CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model.
Under Review, IJHPCA 2011
Special Issue Devoted to the CESM Version 1.

July 31, 2011

J. Dennis, National Center for Atmospheric Research, Boulder, CO
J. Edwards, National Center for Atmospheric Research, Boulder, CO
Katherine J. Evans Oak Ridge National Laboratory, Oak Ridge, TN
O. Guba, Sandia National Laboratories, Albuquerque, NM
P.H. Lauritzen, National Center for Atmospheric Research, Boulder, CO
A.A. Mirin, Lawrence Livermore National Laboratory, Livermore, CA
A. St-Cyr, National Center for Atmospheric Research, Boulder, CO
M. A. Taylor, Sandia National Laboratories, Albuquerque, NM
P. H. Worley, Oak Ridge National Laboratory, Oak Ridge, TN

Running head: CAM-SE

Author addresses:

J. Dennis,
National Center for Atmospheric Research, Boulder Colorado, 80307-3000
email: dennis@ucar.edu, phone: 303-497-1809, fax: 303-497-1286

J. Edwards
National Center for Atmospheric Research, Boulder Colorado, 80307-3000
email: jedwards@ucar.edu

Katherine. J. Evans
Oak Ridge National Laboratory, PO Box 2008, Oak Ridge, TN 37831-6301
email: evanskj@ornl.gov, phone: 865-576-6517

O. Guba
Sandia National Laboratories, PO Box 5800, Albuquerque, NM 87185-0370.
Email: onguba@sandia.gov, Phone: 505-844-7414

P.H. Lauritzen
National Center for Atmospheric Research, Boulder Colorado, 80307-3000
email: pel@ucar.edu, phone: 303-497-1316, fax: 303-497-1314

A. A. Mirin
Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 94551
Email: mirin@llnl.gov, Phone: (925) 422-4020, Fax: (925) 423-2993

Amik St.-Cyr
National Center for Atmospheric Research, Boulder Colorado, 80307-3000
email: amik@ucar.edu, phone: 303-497-1287, fax: 303-497-1286

M. A. Taylor (corresponding author)
Sandia National Laboratories, PO Box 5800, Albuquerque, NM 87185-0370.
Email: mataylo@sandia.gov, Phone: 505-284-1874,

P. H. Worley
Oak Ridge National Laboratory, PO Box 2008, Oak Ridge, TN 37831-6173
email: worleyph@ornl.gov, phone: 865-574-3128

Abstract

The Community Atmosphere Model (CAM) version 5 includes a spectral element dynamical core option from NCAR's High-Order Method Modeling Environment. It is a continuous Galerkin spectral finite element method designed for fully unstructured quadrilateral meshes. The current configurations in CAM are based on the cubed-sphere grid. The main motivation for including a spectral element dynamical core is to improve the scalability of CAM by allowing quasi-uniform grids for the sphere that do not require polar filters. In addition, the approach provides other state-of-the-art capabilities such as improved conservation properties. Spectral elements are used for the horizontal discretization, while most other aspects of the dynamical core are a hybrid of well tested techniques from CAM's finite volume and global spectral dynamical core options. Here we first give a overview of the spectral element dynamical core as used in CAM. We then give scalability and performance results from CAM running with three different dynamical core options within the Community Earth System Model, using a pre-industrial time-slice configuration. We focus on high resolution simulations, using 1/4 degree, 1/8 degree, and T341 spectral truncation horizontal grids.

Keywords: spectral elements, dynamical core, parallel scalability, global circulation model, atmospheric modeling

1 Introduction

Today's petascale computers have hundreds of thousands of processor cores, and the next generation machines could have millions. As we no longer see much increase in single processor core performance, these machines are relying almost entirely on increasing performance through increased parallelism. Translating this to application performance is thus only possible with very scalable applications. Achieving the required level of scalability in modern climate models remains a challenge due to several scalability bottlenecks. The largest bottleneck in many of these models occurs in the dynamical core of the atmosphere component. The dynamical core solves the partial differential equations governing the fluid dynamical aspects of the atmosphere, but does not include the suite of subgrid parametrizations used for unresolved physical processes such as convection, precipitation and radiative forcings. The bottleneck is due to the commonly used latitude/longitude grid, which clusters grid points at the poles. Many numerical methods can successfully deal with this clustering but they result in reduced parallel scalability (Taylor et al., 2008). Thus there is a growing interest in highly scalable dynamical cores based on less structured or unstructured grids with more uniform resolution. Many of these approaches have been surveyed in (Williamson, 2007).

Here we report on the performance of a highly scalable configuration of the Community Earth System Model (CESM) version 1 using a spectral element dynamical core. The CESM is a state of the art climate model with atmosphere, ocean, land and ice component models[1]. The spectral element dynamical core is a new option within the Community Atmosphere Model (CAM), the CESM atmosphere component model (Neale et al., 2010). It comes from HOMME, the High-Order Method Modeling Environment (Dennis et al., 2005). HOMME's spectral element method (SEM) can naturally make use of fully unstructured quadrilateral grids. For quasi-uniform grids that do not cluster grid points at the poles, we use the cubed-sphere grid (Fig. 1) first used in Sadourny (1972). The standalone HOMME model is also used for research into other numerical methods, such as discontinuous-Galerkin (Nair, 2009) and adaptive mesh refinement (St.-Cyr et al., 2008). Here we refer to CAM with the spectral element dynamical core as CAM-SE.

CAM-SE is the first dynamical core in CAM that is capable of using a fully unstructured grid. This integration was made possible by the decoupling of dynamics and physics used in CAM's process split approach (Williamson, 2002), the extensive infrastructure work to support unstructured grids (Worley and Drake, 2005), and the CESM *tri-grid* capability (Craig et al., 2011). The SEM was chosen based on its demonstrated scalability when run as a standalone dry dynamical core (Dennis et al., 2005; Bhanot et al., 2008). In addition to scalability, the SEM is known to produce accurate solutions for atmospheric problems, as demonstrated first with shallow water test cases (Taylor et al., 1997; Thomas and Loft, 2002), three-dimensional dry dynamical test cases (Taylor et al., 1998; Thomas and Loft, 2005; Dennis et al., 2005; Taylor et al., 2007; Lauritzen et al., 2010), idealized multicloud simulations (Khouider et al., 2010), idealized aqua planet experiments which include full physics (Taylor et al., 2008; Mishra et al., 2011b,a) and realistic simulations with CAM2 physics (Wang et al., 2007). A tuned version of CAM-SE simulating the Earth's climate is evaluated in (Evans et al., 2011). The SE method has also been pursued for global forecast modeling in (Giraldo and Rosmond, 2004; Giraldo, 2005; Kim et al., 2008).

In what follows we will first give an overview of the CAM-SE dynamical core. We combine some well proven approaches with several new features added to CAM-SE in order to make it more suitable when running with CAM's full suite of atmospheric physics. We then describe CESM benchmark results using a realistic 1850s pre-industrial time-slice configuration. We present results using CAM-SE, finite volume (CAM-FV) and Eulerian global spectral (CAM-EUL) dynamical cores. CAM-EUL was the original dynamical core in the CCSM (the predecessor to the CESM) and used in the IPCC 4th assessment report (Intergovernmental Panel on Climate Change, 2007).

[1]http://www.cesm.ucar.edu

CAM-FV is the current default dynamical core in the CESM and is being used for the CESM's contributions to the IPCC 5th assessment report. CAM-EUL and CAM-FV are described in (Collins et al., 2004). The computational performance of CAM-FV and CAM-EUL have been extensively reported on, e.g. (Putman et al., 2005; Mirin and Sawyer, 2005; Worley and Drake, 2005; Worley et al., 2006; Mirin and Worley, 2011).

2 The CAM-SE dynamical core

CAM-SE uses a continuous Galerkin spectral finite element method (Taylor et al., 1997; Fournier et al., 2004; Thomas and Loft, 2005; Wang et al., 2007; Taylor and Fournier, 2010), here abbreviated to the spectral element method (SEM). It solves the atmospheric primitive equations. CAM-SE represents a large change in the horizontal grid as compared to the other dynamical cores in CAM, but most other aspects of CAM-SE are based on a combination of well-tested approaches from the Eulerian global spectral and finite volume dynamical cores. For tracer advection, CAM-SE is modeled closely on the FV core. It uses the same conservation form of the transport equation and the same vertically Lagrangian discretization (Lin, 2004). The CAM-SE dynamics are modeled closely on the CAM-EUL core. They share the same vertical coordinate, vertical discretization, hyper-viscosity based horizontal diffusion, top-of-model dissipation, and solve the same moist hydrostatic equations. The main differences are that CAM-SE uses the vector-invariant form of the momentum equation instead of the vorticity-divergence formulation. CAM-SE provides improved scaling in CAM due in part to its ability to efficently use a fixed two-dimensional domain decomposition and the fact that it does not require inherently load-imbalanced polar filtering. In addition to scaling, CAM-SE improves CAM's numerical conservation: it locally conserves both mass and moist total energy without the use of ad-hoc fixers (Taylor, 2011), and the horizontal advection operator locally conserves potential vorticity (PV).

2.1 Time Step Approach

We start by giving a high level overview of the CAM-SE time split approach, which decomposes each time step into components involving dynamics, tracer advection, forcing and dissipation. Dynamics and tracer advection are the processes governed by the atmospheric primitive equations. The forcing terms are computed by the CAM physics (the suite of subgrid physical parametrizations in CAM). Dissipation is the horizontal dissipation that is explicitly added to the model. Vertical dissipation is implemented within the physics component. These components evolve on different time scales, so we use a form of dynamics/tracer/physics subcycling achieved through multi-stage

2nd order accurate Runge-Kutta methods. The tracer and dynamics components are subcycled with respect to the physics. Instead of further subcycling the dynamics component with respect to the tracer component, we use a different Runge-Kutta method for each component. The two methods are chosen so that their stability regions are similar and thus both components will be stable with the same timestep.

Using the spatial discretization described below, the primitive equations can be written as an ODE for a vector U containing all the prognostic state variables and right-hand-side terms representing the forcing (F), dynamics and tracers (A) and dissipation (D),

$$\frac{\partial U}{\partial t} = \mathrm{F} + \mathrm{A} + \mathrm{D}.$$

CAM-SE solves this equation in a time-split fully explicit form. For simplicity we illustrate this using the forward-Euler method:

$$U^* = U^t + \Delta t \mathrm{F}(U^{t_0}) \tag{1}$$

$$U^{**} = U^* + \Delta t \mathrm{A}(U^*) \tag{2}$$

$$U^{t+\Delta t} = U^{**} + \Delta t \mathrm{D}(U^{**}) \tag{3}$$

For the forcing and dissipation steps, (1) and (3), the equations are solved exactly as written. The dissipation term used is a hyper-viscosity operator described in Sec. 2.9. The forcing tendency F (computed by the CAM physics) is typically held fixed over `nsplit` timesteps, where `nsplit` is a CAM input parameter common to all the CAM dynamical cores. Every `nsplit` timesteps, we take $U^{t_0} = U^t$ and recompute $F(U^{t_0})$. The CAM physics package updates the forcings based on a physics timestep of Δt_{phys} minutes. (In our simulations we use $\Delta t_{\mathrm{phys}} = 20$ minutes.) Once Δt_{phys} is chosen, we take $\Delta t = \Delta t_{\mathrm{phys}}/$`nsplit` with `nsplit` determined so that the resulting Δt satisfies the CFL restriction required by (2).

For the dynamics and tracer advection, (2), CAM solves the atmospheric primitive equations. We do not use the forward-Euler approach shown in (2), but instead use a family of N-stage Runge-Kutta methods for better accuracy. The dynamics advances the prognostic variables of the atmospheric primitive equations, while the tracer advection is used for specific humidity, liquid water and ice, as well as additional quantities depending on the CAM configuration. The tracer advection time-step uses a 3 stage, second-order accurate RK-SSP method (Spiteri and Ruuth, 2002). The SSP method ensures that the the timestep will preserve any monotonicity properties preserved by the underlying spatial discretization.

For the dynamics, we use a second order accurate N-stage RK method, where the extra stages are chosen to maximize the stable timestep size. This method allows for a CFL number close to $N - 1$ (normalized so that the Robert-filtered Leapfrog method

has a CFL number of 1.0). It was chosen since it can be implemented as a combination of the RK2 and Leapfrog schemes that were previously used in CAM-SE. The most computationally efficient second order N-stage method for pure imaginary eigenvalues (van der Houwen, 1977; Kinnmark and Gray, 1984) can obtain an advective CFL number of $N-1$, and we hope to evaluate this method in CAM-SE in the future. The leapfrog scheme is more efficient then these methods and has been used successfully in CAM-EUL, but we do not use leapfrog in CAM-SE because we were only able to achieve tracer-mass consistency (per Sec. 2.8) if both dynamics and tracers used RK based methods.

For dynamics, the CFL number is controlled by the maximum gravity wave speed c_0, while for tracer advection it is controlled by the maximum wind speed u_0. We choose N sufficiently large so that the dynamics and tracer advection can both use the same timestep. For typical values $u_0 \approx 140$m/s and $c_0 \approx 340$m/s, $N = 5$ is found to be suitable. The Butcher tableau (Butcher, 2003) for $N = 5$ is given by:

0					
$\frac{1}{8}$	$\frac{1}{8}$				
$\frac{1}{4}$	0	$\frac{1}{4}$			
$\frac{1}{2}$	0	0	$\frac{1}{2}$		
$\frac{3}{4}$	0	$\frac{1}{4}$	0	$\frac{1}{2}$	
	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$

2.2 Continuum Formulation of the Equations

CAM-SE uses a conventional vector-invariant form of the moist primitive equations. For the vertical discretization it uses the hybrid η pressure vertical coordinate system derived from (Kasahara, 1974; Simmons and Burridge, 1981; Simmons and Strüfing, 1981). In order to conserve both mass and energy, surface pressure is used as a prognostic variable as opposed to its logarithm.

In the η-coordinate system, the pressure is given by $p(\eta) = A(\eta)p_0 + B(\eta)p_s$ where p_0 is a constant and p_s is the surface pressure. The hydrostatic approximation $\partial p/\partial z = -g\rho$ with height coordinate z is used to replace the mass density ρ by an η-coordinate pseudo-density $\partial p/\partial \eta$. Following conventional notation (Satoh, 2004), we let $\dot{\eta}$ denote the η-coordinate velocity, $\vec{u}$ the horizontal velocity component tangent to constant z-surfaces (not η-surfaces), $\nabla(), \nabla\cdot()$ and $\nabla\times()$ denote the two-dimensional gradient, divergence and curl operators on constant η-surfaces, and $\partial/\partial\eta$ is the vertical derivative operator. The η-coordinate atmospheric primitive equations, neglecting dissipation and

forcing terms can then be written as

$$\frac{\partial \vec{u}}{\partial t} + (\zeta + f)\,\hat{k}\times\vec{u} + \nabla\left(\frac{1}{2}\vec{u}^2 + \Phi\right) + \dot{\eta}\frac{\partial \vec{u}}{\partial \eta} + \frac{RT_v}{p}\nabla p = 0 \tag{4}$$

$$\frac{\partial T}{\partial t} + \vec{u}\cdot\nabla T + \dot{\eta}\frac{\partial T}{\partial \eta} - \frac{RT_v}{c_p^* p}\omega = 0 \tag{5}$$

$$\frac{\partial}{\partial t}\left(\frac{\partial p}{\partial \eta}\right) + \nabla\cdot\left(\frac{\partial p}{\partial \eta}\vec{u}\right) + \frac{\partial}{\partial \eta}\left(\dot{\eta}\frac{\partial p}{\partial \eta}\right) = 0 \tag{6}$$

$$\frac{\partial}{\partial t}\left(\frac{\partial p}{\partial \eta}q\right) + \nabla\cdot\left(\frac{\partial p}{\partial \eta}q\vec{u}\right) + \frac{\partial}{\partial \eta}\left(\dot{\eta}\frac{\partial p}{\partial \eta}q\right) = 0 \tag{7}$$

These are prognostic equations for $\vec{u}$, the temperature T, density $\frac{\partial p}{\partial \eta}$, and $\frac{\partial p}{\partial \eta}q$ where q is the specific humidity. The prognostic variables are functions of time t, vertical coordinate η and two coordinates describing the surface of the sphere. The unit vector normal to the surface of the sphere is denoted by $\hat{k}$. This formulation has already incorporated the hydrostatic equation and the ideal gas law, $p = \rho R T_v$. There is a no-flux ($\dot{\eta} = 0$) boundary condition at $\eta = 1$ and $\eta = \eta_{\text{top}}$. The vorticity is denoted by $\zeta = \hat{k}\cdot\nabla\times\vec{u}$, f is a Coriolis term, $\omega = Dp/Dt$ is the pressure vertical velocity, and Φ_s is the prescribed surface geopotential. The virtual temperature T_v and variable-of-convenience c_p^* are defined by $RT_v = RT + (R_v - R)\,qT$ and $c_p^* = c_p + (c_{pv} - c_p)\,q$ where R and R_v the ideal gas constants for dry air and water vapor, respectively and c_p, c_{pv} the specific heat at constant pressure for dry air and water vapor, respectively.

To complete the system, we need diagnostic equations for the geopotential height Φ, $\dot{\eta}$ and ω. The diagnostic equation for Φ is given by

$$\Phi = \Phi_s + \int_\eta^1 \frac{RT_v}{p}\frac{\partial p}{\partial \eta}\,d\eta. \tag{8}$$

The remaining diagnostic equations come from integrating (6) with respect to η. For numerical efficiency, (6) can be replaced by a diagnostic equation for $\dot{\eta}\frac{\partial p}{\partial \eta}$ and a prognostic equation for surface pressure p_s

$$\frac{\partial}{\partial t}p_s + \int_{\eta_{\text{top}}}^1 \nabla\cdot\left(\frac{\partial p}{\partial \eta}\vec{u}\right)\,d\eta = 0 \tag{9}$$

$$\dot{\eta}\frac{\partial p}{\partial \eta} = -\frac{\partial p}{\partial t} - \int_{\eta_{\text{top}}}^\eta \nabla\cdot\left(\frac{\partial p}{\partial \eta'}\vec{u}\right)\,d\eta', \tag{10}$$

where (9) is (10) evaluated at the model bottom ($\eta = 1$) after using that $\partial p/\partial t = B(\eta)\partial p_s/\partial t$ and $\dot{\eta}(1) = 0, B(1) = 1$. The diagnostic equation for ω, using (10), is written

$$\omega = \vec{u}\cdot\nabla p - \int_{\eta_{\text{top}}}^\eta \nabla\cdot\left(\frac{\partial p}{\partial \eta'}\vec{u}\right)\,d\eta' \tag{11}$$

Finally, we rewrite (10) as

$$\dot{\eta}\frac{\partial p}{\partial \eta} = B(\eta)\int_{\eta_{\text{top}}}^{1} \nabla\cdot\left(\frac{\partial p}{\partial \eta}\vec{u}\right)\,d\eta - \int_{\eta_{\text{top}}}^{\eta} \nabla\cdot\left(\frac{\partial p}{\partial \eta'}\vec{u}\right)\,d\eta'. \tag{12}$$

We will discretize the prognostic equations (4), (5), (7), (9), and the diagnostic equations (8), (11), (12) exactly as written, with the integrals replaced by sums, the vertical derivatives replaced by finite differences and the horizontal derivatives treated by the SEM.

2.3 Vertical Discretization

The vertical coordinate system uses a Lorenz staggering of the variables as shown in Fig. 2. Let K be the total number of layers, with variables $\vec{u}, T, q, \omega, \Phi$ at layer mid points denoted by $k = 1, 2, \ldots, K$. We denote layer interfaces by $k+\frac{1}{2}, k = 0, 1, \ldots, K$, so that $\eta_{1/2} = \eta_{\text{top}}$ and $\eta_{K+1/2} = 1$. We consider $(\dot{\eta}\frac{\partial p}{\partial \eta})$ a single quantity given at layer interfaces and defined by (12). The no-flux boundary condition is $(\dot{\eta}\frac{\partial p}{\partial \eta})_{1/2} = (\dot{\eta}\frac{\partial p}{\partial \eta})_{K+1/2} = 0$.

There are two types of vertical derivative operators that appear in the equations: (1) derivatives with respect to η of a layer interface value with the result at layer mid points, which are discretized by centered differences, and (2) the $\dot{\eta}\partial/\partial\eta$ operator which acts on quantities defined at layer midpoints and returns a result also at layer midpoints. It is discretized by a more complex formulation specifically constructed in (Simmons and Burridge, 1981) in order to ensure mass and energy conservation. For the integrals with respect to η, we discretize (12) using the midpoint quadrature rule where the integrand is evaluated at layer midpoints. For the remaining indefinite integrals, we first integrate to layer interface $k-\frac{1}{2}$ using the same midpoint quadrature rule, and then add an additional term representing the integral from $k-\frac{1}{2}$ to k. The precise quadrature formulas used are given in (Taylor, 2011).

The vertical discretization quadrature rules and finite difference stencils are chosen to ensure energy conservation (following Simmons and Burridge (1981)), and to ensure that the discrete equations are as consistent as possible, following Williamson and Olson (1994). By consistency, we mean that the discrete diagnostic equations for ω (11) and $\dot{\eta}\frac{\partial p}{\partial \eta}$ (12) are equivalent to a discrete version of the continuity equation (6) and the pressure evolution equation

$$\omega = \frac{Dp}{Dt} = \frac{\partial p}{\partial t} + (\vec{u}\cdot\nabla p) + \dot{\eta}\frac{\partial p}{\partial \eta}. \tag{13}$$

In CAM-SE, if one applies the discrete $\partial/\partial\eta$ operator to the discretized (12) one recovers a natural discretization of (6) at layer midpoints. The discrete diagnostic

equations can be layer averaged and combined to derive a natural discretization of (13) at layer midpoints.

2.4 The Horizontal Discretization

CAM-SE uses the SEM for the horizontal discretization, applied on each η-surface after applying the vertical discretization from Sec. 2.3. In a continuous-Galerkin finite element method such as SEM, instead of constructing discrete approximations to the derivative operators, we construct a discrete functional space. We then find the function in this space that solves the primitive equations in a minimum residual sense. Compared to finite volume methods, there is less choice in how one constructs the discrete derivative operators in this setting, since functions in the discrete space are represented in terms of known basis functions whose derivatives are known analytically. Below, we give an overview of the SEM used by CAM-SE. For complete details, see (Taylor and Fournier, 2010).

To apply the SEM, we first decompose each η-surface, represented by the surface of the sphere Ω, using a quadrilateral finite-element mesh consisting of a $\{\Omega_m\}_{m=1}^M$ of non-overlapping elements,

$$\Omega = \bigcup \Omega_m$$

We assume the mesh is conforming (has no hanging nodes), and that each element can be mapped to the reference element $[-1,1]^2$, with reference coordinates x and y. CAM-SE can work with any such decomposition of the sphere, the only requirements are that the reference element map is differentiable within each element, and the maps for two adjacent elements agree at the common element edge. When uniform resolution is desired at all locations, we utilize the cubed-sphere grid (Fig. 1) with the equal-angle gnomonic projection (Rančić et al., 1996). This grid is the most natural way to generate a highly uniform grid for the sphere with all quadrilateral elements.

The SEM works in the space of globally continuous piecewise polynomials, $\mathcal{V}_d^1$. To describe this space, we first introduce the larger space $\mathcal{V}_d^0$, which is the space of all functions whose values within each element Ω_m are polynomials in x and y of at most degree d. These functions are not required to be continuous at element boundaries. We then take $\mathcal{V}_d^1 = C^0 \cap \mathcal{V}_d^0$, where C^0 is the set of all continuous functions in Ω. The SEM is a Galerkin method with respect to the $\mathcal{V}_d^1$ subspace.

In CAM-SE, the default configuration uses $d = 3$, which achieves a globally 4'th order accurate horizontal discretization (Taylor and Fournier, 2010). Functions in $\mathcal{V}_d^0$ can be represented by coefficients of a polynomial expansion (modal basis approach) within each element, but for the SEM it is more efficient to use an equivalent nodal basis approach, where functions are represented by their grid point values at the $(d+$

$1) \times (d+1)$ tensor product of Gauss-Lobatto-Legendre (GLL) quadrature points (shown in Fig. 3). In CAM-SE, functions in $\mathcal{V}_d^1$ are also represented as functions in $\mathcal{V}_d^0$, with the continuity constraint being enforced by requiring that grid point values at GLL points shared by multiple elements agree.

The horizontal discretization of the primitive equations requires discrete divergence, gradient and curl operators. These are the fundamental operators that appear in CAM-SE. All other derivative operators (such as hyper-viscosity, discussed later) are written in terms of these operators. Non-differential operations such as cross products are computed directly at the GLL grid points. As the method is fully collocated, no spatial averaging is needed.

The divergence, curl and gradient operators are computed by representing their argument in terms of its polynomial expansion and differentiating this expression analytically. The SEM is well known for being quite efficient in computing these types of operations. Because of the tensor product decomposition, the SEM divergence, gradient and curl can all be evaluated at the $(d+1)^2$ GLL nodes within each element in $O(d)$ operations per node using the tensor-product property of these points (Deville et al., 2002; Karniadakis and Sherwin, 2005).

In the case where the Jacobian of the map to the reference element is constant, these operators will be computed exactly for all functions in $\mathcal{V}_d^1$. However, the result will belong to $\mathcal{V}_d^0$, as it will be multi-valued at element boundaries because derivatives computed with the polynomial expansions in adjacent elements will not necessarily agree along their shared boundary. In CAM-SE, we have a non-constant Jacobian and these operators may not be computed exactly by the SEM due to the Jacobian factors in the operators and the Jacobian factors that appear when converting between covariant and contravariant coordinates. In this case there is some freedom in how the operators are constructed. We use the formulation from Thomas and Loft (2000).

2.5 Horizontal Discretization: Discrete Inner-Product

Instead of using exact integration of the basis functions as in a traditional finite-element method, the SEM uses quadrature approximation for the integral over Ω based on the same GLL points used for representing $\mathcal{V}_d^0$. We denote this quadrature approximation by $\langle \cdot \rangle$. It is given by first writing the integral over the sphere as a sum of area-weighted integrals over the set of elements $\{\Omega_m\}_{m=1}^M$,

$$\langle fg \rangle \approx \int fg \, d\mathcal{A} = \sum_{m=1}^{M} \int_{\Omega_m} fg \, d\mathcal{A} = \sum_{m=1}^{M} \iint_{[-1,1]^2} fg J_m \, dx \, dy$$

where J_m is the Jacobian of the map from the reference element to Ω_m. The integral over $[-1,1]^2$ is then approximated by the standard GLL quadrature rule

$$\iint_{[-1,1]^2} fgJ_m \, dx \, dy \approx \sum w_i w_j J_m(i,j) f(i,j) g(i,j)$$

with Gauss-Legendre weights w_i and w_j. When applied to the product of functions $f, g \in \mathcal{V}_d^0$, the quadrature approximation $\langle fg \rangle$ defines a discrete inner-product in the usual manner.

2.6 The SEM Projection Operators

Let $P : \mathcal{V}_d^0 \to \mathcal{V}_d^1$ be the unique orthogonal projection operator from $\mathcal{V}_d^0$ onto $\mathcal{V}_d^1$ with respect to the SEM discrete inner product $\langle \cdot \rangle$. The operation P is essentially the same as the common procedure in the SEM described as *assembly* (Karniadakis and Sherwin, 2005, p. 7), or *direct stiffness summation* (Deville et al., 2002, eq. 4.5.8). Applying the projection operator in a finite element method requires inverting the finite element mass matrix. A remarkable fact about the SEM is that with the GLL-based discrete inner product and the careful choice of global basis functions, the mass matrix is diagonal (Maday and Patera, 1987). The resulting projection operator then has a very simple form: at element interior points, it leaves the nodal values unchanged, while at element boundary points shared by multiple elements it is a Jacobian-weighted average over all redundant values.

2.7 Galerkin Formulation

We can now describe the Galerkin formulation of (2). The SEM solves this equation in integral form with respect to the SEM inner product. That is, for $A \in \mathcal{V}_d^0$, the SEM finds the unique $U^{t+\Delta t} \in \mathcal{V}_d^1$ such that

$$\left\langle \phi U^{t+\Delta t} \right\rangle = \left\langle \phi \left(U^t + \Delta t A(U^t) \right) \right\rangle \qquad \forall \phi \in \mathcal{V}_d^1, \tag{14}$$

where again we use a forward-Euler timestep for simplicity of illustration only. As U^t belongs to $\mathcal{V}_d^1$, the RHS will in general belong to $\mathcal{V}_d^0$ since it contains derivatives of the prognostic variables, resulting in the loss of continuity at the element boundaries. Picking a suitable basis for $\mathcal{V}_d^1$, and requiring (14) to hold for all basis functions results in a system of L equations for the L expansion coefficients of $\frac{\partial U}{\partial t}$, where $L = \dim \mathcal{V}_d^1$. The SEM solves these equations exactly, and the solution can be written in terms of the SEM projection operator P:

$$U^{t+\Delta t} = P\left(U^t + \Delta t A(U^t) \right) = U^t + \Delta t P\left(A(U^t) \right).$$

The projection operator commutes with any time-stepping scheme, so the equations can be solved in a two step process,

- Step 1:
$$U^* = U^t + \Delta t\, A(U^t) \qquad U^* \in \mathcal{V}_d^0$$
- Step 2:
$$U^{t+\Delta t} = P\,(U^*) \qquad U^{t+\Delta t} \in \mathcal{V}_d^1$$

In a traditional finite element method, step 2 would require inverting the finite element mass matrix, usually with an iterative method. But in the SEM, because of the diagonal mass matrix (per Sec. 2.6), step 2 is fully explicit and no solver is needed.

This two step Galerkin solution process represents a natural separation between computation and communication for the implementation of the SEM on a parallel computer. The computations in step 1 are all local to the data contained in a single element. Assuming an element-based decomposition so that each processor contains at least one element, no inter-processor communication is required in step 1. All inter-processor communication in CAM-SE is isolated to the projection operator step, in which element boundary data must be exchanged between adjacent elements.

This decomposition is the main reason the SEM performs so well on parallel computers. The element structure provides a natural cache blocking strategy, as all the data needed for all element computations within the $(d+1) \times (d+1)$ vertical columns belonging to each element will easily fit into cache. Furthermore, as all communication is isolated to a single projection operator, one needs only focus on load balancing strategies and message scheduling designed for this single operator. Significant resources have been devoted to optimizing the SEM projection operator (Loft et al., 2001; Dennis, 2003).

2.8 Tracer Advection

All tracers, including specific humidity, are advected with a discretized version of (7). Due to the need to impose monotonicity constraints, tracer advection is handled within the vertical discretization described above but using the vertically Lagrangian approach from Lin (2004). The main advantage of this approach is that it avoids the difficulty of developing a three-dimensional monotone limiter by separating the problem into horizontal and vertical components, where the horizontal component can use limiters developed specifically for the spectral-element method and the vertical component can use any of the well-proven one-dimensional limiters. The vertically Lagrangian approach can also be used in the dynamics, as in Lin (2004). Currently CAM-SE does not use this approach for the dynamics (Equations (4),(5) and (9)) because we do not

impose any monotonicity constraints on these variables and the Simmons and Burridge (1981) approach is more efficient if limiters are not needed.

In CAM-SE, at the beginning of each timestep the tracers are assumed to be given on the η-coordinate layer midpoints. The tracers are advanced in time on a moving vertical coordinate system η' defined so that $\dot{\eta}' = 0$. At the end of the timestep, the tracers are remapped back to the η-coordinate layer midpoints using the monotone remap algorithm from Zerroukat et al. (2005). The horizontal advection step consists of using the SEM to solve

$$\frac{\partial}{\partial t}\left(\frac{\partial p}{\partial \eta} q\right) = -\nabla_{\mathrm{h}} \cdot \left(\overline{\left(\frac{\partial p}{\partial \eta}\vec{u}\right)} q\right) \tag{15}$$

on the surfaces defined by the η' layer midpoints. The quantity $\overline{\left(\frac{\partial p}{\partial \eta}\vec{u}\right)}$ is the mean flux computed during the dynamics update. The mean flux used in (15), combined with a suitable mean vertical flux used in the remap stage allows CAM-SE to preserve mass/tracer-mass consistency: The tracer advection of $\frac{\partial p}{\partial \eta} q$ with $q = 1$ will be identical to the advection of $\frac{\partial p}{\partial \eta}$ implied from (6).

Simply discretizing (15) with the SEM will result in locally conservative, high-order accurate but oscillatory transport scheme. A limiter is added to reduce or eliminate these oscillations (Taylor et al., 2009). CAM-SE supports both monotone and sign-preserving limiters, but the most effective limiter for CAM-SE has not yet been determined. The default configuration is to use the sign-preserving limiter to prevent negative values of q coupled with a sign-preserving hyper-viscosity operator which dissipates q^2. We use the optimal 3-stage second-order RK-SSP timestepping method from Spiteri and Ruuth (2002) which preserves the monotonicity properties enforced by the limiter.

2.9 Dissipation

A horizontal hyper-viscosity operator, modeled after that used in CAM-EUL is applied to the momentum and temperature equations. It is applied in a time-split manner after each dynamics timestep. The hyper-viscosity step for vectors is written

$$\vec{u}(t + \Delta t) = \vec{u}(t) - \Delta t \nu \nabla^4 \vec{u}. \tag{16}$$

A possible implementation of this operator would be the repeated application of the Laplacian. The Laplacian operator has been used previously in HOMME, in a baroclinic instability test case (Polvani et al., 2004). In that formulation, the Laplacian required two applications of the SEM projection operator, suggesting the hyper-viscosity

operator would require four such applications. As the timestep only requires one application, and all the communication in the SEM are isolated to this operator, this represents a potential 4x increase in the amount of communication per timestep.

Here we present a more efficient hyper-viscosity operator that only requires 2 projection steps. It uses an integral form of the operator, from the mixed finite element approach (following Giraldo (1999)) which writes the equation as a system of equations involving only first derivatives. We start by introduced an auxiliary vector $\vec{f}$ so that hyper-viscosity step (16) can be written

$$\vec{u}(t+\Delta t) = \vec{u}(t) - \Delta t\, \nu \nabla^2 \vec{f} \qquad \vec{f} = \nabla^2 \vec{u}.$$

In order to avoid expressing ∇^2 in spherical coordinates, we use the identity $\nabla^2\vec{u} = \nabla(\nabla\cdot\vec{u}) - \nabla\times(\nabla\times\vec{u})$ to write it in terms of the same operators that appear in the vector-invariant form of the primitive equations, resulting in

$$\vec{u}(t+\Delta t) = \vec{u}(t) - \Delta t\, \nu \left(\nabla(\nabla\cdot\vec{f}) - \nabla\times\hat{k}(\nabla\times\vec{f})\right) \tag{17}$$

$$\vec{f} = \nabla(\nabla\cdot\vec{u}) - \nabla\times(\nabla\times\vec{u})\hat{k}. \tag{18}$$

Applying the Galerkin approach of writing this as a system of integral equations for a family of test functions $\vec{\phi}$, and then integrating the gradient and curl operators by parts gives

$$\iint \vec{\phi}\cdot\vec{u}(t+\Delta t)\, d\mathcal{A} = \iint \vec{\phi}\cdot\vec{u}(t)\, d\mathcal{A} + \nu \iint \left[(\nabla\cdot\vec{\phi})(\nabla\cdot\vec{f}) + (\nabla\times\vec{\phi})\cdot\hat{k}(\nabla\times\vec{f})\right] d\mathcal{A} \tag{19}$$

$$\iint \vec{\phi}\cdot\vec{f}\, d\mathcal{A} = -\iint \left[(\nabla\cdot\vec{\phi})(\nabla\cdot\vec{u}) + (\nabla\times\vec{\phi})\cdot\hat{k}(\nabla\times\vec{u})\right] d\mathcal{A}. \tag{20}$$

These equations are then discretized with the SEM approach from Sec. 2.7. The integrals are approximated with the SEM inner-product $\langle\cdot\rangle$, resulting in a linear system of equations that must be solved. Again due to the diagonal mass matrix, the resulting system can be solved directly with two applications of the projection operator, first to compute the auxiliary vector $\vec{f}$ and then to compute $\vec{u}(t+\Delta t)$.

The hyper-viscosity formulation used for scalars such as T is simpler, since instead of the vector Laplacian identity we use $\Delta T = \nabla\cdot\nabla T$. Otherwise the approach is identical to that used above so we omit the details.

2.10 Conservation and Compatibility

The SEM used in CAM-SE is *compatible*, meaning it has a discrete version of the divergence theorem, Stokes theorem and curl/gradient annihilator properties (Taylor

and Fournier, 2010). The divergence theorem is the key property of the horizontal discretization that is needed to show conservation of both mass and energy. Because of the discrete divergence theorem, a compatible method obtains global energy conservation by mimicking the behavior of the continuum energy dynamics on a term-by-term basis. The discrete form of the terms in the energy equation that are responsible for the transfer between kinetic, internal and potential energy will be in exact balance, while the advection terms will vanish as in the continuum form of the equations.

For an arbitrary scalar h and vector $\vec{u}$ at layer midpoints, the divergence theorem (or the divergence/gradient adjoint relation) can be written

$$\iint_D h \nabla \cdot \vec{u} \, d\mathcal{A} + \iint_D \vec{u} \nabla h \, d\mathcal{A} = \oint_{\partial D} h \vec{u} \cdot \hat{n} \, dS$$

for any domain D with piecewise smooth boundary. The SEM discrete divergence and Stokes theorem apply locally at the element level, $D = \Omega_m$, or any collection of elements. This is the fundamental identity needed to show mass and energy conservation. When combined with the curl/gradient annihilator properties, it is used to show that the horizontal advection operator locally conserves the two-dimensional potential vorticity.

In the vertical, Simmons and Burridge (1981) showed that the discrete $\partial/\partial\eta$ and $\dot{\eta}\partial/\partial\eta$ operators need to satisfy two integral identities to ensure energy and mass conservation. For any layer interface velocity $\dot{\eta}$ that satisfies $\dot{\eta}_{1/2} = \dot{\eta}_{K+1/2} = 0$ and f, g arbitrary functions of layer midpoints, the first identity is the adjoint relationship between the $\dot{\eta}\partial/\partial\eta$ operator and the $\partial/\partial\eta$ operator acting on functions of the form $\dot{\eta} f$. The second identity is a discrete integrated-by-parts analog of $\partial(fg) = f\partial g + g\partial f$. Construction of methods with both properties on a staggered unequally spaced grid is the reason behind the complex definition for $\dot{\eta}\partial/\partial\eta$.

The energy conservation properties of CAM-SE were studied in Taylor (2011) using the aqua planet test case (Neale and Hoskins, 2000a,b). CAM-SE uses

$$E = \left\langle \sum_{i=1}^{K} \Delta\eta_i \left(\frac{\partial p}{\partial \eta}\right)_i \left(\frac{1}{2}\vec{u}^2 + c_p^* T\right)_i \right\rangle + \langle p_s \Phi_s \rangle$$

as the discretization of the total moist energy. The conservation of E is *semi-discrete*, meaning that the only error in conservation is the time truncation error. In the adiabatic case (with no hyper-viscosity and no limiters), running from a fully spun up initial condition, the error in conservation decreases to machine precision at a second-order rate with decreasing timestep. In the full non-adiabatic case with a realistic timestep, $dE/dt \approx 0.013 \mathrm{W/m}^2$, or about 0.04% of the mean solar flux.

3 OpenMP/MPI strategy

CAM-SE uses both MPI and OpenMP for distributed memory and shared memory parallelism, respectively. The collection of elements is horizontally domain-decomposed then mapped to processes based on a space-filling curve ordering and a specified two-dimensional process topology. Subdomains contain entire vertical columns, and the horizontal extent of a given subdomain is disjoint from that of other subdomains with the exception of the subdomain boundary. MPI communicates data between the respective memories. Most of the interprocess communication within CAM-SE occurs at element boundaries (hence with nearby processes in the process topology), but there are several instances of global communication, mostly in the CAM physics for diagnostics. The decomposition paradigm requires at least one element per MPI task.

Within a given MPI task, OpenMP may be applied to utilize multiple cores in shared memory space (e.g., on a processor or node). The present OpenMP paradigm within CAM-SE is with respect to elements. As little as one element may be assigned to a thread, and there is no provision at this time to assign multiple threads to a given element. OpenMP is done so in a static manner. In advance, specific threads are assigned to specific elements. The solution timestep is one large parallel region. There are certain instances though where only the master thread for a given task is active. The most notable of these are during message-passing and the reproducible (parallel decomposition independent) distributed sum computation.

Within CAM-SE, OpenMP and MPI work against one-another. The total number of threads is limited to the number of elements; it is a question of how those threads are divided between MPI and OpenMP. In many instances it is superior to apply more distributed memory parallelism (MPI) than shared memory parallelism (OpenMP). However, at scale, using more OpenMP threads and fewer MPI tasks is often superior because of the high cost of global communication for very large MPI task counts.

4 Performance Benchmarks

For our benchmark results, we evaluate the performance and scalability of CAM-SE, CAM-FV and CAM-EUL using a realistic CESM time-slice configuration. This configuration is often used for short high-resolution simulations driven by prescribed ocean conditions where the prescribed conditions come from a slice in time of a century-long, low-resolution coupled atmosphere-ocean simulation. The time-slice configuration uses the CESM's atmosphere (CAM), land model (CLM), and sea ice model (CICE), but with prescribed sea ice extent and sea surface temperatures (SST). The SST data are provided by the CESM data ocean component, DOCN. For our simulations we used

prescribed data that simulates 1850 pre-industrial conditions, as implemented in the CESM version 1 F_1850 *compset*.

The CAM supports several atmospheric physics configurations. For our simulations, we used the CAM4 physics option (the default physics used with CAM Version 4). CAM4 physics uses 26 vertical levels and advects three tracers (specific humidity, liquid water and ice) in addition to the prognostic variables for temperature, velocity and pressure. All simulations used the CLM model running on a 0.25° (768×1152) latitude/longitude grid (the same horizontal grid employed by the 0.25° CAM-FV). DOCN and CICE used a 0.1° (2400×3600) tri-pole grid. We made this choice in order to mimic a global high-resolution fully coupled (active ocean) simulation. Even though the prescribed SST data comes from a low resolution (1.0°) data set, it is interpolated by the DOCN component to the 0.1° tri-pole grid during the simulation. Thus all the costs of interacting with a high-resolution ocean are included in this configuration; only the cost of performing the ocean simulation is omitted. In the CESM, the ocean component can be run concurrently with the other model components, using an independent set of processors. Thus the timings reported here will be similar to the timings of a fully coupled CESM simulation where the DOCN is replaced with an active ocean model and the ocean model is given additional processors so that its integration rate matches that of the other CESM components.

In the CAM-FV configuration described above, only two different horizontal grids are present. For the CAM-SE and CAM-EUL configurations, CAM and CLM are using different grids. As a result, the CESM coupling will involve three different grids, requiring the CESM *tri-grid* configuration described in (Craig et al., 2011). Our CAM-EUL configuration uses the T341 truncation, with its associated 512×1024 latitude/longitude grid on the sphere. CAM-SE uses a cubed-sphere grid and we present results from both 0.25° and 0.125° resolutions, where the resolution is determined by the average grid spacing at the Equator. The 0.25° and 0.125° grids have 777,602 and 3,110,402 total grid points (in the horizontal) distributed among 86,400 and 345,600 elements, respectively. For all the dynamical cores, we used $\Delta t_{\mathrm{p}hys} = 900$ seconds. For the tracer timestep Δt, we used the largest stable value for which $900/\Delta t$ is an integer (and for CAM-FV, the largest stable dynamics timestep which evenly divides Δt). For the 0.25° CAM-SE, $\Delta t = 81.82$, using a 3 stage (tracers) and 5 stage (dynamics) Runge-Kutta method. For the 0.125° CAM-SE, Δt was reduced by a factor of 2. For T341 CAM-EUL, the tracer and dynamics timestep was $\Delta t = 150.0$, using a semi-Lagrange method (tracers) and a semi-implicit Robert filtered leapfrog scheme (dynamics). For 0.25° CAM-FV, which supports subcycling the dynamics with respect to the tracers, $\Delta t = 112.5$ and $\Delta t_{dyn} = 56.25$. All other parameters were left at their CAM default values.

4.1 Load balancing

The CESM supports a very flexible static load balancing capability. Every model component as well as the flux coupler can be assigned arbitrary subsets of MPI tasks which may or may not overlap. On modern machines with tens of thousands of processor cores, the best performance can be obtained by identifying the components that can run concurrently and assigning them independent MPI tasks. The ocean component is the main example. It needs to communicate with the atmosphere component once per model day, but otherwise runs independently of all other components. In contrast, the ice and land models communicate with the atmosphere every atmosphere model timestep and can not run concurrently with the atmosphere. (However, the land and ice can run concurrently with one-another.) All communication between the component models is via the coupler component.

As our focus is solely on the evaluation and comparison of performance of the three atmospheric dynamical cores, a *stacked* configuration that uses the same processors for all components, running the component models sequentially, one after the other, would do the best job of isolating the performance of the atmosphere. An example CESM stacked configuration is illustrated in Fig. 4, left panel. This can require a significant amount of memory, since a processor shared by several components will need sufficient memory to support all the components. Due to memory limitations on one of our target systems, a stacked configuration was not possible. Instead we ran the ocean model on a separate set of processors and ran the coupler, ice, and land model on a subset of the same processors as the atmosphere model, as shown in Fig. 4, right panel. This configuration is very efficient in terms of both memory usage and allows the maximum amount of concurrent execution. More specifically, given N MPI tasks we assigned 64 independent tasks to DOCN, and assigned the remaining $N - 64$ tasks to CAM. Within these $N - 64$ tasks, we assigned non-overlapping subsets to the remaining three components (flux coupler, CICE, CLM). The land and ice components will then run concurrently with each other, while the coupler, land, and ice components run sequentially with respect to the atmosphere.

Note that we did not attempt to determine a configuration of the other components (land, ice, ocean) that would optimize the throughput of the CESM run, and we report only the execution time of the atmosphere component. To further insulate our CAM benchmark numbers from the load imbalance between CLM and CICE, we added a call to MPI_Barrier (using an MPI communicator that included only the atmosphere processes) at the beginning of each call to the atmosphere model. Without this barrier, any CLM/CICE load imbalance could be counted as time spent running CAM, while with this barrier this time is attributed instead to the coupler. We also timed this barrier to determine whether load imbalances introduced by the other component

models would have impacted performance of the atmosphere model if the barrier were not called. In all cases the time spent in this barrier was insignificant compared to the execution time of the atmosphere model as a whole.

4.2 Benchmark Results

The CAM performance was evaluated on two Leadership Computing Platforms (LCF), the Oak Ridge LCF Cray XT5 JaguarPF system and the Argonne LCF IBM Blue Gene/P Intrepid system. Each of the 18,688 compute nodes in JaguarPF system nodes consists of two hex-core processors for a total of 224,256 cores. Each of the 40,960 nodes of Intrepid consists of a quad core processor for a total of 163,840 cores.

In Fig. 5 we show the approximate 0.25° comparison of CAM running with each of the three dynamical cores on Intrepid. All simulations were made using CAM's hybrid openMP/MPI approach, with one MPI task per node and 4 OpenMP threads per MPI task (1 thread per core on Intrepid). We first note that CAM-EUL is the most efficient in terms of core-hours per simulated day. On two-thousand cores, it is more than twice as fast as CAM-SE and CAM-FV. This is largely due to a recent implementation of subcycled dynamics Taylor et al. (2011) and the fact that it uses a semi-implicit method that allows for a relatively larger timestep. However its scalability is limited due to the spherical harmonic transforms. As the number of cores is increased, the integration rate stagnates at about 0.9 simulated year per day (SYPD). CAM-FV benefits from a larger timestep, while for the same resolution, CAM-SE benefits from fewer degrees of freedom becuase of its more uniform grid. Above eight thousand cores, CAM-FV starts to lose scalability due in part to its inherently load-imbalanced polar filtering. CAM-SE exhibits near perfect strong scaling to 1 element per core (86,400 cores), obtaining 12.2 SYPD. Currently it is not possible to execute CAM-SE using multiple MPI tasks per element. Further, unlike in CAM-FV, the use of the OpenMP/MPI programing model within CAM-SE does not enable additional parallelism. We hope to enable CAM-SE at 0.25° to use more than 86,000 cores by utilizing additional parallelism within the physics component and through OpenMP threading of the vertical dimension.

In Fig. 6 we show results from CAM-SE at 0.25° and 0.125° resolutions on both Intrepid and JaguarPF. For JaguarPF, we tested 1,2,3 and 6 threads per MPI task, with 12,6,4 and 2 MPI tasks per node respectively, and report only the best times. When running on few numbers of cores, the best times were obtained with a pure MPI configuration (1 thread per task); we found the best results were obtained by increasing the number of threads whenever the number of MPI tasks neared about 8000.

Focusing on the 0.25° results in Fig. 6, we see that JaguarPF cores are significantly faster than the Intrepid cores. But CAM-SE does not scale as well on JaguarPF, and at the limit of CAM-SE's scalability, both platforms obtain the same throughput. Pre-

liminary analysis suggests that CAM-SE's space-filling curve decomposition strategy is suboptimal for the 12 cores per node processors used by JaguarPF, so we expect future improvements in domain decomposition to improve the JaguarPF results at high task counts. Similar results are seen for the 0.125° benchmarks. This problem size has 345,600 elements and thus would allow for up to 345,600 MPI tasks. For both machines, we have obtained data out to the largest core counts possible, 2 elements per core on JaguarPF and 3 elements per core on Intrepid. Using 1 element per core on JaguarPF or 2 elements per core on Intrepid would require more cores than currently available on these machines. JaguarPF again exhibits significantly faster single-core performance at the low end, while Intrepid shows excellent scalability at all core counts. JaguarPF obtains a throughput of 4.6 SYPD, while Intrepid achieves 2.83 SYPD. To our knowledge, both of these numbers are a new record for any 0.125° atmospheric model running with full physics.

For completeness, in Fig. 7 we show the same 0.25° dynamical core comparison as in Fig. 5, only on the JaguarPF system. Again CAM-EUL is the most efficient in terms of core-hours per simulated day, but cannot run at more than 3 SYPD due to lack of scalability. Both CAM-SE and CAM-FV do not scale as well on JaguarPF as they do on Intrepid. CAM-SE shows better scaling than CAM-FV at all processor counts, but starts to depart from perfect scalability at around 5000 cores. However, Intrepid's better scaling does not quite allow its CAM-SE performance (12.2 SYPD) to catch up with JaguarPF (12.7 SYPD).

5 High Resolution Simulations

The benchmarks presented above show that the CESM with the CAM-SE atmosphere can now take full advantage of modern petascale platforms and achieve integration rates suitable for long term climate studies at very high resolutions. These performance results, coupled with the 1.0° resolution CESM results (Evans et al., 2011) and the proven high-order accuracy and mesh convergence properties of the spectral element method give us confidence that the software engineering and dynamics in the CESM are now fully capable of ultra high resolution simulations.

A challenging remaining issue is to evaluate the suitability of the CAM physics parametrizations at high resolution. The parametrizations have been developed for use with lower resolutions (typically 1.0° or lower), where tuning parameters are based on assumptions of mean values of quantities such as the vertical velocity and relative humidity. Preliminary runs at 0.25° with all three dynamical cores show they are able to produce a global annually-averaged residual energy flux at the model top and bottom within several tenths of a W/m^2 for relatively short 1-3 year simulations with

properly spun up initial conditions. Annual means of the CAM-SE 0.25° preindustrial configuration 500mb geopotential height, tropopause temperature, sea level pressure, and total precipitable water have a root mean square error (rmse) of 0.20, 4.39. 4.59, and 2.46, respectively with respect to observational data (as computed by the CESM Atmospheric Model Working Group Diagnostics Package[2]). Comparing the same quantities of the CAM-FV 1° preindustrial configuration with rmse of 0.21, 5.41, 3.92, and 2.56, respectively, it is clear that the gross features are reasonable.

Many dynamical features are clearly improved at these resolutions, such as the response to topography, meso-scale variability and the ability of the model to capture tropical cyclones. As an example, a snapshot from a 0.125° simulation is shown in Fig. 8, where many fine scale features of the flow, such as well defined mid-latitude fronts and tropical cyclones, are visible. However, the preliminary high-resolution simulations show some issues with mean climate properties such as moisture levels in the free atmosphere. Improving the results through turning or more fundamental changes to the parametrizations is an active area of research.

6 Conclusions

The finite volume and Eulerian global spectral dynamical cores have been engineered to take advantage of supercomputers with thousands of processor cores. However, due to limitations arising from the nature of the underlying computational grids and the dependencies in the numerical schemes, performance scalability has proven to be poor when using tens of thousands of processor cores. The spectral element dynamical core eliminates these scalability bottlenecks and provides excellent performance on petascale platforms, with the potential for better performance on even larger systems (Carpenter et al., 2011). This performance is obtained within a state-of-the-art compatible numerical method, featuring high-order accuracy, mass, energy and PV conservation and sign-preserving consistent advection on fully unstructured quadrilateral grids. Spectral element methods have traditionally scaled very well on modern parallel computers. That scalability has lead to the excellent performance results given here. We have also further increased the method's throughput for climate applications by the introduction of multi-stage RK methods (which achieve the performance benefit of the older subcycling approach) and a new integrated-by-parts hyper-viscosity operator which reduces the number of projection (communication) steps by a factor of 2.

[2]http://www.cgd.ucar.edu/cms/rneale/tools/amwg_diagnostics.html

7 Acknowledgments

We thank Jamison Daniel at Oak Ridge National Laboratory for producing Fig 8. This work was supported by the DOE BER SciDAC 06-13194, *A Scalable and Extensible Earth System Model*. Work by A.A.M. was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Work by P.H.W. and K.J.E. was sponsored by the Climate and Environmental Sciences Division of the Office of Biological and Environmental Research and by the Office of Advanced Scientific Computing Research, both in the Office of Science, U.S. Department of Energy, under Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC. As this work has been co-authored by contractors of the U.S. Government under contracts No. DE-AC05-00OR22725 and No. DE-AC52-07NA27344, the U.S. Government retains a nonexclusive, royalty free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357 and the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

8 Biographies

John M. Dennis received his PhD in computer science at the University of Colorado at Boulder. He is one of the chief architects of the HOMME framework. John's work has been recognized with an honorable mention Gordon Bell prize in 2001 for excellence in high-performance supercomputing. His interests in include high-performance parallel computing, parallel I/O, mesh partitioning computer architecture, and iterative algorithms.

Jim Edwards is a senior software engineer at the National Center for Atmospheric Research. He received his masters in Atmospheric Sciences from Colorado State University in 1994 and has been developing software for geophysics applications since that time. He began working on the HOMME model in 2004 and began incorporating the SE dynamical core in CAM in 2006.

Katherine J. Evans is a research staff member in the Computational Earth Sciences Group at Oak Ridge National Laboratory. She received her Ph.D. from Georgia Institute of Technology in 2000 and completed a postdoc at Los Alamos National Laboratory studying implicit numerical methods for phase change and fluid flows from 2004

to 2006. Kate's work includes high resolution coupled climate model development and scalable numerical methods development in global climate models. She is a member of the American Meteorological Society, American Geophysical Union, and the Society for Industrial and Applied Mathematics.

Oksana Guba is a post-doctoral researcher at Sandia National Laboratories. She received her Ph.D. from the University of New Mexico in 2008 and was a postdoc there studying and implementing numerical methods for modeling sea ice. Her recent research interests are high-order schemes for modeling atmospheric flows.

Peter H. Lauritzen received his PhD from the University of Copenhagen (Denmark). After his PhD he worked at the Danish Meteorological Institute (DMI) before becoming an ASP postdoctoral fellow, and later scientist, at the National Center for Atmospheric Research (NCAR). At NCAR he is in the Atmospheric Modeling and Predictability group that supports, maintains, and develops the atmospheric component of the Community Earth System Model (CESM). Peter's research interests include dynamical cores for weather and climate models, remapping between spherical grids, advection/transport schemes for unstructured grid, and idealized test cases.

Arthur A. Mirin is the leader of the Scientific Computing Group of the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory. Dr. Mirin received his Ph.D. in mathematics from the University of California, Berkeley in 1974 and has been working at the Laboratory since 1969. He spent the first portion of his career working on large-scale computing applied to magnetic fusion energy, and for much of that time was on the staff of the National Energy Research Supercomputer Center. Since the late 1980's, Dr. Mirin has been focusing on high-performance computing, particularly with respect to climate modeling, atmospheric flow, and numerical hydrodynamics. He coordinated the turbulence simulation that won the Gordon Bell Award for Best Performance in 1999. Dr. Mirin is a member of the Society for Industrial and Applied Mathematics, the American Physical Society, the American Geophysical Union, and the Phi Beta Kappa academic honor society.

Amik St.-Cyr is a scientist at the National Center for Atmospheric Research. He is a former postdoctoral fellow from McGill University's Computational Fluid Dynamics Lab and has a Ph.D. in applied mathematics from the University of Montreal (2002) and a bachelor's in mathematical physics. His principal research interests are high-order methods for numerical modeling, time-stepping, parallel computing and iterative methods for large sparse matrices and mesh adaptation.

Mark A. Taylor received his Ph.D. from New York University's Courant Institute of Mathematical Sciences in 1992. Mark worked as a post-doc and then a software engineer at the NCAR from 1992 to 1998. From 1998 to 2004, he was a staff member in the Computer and Computational Sciences Division at Los Alamos National Labo-

ratory. Since 2004, Mark has been a Principal Member of the Technical Staff at Sandia National Laboratories, in the Computation, Computers, Information and Mathematics Center. Mark works on scalable numerical methods for climate modeling.

Patrick H. Worley is a senior research computer scientist in the Computer Science and Mathematics Division of Oak Ridge National Laboratory. His research interests include parallel algorithm design and implementation (especially as applied to simulation models used in climate, fusion energy, and astrophysics research) and the performance evaluation of parallel applications and computer systems. He is currently a co-chair of the CCSM Software Engineering Working Group, the principal investigator for the Performance Engineering and Analysis Consortium End Station DOE INCITE project, and is an Associate Editor of the journal Parallel Computing. Pat has a Ph.D. in computer science from Stanford University. He is a member of the Association for Computing Machinery and of the Society for Industrial and Applied Mathematics.

References

Bhanot, G., Dennis, J. M., Edwards, J., Grabowski, W., Gupta, M., Jordan, K., Loft, R. D., Sexton, J., St.-Cyr, A., Thomas, S. J., Tufo, H. M., Voran, T., Walkup, R., and Wyszogrodski, A. A. (2008). Early experiences with the 360TF IBM BlueGene/L platform. *International Journal of Computational Methods*, 5:237–253.

Butcher, J. (2003). *Numerical methods for ordinary differential equations*. Wiley, Chichester and New York.

Carpenter, I., Archibald, R., Evans, K., Larkin, J., Micikevicius, P., Rosinski, J., Schwarzmeier, J., and Taylor, M. A. (2011). Progress towards accelerating homme on hybrid multi-core systems. *Int. J. High Perf. Comput. Appl.* under review.

Collins, W. D., Rasch, P. J., Boville, B. A., Hack, J. J., McCaa, J. R., Williamson, D. L., Kiehl, J. T., Briegleb, B. P., Bitz, C. M., Lin, S.-J., Zhang, M., and Dai, Y. (2004). Description of the NCAR Community Atmosphere Model (CAM3.0). NCAR Technical Note NCAR/TN-464+STR, National Center for Atmospheric Research, Boulder, Colorado. 214 pp., available from http://www.ucar.edu/library/collections/technotes/technotes.jsp.

Craig, A., Vertenstein, M., and Jacob, R. (2011). A new flexible coupler for earth system modeling developed for CCSM4 and CESM1. *Int. J. High Perf. Comput. Appl.* under review.

Dennis, J., Fournier, A., Spotz, W. F., St.-Cyr, A., Taylor, M. A., Thomas, S. J., and Tufo, H. (2005). High resolution mesh convergence properties and parallel efficiency

of a spectral element atmospheric dynamical core. *Int. J. High Perf. Comput. Appl.*, 19:225–235.

Dennis, J. M. (2003). Partitioning with space-filling curves on the cubed-sphere. In *Proceedings of the Workshop on Massively Parallel Processing at IPDPS'03*. Nice, France.

Deville, M. O., Fischer, P. F., and Mund, E. H. (2002). *High Order Methods for Incompressible Fluid Flow*. Cambridge University Press, 1 edition.

Evans, K., Lauritzen, P., Mishra, S., Neale, R., Taylor, M., and Tribbia, J. (2011). AMIP simulations with the CAM4 spectral element dynamical core. *J. Climate*. under review.

Fournier, A., Taylor, M., and Tribbia, J. (2004). The spectral element atmosphere model (SEAM): High-resolution parallel computation and localized resolution of regional dynamics. *Mon. Wea. Rev.*, 132:726–748.

Giraldo, F. (2005). Semi-implicit time-integrators for a scalable spectral element atmospheric model. *Quart. J. Roy. Meteor. Soc.*, 131:2431–2454.

Giraldo, F. X. (1999). Trajectory calculations for spherical geodesic grids in cartesian space. *Mon. Wea. Rev.*, 127:1651–1662.

Giraldo, F. X. and Rosmond, T. E. (2004). A scalable spectral element eulerian atmospheric model (SEE-AM) for NWP: Dynamical core tests. *Mon. Wea. Rev.*, 132:133–153.

Intergovernmental Panel on Climate Change (2007). *Fourth Assessment Report: Climate Change 2007: The AR4 Synthesis Report*. Geneva: IPCC.

Karniadakis, G. E. and Sherwin, S. J. (2005). *Spectral/hp Element Methods for Computational Fluid Dynamics (Numerical Mathematics and Scientific Computation)*. Oxford University Press, USA, 2 edition.

Kasahara, A. (1974). Various vertical coordinate systems used for numerical weather prediction. *Mon. Wea. Rev.*, 102:509–522.

Khouider, B., St-Cyr, A., Majda, A. J., and Tribbia, J. (2010). MJO and convectively coupled waves in a coarse resolution GCM with a simple multicloud parametrization. *Journal of the Atmospheric Sciences*, 0(0):null.

Kim, Y., Giraldo, F., Flatau, M., Liou, C., and Peng, M. (2008). A sensitivity study of the kelvin wave and the madden-julian oscillation in aqua-planet simulations by the naval research laboratory spectral element atmospheric model (NSEAM). *J. Geophys. Res.*, 113.

Kinnmark, I. and Gray, W. (1984). One step integration methods with maximum stability regions. *Math. Comput. Simulat.*, 26:87–92.

Lauritzen, P., Jablonowski, C., Taylor, M., and Nair, R. (2010). Rotated versions of the jablonowski steady-state and baroclinic wave test cases: A dynamical core intercomparison. *J. Adv. Model. Earth Syst.*, 2. Art.#15, 34pp.

Lin, S.-J. (2004). A vertically lagrangian finite-volume dynamical core for global models. *Mon. Wea. Rev.*, 132:2293–2397.

Loft, R., Thomas, S., and Dennis, J. (2001). Terascale spectral element dynamical core for atmospheric general circulation models. In *Proceedings of the ACM / IEEE Supercomputing SC'2001 conference*.

Maday, Y. and Patera, A. T. (1987). Spectral element methods for the incompressible Navier Stokes equations. In Noor, A. K. and Oden, J. T., editors, *State of the Art Surveys on Computational Mechanics*, pages 71–143. ASME, New York.

Mirin, A. and Worley, P. (2011). Improving the performance scalability of the community atmosphere model. *Int. J. High Perf. Comput. Appl.* under review.

Mirin, A. A. and Sawyer, W. B. (2005). A scalable implementation of a finite-volume dynamical core in the community atmosphere model. *Int. J. High Perf. Comput. Appl.*, 19:203–212.

Mishra, S. K., Taylor, M. A., Nair, R. D., Lauritzen, P., Tufo, H. M., and Tribbia, J. J. (2011a). Evaluation of the HOMME dynamical core in the aqua-planet configuration of NCAR CAM4: Rainfall. *J. Climate*. to appear.

Mishra, S. K., Taylor, M. A., Nair, R. D., Tufo, H. M., and Tribbia, J. J. (2011b). Performance of the homme dynamical core in the aqua-planet configuration of ncar cam4: equatorial waves. *Annales Geophysicae*, 29(2):221–227.

Nair, R. D. (2009). Diffusion experiments with a global discontinuous galerkin shallow water model. *Mon. Wea. Rev.*, 137:3339–3350.

Neale, R. B. and Hoskins, B. J. (2000a). A standard test case for AGCMs including their physical parametrizations: I: The proposal. *Atmos. Sci. Lett.*, 1:101–107.

Neale, R. B. and Hoskins, B. J. (2000b). A standard test case for AGCMs including their physical parametrizations: II: Results for the Met Office Model. *Atmos. Sci. Lett.*, 1:108–114.

Neale, R. B., Richter, J. H., Conley, A. J., Park, S., Lauritzen, P. H., Gettelman, A., Williamson, D., Rasch, P., Vavrus, S., Taylor, M., Collins, W. D., Zhang, M., and Lin, S. (2010). Description of the ncar community atmosphere model (CAM 4.0). Technical Report TN–???+STR, NCAR.

Polvani, L. M., Scott, R. K., and Thomas, S. (2004). Numerically converged solutions of the global primitive equations for testing the dynamical core of atmospheric gcms. *Mon. Wea. Rev.*, 11:2539–2552.

Putman, W. M., Lin, S.-J., and Shen, B.-W. (2005). Cross-platform performance of a portable communication module and the nasa finite volume general circulation model. *Int. J. High Perf. Comput. Appl.*, 19:213–223.

Rančić, M., Purser, R., and Mesinger, F. (1996). A global shallow-water model using an expanded spherical cube: Gnomonic versus conformal coordinates. *Q. J. R. Meteorol. Soc.*, 122:959–982.

Sadourny, R. (1972). Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids. *Mon. Wea. Rev.*, 100(2):136–144.

Satoh, M. (2004). *Atmospheric Circulation Dynamics and General Circulation Models*. Springer, 1 edition.

Simmons, A. J. and Burridge, D. M. (1981). An energy and angular momentum conserving vertical finite-difference scheme and hybrid vertical coordinates. *Mon. Wea. Rev.*, 109:758–766.

Simmons, A. J. and Strüfing, R. (1981). An energy and angular-momentum conserving finite-difference scheme, hybrid coordinates and medium-range weather prediction. Technical Report 28, European Centre for Medium-Range Weather Forecasts, Reading, U.K. 68 pages.

Spiteri, R. J. and Ruuth, S. J. (2002). A new class of optimal high-order strong-stability-preserving time discretization methods. *SIAM J. Numer. Anal.*, 40(2):469–491.

St.-Cyr, A., Jablonowski, C., Dennis, J. M., Tufo, H. M., and Thomas, S. J. (2008). A comparison of two shallow water models with non-conforming adaptive grids. *Mon. Wea. Rev.*, 136:1898–1922.

Taylor, M. (2011). Conservation of mass and energy for the moist atmospheric primitive equations on unstructured grids. In Lauritzen, P. H., Jablonowski, C., Taylor, M. A., and Nair., R. D., editors, *Numerical Techniques for Global Atmospheric Models*, volume 80 of *Lecture Notes in Computational Science and Engineeering*. Springer, Berlin, Heidelberg, New York.

Taylor, M., Loft, R., and Tribbia, J. (1998). Performance of a spectral element atmospheric model (SEAM) on the HP Exemplar SPP2000. Technical Report TN–439+EDD, NCAR.

Taylor, M., St.-Cyr, A., and Fournier, A. (2009). A non-oscillatory advection operator for the compatible spectral element method. In *Computational Science ICCS 2009 Part II, Lecture Notes in Computer Science 5545*, pages 273–282, Berlin / Heidelberg. Springer.

Taylor, M., Tribbia, J., and Iskandarani, M. (1997). The spectral element method for the shallow water equations on the sphere. *J. Comput. Phys.*, 130:92–108.

Taylor, M. A., Edwards, J., and St.-Cyr, A. (2008). Petascale atmospheric models for the community climate system model: New developments and evaluation of scalable dynamical cores. *J. Phys. Conf. Ser.*, 125(012023).

Taylor, M. A., Edwards, J., Thomas, S., and Nair, R. (2007). A mass and energy conserving spectral element atmospheric dynamical core on the cubed-sphere grid. *J. Phys. Conf. Ser.*, 78(012074).

Taylor, M. A., Evans, K., Hack, J., and Worley, P. (2011). Subcycled dynamics in the spectral community atmosphere model version 4. *J. Phys. Conf. Ser.* to appear.

Taylor, M. A. and Fournier, A. (2010). A compatible and conservative spectral element method on unstructured grids. *J. Comput. Phys.*, 229:5879– 5895.

Thomas, S. and Loft, R. (2000). Parallel semi-implicit spectral element methods for atmospheric general circulation models. *Journal of Scientific Computing*, 15:499–518.

Thomas, S. and Loft, R. (2002). Semi-implicit spectral element atmosphere model. *Journal of Scientific Computing*, 17:339–350.

Thomas, S. and Loft, R. (2005). The NCAR spectral element climate dynamical core: Semi-implicit Eulerian formulation. *J. Sci. Comput.*, 25:307–322.

van der Houwen, P. J. (1977). *Construction of Integration Formulas for Initial Value Problemsa*. North-Holland series in applied mathematics and mechanics, v. 19. North-Holland, Amsterdam, New York.

Wang, H., Tribbia, J. J., Baer, F., Fournier, A., and Taylor, M. A. (2007). A spectral element version of CAM2. *Monthly Weather Review*, 135:3825–3840.

Williamson, D. L. (2002). Time-split versus process-split coupling of parameterization and dynamical core. *Mon. Wea. Rev.*, 130:2024–2041.

Williamson, D. L. (2007). The evolution of dynamical cores for global atmospheric models. *J. Meteor. Soc. Japan*, 85B:242–269.

Williamson, D. L. and Olson, J. G. (1994). Climate simulations with a semi-lagrangian version of the NCAR community climate model. *Mon. Wea. Rev.*, 122:1594–1610.

Worley, P., Mirin, A., Drake, J., and Sawyer, W. (2006). Performance engineering in the community atmosphere model. *Journal of Physics: Conference Series*, 46:356–362. doi: 10.1088/1742-6596/46/1/050.

Worley, P. H. and Drake, J. B. (2005). Performance portability in the physical parameterizations of the community atmospheric model. *International Journal of High Performance Computing Applications*, 19:187 – 201.

Zerroukat, M., Wood, N., and Staniforth, A. (2005). A monotonic and positive-definite filter for a semi-lagrangian inherently conserving and efficient (slice) scheme. *Quart. J. R. Meteorol. Soc.*, 131(611):2923–2936.

9 List of Figures

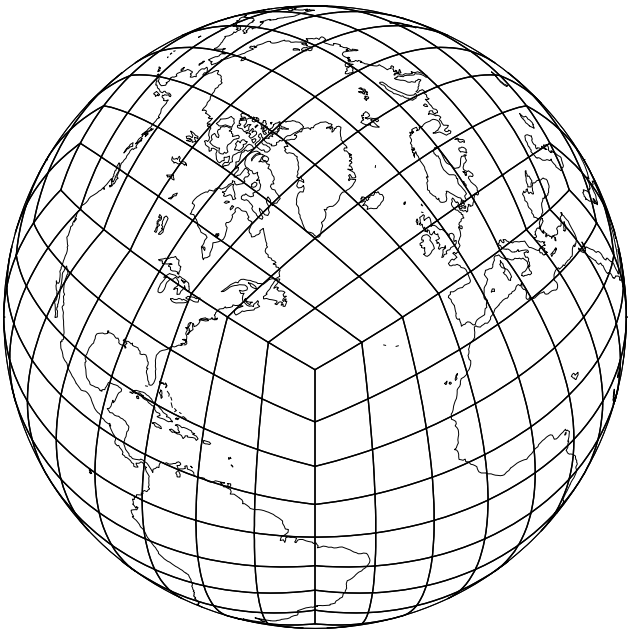

Figure 1: An example quadrilateral element grid for the sphere. An inscribed cube is projected on the surface of the sphere. The faces of the cubed-sphere are further subdivided to form a quadrilateral grid of the desired resolution. Coordinate lines from the gnomonic equal-angle projection are shown.

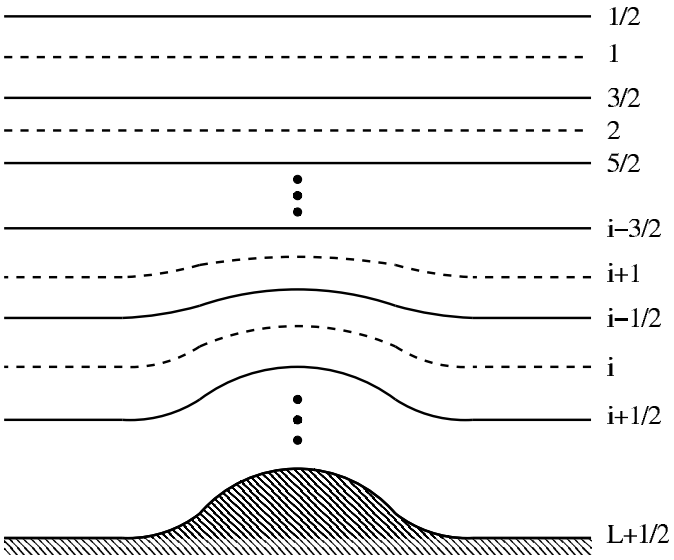

Figure 2: The terrain following η-coordinate layers and layer indexing. There are K layer midpoints denoted by $k = 1, 2, \ldots, K$ and $K + 1$ layer interfaces denoted by $k + \frac{1}{2}, k = 0, 1, \ldots, K$

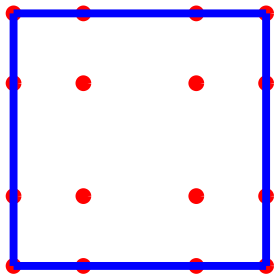

Figure 3: A 4×4 tensor product grid of GLL nodes used within each element, for a degree $d = 3$ discretization. Nodes on the boundary are shared by neighboring elements.

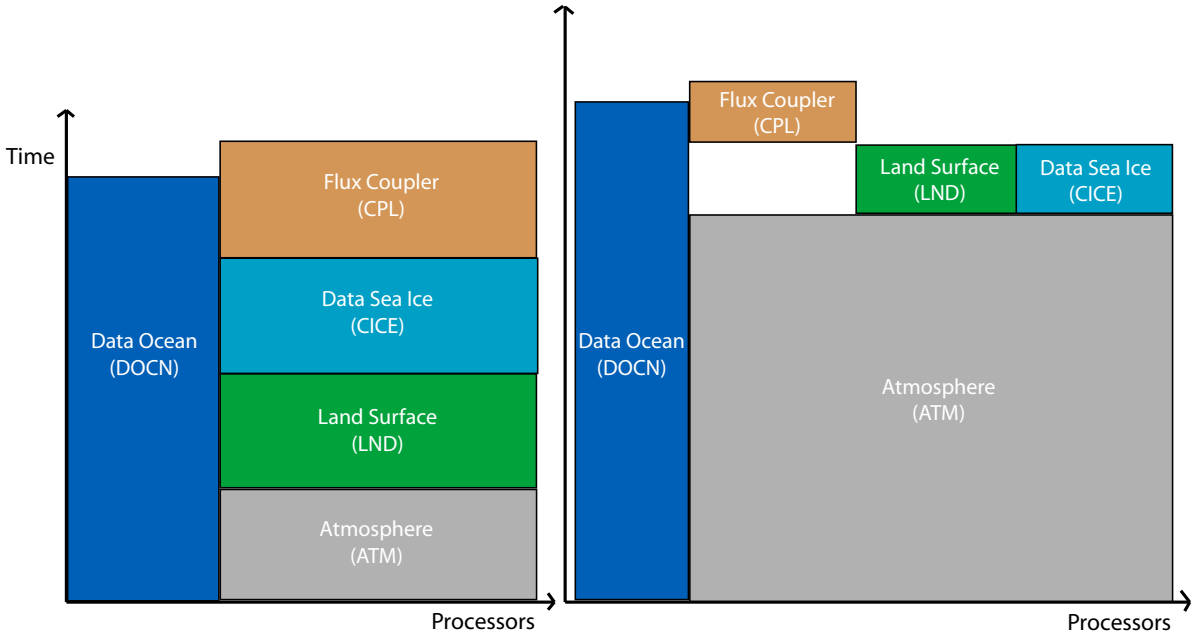

Figure 4: Two example processor layouts supported by the CESM. The position of each model component box in the figure shows (not exactly to scale) the wall clock time and number of processors used for a simulated model day. In both configurations, the ocean runs on an independent set of processors from the rest of the components. In a *stacked* configuration (left panel) all remaining processors are used by the atmosphere, land ice and flux coupler, and these components will then run sequentially. For high-resolution simulations running on thousands of processors, more efficient distributions (right panel) allow the land and sea ice components to run concurrently and reduce the memory requirements.

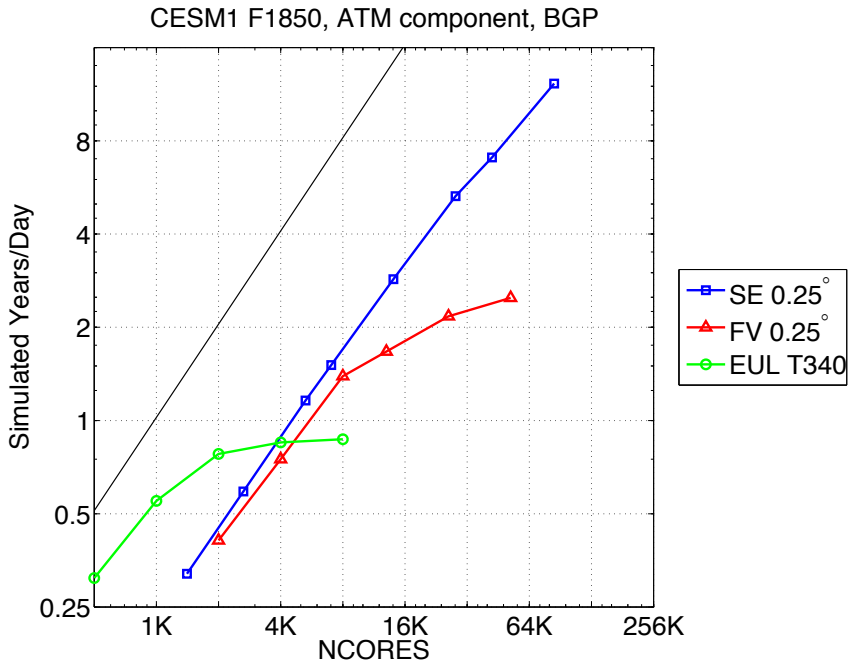

Figure 5: Performance of the CESM atmosphere component model on Intrepid (IBM BG/P) when using the CAM-SE, FV or EUL dynamical core, showing the simulated-years-per-day as a function of the number of processing cores. Atmosphere component times taken from a CESM time-slice simulation, coupling the atmosphere (at 0.25° or T341 resolution), the land model (0.25° resolution), and the sea ice and data ocean model (0.1°). The solid black line shows perfect parallel scalability. When using CAM-SE, the CESM achieves near perfect scalability down to one element per processor, running at 12.2 SYPD on 86,400 cores.

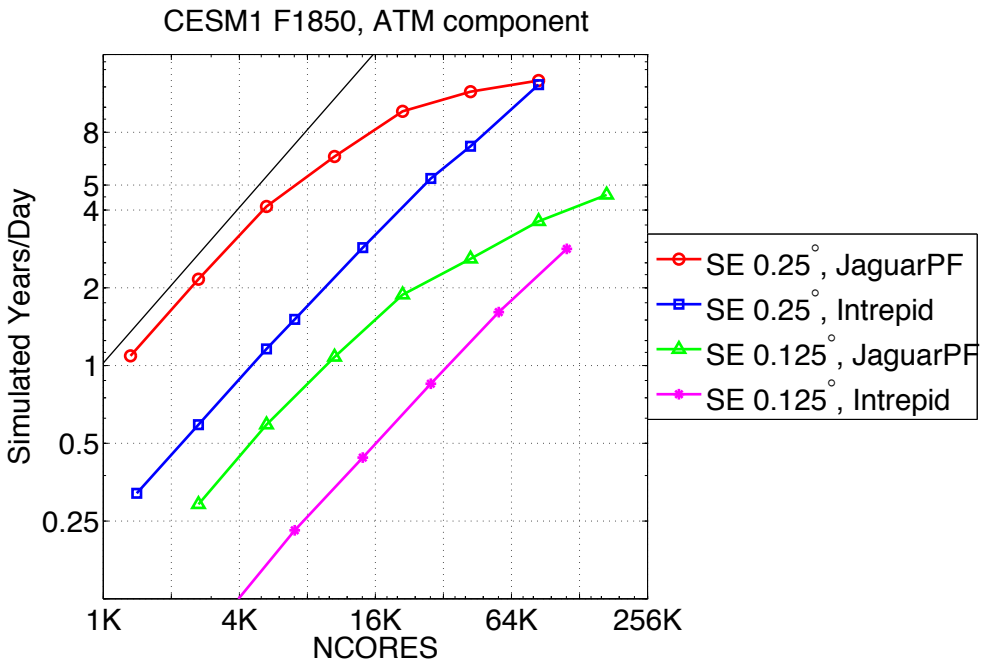

Figure 6: Performance of the CESM atmosphere component model when using CAM-SE atmosphere component, comparing the performance on Intrepid and JaguarPF at both 0.25° and 0.125° resolutions. Atmosphere component times taken from a CESM time-slice simulation, coupling the atmosphere (at 0.25° or 0.125° resolution), the land model (0.25° resolution), and the sea ice and data ocean model (0.1°). The solid black line shows perfect parallel scalability. At 0.125° resolution, the CESM achieves 4.6 SYPD on 172,800 cores of JaguarPF.

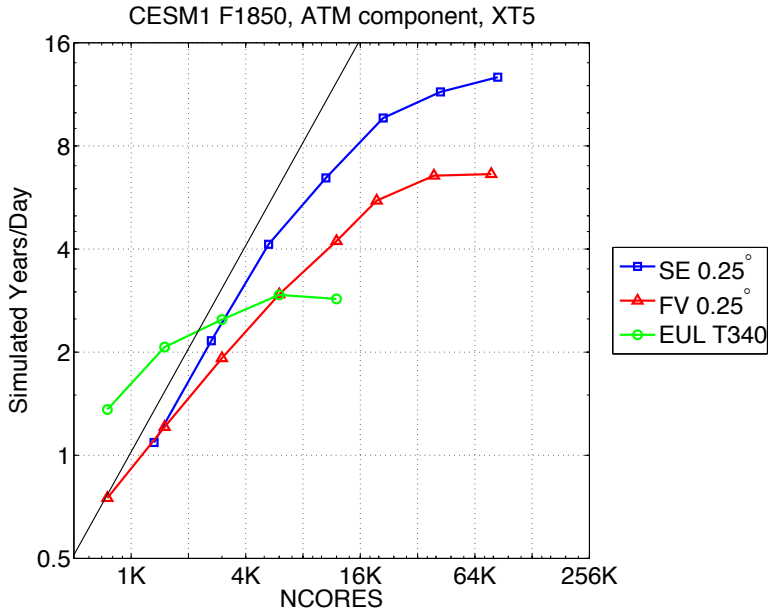

Figure 7: Performance of the CESM atmosphere component model as in Fig. 5, only on the JaguarPF (Cray XT5) system. The best performance is obtained with CAM-SE, running at 12.7 SYPD on 86,400 cores.

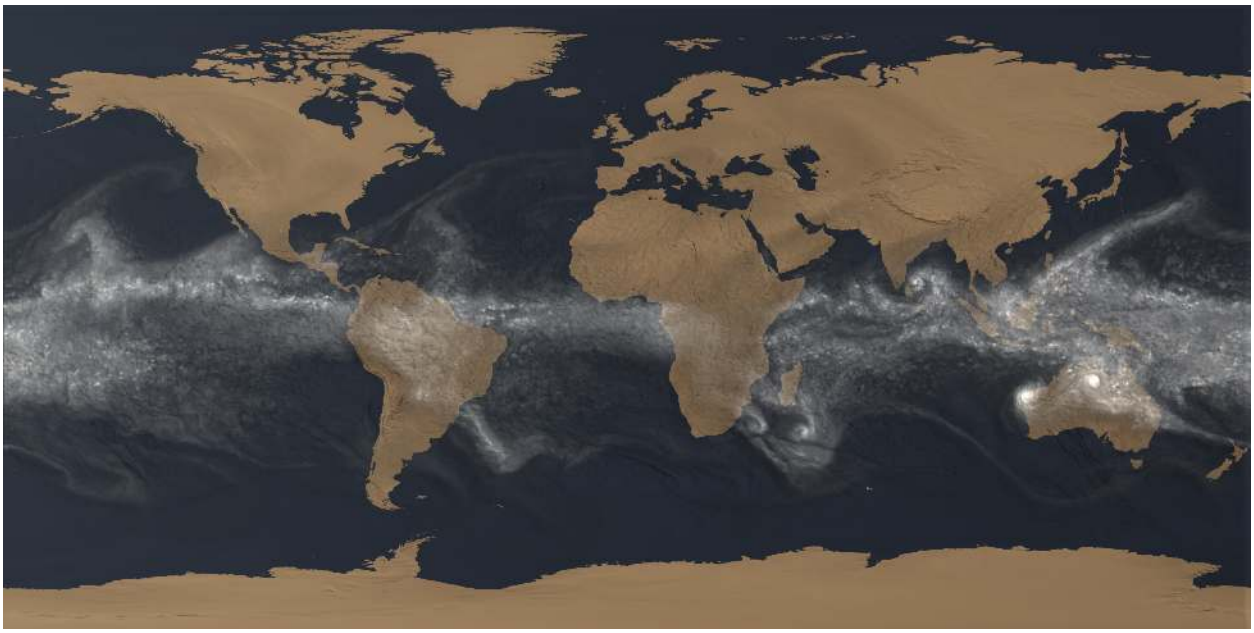

Figure 8: A snapshot of the column integrated precipitable water from a CESM time-slice simulation using CAM-SE at 0.125° resolution. The snapshot is from early January and shows the winter-time cyclone activity in the Indian ocean.