







Efficient exascale discretizations: High-order finite element methods

Tzanio Kolev¹, Paul Fischer^{2,3,4}, Misun Min², Jack Dongarra⁵,
Jed Brown⁶, Veselin Dobrev¹, Tim Warburton⁷,
Stanimire Tomov⁵, Mark S Shephard⁸, Ahmad Abdelfattah⁵,
Valeria Barra⁶, Natalie Beams⁵, Jean-Sylvain Camier¹,
Noel Chalmers⁹, Yohann Dudouit¹, Ali Karakus¹⁰,
Ian Karlin¹, Stefan Kerkemeier², Yu-Hsiang Lan²,
David Medina¹¹, Elia Merzari^{2,12}, Aleksandr Obabko²,
Will Pazner¹, Thilina Rathnayake³, Cameron W Smith⁵,
Lukas Spies³, Kasia Swirydowicz¹³, Jeremy Thompson⁶,
Ananias Tomboulides^{2,14} and Vladimir Tomov¹

The International Journal of High
Performance Computing Applications
1–26

© The Author(s) 2021

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/10943420211020803

journals.sagepub.com/home/hpc



Abstract

Efficient exploitation of exascale architectures requires rethinking of the numerical algorithms used in many large-scale applications. These architectures favor algorithms that expose ultra fine-grain parallelism and maximize the ratio of floating point operations to energy intensive data movement. One of the few viable approaches to achieve high efficiency in the area of PDE discretizations on unstructured grids is to use matrix-free/partially assembled high-order finite element methods, since these methods can increase the accuracy and/or lower the computational time due to reduced data motion. In this paper we provide an overview of the research and development activities in the Center for Efficient Exascale Discretizations (CEED), a co-design center in the Exascale Computing Project that is focused on the development of next-generation discretization software and algorithms to enable a wide range of finite element applications to run efficiently on future hardware. CEED is a research partnership involving more than 30 computational scientists from two US national labs and five universities, including members of the Nek5000, MFEM, MAGMA and PETSc projects. We discuss the CEED co-design activities based on targeted benchmarks, miniapps and discretization libraries and our work on performance optimizations for large-scale GPU architectures. We also provide a broad overview of research and development activities in areas such as unstructured adaptive mesh refinement algorithms, matrix-free linear solvers, high-order data visualization, and list examples of collaborations with several ECP and external applications.

Keywords

High-performance computing, co-design, high-order discretizations, unstructured grids, PDEs

¹ Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA, USA

² Mathematics and Computer Science, Argonne National Laboratory, Lemont, IL, USA

³ Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, USA

⁴ Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA

⁵ Innovative Computing Laboratory, University of Tennessee, Knoxville, TN, USA

⁶ Department of Computer Science, University of Colorado, Boulder, CO, USA

⁷ Department of Mathematics, Virginia Tech, Blacksburg, VA, USA

⁸ Scientific Computation Research Center, Rensselaer Polytechnic Institute, Troy, NY, USA

⁹ AMD Research, Austin, TX, USA

¹⁰ Mechanical Engineering Department, Middle East Technical University, Ankara, Turkey

¹¹ Occalytics LLC, Weehawken, NJ, USA

¹² Department of Nuclear Engineering, Penn State, PA, USA

¹³ Pacific Northwest National Laboratory, WA, USA

¹⁴ Department of Mechanical Engineering, Aristotle University of Thessaloniki, Greece

Corresponding author:

Tzanio Kolev, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA 94550, USA.

Email: tzanio@llnl.gov

1. Introduction

Efficient exploitation of exascale architectures requires rethinking of the numerical algorithms for solving partial differential equations (PDEs) on general unstructured grids. New architectures, such as general purpose graphics processing units (GPUs) favor algorithms that expose ultra fine-grain parallelism and maximize the ratio of floating point operations to energy intensive data movement.

Many large-scale PDE-based applications employ unstructured finite element discretization methods, where practical efficiency is measured by the accuracy achieved per unit computational time. One of the few viable approaches to achieve high performance in this case is to use matrix-free high-order finite element methods, since these methods can both increase the accuracy and/or lower the computational time due to reduced data motion. To achieve this efficiency, high-order methods use mesh elements that are mapped from canonical reference elements (hexahedra, wedges, pyramids, tetrahedra) and exploit, where possible, the tensor-product structure of the canonical mesh elements and finite element spaces. Through matrix-free partial assembly, the use of canonical reference elements enables substantial cache efficiency and minimizes extraneous data movement in comparison to traditional low-order approaches.

The Center for Efficient Exascale Discretizations (CEED) is a focused team effort within the U.S. Department of Energy (DOE) Exascale Computing Project (ECP) that aims to develop the next-generation discretization software and algorithms to enable a wide range of finite element applications to run efficiently on future hardware. CEED is a research partnership involving more than 30 computational scientists from two DOE labs and five universities, including members of the Nek5000, MFEM, MAGMA and PETSc projects (Anderson et al., 2020; Balay et al., 2019; MAGMA; Medina et al., 2014; MFEM; Nek5000). This article provides an overview of the co-design research and development activities in the CEED project based on targeted benchmarks, miniapps and discretization libraries. We also discuss several examples of collaborations with ECP, including ExaSMR, MARBL, Urban, and ExaWind, as well as external applications.

Following the ECP co-design philosophy, CEED is positioned as a computational motif hub between applications, hardware vendors and software technologies projects. As such, the main objectives of the project are to:

1. Help applications leverage future architectures by providing them with state-of-the-art discretization algorithms that better exploit the hardware and deliver a significant performance gain over conventional low-order methods.
2. Collaborate with hardware vendors and software technologies projects to utilize and impact the upcoming exascale hardware and its software stack through CEED-developed proxies and miniapps.

3. Provide an efficient and user-friendly unstructured PDE discretization component for the upcoming exascale software ecosystem.

To address these objectives, the center's co-design efforts are organized in four interconnected research and development thrusts, focused on the following computational motifs and their performance on exascale hardware:

1.1. PDE-based simulations on unstructured grids

CEED is producing a range of software products supporting general finite element algorithms on triangular, quadrilateral, tetrahedral and hexahedral meshes. We target the whole de Rham complex: H^1 , $H(\text{curl})$, $H(\text{div})$ and L^2/DG spaces and discretizations, including conforming and non-conforming unstructured adaptive mesh refinement (AMR).

1.2. High-order/spectral finite elements

Our algorithms and software come with comprehensive high-order support: we provide efficient matrix-free operator evaluation for any order space on any order mesh, including high-order curved meshes and all geometries in the de Rham complex. The CEED software also includes optimized assembly support for low-order methods.

The rest of the paper is organized as follows. In Section 2 we describe our co-design goals and organization. The needs of a complete high-order software ecosystem are then reviewed in Section 3. The CEED high-order benchmarks designed to test and compare the performance of high-order codes are described in Section 4. CEED is developing a variety of miniapps encapsulating key physics and numerical kernels of high-order applications. These are described in Section 5. We deliver performant algorithms to applications via discretization libraries both at low-level, see libCEED described in Section 6, and high-level, see MFEM and Nek described in Section 7. The impact of these CEED-developed technologies in several applications is illustrated in Section 8, followed by conclusions in Section 9.

2. Co-design

CEED's co-design activities are organized in four R&D thrusts described below.

2.1. Applications thrust

The goal of CEED's Applications thrust is to impact a wide range of ECP application teams through focused one-on-one interactions, facilitated by CEED application liaisons, as well as through one-to-many interactions, based on the development of easy-to-use discretization libraries for high-order finite element methods.

2.2. Hardware thrust

The goal of CEED's Hardware thrust is to build a two-way (pull-and-push) collaboration with vendors, where the

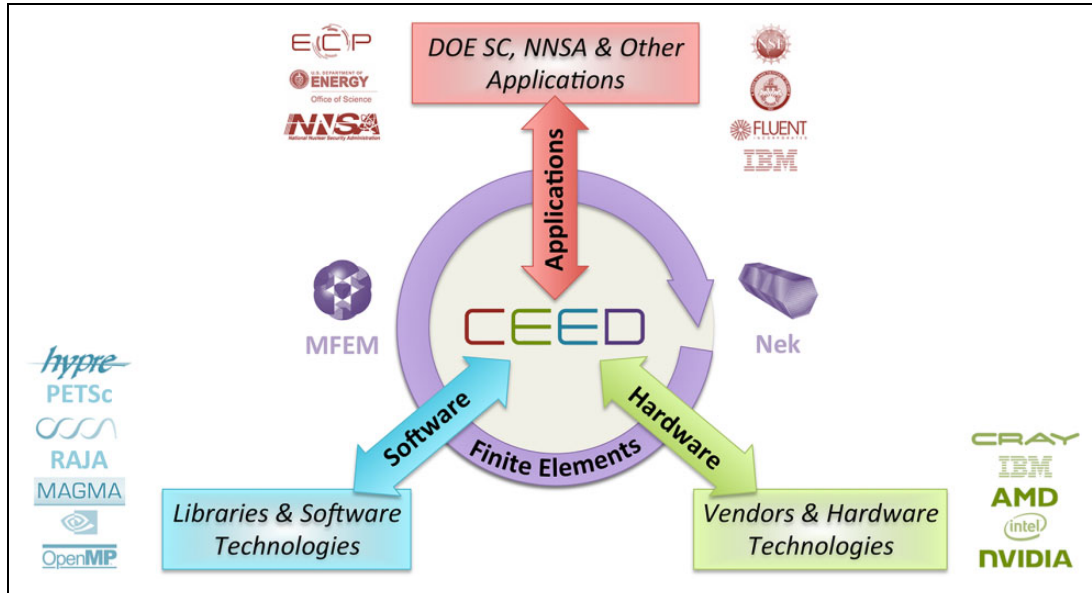


Figure 1. CEED research and development thrusts.

CEED team will develop hardware-aware technologies (pull) to understand performance bottlenecks and take advantage of inevitable hardware trends, and vendor interactions to seek (push) impact and improve hardware designs within the ECP scope.

2.3. Software thrust

The goal of CEED's Software thrust is to participate in the development of software libraries and frameworks of general interest to the scientific computing community, facilitate collaboration between CEED software packages, enable integration into and/or interoperability with overall ECP software technologies stack, streamline developer and user workflows, maintain testing and benchmarking infrastructure, and coordinate CEED software releases.

2.4. Finite elements thrust

The goal of CEED's Finite Element thrust is to continue to improve the state-of-the-art high-order finite element and spectral element algorithms and kernels in the CEED software targeting exascale architectures, connect and contribute to the efforts of the other thrusts, and lead the development of discretization libraries, benchmarks and miniapps.

The CEED co-design approach is driven by applications, and is based on close collaboration between the Applications, Hardware, and Software thrusts, each of which has a two-way, push-and-pull relation with the external application, hardware and software technologies teams.

This is illustrated in Figure 1. CEED's Finite Elements thrust serves as a central hub that ties together, coordinates and contributes to the efforts in all thrusts. For example, the development of discretization libraries in CEED is led by

the Finite Elements thrust but involves working closely with vendors (Hardware thrust) and software technology efforts (Software thrust) to take full advantage of exascale hardware. Making sure that these libraries meet the needs of, and are successfully incorporated in, ECP applications is based on collaboration between the Applications and Finite Elements thrusts.

To facilitate the co-design process, the CEED project is developing a number of benchmarks, libraries of highly performant kernels, and a set of miniapps that are serving multiple roles. One of these roles is to provide a mechanism to test and optimize across the breadth of implementations already developed by team members for a variety of platforms. The CEED bake-off problems (BPs) described in Section 4 were specifically designed for that purpose. They are simple enough to be able to be run in a simulator, but include the key local and global kernels in model problem settings. CEED also provides well-documented miniapps that are simple yet capture application-relevant physics to work with vendors, be used in system procurement, collaborate software technologies projects, and provide test and demonstration cases for application scientists. These miniapps, which are one step above the benchmarks are described in Section 5. One of their uses is to highlight performance critical paths (e.g. size of on package memory, internode latency, hardware collectives) with the goal to impact the design of exascale architectures, and system and application software, for improved portability and performance of the high-order algorithms. All of the optimizations and performance improvements resulting from the benchmarks and miniapps work is made available to applications via the CEED discretization libraries described in Sections 6 and 7.

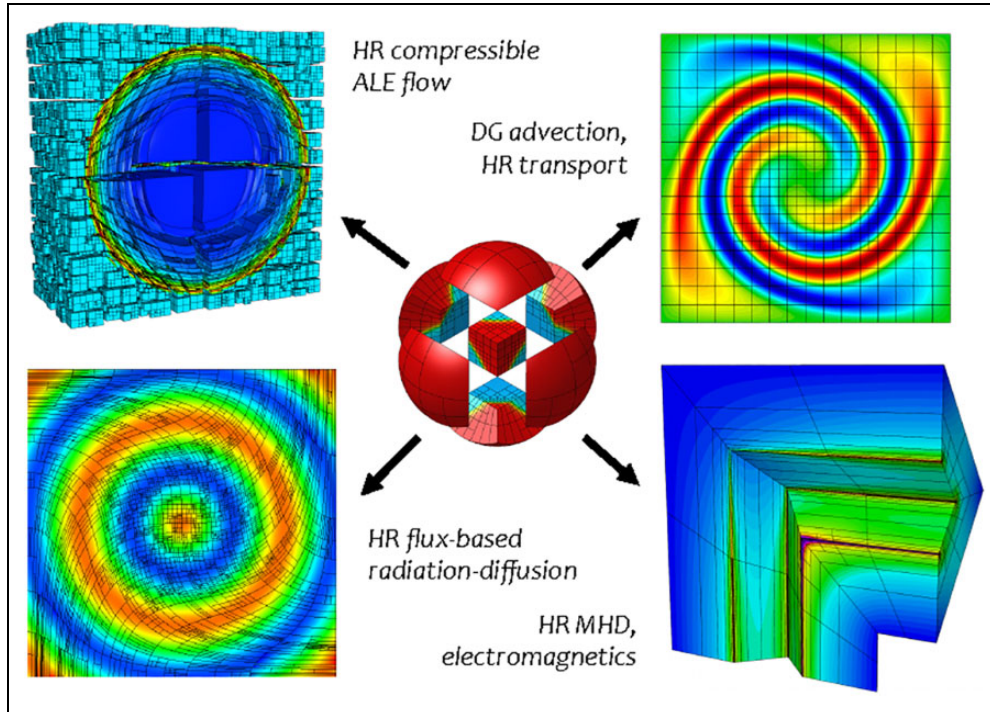


Figure 2. By incorporating AMR at the library level, many MFEM-based applications can take advantage of it with minimal code changes. Examples from high-order (HR) compressible flow, radiation diffusion/transport and electromagnetics.

3. High-order software ecosystem

While the main focus of the CEED effort is the development and improvement of efficient discretization algorithms, a full-fledged high-order application software ecosystem requires many other components: from meshing, to adaptivity, solvers, visualization and more. Therefore, CEED is also engaged in improving the additional components of the overall high-order simulation pipeline. We describe some of these efforts as well as some key enabling technologies in this section to provide a background for the discretization work discussed in the remainder of the paper. Note that some of the components described below (e.g. the MAGMA and OCCA projects) are generally applicable and could be useful in applications that do not use finite elements methods.

2.5. High-order meshing

When applying high-order discretization methods over domains with curved boundaries and/or curved material interfaces, the mesh must maintain a curved mesh geometric approximation, whose order is dictated by the order of the basis functions used to discretize the PDEs to ensure convergence of the solution. In the case when Lagrangian reference frame methods are applied the mesh geometry will naturally become curved to the same order as the elements discretizing the PDEs. Thus, the application of high-order methods requires the ability to generate curved initial meshes and to support curved mesh adaptation whenever adaptive mesh control is applied. To meet these needs the

CEED software supports curved mesh representations and has developed tools for curved mesh adaptation that include non-conforming mesh refinement/derefinement of quadrilateral and hexahedral meshes, and conforming mesh adaptation of triangular and tetrahedral meshes that can refine and coarsen the mesh to match a given anisotropic mesh metric field.

Tensor-product mesh elements (quadrilaterals in 2D and hexahedra in 3D) are attractive in many high-order applications, because their tensor-product structure enables efficient operator evaluation (see e.g. Section 4), as well as refinement flexibility (e.g. *anisotropic* refinement). Unlike the conforming case however, *hanging* nodes that occur after local refinement of quadrilaterals and hexahedra are not easily avoided by further refinement. Therefore, CEED researchers are interested in *non-conforming* (irregular) meshes, in which adjacent elements need not share a complete face or edge and where some finite element degrees of freedom (DOFs) need to be constrained to obtain a conforming solution.

The MFEM finite element library provides general support for such non-conforming adaptive mesh refinement, including anisotropic refinement, derefinement and parallel load balancing. In order to support the entire de Rham sequence of finite element spaces, at arbitrarily high-order, we use a variational restriction approach to AMR described in (Cerveny et al., 2019). This approach naturally supports high-order curved meshes, as well as finite element techniques such as hybridization and static condensation. It is also highly scalable, easy to incorporate into

existing codes, and can be applied to complex (anisotropic, n -irregular) 3D meshes, see Figure 2.

The CEED conforming mesh generation capability builds on the PUMI/MeshAdapt (Ibanez et al., 2016) libraries developed as part of the FASTMath SciDAC applied math institute. Within PUMI the curved mesh entities, edges, faces and regions, are represented as Bezier polynomials (Farin, 2014). The use of the Bezier properties, curve containment in the convex hull of control points, derivatives and products of Bezier functions being Bezier functions, and the existence of efficient degree elevation and subdivision algorithms, simplify the definition of curved mesh entity operations. One critical operation is the conversion of curved mesh Bezier geometry into interpolating geometry that is common input to analysis codes. The MeshAdapt procedures employ cavity based mesh modification operators that include optimization based entity curving, mesh entity refinement, mesh cavity coarsening, and mesh cavity swap operations (Lu et al., 2014; Luo et al., 2004). The input to MeshAdapt is an anisotropic mesh metric field defined over the entities of the current mesh. The mesh metric field can be defined as any combination of sizes as dictated by error estimation/indication procedures, feature based detection operators or other user defined size field information. Given a mesh size field MeshAdapt carries out a series of cavity based operations to modify the local mesh topology and/or geometry to satisfy the requested mesh size field. The current curved mesh adaptation procedures operate on CPUs. Efforts have been initiated to extend the GPU-based Omega_h (Ibanez, 2016a, 2016b) straight edged mesh adaptation procedures to support curved mesh entities and to include additional mesh modification operators used in curved mesh adaptation.

Controlling element shapes for evolving meshes when curved elements are used introduces additional complexity past those encountered when straight edge elements are used. In particular, methods are needed to effectively support the definition of well shaped elements in the application of ALE methods in Lagrangian reference frame simulations when meshes become highly deformed, or in the application of cavity based curved mesh modifications where new curved mesh entities must be defined within a curved mesh cavity. Methods that apply direct curved element shape optimization are being used to address these needs (Dobrev et al., 2019; Feuillet et al., 2018).

2.6. Performance portability

The MFEM, libCEED, NekRS, and libParanumal software packages developed as part of the CEED project all include support for performance portability achieved to varying degrees using the Open Concurrent Compute Abstraction (OCCA) (Medina et al., 2014; OCCA). As pictured in Figure 3 OCCA includes APIs for C, C++, F90, and Python. It provides multiple backends enabling portability to GPUs programmed using CUDA, OpenCL, and HIP. A new

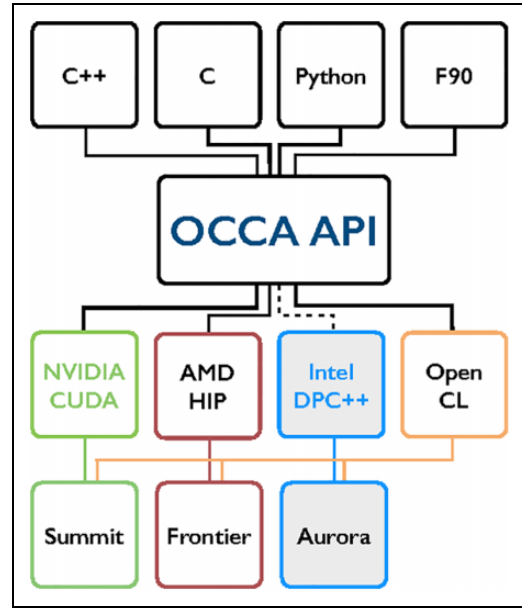


Figure 3. The OCCA portability layer provides a unified API for offloading computation to multiple backends. The Intel OneAPI backend is currently in progress.

DPC++ OCCA backend is in development to provide native support for upcoming Intel discrete GPUs. Several of these programming models also enable cross platform portability providing additional options to achieve cross platform efficiency.

OCCA exposes all performance critical features of the support backends required for high-order finite element calculations, enabling performance tuning of kernels that can achieve performance similar to kernels written to target the backends directly. We take advantage of the OCCA capability to compile compute kernels at run time with just-in-time (JIT) specialization and optimization, which is particularly important for high-order methods where innermost loops have bounds depending on the order.

2.7. Small tensor contractions

The numerical kernels of efficient high-order operator evaluation reduce to many small dense tensor contractions, one for each element of the computational mesh. These contractions can be performed in parallel over the elements and can be implemented as a batch of small matrix-matrix multiplications (DGEMMs, see Figure 4). Vendor-optimized BLAS routines have been successfully used in many areas to provide performance portability across architectures. Similarly, the availability of highly optimized Batched BLAS for various architectures can provide tensor contractions, and consequently high-order applications, performance portability. Therefore, CEED scientists have been working with vendors and the community on defining a Batched BLAS API, and finalized a proposed API for Batched BLAS (Dongarra et al., 2016, 2018).

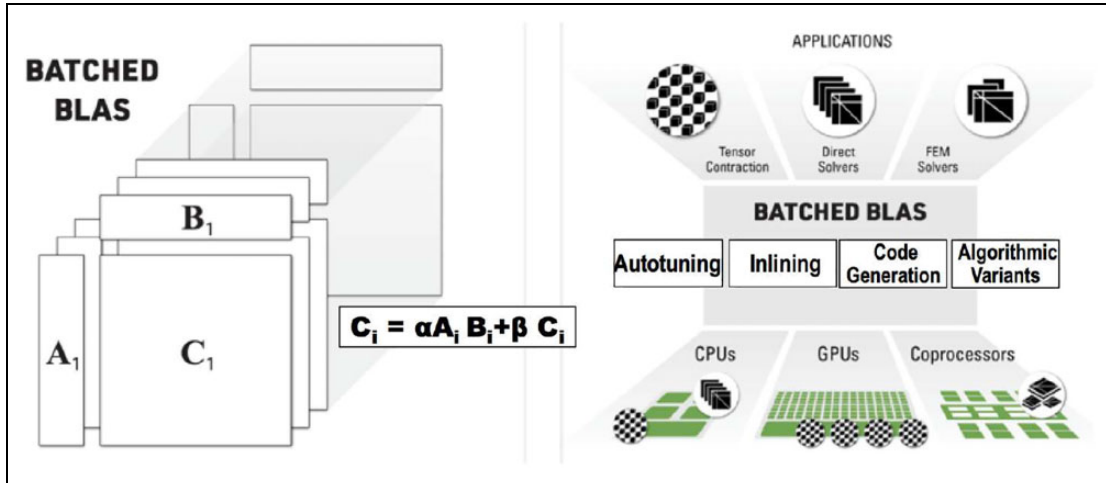


Figure 4. Standardizing a Batched BLAS API, an extension to the BLAS standard, enables users to perform thousands of small BLAS operations in parallel while making efficient use of their hardware.

The MAGMA library provides the most complete set of highly optimized Batched BLAS, including batched DGEMMs on GPUs. Very small batched DGEMMs have been optimized to perform at their theoretical performance upper bounds for a number of architectures (Abdelfattah et al., 2016a; Masliah et al., 2016). Furthermore, the tensor contraction kernels in CEED often require a sequence of batch DGEMMs. Such calls can share the same execution context so that they operate on the fast memory levels of the hardware, thus maximizing the memory bandwidth (Tomov et al., 2019).

In addition, CEED has modes of operation where the elementwise operator evaluation can be recast as standard batch DGEMMs on medium-to-large-sized matrices (Abdelfattah et al., 2016b); the MAGMA backend for libCEED exploits this to improve performance for non-tensor finite elements (Kolev et al., 2020). The use of the batch BLAS operations increases the chances of performance portability, since BLAS is often highly optimized by vendors and other open source numerical software. This was recently illustrated with the MAGMA port and CEED backend for AMD GPUs (Brown et al., 2020a, 2020b; Kolev et al., 2020).

2.8. Matrix-free linear solvers

In addition to efficient discretization and operator evaluation, matrix-free preconditioning is essential in order to obtain highly performant solvers at high order. Solvers based on explicitly formed matrices tend to have low arithmetic intensity, and the memory requirements associated with the system matrices for high-order discretizations are typically too large to be practical on GPUs and accelerator-based architectures. On the other hand, many standard preconditioning techniques rely on the knowledge of the matrix entries. For these reasons, matrix-free preconditioning is both an important and challenging topic.

Multigrid methods provide one promising avenue for the development of matrix-free linear solvers (Kronbichler and

Ljungkvist, 2019). These methods have optimal complexity, and when combined with effective matrix-free smoothers, have the potential to achieve excellent performance (Lottes and Fischer, 2005). Recent work has also studied the matrix-free construction of fast diagonalization smoothers for discontinuous Galerkin methods (Pazner and Persson, 2018). Both h -multigrid, where a sequence of geometrically coarsened meshes is used, and p -multigrid, in which a hierarchy of polynomial degrees is constructed, can be used in conjunction to obtain an efficient solver (Sundar et al., 2015). At the coarsest level, algebraic multigrid (AMG) methods, such as those from the *hypre* software library, are required in order to obtain a truly scalable solver.

An additional technique used to precondition high-order systems is to assemble a spectrally equivalent sparsified system, to which standard matrix-based preconditioning techniques may be applied. One method of obtaining a spectrally equivalent sparse matrix is using a low-order discretization on a refined mesh, and making use of the so-called finite element method–spectral element method (FEM–SEM) equivalence for tensor-product elements (Canuto, 1994; Canuto et al., 2006; Orszag, 1980). Recent work has demonstrated that, when combined with efficient solvers for the sparsified system, this approach can result in highly efficient solvers (Bello-Maldonado and Fischer, 2019; Pazner, 2020). One challenging property of the resulting low-order refined system is that the meshes resulting from the refinement procedure are not shape regular with respect to the polynomial degree p : the aspect ratio of the mesh elements increases with increasing polynomial degree. As a result, algebraic multigrid methods with pointwise smoothers such as point Jacobi result in degraded convergence at high orders. Consequently, the development of specialized matrix-free smoothers for these anisotropic low-order systems is also of interest. Additionally, the extension of these low-order preconditioners to high-order simplex elements is a topic of ongoing research (Chalmers and Warburton, 2018).

Also of interest is the development of efficient matrix-free solvers for $H(\text{curl})$, $H(\text{div})$, and discontinuous Galerkin finite element spaces. It is often the case that efficient solvers for H^1 discretizations can be modified or supplemented to obtain good preconditioners for these more challenging cases. For example, multigrid solvers for diffusion problems can be combined with a discrete gradient operator to obtain uniform preconditioners for definite Maxwell problems discretized using $H(\text{curl})$ finite elements (Kolev and Vassilevski, 2009). Although these solvers were originally developed in the context of matrix-based AMG, the same ideas can be extended to the matrix-free setting. Furthermore, uniform preconditioners for H^1 conforming diffusion problems can be combined with a simple diagonal scaling to obtain uniform preconditioners for DG diffusion problems (Antonietti et al., 2016; Dobrev et al., 2006).

An additional method that is capable of using fast diagonalization methods (for operators that admit separable approximations) is Balancing Domain Decomposition by Constraints (BDDC) (Dohrmann, 2003; Zampini, 2016), which offers more localized smoother construction, faster convergence for additive cycles, and more rapid coarsening than the fast diagonalization technique discussed above. BDDC has been used for high-order elements applied to almost incompressible elasticity (Pavarino et al., 2010), where the condition number of the BDDC-preconditioned operator for single-element smoothing and coarsening was shown to scale as $\kappa \leq C(1 + \log p^2)^2$, where p is the polynomial degree and C is robust to element size/shape and the Poisson ratio. BDDC has also been analyzed as a multigrid method (Brown et al., 2019), and can be composed with other multigrid methods.

2.9. High-order data analysis and visualization

Accurate visualization of general finite element meshes and functions in the de Rham complex requires finite element knowledge that may not be present in visualization tools employed by applications. The visualization needs to account for the orders of the mesh and solution fields, as well as the type of finite element basis used for each of them. Our work in this direction is based on the current capabilities in MFEM, illustrated in its native GLVis visualization tool (GLVis), as well as in the VisIt visualization and data analysis application (VisIt).

An additional challenge for high-order meshes and functions is that there is no common community standard for the description of high-order data at arbitrary order. CEED is working with visualization and application teams to develop a standard called Field and Mesh Specification (FMS) that not only improves visualization capabilities but also enables consistent data transfer between high-order applications. See (FMS) and (Brown et al., 2018).

4. Benchmarks

Application-relevant performance testing and analyses are critical to effective HPC software deployment. One of the foundational components of CEED is a sequence of PDE-motivated bake-off problems (BPs) designed to establish best practices for performant implementations of high-order methods across a variety of platforms. The idea is to pool the efforts of multiple high-order development groups to identify effective code optimization strategies for candidate architectures. In an initial round of tests we compared performance from four software development projects (Nek5000, MFEM, deal.II, and libParanumal) on Mira, the BG/Q at ALCF, and Summit, the NVIDIA V100-based platform at ORNL. The results of this bake-off were documented in (Fischer et al., 2020). We are interested in peak performance (degrees of freedom per second, per node) and in strong-scale performance at a significant fraction of this peak (e.g., 80%), as this regime is frequently of paramount concern to computational scientists. While we consider matrix-free implementations of p -type finite and spectral element methods as the principal vehicle for our study, the performance results are relevant to a broad spectrum of numerical PDE solvers, including finite difference, finite volume, and h -type finite elements, and thus are widely applicable.

The first suite of CEED bake-off problems, BP1–BP6, is focused on simple solver kernels—conjugate gradient (CG) iterations to solve systems of the form $(\alpha A + \beta B)\underline{u}_i = B\underline{f}_i$, which are the discrete equivalents of the constant-coefficient 3D positive-definite Helmholtz problem,

$$-\alpha \nabla^2 u_i + \beta u_i = f_i(\mathbf{x}), \quad \mathbf{x} \in \Omega \subset \mathbb{R}^3,$$

for $i = 1, \dots, m$, with homogeneous Dirichlet conditions, $u_i = 0$ on $\partial\Omega$. The odd-numbered BPs correspond to scalar problems ($m = 1$), whereas the even-numbered cases correspond to (potentially more efficient) vector problems ($m = 3$). An important aspect of using CG is that it involves a mix of local work with both nearest neighbor and global communication (vector reductions), which provides at least moderate stress on the system communication.

The BP discretizations are based on isoparametric \mathcal{Q}_p finite elements (curvilinear bricks) on a tensor-product reference domain, $\mathbf{r} \in \hat{\Omega} = [-1, 1]^3$, which is mapped through a transformation $\mathbf{x}^e(\mathbf{r})$ for each of E elements, Ω^e , $e = 1, \dots, E$. Denoting the underlying C^0 -Lagrangian basis functions as $\phi_i(\mathbf{x})$, $i = 1, \dots, n$, the respective stiffness and mass matrix entries are

$$A_{ij} = \int_{\Omega} \nabla \phi_i \cdot \nabla \phi_j dV, \quad B_{ij} = \int_{\Omega} \phi_i \phi_j dV$$

These matrices are never formed, but instead are applied using fast, low-storage, tensor-product-sum factorization that are at the heart of efficient high-order methods (Deville et al., 2002; Orszag, 1980).

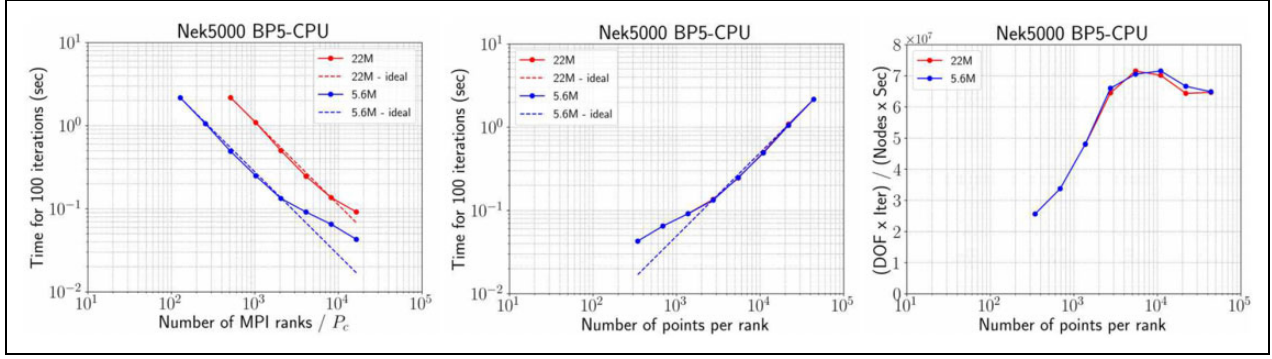


Figure 5. Strong-scaling analysis for BP5 in -c32 mode on Mira: (left) standard strong-scaling plots with increasing number of MPI ranks, P_c for Poisson problems using $n=5.6\text{M}$ and 22M grid points; (center) data collapse manifest when independent variable is n/P_c ; and (right) work-rate (left ordinate) and parallel efficiency (right ordinate) vs n/P_c .

Test problems BP1–BP2 correspond to solving the mass matrix ($\alpha = 0$, $\beta = 1$), while BP3–BP6 correspond to solving the Poisson problem ($\alpha = 1$, $\beta = 0$). For BP1–BP4, integration is performed over each element using Gauss-Legendre quadrature with $q = p + 2$ nodes in each direction in Ω . BP5–BP6 correspond to the spectral element formulation, in which integration is performed on the underlying $(p + 1)^3$ Gauss-Lobatto-Legendre nodal points, thus bypassing interpolation from nodes to quadrature points.

An important question in the development of HPC software is to ensure that testing reflects actual use modalities. On large HPC platforms, users typically use as many nodes as are effective, meaning that they run at the *strong-scale limit*, rather than the work-saturated limit. Figure 5 illustrates these limits for the case of BP5 on up to 16,384 MPI ranks on Mira. On the left we see standard strong-scale plots for two different problem sizes, $n = 5.6$ million points and $n = 22$ million points. The smaller case exhibits perfect linear speedup up to $P_c = 2048$ MPI ranks whereas the larger case sustains linear speedup out to $P_c = 8192$ ranks. For this class of problems with a given code and platform the dominant factor governing parallel efficiency is the number of points per node (or core, or other independent compute resource) (Fischer et al., 2015). Indeed, with this metric we see a perfect data collapse in Figure 5 (center), which shows the time as a function of the number of points per rank, and (right), which shows the work-rate (DOFS=degrees-of-freedom \times number of iteration per second per node) and the parallel efficiency, $\eta = T_1/(PT_P)$, where T_P is the time when running on P MPI ranks.

We make several observations about Figure 5 (right). First, the strong-scale limit is at about 2700 points per rank. Running with more points per rank keeps the efficiency at unity but increases the runtime. Running with fewer points per rank means increasing the total number of cycles (core-hours) to complete the job. Very often, users will trade some degree of inefficiency for decreased runtime. If we choose, for example, 80% efficiency, the value of n/P where this value is realized is denoted by $n_{0.8}$. Second, it is beneficial to increase rate of work (DOFS) because fewer

core-hours are then required to complete the overall task. Wall-clock time, however, may not be reduced if increasing the work-rate implies an increase in n/P to stay above the token (e.g., 80%) efficiency mark. For a problem of size n , the time-to-solution will be $t_\eta = C \frac{n}{\eta P r_{\max}}$, where η is the parallel efficiency, r_{\max} is the saturated work-rate (e.g. per computational node), P is the number of nodes used, and C is a constant that reflects the amount of work per gridpoint. The choice $\eta = 0.8$ implies that $n/P = n_{0.8}$, such that the run time is $t_{0.8} = \frac{C}{0.8} \frac{n_{0.8}}{r_{\max}}$. The problem size n and number of nodes drop out of the run time formula—the only thing that influences *time-to-solution* at the strong-scale limit is the *ratio* of the local problem size to peak processing rate, $n_{0.8}/r_{\max}$. Minimization of this ratio is of paramount importance for reduced run time.

Motivated by the preceding analysis, we routinely collect performance data for BP1–BP6 for varying problem sizes $n = Ep^3$ on a variety of platforms. A comprehensive study entailing more than 2000 trials was reported in (Fischer et al., 2020), which considered $E = 2^{14}$ to $E = 2^{21}$ and $p = 1$ to $p = 16$ using MFEM, Nek5000, and deal.ii on $P = 512$ nodes on ALCF’s Mira (in -c32 mode) and the OCCA-based libParanumal code with $P = 4$ (24 NVIDIA V100 s) on OLCF’s Summit. Typical BP results are shown in Figure 6. The left panel shows BP5 results for Nek5000 on Mira. We see that higher polynomials generally realize a higher DOFS rate and that $n_{0.8} \approx 50,000$, with $r_{0.8} \approx 65$ MDOFS. On Summit, libParanumal realizes a peak of more than 10 GDOFS with $n_{0.8} \approx 10^7$. We also show recent results for BP3 using MFEM on 256 nodes of the V100-based platform, Lassen. Again, higher polynomial orders sustain higher DOFS, peaking at ≈ 2 GDOFS per GPU (8 GDOFS/node) with an $n_{0.8} \approx 3$ million. The corresponding $t_{0.8}$ for these cases are approximately .0008s on Mira, .001 s on Summit, and .0015s on Lassen. We reiterate that the strong-scale limit is the fastest point where users can (and will) run while sustaining their desired efficiency. Thus, performance at other points on Figures 5 and 6 are of far less importance. Performance tuning must focus on moving up and to the left in these plots.

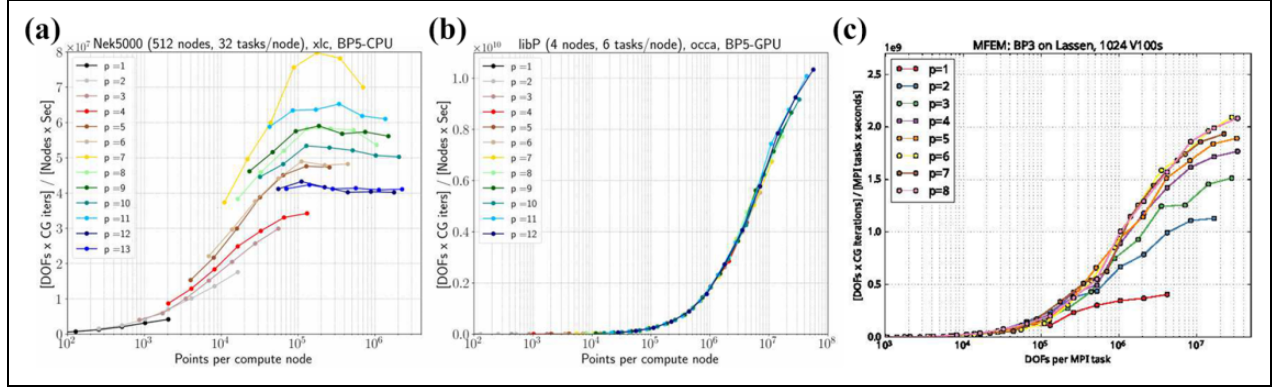


Figure 6. BP results: (a) BP5 with Nek5000 on Mira; (b) BP5 with libParanumal on Summit; (c) BP3 with MFEM on Lassen.

Future BPs will look at dealiased advection kernels, which are compute and memory intensive, and optimal preconditioning strategies for high-order discretizations of elliptic problems that represent computational bottlenecks in several of the ECP applications.

5. Miniapps

CEED is developing a variety of miniapps encapsulating key physics and numerical kernels of high-order applications. The miniapps are designed to be used in a variety of co-design activities with ECP vendors, software technologies projects and external partners. For example, several of the CEED miniapps (Nekbone and Laghos) are used as vendor benchmarks in the DOE’s CORAL-2 and LLNL’s CTS-2 procurements.

5.1. libParanumal

libParanumal (LIBrary of PARAllel NUMerical ALgorithms) is an open source project (Chalmers et al., 2020) under development at Virginia Tech. It consists of a collection of miniapps with high-performance portable implementations of high-order finite-element discretizations for a range of different fluid flow models. The miniapps embedded in libParanumal include solvers for incompressible flows (Karakus et al., 2019b), compressible flows, finite moment Boltzmann gas dynamics models (Karakus et al., 2019a), and Fokker-Planck models. All of these miniapps are accompanied with highly performant GPU kernels for high-order Galerkin (Swirydowicz et al., 2019) and/or discontinuous Galerkin spatial discretizations with a collection of high-order time integrators. libParanumal is constructed as a set of core libraries as shown in Figure 7 including high-performance scalable preconditioned iterative Krylov subspace solvers with optional multigrid preconditioning. All computationally intensive calculations are implemented using kernels compatible with the OCCA portability layer (Medina et al., 2014) and have been analyzed and optimized to guarantee that they achieve a high percentage of the attainable DEVICE memory bandwidth on NVIDIA P100 (Swirydowicz et al., 2019) and V100

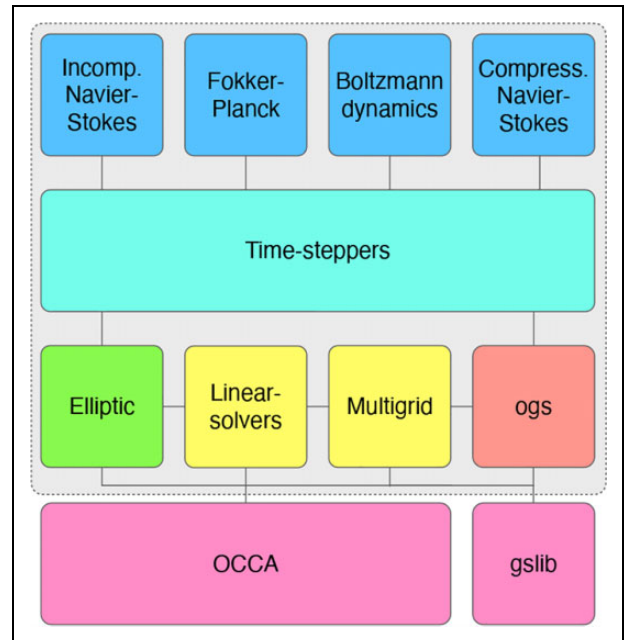


Figure 7. The libParanumal library includes portable GPU accelerated miniapps for solving the incompressible and compressible Navier-Stokes, Fokker-Planck, and finite moment Boltzmann gas dynamics equations among others.

GPUs (Fischer et al., 2020). The libParanumal library provides core GPU-acceleration capabilities to NekRS and algorithms developed for it have been deployed in the MFEM cuda-gen backend.

5.2. NekBench and Nekbone

NekBench is a benchmark suite representing key components of Nek5000/CEM/RS. This miniapp supports a variety of benchmarks for fundamental analysis in different architectures. It supports a single-step driver that delivers timing measurements on CPUs and GPUs for ping-pong (one-to-one and bisection-bandwidth tests), gather-scatter, all-reduce, dot-product, and device-to/from-host memcopy. It also performs weak and strong-scaling tests for BK5–BK6¹ and BP5–BP6, as shown in Figure 8. The figure

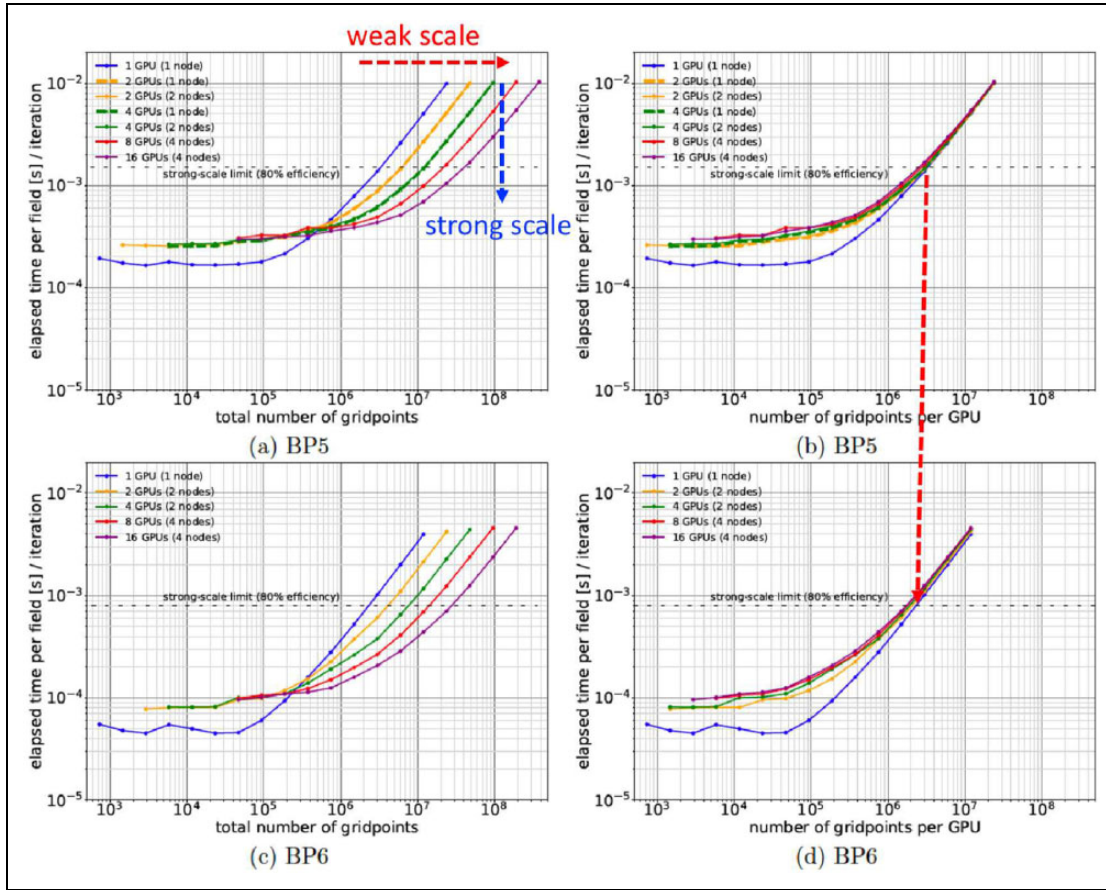


Figure 8. Strong and weak scaling studies of (a) and (b) BP5 and (c) and (d) BP6 using V100s on Lassen.

demonstrates that $n_{0.8}$ is reduced from 3M points to 2M points per V100 when we switch from the scalar (BP5) solver to the vector (BP6) variant of the solver.

Nekbone solves a standard Poisson equation using conjugate gradient iteration with a simple diagonal preconditioner on a block or linear geometry. It encapsulates one of the principal computational kernels pertinent to Nek5000, which includes a mixture of local (near-neighbor) and non-local (vector reduction) communication patterns that are central to efficient multilevel solvers. Nekbone has been updated to include vector solutions, which allows amortization of message and memory latencies. Nekbone has been used for assessment of advanced architectures and for evaluation of lightweight MPI implementations on the ALCF BG/Q, Cetus, in collaboration with Argonne’s MPICH team (Raffenetti et al., 2017).

5.3. Laghos and Remhos

Laghos (LAGrangian High-Order Solver) and Remhos (REMap High-Order Solver) are MFEM-based miniapps developed by the CEED team. The objective of these miniapps is to provide open source implementations of efficient discretizations for Lagrangian shock hydrodynamics (Laghos) and field remap (Remhos) based on high-order finite elements.

Laghos (Laghos) solves the time-dependent Euler equations of compressible gas dynamics in a moving Lagrangian frame. The miniapp is based on the method described in (Dobrev et al., 2012). It exposes the principal computational kernels of explicit time-dependent shock-capturing compressible flow, including the FLOP-intensive definition of artificial viscosity at quadrature points.

Laghos supports two options for deriving and solving its system of equations, namely, the full assembly and the partial assembly methods. Full assembly relies on global mass matrices in CSR format; this option is appropriate for first- or second-order methods. Partial assembly utilizes the tensor structure of the finite element spaces, resulting in less data storage, memory transfers and FLOPs; this option is of interest in terms of efficiency for high-order discretizations. The Laghos implementation includes support for hardware devices, such as GPUs, and programming models, such as CUDA, OCCA, RAJA and OpenMP, based on MFEM 4.1 or later. These device backends are selectable at runtime. Laghos also contains an AMR version demonstrating the use of dynamic adaptive mesh refinement for a moving mesh with MFEM.

Large-scale GPU runs of Laghos were performed on Lassen. All computations were kept on the device, except for the result of the dot-product which is brought back on to the CPU during the iterations of the CG solver. Initial

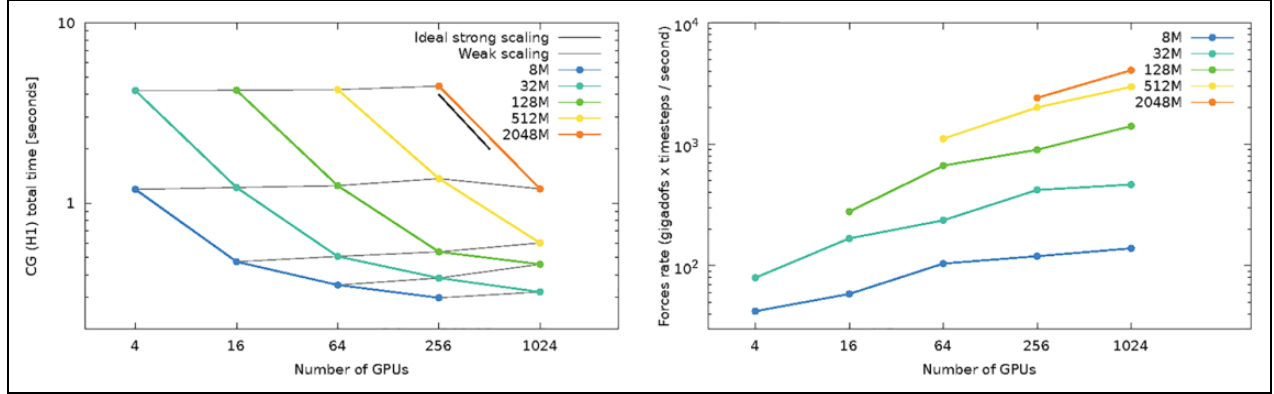


Figure 9. Left: weak and strong-scaling results with Laghos and MFEM-4.1: 2D problem on Lassen, using up to 1024 GPUs. Right: throughput for the Laghos force kernel in (GDOF \times timesteps/second), reaching above four TDOF/s on 1024 GPUs.

results are presented in Figure 9, showing both the weak (gray lines) and strong (colored lines) scaling obtained on four to one thousands of GPUs during the CG iterations of the velocity solver, which corresponds to the BP2 CEED benchmark. Ideal strong scaling is possible for problem size large enough, while weak scaling is more easily reached through all the range of the runs. The bottom panel of Figure 9 presents the throughput in DOFs per second for the Laghos force kernel, reaching more than four TDOF/s on the same configuration.

Remhos solves the pure advection equations that are used to perform conservative and monotonic DG advection-based discontinuous field interpolation, or “remap” (Remhos). Remhos combines discretization methods described in the following articles: (Anderson et al., 2015, 2017, 2018; Hajduk et al., 2020a, 2020b). It exposes the principal computational kernels of explicit time-dependent Discontinuous Galerkin advection methods, including monotonicity treatment computations that are characteristic to FCT (Flux Corrected Transport) methods.

Remhos supports two execution modes, namely, transport and remap, which result in slightly different algebraic operators. In the case of remap, the finite element mass and advection matrices change in time, while they are constant for the transport case. Just like Laghos, Remhos supports full assembly and partial assembly options for deriving and solving its linear system. Support for different hardware devices in Remhos is work in progress.

Other computational motifs supported by both Laghos and Remhos include: domain-decomposed MPI parallelism; support for unstructured 2D and 3D meshes, with quadrilateral and hexahedral elements; moving high-order curved meshes; explicit high-order time integration methods; optional in-situ visualization with GLVis and data output for visualization and data analysis with VisIt.

6. libCEED

libCEED is CEED’s low-level API library that provides portable and performant evaluation of high-order operators (Abdelfattah et al., 2020). It is a C99 library with no

required dependencies, and with Fortran and Python interfaces (see for details on the Python interface, Barra et al., 2020).

One of the challenges with high-order methods is that a global sparse matrix is no longer a good representation of a high-order linear operator, both with respect to the FLOPs needed for its evaluation, as well as the memory transfer needed for a matvec. Thus, high-order methods require a new “format” that still represents a linear (or more generally nonlinear) operator, but not through a sparse matrix.

The goal of libCEED is to propose such a format, as well as supporting implementations and data structures, that enable efficient operator evaluation on a variety of computational device types (CPUs, GPUs, etc.). This new operator description, outlined below and in (Abdelfattah et al., 2020), is based on algebraically factored form, which is easy to incorporate in a wide variety of applications, without significant refactoring of their own discretization infrastructure.

6.1. Finite element operator decomposition

Finite element operators are typically defined through weak formulations of partial differential equations that involve integration over a computational mesh. The required integrals are computed by splitting them as a sum over the mesh elements, mapping each element to a simple *reference* element (e.g. the unit square) and applying a quadrature rule in reference space.

This sequence of operations highlights an inherent hierarchical structure present in all finite element operators where the evaluation starts on *global (trial) degrees of freedom (DOFs) or nodes on the whole mesh*, restricts to *DOFs on subdomains* (groups of elements), then moves to *independent DOFs on each element*, transitions to independent *quadrature points* in reference space, performs the integration, and then goes back in reverse order to global (test) degrees of freedom on the whole mesh.

This is illustrated below for the simple case of symmetric linear operator on third order (Q_3) scalar continuous (H^1) elements, where we use the notions **T-vector**, **L-vector**, **E-vector**, and **Q-vector** to represent the sets corresponding

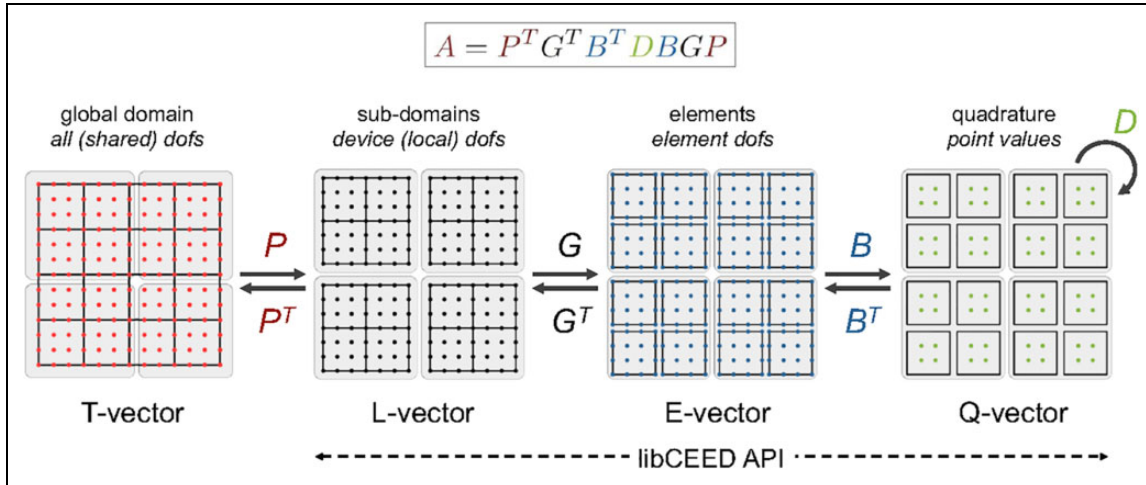


Figure 10. Finite element operator decomposition.

to the (true) degrees of freedom on the global mesh, the split local degrees of freedom on the subdomains, the split degrees of freedom on the mesh elements, and the values at quadrature points, respectively (see Figure 10). We refer to the operators that connect the different types of vectors as:

- P : Subdomain restriction;
- G : Element restriction;
- B : Basis (DOFs-to-Qpts) evaluator;
- D : Operator at quadrature points.

More generally, when the test and trial space differ, they have their own versions of P , G and B .

The libCEED API takes an algebraic approach, where the user describes in the *frontend* the operators G , B , and D and the library provides *backend* implementations and coordinates their action to the original operator on **L-vector** level (i.e. independently on each device/MPI task). The subdomain restriction operation, P is outside of the scope of the current libCEED API.

One of the advantages of this purely algebraic description is that it includes all the finite element information, so the backends can operate on linear algebra level without explicit finite element code. The frontend description is general enough to support a wide variety of finite element algorithms, as well as some other types algorithms such as spectral finite differences. The separation of the front and backends enables applications to easily switch/try different backends and enables backend developers to impact many applications from a single implementation.

The mapping between the decomposition concepts and the code implementation is as follows:

- **L-, E-, and Q-vectors** are represented by *CeedVector* objects²;
- G is represented as a *CeedElemRestriction*;
- B is represented as a *CeedBasis*;

- D is represented as a *CeedQFunction*;
- $G^T B^T D B G$ the local action of the operator is represented as a *CeedOperator*.

Users can provide source code for pointwise application of their weak form in a single source file using mutually supported constructs from C99, C++11, and CUDA.

6.2. CPU and GPU performance

libCEED provides a unified interface for all types of hardware, allowing users to write a single source code and to select the desired backend at run time. Backends differ in the hardware they target, but also in their implementation and algorithmic choices.

libCEED provides backends for CPUs, NVIDIA GPUs, and AMD GPUs implemented in C (with or without AVX intrinsics), CUDA, and HIP respectively. libCEED also provides backends taking advantage of specialized libraries like libXSMM for CPUs, or MAGMA for NVIDIA and AMD GPUs (see Figure 11). The OCCA backend is special in the sense that it aims at supporting all possible hardware in a unified backend.

Backends are interoperable, allowing to use different backends together on heterogeneous architectures. Moreover, each process or thread can instantiate an arbitrary number of backends, this can be used to select the backend with the highest performance for each operator, or to run on mixed meshes.

The best performing CPU backends use the LIBXSMM library. In order to best use modern CPU architectures, the basis application operations are decomposed as small matrix multiplications, using tensor contractions when on tensor-product bases. LIBXSMM provides efficient computation of these small matrix-matrix products, vectorized across the quadrature points for a single element or batches of elements.

The best performing GPU backends on the CEED benchmark problems are `/gpu/cuda/gen` for

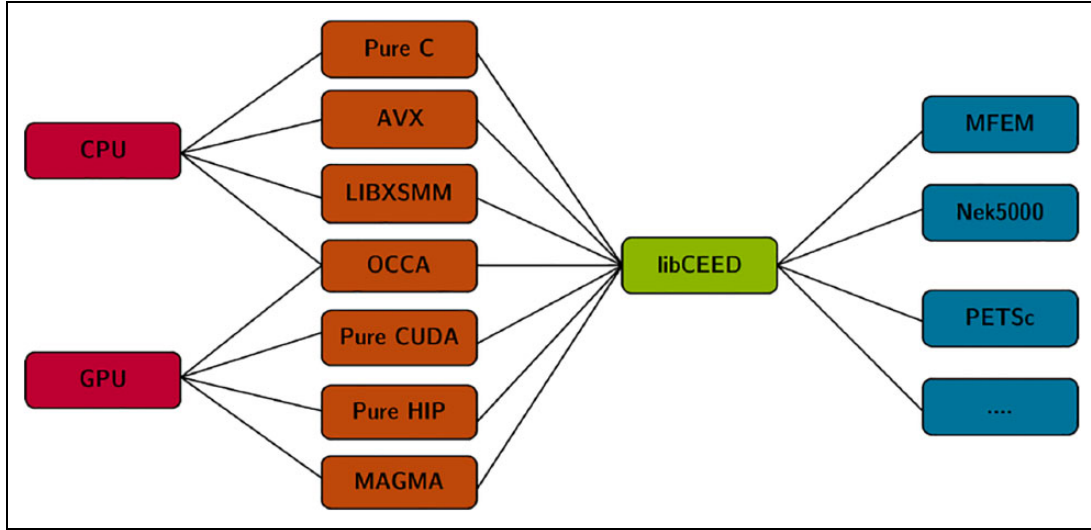


Figure 11. The role of libCEED as a low-level API.

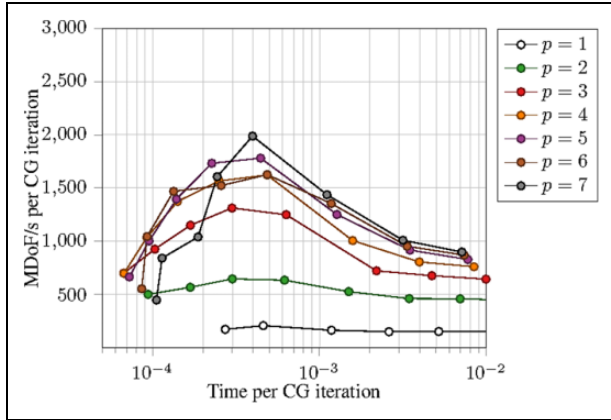


Figure 12. Throughput vs latency for the libCEED /cpu/self/xsmm/blocked backend solving BP3 on a 2-socket AMD EPYC 7452.

quadrilateral and hexahedral elements, and /gpu/magma for simplex elements.

The /gpu/cuda/gen backend is using runtime code generation and JIT compilation to generate a unique optimized GPU kernel for each libCEED operator. The gpu/magma backend is based on the MAGMA library.

In order to explore the weak and strong scaling of these algorithms on CPU and GPU architectures, we consider the throughput in terms of the latency. High throughput for low latencies are good for architectures and problems that need strong scalability. On the other hand, high throughput for high latency corresponds to architectures that weak-scale efficiently. Comparing these measures for a two-socket AMD EPYC 7452 CPU, in Figure 12, with an NVIDIA V100 GPU in Figure 13, using the best CPU and GPU backends of libCEED, we observe a much better ability for the CPU to strong scale over the GPU, but a better weak

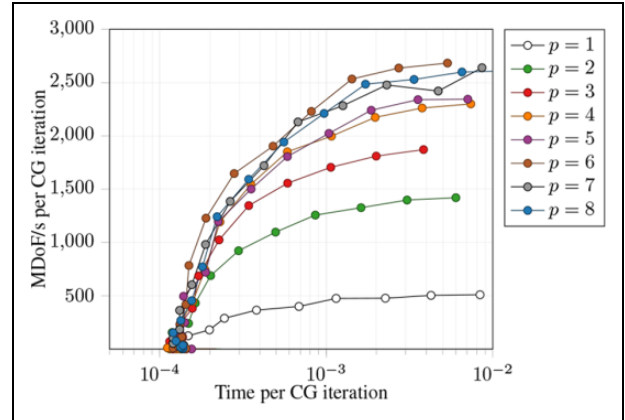


Figure 13. Throughput vs latency for the libCEED /gpu/cuda/gen backend solving BP3 on a NVIDIA V100.

scalability for the NVIDIA V100 over the AMD EPYC 7452 on the benchmark problem BP3 (in both Figures, p is the polynomial order).

For all GPU backends the weak and strong scalability behave similarly. A constant GPU overhead limits dramatically the strong scalability. Any kernel below a million degrees of freedom is dominated by this constant cost, therefore the computation time between one and a million degrees of freedom is roughly the same, which correspond to the clustering on the left of Figure 13. However, above a million degrees of freedom, we can observe a superlinear weak scalability: a higher number of degrees of freedom results in higher performance per degree of freedom. The CPU backends have typical CPU performance profile, with decent strong and weak scalability, and optimal performance achieved in the middle for a number of degrees of freedom that depends on the architecture cache memories, see Figure 12.

7. Nek and MFEM

At a higher level of abstraction, CEED provides a “high-level API” to applications through the MFEM and Nek discretization libraries. This API operates with global discretization concepts, specifying a global mesh, finite element spaces and PDE operators to be discretized with the pointwise physics representing the coefficients in these operators. Given such inputs, CEED provides efficient discretization and evaluation of the requested operators, without the need for the application to be concerned with element-level operations. Internally, the high-level API can make use of CEED’s low-level APIs described in the previous sections. The global perspective also allows CEED packages to provide general unstructured adaptive mesh refinement support, with minimal impact in the application code.

7.1. Nek5000/CEM/RS

Nek5000 is a thermal-fluids code based on the spectral element method (SEM) (Patera, 1984) that is used for a wide range of scientific applications, including reactor thermal-hydraulics, thermal convection, ocean modeling, combustion, vascular flows, and fundamental studies of turbulence. NekCEM supports both an SEM and an SE discontinuous Galerkin (SEDG) formulation for applications in electromagnetics, drift-diffusion, and quantum-mechanical systems. These codes have scaled to millions of MPI ranks using the Nek-based *gslib* communication library to handle all near-neighbor and other stencil type communications (e.g., for algebraic multigrid). Tensor contractions constitute the principal computational kernel, which leads to high CPU performance with a minor amount of tuning.

Initial GPU development and testing was done with NekCEM using OpenACC (Otten et al., 2016). For portability reasons, NekRS—the GPU variant of Nek5000—was built on top of kernels from libParanumal using OCCA. In both cases, node-level parallelism requires kernels written at a higher level than simple tensor contractions. For performance, full operations (e.g., ∇u) are cast into a single kernel call for the GPU. Significant effort has gone into overlapping the gather-scatter operation that is central to matrix-free SEM/FEM operator evaluation. On a CPU platform, where there are only one or two spectral elements per MPI rank, there is no opportunity for communication overlap. However, GPUs such as the NVIDIA V100 require about 2 million gridpoints per V100 for reasonable efficiency, which means that there is enough work on subdomain interior points to cover some internode communication. Strong- and weak-scaling performance for NekRS and Nek5000 on Summit are illustrated in Figure 14. The strong-scaling plots reflect the most recent performance enhancements in NekRS, including communication overlap and improved preconditioners. Time-per-step in this case is less than 0.1 seconds.

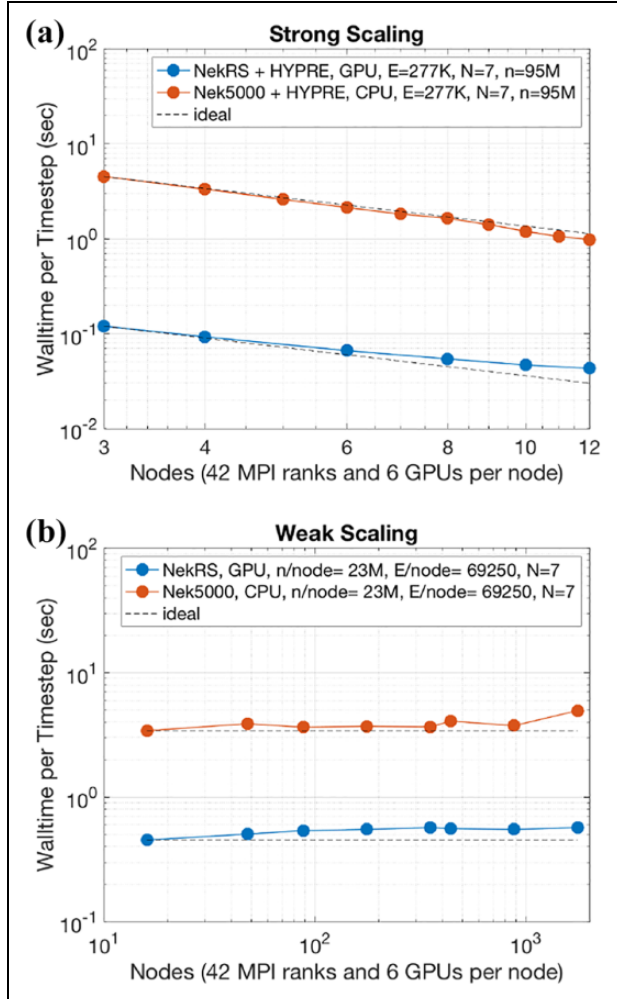


Figure 14. NekRS GPU and Nek5000 CPU (a) strong- and (b) weak-scaling performance on OLCF’s Summit. The plots show the wall-time per step, averaged over 100 steps, for turbulent flow in a 17×17 rod-bundle example coming from ExaSMR (Figure 17(a)).

A major push for Nek5000/CEM/RS applications is in the area of solvers. Effective preconditioners for the Poisson problem are of primary importance for the unsteady Navier-Stokes equations. Steady-state solvers are important for the drift-diffusion equations and for Reynolds-averaged Navier-Stokes (RANS) models used in nuclear engineering. Jacobi-free Newton-Krylov methods (Knoll and Keyes, 2004) are under development for these applications.

7.2. The MFEM finite element library

MFEM is a free, lightweight, scalable C++ library for finite element methods (Anderson et al., 2020; MFEM). Its goal is to enable high-performance scalable finite element discretization research and application development on a wide variety of platforms, ranging from laptops to exascale supercomputers. It also provides a range of features beyond finite elements that allow for rapid prototyping and

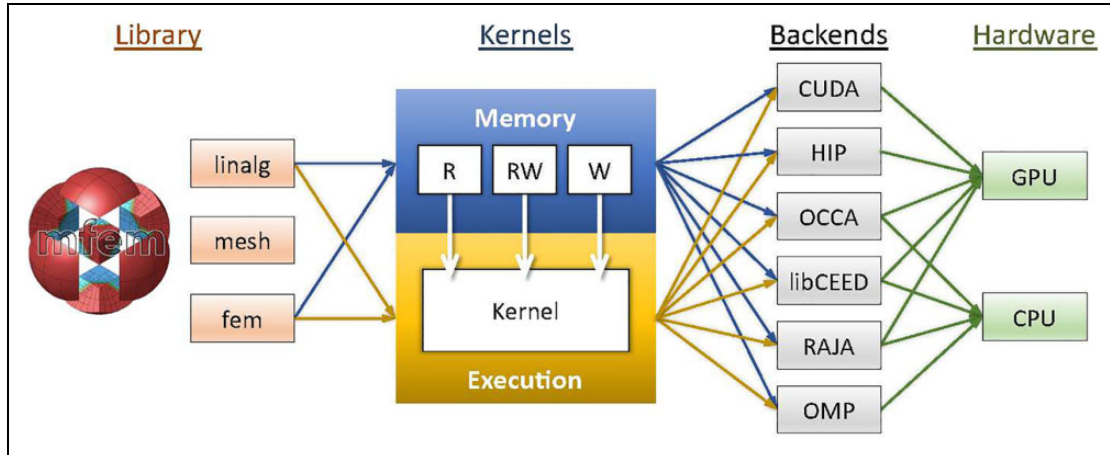


Figure 15. Conceptual diagram of MFEM's portability abstractions.

development of scientific and engineering simulations. In CEED, MFEM is a main component of the efforts in the Applications and Finite Element thrusts.

MFEM includes capabilities for basic linear algebra: vectors, dense and sparse matrices and operations with them; iterative (Krylov) linear solvers; smoothers and preconditioners, including multigrid; nonlinear operators and solvers; and time stepping methods. The library offers support for a wide variety of mesh types and operations on them: arbitrary high-order curvilinear meshes in 1D, 2D (triangles and quads), and 3D (tets, hexes, prisms), including surface and periodic meshes; mesh import from meshing tools such as Gmsh, Netgen, CUBIT; adaptive conforming mesh refinement for simplicial meshes; adaptive non-conforming refinement and derefinement for all mesh types, including parallel rebalancing; mesh optimization via node movement: TMOP (Dobrev et al., 2019). The PDE discretization features include: arbitrary order L^2 - (discontinuous), H^1 - (continuous), $H(\text{div})$ -, and $H(\text{curl})$ -conforming finite elements and discretization spaces; NURBS meshes and discretization spaces (IGA); a large variety of predefined linear, bilinear, and nonlinear forms; support for many discretization approaches including continuous, mixed, DG, DPG, IGA, etc. In terms of parallel programming, MFEM supports MPI-based distributed memory parallelism, OpenMP-based shared memory parallelism on CPUs, and GPU-acceleration through various backends (see below). Last but not least, the source distribution includes many examples and miniapps that can be used as an introduction to the library and its capabilities, as well as templates for developing more complex simulations.

In addition to its built-in capabilities, MFEM provides integration with many other scientific libraries, including ECP software technologies projects such as *hypr*, PETSc, SUNDIALS, PUMI, libCEED, OCCA, etc. Support for the GPU capabilities in some of these libraries is already available (e.g. OCCA, libCEED) and for others it is currently under active development (e.g. *hypr*, SUNDIALS).

Starting with version 4.0 (released in May 2019), MFEM introduced initial support for GPU accelerators. Since then these capabilities are being actively developed to add support in more components of the library while also improving the performance of already existing kernels. The set of examples and miniapps in MFEM that support GPUs is growing and now includes a number of PDE problems: diffusion, advection, definite Maxwell, grad-div, Darcy, etc. Other algorithms like AMR, TMOP (Dobrev et al., 2019), and multigrid are also supported (at least partially) on GPUs.

The support for different hardware (CPUs and GPUs) and different programming models (such as CUDA, OpenMP, HIP, RAJA (Beckingsale et al., 2019; RAJA), OCCA, libCEED) is facilitated by the concept of *backends*, see Figure 15. The selection of the backend happens at runtime at the start of the program which allows code to be developed, tested, and used without the need to recompile the library or the application. The backends currently supported (specified as strings) are: `cpu`, `raja-cpu`, `occa-cpu`, `ceed-cpu`, `omp`, `raja-omp`, `occa-omp`, `debug`, `hip`, `cuda`, `raja-cuda`, `occa-cuda`, and `ceed-cuda`.

To facilitate gradual transition to GPU architectures for users and for the library itself, MFEM introduced two features: a lightweight memory manager to simplify the handling of separate host and device memory spaces, and a set of `MFEM_FORALL` macros for writing portable kernels that can dispatch execution to different backends.

These capabilities are illustrated in Figure 16, where we present results for the BP3 benchmark (implemented as a slightly modified version of MFEM's Example 1) on a single V100 GPU on LLNL's Lassen machine. These results show the performance advantage of libCEED's CUDA-gen backend (exposed in MFEM as the `ceed-cuda` backend) over the `cuda` MFEM backend. The main reason for this improvement is the additional kernel fusion used by CUDA-gen: the action of the operators $G^T B^T D B G$ (see Section 6, and Figure 10) is implemented as one kernel

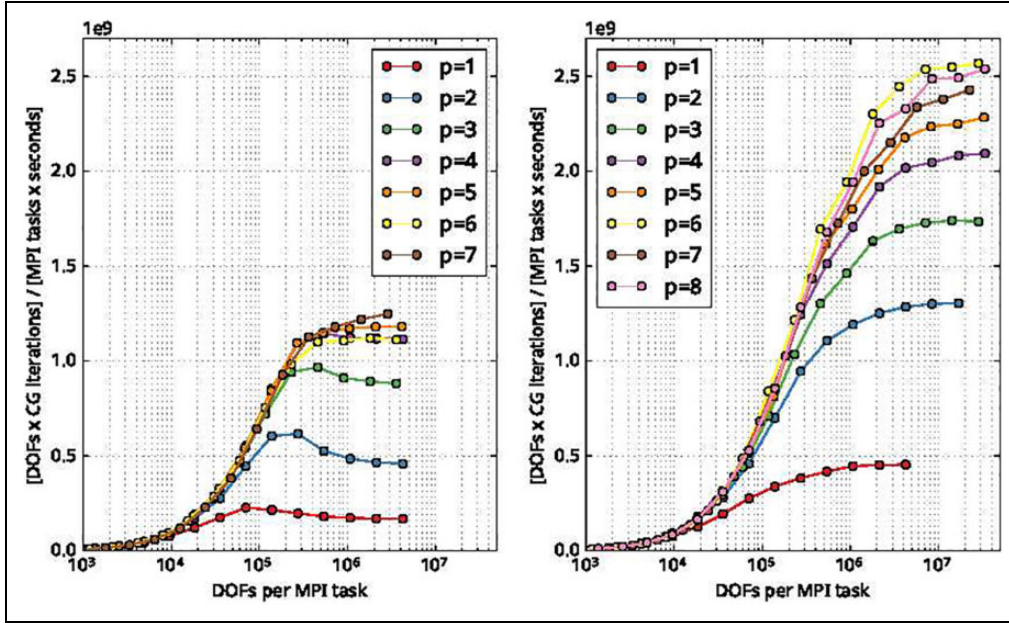


Figure 16. BP3 performance comparison of MFEM's `cuda` (left) and `ceed-cuda` (i.e. `/gpu/cuda/gen` from libCEED, right) backends on a single V100 GPU (Lassen machine at LLNL).

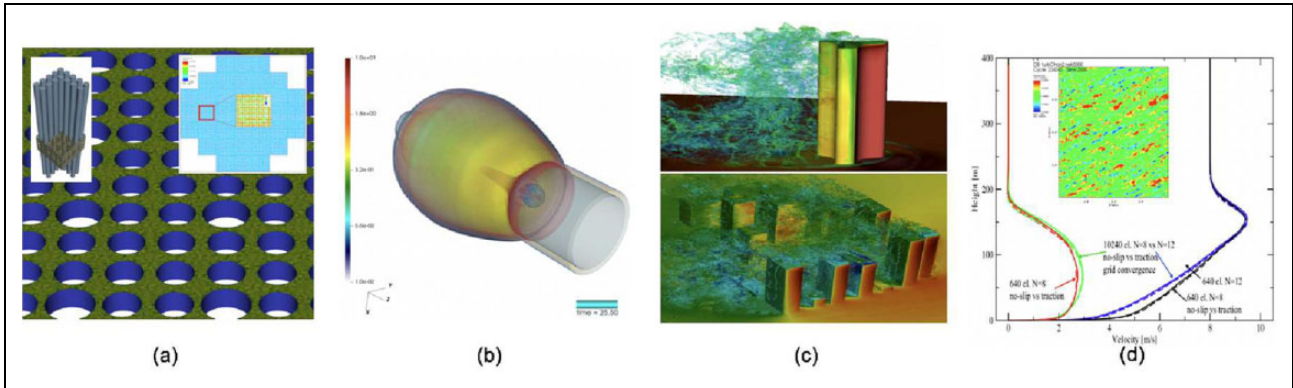


Figure 17. ECP applications: (a) ExaSMR: 17×17 rod-bundle turbulent flow simulation. (b) MARBL: 3D multi-material ALE simulation that is used as a performance benchmark. (c) Urban: LES modeling for vortex flows around Lake Point Tower and 20 buildings in Chicago downtown block. (d) ExaWind: GABLS benchmark studies with no-slip and traction boundary conditions.

whereas the `cuda` backend uses three separate kernels for G^T , $B^T DB$, and G . Another MFEM result using the `ceed-cuda` backend to solve BP3 on 1024 V100 GPUs on Lassen was presented earlier in the right panel of Figure 6.

8. Application integration

The ultimate goal of CEED is to extend state-of-the-art high-order algorithms to DOE-ECP mission applications. This section illustrates the use of CEED-developed technologies for several ECP applications including ExaSMR, MARBL, Urban, and ExaWind. We also demonstrate impact over a range of other important applications including work sponsored by DOE's Nuclear Energy Advanced Modeling and Simulation program, Vehicle Technologies

Office, COVID-19 research, and SciDAC. Applications in these areas present significant challenges with respect to scale resolution, multiphysics, and complex computational domains. The CEED team has focused on developing algorithmic and scientific research at scale to address these issues in collaboration with the application teams. The outcomes of the CEED technologies have been integrated into the open source codes, Nek5000/CEM/RS, MFEM, libCEED and libParanumal. Their impact in various application problems have been demonstrated in the CEED milestone reports (Brown et al., 2017; Kolev et al., 2020; Min et al., 2017, 2019a; Shephard et al., 2019; Tomov et al., 2018, 2019) and other project reports (Ameen et al., 2020; Merzari et al., 2017; Min et al., 2019b). Some of the results are shown in Figures 17 and 18.

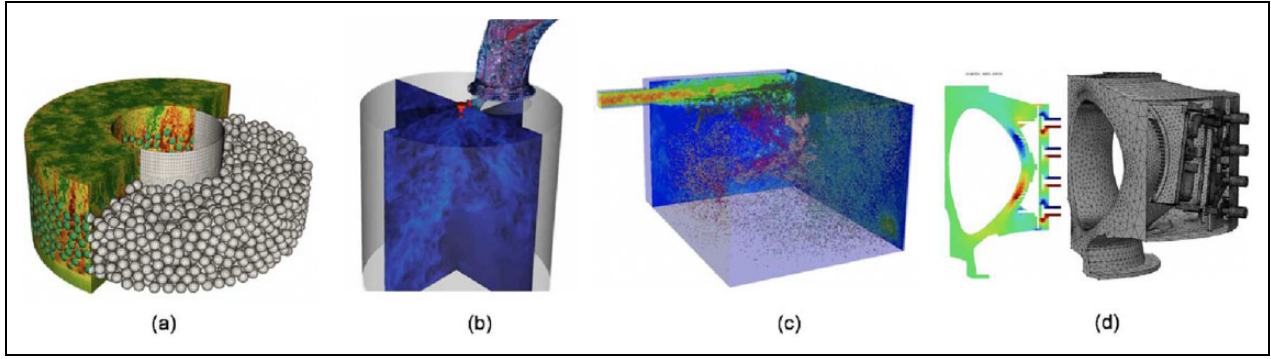


Figure 18. Other applications: (a) DOE NEAMS: Turbulent flows around 3344 pebbles with an all-hex mesh. (b) DOE VTO: Exhaust stroke TCC engine modeling. (c) COVID19: LES Lagrangian particle tracking simulation for 500,000 aerosols. (d) SciDAC RF: EM analysis of the vacuum region from a RF antenna.

8.1. Small modular reactor analysis: ExaSMR

The goal of the ExaSMR project is to combine advanced thermal-hydraulics modeling with scalable neutronics computations to allow reactor-scale modeling. For the thermal-hydraulics analysis, the reactor-core comprises hundreds of thousands of channels supporting turbulent flow with very fine solution scales. The channels are typically hundreds of hydraulic diameters in length. For full reactor-core simulations, Reynolds-Averaged Navier-Stokes (RANS) approach in the majority of the core with more detailed large-eddy simulations (LES) is required in critical regions. In addition, while the turbulence is challenging to resolve, it tends to reach a statistically fully developed state within just a few channel diameters, whereas thermal variations take place over the full core size. This poses a challenge for coupled calculations. It is too expensive to consider performing full large-eddy simulations (LES).

To accelerate the time-to-solution, CEED-developed fully implicit and steady-state solvers for spectral-element-based thermal transport and RANS. For the nonlinear Navier-Stokes and RANS transport, the Jacobian-free Newton Krylov (JFNK) routines from Nek-CEM's drift-diffusion solver (Tsai et al., 2020) have been imported to Nek5000 and tested on various flow problems including vortex flow, Dean's flow, lid-driven cavity, flow past cylinder and flow around rods (Brown et al., 2018). This new steady-state solver uses an inexact (Jacobi-free) formulation based on a first-order Taylor series expansion and converges to the steady-state solutions with a small number of pseudo-time steps. Preconditioning the GMRES routine within the Newton step remains as future work.

While the thermal development time is governed by the long channel length, the velocity rapidly reaches a statistically steady state, proportional to the hydraulic diameter. By freezing the expensive-to-generate velocity field, one can accelerate equilibration of the thermal field without having to laboriously compute tens of thousands of transient turbulent eddies. We have developed preconditioning strategies for steady or implicit advection-diffusion and Navier-Stokes equations using tensor-product-based

spectral element methods. For the advection-diffusion problem, p -multigrid (PMG) is used directly as a preconditioner within a Krylov subspace projection (KSP) method such as GMRES (Brubeck and Fischer, 2019; Pazner and Persson, 2018; Tomov et al., 2018). For Navier-Stokes, PMG is part of a larger preconditioner that includes restriction of velocity search directions to the space of divergence-free fields through a projection technique (Brubeck et al., 2020). These strategies have been applied for turbulent thermal-stress models for rod-bundle simulations (Martinez et al., 2019).

For the unsteady ExaSMR simulations, the GPU-based NekRS code has made significant advances in development. For the 17×17 rod-bundle in Figure 17(a), we have demonstrated improved NekRS simulation capabilities to extend to largest problem size to date, with 175 million spectral elements ($n = 61$ billion grid points) using 3520 nodes (21120 V100s) on Summit as well as strong and weak scaling performance at large scale in (Kolev et al., 2020; Tomov et al., 2019) (see Figure 14).

8.2. Compressible shock hydrodynamics

MARBL is a next-gen multi-physics simulation code being developed at LLNL. The code targets multi-material hydrodynamics and radiation/magnetic diffusion simulations, with applications in inertial confined fusion, pulsed power experiments, and equation of state/material strength analysis. The goal of this application is to enhance LLNL's modular physics simulation capabilities, with increased performance portability and flexibility. One of the central features of MARBL is an ALE formulation based on the MFEM-driven BLAST code (Anderson et al., 2018), which solves the conservation laws of mass, momentum, and energy. The BLAST code utilizes high-order finite element discretizations of several physical processes on a high-order (curved) moving mesh. The method consists of (i) a Lagrangian phase, where the multi-material compressible Euler equations are solved on a moving mesh (Dobrev et al., 2012, 2016), (ii) a remesh phase, which improves the mesh quality (Dobrev et al., 2020), and a field remap

phase that performs a conservative and monotone advection between two meshes (Anderson et al., 2015).

The first major step toward improved efficiency in MARBL was the introduction of matrix-free/partial assembly based methods. The CEED-developed Laghos miniapp played a critical role for that, as it exposed the main computational kernels of BLAST's Lagrangian phase, without the additional overhead of physics-specific code. Laghos introduced partial assembly versions for many of BLAST's specific kernels, which were later directly used by the application. For its more standard finite element operations, BLAST utilized MFEM's tensor-based routines. These included partially assembled bilinear forms for mass, diffusion and advection; tensor-based evaluation of finite element functions and their gradients; matrix-free diagonal preconditioning; and other algorithms as well. These methods were used extensively throughout the application's Lagrangian and remap phases. Furthermore, the CEED team derived a matrix-free version of MFEM's mesh optimization miniapp, which could also be used directly by the remesh phase of the application.

The GPU port of MARBL/BLAST is exclusively based on the partial assembly technology from CEED and the GPU support via the MFEM version 4.0 release. The CEED team developed GPU versions of Laghos and the MFEM's mesh optimization miniapps. GPU kernels from these miniapps, together with general MFEM finite element operations as the ones mentioned above, could be used directly by the MARBL code. Application-specific operations, on the other hand, are implemented in MARBL, making use of the RAJA kernel abstractions and MFEM memory management, GPU-friendly data structures, small dense matrix kernels, use of shared memory, etc.

The current state of MARBL's GPU capability provides around $15\times$ speedup on the main benchmark problem, which is a multi-material ALE simulation on a 3D unstructured mesh, see Figure 17(b). This comparison uses 4 CPU nodes (144 cores) of LLNL's *rzgenie* machine versus 4 GPU nodes (16 GPUs) of LLNL's *rzansel* machine. Broken over the ALE phases, the observed speedups are $16\times$ in the Lagrange phase, $15\times$ in the remap phase, and $6\times$ in the mesh optimization phase.

8.3. Flow in Urban environments

The urban challenge problem considers the assessment of extreme heat events on buildings in dense urban environments, with up to a few 1000 buildings being modeled during an event. This challenge problem involves coupling of WRF (to define initial weather conditions), Nek5000 (to model heat transfer near buildings), and EnergyPlus (to model heat emissions and energy performance). In collaboration with the ECP-Urban team, CEED team built spectral element meshes and performed LES simulations of Lake Point Tower and Chicago downtown block consisting of 20 buildings as shown in Figure 17(c) (Min et al., 2019a; Shephard et al., 2019). The 20-building mesh comprises

$E = 143340$ spectral elements and its simulation with $N = 13$ is performed using 1024 nodes of ALCF/Mira (32768 MPI ranks). This effort has also generated interest from other federal agencies outside of DOE.

8.4. Atmospheric boundary layer flows

Efficient simulation of atmospheric boundary layer flows (ABL) is important for the study of wind farms, urban canyons, and basic weather modeling. In collaboration with the ExaWind team, we identified an atmospheric boundary layer benchmark problem (Churchfield et al., 2000) to serve as a point of comparison for code and modeling strategies. We have addressed cross-verification and validation of our LES results and corresponding wall models. We demonstrated the suitability of high-order methods for a well-documented stably stratified atmospheric boundary layer benchmark problem, the Global Energy and Water Cycle Experiment (GEWEX) Atmospheric Boundary Layer Study (GABLS) as shown in Figure 17(d). This collaboration will be extended to perform scaling studies to compare the performance of several ABL codes on CPU and GPU platforms.

As another component of the ExaWind collaboration we performed RANS simulations to compute the drag and lift forces of wind-turbine and NACA0012 airfoil structures at Reynolds numbers up to $Re = 10$ million (Min et al., 2019a; Tomov et al., 2019). We investigated several models for the boundary layer treatment in the Nek5000 RANS solver including a wall-resolved regularized approach where we have to use adequate resolution inside the very thin log and viscous sub-layers (Tomboulides et al., 2018) and a stability-enhanced wall-resolved $k-\omega$ and $k-\tau$ models where we do not need such high resolution.

8.5. Pebble-bed reactors

Flow through beds of randomly packed spheres is encountered in many science and engineering applications. The meshing challenge is to have a high-quality all-hex mesh with relatively few elements per sphere. Working with the DOE NEAMS project, the CEED team has developed novel scalable meshing strategies for generating high-quality hexahedral element meshes that ensure accurate representation of densely packed spheres for complex pebble-bed reactor geometries. Our target is to capture highly turbulent flow structures in the solution at minimal cost by using relatively few elements (≈ 300 per sphere) of high order ($p = 7$). Algorithmic strategies including efficient edge collapse, tessellation, smoothing, and projection, are presented in (Kolev et al., 2020) along with quality measurements, flow simulations, validation, and performance results for pebble-bed geometries ranging from hundreds to thousands of pebbles. Figure 17(a) shows a case of 3344 pebbles in an annular domain using 1.1M spectral elements.

8.6. Internal combustion

Turbulence in IC engines presents a challenge for computational fluid dynamics due primarily to the broad range of length and time scales that need to be resolved. Specifically, simulations need to predict the evolution of a variety of flow structures in the vicinity of complex domains that are moving. Executing these simulations accurately and in a reasonable amount of time can ultimately lead to engine design concepts with improved efficiency.

The CEED team has been working with researchers at ETH Zurich (Giannakopoulos et al., 2019) and ANL’s Energy Systems Division (under support from DOE’s Vehicle Technologies Office) on detailed studies of turbulence in the IC engine cycle. We developed a characteristic-based spectral element method for moving-domain problems (Patel et al. 2019), and demonstrated it for the TCC III engine model illustrated in Figure 18(b). We also added a significantly enhanced capability for handling complex moving geometries by adding scalable support for overset grids, referred to as NekNek, based on generalized Schwarz overlapping methods (Mittal et al., 2019). The NekNek multimesh coupling is based entirely on the kernels in Nek’s *gslib* communication library, which has scaled to millions of MPI ranks. A newly developed preconditioner based on the SEM/FEM spectral equivalence was shown to be effective for solving the pressure-Poisson systems in these configurations (Bello-Maldonado and Fischer, 2019).

8.7. Aerosol transport modeling

Related to the current COVID-19 pandemic, the Nek5000 team, in collaboration with NVIDIA and Utah State is researching aerosol transport analysis. High-resolution LES coupled with Lagrangian particle tracking is used for predicting the dynamics of virus-laden aerosols in indoor classroom environments (Dutta et al., 2020). Figure 18(c) demonstrates a recent simulation, using 70 million grid points and 500,000 five-micron aerosols with a future target of 1 billion polydisperse aerosols in a full classroom size.

This application uses efficient algorithms for point containment and general interpolation in physical space, *findpts* and *findeval*, which are available on CPU platforms in CEED’s Nek5000 and MFEM codes. Detailed discussion of these methods and their porting to exascale machines is beyond the scope of this paper. Future developments may include synchronous utilization of the CPU or particle tracking on the device.

8.8. Magnetic fusion

Accurate radio-frequency (RF) heating simulations of fusion systems like the ITER tokamak require the definition of analysis domains that include detailed antenna, reactor wall and physics region geometric representations. As is the case with other wave equation simulations, the application of high-order methods, with their higher rates

of convergence and high flop rate to memory access, is critical for the accurate simulation of these classes of problems. The software components being integrated to address this simulation workflow as part of the DOE SciDAC Center for Integrated Simulation of Fusion Relevant RF Actuators (Bonoli, 2020) must support higher order geometry and high-order analysis methods.

These components include complete curved domain definitions based on CAD system produced models of RF antenna geometries and geometry construction tools for the analysis domain that support defeaturing the antenna CAD models as desired, and combining the CAD geometry with the reactor wall geometry and any other “physics.” Historically, ad-hoc methods are employed to execute this time consuming step. Recent efforts have focused on providing a graphical construction tool for tokamak systems building on general geometry manipulation capabilities (Simmetrix, 2020). The user interface to the PetraM (Physics Equation-Translator for MFEM) component (Shiraiwa et al., 2017) supports the association of the RF simulation material properties, loads and boundary conditions (essential and natural) to the analysis domain geometry. With these in place, fully automatic mesh generation of an initial curved tetrahedral mesh using either Gmsh (Geuzaine and Remacle, 2013) or MeshSim (Simmetrix, 2020) is performed. If the curved meshes are not of sufficiently high order for the basis functions to be used, a tool has been developed to increase the order of approximation of mesh edges and faces on curved domain boundaries up to order six. We then execute a high-order MFEM simulation supplemented with PetraM routines to control needed field information to perform the RF analysis. This is followed by error estimation and mesh adaptation using the conforming curved mesh adaptation procedure in PUMI/MeshAdapt and return to the analysis step until acceptable solution accuracy is obtained.

Figure 18(d) shows an example of the application of the basic steps in the workflow of a tokamak geometry with RF antenna geometry inserted. The result shown is for a low-order mesh. Recent results up to order five are showing a clear advantage to use of higher order elements.

9. Conclusion

In this paper we reviewed the co-design activities in the Center for Efficient Exascale Discretizations of the Exascale Computing Project, focused on the computational motif of PDE discretizations on general unstructured grids, with emphasis on high-order methods. We described our co-design approach together with the mathematical and algorithmic foundations for high-order finite element discretizations. We also reviewed other topics that are necessary for a complete high-order ecosystem, such as matrix-free linear solvers and high-order mesh adaptivity, which are still active areas of research. A number of freely available software products are being actively developed and supported by the center, including

benchmarks, miniapps and high- and low-level library APIs. The CEED team is very much interested in collaborations and feedback, please visit our website ceed.exascaleproject.org to get in touch.

Authors' note

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.





Declaration of conflicting interests


The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.


Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research is supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering and early testbed platforms, in support of the nation's exascale computing imperative. The research used resources of the Argonne Leadership Computing Facility, which is supported by the U.S. Department of Energy, Office of Science, under Contract DE-AC02-06CH11357. This research also used resources of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract DE-AC05-00OR22725. Work performed under the auspices of the U.S. Department of Energy under Contract DE-AC52-07NA27344 (LLNL-JRNL-814059).

ORCID iD

Tzanio Kolev  <https://orcid.org/0000-0002-2810-3090>
 Misun Min  <https://orcid.org/0000-0002-5646-5689>
 Valeria Barra  <https://orcid.org/0000-0003-1129-2056>
 Natalie Beams  <https://orcid.org/0000-0001-6060-4082>

Yohann Dudouit  <https://orcid.org/0000-0001-5831-561X>

Cameron W Smith  <https://orcid.org/0000-0001-9258-5226>

Supplemental material

Supplemental material for this article is available online.

Notes

1. The BKs are *bake-off kernels* that involve only the local, elementwise, portion of the matrix-vector products, $A\mathbf{u}$ or $B\mathbf{u}$ associated with the corresponding BPs described in Section 4.
2. A backend may choose to operate incrementally without forming explicit **E**- or **Q**-vectors.

References

- Abdelfattah A, Baboulin M, Dobrev V, et al. (2016a) High-performance tensor contractions for GPUs. In: *International Conference on Computational Science 2016, ICCS 2016*, 6–8 June 2016, San Diego, California, USA, pp. 108–118.
- Abdelfattah A, Barra V, Beams N, et al. (2020) libCEED User Manual. DOI:10.5281/zenodo.4302737.
- Abdelfattah A, Haidar A, Tomov S, et al. (2016b) Performance, design, and autotuning of batched GEMM for GPUs. In: *High Performance Computing—31st International Conference, ISC High Performance 2016*, Frankfurt, Germany, 19–23 June 2016, pp. 21–38.
- Ameen M, Patel S, Colmenares J, et al. (2020) Direct Numerical Simulation (DNS) and high-fidelity large-eddy simulations for improved prediction of in-cylinder flow and combustion processes. Technical report, DOE Vehicle Technologies Office Annual Merit Review.
- Anderson R, Andrej J, Barker A, et al. (2020) MFEM: a modular finite element library. *Computers & Mathematics with Applications* 81: 42–74.
- Anderson RW, Dobrev VA, Kolev TV, et al. (2015) Monotonicity in high-order curvilinear finite element arbitrary Lagrangian–Eulerian remap. *International Journal for Numerical Methods in Engineering* 77(5): 249–273.
- Anderson RW, Dobrev VA, Kolev TV, et al. (2017) High-order local maximum principle preserving (MPP) discontinuous Galerkin finite element method for the transport equation. *Journal of Computational Physics* 334: 102–124.
- Anderson RW, Dobrev VA, Kolev TV, et al. (2018) High-order multi-material ALE hydrodynamics. *SIAM Journal on Scientific Computing* 40(1): B32–B58.
- Antonietti PF, Sarti M, Verani M, et al. (2016) A uniform additive Schwarz preconditioner for high-order discontinuous Galerkin approximations of elliptic problems. *Journal of Scientific Computing* 70(2): 608–630.
- Balay S, Abhyankar S, Adams MF, et al. (2019) PETSc Web page. Available at: <https://www.mcs.anl.gov/petsc> (accessed 26 May 2021).
- Barra V, Brown J, Thompson J, et al. (2020) High-performance operator evaluations with ease of use: libCEED's Python interface. In: Agarwal M, Calloway C, Niederhut D, and Shupe D

- (eds) *Proceedings of the 19th Python in Science Conference*, Austin, Texas, 6–12 July 2020, pp. 85–90.
- Beckingsale DA, Burmark J, Hornung R, et al. (2019) RAJA: portable performance for large-scale scientific applications. In: *IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*, GA, USA, 13 November 2020.
- Bello-Maldonado PD and Fischer PF (2019) Scalable low-order finite element preconditioners for high-order spectral element Poisson solvers. *SIAM Journal on Scientific Computing* 41(5): S2–S18.
- Bonoli P (2020) Center for integrated simulation of fusion relevant RF actuators. Available at: <https://www.rfscidac4.org/home> (accessed 26 May 2021).
- Brown C, Abdelfattah A, Tomov S, et al. (2020a) Design, optimization, and benchmarking of dense linear algebra algorithms on AMD GPUs. Technical Report ICL-UT-20-12, University of Tennessee, USA.
- Brown C, Abdelfattah A, Tomov S, et al. (2020b) hipMAGMA v2.0.0. DOI: 10.5281/zenodo.3928667.
- Brown J, Dobrev V, Dutta S, et al. (2018) Propose high-order mesh/data format. Technical Report CEED-MS18, Exascale Computing Project. DOI: 10.5281/zenodo.2542346.
- Brown J, Dobrev V, Fischer P, et al. (2017) Initial Integration of CEED Software in ECP/CEED Applications. Technical Report CEED-MS8, Exascale Computing Project. DOI: 10.5281/zenodo.2542338.
- Brown J, He Y and MacLachlan S (2019) Local Fourier analysis of BDDC-like algorithms. *SIAM Journal on Scientific Computing* 41: S346–S369.
- Brubeck P and Fischer P (2019) Fast diagonalization preconditioning for nonsymmetric spectral element problems. *ANL/MCS-P9200-0719*.
- Brubeck P, Kaneko K, Lan Y, et al. (2020) Schwarz preconditioned spectral element methods for steady flow and heat transfer. *ANL/MCS-P9199-0719*.
- Canuto C (1994) Stabilization of spectral methods by finite element bubble functions. *Computer Methods in Applied Mechanics and Engineering* 116(1–4): 13–26.
- Canuto C, Hussaini MY, Quarteroni A, et al. (2006) *Spectral Methods: Fundamentals in Single Domains*. Berlin Heidelberg: Springer.
- Cervený J, Dobrev V and Kolev T (2019) Non-conforming mesh refinement for high-order finite elements. *SIAM Journal on Scientific Computing* 41(4): C367–C392.
- Chalmers N and Warburton T (2018) Low-order preconditioning of high-order triangular finite elements. *SIAM Journal on Scientific Computing* 40(6): A4040–A4059.
- Chalmers N, Karakus A, Austin AP, et al. (2020) libParanumal: a performance portable high-order finite element library [Software]. Available at: <https://github.com/paranumal/libparanumal>. Release 0.3.1 (accessed 26 May 2021).
- Churchfield M, Lee S and Moriarty P (2000) Adding complex terrain and stable atmospheric condition capability to the OpenFOAM-based flow solver of the simulator for on/off-shore wind farm application (SOWFA). Technical Report NREL/CP-5000-58539, NREL.
- Dewille M, Fischer P and Mund E (2002) *High-Order Methods for Incompressible Fluid Flow*. Cambridge: Cambridge University Press.
- Dobrev V, Knupp P, Kolev T, et al. (2019) The target-matrix optimization paradigm for high-order meshes. *SIAM Journal on Scientific Computing* 41(1): B50–B68.
- Dobrev VA, Knupp P, Kolev TV, et al. (2020) Simulation-driven optimization of high-order meshes in ALE hydrodynamics. *Computers & Fluids* 208: 104602. DOI: 10.1016/j.compfluid.2020.104602
- Dobrev VA, Kolev TV and Rieben RN (2012) High-order curvilinear finite element methods for Lagrangian hydrodynamics. *SIAM Journal on Scientific Computing* 34(5): B606–B641.
- Dobrev VA, Kolev TV, Rieben RN, et al. (2016) Multi-material closure model for high-order finite element Lagrangian hydrodynamics. *International Journal for Numerical Methods in Engineering* 82(10): 689–706.
- Dobrev VA, Lazarov RD, Vassilevski PS, et al. (2006) Two-level preconditioning of discontinuous Galerkin approximations of second-order elliptic equations. *Numerical Linear Algebra with Applications* 13(9): 753–770.
- Dohrmann C (2003) A preconditioner for substructuring based on constrained energy minimization. *SIAM Journal on Scientific Computing* 25: 246.
- Dongarra J, Duff I, Gates M, et al. (2016) A proposed API for Batched Basic Linear Algebra Subprograms. MIMS EPrint 2016.25, Manchester Institute for Mathematical Sciences, The University of Manchester. Available at: <http://eprints.ma.man.ac.uk/2464/> (accessed 26 May 2021).
- Dongarra J, Duff I, Gates M, et al. (2018) Batched BLAS (basic linear algebra subprograms) 2018 specification. ICL-UTK technical report. Available at: <https://www.icl.utk.edu/files/publications/2018/icl-utk-1170-2018.pdf> (accessed 26 May 2021).
- Dutta S, Fischer P, Shyuan C, et al. (2020) On turbulence and particle transport in closed rooms. *American Physical Society, Division of Fluid Dynamics* submitted for publication. Under review.
- Farin G (2014) *Curves and Surfaces for Computer-Aided Geometric Design: A Practical Guide*. Amsterdam: Elsevier.
- Feuillet R, Loseille A, Marcum D, et al. (2018) Connectivity-change moving mesh methods for high-order meshes: toward closed advancing-layer high-order boundary layer mesh generation. In: *2018 Fluid Dynamics Conference*, Atlanta, Georgia, 25–29 June 2018. pp. 4167. DOI: 10.2514/6.2018-4167.
- Fischer P, Heisey K and Min M (2015) Scaling limits for PDE-based simulation (invited). In: *22nd AIAA Computational Fluid Dynamics Conference, AIAA Aviation*. Dallas, TX, 22–26 June 2015, AIAA, pp. 2015–3049.
- Fischer P, Min M, Rathnayake T, et al. (2020) Scalability of high-performance PDE solvers. *The International Journal of High Performance Computing Applications* 34(5): 562–586.
- FMS (2020) FMS: High-order field and mesh specification [Software]. Available at: <https://github.com/CEED/FMS> (accessed 26 May 2021).
- Geuzaine C and Remacle JF (2013) Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-

- processing facilities [Software]. Available at: <http://gmsh.info/> (accessed 26 May 2021).
- Giannakopoulos G, Frouzakis C, Fischer P, et al. (2019) LES of the gas-exchange process inside an internal combustion engine using a high-order method. *Flow, Turbulence and Combustion* 104: 673–692.
- GLVis (2020) GLVis: OpenGL finite element visualization tool [Software]. Available at: <https://glvis.org> (accessed 26 May 2021).
- Hajduk H, Kuzmin D, Kolev TV, et al. (2020a) Matrix-free sub-cell residual distribution for Bernstein finite element discretizations of linear advection equations. *Computer Methods in Applied Mechanics and Engineering* 359: 112658.
- Hajduk H, Kuzmin D, Kolev TV, et al. (2020b) Matrix-free sub-cell residual distribution for Bernstein finite elements: monolithic limiting. *Computers & Fluids* 200: 104451.
- Ibanez D (2016a) Omega_h GitHub repository [Software]. Available at: https://github.com/ibaned/omega_h (accessed 26 May 2021).
- Ibanez D, Seol E, Smith C, et al. (2016) Pumi: parallel unstructured mesh infrastructure. *ACM Transactions on Mathematical Software (TOMS)* 42(3): 17.
- Ibanez DA (2016b) *Conformal Mesh Adaptation on Heterogeneous Supercomputers*. Troy, NY: Rensselaer Polytechnic Institute.
- Karakus A, Chalmers N, Hesthaven JS, et al. (2019a) Discontinuous Galerkin discretizations of the Boltzmann–BGK equations for nearly incompressible flows: semi-analytic time stepping and absorbing boundary layers. *Journal of Computational Physics* 390: 175–202.
- Karakus A, Chalmers N, Swirydowicz K, et al. (2019b) A gpu accelerated discontinuous Galerkin incompressible flow solver. *Journal of Computational Physics* 390: 380–404.
- Knoll D and Keyes D (2004) Jacobian-free Newton-Krylov methods: a survey of approaches and applications. *Journal of Computational Physics* 193: 357–397.
- Kolev T, Fischer P, Abdelfattah A, et al. (2020) Improve performance and capabilities of CEED-enabled ECP applications on Summit/Sierra. Technical Report CEED-MS34, Exascale Computing Project. DOI: 10.5281/zenodo.3860804.
- Kolev TV and Vassilevski PS (2009) Parallel auxiliary space AMG for $h(\text{curl})$ problems. *Journal of Computational Mathematics* 27(5): 604–623.
- Kronbichler M and Ljungkvist K (2019) Multigrid for matrix-free high-order finite element computations on graphics processors. *ACM Transactions on Parallel Computing* 6(1): 1–32.
- Laghos (2020) Laghos: High-order Lagrangian hydrodynamics miniapp [Software]. Available at: <https://github.com/ceed/Laghos> (accessed 26 May 2021).
- Lottes JW and Fischer PF (2005) Hybrid multigrid/Schwarz algorithms for the spectral element method. *Journal of Scientific Computing* 24(1): 45–78.
- Lu Q, Shephard MS, Tendulkar S, et al. (2014) Parallel mesh adaptation for high-order finite element methods with curved element geometry. *Engineering with Computers* 30(2): 271–286.
- Luo XJ, Shephard MS, O’bara RM, et al. (2004) Automatic p-version mesh generation for curved domains. *Engineering with Computers* 20(3): 273–285.
- MAGMA (2020) MAGMA: Matrix algebra on gpu and multicore architectures [Software]. Available at: <https://icl.utk.edu/magma> (accessed 26 May 2021).
- Martinez J, Lan YH, Merzari E, et al. (2019) On the use of LES-based turbulent thermal-stress models for rod bundle simulations. *International Journal of Heat and Mass Transfer* 142: 118399.
- Masliyah I, Abdelfattah A, Haidar A, et al. (2016) High-Performance Matrix-Matrix Multiplications of Very Small Matrices. In: *Euro-Par 2016: Parallel Processing—22nd International Conference on Parallel and Distributed Computing*, Grenoble, France, 24–26 August 2016, pp. 659–671.
- Medina DS, St-Cyr A and Warburton T (2014) OCCA: a unified approach to multi-threading languages. *arXiv preprint arXiv:1403.0968*.
- Merzari E, Rahaman R, Patel S, et al. (2017) Cfd smr assembly performance baselines with nek5000. Technical Report ECP-SE-08-47, DOE ECP ExaSMR Milestone Report.
- MFEM (2020) MFEM: Modular finite element methods [Software]. Available at: <https://mfem.org> (accessed 26 May 2021).
- Min M, Camier JS, Fischer P, et al. (2019a) Engage second wave ECP/CEED applications. Technical Report CEED-MS23, Exascale Computing Project. DOI: 10.5281/zenodo.2542359.
- Min M, Fischer P, Tomov V, et al. (2017) Engage first wave ECP/CEED applications. Technical Report CEED-MS1, Exascale Computing Project. DOI: 10.5281/zenodo.2542292.
- Min M, Tomboulides A, Fischer P, et al. (2019b) Nek5000 enhancements for faster running analysis. Technical Report ANL.MCS-TM-384, ANL NEAMS Report.
- Mittal K, Dutta S and Fischer P (2019) Nonconforming Schwarz-spectral element methods for incompressible flow. *Computers and Fluids* 191: 104237.
- Nek5000 (2020) Nek: Open source, highly scalable and portable spectral element code [Software]. Available at: <https://nek5000.mcs.anl.gov> (accessed 26 May 2021).
- OCCA (2020) OCCA: lightweight performance portability library [Software]. Available at: <https://libocca.org/> (accessed 26 May 2021).
- Orszag S (1980) Spectral methods for problems in complex geometry. *Journal of Computational Physics* 37: 70–92.
- Otten M, Gong J, Mametjanov A, et al. (2016) An MPI/OpenACC implementation of a high order electromagnetics solver with GPUDirect communication. *The International Journal of High Performance Computing Applications* 30: 320–334.
- Patel S, Fischer P, Min M, et al. (2019) A characteristic-based, spectral element method for moving-domain problems. *Journal of Scientific Computing* 79: 564–592.
- Patera A (1984) A spectral element method for fluid dynamics: laminar flow in a channel expansion. *Journal of Computational Physics* 54: 468–488.
- Pavarino L, Widlund O and Zampini S (2010) BDDC preconditioners for spectral element discretizations of almost incompressible elasticity in three dimensions. *SIAM Journal on Scientific Computing* 32: 3604.

- Pazner W (2020) Efficient low-order refined preconditioners for high-order matrix-free continuous and discontinuous Galerkin methods. *SIAM Journal on Scientific Computing* 42(5): A3055–A3083.
- Pazner W and Persson PO (2018) Approximate tensor-product preconditioners for very high order discontinuous Galerkin methods. *Journal of Computational Physics* 354: 344–369.
- Raffenetti K, Amer A, Oden L, et al. (2017) Why is MPI so Slow? Analyzing the Fundamental Limits in Implementing MPI-3.1. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. Denver, Colorado, New York, NY, USA: Association for Computing Machinery. DOI: 10.1145/3126908.3126963.
- RAJA (2020) RAJA performance portability layer [Software]. Available at: <https://github.com/LLNL/RAJA> (accessed 26 May 2021).
- Remhos (2020) Remhos: High-order remap miniapp [Software]. Available at: <https://github.com/ceed/Remhos> (accessed 26 May 2021).
- Shephard M, Barra V, Brown J, et al. (2019) Improved Support for Parallel Adaptive Simulation in CEED. Technical Report CEED-MS29, Exascale Computing Project. DOI: 10.5281/zenodo.3336420.
- Shiraiwa S, Wright J, Bonoli P, et al. (2017) Rf wave simulation for cold edge plasmas using the mfem library. In: *EPJ Web of Conferences*, Vol.157. Les Ulis: EDP Sciences, p. 03048.
- Simmetrix (2020) Simmetrix: Enabling simulation-based design. Available at: <http://www.simmetrix.com/> (accessed 26 May 2021).
- Sundar H, Stadler G and Biros G (2015) Comparison of multigrid algorithms for high-order continuous finite element discretizations. *Numerical Linear Algebra with Applications* 22(4): 664–680.
- Swirydowicz K, Chalmers N, Karakus A, et al. (2019) Acceleration of tensor-product operations for high-order finite element methods. *The International Journal of High Performance Computing Applications* 33(4): 735–757.
- Tomboulides A, Aithal M, Fischer P, et al. (2018) A novel numerical treatment of the near-wall regions in the k - ω class of the rans models. *International Journal of Heat and Fluid Flow* 72: 186–199.
- Tomov S, Abdelfattah A, Barra V, et al. (2019) Performance tuning of CEED software and 1st and 2nd wave apps. Technical Report CEED-MS32, Exascale Computing Project. DOI: 10.5281/zenodo.3477618.
- Tomov S, Bello-Maldonado P, Brown J, et al. (2018) Performance tuning of CEED software and first wave apps. Technical Report CEED-MS20, Exascale Computing Project. DOI:10.5281/zenodo.2542350.
- Tsai PH, Lan YH, Fisher P, et al. (2020) Drift-diffusion solvers. Part II: STEADY PROBLEMS. *ANL/MCS-P9295-0420*.
- VisIt (2020) VisIt: a distributed, parallel visualization and analysis tool [Software]. Available at: <https://visit.llnl.gov>. DOI: 10.11578/dc.20171025.on.1019 (accessed 26 May 2021).
- Zampini S (2016) PCBDDC: a class of robust dual-primal methods in PETSc. *SIAM Journal on Scientific Computing* 38(5): S282–S306.

Author biographies

Tzanio Kolev is a computational mathematician at the Center for Applied Scientific Computing (CASC) in Lawrence Livermore National Laboratory, where he works on finite element discretizations and solvers for problems in compressible shock hydrodynamics, multi-material arbitrary Lagrangian–Eulerian methods, radiation hydrodynamics, and computational electromagnetics. He won an R&D100 award as a member of the *hypre* team. Tzanio is leading the high-order finite element discretization research and development efforts in the MFEM and BLAST projects in CASC and is the director of the Center for Efficient Exascale Discretizations (CEED) in DOE’s Exascale Computing Project.

Paul Fischer is a Blue Waters Professor of Computer Science and Mechanical Science and Engineering at the University of Illinois at Urbana-Champaign and an Argonne senior scientist. He is a Fellow of the American Association for the Advancement of Science (AAAS). He is the chief architect of the fluid thermal simulation code Nek5000, which scales to over a million processors and has been recognized with the Gordon Bell Prize in high-performance computing. Nek5000 is used by over 400 researchers worldwide in a variety of thermal-fluids applications. He was the deputy director for the DOE-ASCR Co-Design Center for Exascale Simulation of Advanced Reactor (CESAR) and is currently the deputy director for the DOE-ECP Co-Design Center for Efficient Exascale Discretizations (CEED).

Misun Min is a computational scientist in the Mathematics and Computer Science Division at Argonne National Laboratory. Her research focuses on developing scalable high-order algorithms and software for solving electromagnetic, drift-diffusion, and fluid systems on advanced high-performance computing architectures. She is the author of the spectral element discontinuous Galerkin (SEDG) electromagnetics simulation code, NekCEM, which scales to over million CPU cores and tens of thousands GPUs. She won an R&D100 award on “NekCEM/Nek5000: Scalable High-Order Simulation Codes.” She is a PI on the DOE Applied Mathematics Research project “High-Order Methods for High-Performance Multiphysics Simulations” and the Argonne PI of the DOE-ECP Co-Design project, CEED.

Jack Dongarra is University Distinguished Professor of Computer Science in the Electrical Engineering and Computer Science Department at the University of Tennessee. He holds the position of a Distinguished Research Staff member in the Computer Science and Mathematics Division at Oak Ridge National Laboratory, Turing Fellowship

in the School of Mathematics at the University of Manchester, and is an adjunct professor in the Computer Science Department at Rice University. He served as a faculty fellow at Texas A&M University's institute for advanced study (2014–2018). Dongarra is the founding director of Innovative Computing Laboratory.

Jed Brown is an assistant professor of computer science at the University of Colorado Boulder. He is a developer of PETSc and specializes in multiscale and high-order numerical methods for geoscience and engineering applications. His work has been recognized by the 2014 SIAG/SC Junior Scientist Prize and the 2014 IEEE TCSC Young Achiever award, and he was co-recipient of the 2015 SIAM/ACM Prize in Computational Science and Engineering.

Veselin Dobrev is a computational mathematician in the numerical analysis and simulations group in the Center for Applied Scientific Computing. His research interests are in the areas of numerical methods for solving PDEs, which include finite element and discontinuous Galerkin methods, shock hydrodynamics simulations, and iterative and multi-grid methods. Veselin received his Ph.D. in mathematics from Texas A&M University in 2007. He is currently working on high-order curvilinear finite elements for Lagrangian hydrodynamics (BLAST project).

Tim Warburton is the John K Costain Faculty Chair in the College of Science and a professor of mathematics at Virginia Tech. He also currently holds an appointment in the Department of Computational and Applied Mathematics at Rice University. He developed the first high-order nodal discontinuous Galerkin solver for time-domain electromagnetics on unstructured grids and led a decadal project to accelerate these methods by devising parallel local time-stepping methods, GPU-acceleration, co-volume filtering techniques, novel rational bases for curvilinear elements, and numerous other innovations. He co-authored the first comprehensive book on discontinuous Galerkin methods. He created the Open Concurrent Compute Abstraction (OCCA) as part of the CESAR co-design center at Argonne. The OCCA framework enables domain scientists and computational scientists to write portable threaded code. The OCCA library has been used as a foundational layer for higher order finite element, spectral element, discontinuous Galerkin, and finite difference PDE solvers for industrial applications and lab miniapps.

Stanimire Tomov is Research Assistant Professor at the University of Tennessee, Knoxville. He specializes in parallel algorithms, data analytics, and high-performance scientific computing. His current work is concentrated on the

development of numerical linear algebra software, and in particular the MAGMA libraries for a modernized LAPACK on new architectures. He has lead and contributed to numerous NSF, DOE, and DOD-funded HPC projects and industry collaborations.

Mark S Shephard is the Samuel A. and Elisabeth C. Johnson, Jr. Professor of Engineering, and the director of the Scientific Computation Research Center at Rensselaer Polytechnic Institute. He holds joint appointments in the departments of Mechanical, Aerospace and Nuclear Engineering and Computer Science. His research has led to well recognized contributions in the areas of automatic mesh generation, automated/adaptive methods, and parallel simulation technologies. He is a past President of the US Association for Computational Mechanics. He is editor of Engineering with Computers, an associate editor for the SIAM Journal on Scientific Computing. He was a co-founder of Simmetrix Inc., a company dedicated to the technologies that enable simulation-based engineering.

Ahmad Abdelfattah is a research scientist at the Innovative Computing Laboratory, the University of Tennessee. He received his PhD in computer science from King Abdullah University of Science and Technology (KAUST) in 2015, where he was a member of the Extreme Computing Research Center (ECRC). His research interests include numerical linear algebra, parallel algorithms, and performance optimization on massively parallel processors. He received his BSc. and MSc. degrees in computer engineering from Ain Shams University, Egypt.

Valeria Barra is a postdoctoral research associate at the University of Colorado Boulder. She received her PhD in Applied Mathematics from NJIT in 2018 with a dissertation on computational fluid dynamics applications for viscoelastic media. She is fond of discretizations and numerical methods for PDEs, and computational geometry. She is now part of the libCEED team, mainly focusing on the development of miniapps. Valeria was selected as a Rising Star in Computational and Data Sciences from the Oden Institute of the University of Texas at Austin in 2020, and she is going to join Caltech soon as a Research Software Engineer working in the CliMA project.

Natalie Beams is a research scientist at the Innovative Computing Laboratory of the University of Tennessee, Knoxville, where she is currently a member of the CEED and PEEKS teams for the Exascale Computing Project. She is interested in all aspects of numerical methods for PDEs, from algorithm creation to efficient implementation, with a particular emphasis on high-order methods.

Jean-Sylvain Camier is a research scientist in the numerical analysis and simulations group at the Center for Applied Scientific Computing in Lawrence Livermore National Laboratory. His current research focus is on parallel computer architectures, performance portability, and computing with high-order finite element methods.

Noel Chalmers received his PhD in Applied Mathematics from the University of Waterloo. He currently works at Advanced Micro Devices (AMD) Research, developing high-performance computing scientific software for exascale computing. His research centers on GPU-acceleration of scientific computing applications, high-order finite element methods, and scalable multilevel preconditioning algorithms.

Yohann Dudouit is a computational mathematician at the Center for Applied Scientific Computing in Lawrence Livermore National Laboratory. His research currently focuses on high performance computing for high-order finite element methods, with an emphasis on the implementation of matrix-free finite element methods and discontinuous Galerkin methods on GPUs.

Ali Karakus is an assistant professor of mechanical engineering at Middle East Technical University. His research focuses on high-order finite element methods, multi-phase flows, and GPU-accelerated scientific computing. He worked on CEED as a postdoctoral researcher at Argonne and Virginia Tech.

Ian Karlin is currently a Principal HPC Strategist with the Livermore Computing Advanced Technology Office (ATO) at Lawrence Livermore National Laboratory (LLNL). He received a PhD degree in computer science from the University of Colorado, Boulder, CO, USA, in 2011. He leads LLNL procurement benchmarking efforts, is the ATO lead for machine learning hardware, assists with preparedness activities for future large-scale machines, and provides user support. His research interests include system design, performance optimization, and workload characterization.

Stefan Kerkemeier is the lead developer for NekRS and a principal developer of Nek5000. As an Argonne affiliate, he has been collaborating with the Mathematics and Computer Science and Nuclear Engineering divisions since 2008. He holds a Ph.D. in Mechanical Engineering from ETH Zurich, where he focused on direct numerical simulation of turbulent combustion with detailed chemistry.

Yu-Hsiang Lan is a predoctoral appointee at the Mathematics and Computer Science Division of Argonne National Laboratory. He works on high performance computing, high-order methods, and scalable high-order

meshing. He received an MS in the Department of Mathematics at National Taiwan University.

David Medina received his PhD in Computational and Applied Mathematics from Rice University. He is the founder of Occalytics LLC, where he works on the OCCA library and other related high performance computing projects.

Elia Merzari is an associate professor in the Department of Nuclear Engineering at Penn State. He has published extensively in the field of massively parallel Computational Fluid Dynamics simulations of Turbulence applied to nuclear engineering. He has received allocations for over 1 billion core-hours on various supercomputers in the last decade including Summit and Mira.

Aleksandr Obabko is a principal computational engineer in the Computational Science (CPS) Division at Argonne. He is actively working on advanced reactor thermal-hydraulic modeling focusing on verification and validation. He is currently the PI on the thermal-hydraulic component of the NEAMS reactor suite, and has been active in development of MHD capabilities of Nek5000 and in multiphysics code coupling and simulations. His research has been recognized with the HPC Innovation Excellence Award and an R&D100 award.

Will Pazner is a computational mathematician at the Center for Applied Scientific Computing (CASC) in Lawrence Livermore National Laboratory, where his research focuses on discretizations and solvers for high-order finite element methods, with a focus on matrix-free discontinuous Galerkin methods and compressible fluid flow applications. Previously, he was CASC's Sidney Fernbach Postdoctoral Fellow.

Thilina Rathnayake is a PhD candidate in computer science at UIUC working on scalable partitioners and solvers and on algorithms for high-order methods targeting GPU platforms. He was a Givens Associate at Argonne, summer intern at LLNL, and most recently an intern with Intel.

Cameron W Smith is a Senior Research Scientist at Rensselaer Polytechnic Institute's Scientific Computation Research Center (SCOREC) with expertise in simulation automation, parallel computation, load balancing, unstructured meshing, and heterogeneous architectures. His contributions to the software libraries PUMI, ParMA, EnGPar, and PUMIPic, have advanced the analysis capabilities and performance of applications in computational fluid dynamics, electromagnetics, and fusion plasma physics. Collaborators on these projects have come from multiple universities, United States Department of Energy

Laboratories, NASA, and several companies from the software development, engineering, manufacturing, and information technology sectors.

Lukas Spies is a graduate student in computer science at UIUC. He conducted the NekBench test suite development as a Givens Associate in the Mathematics and Computer Science Division at Argonne National Laboratory.

Kasia Swirydowicz received her PhD from the Department of Mathematics at Virginia Tech. She is currently working as a computational scientist at the Pacific Northwest National Laboratory. Her interests include numerical linear algebra, Krylov subspace solvers and preconditioners, and using high performance computing to create highly optimized scientific software.

Jeremy Thompson is a graduate student in Applied Mathematics at the University of Colorado Boulder. He is a developer of libCEED and a former U.S. Air Force Operations Research Analyst. His research focuses include performance portability and preconditioning for high-order numerical methods for PDEs. He won the U.S. Air Force

Air Combat Command Scientist of the Year, Junior Military Category in 2011.

Ananias Tomboulides is Professor of Mechanical Engineering at the Aristotle University of Thessaloniki and an Argonne affiliate. He has a long-standing collaboration with the Mathematics and Computer Science Division at Argonne National Laboratory and is one of the main developers of the open source CFD code Nek5000, with contributions in the areas of low Mach number combustion, multi-phase boiling flows, high-order RANS models and moving geometries in ICEs. He has served as Chairman of ERCOFTAC (the European Research Community On Flow, Turbulence And Combustion) during 2012-2019 and is member of the Editorial Board of the Journal Flow Turbulence and Combustion (FTAC).

Vladimir Tomov is a computational mathematician at the Center for Applied Scientific Computing in Lawrence Livermore National Laboratory. His research interests include finite element methods for multi-material ALE and radiation hydrodynamics and high-order mesh optimization.