



EXASCALE COMPUTING PROJECT

ECP Milestone Report

Public release of CEED 2.0

WBS 2.2.6.06, Milestone CEED-MS25

Jed Brown
Ahmad Abdelfattah
Valeria Barra
Veselin Dobrev
Yohann Dudouit
Paul Fischer
Tzanio Kolev
David Medina
Misun Min
Thilina Ratnayaka
Cameron Smith
Jeremy Thompson
Stanimire Tomov
Vladimir Tomov
Tim Warburton

April 3, 2019

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website <http://www.osti.gov/scitech/>

Reports produced before January 1, 1996, may be purchased by members of the public from the following source:

National Technical Information Service

5285 Port Royal Road

Springfield, VA 22161

Telephone 703-605-6000 (1-800-553-6847)

TDD 703-487-4639

Fax 703-605-6900

E-mail info@ntis.gov

Website <http://www.ntis.gov/help/ordermethods.aspx>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information

PO Box 62

Oak Ridge, TN 37831

Telephone 865-576-8401

Fax 865-576-5728

E-mail reports@osti.gov

Website <http://www.osti.gov/contact.html>

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express 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 any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ECP Milestone Report
Public release of CEED 2.0
WBS 2.2.6.06, Milestone CEED-MS25

Office of Advanced Scientific Computing Research
Office of Science
US Department of Energy

Office of Advanced Simulation and Computing
National Nuclear Security Administration
US Department of Energy

April 3, 2019

ECP Milestone Report
Public release of CEED 2.0
WBS 2.2.6.06, Milestone CEED-MS25

Approvals

Submitted by:

Tzanio Kolev, LLNL
CEED PI

Date

Approval:

Andrew R. Siegel, Argonne National Laboratory
Director, Applications Development
Exascale Computing Project

Date

Revision Log

Version	Creation Date	Description	Approval Date
1.0	April 3, 2019	Original	

EXECUTIVE SUMMARY

In this milestone, we created and made publicly available the second full CEED software distribution, release CEED 2.0, consisting of software components such as MFEM, Nek5000, PETSc, MAGMA, OCCA, etc., treated as dependencies of CEED. The release consists of 12 integrated Spack packages for libCEED, mfem, nek5000, nekcec, laghos, nekbone, hpgmg, occa, magma, gslib, petsc and pumi plus the CEED *meta-package*.

The artifacts delivered include a consistent build system based on the above 13 Spack packages, documentation and verification of the build process, as well as improvements in the integration between different CEED components. As part of CEED 2.0, we also released the next version of libCEED (v0.4), which contains four new CPU backends, two new GPU backends, CPU backend optimizations, initial support for operator composition, performance benchmarking, and a Navier-Stokes demo. See the CEED website, <http://ceed.exascaleproject.org/ceed-2.0/> and the CEED GitHub organization, <http://github.com/ceed> for more details.

In addition to details and results from these efforts, in this document we report on other project-wide activities performed in Q2 of FY19, including: a version of the Laghos miniapp for CTS-2 procurement, extensive GPU performance baseline benchmarking for the Urban and ExaSMR applications, papers, minisymposium presentations, and other outreach efforts.

TABLE OF CONTENTS

Executive Summary	vi
List of Figures	viii
List of Tables	ix
1 Introduction	1
2 CEED 2.0 Packages	1
2.1 libCEED 0.4	1
2.1.1 New libCEED Backends	1
2.1.2 Operator Composition	1
2.1.3 Benchmarking libCEED Performance	1
2.1.4 Navier-Stokes example	4
2.2 MFEM 3.4	7
2.3 Nek5000 version 19.0-rc2	9
2.4 PETSc 3.11	10
2.5 OCCA 1.0.8	10
2.6 PUMI 2.2.0	11
2.7 MAGMA 2.5.0	11
2.8 Laghos 2.0	11
2.9 HPGMG 0.4	11
2.10 gslib v.1.0.4	12
2.11 NekCEM	12
3 CEED 2.0 Testing and Distribution	12
3.1 Mac	12
3.2 Linux RHEL7	13
3.3 Linux Ubuntu	14
3.4 Cori	14
3.5 ALCF Theta	16
3.6 Other machines: OLCF Summit, LLNL TOSS3 and Lassen	17
4 GPU Performance Baselines for CEED Applications	17
4.1 Urban: libParanumal performance on Summit vs. Titan	17
4.2 Urban: NekCEM performance on Summit vs. Titan	17
4.3 ExaSMR: Speedup of NekRS on Summit vs. Nek5000 on Titan	19
4.4 ExaSMR: Speedup of Steady Thermal Solver vs. Transient Calculation	19
5 Other Project Activities	20
5.1 Laghos CTS-2 benchmark	20
5.2 Outreach	21
6 Conclusion	21

LIST OF FIGURES

1	Pure C Reference Backends	2
2	AVX Backends	2
3	LIBXSMM Backends	3
4	Pure C Reference Backends	3
5	AVX Backends	4
6	LIBXSMM Backends	4
7	Advection problem. A blob of energy is advected by a uniform circular flow, with constant speed $ \mathbf{u} = 0.5$ m/s. On the left, the initial condition at $t = 0$ s; on the right, the blob has been transported for half a circle, at $t = 12$ s. This simulation has been run with degree 6 polynomials, $Q = 8$ number of quadrature points along one dimension, 512 elements of $[0, 500]^3$ m dimension.	6
8	Density current problem. A bubble of cold air in the stratified atmosphere falls under its own weight and is transported towards the ground by convective currents. In this figure we display the volume density of the bubble. On the left, the initial condition at $t = 0$ s; on the right, the blob has dropped at $t = 73$ s. This simulation has been run with degree 9 polynomials, $Q = 13$ number of quadrature points along one dimension, 864 elements of $[0, 1000]^3$ m dimension.	7
9	Nek5000+libParanumal (NekRS): Taylor-Green vortex.	18
10	Nek5000+libParanumal (NekRS) performance baseline on Summit.	18
11	NekCEM performance baseline on Summit.	19
12	NekRS vs. Nek5000 performance baseline.	20

LIST OF TABLES

1	Experimental code NekRS run for 1000 timesteps with 8 iterations fixed for velocity and pressure. The ratio of data size 2.666 (8000/3000). Performance is shown in Figure 10.	17
2	NekCEM run for 1000 timesteps with 5-stage 4th-order Runge-Kutta time stepping. The ratio of data size 2.777 (5000/1800). Performance is shown in Figure 11.	19
3	Nek5000 OpenACC version run for 100 timesteps of eddy simulations, without fixing iterations. The ratio of data size 3.222 (1600/500). Performance is shown in Figure 12.	20
4	Time and # iterations required to reach L_2 errors at the level of 10^{-7} , using 8 KNL nodes (512 MPI ranks) on Argonne's LCRC/Bebop.	20

noise in the results. The first six plots show the results of Benchmark Problem 1 using reference backends in pure C and optimized backends using AVX instructions and LIBXSMM. The plots on the left show the backends using internal vectorization for each element and the plots on the right show the backends using external vectorization across blocks of eight elements.

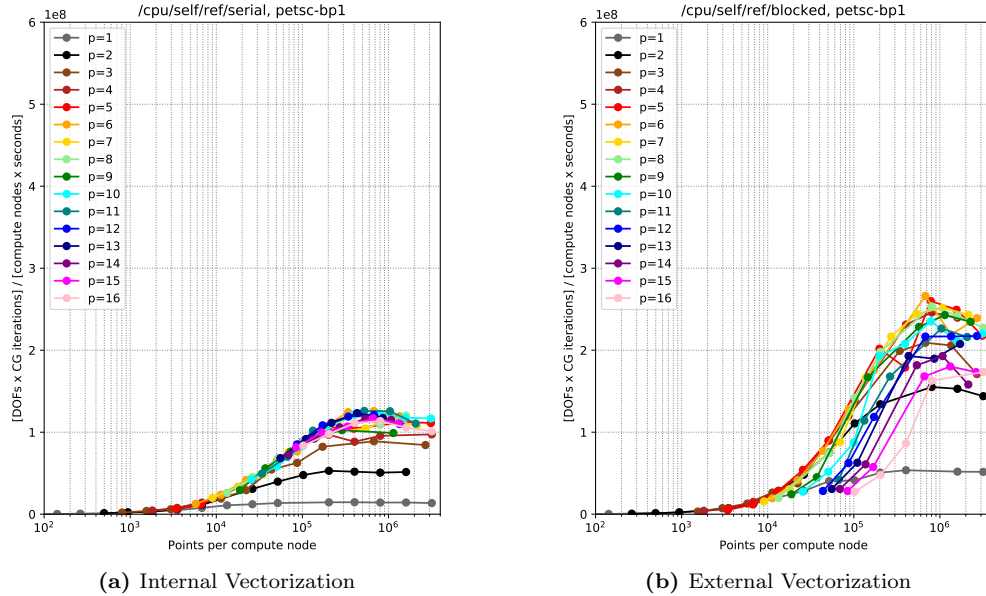


Figure 1: Pure C Reference Backends

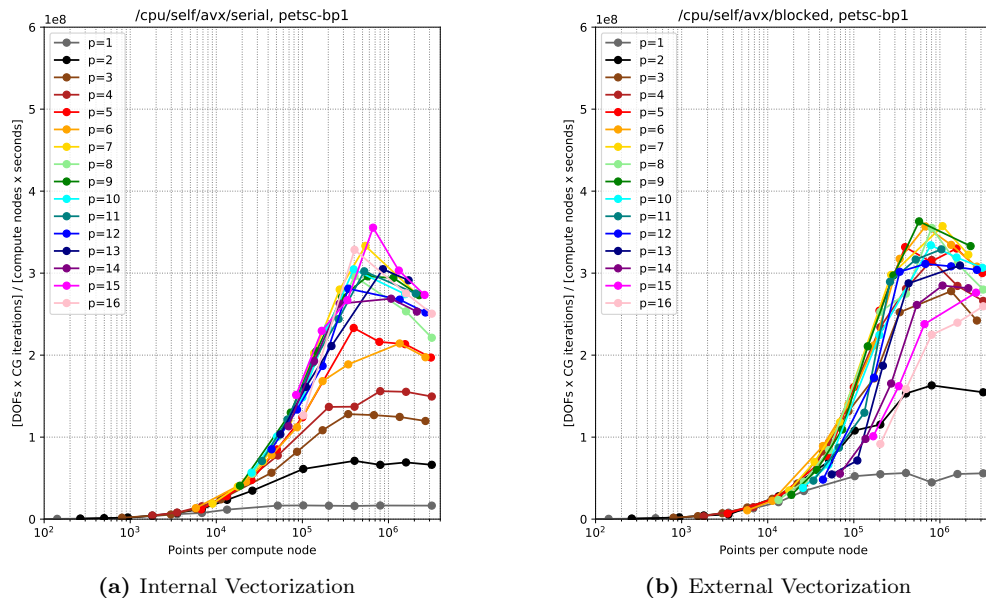


Figure 2: AVX Backends

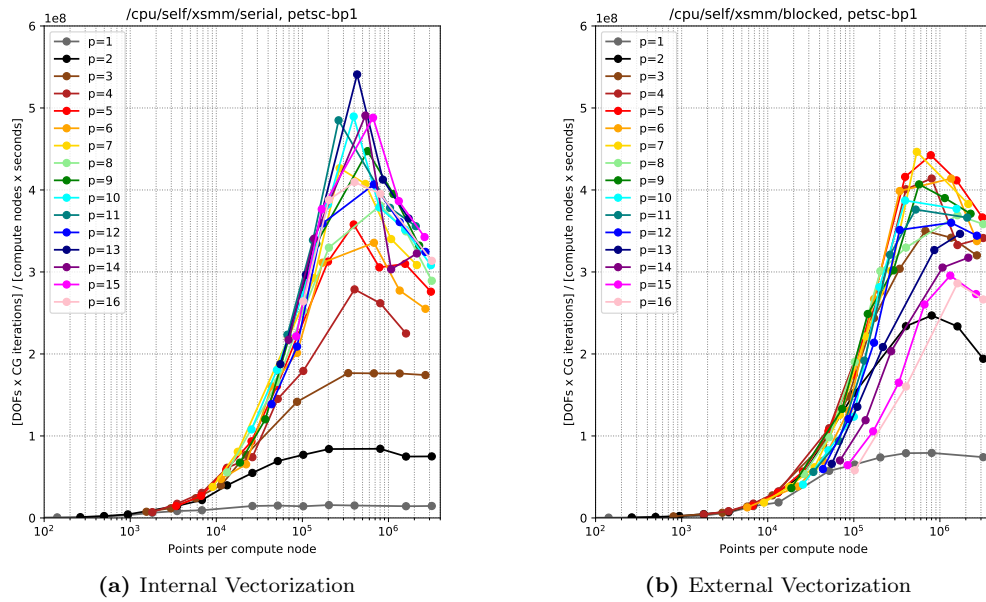


Figure 3: LIBXSMM Backends

The next six plots show the results of Benchmark Problem 3 for the same six backends described above.

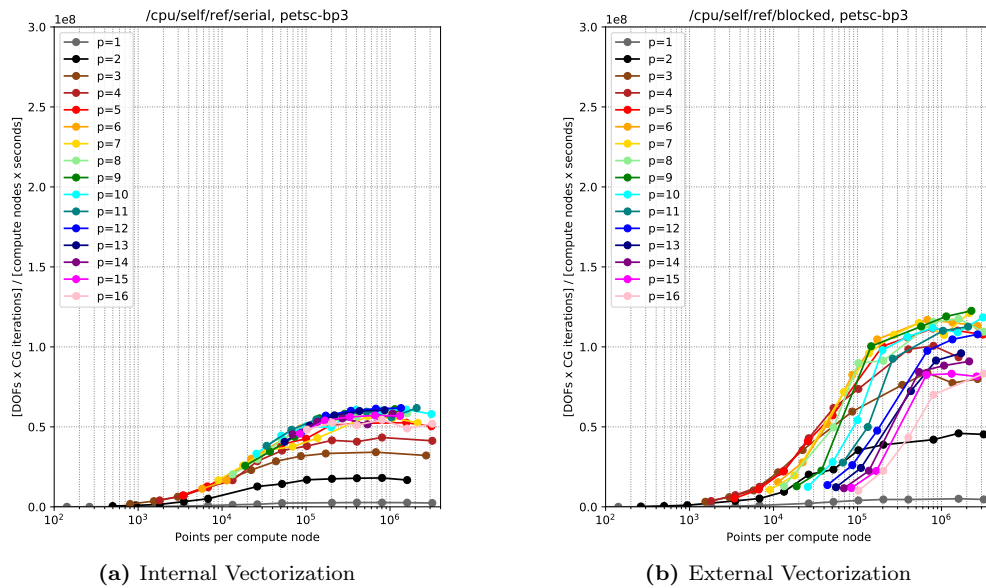


Figure 4: Pure C Reference Backends

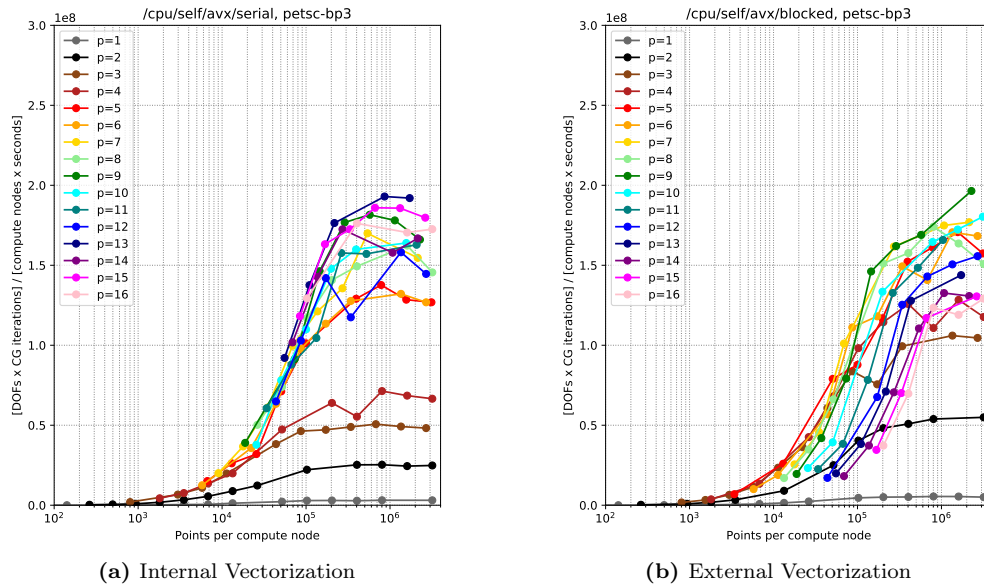


Figure 5: AVX Backends

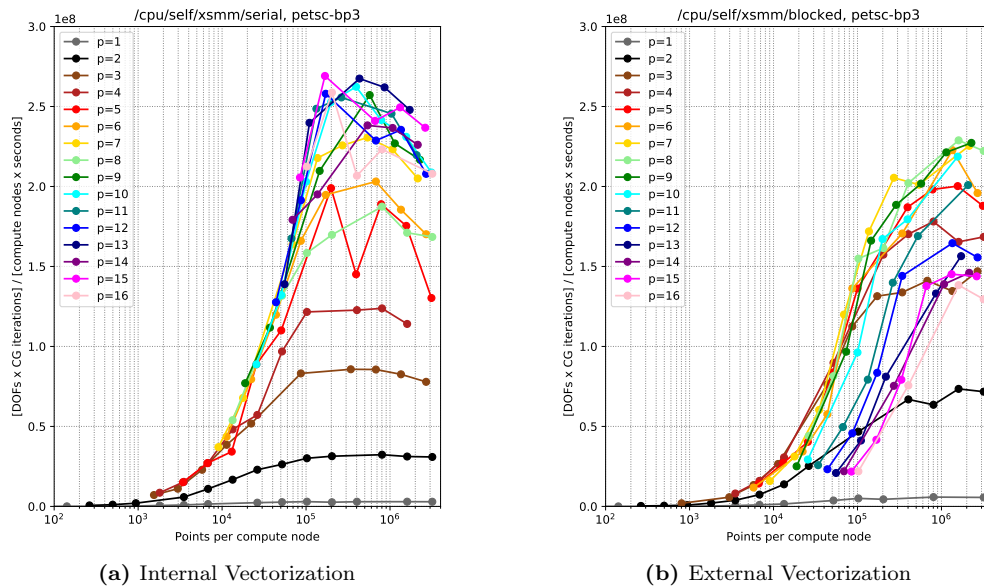


Figure 6: LIBXSMM Backends

2.1.4 Navier-Stokes example

A new explicit Navier-Stokes example was added to the family of libCEED examples. This example represents work being done towards a miniapp for a Navier-Stokes solver, that fully exploits libCEED capabilities. The goal of the libCEED team was to provide users with an example that is more involved in terms of the physics solved for (therefore, also in terms of the corresponding mathematical description and QFunction implementation), compared to the existing CEED Bake-off Problems (BPs) (see <https://ceed.exascaleproject.org/bps/>), so that more complex problems of interest to the scientific community

can be efficiently solved by using libCEED.

This example solves the time-dependent Navier-Stokes equations of compressible gas dynamics in a static Eulerian three-dimensional frame using structured high-order finite element/spectral element spatial discretizations and explicit high-order time-stepping. The Navier-Stokes example has been developed using PETSc, so that the pointwise physics (defined at quadrature points) is separated from the parallelization and meshing concerns.

The mathematical formulation is given in what follows. The compressible Navier-Stokes equations in conservative form are

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{U} = 0, \quad (1a)$$

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \left(\frac{\mathbf{U} \otimes \mathbf{U}}{\rho} + P \mathbf{I}_3 \right) + \rho g \hat{\mathbf{k}} = \nabla \cdot \boldsymbol{\sigma}, \quad (1b)$$

$$\frac{\partial E}{\partial t} + \nabla \cdot \left(\frac{(E + P)\mathbf{U}}{\rho} \right) = \nabla \cdot (\mathbf{u} \cdot \boldsymbol{\sigma} + k \nabla T), \quad (1c)$$

where $\boldsymbol{\sigma} = \mu(\nabla \mathbf{u} + (\nabla \mathbf{u})^T) + \lambda(\nabla \cdot \mathbf{u})\mathbf{I}_3$ is the Cauchy (symmetric) stress tensor, with μ the dynamic viscosity coefficient, and $\lambda = -2/3$ the Stokes hypothesis constant. In equations (1), ρ represents the volume mass density, \mathbf{U} the momentum density (defined as $\mathbf{U} = \rho \mathbf{u}$, where \mathbf{u} is the vector velocity field), E the total energy density (defined as $E = \rho e$, where e is the total energy), \mathbf{I}_3 represents the 3×3 identity matrix, g the gravitational acceleration constant, $\hat{\mathbf{k}}$ the unit vector in the z direction, k the thermal conductivity constant, T represents the temperature, and P the pressure, given by the following equation of state

$$P = (c_p/c_v - 1)(E - \mathbf{U} \cdot \mathbf{U}/(2\rho) - \rho g z), \quad (2)$$

where c_p is the specific heat at constant pressure and c_v is the specific heat at constant volume (that define $\gamma = c_p/c_v$, the specific heat ratio).

The system (1) can be rewritten in vector form

$$\frac{\partial \mathbf{q}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{q}) = S(\mathbf{q}), \quad (3)$$

for the state variables 5-dimensional vector

$$\mathbf{q} = \begin{pmatrix} \rho \\ \mathbf{U} \equiv \rho \mathbf{u} \\ E \equiv \rho e \end{pmatrix} \begin{array}{l} \leftarrow \text{volume mass density} \\ \leftarrow \text{momentum density} \\ \leftarrow \text{energy density} \end{array} \quad (4)$$

where the flux and the source terms, respectively, are given by

$$\mathbf{F}(\mathbf{q}) = \begin{pmatrix} \mathbf{U} \\ (\mathbf{U} \otimes \mathbf{U})/\rho + P \mathbf{I}_3 - \boldsymbol{\sigma} \\ (E + P)\mathbf{U}/\rho - (\mathbf{u} \cdot \boldsymbol{\sigma} + k \nabla T) \end{pmatrix}, \quad (5a)$$

$$S(\mathbf{q}) = - \begin{pmatrix} 0 \\ \rho g \hat{\mathbf{k}} \\ 0 \end{pmatrix}. \quad (5b)$$

For the spatial discretization, we use high-order finite elements/spectral elements, namely, the high-order Lagrange polynomials defined over non-uniformly spaced nodes, the Legendre-Gauss-Lobatto (LGL) points (roots of the p^{th} -order Legendre polynomial P_p). We discretize the domain, $\Omega \subset \mathbb{R}^3$, by letting $\Omega = \bigcup_{e=1}^{N_e} \Omega_e$, with N_e disjoint hexahedral elements. The physical coordinates are $\mathbf{x} = (x, y, z) \in \Omega_e$, while the reference coordinates are $\boldsymbol{\xi} = (\xi, \eta, \zeta) \in \mathbf{I} = [-1, 1]^3$. Let the discrete solution be

$$\mathbf{q}_N(\mathbf{x}, t)^{(e)} = \sum_{k=1}^P \psi_k(\mathbf{x}) \mathbf{q}_k^{(e)} \quad (6)$$




Figure 7: Advection problem. A blob of energy is advected by a uniform circular flow, with constant speed $|\mathbf{u}| = 0.5$ m/s. On the left, the initial condition at $t = 0$ s; on the right, the blob has been transported for half a circle, at $t = 12$ s. This simulation has been run with degree 6 polynomials, $Q = 8$ number of quadrature points along one dimension, 512 elements of $[0, 500]^3$ m dimension.

with P the number of nodes in the element (e). We use tensor-product bases $\psi_{kji} = h_i(\xi)h_j(\eta)h_k(\zeta)$.

For the time discretization, we use an explicit formulation

$$\frac{\mathbf{q}_N^{n+1} - \mathbf{q}_N^n}{\Delta t} = -[\nabla \cdot \mathbf{F}(\mathbf{q}_N)]^n + [S(\mathbf{q}_N)]^n, \quad (7)$$

solved with the adaptive Runge-Kutta-Fehlberg (RKF4-5) method by default (any explicit time-stepping scheme available in PETSc can be chosen at runtime).

The Navier-Stokes solver can be built by moving from the CEED main directory

```
cd libCEED/examples/navier-stokes
```

and by running

```
make
```

For testing and validation purposes the libCEED team has identified a set of problems for the Navier-Stokes solver (see <https://github.com/CEED/libCEED/tree/master/examples/navier-stokes>).

The Advection problem. A simplified version of system (1), only accounting for the transport of total energy, is given by

$$\frac{\partial E}{\partial t} + \nabla \cdot (\mathbf{u}E) = 0, \quad (8)$$

with \mathbf{u} the vector velocity field. In this particular test case, a blob of total energy (defined by a characteristic radius r_c) is transported by a uniform circular velocity field. We have solved (8) with no-slip and non-penetration boundary conditions for \mathbf{u} , and no-flux for E . This problem can be run with

```
./navierstokes -problem advection
```

The size of the domain can be specified by the command line options `lx`, `ly`, `lz`, the element resolution with `resx`, `resy`, `resz`, and the characteristic radius of the energy blob by `rc`. For instance, a domain $\Omega = [0, 3000]^3$ m can be discretized with elements of 1000^3 m size, with

```
./navierstokes -problem advection -lx 3000 -ly 3000 -lz 3000 -resx 1000 -resy 1000 -resz 1000 -rc 800
```

See Figure 7 for an example of results obtained for this problem type.

The Density Current problem. For this test problem, a cold air bubble (of radius r_c) drops by convection in a neutrally stratified atmosphere. Its initial condition is defined in terms of the Exner pressure, $\pi(\mathbf{x}, t)$,




Figure 8: Density current problem. A bubble of cold air in the stratified atmosphere falls under its own weight and is transported towards the ground by convective currents. In this figure we display the volume density of the bubble. On the left, the initial condition at $t = 0$ s; on the right, the blob has dropped at $t = 73$ s. This simulation has been run with degree 9 polynomials, $Q = 13$ number of quadrature points along one dimension, 864 elements of $[0, 1000]^3$ m dimension.

and potential temperature, $\theta(\mathbf{x}, t)$, that relate to the state variables via

$$\rho = \frac{P_0}{(c_p - c_v)\theta(\mathbf{x}, t)} \pi(\mathbf{x}, t)^{\frac{c_v}{c_p - c_v}}, \quad (9a)$$

$$e = c_v\theta(\mathbf{x}, t)\pi(\mathbf{x}, t) + \mathbf{u} \cdot \mathbf{u}/2 + gz, \quad (9b)$$

where P_0 is the atmospheric pressure. For this problem, we have used no-slip and non-penetration boundary conditions for \mathbf{u} , and no-flux for mass and energy densities. This problem can be run with

```
./navierstokes -problem density_current
```

Similarly to the advection example, the user can specify the size of the physical domain, the element resolution, and the characteristic radius of the thermal bubble, together with all other physical parameters involved, at runtime. See Figure 8 for an example of results obtained for this problem type.

For both problems, the user can specify how often to print outputs (in terms of number of steps), by adding the command line option `-output_freq`. For instance, to indicate to print outputs every 1000 steps

```
./navierstokes -output_freq 1000
```

By doing so, outputs for the solution in `vts` format will be generated. Moreover, the libCEED team has developed checkpoints by printing the last calculated solution and the time stamp at which this solution was printed in binary format, so that the user can subsequently continue an existing simulation via the command line option `-continue`, followed by the step number of the solution checkpoint that is desired to be continued. For instance, to continue a simulation from the output number 2000

```
./navierstokes -continue 2000
```

For a list of all available options see <https://github.com/CEED/libCEED/tree/master/examples/navier-stokes/README.md>.

2.2 MFEM 3.4

The CEED 2.0 distribution includes version 3.4 of MFEM, which adds the following new features compared to CEED 1.0:

- More general and efficient mesh adaptivity

- Added support for PUMI, the Parallel Unstructured Mesh Infrastructure from <https://scorec.rpi.edu/pumi>. PUMI is an unstructured, distributed mesh data management system that is capable of handling general non-manifold models and effectively supports automated adaptive analysis. PUMI enables for the first time support for parallel unstructured modifications of MFEM meshes.
 - Significantly reduced MPI communication in the construction of the parallel prolongation matrix in `ParFiniteElementSpace`, for much improved parallel scaling of non-conforming AMR on hundreds of thousands of MPI tasks. The memory footprint of the `ParNCMesh` class has also been reduced.
 - In `FiniteElementSpace`, the fully assembled refinement matrix is now replaced by default by a specialized refinement operator. The operator option is both faster and more memory efficient than using the fully assembled matrix. The old approach is still available and can be enabled, if needed, using the new method `FiniteElementSpace::SetUpdateOperatorType()`.
- Discretization improvements
 - Added support for a general *high-order-to-low-order refined* transfer of `GridFunction` and true-dof data from a *high-order* finite element space defined on a coarse mesh, to a *low-order refined* space defined on a refined mesh. The new methods, `GetTransferOperator` and `GetTrueTransferOperator` in the `FiniteElementSpace` classes, work in both serial and parallel and support matrix-based as well as matrix-free transfer operator representations. They use a new method, `GetTransferMatrix`, in the `FiniteElement` class similar to `GetLocalInterpolation`, that allows the coarse `FiniteElement` to be different from the fine `FiniteElement`.
 - Added class `ComplexOperator`, that implements the action of a complex operator through the equivalent 2×2 real formulation. Both symmetric and antisymmetric block structures are supported.
 - Added classes for general block nonlinear finite element operators (deriving from `BlockNonlinearForm` and `ParBlockNonlinearForm`) enabling solution of nonlinear systems with multiple unknowns in different function spaces. Such operators have assemble-based action and also support assembly of the gradient operator to enable inversion with Newton iteration.
 - Added variable order NURBS: for each space each knot vector in the mesh can have a different order. The order information is now part of the finite element space header in the NURBS mesh output, so NURBS meshes in the old format need to be updated.
 - In the classes `NonlinearForm` and `ParNonlinearForm`, added support for non-conforming AMR meshes.
 - New specialized time integrators: symplectic integrators of orders 1-4 for systems of first order ODEs derived from a Hamiltonian and generalized-alpha ODE solver for the filtered NavierStokes equations with stabilization. See classes `SIASolver` and `GeneralizedAlphaSolver` in `linalg/ode.hpp`.
 - Inherit finite element classes from the new base class `TensorBasisElement`, whenever the basis can be represented by a tensor product of 1D bases.
 - Added support for elimination of boundary conditions in block matrices.
 - New and updated examples and miniapps
 - Added a new serial and parallel example (`ex19`) that solves the quasi-static incompressible hyperelastic equations. The example demonstrates the use of block nonlinear forms as well as custom block preconditioners.
 - Added a new electromagnetics miniapp, Maxwell, for simulating time-domain electromagnetics phenomena as a coupled first order system of equations.
 - A simple local refinement option has been added to the mesh-explorer miniapp (menu option `'r'`, sub-option `'l'`) that selects elements for refinement based on their spatial location - see the function `region()` in the source file.

- Added a set of miniapps specifically focused on Isogeometric Analysis (IGA) on NURBS meshes in the `miniapps/nurbs` directory. Currently the directory contains variable order NURBS versions of examples 1, 1p and 11p.
 - Added PUMI versions of examples `ex1`, `ex1p`, `ex2` and `ex6p` in a new `examples/pumi` directory. The new examples demonstrate the PUMI APIs for parallel and serial mesh loading (`ex1` and `ex1p`), applying BCs using classification (`ex2`), and performing parallel mesh adaptation (`ex6p`).
 - Added two new miniapps related to `DataCollection` I/O in `miniapps/tools`: `load-dc.cpp` can be used to visualize fields saved via `DataCollection` classes; `convert-dc.cpp` demonstrates how to convert between MFEM's different concrete `DataCollection` options.
 - Example 10p with its SUNDIALS and PETSc versions have been updated to reflect the change in the behavior of the method `ParNonlinearForm::GetLocalGradient()` and now works correctly on non-conforming AMR meshes. Example 10 and its SUNDIALS version have also been updated to support non-conforming AMR meshes.
- Miscellaneous
 - Documented project workflow and provided contribution guidelines in the new top-level file, `CONTRIBUTING.md`.
 - Added (optional) Conduit Mesh Blueprint support of MFEM data for both in-core and I/O use cases. This includes a new `ConduitDataCollection` that provides json, simple binary, and HDF5-based I/O. Support requires Conduit \geq v0.3.1 and VisIt \geq v2.13.1 will read the new Data Collection outputs.
 - Added a new developer tool, `config/sample-runs.sh`, that extracts the sample runs from all examples and miniapps and runs them. Optionally, it can save the output from the execution to files, allowing comparison between different versions and builds of the library.
 - Support for building a shared version of the MFEM library with GNU make.
 - Added a build option, `MFEM_USE_EXCEPTIONS=YES`, to throw an exception instead of calling `abort` on MFEM errors.
 - When building with the GnuTLS library, switch to using X.509 certificates for secure socket authentication. Support for the previously used OpenPGP keys has been deprecated in GnuTLS 3.5.x and removed in 3.6.0. For secure communication with the visualization tool GLVis, a new set of certificates can be generated using the latest version of the script `glvis-keygen.sh` from GLVis.
 - Upgraded MFEM to support Axom 0.2.8. Prior versions are no longer supported.

2.3 Nek5000 version 19.0-rc2

The CEED 2.0 distribution includes version 19.0-rc2 of Nek5000 adds the following new features compared to CEED 1.0:

- Enhanced scalability
 - Uncoupled multi-session (`neknek`) simulations
 - Gather scatter operations across sessions
 - Gather scatter options across gtp-planes
 - `par` file support for postnek
 - `mkSIZE` to automatically create SIZE file
 - Object and boundary handling
 - Support `ParMETIS` partitioner
- Enhanced algorithmic supports
 - RANS $k-\omega$ and $k-\omega$ -SST (experimental)

- Online mesh-smoother (experimental)
- **FEM_AMG** preconditioner (experimental) **p40=3**.
- **SEM_AMG_HYPRE** preconditioner (experimental) **p40=2**.
- Lagrangian phase model - LPM (experimental)
- Miscellaneous
 - ElapsedTime option for writeControl (in par)
 - Print runtime-statistics every 100 steps
 - Support for GNU 8.x compilers
 - Support for Cray compilers
 - Support for ARM compilers
 - **CHT** support for generic fld reader
 - Overwrite core routines in usr

2.4 PETSc 3.11

PETSc 3.11 includes contributions from nearly 90 contributors since v3.8.3 (included in CEED-1.0). Some notable features for CEED include

- **VecScatter** support for use of shared memory within a node and for using **PetscSF** (heavily used for unstructured mesh operations),
- AVX-512 optimizations and new matrix formats (sliced ELLPACK and MKL inspector-executor),
- higher performance scalable sparse matrix-matrix products and $P^T AP$ operations (used by algebraic multigrid setup),
- improved parallel mesh IO and hierarchical operations,
- pipelined Krylov methods with improved stability properties, and
- Python-3.4+ support (in addition to Python-2.6+).

2.5 OCCA 1.0.8

OCCA has released versions 1.0.0 through 1.0.8 thanks to 6 contributors, which include the following features and updates

- OKL parser improvements
- Enabled kernel compilation using AMD's HIP
- Added kernel type safety by giving **occa::memory** (C++) and **occaMemory** (C) optional type hints
- Test coverage improvement (57.9% → 70.9%)
- Enabled OKL **@attributes** through **#pragma occa attributes** which can be used in regular C / C++ code
- Released initial Python API which supports building OKL kernels through Python's native function syntax

2.6 PUMI 2.2.0

The principal feature of PUMI 2.2.0 is inclusion in the December 2018 XSDK-0.4.0 release. Several additional developments are listed below.

- Travis CI nightly testing (3450a945)
- User documentation of mesh and geometric model support (<https://github.com/SCOREC/core/wiki/Geometric-Models>, <https://github.com/SCOREC/core/wiki/Mesh-Generation>)
- Field cloning (6f3c3926)
- Generalizations of field accumulation (d783b19e)
- Compatibility with Simmetrix SimModSuite 12.0.180811 (06df5a0f)
- Support for discontinuous Galerkin methods with Simmetrix meshes (1727b32b)

2.7 MAGMA 2.5.0

The CEED 2.0 distribution includes developments from the MAGMA 2.4 release (on 06/25/18) and the latest MAGMA 2.5 release (on 01/02/2019). This includes routines for batched computations and support for the fast FP16 Tensor Cores arithmetic in Nvidia V100 GPUs, namely:

- Batched device interface routines for use in tensor contractions;
- Performance improvements across many batch routines, including batched TRSM, batched LU, batched LU-nopiv, and batched Cholesky;
- Batched GEMM with support for strided access of the matrices;
- GEMM in FP16 arithmetic (HGEMM) as well as auxiliary functions to cast matrices from FP32 to FP16 storage (`magmablas_slag2h`) and from FP16 to FP32 (`magmablas_hlag2s`);
- Mixed-precision solver that is able to provide an FP64 solution with up to 4X speedup using the fast FP16 Tensor Cores arithmetic in the Nvidia V100 GPUs.

2.8 Laghos 2.0

The CEED 2.0 distribution includes version 2.0 of Laghos, which introduces four new variants of the miniapp, namely,

- pure CUDA version, where the major computation kernels are rewritten in CUDA without utilizing any intermediate abstraction layers;
- OCCA version, based on the OCCA abstraction layer;
- RAJA version, based on the RAJA abstraction layer;
- AMR version, demonstrating dynamic adaptive mesh refinement for moving curved meshes of arbitrary order.

2.9 HPGMG 0.4

HPGMG 0.4 includes several improvements to the FV and FE variants, including

- AVX-512 support for FE,
- CEED BP solver configurations and extra quadrature points for FE,
- removal of FE process grid restriction for larger prime factors, and
- layout flexibility and padding/alignment support for FV.

2.10 gslib v.1.0.4

gslib v1.0.4 includes major features and improvements:

- Added support for `gs.float` in Fortran `gs`.
- Updated testing unit with new header files.

2.11 NekCEM

Newly supported features of NekCEM include the following:

- Multi-session run support for large-scale parameter studies that represent structural differences between optical devices. This has been tested up to millions of parameter studies for phase analysis of metals.
- Structural support for OLCF/Summit runs; Measure performance on Summit, in comparison to Titan. (See Section 4.2)

3. CEED 2.0 TESTING AND DISTRIBUTION

The CEED 2.0 distribution consists of 12 Spack packages (plus a CEED meta-package), documented at <https://ceed.exascaleproject.org/ceed-2.0/>. In this section we review the configuration settings for several common and leadership computing architectures.

3.1 Mac

The configuration file below, also included as `ceed2-darwin-highsierra-x86.64-packages.yaml` in the CEED 2.0 distribution, provides a sample `packages.yaml` file based on Homebrew, that should work on most Macs. (You can use MacPorts instead of Homebrew if you prefer.)

```

1 packages:
2   all:
3     compiler: [clang]
4     providers:
5       blas: [veclibfort]
6       lapack: [veclibfort]
7       mpi: [openmpi]
8   openmpi:
9     paths:
10      openmpi@3.0.0: ~/brew
11     buildable: False
12
13   cmake:
14     paths:
15      cmake@3.10.2: ~/brew
16     buildable: False
17   cuda:
18     paths:
19      cuda@9.1.85: /usr/local/cuda
20     buildable: False
21   libx11:
22     paths:
23      libx11@system: /opt/X11
24     version: [system]
25     buildable: False
26   libxt:
27     paths:
28      libxt@system: /opt/X11
29     version: [system]
30     buildable: False
31   xproto:
32     paths:
33      # see /opt/X11/lib/pkgconfig/xproto.pc
34      xproto@7.0.31: /opt/X11

```

```

35     version: [7.0.31]
36     buildable: False
37   python:
38     paths:
39       python@2.7.10: /usr
40     buildable: False
41   zlib:
42     paths:
43       zlib@1.2.11: /usr
44     buildable: False

```

The packages in `/brew` were installed with `brew install package`. If you don't have Homebrew, you can install it and the needed tools with:

```

git clone https://github.com/Homebrew/brew.git
cd brew
bin/brew install openmpi cmake python zlib

```

The packages in `/usr` are provided by Apple and come pre-built with Mac OS X. The `cuda` package is provided by NVIDIA and should be installed separately by downloading it from NVIDIA. We are using the Clang compiler, Open MPI, and Apple's BLAS/LAPACK accelerator library.

3.2 Linux RHEL7

The configuration file below, also included as `ceed2-linux-rhel7-x86_64-packages.yaml` in the CEED 2.0 distribution, provides a sample 'packages.yaml' file that can be adapted to work on a RHEL7 Linux desktop

```

1 packages:
2   all:
3     compiler: [gcc]
4     providers:
5       mpi: [openmpi]
6       blas: [netlib-lapack]
7       lapack: [netlib-lapack]
8   netlib-lapack:
9     paths:
10      netlib-lapack@system: /usr/lib64
11     buildable: False
12   openmpi:
13     paths:
14      openmpi@3.0.0: ~/local
15     buildable: False
16
17   cmake:
18     paths:
19      cmake@3.10.2: ~/local
20     buildable: False
21   cuda:
22     paths:
23      cuda@9.1.85: ~/local/cuda
24     buildable: False
25   libx11:
26     paths:
27      libx11@system: /usr
28     version: [system]
29     buildable: False
30   libxt:
31     paths:
32      libxt@system: /usr
33     version: [system]
34     buildable: False
35   xproto:
36     paths:
37      xproto@7.0.32: /usr
38     version: [7.0.32]
39     buildable: False
40   python:

```

```

41     paths:
42         python@2.7.14: /usr
43     buildable: False
44     zlib:
45         paths:
46             zlib@1.2.11: /usr/lib64
47     buildable: False

```

The above file uses user-installed Open MPI, CMake and CUDA packages, with the rest of the CEED prerequisites installed via the *yum* package manager.

3.3 Linux Ubuntu

A very similar file, [ceed2-ubuntu18.10-packages.yaml](#) provides Spack configuration for the Ubuntu distribution:

```

1 packages:
2   all:
3     compiler: [gcc]
4     providers:
5       mpi: [mpich]
6       blas: [openblas]
7       lapack: [openblas]
8   openblas:
9     paths:
10        openblas@system: /usr/lib
11    buildable: False
12   mpich:
13     paths:
14        mpich@3.3: /usr/local
15    buildable: False
16
17   cmake:
18     paths:
19        cmake@3.12.1: /usr
20    buildable: False
21   libx11:
22     paths:
23        libx11@system: /usr
24    version: [system]
25    buildable: False
26   libxt:
27     paths:
28        libxt@system: /usr
29    version: [system]
30    buildable: False
31   xproto:
32     paths: # See /usr/share/pkgconfig/xproto.pc for version
33        xproto@7.0.32: /usr
34    buildable: False
35   python:
36     paths:
37        python@3.6.7: /usr
38    buildable: False
39   zlib:
40     paths:
41        zlib@1.2.11: /usr/lib
42    buildable: False

```

In this case we use GCC and other development packages via `apt install` and with MPICH installed separately (as needed to use containerized HPC environments like Shifter and Singularity). You can use `docker pull jedbrown/ceed-base` to get a build environment that is ready for `spack install ceed`.

3.4 Cori

The CEED2.0 installation on Cori (NERSC) via the Spack package manager can be obtained by downloading the Spack configuration file for Cori, that can be found on the CEED website <https://ceed.exascaleproject.org/ceed-2.0/>.

After downloading the file, it needs to be renamed from `ceed2-cori-packages.yaml` to `packages.yaml` and placed in the folder `/.spack/` or `/.spack/<platform>/` on Cori. This file contains information on the packages that the CEED2.0 installation requires, so that the user does not need to install them from scratch, which can take a long time. The Spack configuration file for Cori is the following

```

1 packages:
2   all:
3     compiler: [gcc@7.3.0, intel@18.0.5.274]
4     providers:
5       mpi: [mpich]
6       mkl: [intel-mkl]
7       blas: [intel-mkl, cray-libsci]
8       scalapack: [intel-mkl, cray-libsci]
9       pkgconfig: [pkg-config]
10    mpich:
11      modules:
12        mpich@3.2%gcc@7.3.0 arch=cray-cn19-haswell: cray-mpich
13        mpich@3.2%intel@18.0.5.274 arch=cray-cn19-haswell: cray-mpich
14      buildable: False
15    intel-mkl:
16      buildable: false
17      paths:
18        intel-mkl@2018.3.222%intel: /opt/intel
19        intel-mkl@2018.3.222%gcc: /opt/intel
20    pkg-config:
21      buildable: false
22      paths:
23        pkg-config@0.28: /usr
24    cmake:
25      modules:
26        cmake@3.14.0%gcc@7.3.0 arch=cray-cn19-haswell: cmake
27        cmake@3.14.0%intel@18.0.5.274 arch=cray-cn19-haswell: cmake
28      buildable: False
29    libx11:
30      paths:
31        libx11@system: /usr
32      version: [system]
33      buildable: False
34    libxt:
35      paths:
36        libxt@system: /usr
37      version: [system]
38      buildable: False
39    xproto:
40      paths: # See /usr/lib64/pkgconfig/xproto.pc for version
41        xproto@7.0.28: /usr
42      buildable: False
43    python:
44      paths:
45        python@2.7.13: /usr
46      buildable: False
47    boost:
48      modules:
49        boost@1.69.0%gcc@7.3.0 arch=cray-cn19-haswell: boost
50        boost@1.69.0%intel@18.0.5.274 arch=cray-cn19-haswell: boost
51      buildable: False
52    m4:
53      modules:
54        m4@1.4.17%gcc@7.3.0 arch=cray-cn19-haswell: m4
55        m4@1.4.17%intel@18.0.5.274 arch=cray-cn19-haswell: m4
56      buildable: False
57    openssl:
58      modules:
59        openssl@1.1.0a%gcc@7.3.0 arch=cray-cn19-haswell: openssl
60        openssl@1.1.0a%intel@18.0.5.274 arch=cray-cn19-haswell: openssl
61      buildable: False
62    perl:
63      paths:

```



```

64     perl@5.18.2%gcc@7.3.0 arch=cray-cn19-haswell: /usr
65     perl@5.18.2%intel@18.0.5.274 arch=cray-cn19-haswell: /usr
66     buildable: False
67     autoconf:
68     modules:
69         autoconf@2.69%gcc@7.3.0 arch=cray-cn19-haswell: autoconf
70         autoconf@2.69%intel@18.0.5.274 arch=cray-cn19-haswell: autoconf
71     buildable: False
72     automake:
73     modules:
74         automake@1.15%gcc@7.3.0 arch=cray-cn19-haswell: automake
75         automake@1.15%intel@18.0.5.274 arch=cray-cn19-haswell: automake
76     buildable: False

```

3.5 ALCF Theta

The configuration file below, also included as `ceed2-theta-packages.yaml` in the CEED 2.0 distribution, provides a sample ‘packages.yaml’ for the Theta system at ALCF.

```

1 packages:
2   cmake:
3     paths:
4       cmake@3.5.2%gcc@8.2.0 arch=cray-CNL-mic_knl: /usr
5       cmake@3.5.2%intel@16.0.3.210 arch=cray-CNL-mic_knl: /usr
6     buildable: False
7   python:
8     paths:
9       python@2.7.13%gcc@8.2.0 arch=cray-CNL-mic_knl: /usr
10      python@2.7.13%intel@16.0.3.210 arch=cray-CNL-mic_knl: /usr
11    buildable: False
12  pkg-config:
13    paths:
14      pkg-config@0.28%gcc@8.2.0 arch=cray-CNL-mic_knl: /usr
15      pkg-config@0.28%intel@16.0.3.210 arch=cray-CNL-mic_knl: /usr
16    buildable: False
17  autoconf:
18    paths:
19      autoconf@2.69%gcc@8.2.0 arch=cray-CNL-mic_knl: /usr
20      autoconf@2.69%intel@16.0.3.210 arch=cray-CNL-mic_knl: /usr
21    buildable: False
22  automake:
23    paths:
24      automake@1.13.4%gcc@8.2.0 arch=cray-CNL-mic_knl: /usr
25      automake@1.13.4%intel@16.0.3.210 arch=cray-CNL-mic_knl: /usr
26    buildable: False
27  libtool:
28    paths:
29      libtool@2.4.2%gcc@8.2.0 arch=cray-CNL-mic_knl: /usr
30      libtool@2.4.2%intel@16.0.3.210 arch=cray-CNL-mic_knl: /usr
31    buildable: False
32  m4:
33    paths:
34      m4@1.4.16%gcc@8.2.0 arch=cray-CNL-mic_knl: /usr
35      m4@1.4.16%intel@16.0.3.210 arch=cray-CNL-mic_knl: /usr
36    buildable: False
37  intel-mkl:
38    paths:
39      intel-mkl@16.0.3.210%intel@16.0.3.210 arch=cray-CNL-mic_knl: /opt/intel
40    buildable: False
41  mpich:
42    modules:
43      # requires 'module load cce' otherwise gives parsing error
44      mpich@7.6.3%gcc@8.2.0 arch=cray-CNL-mic_knl: cray-mpich/7.6.3
45      mpich@7.6.3%intel@16.0.3.210 arch=cray-CNL-mic_knl: cray-mpich/7.6.3
46    buildable: False
47  boost:
48    paths:

```

```

49 boost@1.64.0%gcc@8.2.0 arch=cray-CNL-mic_knl: /soft/libraries/boost/1.64.0/gnu
50 boost@1.64.0%intel@16.0.3.210 arch=cray-CNL-mic_knl: /soft/libraries/boost/1.64.0/
intel
51 buildable: False
52 all:
53 providers:
54 mpi: [mpich]
55 compiler: [gcc@8.2.0]

```

3.6 Other machines: OLCF Summit, LLNL TOSS3 and Lassen

For Spack configurations tested on other machines, please consult the CEED 2.0 release website <https://ceed.exascaleproject.org/ceed-2.0>.

4. GPU PERFORMANCE BASELINES FOR CEED APPLICATIONS

In responding to the needs of ECP Urban and ExaSMR teams, we measured baseline performance of our current versions of Nek and libParanumal codes on Summit, in comparison to Titan.

4.1 Urban: libParanumal performance on Summit vs. Titan

A number of performance improvements were recently implemented in libParanumal’s incompressible Navier-Stokes flow solver, including improved JIT compilation of OCCA for 1000s of MPI ranks, experiments with overlapping additive Schwarz and communication avoiding, pipelined preconditioned conjugate gradient solvers. With these changes, the libParanumal solver matched the NekBone miniapp within 10% in CPU mode. In GPU mode, the libParanumal solver demonstrated weak GPU scaling at approximately 90% on 1/2 of both Summit and Titan.

An experimental version (NekRS) has been under development that is based on integration of Nek5000 + libParanumal. As a part of collaboration with ECP Urban team, we measured the speedup of NekRS on Summit (vs. Titan) for incompressible Navier-Stokes solver for Taylor-Green vortex simulations shown in Figure 9 and demonstrated the baseline performance in Figure 10.

Our testing setups for baseline performance on Summit and Titan are summarized in Table 1. We performed simulations for 1000 timesteps with 8 iterations fixed for velocity and pressure. We measured maximum size per GPU, $n = EN^3$ with $N = 7$, keeping the ratio 2.666 (8000/3000) which is equivalent to the ratio of the memory size per GPU, 16GB on Summit and 6GB on Titan. While the total #GPUs are 27648 on Summit and 18688 on Titan, we keep the same ratio with 1.479. We achieve $4.92\times$ speedup in simulation time and $7.2\times$ speedup in DOFs/sec per step as shown in Figure 10. From the ratio of the memory bandwidth (900GB/s on Summit and 250GB/s on Titan), the speedup in timing would be ~ 5.3 and the theoretical speedup would be $8.8x$ (5.3×1.5) on Summit with $1.5\times$ more resources.

Table 1: Experimental code NekRS run for 1000 timesteps with 8 iterations fixed for velocity and pressure. The ratio of data size 2.666 (8000/3000). Performance is shown in Figure 10.

system	E per GPU	N	$n = EN^3$	$DOFs = 4 \times n$ per GPU
Summit	8000	7	2.7 M	10.8 M
Titan	3000	7	1.0 M	4 M

4.2 Urban: NekCEM performance on Summit vs. Titan

As a part of collaboration with ECP Urban team, we also measured the speedup of NekCEM on Summit (vs. Titan). NekCEM’s GPU implementation is based on OpenACC and the code is fully explicit solver which is highly optimized. We use the similar experimental setup and demonstrate the baseline performance for periodic traveling wave simulations shown in Figure 11. Our testing setups for baseline performance on Summit and Titan are summarized in Table 2. We performed 1000 timesteps, involving stage 4th-order

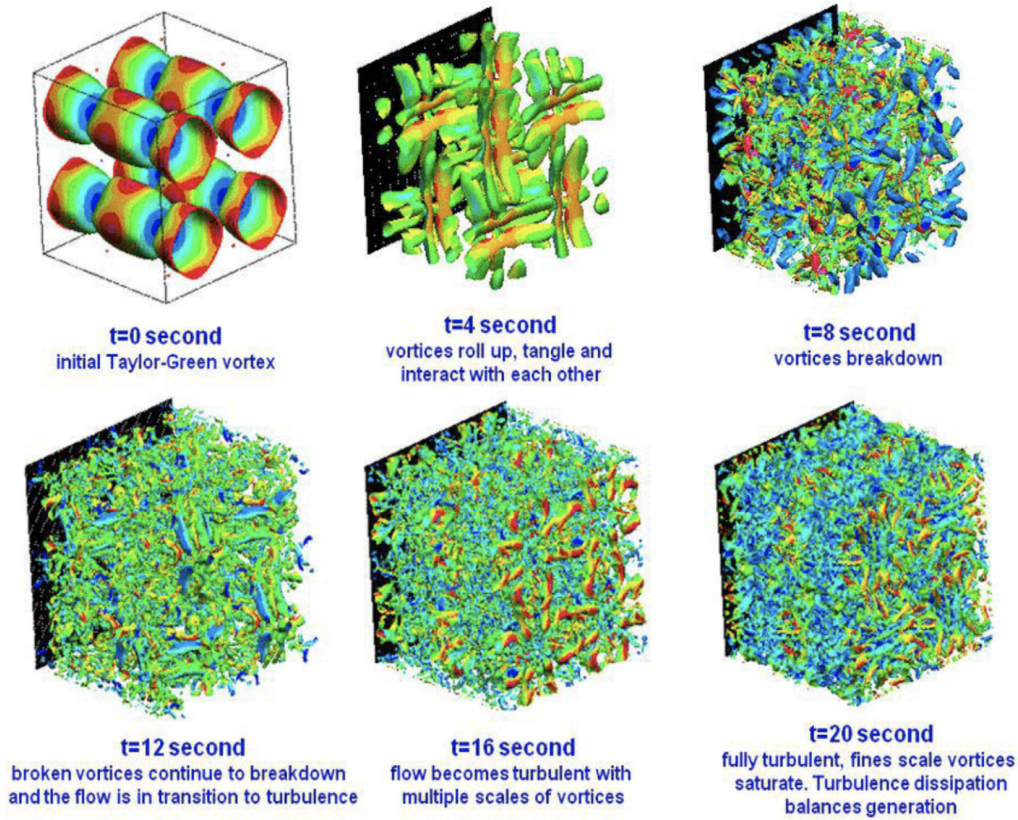


Figure 9: Nek5000+libParanumal (NekRS): Taylor-Green vortex.

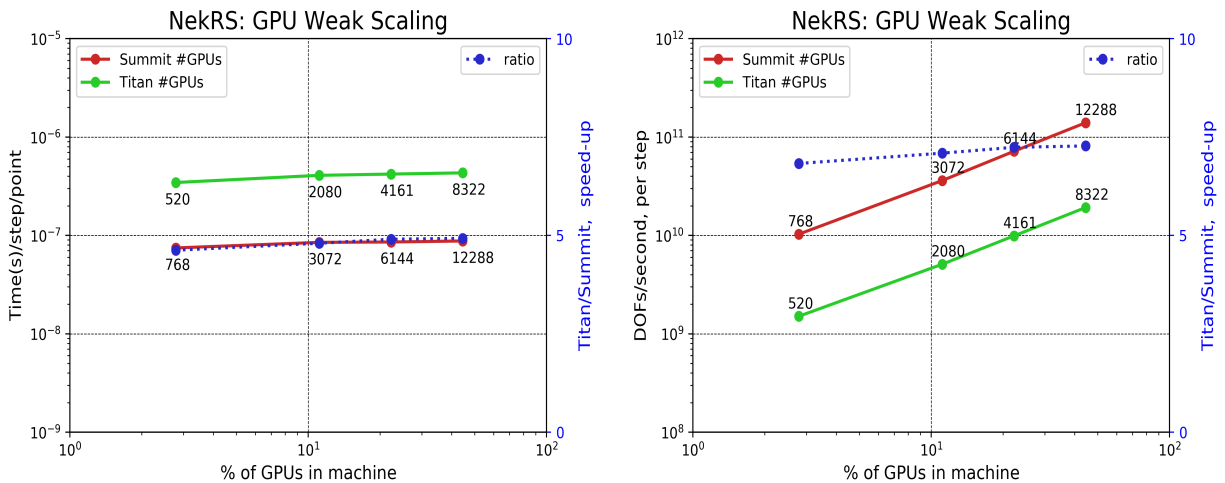


Figure 10: Nek5000+libParanumal (NekRS) performance baseline on Summit.

Runge-Kutta scheme and measured maximum size per GPU, $n = E(N + 1)^3$ with $N = 11$, keeping the ratio 2.777 (5000/1800). We had $n = 5000 \times 12^3 = 8.6$ million grid points per GPU on Summit and $n = 1800 \times 12^3 = 3.1$ million grid points per GPU on Titan with the total degree of freedoms $DOFs = 6 \times n \times (\#GPUs)$ for the 6 components from electric and magnetic fields. Titan's #GPUs is chosen 76% of Summit,

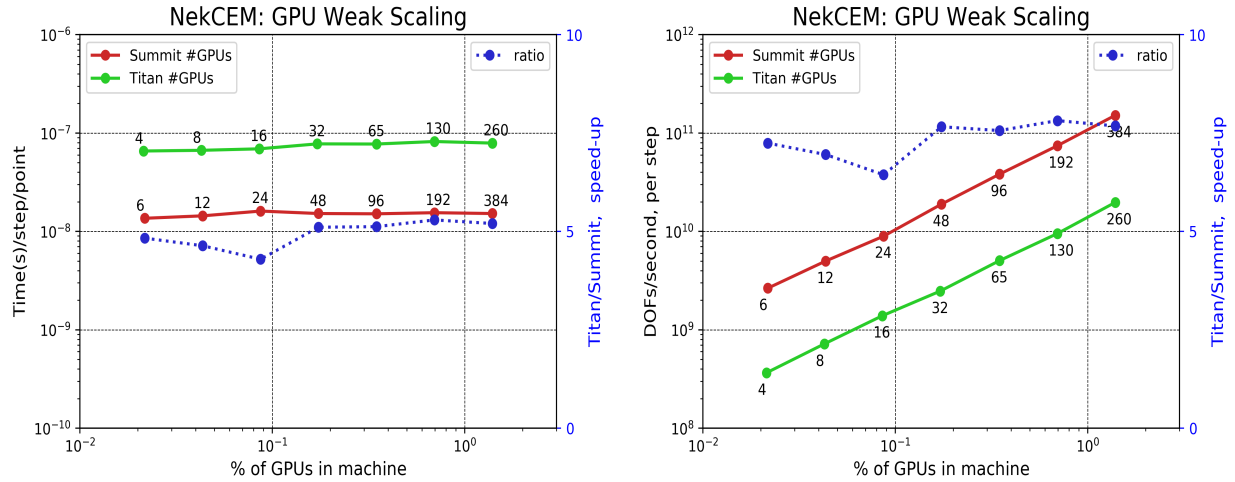


Figure 11: NekCEM performance baseline on Summit.

reflecting full machine GPU ratios. We achieve $5.2\times$ speedup in simulation time and $7.6\times$ speedup in DOFs/sec per step as shown in Figure 11.

Table 2: NekCEM run for 1000 timesteps with 5-stage 4th-order Runge-Kutta time stepping. The ratio of data size 2.777 (5000/1800). Performance is shown in Figure 11.

system	E per GPU	N	$n = E(N + 1)^3$	$DOFs = 6 \times n$ per GPU
Summit	5000	11	8.6 M	51.6 M
Titan	1800	11	3.1 M	18.1 M

4.3 ExaSMR: Speedup of NekRS on Summit vs. Nek5000 on Titan

In responding to the need of ExaSMR team, we measured performance baseline the speedup from the newly developed NekRS vs the previously developed OpenACC-based Nek5000. Table 3 shows the experimental setup for Nek5000 OpenACC version.

Figure 12 demonstrates the current baseline measurement showing $\sim 11\times$ speedup from newer version of NekRS on Summit, compared to the old Nek5000 OpenACC version on Titan.

However we used less than 10% for the Nek5000 OpenACC version. With the recent AMG support enabled in Nek5000 OpenACC version, CEED team plans to measure with larger elements counts using 50% of Titan and Summit for 17×17 rods geometry, in collaboration with ExaSMR team.

Our testing setups for baseline performance of Nek5000 on Summit and Titan are summarized in Table 3. We performed Nek5000 OpenACC version running 100 timesteps of eddy simulations without fixing iterations and measured maximum size per GPU, $n = EN^3$ with $N = 14$, keeping the ratio 3.222 (1600/500). We used $n = 1600 \times 14^3 = 3.6$ million grid points per GPU on Summit and $n = 500 \times 14^3 = 3.1$ million grid points per GPU on Titan. The total degree of freedoms is $DOFs = 4 \times n \times (\#GPUs)$ with 4 fields (3 velocity components and 1 pressure). We chose Titan's #GPUs as 76% of Summit.

4.4 ExaSMR: Speedup of Steady Thermal Solver vs. Transient Calculation

In collaboration with ExaSMR team, we performed the spectral element simulations with Nek5000 for all LES, RANS, and pseudo-RANS simulations. The recently proposed Nek5000 steady-state solver has been used for solving the temperature field in the pseudo-RANS approach and has proved significantly faster than transient schemes. Prediction of thermal quantities is compared with classical linear and nonlinear RANS

Table 3: Nek5000 OpenACC version run for 100 timesteps of eddy simulations, without fixing iterations. The ratio of data size 3.222 (1600/500). Performance is shown in Figure 12.

system	E per GPU	N	$n = EN^3$	$DOFs = 4 \times n$ per GPU
Summit	1600	14	8.6 M	34.4 M
Titan	500	14	3.1 M	12.4 M




Figure 12: NekRS vs. Nek5000 performance baseline.

models. LES for the full-length rods has also been performed and are used as a reference. Results of the proposed method are demonstrated in Table 4 with significant improvements ($\sim 7\times$) with respect to those obtained with RANS.

Table 4: Time and # iterations required to reach L_2 errors at the level of 10^{-7} , using 8 KNL nodes (512 MPI ranks) on Argonne’s LCRC/Bebop.

Case	# iter			CPU time [min]		
	$tol = 10^{-2}$	$tol = 10^{-4}$	$tol = 10^{-6}$	$tol = 10^{-2}$	$tol = 10^{-4}$	$tol = 10^{-6}$
BDF1	75275	34456	23026	154	97	114
BDF2	63854	31854	22371	130	90	106
BDF3	57163	30063	22095	117	85	103
SS	60			12		

5. OTHER PROJECT ACTIVITIES

5.1 Laghos CTS-2 benchmark

The CEED team at LLNL developed a version of the Laghos miniapp to be used in the second edition of the Commodity Technology Systems procurement process. These systems leverage industry advances and open source software standards to build, field, and integrate Linux clusters of various sizes into production service. The CTS-2 version of Laghos features RAJA backend, improved robustness and figure of merit computation, and is available at <https://github.com/CEED/Laghos/releases/tag/cts2>.

5.2 Outreach

CEED researchers were involved in a number of outreach activities, including participation in the SIAM CSE19 conference, where CEED organized two minisymposiums (16 talks total) on Exascale Software for High-Order Methods and Exascale Applications with High-Order Methods and CEED researchers participated in various additional minisymposiums related to Batched BLAS, meshing, DG methods, matrix-free solvers, and more. CEED organized a successful breakout session on high-order methods and application at the ECP third annual meeting in Houston and we began the planning for CEED’s 3rd annual meeting to be held at Virginia Tech in August 2019.

6. CONCLUSION

The deliverable for this milestone was the release of CEED 2.0 – the CEED software distribution consisting of integrated libraries and software packages enabling efficient high-order discretizations on unstructured grids provided through the CEED website, the CEED GitHub organization, and Spack.

The artifacts delivered include a consistent build system based on 12 Spack packages + a the CEED meta-package, documentation and verification of the build process, as well as improvements in the integration between different CEED components. See the CEED website, <http://ceed.exascaleproject.org/ceed-2.0/> and the CEED GitHub organization, <http://github.com/ceed> for more details. As part of CEED 2.0, we also released the next version of libCEED (v0.4), which contains four new CPU backends, two new GPU backends, CPU backend optimizations, initial support for operator composition, performance benchmarking, and a Navier-Stokes demo.

In addition to details and results from these efforts, in this document we report on other project-wide activities performed in Q2 of FY19, including: a version of the Laghos miniapp for CTS-2 procurement, extensive GPU performance baseline benchmarking for the Urban and ExaSMR applications, papers, minisymposium presentations, and other outreach efforts.

ACKNOWLEDGMENTS

This research was supported by the Exascale Computing Project (ECP), Project Number: 17-SC-20-SC, a collaborative effort of two DOE organizations—the 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—to support the nation’s exascale computing imperative.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, LLNL-TR-771107.