ECP Milestone Report
Engage first wave ECP/CEED applications
WBS 1.2.5.3.04, Milestone CEED-MS1

Misun Min
Paul Fischer
Vladimir Tomov
Elia Merzari
Robert Rieben
Tzanio Kolev

March 31, 2017

ECP Milestone Report
Engage first wave ECP/CEED applications
WBS 1.2.5.3.04, Milestone CEED-MS1

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

March 31, 2017

ECP Milestone Report
Engage first wave ECP/CEED applications
WBS 1.2.5.3.04, Milestone CEED-MS1

Approvals

Submitted by:

Tzanio Kolev, LLNL
CEED PI

Date

Approval:

Douglas B. Kothe, Oak Ridge National Laboratory
Director, Applications Development
Exascale Computing Project

Date

Revision Log

Version	Creation Date	Description	Approval Date
1.0	March 31, 2017	Original	

EXECUTIVE SUMMARY

The CEED co-design center aims to impact a wide range of ECP application teams through focused 1-to-1 interactions, facilitated by CEED application liaisons, as well as through 1-to-many interactions, based the development of easy-to-use discretization libraries for high-order finite element methods.

In this milestone we identified the first wave of ECP applications that can potentially give a highly impactful outcome within short time frame for the first year of the project. For this, we focused on the following applications that already use the CEED-developed Nek and MFEM high-order software:

- Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors - SMRs (ORNL): Nek5000
- Next-gen Multi-physics Simulation Code - MARBL (LLNL): MFEM

The project goal of the small modular reactors (SMRs) is to analyze a full-core system with 40 sub-assemblies, where each sub-assembly consists of 17×17 pins. The target simulations require scalable simulations up to one billion spectral elements and trillions of grid points. The computations require complex mesh structures, external and internal fluid mixing, advanced RANS/LES models, efficient time stepping, and improved parallel mesh I/O support. The thermal-hydraulics simulations are based on high-order discretizations being developed in CEED (initially, with Nek5000) and are coupled to Monte-Carlo-based neutronics calculations.

MARBL is targeting high-energy-density physics problems, including single fluid multi-material hydrodynamics and radiation/magnetic diffusion simulation, with applications in inertial confined fusion, pulsed power experiments, and equation of state/material strength analysis. The goal of the project is to enhance modular physics simulation capabilities with increased performance portability and flexibility. The code uses high-order methods based on arbitrary Lagrangian-Eulerian, direct Eulerian, and unstructured adaptive mesh refinement developed within CEED (initially, using MFEM).

As part of the milestone, we also appointed application liaisons on the CEED team that will be responsible for engaging each application, identifying its challenging discretization needs, outlining a plan for interaction with CEED. Working with all CEED thrusts, the liaisons will be ultimately responsible for ensuring the adoption of CEED technologies in each application as well as for demonstrating their benefits.

The selected liaisons were:

- Vladimir Tomov (LLNL) for MARBL
- Elia Merzari (ANL) for SMRs.

In this document we are also reporting on additional CEED activities performed in Q2 of FY17, including initial engagements with other ECP applications and multiple team integration tasks.

TABLE OF CONTENTS

Executive Summary **vi**

List of Figures **viii**

1 Introduction **1**

2 SMR project **1**
2.1 Overview . . . 1
2.2 Application liaison . . . 3
2.3 Plan for CEED interactions . . . 3
2.4 Q1 Accomplishments . . . 3

3 MARBL project **5**
3.1 Overview . . . 5
3.2 Application liaison . . . 6
3.3 Plan for CEED interactions . . . 6

4 Exploration of other ECP application engagements **6**
4.1 Multiscale Coupled Urban System (ANL) . . . 6
4.2 Transforming Combustion Science and Technology with Exascale Simulations (SNL) . . . 6
4.3 Cloud-Resolving Climate Modeling of the Earths Water Cycle (SNL) . . . 6

5 Other project activities **7**
5.1 Team Integration: kickoff, collaboration hub, funding, staffing . . . 7

6 Conclusion **7**

LIST OF FIGURES

1 ECP Application Project: Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors using Nek5000. 2
2 Comparison: lightweight MPI (left)vs. regular MPI (right) for scalar mass-matrix inversion performed on ALCF Cetus using 16384 MPI ranks (512 nodes with 2 MPI ranks per core). . 4
3 Mass-matrix inversion performance for Nek5000 on ALCF Cetus: (left) performance vs. n/P for the scalar case and (center) the vector case; (right) the ratio of the vector vs. scalar performance for each data point from the left and center panels. 4
4 ECP Application Project: MARBL Next-gen Multi-physics Simulation Code using MFEM. . 5

1. INTRODUCTION

The CEED co-design center aims 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.

In this milestone, we identified the first wave of ECP applications that can potentially give a highly impactful outcome within the first year of the project. To this end, we focus on the following applications that have already identified CEED project codes (Nek and MFEM) to provide efficient high-order discretizations:

- Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors - SMRs (ORNL): Nek5000
- Next-gen Multi-physics Simulation Code - MARBL (LLNL): MFEM

As part of the milestone, we appointed application liaisons on the CEED team that are responsible for engaging each application, identifying its challenging discretization needs, and outlining a plan for interaction with CEED. Working with all CEED thrusts, the liaisons will be ultimately responsible for ensuring the adoption of CEED technologies in each applications as well as for demonstrating their benefits.

In addition, we explored opportunities for future engagements with applications in the second and third waves of the ECP/CEED applications that include Multiscale Coupled Urban System (ANL), Transforming Combustion Science and Technology with Exascale Simulations (SNL), and Cloud-Resolving Climate Modeling of the Earths Water Cycle (SNL). In this group, the Urban project (which received FY17 seed funding), was identified as a potential early-impact application because their simulations can readily exploit existing CEED codes.

2. SMR PROJECT

2.1 Overview

The SMR project involves coupled thermal-hydraulics/neutronics analysis, with the latter based on Monte Carlo methods and the former based on high-order discretizations of the Navier-Stokes and energy equations for accurate simulation of turbulent transport.

The objective of this project is to use exascale compute platforms to carry out extreme-fidelity simulations of SMRs, namely to enable verification and validation of current analysis methods by generating detailed power distributions during startup conditions, full cycle depletion analysis, and fast transient behavior characteristic of accident scenarios (see Figure 1). The project requires fast and efficient turbulence simulation capabilities that will scale to exascale platforms. Work within CEED will ensure that all aspects of the Nek5000 workflow scale to the target problem space, which involves $n > 10^{10}$ degrees of freedom in typical production runs. CEED is to provide implementations that perform with maximum attainable efficiency and deliver low turn-around times. A shared milestone between SMR and CEED is to have high-performance GPU-based variant of Nek5000, as part of milestones CEED-MS8 ("Initial integration of CEED software in 1st wave apps, explore future apps") and CEED-MS20 ("Performance tuning of CEED software and 1st wave apps").

Reactor thermal-hydraulics analysis is based on simulation of incompressible flow governed by the Navier-Stokes equations and energy transport,

$$\begin{aligned}\frac{\partial \mathbf{u}}{\partial t} &+ \mathbf{u}\cdot\nabla\mathbf{u} = -\nabla p + \tfrac{1}{Re}\nabla^2\mathbf{u}, \qquad \nabla\cdot\mathbf{u} = \mathbf{0}, \\ \frac{\partial T}{\partial t} &+ \mathbf{u}\cdot\nabla\mathbf{T} = \tfrac{1}{Re\,Pr}\nabla^2 T,\end{aligned} \tag{1}$$

where $\mathbf{u}(\mathbf{x},t) = (u_1,\, u_2,\, u_3)$ represents the fluid velocity components as a function of space and time, p is the pressure, and T is the temperature. The Reynolds number based on fuel pin diameter will be $Re \approx 10^4$–10^5, indicating that the flow is fully turbulent. The nondimensional parameter Pr is the Prandtl number and is typically in the range $Pr = .01$–10, depending on the coolant.

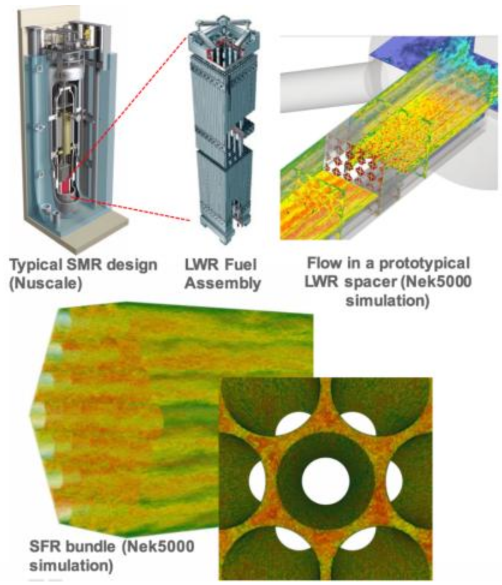

Figure 1: ECP Application Project: Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors using Nek5000.

Heat transfer from the fuel-pins to the coolant is of major importance in determining peak fuel temperatures and overall thermal efficiency. Enhancing heat transfer and ensuring a uniform thermal distribution are key to realizing improved reactor performance. Effective realization of these objectives requires detailed analysis of turbulent thermal transport in the large complex domains that constitute a reactor core. Even for an SMR (which is relatively small), one has tens-of-thousands of passages, with length to diameter ratios on the order of 100-200. Figure 1 gives some indication of the range of scales to be considered. The simulation depicted in the upper right panel of Figure 1 involves a 5×5 pin bundle and required 200 million grid points to capture the turbulence (which, crucially, is different before and after the grid-spacer, as indicated in the figure). Scaling to an SMR core comprising forty 17×17 arrays would lead to $n = 10^{11}$–10^{12} gridpoints, which is clearly an exascale problem.

Accurate simulations of turbulent thermal transport over such a large range of scales calls for minimal numerical dissipation and dispersion, for which high-order methods are ideally suited. While higher order can significantly reduce the required number of gridpoints, n, it does not immediately follow that a reduced n yields a savings in simulation time. Spectral element methods derive significant performance benefits from their local structure, in which solution and data within each of E deformable 3D hexahedral (brick) elements are represented in terms of Nth-order tensor-product polynomials, that is,

$$u(\mathbf{x})|_{\Omega^e} \quad = \quad \sum_{i=0}^{N}\sum_{j=0}^{N}\sum_{k=0}^{N} u^e_{ijk}\, h_i(\xi)\, h_j(\eta)\, h_k(\zeta),\ (\xi,\eta,\zeta) \in [-1,1]^3,\ e \in [1,\ldots,E]. \tag{2}$$

Here, $h_i(\xi)$ is the ith cardinal Lagrange interpolant on the Gauss-Lobatto-Legendre (GLL) quadrature points, part of the underlying SE basis functions on the reference element, $(\xi,\eta,\zeta) \in \hat{\Omega} := [-1,1]^3$, and $\{u^e_{ijk}\}$ are the tensor-product basis coefficients on element Ω^e. Geometric deformations are based on isoparametric mappings of the geometry,

$$\mathbf{x}|_{\Omega^e} \quad = \quad \mathbf{x}^e(\xi,\eta,\zeta) = \sum_{i=0}^{N}\sum_{j=0}^{N}\sum_{k=0}^{N} \mathbf{x}^e_{ijk}\, h_i(\xi)\, h_j(\eta)\, h_k(\zeta). \tag{3}$$

Spatial derivatives are computed by differentiating the h_is in $\hat{\Omega}$ and using the chain rule to map to Ω^e.

Ostensibly, (2) requires $O(N^3)$ operations per gridpoint. However, evaluation of (2) or any of its derivatives on any tensor-product set (μ_p, μ_q, μ_r) takes the form $\underline{\tilde{u}} = D^{lmn}\underline{u} = D^l \otimes D^m \otimes D^n \underline{u}$, with $D^l_{ij} := h^{(l)}_j(\mu_i)$ being the lth-derivative of the basis function [1]. This form can be implemented as fast **matrix-matrix products** (in fact, tensor-contractions, which are central to many machine-learning algorithms), such that operator evaluation costs scale essentially as the number of gridpoints $n \approx EN^3$, corresponding to the total data movement, rather than some higher power in N. The local matrix operators have $(N+1)^6$ nonzeros, the tensor-product forms can be evaluated in only $12(N+1)^4$ operations with only $6(N+1)^3$ storage. Moreover, the local structure allows for fast block inversion, which is useful for preconditioning.

As a part of the code suite forming the basis for CEED, Nek5000 will be used as a common framework for high-performance SMR simulations on next generation HPC platforms with thorough performance comparisons to identify optimal approaches for the variety of substeps that constitute time-advancement of a multi-physics code.

2.2 Application liaison

The application liaison for SMR is Elia Merzari from ANL. Elia has over 7 years of experience working on Nek5000 with a particular focus on applications in nuclear engineering. He has received several DOE ALCC allocations and applied Nek5000 at scale for a variety of applications related to nuclear reactors. In particular he has extensive experience with rod bundle flows, of particular relevance for the SMR app.

2.3 Plan for CEED interactions

- Provide spectral elements on GPUs and/or Phis in Nek5000 for SMRs in FY17 Q4 (CEED-MS8).
- Demonstrate improved performance on GPUs/Phis for SMRs in Nek5000 in Fy18 Q4 (CEED-MS20).
- Explore the use of nonconforming AMR mesh capabilities developed in CEED, in conjunction with hybrid RANS/LES models in SMRs in FY19 Q3 (CEED-MS29).
- Better AMG solvers for high-order discretizations in the SMRs project in FY20 Q3 (CEED-MS40).

2.4 Q1 Accomplishments

- Identified SMR needs from CEED:
  Scalable simulations for the full core system analysis, requiring complex mesh structures, external and internal fluid mixing, advanced RANS/LES models, efficient time stepping, and improved parallel mesh I/O support.
- Established close interactions with the reactor app team to initiate GPU-enabled core compute kernels of Nek5000: Activities include
  - Setup of GPU-enabled Nek5000 branch "`openacc`" at `https://github.com/Nek5000/`.
  - Collaborate implementing OpenACC version of Nek5000.
  - Completed OpenACC implementation for the spectral element multigrid (SEMG) routines.
  - Will continue OpenACC porting for other routines to be completed by CY17/Q3.
- Established close interactions with vendors to release GPU-enabled Nek5000 miniapp (Nekbone): Activities include
  - Setup of Nekbone branch "`cuda-openacc`" at `https://github.com/Nek5000/Nekbone` that is based on an existing GPU-enabled Nekbone version.
  - Restructured and tuned an existing OpenACC- and OpenACC+CUDA-enabled Nekbone based on our work [2, 3].
  - Tested on OLCF Titan Nvidia K20X & ANL JLSE with Nvidia P100.
  - To be released to the vendors by CY17/Q1.
- Major implementation in Nek5000 kernel (discretizations, operator evaluation and communication): Activities include tuned core kernel of operator evaluations in Nek5000 pressure solver in order reduced communication latency.

- Established close interactions with MPICH team: Activities include
  - Interaction with MPICH team to use light weight MPI and performance tested on Cetus on 512 nodes (32 MPI ranks per node) up to strong scaling limit with only one element with degrees of freedom ranging from $2^3 \sim 12^3$ per MPI rank. (see Figure 2.)
  - Performed tests also on Theta/KNL. (see Figure 3.)
  - Paper to be submitted in CY17/Q1.

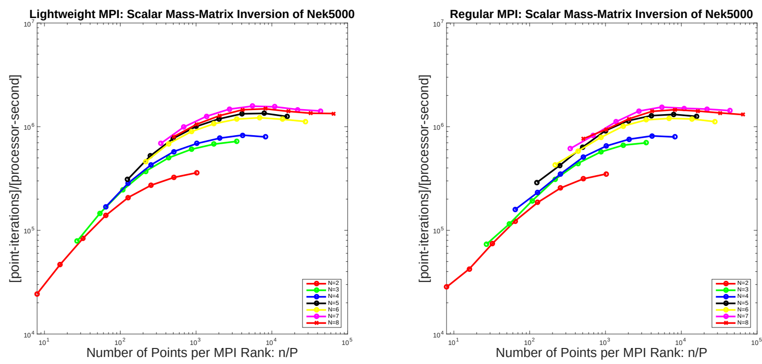

Figure 2: Comparison: lightweight MPI (left)vs. regular MPI (right) for scalar mass-matrix inversion performed on ALCF Cetus using 16384 MPI ranks (512 nodes with 2 MPI ranks per core).

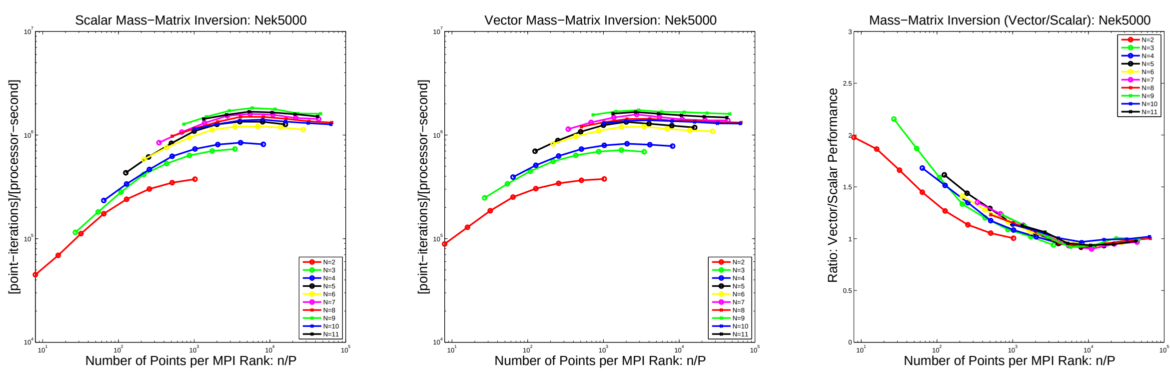

Figure 3: Mass-matrix inversion performance for Nek5000 on ALCF Cetus: (left) performance vs. n/P for the scalar case and (center) the vector case; (right) the ratio of the vector vs. scalar performance for each data point from the left and center panels.

3. MARBL PROJECT

3.1 Overview

The project goal of MARBL is to enhance modular physics simulation capabilities with increased performance portability and flexibility for single fluid multi-material hydrodynamics and radiation/magnetic diffusion simulation, with applications in inertial confined fusion, pulsed power experiments, and equation of state/material strength analysis (see Figure 4). The initial targets of the CEED team are the optimization of MARBL's (i) Lagrangian hydrodynamics and (ii) remap phases.

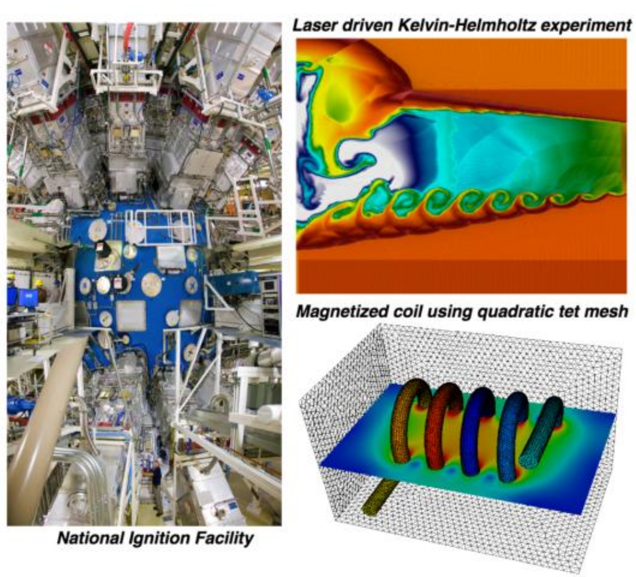

Figure 4: ECP Application Project: MARBL Next-gen Multi-physics Simulation Code using MFEM.

The Lagrangian hydrodynamics phase [4, 5] of MARBL is based on high-order finite element discretization on moving curved meshes. Using appropriate quadrature rules, the following matrices are assembled:

$$(M_v)_{ij} = \int \rho w_j w_i\,, \quad [(M_e)_{ij}]_k = \int \rho_k \phi_j \phi_i\,, \quad [F_{ij}]_k = \int (\sigma_k : \nabla w_i)\phi_j\,,$$

where k is the material index, w_i is an H^1 basis function, ϕ_j is an L^2 basis function, and σ_k is a material-dependent stress tensor. The M_v and $[M_e]_k$ matrices need to be assembled only once, while the assembly of F_k is performed at each time step, and hence is the most computationally-intensive section of the Lagrangian code. The global M_v and zone-based $[M_e]_k$ mass matrices are then inverted at each time step to solve the following system of linear equations:

$$M_v \frac{dv}{dt} = -\sum_k F_k \cdot 1\,, \quad [M_e]_k \frac{de_k}{dt} = F_k^T \cdot v\,.$$

MARBL's remap phase [6, 7] is based on a high-order discontinuous Galerkin finite element formulation. In this case, the mass matrix M_v is assembled at every time step, along with the matrices

$$M_{ij} = \int \phi_j \phi_i\,, \quad (K_v)_{ij} = \int u \cdot \nabla w_j w_i\,, \quad K_{ij} = \sum_z \int_z u \cdot \nabla \phi_j \phi_i - \sum_j \int_f (\phi_i u \cdot n)_d [\phi_j]\,.$$

These matrices are then used to solve the linear systems

$$M_v \frac{dv}{d\tau} = K_v v\,, \quad M \frac{d\eta}{d\tau} = K\eta\,,$$

where the second type of equation is solved 3 times for every material, and includes non-variational monotonicity manipulations.

3.2 Application liaison

The application liaison for MARBL is Vladimir Tomov from LLNL. With over 3 years of experience in developing different modules of MARBL, Tomov is well familiar with its code and performance requirements. Tomov is also a contributor to the MFEM library and member of the CEED team.

3.3 Plan for CEED interactions

As many parts of MARBL are based on MFEM, the CEED and MARBL teams will collaborate to improve the MARBL-relevant MFEM technologies, as well as to develop an MFEM-based Lagrangian hydro miniapp, which will closely resemble MARBL's computationally-intensive kernels. Some specific tasks include:

- Develop an initial Lagrangian hydro miniapp will be in FY17 Q3 (CEED-MS6: "Identify initial kernels, BPs and miniapps; engage vendors and STs").
- Initial integration of MFEM's efficient partial assembly capabilities in MARBL's BLAST component in FY17 Q4 (CEED-MS8: "Initial integration of CEED software in 1st wave apps, explore future apps").
- Optimization of the Lagrangian hydro miniapp, including for example, capabilities for partial assembly, solvers for partially assembled high-order operators, matrix free computations, along with GPU implementations in FY18 Q2 (CEED-MS15: "Update kernels, BPs and miniapps; push/pull with vendors and STs").
- The MFEM and MARBL teams will work together to apply the developed techniques in the application's Lagrangian and/or remap phases as part of the FY18 Q4 milestone CEED-MS32: "Performance tuning of CEED software and 1st wave apps".
- Efficient implicit assembly-free high-order algorithms for scalar and/or electromagnetic diffusion in MARBL in FY20 Q3 (CEED-MS40: "Efficient solvers for high-order discretizations").

4. EXPLORATION OF OTHER ECP APPLICATION ENGAGEMENTS

4.1 Multiscale Coupled Urban System (ANL)

One of the goals of the Urban exascale project is to accurately simulate the urban turbulence. Urban turbulence involves both convection, shear and high Reynolds number flows. CEED team will be engaged to improve its high-fidelity RANS/LES simulations with efficient time stepping for the test site of city of Chicago. The initial effort has been on developing CAD models and meshes for the city. We have developed models for a few buildings around Lake Michigan, Chicago and are performing steady-state solution to provide initial conditions for RANS simulations.

4.2 Transforming Combustion Science and Technology with Exascale Simulations (SNL)

The combustion project is targeting simulation of internal combustion engines. Proposed work includes multicycle simulations of full cylinder geometries with moving pistons and valves, fuel spray modeling, and detailed chemical kinetics. The Reynolds numbers in typical intake ports ranges from $Re = 10^4$–10^6, depending on RPM, so the flow coming into the cylinder is fully turbulent. Accurate simulation of detailed turbulence at these levels is most efficiently realized with high-order methods. Members of the CEED team have been working on the intake-port problem and on the fundamental combustion phenomena and are in contact with the Combustion ECP team.

4.3 Cloud-Resolving Climate Modeling of the Earths Water Cycle (SNL)

The climate modeling project is targeting to explore its full potential to scientifically and computationally advance climate simulation and prediction with a cloud-resolving convective parametrization into the DOE ACME Earth System model using the Multiscale Modeling Framework (MMF). Their superparametrization will be designed to make full use of GPU systems and they will refactor and port other key components of the ACME model for GPU systems. CEED team's efforts on GPU and communication library will have an opportunity to find overlapping interests in optimizing their programming models and scalable performance for simulating the performance-critical atmosphere, ocean and ice components.

5. OTHER PROJECT ACTIVITIES

5.1 Team Integration: kickoff, collaboration hub, funding, staffing

The CEED team consists of 30+ researchers from 7 different institution, geographically dispersed in the US. To achieve a focused common effort, we organized the project in four inter-connected thrusts: Application (AP), Hardware (HW), Software (SW) and Finite Element (FE), cross-cutting between the different institutions:

- Management: Director (Tzanio Kolev, PI), Deputy Director (Paul Fischer)
- Thrust Leads: Misun Min (AP), Veselin Dobrev (FE), Jed Brown (SW), Jack Dongarra (HW)

To enable efficient collaboration, we picked GitHub, github.com/ceed, as an online collaboration hub for the project, based on a combination of public and private repositories in the CEED organization on GitHub. We also started holding online weekly meetings for understanding CEED key components, and organized a project kickoff meeting for planning and identifying technical strategies in detail. Project requests for ALCF/OLCF Director's Discretionary allocations were made at the early stage of the CEED project during the period of November–December 2016. These activities are summarized as:

- Set up project-wide and thrust-specific mailing lists.
- Establish a workflow for using GitHub as online developer center for the project.
- Establish weekly CEED hackatons on Friday afternoons by Gitter chat.
- Initiate CEED kickoff meeting including two days of talks and a hackaton on last day.
- Request ALCF/ORNL projects for CPU/GPU hours allocations.

In addition to the 2 DOE labs: LLNL and ANL, CEED includes subcontracts between LLNL and the 5 university partners: UTK, VTech, RPI, UIUC and CU Boulder. The current status of the CEED funds is:

- Funding received at LLNL and ANL
- UIUC, UTK, VTech and RPI contracts fully executed
- CU (200K) contracts in final stages of preparation

As of the end of Q2 FY17, all open positions on the project (staff, postdocs, graduate students) are either filled in, or we are in the process of extending offers.

6. CONCLUSION

We have appointed two application liaisons on the CEED team: Vladimir Tomov (LLNL) for MARBL and Elia Merzari (ANL) for SMRs. For both applications, we have developed plans for CEED interactions to ensure that our activities are motivated by and beneficial for these applications.

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-728221.

REFERENCES

[1] M.O. Deville, P.F. Fischer, and E.H. Mund. *High-order methods for incompressible fluid flow.* Cambridge University Press, Cambridge, 2002.

[2] Matthew Otten, Jing Gong, Azamat Mametjanov, Aaron Vose, John Levesque, Paul Fischer, and Misun Min. An MPI/OpenACC implementation of a high-order electromagnetics solver with GPUDirect communication. *Int. J. High Perform. Comp. Apps.*, 30(3):320–334, 2016.

[3] Jing Gong, Stefano Markidis, Erwin Laure, Matthew Otten, Paul Fischer, and Misun Min. Nekbone performance on gpus with openacc and cuda fortran implementations. *Special issue on Sustainability on Ultrascale Computing Systems and Applications: Journal of Supercomputing*, 72:4160–4180, 2016.

[4] V. A. Dobrev, T. V. Kolev, and R. N. Rieben. High-order curvilinear finite element methods for Lagrangian hydrodynamics. *SIAM J. Sci. Comp.*, 34(5):B606–B641, 2012.

[5] V. A. Dobrev, T. V. Kolev, R. N. Rieben, and V. Z. Tomov. Multi-material closure model for high-order finite element Lagrangian hydrodynamics. *Internat. J. Numer. Methods Engrg.*, 82(10):689–706, 2016.

[6] R. W. Anderson, V. A. Dobrev, T. V. Kolev, and R. N. Rieben. Monotonicity in high-order curvilinear finite element arbitrary Lagrangian–Eulerian remap. *Internat. J. Numer. Methods Engrg.*, 77(5):249–273, 2015.

[7] R. W. Anderson, V. A. Dobrev, T. V. Kolev, D. Kuzmin, M. Quezada de Luna, R. N. Rieben, and V. Z. Tomov. High-order local maximum principle preserving (MPP) discontinuous Galerkin finite element method for the transport equation. *J. Comput. Phys.*, 334:102–124, 2017.