firedrake: automating the finite element method by composing abstractions, *ACM*

Transactions on Mathematical Software 43 (3) (2015) 24:1 – 24:27.  
doi:10.1145/2998441.

- [20] C. T. Jacobs, S. P. Jammy, N. D. Sandham, OpenSBLI: A framework for the automated derivation and parallel execution of finite difference solvers on a range of computer architectures, *Journal of Computational Science* 18 (2017) 12–23. doi:10.1016/j.jocs.2016.11.001.
- [21] The Open Group, POSIX Threads (1997).  
URL <http://pubs.opengroup.org/onlinepubs/007908799/xsh/threads.html>
- [22] R. Hornung, H. Jones, J. Keasler, R. Neely, A. Kunen, O. Pearce, RAJA Overview, Tech. Rep. LLNL-TR-677453, Lawrence Livermore National Laboratory (2015).  
URL <https://github.com/LLNL/RAJA>
- [23] D. S. Medina, A. St-Cyr, T. Warburton, OCCA: A unified approach to multi-threading languages, *ArXiv* (2014) 1–25.  
URL <http://arxiv.org/abs/1403.0968>
- [24] Khronos Group, SYCL Provisional Specification (2016).  
URL <https://www.khronos.org/registry/sycl/specs/sycl-2.2.pdf>
- [25] S. A. Orszag, Spectral methods for problems in complex geometries, *Journal of Computational Physics* 37 (1) (1980) 70–92. doi:10.1016/0021-9991(80)90005-4.
- [26] D. Moxey, C. D. Cantwell, R. M. Kirby, S. J. Sherwin, Optimising the performance of the spectral/hp element method with collective linear algebra operations, *Computer Methods in Applied Mechanics and Engineering* 310 (2016) 628–645. doi:10.1016/j.cma.2016.07.001.
- [27] A. Heinecke, G. Henry, M. Hutchinson, H. Pabst, LIBXSMM: Accelerating Small Matrix Multiplications by Runtime Code Generation, *International Conference for High Performance Computing, Networking, Storage and Analysis, SC* (2017) 981–991doi:10.1109/SC.2016.83.
- [28] C. D. Cantwell, S. J. Sherwin, R. M. Kirby, P. H. J. Kelly, From h to p efficiently: Strategy selection for operator evaluation on hexahedral

- and tetrahedral elements, *Computers and Fluids* 43 (1) (2011) 23–28. doi:10.1016/j.compfluid.2010.08.012.
- [29] T. Dong, A. Haidar, P. Luszczek, S. Tomov, A. Abdelfattah, J. Dongarra, MAGMA Batched: A Batched BLAS Approach for Small Matrix Factorizations and Applications on GPUs, ICL Tech Report 08/2016, Innovative Computing Laboratory, University of Tennessee, Knoxville, TN, 37996 (2016).
- [30] Y. Shi, U. N. Niranjan, A. Anandkumar, C. Cecka, Tensor Contractions with Extended BLAS Kernels on CPU and GPU, in: *Proceedings - 23rd IEEE International Conference on High Performance Computing, HiPC 2016, 2017*, pp. 193–202. doi:10.1109/HiPC.2016.031.
- [31] G. Alzetta, D. Arndt, W. Bangerth, V. Boddu, B. Brands, D. Davydov, R. Gassmoeller, T. Heister, L. Heltai, K. Kormann, M. Kronbichler, M. Maier, J.-P. Pelteret, B. Turcksin, D. Wells, The deal.II Library, Version 9.0, *Journal of Numerical Mathematics*.  
URL <https://doi.org/10.1515/jnma-2018-0054>
- [32] J. Eichstädt, M. Green, M. Turner, J. Peiró, D. Moxey, Accelerating high-order mesh optimisation with an architecture-independent programming model, *Computer Physics Communications* 229 (2018) 36–53. doi:10.1016/j.cpc.2018.03.025.