Proposal to the National Science Foundation's
Program for Cyberinfrastructure for Sustained Scientific Innovation for

Collaborative Research: Frameworks:
Community-Based Weather and Climate Simulation With a Global Storm-Resolving Model

Principal Investigator:
David A. Randall, Colorado State University

Co-Principal Investigators:
James Hurrell, Colorado State University
Andrew Gettelman, National Center for Atmospheric Research
Richard Loft, National Center for Atmospheric Research
William Skamarock, National Center for Atmospheric Research

tel: (970) 222-4983
email: david.randall@colostate.edu

To be submitted November 1, 2019

Proposal Summary

Overview

The open-source Community Earth System Model (CESM) is both developed and applied to scientific problems by a large community of researchers. The CESM includes sub-models of the atmosphere, ocean, land surface, and sea ice. It is principally funded by the National Science Foundation (NSF) and managed by the U. S. National Center for Atmospheric Research (NCAR). *The CESM is critical infrastructure for the U.S. climate research community*.

In the atmosphere and ocean components of the CESM, the adiabatic terms of the partial differential equations that express conservation of mass, momentum, and thermodynamic energy are solved numerically using what is called a ''dynamical core.'' Atmosphere and ocean models also include parametric representations, called ''parameterizations,'' that are designed to include the effects of processes that occur on scales too small to be represented on the model's grid. Despite decades of work by many scientists, today's parameterizations are still problematic and limit the utility of earth system models for many applications. However, recent advances in computer power have made it possible to parameterize less, by using grid spacings on the order of a few kilometers over the entire globe. These "global storm-resolving models" (GSRMs) can only be run on today's fastest computers. GSRMs are under very active development at a dozen or so modeling centers around the world. Unfortunately, the current formulation of the CESM prevents it from being run as a GSRM.

Our project, called EarthWorks, will create a new, openly available model by leveraging CESM but extensively modifying it so that it can be run as a GSRM. The model we envision could, by mid next decade, be applied to both weather forecasting and climate simulation. To accomplish this goal, we will use closely related dynamical cores for the atmosphere and ocean that have been developed at NCAR and at the Los Alamos National Laboratory, respectively, and are well-suited for applications that require very high spatial resolution. All components of our model will use the same very high-resolution grid. The use of kilometer-scale resolution makes it possible to eliminate the particularly troublesome parameterization of deep cumulus convection (i.e., thunderstorms). The remaining parameterizations (e.g., for turbulence, cloud microphysics, and radiation) are suitable for use with the fine grid. The component model codes are close to completion and are currently being tested on graphic processor units (GPUs). The architecture of the completed model will be simple but powerful.

We will apply the model to pressing scientific problems in both numerical weather prediction and climate simulation. The model will be particularly well positioned to study important problems at the weather/climate interface, and our team has the breadth of expertise to address both.

We will make the open-source model and input datasets (topography, etc.) available via GitHub. We will create a detailed model documentation, and seek official CESM support. Pending such support, the current investigator team will use the agency proposal process to manage support of the model.

Intellectual merit

The intellectual merit of our project lies in combining well matched model components with emerging trends in exascale computing to create a new and uniquely capable Earth System Model.

Broader impacts

The broader impacts of this project include the benefits to the national and international research community from free access to a well documented GRSM. In addition, we will train undergraduates, graduate students and postdoctoral fellows in the nuts and bolts of exascale programming and Earth System Model development. Specifically, we will leverage university partners at EPSCoR states and a long-established and successful computational science program (Summer Internships in Parallel Computational Science) at NCAR to build a pipeline of exascale model-development talent. Finally, by providing access to previously unresolved scales, EarthWorks will advance the rapidly merging sciences of weather forecasting and climate simulation, both of which are of critical importance to society.

Brief Synopsis

(Aimed at non-scientists, and limited to 500 characters)

We will modify an existing computer model of the atmosphere, ocean, and land surface so that it can predict weather and climate on fine spatial scales, comparable to the size of a small town. The modified model will represent the atmosphere, the ocean, and the land surface all on the same “grid” of points. We will use the model to do both weather forecasting and climate simulation. We will make the model freely available to scientists everywhere.

Project Description

1. Results from Prior Support

PI David Randall was the Director of the Center for Multiscale Modeling of Atmospheric Processes (CMMAP), an NSF Science and Technology Center (ATM-0425247, July 1, 2006–June 30, 2016, $37,505,835). CMMAP's focus was on development of improved methods for simulation of the global atmospheric circulation, with a special focus on cloud processes. CMMAP's Broader Impacts included an extensive education and outreach program that worked with K12 students, supported undergraduate interns, and ran an annual teacher-training course. Randall is currently PI or co-PI on several NSF grants: 1. "Implementation and evaluation of the unified parameterization in NCAR Community Atmospheric Model" (AGS-1538532, 07/01/2016 – 06/30/2020, $449,437), which is testing a generalized version of the Unified Parameterization developed by Arakawa and Wu. 2. "Development and Testing of a Global Quasi-3-D Multi-scale Modeling Framework" (AGS-1500187, 02/01/2016 – 01/31/2020, $656,746), which is testing a second-generation" super-parameterization in the CAM. 3. "Air-sea Interaction and Island Geography Impacts on MJO Initiation and Propagation Through the Maritime Continent" (AGS-1445191, 03/01/2015 – 02/29/2020, $513,084) which is aimed at understanding the mechanisms through which air-sea interaction affects convection within the MJO. 4. "Collaborative research: A teleconnection between the tropical Madden-Julian Oscillation and Arctic Sudden Stratospheric Warming events in warm climates" (AGS-1826643, 07/01/2018 – 06/30/2021, $418,199)which studying how the tropical troposphere influences the Arctic stratosphere. Results from these four projects have been reported in refereed publications. The Broader Impacts of the four projects include educating graduate students and creating improved models that are useful to society.

2. Introduction

Numerical weather prediction and climate simulation are based on global models of the atmosphere, land surface, oceans, and cryosphere. The conceptual development of such Earth System Models (ESMs) can be traced back to the beginning of the 20th century (Randall et al., 2019), and precursors of ESMs were among the first applications of electronic digital computers, starting around 1950. Over the decades, the evolving models have always needed the most powerful available computing platforms. ESMs of various complexities are now developed and used at dozens of centers around the world, including several in the United States. They are powerful tools for meeting the intellectual challenge of understanding the climate and the Earth system, and they are the only scientific tools capable of simulating the interactions of the myriad physical, chemical and biological processes that determine past, present and future climate. ESMs are also essential for making predictions of use to society and policymakers.

The open-source Community Earth System Model (CESM) is unique in that it is both developed and applied to scientific problems by a large community of researchers. *The CESM is critical infrastructure for the U.S. climate research community*. It is principally funded by the National Science Foundation (NSF) and managed by the U. S. National Center for Atmospheric Research (NCAR). Simulations performed with the CESM have made many significant contributions to climate research, ranging from paleoclimate applications (e.g., Otto-Bliesner et al. 2016) to the North American Multi-Model Ensemble (NMME, Kirtman et al. 2014) seasonal forecasting effort, which is led by the National Oceanic and Atmospheric Administration (NOAA) but supported by NSF and other U.S. agencies. Simulations with CESM have also been used extensively in both national and international assessments of climate science, such as those of the Intergovernmental Panel on Climate Change (IPCC) and the U.S. Global Change Research Program (USGCRP). The salient point is that CESM provides the broader academic community with a core modeling system that is used to investigate a diverse set of earth system interactions across multiple time and space scales.

Over three thousand peer-reviewed journal articles have been published based on work with the CESM, and several hundred scientists are actively involved in its development and application. Between 2012 and 2017, there were a total of over 4400 downloads of the CESM software. Since August of 2018, downloads of CESM-2 have averaged about 100 per week, according to the model's development team leader. The annual CESM workshop is typically attended by 300 scientists.

A hallmark of the CESM project is community governance of all its activities. Accordingly, development and production objectives and priorities emanate directly from the community of scientists who participate in the management of the CESM project. This includes 12 CESM model development and application working groups (Hurrell et al. 2013), which are teams of scientists that contribute to the development of individual component models and applications of the model to questions of interest. More information about the CESM management structure, the roles and responsibilities of the aforementioned groups, terms of use, and meeting schedules can be found on the CESM web page.

ESMs are essential to modern weather and climate science because the global coupled atmosphere-ocean-land-cryosphere system exhibits a wide range of physical and dynamical phenomena associated with physical, biological, and chemical feedbacks. Collectively, these interacting processes result in a continuum of temporal and spatial variability that, in theory, can be captured in ESMs. Achieving accurate results is an extremely challenging goal, however. Fundamental barriers to advancing weather and climate prediction on time scales from days to years, as well as long standing systematic errors in weather and climate models, are partly attributable to our limited understanding and capability for simulating the complex, multi-scale interactions intrinsic to atmospheric, ocean, and cryospheric motion (Hurrell et al. 2009).

In the atmosphere and ocean components of ESMs, the adiabatic terms of the equations that express conservation of mass, momentum, and thermodynamic energy are represented using what is called a ''dynamical core.'' Atmosphere and ocean models also include parametric representations, called ''parameterizations,'' that are designed to include the transports of energy, momentum, and other quantities by the unresolved or ''subgrid scale'' motions of the air and water, as well as by radiation and precipitation. For many ocean models, the subgrid-scale motions include very energetic "mesoscale" eddies and oceanic deep convection, which are critically important for the climate system. In both atmosphere and ocean models, all of the parameterized processes are formulated in terms of the fields that are directly simulated by the dynamical core.

Despite decades of work by many scientists, today's parameterizations are still problematic. This is especially true for parameterizations of cumulus convection, which have been needed up to now because cumulus clouds are very small compared to the grid cells of low-resolution global atmosphere models. Continuing increases in computer power have recently made it possible to avoid such parameterizations: Since about 2005, it has been possible to run a global atmospheric model with a horizontal grid spacing fine enough (about 5 km or less) to explicitly simulate the growth and decay of large and deep cumulus clouds (e.g., Satoh et al., 2005; Tomita et al., 2005). Such models do not need cumulus parameterizations. They are called "cloud-permitting" or "storm-resolving" models, rather than cloud-resolving models, because their grid spacing is still not fine enough to allow detailed simulations of individual large clouds, and the effects of small clouds must still be parameterized. Despite this limitation, numerous studies have shown that such global storm-resolving models (GSRMs) are able to realistically simulate important atmospheric processes that lower-resolution models miss (e.g., Stevens and Bony 2013).

A recent paper (Stevens et al., 2019) presents results from DYAMOND (The DYnamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains), an intercomparison of nine atmospheric GSRMs. Several of the models that participated in DYAMOND were developed in the U.S. Unfortunately, the Community Atmosphere Model (CAM; Neale et al. 2013), which is the atmosphere sub-model of the CESM, did not participate because *the current version of the CAM cannot be run as a GSRM*. Moreover, at the summer 2019 CESM Workshop held in Boulder, Colorado, it was argued that CAM development should focus on lower-resolutions for the indefinite future. The limited resources available for CESM development are a key reason for this strategy. The CESM and its user community of university and national laboratory researchers are in danger of being left behind as the international modeling community moves decisively towards more extensive use of GSRMs. This proposal represents our plan to do something about that.

The CESM is developed through a community-based system in which investigators both inside and outside NCAR can implement and test changes to a branched version of the model, and then, on the basis of those results, offer the changes for possible incorporation into the officially supported version. This system is both a strength and a weakness. It is a strength because new ideas can be gathered from the broader scientific community, rather than just from NCAR. It is a weakness because the quasi-democratic machinery needed for community-based model development is unwieldy. It involves multiple working groups, a Scientific Steering Committee, and an Advisory Board. Experience has shown that, in practice, this complicated system does not give rise to an architecturally coherent vision for development of the CESM as a whole.

We are proposing to create a GSRM capability for community use within the framework of the CESM, with a horizontal grid spacing of 4 km or finer, by carrying out the following steps:

- spin off a copy of the most recent version of CESM, including atmosphere, ocean, sea ice, and land surface models;
- implement, test, and document the various modifications described in this proposal; and then
- make the new GSRM available for applications and further development by the CESM community, as part of the open-access CESM infrastructure.

The result will be a GSRM that uses many standard CESM components. It will not do everything for everybody, but it will do some very important things extremely well, and it will significantly advance the scope of community modeling. We will work with NCAR software engineers to ensure that the model is compatible with the evolving CESM code base.

Our project will produce radical innovation on a short time scale, to the benefit of the broader research community. It will be led by a university PI and primarily carried out with heavy student and post-doc involvement, with technical support provided by NCAR. This is a novel approach designed to accelerate the development of advanced GSRM capabilities. We call our project *EarthWorks*.

3. The *EarthWorks* vision

a) Common component grids and new coupling infrastructure

The components (or sub-models) of the current version of CESM are as follows:

- CMEPS, the Community Mediator for Earth Predictive Systems (Danabasoglu et al., 2019), which is more fully described below;
- CAM6, the sixth version of the Community Atmosphere Model;
- POP, an ocean model developed at the Los Alamos National Laboratory, which is expected to be replaced by MOM6, an ocean model that is under development at the Geophysical Fluid Dynamics Laboratory (GFDL) of the National Oceanic and Atmospheric Administration (NOAA);
- CTSM, the Community Terrestrial System Model that represents the vegetated land surface;
- CICE, the Community Ice CodE, which represents sea ice and was developed at the Los Alamos National Laboratory;
- CISM, the Community Ice Sheet Model, which represents terrestrial ice sheets and was developed at the Los Alamos National Laboratory;
- MOSART, the Model for Scale Adaptive River Transport, which carries to the oceans the precipitation that falls on the continents; and
- WW3, an ocean surface wave model that was developed by the Environmental Modeling Center of the National Oceanic and Atmospheric Administration.

Further description of the components is given by Danabasoglu et al. (2019).

The current CESM defines one grid for the ocean and sea ice, and another for the atmosphere and land surface. There is also an "internal land ice grid." Finally, there is a river grid. These multiple grids (Hurrell et al., 2013; Danabasoglu et al., 2019) are necessary for several reasons. One of the most important is that eddies in the ocean are quite a bit smaller than the corresponding eddies in the atmosphere; as a result, the ocean model needs a finer grid than the atmosphere model.

Fortunately, the horizontal grid spacing of a storm-resolving atmosphere model is also suitable for use with an "eddy-resolving" ocean model that can explicitly simulate both mesoscale ocean eddies and oceanic deep convection. The same fine grid spacing is good for the land model, in part because it allows coastlines, mountains, lakes, rivers, and cities to be realistically represented. *Our plan is to use the same horizontal grid for all components of EarthWorks*. This "common grid" strategy is made possible by our intention to use a *horizontal grid spacing of 4 km or finer.* Relative to versions of CESM that use diverse grids, the common-grid architecture of *EarthWorks,* combined with CMEPS, will enable a lower operation count, less message-passing overhead, and reduced memory usage. It will simplify the logical structure of the model and make it easier to modify going forward.

In any ESM a "mediator" (sometimes called a "coupler") is needed to route data from source components to destination components, and to merge data from more than one source component to a particular destination component. As an example, surface temperatures from land, ocean and ice have to be merged into a single field to be sent to the atmosphere. A mediator can also carry out calculations. For instance, the surface moisture flux due to evaporation, which links the atmosphere and ocean models, is computed in the mediator. The mediator interpolates data from the ocean, ice and land grids to the atmosphere grid, and subsequently averages that data before sending it to the atmospheric component. Similar interpolations and averages are needed to couple the other components. The interpolations and averages must satisfy important physical constraints. For example, to conserve water mass the total amount of precipitation falling out of the atmosphere must match the total amount of precipitation received by the Earth's surface. Conservation of energy imposes similar constraints.

CESM is in the process of moving to a new coupling infrastructure based on widely used community standards, and *EarthWorks* will take advantage of this development. The new infrastructure is called CMEPS (Community Mediator for Earth Prediction Systems). It combines the Earth System Modeling Framework (ESMF; Hill et al., 2004) with the National Unified Operational Prediction Capability (NUOPC; Sandgathe et al., 2011, Theurich et al., 2016). It is being implemented through a partnership among NCAR, NOAA's Environmental Modeling Center (EMC) and NOAA's GFDL. CMEPS is a highly flexible tool that can replicate all coupling functions currently used in CESM. It facilitates sharing coupling code across organizations, technology transfer from research to operations, controlled experimentation, and simpler code through reusable abstractions such as field mapping and merging.

ESMF is a high-performance, flexible software infrastructure for building and coupling weather, climate, and related Earth science applications. It includes data structures and utilities that are designed to be used for developing individual models. The basic idea behind ESMF is that complicated applications should be broken up into coherent pieces, or components, with standard calling interfaces. In ESMF, a component may be a physical domain, or a function such as a coupler or I/O system. ESMF also includes toolkits for building components and applications, such as regridding, calendar management, logging and error handling, and parallel communications. NUOPC complements ESMF by defining conventions and templates for using ESMF to couple model components.

CMEPS will provide significant advantages to the CESM community, relative to the aging and unsupported coupling infrastructure that it replaces. CMEPS will make it easier to introduce new candidate components, and couple them to existing CESM components.

CMEPS is a highly flexible tool that can replicate all coupling functions currently used in CESM and serve as the primary coupling infrastructure for CESM. It facilitates sharing coupling code across organizations, technology transfer from research to operations, controlled experimentation, and simpler code through reusable abstractions such as field mapping and merging.

CMEPS has been implemented in the CESM by by the CESM Software Engineering Group, and is currently being tested. In order to be compatible with CMEPS, a component must implement a translation layer (called a "NUOPC cap") that maps model data to ESMF/NUOPC data structures and vice versa. All CESM components currently have such caps. CMEPS is targeted to be operational in CESM before the proposed May 1, 2020 start date of *EarthWorks*. We will use CMEPS to couple the common-grid components of ~~our model~~*EarthWorks*. The resulting architecture will offer some interesting opportunities for faster computational performance~~al~~ ~~speed up~~ by taking advantage of the common grid. This is one of the many benefits of kilometer-scale grid spacing.

b) Closely related dynamical cores for the atmosphere, ocean, and sea ice models

The atmosphere and ocean components of *EarthWorks* will both be based on the MPAS ("Model for Prediction Across Scales") dynamical core. The MPAS atmosphere dynamical core, which is called MPAS-A, was developed at NCAR by W. Skamarock and colleagues (Thuburn et al., 2009; Ringler et al., 2010; Skamarock et al., 2012; Park et al., 2013; Heinzeller et al., 2016; Judt, 2018). MPAS ocean (MPAS-O) was developed at the Los Alamos National Laboratory by T. Ringler and colleagues (Ringler et al., 2013; Petersen et al., 2015). Los Alamos has also developed an MPAS sea ice model, which is called MPAS-I. It is based on the CICE model used in CESM, and has been modified to work with MPAS-O (Turner, 2018). *A key point is that MPAS-A, MPAS-O, and MPAS-I all use spherical geodesic grids, so that it is straightforward to configure them on a common grid.*

c) Atmosphere model

MPAS-A is a nonhydrostatic atmospheric model that solves the fully compressible nonhydrostatic equations. The model uses finite-volume numerics discretized on centroidal Voronoi (nominally hexagonal) meshes using C-staggering of the prognostic variables, based on the work of Ringler et al. (2010). The MPAS spherical centroidal Voronoi mesh is unstructured, and allows for both quasi-uniform and variable horizontal resolution. *EarthWorks* will employ quasi-uniform meshes.

The approach that MPAS-A uses to solve the nonhydrostatic equations can be considered an extension of techniques that were developed over a period of decades for use in regional nonhydrostatic models Tests show that MPAS gives solutions as accurate as the widely-used Weather Research and Forecasting (WRF) model and similar models at both hydrostatic and nonhydrostatic scales. These techniques represent the state-of-the-art in mesoscale and cloud-scale modeling. MPAS-A is the first model to use the C-staggering on a geodesic grid. This allows MPAS-A to retain the accuracy of the many C-staggered mesoscale and cloud-scale models that preceded it.

The quasi-uniform centroidal Voronoi mesh used by MPAS-A is similar to the icosahedral (hexagonal) meshes used in other nonhydrostatic icosahedral atmospheric models, including NICAM (Satoh et al. 2005). These meshes provide nearly uniform and quasi-isotropic resolution over the globe (Heikes et al., 2013), and allow for good scaling performance on massively parallel architectures.

MPAS-A has undergone extensive testing and evaluation. The basic tests of the solver are described by Skamarock et al. (2012) and further tests in idealized flow configurations on the sphere were reported by Park et al (2013, 2014) and Klemp et al (2015). Tests in numerical weather prediction configurations are discussed by Davis et al (2016), which focuses on tropical cyclone forecasts with 15-km mesh spacing, and Wong and Skamarock (2016), which examines MPAS capabilities related to forecasting springtime convection over the central US employing convective-permitting resolution in a variable-resolution configuration. MPAS-A has also been used within the Data Assimilation Research Testbed (DART) ensemble data assimilation system (Ha et al., 2018) which demonstrated the robustness of MPAS-A in a cycling ensemble-DA framework. As will be discussed later, the MPAS-A dynamical core already runs on graphic processor units (GPUs). MPAS-A was also a participant in the DYAMOND inter-comparison mentioned in the introduction and summarized by Stevens et al. (2019).

We will implement the GPU version of the MPAS-A dynamical core into the CAM using the new System for Integrated Modeling of the Atmosphere (SIMA) infrastructure that is being developed by NCAR and community partners (Skamarock et al., 2019). Co-Investigators Gettelman and Skamarock have been instrumental in the development of SIMA. One goal of SIMA is to unify different atmospheric modeling efforts across NCAR. SIMA, a priority for NCAR, will be composed of a series of interoperable atmospheric components: dynamical cores, physical parameterizations, suites of parameterizations and even chemical models. It will be possible to configure these components in a variety of ways to satisfy different atmospheric application and workflow requirements. This framework will facilitate the changes needed to run CESM at high resolutions, and with a non-hydrostatic dynamical core. The SIMA framework should be ready for testing by *EarthWorks* in mid 2020. *EarthWorks* will assist in development of the SIMA framework with a leading-edge application that is representative of future architectures and use cases for SIMA.

SIMA also brings a flexible Application Programming Interface (API) called the Common Community Physics Package (CCPP). The CCPP is an API for physical parameterizations, and is being shared between CESM, MPAS-A, the Weather Research and Forecasting (WRF) model and NOAA. Thus EarthWorks will have access to physical parameterization "suites" (sets of interoperable parameterizations, including clouds, turbulence and radiation) developed for both climate and high-resolution short-term weather prediction, that can be tested with minimal software re-engineering. SIMA and the CCPP are being implemented now, with NCAR and community participation, and the weather and climate components are being led by Co-Investigators Skamarock and Gettelman.

We will configure *EarthWorks* to include a resolved stratosphere, by locating the top of the atmosphere model about 80 km above the Earth's surface (well above the stratopause) and using enough layers to resolve the vertical structure of the troposphere and stratosphere (Skamarock et al., 2019). This will enable simulation of wave propagation and other dynamical processes that govern the stratospheric circulation. It will also allow simulation of low-frequency variability of the high latitude tropospheric circulation (Thompson et al., 2002).

EarthWorks will use a suite of atmospheric physical parameterizations appropriate for 4-km grid spacing. We will simply omit any parameterization of deep cumulus convection. The physical parameterizations that we intend to use consist mainly of cloud microphysics, unified moist turbulence and radiation. Specifically, a new version of the CAM6 cloud microphysics with rimed ice (graupel or hail) will be used to handle higher updraft speeds found in deep convective systems (Gettelman et al., 2019). Other CAM parameterizations will be used with minimal modification, including the radiation parameterization (Rapid Radiative Transfer Model-GCMs or RRTMG) and a unified higher order closure scheme for moist turbulence (CLUBB, Golaz and Larson, 2002) that has been extended to all but deep convective cloud types and integrated with the CAM cloud microphysics for CAM6 (Bogenschutz et al., 2013). Coding modifications to those parameterizations will be strictly minimized.

These CAM6 physical parameterizations have been tested with grid spacings similar to those planned for *EarthWorks* (Gettelman et al 2019). Importantly, they are improved over similar schemes commonly used in mesoscale models because of their ability to conserve mass and energy, their numerical stability across a range of time step sizes, and their ability to efficiently include complex interactions with aerosols. That said, it is important to note that we intend to combine these parameterizations to run with 4-km grid spacing for EarthWorks. This will be a key test of the CCPP infrastructure and the cross-scale readiness of the physical parameterizations.

Because of the use of the CCPP API, *EarthWorks* can also use cloud, turbulence and boundary layer parameterizations developed for mesoscale models and used for DYAMOND experiments with little infrastructure work. This includes microphysics with a unified ice and snow phase (Eidhammer et al 2016), and cloud schemes being used for high-resolution weather forecast models (Thompson and Eidhammer 2014). While CAM schemes have not been used for storm resolving simulations, mesoscale schemes have rarely been tested in coupled climate mode.

d) Ocean model

MPAS-O is currently being used as the ocean model of E3SM (the "Energy Exascale Earth System Model"), which is a new open-source climate model supported by the U.S. Department of Energy (Golaz et al., 2019). Although MPAS-A is non-hydrostatic, MPAS-O is currently hydrostatic. The Los Alamos National Laboratory is committed to providing the GPU-enabled source for MPAS-O, including advanced features in developmental branches, together with advice/support in mesh generation, model configuration and analysis. See the Letter of Collaboration from Dr. Philip Jones.

e) Land-surface model

EarthWorks will follow CESM in using CTSM, which is an open-source (github) community model that represents a unification of the widely-used Community Land Model version 5 (CLM5, Lawrence et al., 2019;) and Noah-MP (Niu et al., 2011) models. CTSM is a process-based model that represents many hydrological, biogeophysical and biochemical processes. Historic and future land-cover and land-use changes can be prescribed, and urban climates are represented with an urban module. An active crop model with irrigation, fertilization, and other agricultural management practices can be used over agricultural land. Recently, CTSM has been extended to enable simulation of small-scale lateral transfers of snow, water, and heat. EarthWorks will benefit from the work of LILAC (Light-weight Infrastructure for Land-Atmosphere Coupling), an ongoing project to create a simple, portable interface for CTSM. LILAC is supported by NSF's Program for Cyberinfrastruction for Sustained Scientific Innovation. LILAC will simplify the CTSM toolchain for creating new surface datasets, which will be needed for the target resolution of EarthWorks.

f) Machine Learned Model Physics

Replacing human-crafted routines offers an opportunity to reduce model complexity, improve model predictions and better map models onto computer architecture trends, which are now being strongly influenced by the requirements of deep learning algorithms (deep neural networks). Recently, there has been success at NCAR, ECMWF and elsewhere in training emulators to replace or even upgrade complex physics code. For example, a neural network that emulates a prohibitively expensive binned microphysics autoconversion (rain formation) routine has run nine simulated years in the Community Earth System Model (CESM), producing a climate that, on initial inspection, looks reasonable. Neural net and random forest algorithms that have been trained on observational tower data have been shown to out-perform Monin-Obukhov similarity theory when compared with observations.

These investigations are proving useful but require more research before they can be put into production. EarthWorks will develop a machine learning (ML) emulator framework (MLEF) that will reduce the barriers to experimentation with parameterization emulators that may need speeding up. The MLEF will address an impediment to implementation of emulators in legacy HPC models, namely the integration of the ML algorithm into the host model, where differing data structures, languages and the lack of ML debugging tools make integrating the two worlds a tedious process. This MLEF could take the form of a Keras-Fortran API that could import existing Keras models and run them, and have the flexibility to interface directly with one or more C++ deep learning libraries without having to go through the Python layer first. This framework will be an important practical outcome of our project, because it can be adopted by other weather and climate model development teams.

4. Computational performance

a) Overall Approach and Performance Goals

As discussed earlier, we propose to combine elements of GPU-enabled MPAS atmosphere (MPAS-A) and MPAS ocean (MPAS-O) models with parts of the CAM/CESM modeling framework to create EarthWorks, an exascale-enabled, highly scalable adjunct to the CESM. Besides targeting GPUs, other innovative elements incorporated into EarthWorks will include the extensive use of task parallelism, a simplified coupling approach, and a smart, modular approach to integrating physics suites. For the end of our five-year project, our performance goals for a version of EarthWorks with 4 km global grid spacing are: 1) to demonstrate 0.5 simulated years per day (SYPD) for atmosphere-only simulations with high vertical resolution; and 2) 1 SYPD ~~for year~~ for ~~climate~~coupled simulations with fewer layers. Later in this proposal we present evidence that these goals are achievable given what we know about future exascale systems and the initial state of the software components. We expect to gain access to future exascale systems through separate allocation requests (e.g. XSEDE or INCITE).

The scope of this project's Data Management Plan only covers hosting the source code and initial datasets, which are quite modest in scale. The project team recognizes that the data volumes from EarthWorks experiments will be extreme (order 10s of petabytes). The management of these data will be dealt with in the allocation request proposals to various large-scale computing facilities. In the course of these experiments, EarthWorks will take advantage of lossy compression tools and methodologies (Baker, 2014 & Baker, 2018), in situ analysis in memory or on NVRAM, and parallel off-line analysis at the generating center, to cope with the data volumes.

The initial state of the required EarthWorks components is as follows. An open source version of the standalone GPU-enabled MPAS-A model (finite volume dynamical core with a WRF-like physics suite) is nearing completion. This work was performed, in part, by members of the EarthWorks team, through a partnership between NCAR and IBM (the International Business Machines Corporation), with support from students at the University of Wyoming and the Korea Institute for Science and Technology Information (KISTI). A version of this model is expected to go into production at The Weather Company as a component of the Global High-Resolution Atmospheric Forecasting System (GRAF) by the end of calendar year 2019. The

MPAS ocean dynamical core is currently being ported to GPUs by U.S. Department of Energy scientists. An initial version of MPAS-O is also expected to be completed by then end of calendar 2019 (personal communication, Phil Jones of Los Alamos National Laboratory). The EarthWorks project will create GPU-enabled versions of the atmospheric physics parameterizations. First, we intend to leverage available GPU ports of the radiation code. We have used estimates of the number of source lines of code (SLOC) for MG3 (8k SLOC) and CLUBB (32k SLOC) and have translated the code size into level of effort by referencing the team's experience refactoring MPAS-A physics. Other minor modifications for GPUs to the CESM and CAM infrastructure will also be required to accomplish this task.

b) Future Exascale Hardware and Software Roadmaps

Building ESMs capable of executing at global resolutions finer than 4 km requires large-scale high-performance computing (HPC) systems that allow a diversity of experiments and can be affordably built. A survey of the published results (Fuhrer, 2018, Govett, 2017) from some of the atmospheric models that have been successfully ported to graphics processing units (GPUs) shows that GPU acceleration can reduce power and cost requirements by 2.5-4 times: factors that drive the design of exascale systems toward GPUs.

The announced portions of the DoE pre-exascale and exacale computing roadmap give us a preview of what our target architectures look like in the 2020s: they will all include heterogeneous nodes with graphic processing units (GPU) acceleration, albeit from a more diverse set of suppliers, and the design of these architectures will be increasingly driven by deep learning use cases, another area where GPUs excel. For example, according to press-releases in March of 2019, the Aurora system at Argonne National Laboratory, due out in 2021, will be composed of heterogeneous computing nodes that team next-generation Intel Xeon CPUs and first-generation Intel GPUs – codenamed "Xe". The follow-on to the pre-exascale IBM Summit system at Oak Ridge National Laboratory will be called Frontier, and will have nodes with both CPUs and GPUs made by Advanced Micro Devices (AMD), Inc. Both Frontier and Aurora will be exascale systems.

With a diverse set exascale architectures in the offing it will be critical for ESM software to focus on performance and portability. Approaches that bridge the gap between hardware and complex ESM legacy software complexity are required: see Lawrence et al. (2018) for a detailed discussion. We find massive rewrites of ESMs (e.g. via GridTools, Kokkos, or OCCA) to be unattractive for a couple of reasons: First, the excessive porting time is inconsistent with the required rapid pace of model development; and second, the required skillsets do not mesh well with the community development paradigm of CESM, a core principle of this proposal.

In contrast, a directive-based approach leaves nearly all of the code unmodified, while providing legacy code access to new hardware without necessarily sacrificing performance on existing CPU platforms. The path forward for the directive approach on future exascale systems has been clarified in recent years. In September 2016, the DoE Exascale Computing Project (ECP) funded the Scaling OpenMP via Low-Level Virtual Machine for Exascale Performance and Portability (SOLLVE) project to provide the needed OpenMP support for exascale architectures. With the development and maturation of OpenMP directives for accelerators, we anticipate having to eventually migrate OpenACC directives to OpenMP to exploit these new systems. The EarthWorks GPU project team has budgeted for this eventuality, including dealing with compiler issues that are sure to arise in supporting these new technologies.

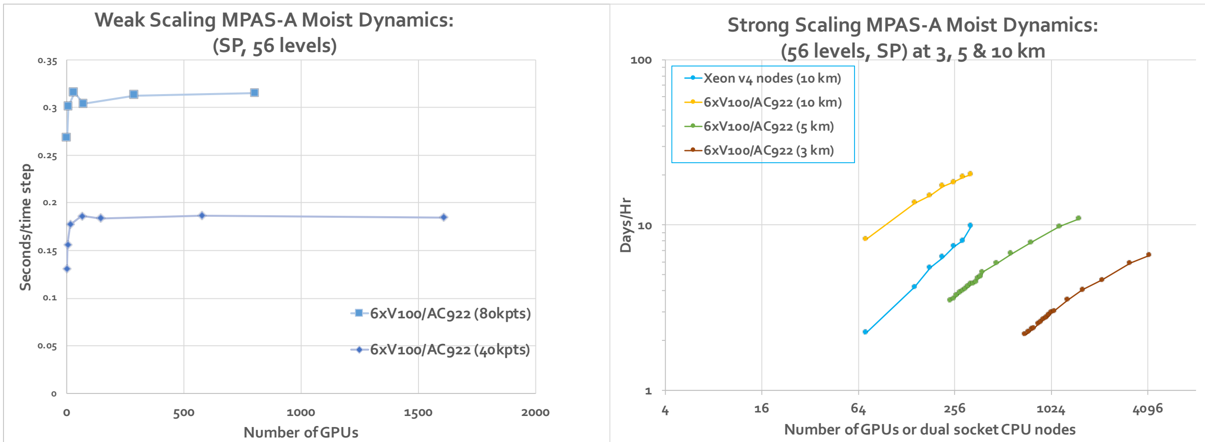

Fig 1: On the left we show weak scaling on an IBM AC922 cluster with 6 NVIDIA V100 "Volta" GPUs per node. Here the horizontal patch size is held to 40,960 horizontal points per node (diamonds) and 81,920 points per node (squares). The right panel shows"strong scaling curves" for MPAS-A on CPU and GPU-based architectures for single precision, 56 levels. NVIDIA V100 GPUs (10 km – orange; 5 km – green; 3 km – brown); and for 18-core Intel Xeon v4 E2697 CPUs (10 km - blue). The largest configuration used 4200 GPUs with a 3 km grid spacing.

c) *Performance Design Strategy*

Our strategy for achieving this performance in a complex ESM is based on the following three key elements:

- **Task Parallelism:** Heterogeneous nodes provide an opportunity to add task parallelism to the data parallelism of existing models in creative ways. Task parallelism provides a way to increase concurrency beyond the limits imposed by data parallism, and can serve to reduce communication and I/O overheads and load imbalances. CPUs can asynchronously execute I/O calls and certain model components while the GPUs continue uninterrupted computation. An example of the latter, discussed in more detail below, was demonstrated in the MPAS-A GPU project, where CPU-based radiative transport calculations were successfully overlapped with GPU-based model dynamics and physics components. Also, although architecture dependent, multiple tasks can share GPU resources, allowing thread occupancy to be boosted. This will be explored as a mechanism for obtaining GPU-resident concurrent execution of the ocean and atmospheric components.
- **Simplified Coupling Approach:** Mediator scalability is an issue for ESMs, like CESM, which make no assumptions about the grids used by different component models, and must provide distributed regridding and other functions that require substantial data motion. As discussed earlier, EarthWorks will focus on model configurations in which all components are on identical grids. Together with CMEPS, this will allow *EarthWorks* to achieve a lower operation count and significantly less communication overhead relative to models with more diverse grids. It will also simplify the logical structure of the model, and make it easier to modify going forward.

- **Smart, Modular Physics:** It is important that a flexible, modular approach be taken to allow experimentation with different physics software suites interacting with dynamical cores. To ensure adequate performance, however, the model software needs to be smart enough to avoid unnecessary data motion between CPU and GPU. EarthWorks will leverage the modular physics/dynamics framework, CCPP. For performance reasons, the CCPP framework will need to be generalized to support connecting host models to physics, where either may be located on CPUs or GPUs.. Currently there is no plan to create GPU versions of CCPP, so this work may fall to EarthWorks.

d) MPAS-A Background

MPAS-A consists of a set of nonlinear partial differential equations that govern the motion of the atmosphere - termed the dynamical core, combined with a set of forcing terms – typically physics parameterizations which describe the sub-grid-scale phenomenology driving the atmosphere. In terms of the number of lines of code, the software complexity of MPAS-A is fairly typical of other similar atmospheric models used in operational NWP. In terms of source lines of code (SLOC) for its various subcomponents, the totals are: dynamics (12 kSLOC), radiative transport (37 kSLOC), NOAH land surface model (21 kSLOC), and other physics (42 kSLOC)), totalling of 112 kSLOC.

e) The MPAS-A GPU Implementation.

The project to create a GPU-enabled version of MPAS-A sought to achieve good performance on GPUs without compromising CPU performance, while retaining the readability and maintainability of the source. The implementation of the MPAS-A execution model began with MPAS's finite-volume dynamics. Numerically, MPAS's dynamical core is fairly similar to NOAA's Nonhydrostatic Icosahedral Model (NIM), a global meteorological model targeted at operating with convection-permitting grids (Govett, 2017). NIM is important because of the pioneering work on performance portability of the model on CPUs and GPUs performed by Mark Govett and his team (Govett, 2017). NIM was ported across Intel Xeon, Xeon Phi, and NVIDIA GPUs using an in-house tool called F2C-ACC. Govett, (2017) found that it ran 2.0 and 2.5 times faster on the Intel Xeon Phi and NVIDIA K80 GPU systems, respectively, compared to contemporaneous, dual-socket Intel Xeon v3 "Haswell" systems.

After reviewing the available performance portability approaches, the OpenACC directive-based parallelization language was selected as being closest to matching the project objectives. Using standard OpenACC directives in the commercial PGI compiler, the MPAS-A team found that the moist dynamical core achieves approximately 2.5 to 3 times the performance of the 18-core dual-socket Xeon Broadwell node, using later generation NVIDIA Volta architectures, a result quite comparable to those obtained by Govett (2017) using F2C-ACC. See the 10 km CPU and GPU results in Figure 1 (right).

For maintainability and readability, by using OpenACC directives, the project avoided duplicating source or toggling architecture-specific code in and out with "ifdef" statements. In the end, after porting about 54 kSLOC of MPAS-A, the source size had grown by only 5-10%.

f) MPAS-A Performance

The performance data for MPAS-A was collected on the Summit supercomputer at the ORNL Leadership Class Facility (OLCF) using version 19 PGI compiler. Summit's AC922 nodes consist of six, NVIDIA Volta GPUs, each with 16 GB of onboard memory and two IBM Power 9 processors. The reference performance was obtained using the Intel 19 compiler running on NCAR Cheyenne supercomputer, which has nodes with dual 18-core Intel Xeon v4 Broadwell processors. The MPAS-A model test case configuration consisted of single precision, 56 levels, with 6 moist tracers.

The MPAS-A project also implemented physics-based task parallelism as a way to further accelerate the model on heterogeneous CPU-GPU architectures. Leveraging an approach first proposed and evaluated by scientists at Geophysical Fluid Dynamics Laboratory (GFDL), in which the radiative transport (RT) calculations can be run concurrently with the dynamics and other physics (Balaji, et al, 2016), this strategy had the dual advantages of reducing the amount of code to port to GPUs, and speeding up model throughput. The CPU-based RT domains coincide with the GPU domains, ensuring that the communication between the GPUs-based portion of the model and the RT scheme is confined on node.

g) Scaling Existing Performance to EarthWorks Goals

Referring to the right panel in Figure 1, we see the 3 km, 56 level model obtains 6.55 days/hour on 700 nodes of Summit. For comparison, NOAA established a forecast model throughput goal of approximately 7 days/hour in the benchmark competition for the Next Generation Global Prediction System (NGGPS). This threshold was tested for global model dynamical cores running at 13 km and 3 km resolution and 128 levels (Michalakes, 2015). Scaling the observed MPAS-A performance to 3 km to 4 km (1.33 times speedup) and from 56 to 100 levels (1.79 times slowdown) we project the current moist dynamical core would run in SP, with 100 levels on 400 nodes of Summit at 0.32 SYPD. Combining further optimizations with edging further on the scaling curve, we estimate that 0.5 SYPD throughput could be achieved. Adding in the cost of physics, excluding lagged radiation, (a 1.5 times slowdown), indicates a MPAS-A throughput of 0.35 SYPD for the full MPAS-A model today, using Summit's AC922 nodes. Thus, a minimum 5.7x increase in sustained hardware throughput is required to achieve 2 SYPD. Looking at the improvement of NVIDIA Tesla GPU sustained performance generation to generation from P100 to V100 architectures, it is reasonable to expect about a factor of 5-6x throughput improvement in NVIDIA products during this project's lifespan. With renewed competition in the GPU architectural space from other vendors, the throughput improvement may even be greater.

h) MPAS ocean performance

The MPAS-Ocean model development team at Los Alamos National Laboratory (LANL) have provided baseline CPU-based throughput data taken for a G-case (MPAS Ocean-Ice model forced by data atmosphere) using version 1 of the code. This case is a high-res configuration with a variable (18 km-to-6 km) resolution grid. This grid has 3.7 million horizontal cells and 80 vert layers, runs at about 1.05 SYPD on 4830 IBM Power-9 cores. This test case has about 10 times fewer horizontal points than a global, uniform 4 km grid would have (36.8 million cells) that we are interested in in EarthWorks.

The GPU throughput can be estimated for a uniform 4 km grid by weak scaling the horizontal points and applying the CFL scaling of the timestep by 2/3. This yields throughput of 0.7 SYPD on 48,200 Power-9 cores (over 1100 AC922 nodes) for this case. For comparison, recall that the MPAS-A model with 56 levels is expected to achieve about 0.25 years per day on 400 nodes of Summit. Scaling the CPU MPAS-O/I CPU throughput to 400 nodes, we find the CPU version of MPAS-O/I could just keep up the GPU-based MPAS-A, or even be faster if more levels are added to MPAS-A. This throughput disparity is to be expected because the atmosphere's time step is five times smaller than the ocean's at 4 km (24 sec vs 120 sec). A first step in integration of the Ocean/Ice/Atmosphere system should involve testing this hybrid GPU-atmosphere – CPU ocean configuration.

i) GPU execution of EarthWorks

Obviously, converting to a fully GPU-based EarthWorks depends on its achieving a breakeven speed up over the CPU version. For Summit, we estimate this to occur for seven CPU cores per GPU. Very preliminary benchmark data from the partial port of MPAS-O to Summit nodes indicates that this breakeven has already been achieved. Correspondence with the team leader suggest that substantially better GPU performance can be expected when the port is complete, and host-device transfer overheads are eliminated. Regardless, we assume that for EarthWorks, comparatively fewer GPU-resources will need to be devoted to the MPAS Ocean/ Ice model to match the atmospheric integration rate.

Finally, because the ocean and atmosphere grids coincide, we expect to collocate GPU-based ranks for each component on the same compute nodes. Whether the components share individual GPU resources or are distributed on separate GPU devices is a subject of experimentation. Another issue that will need to be addressed is developing a load-balancing the overall model that is portable across multiple systems. Obviously not all cells have ocean or sea ice present, and not all systems will have the same GPUs per node. We will experiment with both static domain decomposition, task placement, and dynamic load balancing on the GPUs to optimized overall model throughput.

We suspect that all or a large portion of EarthWorks would work just fine in single precision, which would speed the model up and also cut memory and I/O requirements. This possibility will be explored as part of the proposed research.

j) Ocean/Ice Mesh generation

The Los Alamos ocean modeling team has agreed to feed the quasi-uniform atmospheric mesh through their cell culling (land removal) tools and topography generator create a compatible mesh for use by EarthWorks. This will greatly simplify the process of generating the grid that will be used by the ocean model, including the accompanying bottom topography data set.

k) I/O performance

MPAS-A has demonstrated I/O with a 3 km grid, running on up to 700 nodes of Summit using PIO2 and pNETCDF. Modifications to help PIO2 operate high resolutions have been fed back into the PIO2 developer at NCAR. As part of ongoing SIMA work on CAM I/O should be ready by the project's start.

l) Summary of how it will work

Our vision for how EarthWorks will execute is shown in Figure 2. The strawman exascale architecture (center) consists of a collection of heterogeneous CPU/GPU nodes (top), and a storage hierarchy consisting of (middle) solid-state devices and (bottom) rotating disk. The ESM components are distributed across CPU (land model, etc…) and GPU (atmosphere, ocean and sea ice, the latter two executing sequentially. Radiative transport, atmosphere and ocean/sea ice execute concurrently on GPU. Asynchronous I/O is performed by dedicated CPU ranks. Model output lands first on SSD storage, where it is either saved for immediate analysis or compressed for later study. Lossy compression has demonstrated 5:1 decrease in climate model output

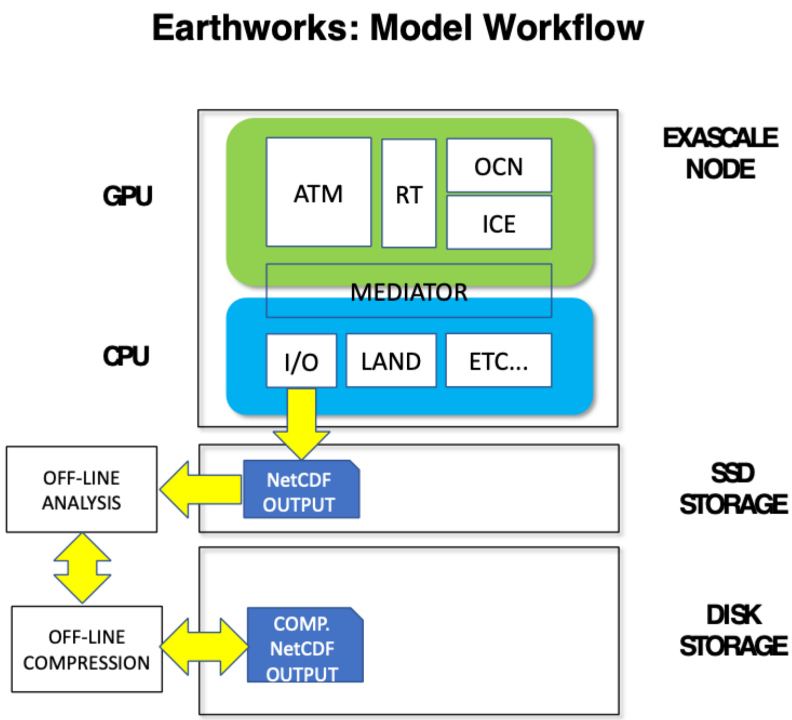

Figure 2. Coupled EarthWorks model task placement on a strawman exascale architecture. In its most general form ESM components run on both CPU and GPU components. The mediator (CMEPS) couples the components, which all share the same grid. The yellow arrows mark the data pathways from the model to analytics and compression utilities. Files are stored in machine portable NetCDF or compressed NetCDF formats.

(Baker, 2014). In situ analysis and compression are capabilities that can be added in the future. Parallel Python tools (PyReshaper, etc.) have demonstrated the ability to manipulate data in parallel (Paul, 2015), and parallel workflow system like PanGeo are developing into a more generalizable analytics environment (Hamman, 2018).

5. Goals and deliverables

a) Software engineering goals

As stated earlier, our software engineering goals are to create the documented EarthWorks model, and to demonstrate with 4 km or finder global grid spacing, and 1) 0.5 SYPD for atmosphere-only simulations with high vertical resolution, and 2) 1 SYPD for year for climate simulations with fewer layers.

b) Scientific goals

The early tests of EarthWorks will be carried out with coarse grid spacing, on the order of 100 km, as we work to confirm that the various components are working together as intended.

We will also test with coarse (~100 km) grid spacing in coupled (land-atmosphere-ocean) climate mode to ensure that the overall system with the same parameterizations and parameter settings reproduces observed global climate metrics. This testing strategy is efficient because it does not require long runs of a fully coupled high resolution model. It is similar to the strategy that has been very successfully used used by the European Center for Medium Range Weather Forecasting (ECMWF) to develop their extended-range weather forecasting system.

Our team includes participants from NCAR's weather, climate, and computational laboratories. We (and the broader community) will use EarthWorks for both weather and climate science. The second phase of testing will be based on the evaluation of ten-day weather forecasts. Weather forecasting is an excellent way to test atmospheric models that are intended for use in climate simulation (e.g., Palmer et al., 2008). The forecasts will be initialized with existing weather analyses from the operational centers, and the predicted weather will be compared to the observed weather. Forecast cases will be selected based on weather regimes that high-resolution global models are needed to simulate. These include tropical cyclone intensification and movement, severe weather events over the Great Plains of North America in summer, and extreme precipitation events associated with mountain ranges.

In free-running (non-forecast) mode, we will study the ability of the full-resolution version of EarthWorks to simulate the observed structure of the Inter-Tropical Convergence Zone and wave propagation from the troposphere into the stratosphere, in both the tropics and higher latitudes.

We will also compare the results from EarthWorks with results from other models by participating in DYAMOND or similar community-based intercomparison projects.

Finally, within the five-year horizon of the proposed project, we plan to conduct decadal simulations with EarthWorks to understand coupled behavior, land surface-weather interactions, and low-frequency (multi-year) modes of climate variability at high resolution.

c) What we will not do

To achieve our ambitious goals on the proposed five-year time scale, EarthWorks has to be highly focused. It is therefore important to identify things that EarthWorks will *not* try to do within its five-year horizon. MPAS for both the atmosphere and ocean can be used with variable resolution over the sphere, but because EarthWorks will have high resolution everywhere, we do not plan to make use of this capability. We do not plan to implement the land ice component of CESM. We do not plan to implement the river component of CESM. MPAS has been tested with data assimilation systems (e.g., Ha et al., 2017), but we do not plan to implement data assimilation in EarthWorks. Finally, we do not plan to perform simulations with intensive (and expensive) atmospheric chemistry. With respect to the five omissions listed above, we will avoid any modeling choices that would impede future efforts to implement these capabilities, each of which is important to a portion of the CESM community. Our collaboration with CESM software engineers will be critical in this regard.

d) Deliverables

EarthWorks has four major software deliverables. They are: EarthWorks coupled ESM software (i.e., the model code), PanGeo-based modeling scripts, lossy compression tools, and the machine learning emulator framework. There are two major dataset deliverables: initial datasets (simulated testcases) and data sets used to evaluate the model results. Further discussion of deliverables, delivery mechanisms, and related topics is given in the Supplemental Document on "Delivery Mechanism and Community Usage Metrics."

6. Summary of intellectual merit

The intellectual merit of our project lies in combining well matched model components with emerging trends in exascale computing to create a new and uniquely capable Earth System Model. The development of such a version of the community-based CESM is not planned in the foreseeable future. EarthWorks will allow us and the community to explore high-resolution storm resolving simulation capabilities in a CESM configuration using many standard CESM components. The move to storm-resolving scale is a major step forward for climate applications, eliminating the need for parameterizations of deep convection. This will be computing at the very edge of our capabilities. The resulting modeling system will be GPU/CPU capable, and hence usable on the latest generation of supercomputers.

7. Summary of broader impacts

The broader impacts of this project include the benefits to the national and international research community from free access to a well documented GRSM. In addition, we will train undergraduates, graduate students and postdoctoral fellows in the nuts and bolts of exascale programming and Earth System Model development. Specifically, we will leverage university partners at EPSCoR states and a long-established and successful computational science program (Summer Internships in Parallel Computational Science) at NCAR to build a pipeline of exascale model development talent. Finally, by opening up access to previously unresolved scales, the framework will advance the rapidly merging sciences of weather forecasting and climate simulation, both of which are of critical importance to society.

8. Summary with respect to additional review criteria

Finally, we briefly summarize how our proposal is responsive to each of the supplementary review criteria (beyond intellectual merit and broader impacts), as listed in the Program Solicitation.

- EarthWorks is *science-driven*. It is motivated by the scientific need to study problems at the weather/climate interface and improving predictions of weather and climate, including extreme events.
- EarthWorks is *innovative*: EarthWorks will provide simulation capabilities of the earth system that do not exist anywhere else.
- EarthWorks is based on *close collaborations among stakeholders*: EarthWorks will be developed by both university and national laboratory researchers.
- EarthWorks *builds on existing, recognized capabilities*. Specifically, we are building on the CESM, which is a major component of U.S. climate-modeling infrastructure.
- *Project plans, and system and process architecture*: EarthWorks will make it possible to deliver important new services to the climate and weather communities. The EarthWorks model, along with the required input data and user documentation, will be made available on GitHub. We will also propose to CESM management that the model be officially supported as a CESM configuration.
- *Deliverables* and *Metrics* are discussed in the Supplemental Document on "Delivery Mechanism and Community Usage Metrics."
- *Sustained and sustainable impacts*: EarthWorks and the simulations performed with it by our project team and the wider scientific community will lead to important practical advances in our ability to predict the weather and simulate the climate.

References

Bullock, J.O., Foroutan, H., Gilliam, R.C. and Herwehe, J.A., 2018. Adding four-dimensional data assimilation by analysis nudging to the Model for Prediction Across Scales–Atmosphere (version 4.0). *Geoscientific model development*, **11**, 2897-2922.

Danabasoglu, G., Lamarque, J. -F., Bachmeister, J., Bailey, D. A., DuVivier, A. K., Edwards, J., Emmons, L. K., Fasullo, J., Garcia, R., Gettelman, A., Hannay, C., Holland, M. M., Large, W. G., Lawrence, D. M., Lenaerts, J. T. M., Lindsay, K., Lipscomb, W. H., Mills, M. J., Neale, R., Oleson, K. W., Otto-Bliesner, B., Phillips, A. S., Sacks, W., Tilmes, S., van Kampenhout, L., Vertenstein, M., Bertini, A., Dennis, J., Deser, C., Fischer, C., Fox-Kember, B., Kay, J. E., Kinnison, D., Kushner, P. J., Long, M. C., Mickelson, S., Moore, J. K., Nienhouse, E., Polvani, L., Rasch, P. J., Strand, W. G. The Community Earth System Model version 2 (CESM2). Manuscript submitted for publication to *Journal of Advances in Modeling Earth Systems*. Manuscript available at http://www.cesm.ucar.edu/publications/papers/Danabasoglu_etal_JAMES2019.pdf.

Davis, C. A., D. A. Ahijevych, W. Wang and W. C. Skamarock, 2016: Evaluating Medium-Range Tropical Cyclone Forecasts in Uniform and Variable-Resolution Global Models, *Mon. Wea. Rev*., **144**, 4141-4160.

Eidhammer, T., H. Morrison, D. Mitchell, A. Gettelman, and E. Erfani, 2016: Improvements in Global Climate Model Microphysics Using a Consistent Representation of Ice Particle Properties. *J. Climate*, 609–29.

Gettelman, A., Morrison, H., Thayer‐Calder, K. and Zarzycki, C.M., 2019. The impact of rimed ice hydrometeors on global and regional climate. *Journal of advances in modeling earth systems*, **11**, 1543-1562.

Golaz, J.-C., V. E. Larson, and W. R. Cotton, 2002: A PDF-Based Model for Boundary Layer Clouds. Part I: Method and Model Description. *Journal of the Atmospheric Sciences,* **59**, 3540–51.

Golaz, J.C., Caldwell, P.M., Van Roekel, L.P., Petersen, M.R., Tang, Q., Wolfe, J.D., Abeshu, G., Anantharaj, V., Asay‐Davis, X.S., Bader, D.C. and Baldwin, S.A., 2019. The DOE E3SM coupled model version 1: Overview and evaluation at standard resolution. *Journal of Advances in Modeling Earth Systems*, **11**, 2089-2129.

Ha, S., C. Snyder, W. C. Skamarock, J. Anderson, and N. Collins, 2017: Ensemble Kalman Filter Data Assimilation for the Model for Prediction Across Scales (MPAS). *Mon. Wea. Rev*., **145**, 4673-4692.

Heikes, R. P., D. A. Randall, and C. S. Konor, 2013: Optimized icosahedral grids: Performance of finite-difference operators and multigrid solver. *Mon. Wea. Rev.*, **141**, 4450-4469.

Heinzeller, D., Duda, M.G. and Kunstmann, H., 2016. Towards convection-resolving, global atmospheric simulations with the Model for Prediction Across Scales (MPAS) v3. 1: An extreme scaling experiment. *Geosci. Model Dev*., **9**, 77–110

Hill, C., C. DeLuca, V. Balaji, M. Suarez, and A. da Silva, 2004: The architecture of the Earth System Modeling Framework. *IEEE Comput. Sci. Eng*., 6, 18–28, doi:https://doi.org/10.1109/MCISE.2004.1255817.

Hurrell, J. W., G. A. Meehl, D. Bader, T. Delworth, B. Kirtman, and B.Wielicki, 2009: A Unified Modeling Approach to Climate System Prediction. *Bulletin of the American Meteorological Society*, **90**, 1819-1832

Hurrell, J.W., Holland, M.M., Gent, P.R., Ghan, S., Kay, J.E., Kushner, P.J., Lamarque, J.F., Large, W.G., Lawrence, D., Lindsay, K. and Lipscomb, W.H., 2013. The community earth system model: a framework for collaborative research. *Bulletin of the American Meteorological Society*, **94**, 1339-1360.

Judt, F., 2018. Insights into atmospheric predictability through global convection-permitting model simulations. *Journal of the Atmospheric Sciences*, **75**, 1477-1497.

Kirtman B P, Min D, Infanti J M, Kinter J L, Paolino D A, Zhang Q, van den Dool H, Saha S, Mendez M P, Becker E, Peng P, Tripp P, Huang J, DeWitt D G, Tippett M K, Barnston A G, Li S, Rosati A, Schubert S D, Rienecker M, Suarez M, Li Z E, Marshak J, Lim Y-K, Tribbia J, Pegion K, Merryfield W J, Denis B and Wood E F 2014: *Bull. Am. Meteorol. Soc.*, **95**, 585– 601

Klemp, J. B., W. C. Skamarock, W. C., and S.-H. Park, 2015: Idealized Global Nonhydrostatic Atmospheric Test Cases on a Reduced Radius Sphere. *Journal of Advances in Modeling Earth Systems*. doi:10.1002/2015MS000435

Lawrence, B.N., Rezny, M., Budich, R., Bauer, P., Behrens, J., Carter, M., Deconinck, W., Ford, R., Maynard, C., Mullerworth, S. and Osuna, C., 2018. Crossing the chasm: how to develop weather and climate models for next generation computers? *Geosci. Model Dev.*, **11**, 1799–1821.

Lawrence, D. M. R. A. Fisher, C. D. Koven, K. W. Oleson, S. C. Swenson, G. Bonan, N. Collier, B. Ghimire, L. van Kampenhout, D. Kennedy, E. Kluzek, P. J. Lawrence, F. Li, H. Li, D. Lombardozzi, W. J. Riley, W. J. Sacks, M. Shi, M. Vertenstein, W.R. Wieder,, C. Xu, A.A. Ali, A.M. Badger, G. Bisht, M.A. Brunke, S.P. Burns,, J. Buzan, M. Clark, A. Craig, K. Dahlin, B. Drewniak, J.B. Fisher, M. Flanner, A.M. Fox, P. Gentine, F.Hoffman, G. Keppel-Aleks, R., Knox, S. Kumar, J. Lenaerts, L.R. Leung, W.H. Lipscomb, Y. Lu, A., Pandey, J.D. Pelletier, J. Perket,, J.T. Randerson, D.M. Ricciuto, B.M., Sanderson, A. Slater, Z.M. Subin, J. Tang, R.Q. Thomas, M. Val Martin, and X. Zeng, 2019. The Community Land Model version 5: Description of new features, benchmarking, and impact of forcing uncertainty. Accepted by *Journal of Advances in Modeling Earth Systems*..

Neale, R. B., J. Richter, S. Park, P. H. Lauritzen, S. J. Vavrus, P. J. Rasch, and M. Zhang, 2013: The mean climate of the Community Atmosphere Model (CAM4) in forced SST and fully coupled experiments. *J. Climate*, **26**, 5150–5168.

Niu, G.-Y., et al. , 2011. The community Noah land surface model with multiparameterization options (Noah-MP): 1. Model description and evaluation with local-scale measurements. *J. Geophys. Res.*, **116**, D12109, doi: 10.1029/2010JD015139.

Otto-Bliesner, B.L., E.C. Brady, J. Fasullo, A. Jahn, L. Landrum, S. Stevenson, N. Rosenbloom, A. Mai, and G. Strand, 2016: Climate Variability and Change since 850 CE: An Ensemble Approach with the Community Earth System Model. *Bull. Amer. Meteor. Soc.*, ***97***, 735–754.

Palmer, T.N., Doblas-Reyes, F.J., Weisheimer, A. and Rodwell, M.J., 2008. Toward seamless prediction: Calibration of climate change projections using seasonal forecasts. *Bulletin of the American Meteorological Society*, **89**, pp.459-470.

Park, S.H., Skamarock, W.C., Klemp, J.B., Fowler, L.D. and Duda, M.G., 2013. Evaluation of global atmospheric solvers using extensions of the Jablonowski and Williamson baroclinic wave test case. *Monthly Weather Review*, **141**, 3116-3129.

Park, S.-H., J. B. Klemp and W. C. Skamarock, 2014: A Comparison of Mesh Refinement in the Global MPAS-A and WRF Models Using an Idealized Normal-Mode Baroclinic Wave Simulation. *Mon. Wea. Rev.*, **142**, 3614-3634.

Petersen, M.R., Jacobsen, D.W., Ringler, T.D., Hecht, M.W. and Maltrud, M.E., 2015. Evaluation of the arbitrary Lagrangian–Eulerian vertical coordinate method in the MPAS-Ocean model. *Ocean Modelling*, **86**, 93-113.

Petersen, M.R., Asay−Davis, X.S., Berres, A.S., Chen, Q., Feige, N., Hoffman, M.J., Jacobsen, D.W., Jones, P.W., Maltrud, M.E., Price, S.F. and Ringler, T.D., 2019. An Evaluation of the Ocean and Sea Ice Climate of E3SM Using MPAS and Interannual CORE−II Forcing. *Journal of Advances in Modeling Earth Systems*.

Randall, D. A., C. M. Bitz, G. Danabasoglu, A. S. Denning, P. R. Gent, A. Gettelman, S. M. Griffies, P. Lynch, H. Morrison, R. Pincus, and J. Thuburn, 2018: 100 Years of Earth System Model Development. *Meteorological Monographs*, **59**, 12.1–12.66.

Reckinger, S.M., Petersen, M.R. and Reckinger, S.J., 2015. A study of overflow simulations using MPAS-Ocean: Vertical grids, resolution, and viscosity. *Ocean Modelling*, **96**, 291-313.

Ringler, T.D., Thuburn, J., Klemp, J.B. and Skamarock, W.C., 2010. A unified approach to energy conservation and potential vorticity dynamics for arbitrarily-structured C-grids. *Journal of Computational Physics*, **229**, 3065-3090.

Ringler, T., Petersen, M., Higdon, R.L., Jacobsen, D., Jones, P.W. and Maltrud, M., 2013. A multi-resolution approach to global ocean modeling. *Ocean Modelling*, **69**, 211-232.

Sandbach, S., Thuburn, J., Vassilev, D. and Duda, M.G., 2015. A semi-implicit version of the MPAS-Atmosphere dynamical core. *Monthly Weather Review*, **143**, 3838-3855

Sandgathe, S., O'Connor, W., Lett, N., McCarren, D. and Toepfer, F., 2011. National Unified Operational Prediction Capability Initiative. *Bulletin of the American Meteorological Society*, *92*(10), 1347-1351.

Satoh, M., Tomita, H, Miura, H., Iga, S., and Nasuno, T. , 2005: Development of a global cloud resolving model -- a multi-scale structure of tropical convections. *J. Earth Simulator*, **3**, 11-19.

Skamarock, W.C., Klemp, J.B., Duda, M.G., Fowler, L.D., Park, S.H. and Ringler, T.D., 2012. A multiscale nonhydrostatic atmospheric model using centroidal Voronoi tesselations and C-grid staggering. *Monthly Weather Review*, **140**, 3090-3105.

Skamarock, W. C., C. Snyder, J. B. Klemp and S. Park, 2019: Vertical resolution requirements in atmospheric simulation. *Mon. Wea. Rev.*, **147**, 2641-2656.

Skamarock, W. C., M. Barth, A. Gettleman, and H. Liu, 2019: System for Integrated Modling of the Atmosphere (SIMA). Paper presented at the *2019 CESM Workshop* in Boulder, Colorado. Available at http://www.cesm.ucar.edu/events/workshops/ws.2019/presentations/cross/skamarock.pdf.

Sinkovits, R.S. and M.G. Duda, 2016, July. Optimization and parallel load balancing of the MPAS Atmosphere Weather and Climate Code. In *Proceedings of the XSEDE16 Conference on Diversity, Big Data, and Science at Scale* (p. 5). ACM.Vancouver

Stevens, B. and S. Bony, 2013: What are climate models missing?. *Science*, **340**, 1053-1054.

Stevens, B., Satoh, M., Auger, L., Biercamp, J., Bretherton, C.S., Chen, X., Dueben, P., Judt, F., Khairoutdinov, M., Klocke, D. and Kodama, C., 2019. DYAMOND: the DYnamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains. Submitted to the *Journal of the Meteorological Society of Japan*.

Theurich, G., DeLuca, C., Campbell, T., Liu, F., Saint, K., Vertenstein, M., Chen, J., Oehmke, R., Doyle, J., Whitcomb, T. and Wallcraft, A., 2016. The earth system prediction suite: toward a coordinated US modeling capability. *Bulletin of the American Meteorological Society*, *IB*, 1229-1247.

Thompson, David W. J., Mark P. Baldwin, and John M. Wallace. “Stratospheric Connection to Northern Hemisphere Wintertime Weather: Implications for Prediction.” *J. Climate,* **15**, 1421–28.

Thompson, Gregory, and Trude Eidhammer. "A Study of Aerosol Impacts on Clouds and Precipitation Development in a Large Winter Cyclone." *Journal of the Atmospheric Sciences,* **71**, 3636–58.

Thuburn, J., Ringler, T.D., Skamarock, W.C. and Klemp, J.B., 2009. Numerical representation of geostrophic modes on arbitrarily structured C-grids. *Journal of Computational Physics*, **228**, 8321-8335.

Tomita, H, Miura, H., Iga, S., Nasuno, T., and Satoh,M. (2005) : A global cloud-resolving simulation: preliminary results from an aqua planet experiment. *Geophys. Res. Lett.*, **32**, L08805, doi:10.1029/2005GL022459.

Turner, A.K., 2018. *MPAS-Seaice Model User's Guide* (No. LA-UR-18-23512). Los Alamos National Lab.(LANL), Los Alamos, NM (United States).

Wolfram Jr, P.J., Brus, S.R. and Engwirda, D., 2019. *MPAS-O mesh characteristics, generation, and regional refinement* (No. LA-UR-19-22631). Los Alamos National Lab.(LANL), Los Alamos, NM (United States).

Wong, M. and W. C. Skamarock, 2016: Spectral characteristics of convective-scale precipitation observations and forecasts, *Mon. Wea. Rev.*, **144**, 4183-4195. doi: 10.1175/MWR-D-16-0183.1

Supplemental Documents

Supplemental Document: Delivery Mechanism and Community Usage Metrics

(2 page limit)

1. Deliverables

There are four major software deliverables from the EarthWorks project. Listed in Table 1, they are: EarthWorks coupled ESM software, PanGeo-based modeling scripts, lossy compression tools, and the machine learning emulator framework. There are two major dataset deliverables: initial datasets (simulated testcases) and verification data sets (validation data).

2. Distribution for hi-end, critical Earth System research

EarthWorks is laser-focused on providing a high-resolution, storm-resolving, coupled ESM capability to the weather and climate community, consisting of two simulated years per wall-clock day (SYPD). This focus is driven by the intellectual merit of our proposal: addressing fundamental science questions, such as the persistent statistical biases observed in current generation of ESMs, particularly related to precipitation and clouds. Current models run on CPU systems at low-resolutions (25-100 km) too coarse to resolve individual storms. The science objective of our proposal suggests that the initial users of EarthWorks will be a small set of elite users of exascale Leadership Class Facilities, such as those operated by the DOE, and others. For this reason, we do not expect large numbers of downloads in our early releases for EarthWorks.

3. Distribution of software and tools through community modeling

In order to have the broadest long-term impact on the ES research community, EarthWorks modeling software will strive to be compatible, and will hopefully be included in, future releases of the Community Earth System Model (CESM). This will enable EarthWorks to leverage the long-standing distribution, training and support mechanisms which have been in place, and benefit a much wider audience. CESM is currently distributed via GitHub.com at ESSCOMP/CESM. In addition to the model source, the initial and verification datasets will be distributed to assist science teams with their science objectives. We expect these datasets, in total to be less than 10 terabytes. Table 1 below shows how the deliverables will be distributed and preserved.

Mode / Deliverable	Pangeo.io	GitHub	Globus endpoint	Archive Backup
Modeling software		Download, open development		X
initial and verification datasets			Available for download	XX (<10 TB)
PanGeo analysis scripts	Example scripts @ pangeo.io	Download, open development		X
Lossy compression tools		Download, open development		X
ML emulator framework		Download, open development		X

Table 1. An enumeration of deliverable EarthWorks components and how they will be distributed and preserved.

4. Metrics

a) *Timeline*

We present below the release timeline for the project software, what success looks like, our approach to collecting metrics (how we will track downloads and usage).

Years 1-2:	Software development
End of Year 2:	Release of individual component models running with CMEPS.
End of Year 3:	The first full release of the full system (coupled model and related toolchain) on GitHub
End of Year 4:	Second release of full system software on GitHub
End of Year 5/project:	Third release of full system software on GitHub

b) *Success:*

Success at the end of the project looks like this:

- Achieving our science-driven ESM throughput target (2 SYPD).
- At least five ESM science teams, besides the PI and this team, using the modeling suite
- At least ten ESM teams around the world using the PanGeo analysis, lossy compression, and machine learning tools developed by EarthWorks, and substantial evidence of open development (pull-requests, etc.).
- Integration of EarthWorks software into CESM.

c) *Tracking Downloads*

The project will use GitHub for all of its software development projects. GitHub allows contributors to collect metadata about their releases via its API, which will enable the team to compile download statistics.

d) *Tracking Usage*

The project's major digital artifacts will be assigned Data Object Identifiers (DOIs), thus enabling EarthWorks to track through citations, the downstream usage of the models, data and tools. We will follow DOI best practices for datasets, namely: we will assign DOIs for each version of datasets to track any changes; we will include the DOI in the metadata record that describes the data (e.g. in the NetCDF header file); we include the DOI in the citation of the data set. We will use standard DOI tools labeling our data and for correcting bad links.

Supplemental Document: Management and Coordination Plan

1. Roles of personnel

Dr. David Randall is University Distinguished Professor of Atmospheric Science at Colorado State University. He has been working with global atmospheric models continuously since 1972. He has published over 230 refereed journal articles. His research deals with the roles of clouds and turbulence in the climate system, and also on numerical methods for simulation of the atmosphere. From 2006 to 2016, he served as Director of a National Science Foundation Science and Technology Center for Multiscale Modeling of Atmospheric Projects. He will be responsible for overall direction of the *EarthWorks* project, including meetings of the investigator team and meetings of the External Advisory Panel. He will supervise the postdoctoral fellow and two graduate students. He will take the lead on at least some of the publications produced by the project.

Dr. James W. Hurrell is an expert on large-scale climate variability and the evaluation of Earth system models. He holds the Scott Presidential Chair in Environmental Science and Engineering at Colorado State University and is a Professor in the Department of Atmospheric Science. His research has centered on empirical and modeling studies and diagnostic analyses to better understand climate, climate variability, and climate predictability. Dr. Hurrell is the former Director of NCAR and a former Chief Scientist of the Community Earth System Model (CESM). In the latter role, he was responsible for the overall coordination of model development activities, open releases of the model, and the production of CESM simulations for community research and contributions to national and international assessments of climate variability and change. Dr. Hurrell will be responsible for overall coordination with NCAR on the EarthWorks project. He will also serve as a lead in the analysis of simulations to test, evaluate and analyze EarthWorks.

Dr. Andrew Gettelman is an expert in cloud physics and a principal developer of CAM, as well as the climate lead for the SIMA project at NCAR. Dr. Gettelman will manage software engineering support for the atmosphere model at NCAR related to assembly and construction of the EarthWorks atmosphere with MPAS and the physical parameterization suites. This will include managing software engineering in conjunction with Vertenstein, and also co-supervising an Associate Scientist (with Skamarock) doing simulations and analysis of different options for EarthWorks physical parameterizations running with the MPAS atmosphere under the SIMA infrastructure.

Dr. Richard Loft is the Director of Technology Development in the Computational and Information Systems Laboratory (CISL). He has over twenty years of experience in the design, implementation, and optimization parallel atmospheric models. Dr. Loft has led the three-year collaborative effort to port MPAS-A to GPUs, and for the past two years, has led the Analytics and Integrative Machine Learning (AIML) group in CISL. He will manage three aspects of the EarthWorks project: 1) the software engineering required to port EarthWorks model components to GPUs; 2) the HPC performance testing and optimization of the EarthWorks modeling system; and 3) the development of machine-learned emulator framework (MLEF) to facilitate incorporating physics emulators; 4) the development and deployment of PanGeo analytics scirpts; and 5) the development of off-line lossy compression tools.

Dr. William Skamarock is the lead scientist for MPAS-Atmosphere and is one of the main developers of MPAS since its inception. His expertise is in development of fluid-flow solvers for atmospheric models and in atmospheric dynamics, specifically atmospheric convection. Dr. Skamarock will manage NCAR's Mesoscale and Microscale Meteorology Laboratory's participation in EarthWorks, and he will be working on the scientific configurations of MPAS-Atmosphere in CESM, including testing the fluid-flow solver and the physical parameterizations in the EarthWorks configurations. He will participate in the scientific evaluation and analysis of the model results and oversee the MMM-based effort in this area. Dr. Skamarock will also assist in managing and coordinating the software engineering efforts needed in EarthWorks that reside in MMM, particularly that associated with the parallelization and optimization of MPAS-Atmosphere and its CESM/CIME interfaces.

Postdoctoral fellow will participate in the optimization of the model for use with GPUs. This work will be performed in collaboration with Dr. Richard Loft and others at NCAR.

The three CSU graduate students will participate in the model development and analysis of the model results. This work will be performed in collaboration with senior project personnel.

A software engineer in the Climate and Global Dynamics Laboratory at NCAR will provide expertise in GPUs, CCPP and other CESM infrastructure. An associate scientist in the same lab will provide testing support for the atmosphere.

The four NCAR summer student assistantships (2 graduate and 2 undergraduate) will work on GPU-porting, optimization, and HPC performance benchmarking projects during their 11-week visit. There are two sources for these summer students: 1) qualified summer interns admitted into NCAR's 11-week Summer Internships in Parallel Computational Science (SIParCS) program to work on EarthWorks projects; 2) student hired as summer visitors to work on projects coordinated with partner universities.

2. Project management

The project will be organized into three project teams: a *Science Team* (lead by Hurrell); a *Software Engineering Team* (lead by.Gettelman); and a *Management Team* (lead by Randall).

EarthWorks will set up an External Advisory Panel that will review our progress every six months in parallel with the all-hands project meetings. External Advisory Panel members will include a representative from the CESM Software Engineering Group, and also a representative from the Ocean Modeling Group at the Los Alamos National Laboratory. Drs. Mariana Vertenstein of NCAR and Philip Jones of LANL have agreed to this arrangement, as shown in their Letters of Collaboration.

EarthWorks will use software project management tools; the Software Engineering Team will use bug tracking tools and version control systems (GitHub) in accordance with software engineering best practices. The final selection of specific tools will be made at project inception.

3. Coordination mechanisms

All-hands project meetings (including personnel from CSU and NCAR) will occur at least every six months. The External Advisory Panel meetings will be held in concert with the all-hands meetings.

Dissemination Mechanisms and Community Usage Metrics are discussed in detail in the eponymous Supplemental Document in this proposal.

4. Graduate and Undergraduate Student Training

EarthWorks will support three CSU graduate students and one CSU postdoctoral fellow. All four of them will receive hands-on training in Earth System Model development. The students will also receive intensive classroom training that will provide a strong foundation for future research careers. There has been a chronic shortage of young scientists with such training, so this is an important contribution to the welfare of our field.

In addition, EarthWorks will support participation in EarthWorksForce by graduate students at the University of Delaware, the University of Wyoming, and Boise State University. EarthWorksForce will provide students with employable HPC skills and fill a growing need for experts in optimization and porting of scientific applications to GPUs. We propose to grow the next-generation workforce at partner universities by training computer science and electrical and computer engineering students in the porting and optimization of GPU-enabled model components required by the EarthWorks project. This aspect of EarthWorks, called *EarthWorksForce*, will provide students with employable HPC skills and fill a growing need for experts in optimization and porting of scientific applications to GPUs. EarthWorksForce works in the following way. NCAR's role will be to identify qualified partner university faculty as well as the specific tasks for students to work on. Faculty partners will develop/expand student GPU/HPC training pipelines as part of the curriculum at their respective institutions and oversee the training and work of the students. Students that are entrained in these pipelines can be graduate or undergraduate level, and the work is performed as part of a paid research position and typically provides partial credit towards their degree. The proposed engagement model has been prototyped successfully in collaborations with faculty at the University of Wyoming and the University of Delaware over the past three years during the port of the MPAS-A model to GPUs. To date a total of seven university students have been trained in this way, and five have obtained HPC-related employment in industry or at U.S. national laboratories as a result of their participation. The work performed by these students has been of an excellent quality and has helped NCAR provide GPU-CPU portable models, which benefits the entire scientific user community. For EarthWorks, we have obtained three university faculty letters of collaboration: Professor Suresh Muknahallipatna at the University of Wyoming, Professor Sunita Chandrasekaran, at the University of Delaware, and Professor Catherine Olschanowsky at Boise State University. We note that all three universities reside in Experimental Program to Stimulate Competitive Research (EPSCoR) states – thus a broader impact will be to advance the research competitiveness of these states by providing such training.

5. Budget mapping to tasks

a) CSU

CSU will work on assembling, validating, testing and documenting EarthWorks – a new, global storm-resolving model – and make it available for community use as open source, via GitHub. This work will be completed in close collaboration with the NCAR co-PIs and the partner organizations.

- The CSU PI will be responsible for the overall direction of the EarthWorks project, including meetings of the investigator team and meetings of the External Advisory Panel. He will supervise the postdoctoral fellow and two graduate students.
- The CSU co-PI will be responsible for overall coordination with NCAR, and he will lead the analysis of simulations to test, evaluate and analyze the EarthWorks model. He will supervise one graduate student and one research scientist.

- The CSU postdoctoral fellow will participate in the optimization of the model for use with GPUs.
- The three CSU graduate students will participate in the development of the model and analysis of the model results, in collaboration with senior project personnel.

CSU has budgeted for travel of the CSU team to attend and present at one domestic conference per year and present results from EarthWorks. Travel funds are also budgeted for two CSU team members to attend and present at one international conference per year, and for the PI to attend the required NSF PI meeting each year.

b) NCAR

NCAR will work on assembling an EarthWorks Atmospheric component based on MPAS dynamical core, CCPP physics and a selected suite of Physical Parameterizations from WRF and/or CESM/CAM. An associate scientist will help CSU with testing and evaluation. A Software engineer in CGD will help with assembling the atmospheric component in the coupled CESM/CIME system, and working to ensure that the infrastructure for this atmosphere (CAM/SIMA based) are GPU compatible.

- CISL will take the lead with a software engineer and student assistants in porting critical physical parameterizations to GPUs, and advising CGD and MMM on GPU porting issues and assembly of the model.
- MMM will assist CSU and CGD with evaluation of MPAS dynamical core and WRF and CESM physics in EarthWorks model (MMM/CGD AS1).
- CGD will assist CISL with GPU ports of CAM Physics (CGD/MMM-SE2) years 1-3 and assist CSU and help test EarthWorks atmosphere model (MPAS + CCPP + physics suite); resources: 1 FTE AS1 for years 1-5 For AS1; 0.5 FTE SE2 for years 1-5.
- CISL will port Key CAM physical parameterizations to GPUs (Years 1-3); resources: CISL SE2 (0.3 FTE will assist CSU CGD to optimize & Integrate EarthWorks Atmosphere (& Ocean); resources: CISL SE 2 0.15 FTE in years 4-5. Additional FTE input: 2 undergraduate and 2 graduate students with 11 weeks summer support over years 1-5.

NCAR has budgeted for travel based on 2 person-trips to domestic conferences and 1 person trip to a foreign conference each year, over the course of the project.

A five-year total budget of $308,000 will provide summer support for EarthWorkForce development projects.. NCAR (CO-PI Loft) will administer these funds to allow flexibility in developing and assigning student projects.

Supplemental Document: Letters of Collaboration

Letter from Dr. Mariana Vertenstein

Dr. Mariana Vertenstein is the Head of the CESM Software Engineering Group at the National Center for Atmospheric Research.

Letter from Dr. Philip Jones

Dr. Philip Jones is the Project Leader of the Climate Ocean and Sea Ice Modeling (COSIM) Group at the Los Alamos National Laboratory.

Letter from Prof. Sunita Chandrasekaran

Professor Sunita Chandrasekaran is at the University of Delaware and will participate in the EarthWorksForce initiative.

Letter from Prof. Catherine Olschanowsky

Professor Catherine Olschanowsky is at Boise State University and will participate in the EarthWorksForce initiative.

Letter from Prof. Suresh Muknahallipatna

Professor Suresh Muknahallipatna at the University of Wyoming is at the University of Wyoming and will participate in the EarthWorksForce initiative.

Supplemental Document: Project Personnel and Partner Organizations

1. David Randall, Colorado State University, PI
2. James Hurrell, Colorado State University, Co-PI
3. Andrew Gettelman, National Center for Atmospheric Research, Co-PI
4. Richard Loft, National Center for Atmospheric Research, Co-PI
5. William Skamarock, National Center for Atmospheric Research, Co-PI
6. Lantao Sun, Colorado State University, Research Scientist
7. Suresh Muknahallipatna; University of Wyoming, Professor; Unfunded collaborator and author of letter of collaboration
8. Sunita Chandrasekaran; University of Delaware; Professor; Unfunded collaborator and author of letter of collaboration
9. Catherine Olschanowsky; Boise State University; Professor; Unfunded collaborator and author of letter of collaboration
10. Marianna Vertenstein, National Center for Atmospheric Research, Unfunded collaborator and author of letter of collaboration and External Advisory Panel Member representing CESM software engineering
11. Phillip Jones, Los Alamos National Laboratory, Unfunded collaborator and author of letter of collaboration and External Advisory Panel Member representing MPAS ocean
12. TBD, Colorado State University, Postdoctoral fellow
13. TBD, Colorado State University, M.S. Graduate student
14. TBD, Colorado State University, M.S. Graduate student
15. TBD, Colorado State University, M.S. Graduate student
16. TBD; National Center for Atmospheric Research; Software Engineer/Programmer II
17. TBD; National Center for Atmospheric Research; Software Engineer/Programmer II
18. TBD; National Center for Atmospheric Research; Associate Scientist I
19. TBD; National Center for Atmospheric Research; Student Assistant III
20. TBD; National Center for Atmospheric Research; Student Assistant III
21. TBD; National Center for Atmospheric Research; Student Assistant II
22. TBD; National Center for Atmospheric Research; Student Assistant II

Supplemental Document: Data Management Plan

1. Types of Data Produced and Managed

EarthWorks proposal is primarily a software development project. Thus, the scope of this Data Management Plan covers the documentation, configuration tools, source code, build scripts, verification tools and initial and verification datasets required to configure, compile, execute, and verify EarthWorks results. The volume of EarthWorks project software and data is expected, in total, to be on the order of 1-10 terabytes. Storing this data volume will be a small, incremental cost to the project.

The project team recognizes that the data volumes from EarthWorks modeling experiments could be extreme (order 10s of petabytes). *The management of these experimental data is outside the scope of this plan*, and will be detailed in the supporting allocation requests to various large-scale computing facilities where high-resolution experiments will occur. In the course of such experiments, EarthWorks will leverage or create additional software elements, to cope with the data volumes, for example: lossy compression tools and methodologies (Baker, 2014; Baker, 2018), in situ analysis capabilities, and parallel off-line analysis tools. *These software elements will be included in this plan.*

2. File Formats and MetaData

Model input and output files will be stored in community-standard Network Common Data Format (NetCDF) with descriptive metadata conforming to standard Climate and Forecast (CF) metadata conventions. Where compression is used, tools will be provided to compress files from, and uncompress files back into, NetCDF format.

3. Access, Sharing and Reuse

Our intent is to provide free and open access to the EarthWorks software and data for the use of the climate and weather research communities. Access to the project's software and will be through secure servers using software development platforms like GitHub and data access points such as NCAR's Globus GridFTP servers, the latter being subject to a NCAR request/approval process. There are no ethical and privacy issues with these data. To preserve IP rights EarthWorks software, supporting tools will be provided using terms of use (including provisions for re-use, re-distribution and production of derivatives) and disclaimers consistent with the provisions of the pre-existing software, such as CESM, MPAS-A/O/I, each of which, in turn, includes third party tools and libraries that, in turn, may have their own copyright notices and terms.

4. Archival

We have plans for archiving and for preservation of access to EarthWorks project data. The long-term strategy for maintaining, curating and archiving the data is via integration of EarthWorks with CESM and thus leveraging its regular backup of long-term storage at NCAR. Documentation with the software, tools and data will make it reusable for the foreseeable future. Archiving costs for the anticipated data volumes (1-10 TB) are incremental, small, and expected to decrease even further in the future.

Supplemental Document: Postdoctoral Researcher Mentoring Plan

The postdoctoral period is a critical stage in the career of an energetic young researcher, because it is during this time that she/he learns independence, knowledge and technical skills continue to rapidly increase, she/he may be starting the search for permanent positions such as faculty positions. Our mentoring plan for the postdoc will be as follows:

1) We will meet with the newly hired postdoc to conduct an “orientation.” In this meeting we can deal with the various practical problems of a new hire, and get to know each other better.
2) We will meet with the postdoc at least once a week to discuss research results and more general topics related to building a career in atmospheric science. The guiding hand played by a postdoctoral advisor should be less intrusive than that of a graduate advisor, given that a postdoctoral fellow is growing as an independent researcher. However, guidance at this stage is still critical since the postdoctoral researcher is still growing as a scientist.
3) The postdoc will be encouraged to attend at least two scientific conferences or workshops per year to both expose her/his ideas to the rest of the scientific community (to gain new collaborators and contacts) and to gain a broader perspective on her/his research endeavor.
4) We will encourage the postdoc to publish regularly. Unfortunately, the field now seems to be characterized by an escalation of publication rates with quality often sacrificed at the expense of quality. Maybe this is inevitable in an era of declining funding. We will encourage the postdoc to be a relatively prolific author out of necessity for her/his career, but without sacrificing the quality of her research in the process.
5) We will encourage the postdoc to participate in writing grant proposals. CSU’s Department of Atmospheric Science does not permit postdocs to serve as PIs on proposals.
6) We will provide career counseling. In particular, we will encourage the postdoc to start early in a faculty job search, assuming that that is a career goal. That first faculty interview is an eye-opening experience. Given today’s tight job market, it is important to apply for some positions that are not the “perfect” job, since they may turn out to be a great (and unexpected?) fit in the interview process, or/and they can serve as a stepping-stone to something better in the future. In addition, the postdoc will enroll in a 1-credit professional development class offered by CSU’s Department of Atmospheric Science.
7) The postdoc will be fully entrained into our ongoing weekly informal group meetings here at CSU. In these meetings, frequent presentations and questioning of the presenting speakers leads to finely honed presentation styles and an improved ability to answer difficult questions in a group environment. Experience shows that the postdoc will also benefit from excellent ideas contributed by the other meeting participants.
8) The success of this mentoring plan will be assessed by tracking the Postdoctoral Researcher’s progress toward his/her research and career goals.