



RESEARCH ARTICLE

10.1029/2019MS001684

Special Section:

Community Earth System Model
version 2 (CESM2) Special
Collection

Key Points:

- A lower-resolution finite-volume physics grid is implemented into CAM-SE-CLAM, containing $\frac{5}{9}$ fewer grid columns than the dynamical core
- Grid imprinting from the spectral-element method is mitigated in regions with steep terrain, using the coarser physics grid
- The coarser physics grid does not degrade the effective resolution of the model

Correspondence to:

A. R. Herrington,
adam.herrington@stonybrook.edu

Citation:

Herrington A. R., Lauritzen, P. H.,
Reed, K. A., Goldhaber, S., &
Eaton, B. E. (2019). Exploring a
lower-resolution physics grid in
CAM-SE-CSLAM. *Journal of Advances
in Modeling Earth Systems*, 11,
1894–1916. [https://doi.org/10.1029/
2019MS001684](https://doi.org/10.1029/2019MS001684)

Received 8 MAR 2019

Accepted 23 MAY 2019

Accepted article online 29 MAY 2019

Published online 1 JUL 2019

©2019. The Authors.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

Exploring a Lower-Resolution Physics Grid in
CAM-SE-CSLAMAdam R. Herrington¹ , Peter H. Lauritzen² , Kevin A. Reed¹ , Steve Goldhaber²,
and Brian E. Eaton² ¹School of Marine and Atmospheric Sciences, Stony Brook University, Stony Brook, NY, USA, ²National Center for
Atmospheric Research, Boulder, CO, USA

Abstract This paper describes the implementation of a coarser-resolution physics grid into the Community Atmosphere Model (CAM), containing $\frac{5}{9}$ fewer grid columns than the dynamics grid. The dry dynamics is represented by the spectral element dynamical core, and tracer transport is computed using the Conservative Semi-Lagrangian Finite Volume Method (CAM-SE-CSLAM). Algorithms are presented that map fields between the dynamics and physics grids while maintaining numerical properties ideal for atmospheric simulations such as mass conservation and mixing ratio shape and linear-correlation preservation. The results of experiments using the lower-resolution physics grid are compared to the conventional method in which the physics and dynamical grids coincide. The lower-resolution physics grid provides a volume mean state to the physics computed from an equal sampling of the different types of nodal solutions arising in the spectral-element method and effectively mitigates grid imprinting in regions with steep topography. The impact of the coarser-resolution physics grid on the resolved scales of motion is analyzed in an aquaplanet configuration, across a range of dynamical core grid resolutions. The results suggest that the effective resolution of the model is not degraded through the use of a coarser-resolution physics grid. Since the physics makes up about half the computational cost of the conventional CAM-SE-CSLAM configuration, the coarser physics grid may allow for significant cost savings with little to no downside.

1. Introduction

Global atmospheric models fundamentally consist of two components. The dynamical core (*dynamics*), which numerically integrates the adiabatic equations of motion and tracer advection, and the physical parameterizations (*physics*), which compute the effects of diabatic and subgrid-scale processes (e.g., radiative transfer and moist convection) on the grid scale. More out of convenience than anything else, the physics are evaluated on the dynamics grid; that is, the physics and dynamics grids coincide. From linear stability and accuracy analysis of numerical methods, it is a common result that the shortest simulated wavelengths are not accurately represented by the dynamical core. Additionally, simulated downscale cascades result in an unrealistic collection of energy and/or enstrophy near the truncation scale, which may be observed from kinetic energy spectra in model simulations (Skamarock, 2011). Some form of dissipation must be incorporated into models to mitigate these numerical artifacts near the grid scale (Jablonowski & Williamson, 2011). The unrealistic nature of the grid scale led Lander and Hoskins (1997) to speculate whether the physics should be evaluated on a grid that is more reflective of the scales actually resolved by the dynamical core.

Exploring the impact of different physics grid resolutions has so far been limited to models employing the spectral transform method (Lander & Hoskins, 1997; Wedi, 2014; Williamson, 1999). Lander and Hoskins (1997) argued that passing underresolved states to the physics may be especially problematic in spectral transform models, since the physics are evaluated on a latitude-longitude transform grid, and contains more degrees of freedom than the spectral representation to prevent aliasing of quadratic quantities. However, Lander and Hoskins (1997) found that the spectral truncation of the physics tendencies damps errors that may result from passing an underresolved state to the physics, although the extent to which these errors may still be present in the model is difficult to address.

Another class of spectral transform models evaluate the quadratic terms using semi-Lagrangian methods, which are implicitly diffusive, relaxing constraints on the resolution of the transform grid. Wedi (2014) experimented with different transform grid resolutions and concluded that the standard high-resolution

quadratic grid actually improves forecast skill over the use of a lower-resolution transform grid. They suggest that increasing the resolution of the transform grid simulates a kind of subgrid variability on the spectral state, which is thought to be underrepresented in global atmospheric models (Shutts, 2005). This is in principle the purpose of “superparameterization,” in which a cloud-resolving model is embedded in each grid cell to approximate subgrid variability and improves both diurnal and subseasonal variability in the model (Randall et al., 2003).

After the physics tendencies are transformed into spectral space, it is possible to truncate the tendencies at any particular wave number in global spectral transform models. Williamson (1999) conducted a pair of convergence tests using a spectral transform model: a conventional convergence test and one in which the spectral truncation of the physics tendencies is held fixed and the resolution of the dynamical core increased. In contrast to the realistic weather forecasts of Wedi (2014), Williamson (1999) ran their model to equilibrium in an idealized climate configuration. When the physics and dynamics resolutions increase together, as in more typical convergence studies, the strength of the Hadley cell increases monotonically with resolution. This sensitivity of Hadley cell strength to horizontal resolution is a common result of global models at hydrostatic resolutions (see Herrington & Reed, 2017, and references therein). But with the truncation wave number of physics tendencies held fixed, the Hadley cell showed very little sensitivity to dynamical core resolution, resembling the solution for which the dynamics truncation wave number is equal to that of the lower-resolution physics. Herrington and Reed (2017) speculated that these results suggest the scales of motion resolved by the dynamical core may be aliased to the lower-resolution physics.

Global spectral transform models, while remarkably efficient at small processor counts, do not scale well on massively parallel systems. High-order Galerkin methods are becoming increasingly popular in climate and weather applications due to their high-parallel efficiency, high-processor efficiency, high-order accuracy (for smooth problems), and geometric flexibility facilitating mesh-refinement applications (e.g., Brdar et al., 2013; Giraldo & Restelli, 2008; Nair et al., 2009, and the Energy Exascale Earth System Model; <https://e3sm.org/>). High-resolution climate simulations with NCAR's Community Atmosphere Model (CAM; Neale et al., 2012) are typically performed using a continuous Galerkin dynamical core referred to as CAM-SE (CAM Spectral Elements; Dennis et al., 2012; Lauritzen et al., 2018; Taylor et al., 2008). CAM-SE may be optionally coupled to a conservative, semi-Lagrangian tracer advection scheme for accelerated multitracer transport (CAM-SE-CSLAM; Lauritzen et al., 2017). Tracer advection then evolves on an entirely separate, finite-volume grid which contains the same degrees of freedom as CAM-SE's quadrature node grid.

Element-based Galerkin methods are susceptible to grid-imprinting and may need to be considered when contemplating a particular physics grid (Herrington et al., 2018, hereafter referred to as H18). Grid imprinting manifests at the element boundaries, since the global basis is least smooth (C^0 ; all derivatives are discontinuous) for quadrature nodes lying on the element boundaries, and the gradients (e.g., pressure gradients) are systematically tighter producing local extremes. Through computing the physics tendencies at the nodal points, element boundary extrema is also observed in the physics tendencies.

H18 has shown that through evaluating the physics on the finite-volume tracer advection grid in CAM-SE-CSLAM, element boundary errors are substantially reduced, although still problematic in regions of steep terrain, at low latitudes. Through integrating CAM-SE's basis functions over the control volumes of the finite-volume grid, element boundary extrema is additionally weighted by the C^∞ solutions (i.e., the basis representation is infinitely smooth and all derivatives are continuous) that characterize the interior of the element, and the state is smoother. Additionally, in defining an area averaged state, the finite-volume physics grid is made consistent with assumptions inherent to the physics and is more appropriate for coupling to other model components (e.g., the land model), which is typically performed using finite-volume-based mapping algorithms.

The CAM-SE-CSLAM finite-volume grid is defined through dividing the elements of CAM-SE's gnomonic cubed-sphere grid with equally spaced, equi-angular coordinate lines parallel to the equi-angular element boundaries, such that there are 3×3 control volumes per element (hereafter referred to as *pg3*; see Figure 1). While the physics grid in H18 is *pg3*, that is, the physics and dynamics grids have the same degrees of freedom, the control volumes in *pg3* encompass a region of the element in which their proximity to the element boundaries is not equal. Therefore, not every control volume in an element has the same smoothness properties. This may be avoided through defining a physics grid in which the elements are instead divided into 2×2 control volumes (hereafter referred to as *pg2*; see Figure 1). The control volumes of the *pg2* grid all

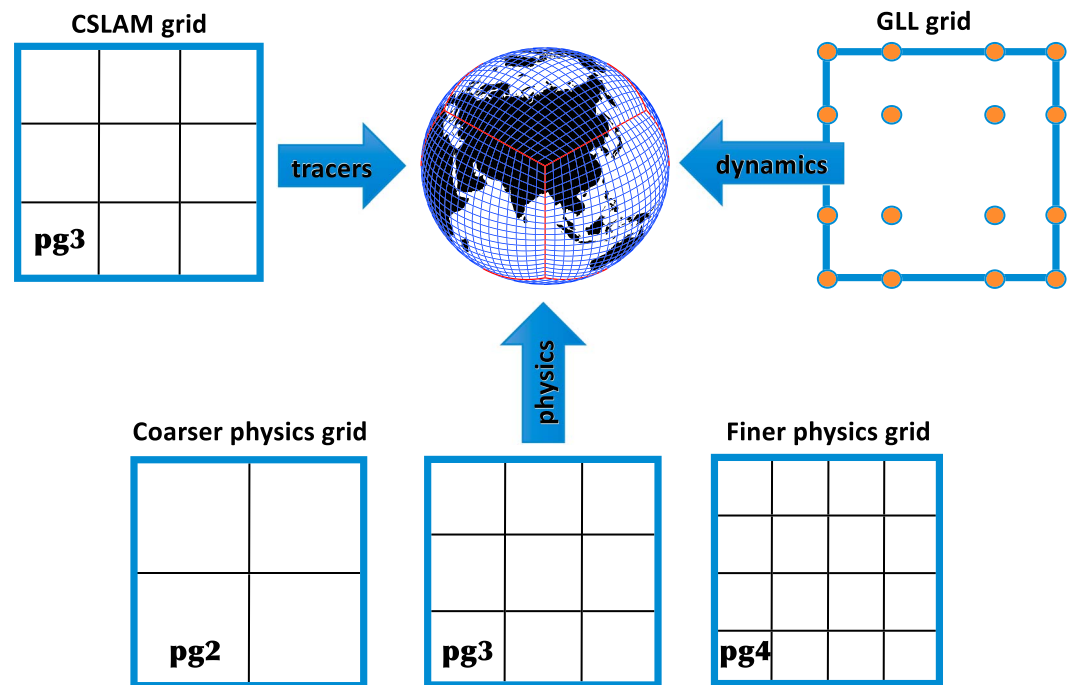


Figure 1. An overview of the different grids in CAM-SE-CSLAM.

have the same proximity to the element boundaries and should mitigate the element boundary noise that remains in the *pg3* grid and is shown in H18.

In this study, we test the hypothesis that the coarser, *pg2* physics grid is effective at reducing spurious noise at element boundaries, particularly over regions of rough topography. In addition, the recent trend toward running models at ever higher resolutions is an almost prohibitive computational burden. As the physics are responsible for over half of the computational cost in CAM-SE (Lauritzen et al., 2018), the improvement in computational performance using a coarser-resolution physics grid is potentially significant. However, any advantages of using a coarser physics grid need to be weighed against any potential reduction in simulation quality, for example, possible aliasing of the resolved scales of motion by the coarser grid, as suggested by the results of Williamson (1999). Section 2 describes the implementation of the *pg2* grid into CAM-SE-CSLAM and the idealized model configurations used throughout this study. Section 3 provides results of model simulations designed to test the mapping algorithms, identify grid imprinting and assess the range of scales resolved by the model. Section 4 provides a discussion of the results and conclusions.

2. Methods

Separating dynamics, tracer, and physics grids introduces the added complexity of having to map the state from dynamics and tracer grids to the physics grid and mapping physics tracer increments back to the tracer grid and physics increments needed by the dynamical core to the dynamics grid (see Figure 1). The dynamics grid in the case of CAM-SE-CSLAM refers to the Gauss-Lobatto-Legendre (GLL) quadrature nodes used by the spectral-element method to solve the momentum equations for the momentum vector (u, v), thermodynamics equation for temperature (T), continuity equation for dry air mass ($\frac{1}{g}p$), and continuity equations for water vapor and thermodynamically and inertially active condensates (see, e.g., Lauritzen et al., 2018, for details). By tracer grid we refer to the *pg3* grid on which CSLAM performs tracer transport of water vapor, condensates, and other tracers. Although water vapor and condensates are being advected by the CSLAM scheme on the *pg3* grid, these quantities are also needed on the GLL grid for the momentum equations and thermodynamic equation. Transport of water variables is also performed by the spectral-element method on the GLL grid. To avoid decoupling of water species on the CSLAM and GLL grids, the GLL water species are overwritten by the CSLAM values every physics time step. This is explained in detail in H18.




Figure 2. Indices notation for (a) the $pg2$ grid, (b) the $pg3$ grid, and (c) their exchange grid.

Similarly to the CAM-SE-CSLAM $pg3$ configuration, the dynamics state (momentum vector, temperature, dry pressure) must be mapped from the GLL grid to the physics grid. Exactly the same algorithms as used in the $pg3$ configuration apply; that is, momentum components are interpolated by evaluating the internal Lagrange basis functions (used in the spectral-element method) at the equi-angular (gnomonic) center of the $pg2$ cells and the Lagrange basis function representations of temperature and pressure are integrated over the $pg2$ control volumes. See H18 for details.

As compared to the $pg3$ configuration, the extra complication with the $pg2$ setup is that the tracer grid does not coincide with the physics grid; that is, the tracer state needs to be mapped from the CSLAM grid ($pg3$) to the physics grid ($pg2$), and tracer increments computed by physics must be mapped from the physics grid back to the CSLAM grid. In order to describe the mapping algorithms between the grids some notation needs to be introduced.

The mapping algorithms are applied to each element Ω (with spherical area $\Delta\Omega$) so without loss of generality consider one element. Let $\Delta A_k^{(pg2)}$ and $\Delta A_\ell^{(pg3)}$ be the spherical area of the physics grid cell $A_k^{(pg2)}$ and CSLAM control volume $A_\ell^{(pg3)}$, respectively. The physics grid cells and CSLAM cells, respectively, span the element, Ω , without gaps or overlaps

$$\bigcup_{k=1}^{nphys} A_k^{(pg2)} = \Omega \quad \text{and} \quad A_k^{(pg2)} \cap A_\ell^{(pg2)} = \emptyset \quad \forall k \neq \ell, \quad (1)$$

$$\bigcup_{k=1}^{nc} A_k^{(pg3)} = \Omega \quad \text{and} \quad A_k^{(pg3)} \cap A_\ell^{(pg3)} = \emptyset \quad \forall k \neq \ell, \quad (2)$$

where $nc = 3$ is the CSLAM grid resolution parameter and $nphys = 2$ is the physics grid resolution parameter (following the Fortran code base), although the methods described here are valid for any arbitrary integer $nphys$ (e.g., $nphys = 4$ is shown in Figure 1). The overlap areas between the k th physics grid cell and ℓ th CSLAM cell are denoted

$$A_{k\ell} = A_k^{(pg2)} \cap A_\ell^{(pg3)}, \quad (3)$$

(see Figure 2) so that

$$A_k^{(pg2)} = \bigcup_{\ell=1}^{nc} A_{k\ell}. \quad (4)$$

This overlap grid is also referred to as the *exchange grid*.

2.1. Mapping Tracers From $A^{(pg3)}$ to $A^{(pg2)}$ (CSLAM to Physics Grid)

The CSLAM and physics grids are both finite-volume grids so existing CSLAM technology can be used to map the tracer state from CSLAM to physics grid. That is, compute a high-order shape-preserving reconstruction of mixing ratio m and dry air mass $\frac{1}{g}\Delta p$ per unit area in each CSLAM control volume and integrate those reconstruction functions over the overlap areas (Lauritzen et al., 2010; Nair & Lauritzen, 2010). This algorithm retains the properties of CSLAM: inherent mass conservation, consistency (constant mixing ratio is preserved), mixing ratio shape preservation, and linear-correlation preservation.

Denote the known cell-averaged values of dry pressure-level thickness and mixing ratio as $\overline{\Delta p}^{(pg3)}$ and $\overline{m}^{(pg3)}$, respectively. We consider a particular layer and for simplicity drop the layer subscript. The same procedure is applied to each layer in a column. The unknowns we would like to compute are the cell-averaged values of the same quantities on the physics grid, $\overline{\Delta p}^{(pg2)}$ and $\overline{m}^{(pg2)}$, respectively. The dry pressure level thickness integrated over the k th physics grid cell is given by

$$\overline{\Delta p}_k^{(pg2)} = \frac{1}{\Delta A_k^{(pg2)}} \sum_{\ell=1}^{nc^2} \langle \delta p \rangle_{k\ell}, \quad (5)$$

where $\langle \delta p \rangle_{k\ell}$ is the dry mass in a layer over overlap area $A_{k\ell}$. It is computed by integrating a high-order (2-D polynomial of degree 2) reconstruction of pressure-level thickness in each CSLAM cell over the overlap area $A_{k\ell}$

$$\langle \delta p \rangle_{k\ell} = \int_{A_{k\ell}} \left[\sum_{i+j \leq 2} \mathcal{P}_\ell^{(ij)} x^i y^j \right] dA. \quad (6)$$

The reconstruction coefficients $\mathcal{P}_\ell^{(ij)}$ in CSLAM cell ℓ are computed from the cell average pressure level thicknesses on the CSLAM grid $\overline{\Delta p}^{(pg3)}$, and the numerical integration over overlap areas is done by line integrals. The details of that are given in Lauritzen et al. (2010) and not repeated here.

The average tracer mass per unit area on the physics grid is given by

$$\overline{m \Delta p}_k^{(pg2)} = \frac{1}{\Delta A_k^{(pg2)}} \sum_{\ell=1}^{nc^2} \langle m \delta p \rangle_{k\ell}, \quad (7)$$

where $\langle m \delta p \rangle_{k\ell}$ is the tracer mass over $A_{k\ell}$ resulting from integrating a high-order reconstruction of Δp and m combined using the approach outlined in Appendix B of Nair and Lauritzen (2010) over the overlap area $A_{k\ell}$

$$\langle m \delta p \rangle_{k\ell} = \int_{A_{k\ell}} \left[\overline{\Delta p}_\ell^{(pg3)} \sum_{i+j \leq 2} \mathcal{M}_\ell^{(ij)} x^i y^j + \overline{m}_\ell^{(pg3)} \sum_{i+j \leq 2} \tilde{\mathcal{P}}_\ell^{(ij)} x^i y^j \right] dA, \quad (8)$$

where $\tilde{\mathcal{P}}_\ell^{(00)} = \mathcal{P}_\ell^{(00)} - \overline{\Delta p}_\ell^{(pg3)}$ and $\tilde{\mathcal{P}}_\ell^{(ij)} = \mathcal{P}_\ell^{(ij)}$ for $i, j > 0$, and $\mathcal{M}_\ell^{(ij)}$ are the reconstruction coefficients for the mixing ratio in CSLAM cell $A_\ell^{(pg3)}$. A shape-preserving limiter is applied to the reconstruction of mixing ratio m (Barth & Jespersen, 1989) and not Δp . This way of combining the reconstruction function for Δp and m in (8) ensures that a constant mixing ratio is preserved (consistency), tracer mass is conserved, linear correlations are preserved, and tracer shape preservation is retained. The mixing ratio on the physics grid is then

$$\overline{m}_k^{(pg2)} = \frac{\overline{(m \Delta p)}_k^{(pg2)}}{\overline{\Delta p}_k^{(pg2)}}, \quad (9)$$

where $\overline{\Delta p}_k^{(pg2)}$ is given in (5).

Perhaps surprisingly a much more challenging problem is to map tracer increments (or state) from the physics grid to the CSLAM grid while retaining important properties such as mass conservation, consistency, and correlation preservation. Why this mapping problem is challenging is explained in detail in section 2.2.1 after having defined important properties for mapping physics increments/tendencies.

2.2. Mapping Tracer Increments From $A^{(pg2)}$ to $A^{(pg3)}$ (Physics to CSLAM Grid)

The increments from the parameterizations are computed on the physics grid. The tracer increment in physics grid cell k is denoted $\overline{f}_k^{(pg2)}$ so that the updated mixing ratio on the physics grid is $\overline{m}_k^{(pg2)} + \overline{f}_k^{(pg2)}$. The problem is how to map $\overline{f}_k^{(pg2)}$ to the CSLAM control volumes, to obtain $\overline{f}^{(pg3)}$, satisfying the following constraints:

1. *Local mass conservation*: At a minimum total physics mass forcing on an element computed on the physics grid should equal the element physics mass forcing on the CSLAM grid

$$\overline{f}_k^{(pg2)} \overline{\Delta p}_k^{(pg2)} \Delta A_k^{(pg2)} = \sum_{\ell=1}^{nc^2} \left[\overline{\Delta p}_{\ell}^{(pg3)} \overline{f}_{\ell}^{(pg3)} \Delta A_{k\ell} \right], \quad (10)$$

where $\overline{\Delta p}_k^{(pg2)}$ is the pressure level thickness in physics grid cell k and similarly for $\overline{\Delta p}^{(pg3)}$. We enforce a more local constraint in which only mass increments overlapping with a particular CSLAM cell contributes to the mass increment in that CSLAM cell.

2. *Local shape preservation in mixing ratio*: The increments mapped to the CSLAM grid and added to the previous CSLAM state should not produce values smaller than the updated physics grid mixing ratios, $\overline{m}_k^{(pg2)} + \overline{f}_k^{(pg2)}$, or values smaller than the existing CSLAM mixing ratios that overlap with physics grid cell A_{ℓ}

$$\overline{m}_{\ell}^{(pg3)} + \overline{f}_{\ell}^{(pg3)} \geq \overline{m}_k^{(\min)} = \min \left(\overline{m}_k^{(pg2)} + \overline{f}_k^{(pg2)}, \{ \overline{m}_{k\ell} | \ell = 1, nc^2 \} \right), \quad (11)$$

where

$$\overline{m}_{k\ell} = \frac{\langle m \delta p_{k\ell} \rangle}{\langle \delta p_{k\ell} \rangle}. \quad (12)$$

The numerator and denominator in (12) are defined in (6) and (8), respectively. In particular, this means that an increment, when mapped to the $pg3$ grid, should not drive the state negative (described in detail below as the “negativity” problem).

A similar definition apply for maxima

$$\overline{m}_{\ell}^{(pg3)} + \overline{f}_{\ell}^{(pg3)} \leq \overline{m}_k^{(\max)} = \max \left(\overline{m}_k^{(pg2)} + \overline{f}_k^{(pg2)}, \{ \overline{m}_{k\ell} | \ell = 1, nc^2 \} \right). \quad (13)$$

3. *Linear correlation preservation*: The physics forcing must not disrupt linear tracer correlation between species on the CSLAM grid (see, e.g., Lauritzen & Thuburn, 2012); that is, if two tracers are linearly correlated and the physics increment preserves linear correlations on the physics grid, then the tracer increment on the CSLAM grid must not disrupt linear correlations.
4. *Consistency*: A nonzero constant mixing ratio increment from physics, $cnst$, on the physics grid, $\overline{f}_k^{(pg2)} = cnst \forall k$, must result in the same (constant) forcing on the CSLAM grid, $\overline{f}_{\ell}^{(pg3)} = \overline{f}_k^{(pg2)} = cnst \forall \ell$.

To motivate the algorithm that will simultaneously satisfy 1–4, it is informative to discuss how “standard” mapping algorithms will violate one or more of the constraints.

2.2.1. Why “Conventional” Conservative Remapping Will Not Work

It is helpful to analyze in detail why conventional remapping cannot satisfy properties 1–4 above. Assume that one remaps the mass increments in exactly the same way as the mapping of mixing ratio state from the CSLAM grid to the physics grid described in section 2.1. That is, replace m with f and map from physics grid to the CSLAM grid instead of the other way around. Denote the mapped mass increment $\overline{f \Delta p}^{(pg3)}$, and due to the properties of the mapping algorithm the mass increment is conserved, linear correlation between mass increments are conserved, and shape in mass increment is preserved. The problems arise when converting from mass to mixing ratio.

2.2.1.1. Conserve Mass but Not Consistency

If ones uses the known pressure-level thickness on the CSLAM grid $\overline{\Delta p}_k^{(pg3)}$ to convert from mass increment to mixing-ratio increment

$$\overline{m}_k^{(pg3)} = \frac{\overline{f \Delta p}_k^{(pg3)}}{\overline{\Delta p}_k^{(pg3)}}, \quad (14)$$

a constant mixing ratio increment is not conserved. Basically, the constant increment mapped to the CSLAM grid and converted to mixing ratio increment through (14) will, rather than being constant, reflect the spurious discrepancy between $\overline{\Delta p}_k^{(pg3)}$ and $\overline{\Delta p}_k^{(pg3)}$, where $\overline{\Delta p}_k^{(pg3)}$ is the pressure-level thickness mapped from the

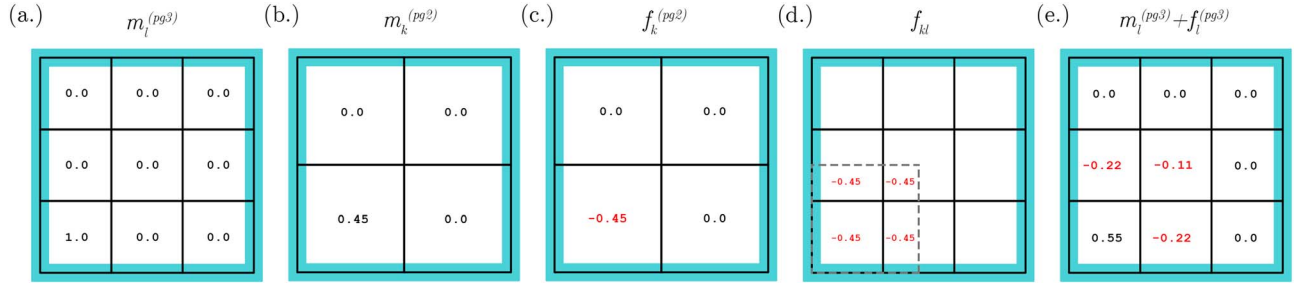


Figure 3. Schematic illustration of the “negativity problem” in a single element. (a) Initial CSLAM tracer values, (b) mapped to $pg2$, (c) produces a tracer increment on $pg2$, (d) with negative increments on the exchange grid overlying CSLAM cells in (a) that were initially 0 and (e) driving those mixing ratios negative.

$pg2$ grid to the $pg3$ grid. That said, mass will be conserved since the dynamical core state has $\overline{\Delta p_k^{(pg3)}}$ (unless the increment drives the mixing ratio negative—described in detail below).

2.2.1.2. Consistent but Not Mass-Conserving

Rather than converting to mixing ratio using $\overline{\Delta p_k^{(pg3)}}$, a constant increment can be preserved by using

$$\overline{m_k^{(pg3)}} = \frac{\overline{f \Delta p_k^{(pg3)}}}{\overline{\Delta p_k^{(pg3)}}}, \quad (15)$$

instead. But now mass conservation is lost since, again, $\overline{\Delta p_k^{(pg2)}} \neq \overline{\Delta p_k^{(pg3)}}$. This issue is similar to the mass-wind inconsistency found in specified dynamics applications (e.g., Jöckel et al., 2001; Lauritzen et al., 2011).

2.2.1.3. The Negativity Problem and Linear Correlations

Even if one could derive a reversible map for mapping $\overline{\Delta p_k^{(pg2)}}$ from the physics grid to the CSLAM grid, there could still be problems if the increment drives the mixing ratios negative (or overshooting occurs) on the CSLAM grid. This can easily happen for tracers, such as cloud liquid amount and cloud ice amount, that are 0 in most of the domain and nonzero in localized areas/points (where there are clouds). We refer to this as the “negativity problem.” This problem is depicted schematically in Figure 3. Consider a single element of CSLAM control volumes, containing only a single cell with mixing ratio 1.0, and 0.0 everywhere else ($\overline{m_k^{(pg3)}}$; Figure 3a). The mixing ratios are mapped to the $pg2$ grid using, for simplicity, the piecewise constant method where a constant value inside the $pg2$ cells is used during the integration over overlap cells ($\overline{m_k^{(pg2)}}$; Figure 3b). Now consider the case in which physics removes all the mass from the physics cell k : $\overline{f_k^{(pg2)}} = -\overline{m_k^{(pg2)}}$ (Figure 3c). The tracer increment is mapped from $pg2$ to $pg3$ using the piecewise constant method. Some of the nonzero increments are now in overlap areas where the original CSLAM grid cells have mixing ratio 0 ($\overline{f_{kl}}$; Figure 3d), and hence, the state is driven negative when adding the overlap increment to the CSLAM state (Figure 3e). This is referred to as the negativity problem although it can also happen for maxima.

The negativity issue could be avoided if one remaps the physics updated state instead of mapping increments/tendencies. In that case a shape-preserving filter will make sure that the state on the CSLAM grid is not negative (and does not overshoot). That said, if physics does not change the state and it is mapped back to the CSLAM grid then spurious tendencies (proportional to the errors introduced by mapping state from the CSLAM grid to the physics grid and back again) are introduced. Hence, it is advantageous to map increments/tendencies since any reasonable algorithm will preserve a zero function.

As illustrated above a standard remapping method will *not* simultaneously satisfy 1–4, and hence, a new algorithm has been derived.

2.3. New Tendency Mapping Algorithm

The problem is how to map the mass increment on the physics grid, $\overline{f_k^{(pg2)}} \Delta A^{(pg2)}$, to the CSLAM cells that overlap with $\Delta A^{(pg2)}$. To maintain shape preservation and linear correlations and to avoid the negativity problem locally, it is advantageous to define a mass excess function on the exchange grid $\Delta m_{kl}^{(excess)}$. It is basically

the maximum amount of mixing ratio that can be removed (in the case $\bar{f}^{(pg2)} < 0$) without producing new minima in the exchange grid mixing ratio $m_{k\ell}$

$$\Delta m_{k\ell}^{(\text{excess})} = \bar{m}_{k\ell} - \bar{m}_k^{(\min)}, \quad (16)$$

where $\bar{m}_{k\ell}$ is defined in (12). So the maximum amount of mass that we can be removed from the exchange grid cells that span physics grid cell A_k without violating the shape-preservation constraint ((11) and (13)) is

$$\sum_{\ell} \Delta m_{k\ell}^{(\text{excess})} \bar{\Delta p}_{k\ell} \delta A_{k\ell}. \quad (17)$$

If physics is designed not to remove more mass than available in A_k (which should be the case for a carefully designed physics package), then it is guaranteed that

$$\sum_{\ell} \Delta m_{k\ell}^{(\text{excess})} \bar{\Delta p}_{k\ell} \delta A_{k\ell} \geq \bar{f}^{(pg2)} \Delta p_k \Delta A_k^{(pg2)}. \quad (18)$$

We distribute the physics mass forcing (assuming $\bar{f}^{(pg2)} < 0$) according to the mass excess in each overlap area by solving this equation for γ_k

$$\Delta A_k^{(pg2)} \bar{\Delta p}_k^{(pg2)} \bar{f}^{(pg2)} = \gamma_k \sum_{\ell} \left[\Delta m_{k\ell}^{(\text{excess})} \bar{\Delta p}_{k\ell} \delta A_{k\ell} \right], \quad (19)$$

and add mass increment (which in this case is negative)

$$\gamma_k \Delta m_{k\ell}^{(\text{excess})} \bar{\Delta p}_{k\ell} \delta A_{k\ell}, \quad (20)$$

to the ℓ th CSLAM cell state $\bar{m}^{(pg3)} \bar{\Delta p}_{\ell}^{(pg3)} \Delta A_{\ell}^{(pg3)}$. This process is repeated for all physics cells. Note that this problem is well posed, that is, $\gamma_k > 0$, since physics will not remove more mass than is locally available (18). The way in which the mass forcing is distributed to the CSLAM cells using the excess function insures that the negativity problem is avoided. Mass is conserved by design, and shape preservation is obtained by using the excess function.

If the physics increment is positive (assuming $\bar{f}^{(pg2)} > 0$), we define a “lack” function

$$\Delta m_{k\ell}^{(\text{lack})} = \bar{m}_{k\ell} - \bar{m}_k^{(\max)}, \quad (21)$$

and solve

$$\bar{\Delta p}_k^{(pg2)} \bar{f}^{(pg2)} \Delta A_k^{(pg2)} = \gamma_k \sum_{\ell} \left[\Delta m_{k\ell}^{(\text{lack})} \bar{\Delta p}_{k\ell} \delta A_{k\ell} \right], \quad (22)$$

for γ_k and follow the same procedure as for mass excess. Since positive and negative forcing is treated in exactly the same way, linear correlations are preserved. Note how the definition of the excess/lack function insures linear correlation preservation; for example, if one would prevent negative values and not do anything about overshoots, then linear correlations would not be preserved since the minima and maxima are not treated in the same way.

While the above algorithm satisfies properties 1–4 in section 2.2, it is not a high-order algorithm in terms of formal accuracy. This is illustrated in Figure 4 where a smooth analytical tendency (approximate spherical harmonic of order 32 and azimuthal wave number 16; Jones, 1999)

$$f^{(pg2)} = \frac{1}{2} + \frac{1}{2} \cos(16\lambda) \sin(2\theta)^{16}, \quad (23)$$

where (λ, θ) is latitude-longitude and is mapped from $pg2$ to $pg3$ grid using this algorithm assuming $m_{\ell}^{(pg3)} = 0, \forall \ell$. The errors in the mapping are not always aligned with large gradients in the analytical function as would be expected for a “traditional” interpolation algorithm. The errors are maximum on the order of 60%. To reduce errors, we therefore perform a higher-order pre-allocation of tendencies that is not mass conserving but satisfies properties 2, 3, and 4 in section 2.2.

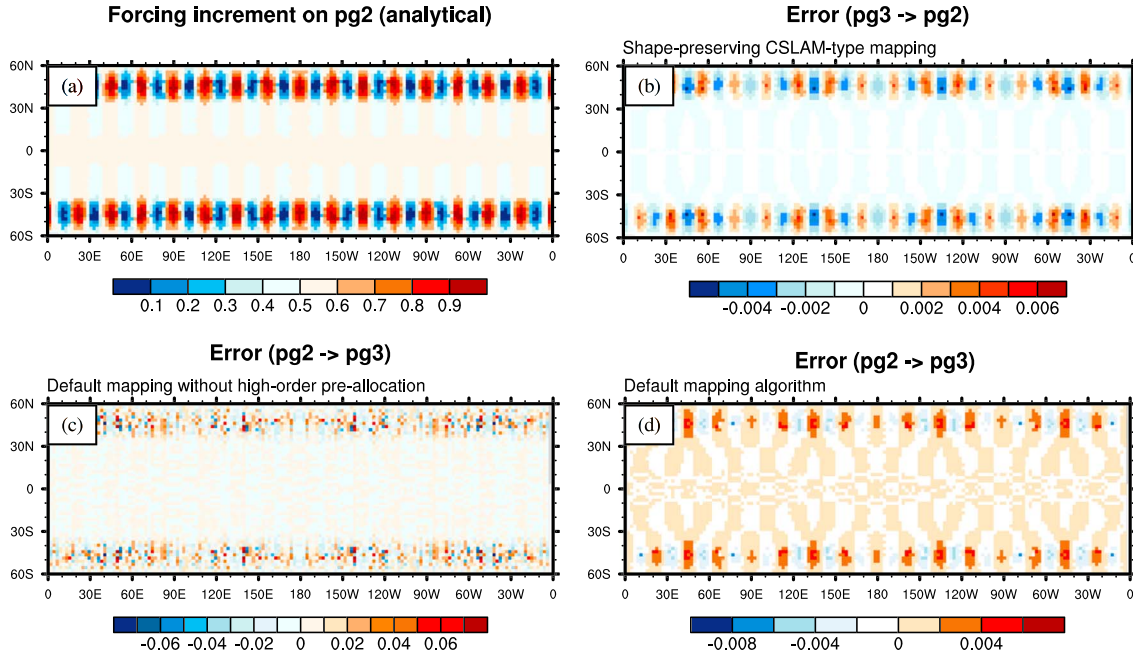


Figure 4. (a) Smooth analytical function (approximate spherical harmonic of order 32 and azimuthal wave number 16 referred to as Y_{32}^{16} ; Jones, 1999) plotted using raster contours on the $pg2$ grid. Remaining plots show errors (mapped field minus analytical function evaluated at centers of grid cells). (c and d) A Y_{32}^{16} tracer tendency is mapped from $pg2$ to $pg3$ grid using the default mapping algorithm without high-order pre-allocation algorithm (c) and the default algorithm that uses high-order pre-allocation (d). The Figure shows that the high-order pre-allocation algorithm used to map tendencies from the $pg2$ to $pg3$ grid drastically reduces mapping errors for a smooth function (compare plots (c) and (d)). (b) For comparison the errors for the high-order CSLAM algorithm used to map state from $pg3$ to $pg2$ are shown. The errors for mapping the state from the CSLAM grid to the physics grid are similar in magnitude compared with the errors arising from mapping tendencies from the physics grid to the CSLAM grid.

2.4. High-Order (Nonconservative) Pre-Allocation of Tracer Tendencies

A high-order tracer mass increment in overlap area $A_{k\ell}$ can be computed using the following formula

$$\langle f \delta p \rangle_{k\ell} = \int_{A_{k\ell}} \left[\overline{\Delta p}_{\ell}^{(pg3)} \sum_{i+j \leq 2} \mathcal{F}_k^{(ij)} x^i y^j + \overline{f}_k^{(pg2)} \sum_{i+j \leq 2} \tilde{\mathcal{P}}_{\ell}^{(ij)} x^i y^j \right] dA, \quad (24)$$

where $\mathcal{F}_k^{(ij)}$ is the forcing increment reconstruction coefficients in the k th physics grid cell and $\overline{f}_k^{(pg2)}$ is the average physics increment in the k th physics grid cell. Note that we are using the known dry pressure reconstruction coefficients on the $pg3$ grid instead of reconstructing subgrid-scale pressure variations from the physics grid cell-averaged values. We can do that since the dry pressure is not modified by physics. This highlights the importance of a dry-pressure formulation of the dynamical core when separating physics and dynamics grids (Lauritzen et al., 2018). If the physics forcing is constant, then $\langle f \delta p \rangle_{k\ell}$ exactly equals $\langle \delta p \rangle_{k\ell}$ from (6); in other words, the mapping is designed to be reversible in dry pressure. The physics increment in terms of mixing ratio change is given by

$$\overline{f}_{k\ell} = \frac{\langle f \delta p \rangle_{k\ell}}{\langle \delta p \rangle_{k\ell}}, \quad (25)$$

where the denominator is given by (6).

Shape preservation, as defined by (11) and (13), is enforced by eliminating undershoots and overshoots on the exchange grid by modifying the forcing increment $\overline{f}_{k\ell}$ so that shape preservation is not violated in the overlap areas (In the computation of $\overline{m}_{k\ell}$ there can be small overshoots and undershoots [due to numerical integration errors] compared to the CSLAM cell average values $\overline{m}_{\ell}^{(pg3)}$ that it overlaps with so we set

$$\overline{m}_k^{(\min)} = \min \left(\overline{m}_k^{(\min)}, \left\{ \overline{m}_{\ell}^{(pg)} \mid \ell = 1, nc^2 \right\} \right) \quad (26)$$

)

$$\overline{m}_k^{(\min)} \leq \overline{m}_{k\ell} + \overline{f}_{k\ell} \leq \overline{m}_k^{(\max)}. \quad (27)$$




Figure 5. This figure is similar to Figure 4 but for a rough distribution (slotted cylinders). (a) The slotted cylinder distributions on the $pg2$ grid plotted using raster contours. The slotted cylinders after being mapped to the $pg2$ grid using the default algorithm without (b) and with (c) the high-order pre-allocation algorithm. These plots show that the mapping algorithms used to map from $pg2$ to the $pg3$ grid are shape preserving; that is, no new extrema are created in the mapping process, and they preserve a constant (the background nonzero field remains constant).

While this algorithm preserves linear correlations and shape and is consistent, is it not mass conservative. Hence, the remaining physics increment not allocated in the algorithm above is allocated using the new tendency algorithm described in section 2.3.

Combining the high-order pre-allocation algorithm with the new tendency algorithm (which in this case can also be considered as a mass-fixer that does not disrupt correlation-preservation, shape, and consistency) leads to an order-of-magnitude reduction in mapping errors for a smooth function (see Figure 4) while fulfilling the mass conservation, shape preservation, linear correlation, and consistency constraint. Mass and linear correlation preservation is illustrated in the baroclinic wave test with terminator chemistry test in section 3.1. Shape preservation and consistency is demonstrated in an idealized mapping test where a smooth function—see (23)—and a slotted-cylinder (see equation 12 in Lauritzen et al., 2012) are mapped to/from the $pg2$ and $pg3$ grids (Figures 4 and 5). Since the background value in the mapping of the slotted-cylinder field is preserved, the mapping algorithm is consistent. Since no new overshoots and undershoots are produced (particularly obvious in the mapping of the slotted cylinders) the mapping is shape preserving. We also note that the mapping errors with the default algorithm (higher-order pre-allocation with new tendency algorithm) are similar to the errors in mapping the same field from $pg3$ to $pg2$ using traditional remapping with CSLAM technology (Figure 4).

2.5. Model Configurations

All simulations in this study are run on the Cheyenne supercomputer hosted at the NCAR-Wyoming Supercomputer Center (Computational and Information Systems Laboratory, 2017). Three model component sets (*compsets*) in the Community Earth System Model, version 2.1 (CESM2.1; <https://doi.org/10.5065/D67H1H0V>) are chosen to carry out the objectives discussed in section 1. The least complex compset is a moist baroclinic wave test using a simple, Kessler microphysics scheme (*FKESSLER* compset; Lauritzen et al., 2018). The baroclinic wave setup is primarily used to evaluate the new mapping algorithms and their ability to preserve linear correlations between two reactive tracers. The role of topography is investigated using a dry Held-Suarez configuration (*FHS94* compset; Held & Suarez, 1994) modified to include real world topography. H18 indicate that this configuration tends to have more grid noise over steep terrain than in a more complex configuration using CAM version 6 physics [CAM6; https://ncar.github.io/CAM/doc/build/html/users_guide/index.html] and is therefore a conservative choice for evaluating any change in grid imprinting between $pg3$ and $pg2$.

To understand whether the resolved scales of motion are influenced by a coarser-resolution physics grid, a suite of aquaplanet simulations (Medeiros et al., 2016; Neale & Hoskins, 2000) are carried out over a range of spectral-element grid resolutions, using CAM6 physics (*QPC6* compset). The aquaplanet is an ocean-covered planet in perpetual equinox, with fixed, zonally symmetric sea surface temperatures idealized after present-day Earth (*QOBS* in Neale & Hoskins, 2000). While the dynamics time step, Δt_{dyn} , varies with resolution according to a Courant-Friedrichs-Lewy (CFL) criterion, there is no established standard for how the physics time step, Δt_{phys} , should vary across resolutions. This is further complicated by several studies indicating a high sensitivity of solutions to Δt_{phys} in CAM (Herrington & Reed, 2018; Wan et al., 2015; Williamson, 2013; Williamson & Olson, 2003).

Here, a scaling for Δt_{phys} across resolutions is proposed, based on results of the moist bubble test (Herrington & Reed, 2018) using CAM-SE-CSLAM and detailed in Appendix A. The basis for the scaling is to alleviate truncation errors that arise in the moist bubble test when Δt_{phys} is too large. The scaling is linear in grid

spacing,

$$\Delta t_{\text{phys}} = \Delta t_{\text{phys},0} \times \frac{N_{e,0}}{N_e} \text{ s}, \quad (28)$$

where $\Delta t_{\text{phys},0}$ is taken to be the standard 1,800 s used in CAM-SE-CSLAM at low resolution, $N_{e,0} = 30$ (equivalent to a dynamics grid spacing of 111.2 km). N_e refers to the horizontal resolution of the grid; each of the six panels of the cubed sphere are divided into $N_e \times N_e$ elements. Throughout the paper, spectral-element grid resolutions are denoted by an *ne* followed by the quantity N_e , for example, *ne30*.

CAM-SE-CSLAM uses a hybrid- σ pressure coordinate in the vertical. The QPC6 (FKESLLER, FHS94) compset uses 32 (30) unequally spaced layers disproportionally clustered within the boundary layer and near the tropopause. The same vertical grid is used in all model configurations in this study. The dynamics uses a vertically Lagrangian approach (Lin, 2004), in which the horizontal dynamics evolve within floating Lagrangian layers that are periodically mapped back to a fixed eulerian reference vertical coordinate (subcycled twice within in each Δt_{dyn}).

The only other parameter varied across resolutions modulates the strength of explicit numerical dissipation. The spectral element method is not implicitly diffusive, so fourth-order hyperviscosity operators are applied to the state to suppress numerical artifacts. The scaling of the hyperviscosity coefficients, ν , across resolutions is defined as

$$\nu_T = \nu_{\text{vor}} = 0.30 \times \left(\frac{30}{N_e} 1.1 \times 10^5 \right)^3 \frac{m^4}{s}, \quad (29)$$

$$\nu_p = \nu_{\text{div}} = 0.751 \times \left(\frac{30}{N_e} 1.1 \times 10^5 \right)^3 \frac{m^4}{s}, \quad (30)$$

where subscripts *T*, *vor*, *p*, and *div* refer to state variables the operators are applied to, temperature, vorticity, pressure, and divergence, respectively. The exponent in equation (30) reduces the coefficient by about (This is approximate. To reduce the coefficients by exactly an order of magnitude for each doubling of the resolution, the exponent should be $\frac{\ln 2}{\ln 10} \approx 3.01029$, which it has been updated to in the most recent version of CESM2.1) an order of magnitude for each doubling of the resolution (as in Lauritzen et al., 2018). No explicit dissipation of tracers (e.g., water vapor) is required since the semi-Lagrangian numerics in CSLAM are adequately diffusive.

3. Results

3.1. Mass Conservation and Linear-Correlation Preservation

To illustrate how different the solutions look using the coarser-resolution physics grid, Figure 6 shows a snapshot of the cloud liquid field of the moist baroclinic wave test on day 10, in the *ne30pg3* and *ne30pg2* configurations. The cloud liquid fields show in detail clouds forming at wave fronts. As expected, the *pg2* grid looks slightly coarser than *pg3* due to its larger control volumes. Despite this, the details of the wave patterns look reasonably similar to one another.

The model's ability to preserve linear correlations is assessed using the idealized Terminator “Toy” Chemistry test (Lauritzen, Conley, et al., 2015; Lauritzen et al., 2017). The tests consists of two reactive species undergoing photolysis as they are advected over the terminator line. The flow field is provided by the moist baroclinic waves test. The model is initialized with species such that their weighted sum CL_y is a constant, that is, $CL_y = CL + 2CL_2 = 4 \times 10^{-6}$ kg/kg. If linear correlations are preserved, then the column integrated weighted sum of the species, $\langle CL_y \rangle$, is constant.

H18 had shown that in the *ne30pg3* configuration, $\langle CL_y \rangle$ on day 15 of the terminator test is everywhere 4×10^{-6} kg/kg, to within machine precision. While the *pg3* to *pg2* mapping algorithm in theory preserves linear correlations to machine precision, we found larger than round-off errors in *pg2*, likely due to *if* logic with machine dependent thresholds in the implementation of the algorithm. Figure 7 shows $\langle CL_y \rangle$ on day 15 in the *ne30pg2* configuration, which has a minimum value of 3.99896×10^{-6} kg/kg, corresponding to a maximum relative error of 0.026% (The maximum relative error in a *pg2* simulation using a different compiler is 0.016%). For comparison, another terminator test is performed with the equivalent dynamics

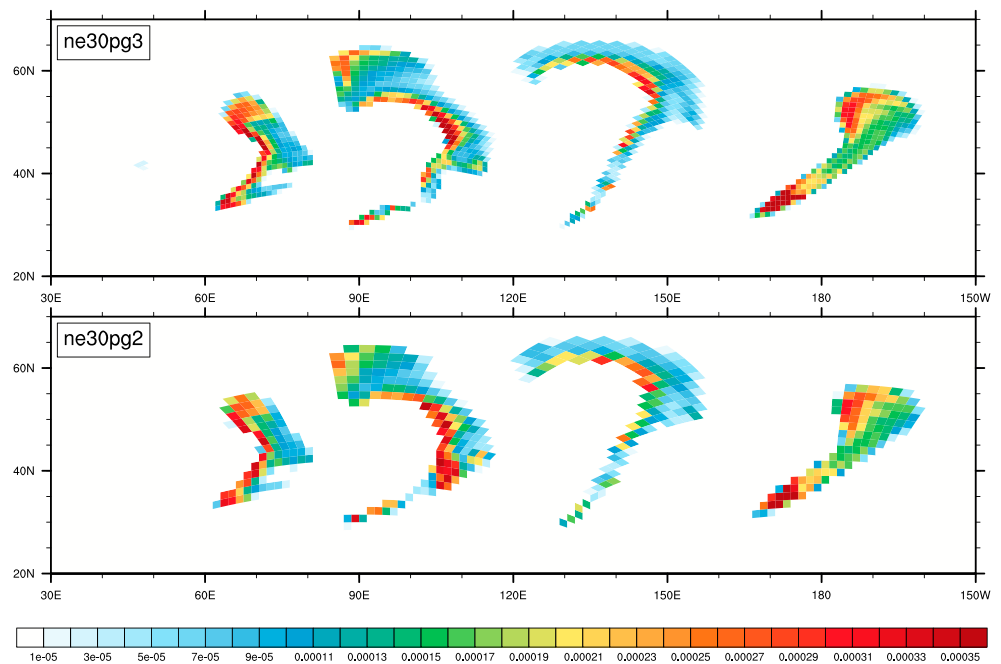


Figure 6. Snapshot of the cloud liquid field in kilograms per kilogram near the 700-hPa level, on day 10 of the moist baroclinic wave test in the *ne30pg3* and *ne30pg2* configurations, displayed on the upper and lower panels, respectively. The fields are shown as a raster plot on their respective physics grids.

grid resolution using CAM-SE (*ne30np4*), in which tracers are advected using the spectral element method. The maximum relative error in this configuration is 31.6%, 3 orders of magnitude greater error than the *ne30pg2* configuration.

Tracer mass conservation is analyzed in a pair of *ne30pg2* and *ne30pg3* aquaplanet simulations, following the method of Lauritzen and Williamson (2019). Energy and mass conservation due to a particular model process is assessed by model state I/O before and after each subprocess in the model. The loss of water vapor mass due to the mapping algorithms in the *ne30pg2* configuration is estimated as $1.184\text{E}-16$ Pa per time step, computed as the difference between the column-integrated, global mean climatological water vapor pressure increment on the physics grid and on the tracer grid. This small error is effectively 0 to within machine precision and similar to an equivalent calculation in the *ne30pg3* simulation of $2.171\text{E}-17$ Pa per time step, which contains no mapping errors since the physics and tracer grids coincide. Negligible mapping error in the *ne30pg2* configuration is primarily a result of solving equations (19) and (22) for γ_k to circumvent the negativity problem. Rerunning the *ne30pg2* aquaplanet simulation without this mass fixer, for example, through setting $\gamma_k = 1$ and $\Delta m_{k\ell}^{(\text{excess})} = \bar{m}_{k\ell}$ in the mass increment (20), results in a spurious loss of water vapor mass of $2.424\text{E}-07$ Pa per time step; the mass fixer is necessary for conserving tracer mass in *ne30pg2*.

3.2. Grid Imprinting

Flow over topography can result in significant grid imprinting using the spectral element method (Lauritzen, Bacmeister, et al., 2015, H18). Figure 8 shows the results of the Held-Suarez with topography simulations. The middle panel is the vertical pressure velocity, ω , averaged over 2 years, over the Andes and Himalayan region at two different levels in the midtroposphere, using the *ne30pg3* grid. The fields are displayed as a raster plot on the physics grid, so that individual extrema, which characterize the flow over the Andes between about 10° and 20°S , may be identified as spurious. Near the foot of the Himalayas, between about 20° and 30°N , there are parallel stripes of extrema aligned with the mountain front that appear to be spurious $2\Delta x$ oscillations.

As discussed in H18, grid imprinting over mountainous terrain tends to occur in regions of weak gravitational stability, causing extrema to extend through the full depth of the troposphere as resolved updrafts and downdrafts. Thus, grid imprinting over mountains may be alleviated through increasing the divergence damping in the model. Figure 8 (right panel) repeats the *ne30pg3* simulation through increasing v_{div} by

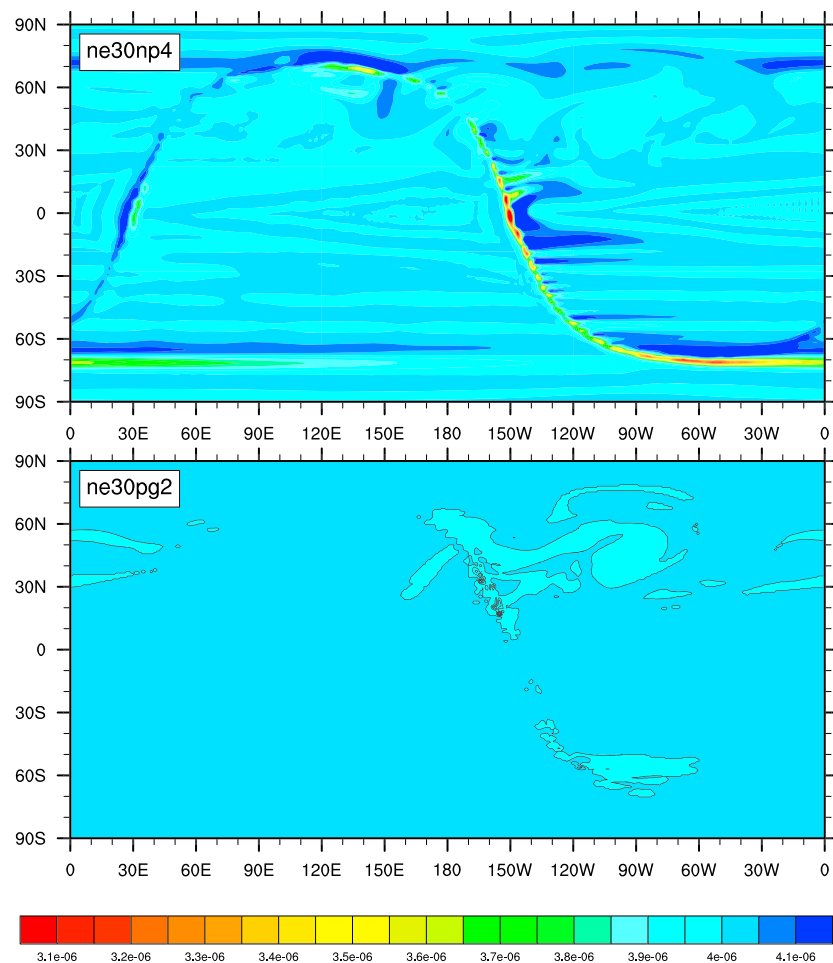


Figure 7. $\langle CLy \rangle$ in kilograms per kilogram on day 15 of the moist baroclinic wave test in the *ne30np4* and *ne30pg2* configurations, displayed on the upper and lower panels, respectively. The lower panel has a single contour level of $3.999\text{E}-6$ kg/kg corresponding to a relative error of 0.025%.

an order of magnitude. The spurious noise over the Andes and the Himalayas are damped, and grid point extrema tend to diffuse into neighboring grid cells. The wavenumber-power spectrum of the kinetic energy due to divergent flow (Figure 9) confirms that divergent modes are damped at higher wavenumbers (greater than 30), by about an order of magnitude relative to the default *ne30pg3* simulation.

The ω field of the *ne30pg2* simulation is provided in Figure 8 (left panel). Grid cell extrema over the Andes is less prevalent than in the *ne30pg3* simulation, as seen by the reduction in large-magnitude ω (e.g., red grid cells). The spurious oscillations at the foot of the Himalayas appear to have been entirely eliminated. This improvement in grid imprinting is due to the consistent smoothness properties of the control volumes in the *pg2* grid compared with the *pg3* grid discussed in section 1, and these results are consistent with our hypothesis. The divergent modes are marginally damped relative to *ne30pg3* for wavenumbers greater than about 50 but are an order of magnitude larger than in the enhanced divergence damping *ne30pg3* run (Figure 9). From a scientific and model development perspective, the *pg2* configuration is preferable to the *pg3* configuration, since it eliminates grid imprinting without placing any additional constraints on v_{div} .

3.3. Impact on Resolved Scales of Motion

Tropical regions are very sensitive to horizontal resolution, primarily due to the scale dependence of resolved updrafts and downdrafts at hydrostatic scales (Herrington & Reed, 2017; 2018; Jeevanjee, 2017; Pauluis & Garner, 2006; Weisman et al., 1997). The vertical velocity of updrafts and downdrafts is related to the horizontal length scales of buoyancy the model is able to support. This can be demonstrated through a scale analysis of the Poisson equation (Jeevanjee & Romps, 2016) valid for hydrostatic scales, showing that the

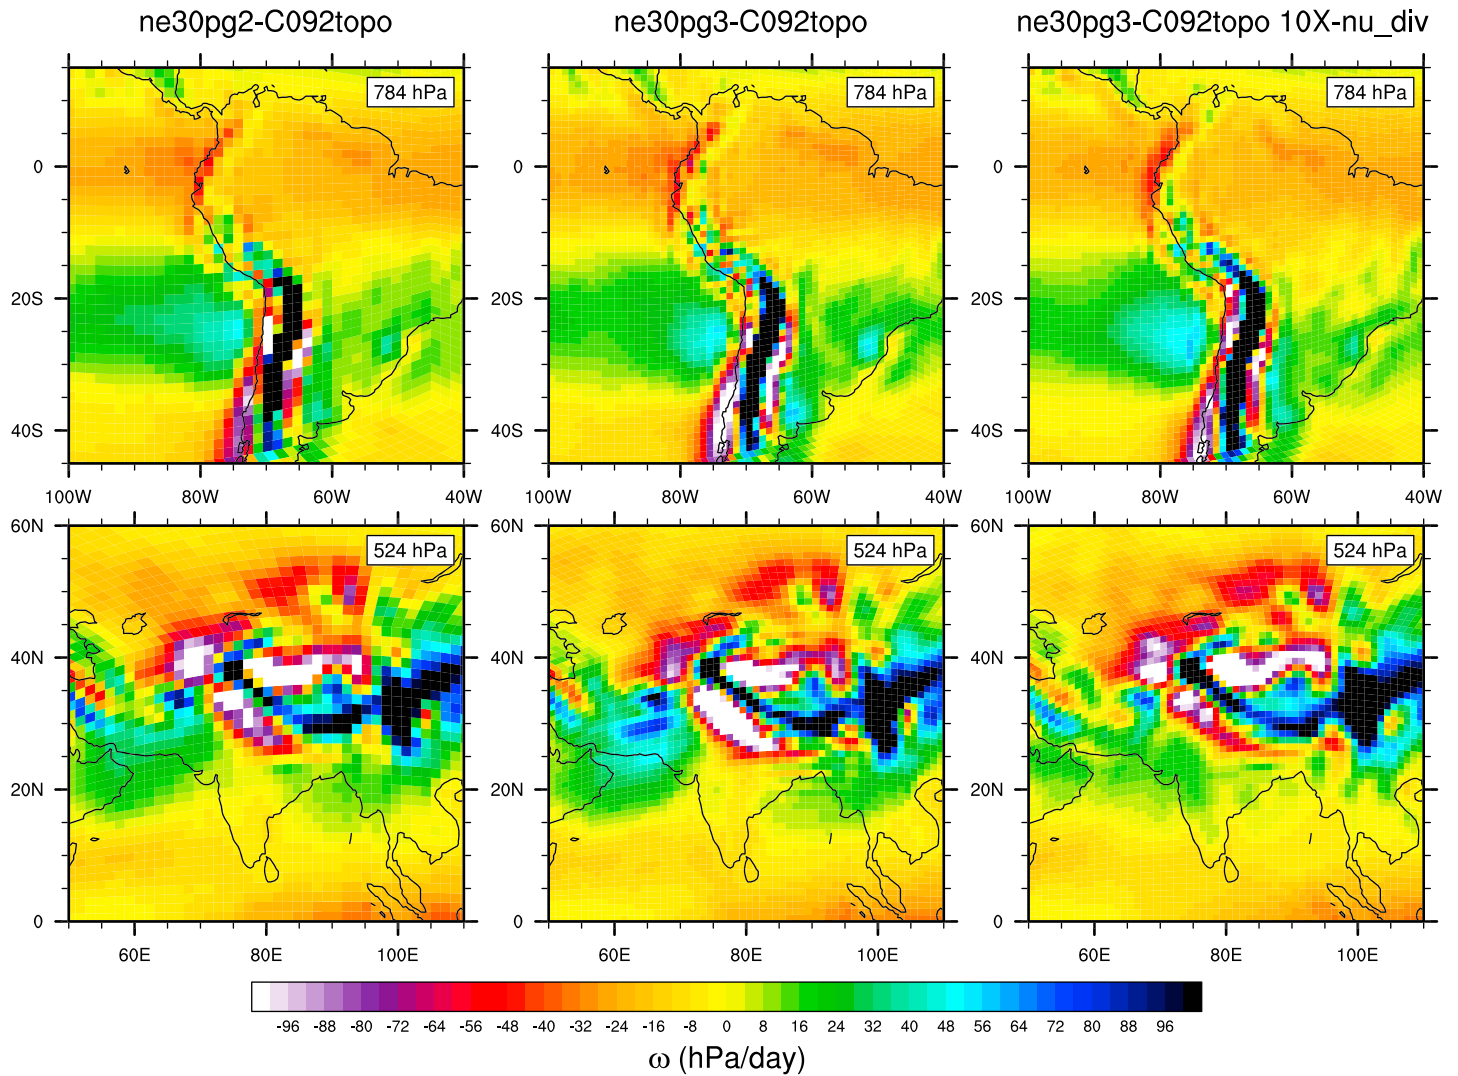


Figure 8. Mean ω at two model levels in the middle troposphere, in a Held-Suarez configuration outfitted with real world topography. (left column) *ne30pg2* (middle column) *ne30pg3* and (right column) *ne30pg3* with the divergence damping coefficient, ν_{div} , increased by an order of magnitude. The ω fields are computed from a 2-year simulation. The data are presented on a raster plot in order to identify individual grid cells

ratio of the scale of ω at two resolutions, due to their respective buoyancies, is

$$\frac{\omega_{\Delta x_1}}{\omega_{\Delta x_2}} = \frac{D_{\Delta x_2}}{D_{\Delta x_1}}, \quad (31)$$

where $D_{\Delta x}$ is a characteristic horizontal buoyancy length scale for grid-spacing Δx (hereafter referred to as the *forcing scale*), and it is presumed that the magnitude of the buoyancy and the vertical scale of the buoyancy is unchanged or compensating across the two resolutions. Equation (31) indicates that the magnitude of the vertical velocity scales like the inverse of the forcing scale, which was verified in a simple moist bubble configuration using CAM-SE and the CAM finite-volume dynamical core (Herrington & Reed, 2018), as well as using CAM-SE-CSLAM as configured in the present study (Appendix A).

In aquaplanet simulations using CAM-SE, the forcing scale is grid limited, varying with resolution in the range of 5 to 10 times the grid spacing (Herrington & Reed, 2018). Through setting the forcing scale proportional to Δx , equation (31) quantifies the sensitivity of updrafts and downdrafts to horizontal resolution. The forcing scale is analogous to an effective resolution, which is the characteristic length scale below which features are overly damped by numerical dissipation and largely absent from the solution. The effective resolution may be inferred from kinetic energy spectra as the wavenumber where the slope of the spectrum

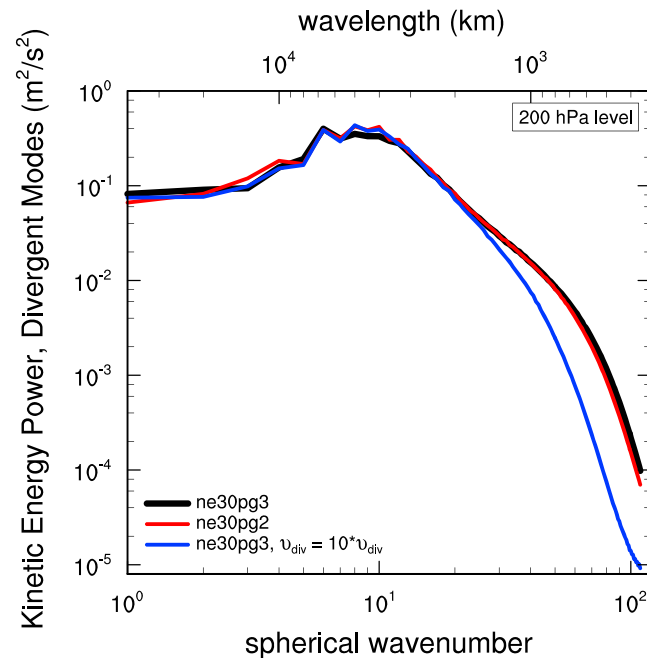


Figure 9. Kinetic energy power spectrum arising from divergent modes in *ne30pg3*, *ne30pg2*, and *ne30pg3* with the divergence damping coefficient, v_{div} , increased by an order of magnitude, in the Held-Suarez with topography simulations. Spectra computed from 5 months of 6-hourly winds.

becomes steeper than the observationally determined slope (Skamarock, 2011). In the CESM2 release of CAM-SE, this criterion occurs near wavenumber 60 (see Figure 6 in Lauritzen et al., 2018), a length scale of about 6 times the grid spacing and overlapping with the estimated forcing scale.

When the physics and dynamics grids are of different resolutions, which grid determines the models characteristic forcing scale? The remainder of this section will show that the dynamics grid spacing largely determines the forcing scale in the *pg2* physics grid configuration. The dynamical core's control over the forcing scale is illustrated at low resolution (section 3.3.1), high resolution (section 3.3.2), and across a range of resolutions typical of present day climate models (section 3.3.3).

3.3.1. Low Resolution

The question posed above may be addressed through comparing *ne30pg2*, where $\Delta x_{phys} = 166.8$ km (hereafter Δx is expressed as the average equatorial grid spacing), $\frac{3}{2}$ times larger than the dynamics grid spacing, $\Delta x_{dyn} = 111.2$ km, to a simulation where both are equal to the physics grid spacing, $\Delta x_{dyn} = \Delta x_{phys} = 166.8$ km (*ne20pg3*), and another simulation where both are equal to the dynamics grid spacing, $\Delta x_{dyn} = \Delta x_{phys} = 111.2$ km (*ne30pg3*). The resolvable scales in the *ne30pg2* solution are expected to be bounded by the *ne20pg3* and *ne30pg3* solutions. Although according to equation (28), Δt_{phys} for *ne20* grids should be different from *ne30* grids, here it is set to the *ne30* value (see Table 1), in order to reduce the differences between the three configurations, and justified because lower-resolution runs are not very sensitive to this range of Δt_{phys} (Figure A2).

Table 1

Δx and Δt for the Physics and Dynamics in the Low-Resolution Simulations

Grid name	Δx_{dyn}	Δt_{dyn}	Δx_{phys}	Δt_{phys}
ne20pg3	166.8 km	300 s	166.8 km	1,800 s
ne30pg2	111.2 km	300 s	166.8 km	1,800 s
ne30pg3	111.2 km	300 s	111.2 km	1,800 s

Note. Δx is computed as the average equatorial grid spacing.

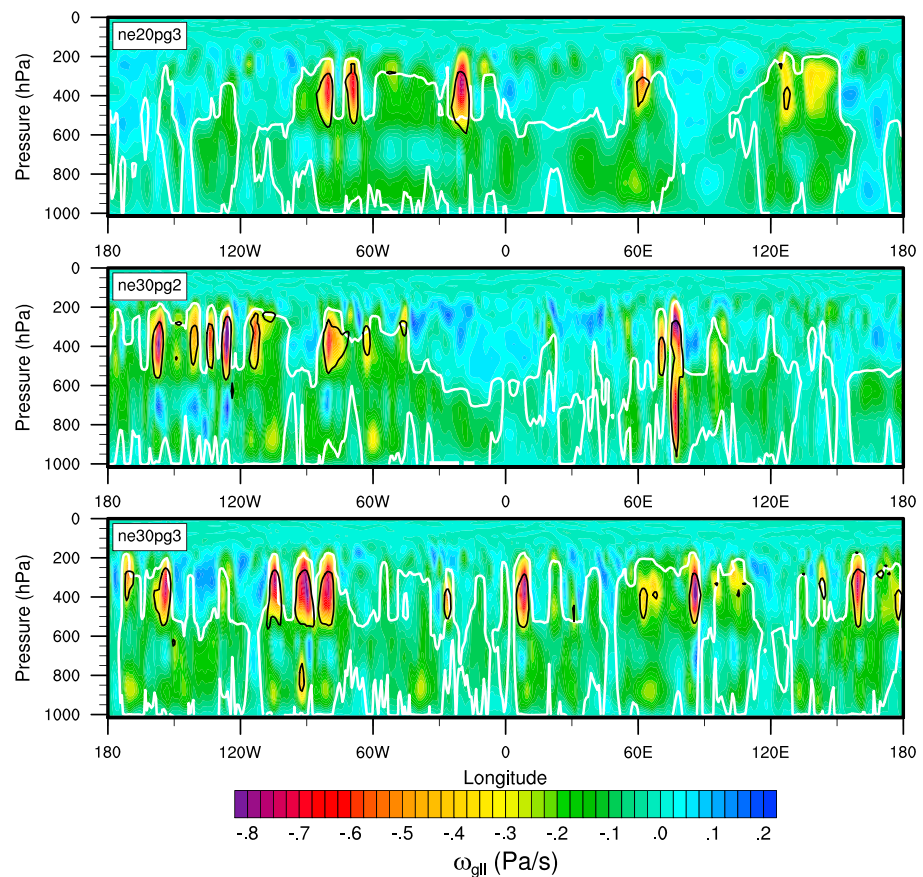


Figure 10. Snapshots in the longitude-pressure plane of $\omega^{(gll)}$ through the ITCZ region in the *ne20pg3*, *ne30pg2*, and *ne30pg3* configurations, in the upper, middle and lower panels, respectively. Black is the ± 15 -K/day contour of the physics tendencies, and the white contour is the $0.0075\text{-kg}\cdot\text{m}^{-2}\cdot\text{s}^{-1}$ contour of the parameterized deep convective mass fluxes.

Figure 10 is a snapshot of the ω field in the Intertropical Convergence Zone (ITCZ) in the pressure-longitude plane, in the three simulations. The ω field is overlaid with the ± 15 K/day contour of the physics temperature tendencies (black), which are primarily due to stratiform cloud formation. Since the component of ω due to buoyancy is determined by the physics temperature tendencies mapped to the GLL grid, the tendencies and ω are shown on the *GLL* grid, $f_T^{(gll)}$, and $\omega^{(gll)}$, respectively. The white contour is intended to outline regions where the deep convection scheme is fairly active, set to the $0.0075\text{ kg}\cdot\text{m}^{-2}\cdot\text{s}^{-1}$ value of the convective mass fluxes (note the convective mass fluxes have not been mapped to the *GLL* grid and are instead shown on the *pg* grid). The figure indicates that large regions of the ITCZ are composed of upward ω that balance the warming due to compensating subsidence produced by the deep convection scheme. Much larger magnitude ω are composed of resolved updrafts driven by the buoyancy of stratiform clouds and resolved downdrafts due to evaporation of condensates produced by overlying clouds (Herrington & Reed, 2018). These large buoyancy stratiform clouds tend to form in the middle-to-upper troposphere due to detrainment of moisture from the deep convection scheme (Zhang & McFarlane, 1995).

It is not obvious from the snapshots in Figure 10 whether the length scales of the stratiform clouds, that is, the models characteristic forcing scale, are any different across the three simulations. Analogous to determining the effective resolution (Skamarock, 2011), the forcing scale may be inferred from the wave-number power spectrum of $f_T^{(gll)}$ as the maximum wavenumber prior to the steep, unphysical decline in power that characterizes the near-grid scale (hereafter $f_T^{(gll)}$ is referred to as the *forcing*). The wave-number power spectrum of the forcing in the middle-to-upper troposphere is shown in Figure 11a. Unlike kinetic energy spectra, the decline in power near the models effective resolution is more gradual, making it difficult to determine a characteristic forcing scale from the spectra. However, it is clear that the slope of the *ne20pg3* spectrum begins to steepen at smaller wavenumbers than in the *ne30pg3* spectra. Additionally, the *ne30pg2* spectra




Figure 11. (a) Wavenumber-power spectrum of the temperature tendencies from the moist physics, near the 369-hPa level, (b) probability density distribution, and (c) the scaled probability density distribution of upward ω everywhere in the model. The scaled distributions are scaled to ne30pg3 using Δx_{phys} .

is remarkably similar to the ne30pg3 spectra, for all wavenumbers. These spectra indicate that the characteristic forcing scale in the ne30pg2 and ne30pg3 simulations are similar and that both are smaller than the ne20pg3 forcing scale. From equation (31), it is expected that the magnitude of the vertical motion is greater in both the ne30pg2 and ne30pg3 simulations.

The probability density function (PDF) of upward $\omega^{(gll)}$ everywhere in the simulations is shown in Figure 11b. Large-magnitude $\omega^{(gll)}$ are more frequent in the ne30pg2 run, compared to ne20pg3, and the PDF is actually more similar to the ne30pg3 distribution, consistent with their similar forcing scales. This may be further illustrated through scaling the PDFs,

$$P_s(\omega) = \alpha \times P(\omega/\alpha), \quad (32)$$

where $P_s(\omega)$ is the scaled PDF of ω and α is the ratio of ω to ω_{target} , the ω associated with the target grid resolution, Δx_{target} . Making the assumption that the forcing scale is linear in Δx , then from equation (31), $\alpha = \Delta x_{target}/\Delta x$. The target resolution is taken here to be equal to the ne30pg3 grid resolution.

If the forcing scale of ne30pg2 is in fact determined by Δx_{phys} , then one sets $\Delta x = \Delta x_{phys}$ in α . This scaled PDF, however, severely overestimates the frequency of upward ω of the target resolution, ne30pg3 (Figure 11c). It is clear from the similarity of the unscaled PDFs of ne30pg2 and ne30pg3 (Figure 11b), and their forcing spectra (Figure 11a), that the characteristic forcing scale in these two configurations are approximately the same. It follows that the forcing scales in ne30pg2 and ne30pg3 are determined by their common grid, Δx_{dyn} , rather than Δx_{phys} , which are different. And one can be reasonably confident in the linear framework used to approximate α —the scaled ne20pg3 PDF fits the ne30pg3 distribution quite well. It then follows that the forcing scale of ne20 simulations is about $\frac{3}{2}$ times that of ne30 simulations, the ratio of their grid spacings.

There are two reasons the pg2 forcing scale is determined by the GLL grid. The first is that the hyperviscosity coefficients are a function of the GLL grid resolution (equation (30)), and the second is that the physics tendencies are mapped to the pg3 and GLL grids using high-order mapping, which reconstructs scales the pg2 grid is unable to support (see Appendix B). The impact of only using low-order mapping or only using ne20 viscosity in a ne30pg2 simulation results in a forcing spectra that lies in between the default ne30pg2 and ne20pg3 runs (not shown). The combined effect of both factors on the forcing scale is illustrated through an ne30pg2 simulation that uses low-order mapping, and with hyperviscosity coefficients set to ne20

Table 2

Δx and Δt for The Physics and Dynamics in the High-Resolution Simulations

Grid name	Δx_{dyn}	Δt_{dyn}	Δx_{phys}	Δt_{phys}
ne80pg3	41.7 km	112.5 s	41.7 km	675 s
ne120pg2	27.8 km	75 s	41.7 km	450 s
ne120pg3	27.8 km	75 s	27.8 km	450 s

Note. Δx is computed as the average equatorial grid spacing.




Figure 12. As in Figure 11, but for the high-resolution simulations. Asterisks indicate that $\Delta t_{\text{phys}} = 675$ s, which is larger than that used for the default *ne120* runs (see Table 2).

values (*ne30pg2-ne20visc-loworder* in Figure 11). The PDF of $\omega^{(gll)}$ and the forcing spectrum more closely resemble the *ne20pg3* run. In the *ne30pg2-ne20visc-loworder* configuration, the forcing scale is more accurately determined by Δx_{phys} since the scaled PDF is in fairly good agreement with the *ne30pg3* simulation (Figure 11c).

3.3.2. High Resolution

The experiment described in the previous section is repeated here for a *ne120pg2* aquaplanet simulation, corresponding to an approximate grid spacing of $\Delta x_{\text{dyn}} = 27.8$ km and $\Delta x_{\text{phys}} = 41.7$ km. *ne80pg3* refers to the grid in which the physics and dynamics are the same resolution as the physics of the *ne120pg2* grid, and *ne120pg3* the grid in which the physics and dynamics are equal to the resolution of the dynamics of *ne120pg2*. At these higher resolutions, the solutions are sensitive to Δt_{phys} (Figure A2), and so the *ne80* grid uses a larger time step than that of the *ne120* grids (Table 2), following equation (28).

Figure 12 is the same as Figure 11 but for the high-resolution simulations. While the *ne80pg3* forcing spectra begins to drop off near wavenumber 100, the *ne120pg2* and *ne120pg3* drop off closer to wavenumber 200, and their spectra lie on top of one another (Figure 12a). The PDFs of (upward) $\omega^{(gll)}$ show that the *ne120* distributions lie on top of one another, and while not a perfect match, both *ne120* runs have substantially more frequent large-magnitude vertical motion than in the *ne80pg3* run (Figure 12b). As in the low-resolution runs, the similarity of the *ne120* forcing spectra and $\omega^{(gll)}$ distributions indicate that the forcing scale of the *ne120pg2* run is not determined by the physics grid spacing but rather the dynamics grid spacing. This is also evident from the overprediction of the frequency of large-magnitude $\omega^{(gll)}$ compared with the *ne120pg3* run, through scaling the *ne120pg2* PDF and setting the forcing scale proportional to Δx_{phys} in equation (32) (Figure 12c).

In the *ne120pg2* simulation, the dynamics grid determines the forcing scale for the same two reasons found in the low-resolution runs. The high-order mapping of the physics to the dynamics is important for reconstructing scales not supported on the *pg2* grid, and scaling the viscosity coefficients by the dynamics grid spacing is also important. But in order to recreate the *ne80pg3* solution using the *ne120pg2* grid, the physics time steps must be the same for these two grids. Combining all three modifications leads to an *ne120pg2* solution that resembles the *ne80pg3* run (*ne120pg2-ne80visc-loworder* * in Figure 12). The forcing spectrum and distribution of $\omega^{(gll)}$ match that of the *ne80pg3* run, and scaling the PDF by Δx_{phys} closely resembles the *ne120pg3* distribution.

Table 3

Δx and Δt for the Physics and Dynamics in the Intermediate Resolution Simulations

Grid name	Δx_{dyn}	Δt_{dyn}	Δx_{phys}	Δt_{phys}
ne40pg3	83.4 km	222.5 s	83.4 km	1,350 s
ne60pg2	55.6 km	150 s	83.4 km	900 s
ne60pg3	55.6 km	150 s	55.6 km	900 s

Note. Δx is computed as the average equatorial grid spacing.




Figure 13. Global mean, time-mean (a) upward ω , (b) Cloud Layers Unified by Binormals precipitation rate, and (c) parameterized deep convective precipitation rate. All means computed from the final 11 months of 1-year simulations, and upward ω is computed using 6-hourly output.

3.3.3. Across Resolutions

Three intermediate resolution aquaplanets are run to provide a continuous representation of the solution spanning from low to high resolution (Table 3). Figure 13 are scatter plots of the climatological global mean state versus Δx_{dyn} for all model configurations listed in Tables 1–3. The fields plotted in the figure, upward ω , and the two components of precipitation, stratiform precipitation rate (Cloud Layers Unified by Binormals; CLUBB) and deep convective precipitation rate (ZM), are all sensitive to resolution. Upward ω and CLUBB precipitation decreases, and ZM precipitation increases monotonically with Δx_{dyn} . The *pg2* solutions have very similar values to the *pg3* solutions, although they are slightly offset toward the lower-resolution side of the plots. The differences between the *pg2* and *pg3* solutions are much less than the differences between *pg2* and configurations where the physics and dynamics grids are both equal to the *pg2* physics grid resolution (e.g., *ne40pg3* compared with *ne60pg2*). The mean state of the configurations resembles that of the transients discussed in the previous sections; the coarser *pg2* physics grid does not appear to degrade the resolved scales of motion, which are primarily determined by the dynamics grid resolution.

4. Conclusions

This study documents the implementation of a coarser-resolution physics grid into the CAM, with spectral element dynamics (based on a dry-mass vertical coordinate) and conservative semi-Lagrangian advection of tracers (CAM-SE-CSLAM). The spectral-element and tracer advection grids are mapped to a finite-volume physics grid after Herrington et al. (2018) but containing $\frac{2}{3}$ fewer degrees of freedom in each horizontal direction. Mapping from the coarser physics grid to the dynamics and tracer grids is performed with high-order reconstructions, and a tendency mapping algorithm is developed to ensure shape preservation, consistency, linear-correlation preservation, and mass conservation. These numerical properties are verified to a high degree of precision through idealized tests.

The coarser-resolution physics grid is designed to remove grid imprinting that manifests for nonsmooth problems using element-based high-order Galerkin methods. The lower-resolution physics grid provides a volume mean state to the physics that is computed from an equal sampling of the different types of nodal solutions arising from the spectral-element method, and it was hypothesized that this method eliminates grid imprinting from the element boundaries. Using a Held-Suarez configuration modified with real-world topography, it was shown that element boundary noise over steep topography is nearly, if not entirely, eliminated from the coarser physics grid solution, consistent with our hypothesis.

Physical parameterizations make up a significant fraction of the total computational cost of atmosphere models, and the coarser physics grid may be used to reduce this overhead. The cost savings is due to the factor $\frac{5}{9}$ fewer grid columns in which the physics need be computed, cutting the physics costs by at least half. In CESM2.0, the CAM6 physics makes up about half the cost of the overall model (Lauritzen et al., 2018), and so the total cost savings is potentially large, but note that additional mappings between the tracer advection and physics grids increase the cost of the dynamical core. The authors sought to understand whether the reduction in computational cost occurs at the expense of a degraded solution, through aliasing the dynamics to the coarser-resolution physics. An exhaustive number of grids were developed and run in




Figure A1. Minimum ω from a series of 1-day long moist bubble simulations in which the dynamics grid spacing (top x axis) and bubble radius (bottom x axis) are varied by the same factor. The circles are for experiments where Δt_{phys} varies according to equation (28), and the crosses for simulations where Δt_{phys} is fixed at 1,800 s. The gray lines are the analytically predicted ω , after equation (31) scaled to the lowest-resolution solutions. The magnitude of ω in the pg3 solutions are systematically larger than the pg2 solutions, which is primarily a result of the damping effect of integrating the basis functions over a larger control volume.

an aquaplanet configuration and confirm that the resolved scales of motion are not degraded through the use of a coarser-resolution physics grid. It was found that the resolved scales are primarily determined by the effective resolution of the dynamical core. This was attributed to two factors: (1) explicit numerical dissipation by the dynamics blurs the distinction between solutions on the physics, dynamics, or tracer grids and (2) high-order mapping of the physics tendencies to the dynamics and tracer grids reconstructs scales that are not supported on the coarser physics grid.

The coarser physics grid in CAM-SE-CSLAM provides significant cost savings with little to no downside. The coarser physics grid replicates solutions from the conventional method of evaluating the physics at the same resolution as the dynamical core, mitigates grid imprinting in the solution, and runs efficiently on massively parallel systems. The coarser physics grid may be leveraged to reduce the computational burden as a component of increasingly expensive Earth System Models or permit once unattainable throughputs for high-resolution climate simulations. This approach may also be useful in the computationally burdensome superparameterization approach (Randall et al., 2003), in which a cloud-resolving model could instead be embedded into the lower-resolution physics grid, reducing computational overhead. The coarser physics grid configuration of CAM-SE-CSLAM is well positioned to address the scientific challenges ahead, as a formidable next-generation climate model.

Appendix A: Defining Δt_{phys} Across Resolutions

Herrington and Reed (2018) developed a moist bubble test, which indicate that time-truncation errors are large at high resolution (about 50 km or less) using more conventional values for the physics time step. The test may be able to provide insight on a reasonable scaling of Δt_{phys} across resolutions in more complex configurations. In the test a set of nonrotating simulations are initialized with a warm, supersaturated moist bubble, and the grid spacing and bubble radius are simultaneously reduced by the same factor in each run through varying the planetary radius. The test was designed to mimic the reduction in buoyancy length scales that occur when the model resolution is increased in more complex configurations (Hack et al., 2006; Herrington & Reed, 2018).

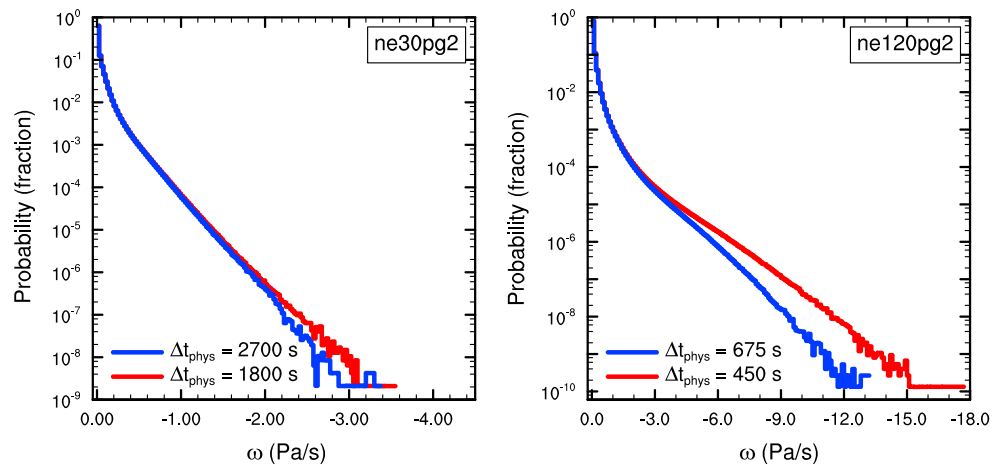


Figure A2. Probability density distribution of upward ω everywhere in the model in the aquaplanets using the *ne30pg2* grid (left) and the *ne120pg2* grid (right). Figure computed for 1 year of 6-hourly data. The different colors indicate the physics time steps used in the runs.

The moist bubble test is performed with CAM-SE-CSLAM and coupled to the simple condensation routine of Kessler (1969) across five different resolutions (pertaining to the *ne30*, *ne40*, *ne60*, *ne80*, and *ne120* grids). The results are expressed as the minimum ω throughout each 1-day simulation and shown in Figure A1. Two sets of simulations are performed with both *pg3* and *pg2*, one with Δt_{phys} determined by equation (28) and an equivalent set of simulations with $\Delta t_{\text{phys}} = 1,800$ s for all resolutions.

With the diameters of the bubbles set proportional to Δx_{dyn} , Herrington and Reed (2018) has shown that ω converges to the scaling of equation (31) in the limit of small Δt_{phys} , where small Δt_{phys} refers to the CFL limiting time step used by the dynamics. Equation (31) is overlain as gray lines in Figure A1, with *ne30* being the reference resolution. The solutions using Δt_{phys} from equation (28) follow the scaling, whereas fixing $\Delta t_{\text{phys}} = 1,800$ s across resolutions damps the solution relative to the analytical solution, progressively more so at higher resolutions. If Δt_{phys} is too large, the solution has nonnegligible error, which is avoided through scaling Δt_{phys} according to equation (28).

To get a handle on whether the test is useful for understanding more realistic configurations, four aquaplanet simulations are performed using the CAM6 physics package: a pair of *ne30pg2* simulations, one in which Δt_{phys} is set to the appropriate value from equation (28) (1,800 s), and another where it is set to the Δt_{phys} corresponding to the *ne20* resolution (2,700 s). Similarly, a pair of *ne120pg2* simulations are performed, one with Δt_{phys} set to the value from equation (28) (450 s) and one with Δt_{phys} set to the *ne80* value (675 s).

Figure A2 shows the PDFs of upward ω computed from a year of six-hourly data in the simulations. At lower resolution, Δt_{phys} has only a very small effect on the solution, near the tail end of the distributions. At high resolution, values of ω less than about -3 Pa/s are more frequent in the small Δt_{phys} run, with the discrepancy growing more for larger magnitudes of ω . The progressively larger errors with increasing resolution also manifests in the moist bubble tests, indicating that truncation errors arising from large Δt_{phys} do exist in more complex configurations.

Appendix B: The Impact of High-Order Mapping to the Dynamics Grids

Figure B1 (left) shows a close-up of the wavenumber power spectrum of the forcing on the *pg* grid (dotted), where it is computed, and on the *GLL* grid (solid), where it has been mapped. In *ne30pg3*, the magnitudes are similar on both grids, except the mapping tends to damp the high wavenumbers of the forcing on the *GLL* grid (greater than 60), but these scales are primarily below the effective resolution of the model and should not effect the solution. For *ne30pg2*, the magnitude of the forcing is actually greater after mapping to the *GLL* grid and more similar to the forcing in the *ne30pg3* simulations. The high-order mapping can therefore replicate the scales of the physics tendencies that occur in the *pg3* simulation, even though the physics are evaluated on a coarser *pg2* grid.




Figure B1. (left) Wavenumber-power spectrum of the temperature tendencies from the moist physics, at the 369-hPa level, and (right) probability density distribution of upward ω , everywhere in the model, for 3-year-long aquaplanet simulations. Solid lines refer to values of on the *GLL* grids, and dashed lines the fields on the *pg* grids. See text for details regarding the three simulations.

The importance of the high-order mapping can be shown with an additional *ne30pg2* simulation, using low-order mapping (*ne30pg2-loworder* in Figure B1). Specifically, low-order mapping refers to piecewise constant mapping between the *pg2* and *CSLAM* grids, and bilinear mapping from *pg2* to the *GLL* grid. The forcing spectrum is now similar on both the *pg2* and *GLL* grids, although the low-order mapping tends to damp the forcing on the *GLL* grid for wavenumbers greater than about 60, scales smaller than the models effective resolution (Figure B1, left). A close-up of the PDF of $\omega^{(gll)}$ is provided in Figure B1 (right, solid lines). As expected, the frequency of large-magnitude $\omega^{(gll)}$ in the low-order run is less compared to the default *ne30pg2* simulation.

The dotted lines in Figure B1 (right) show the PDF of ω on the *pg* grids. The frequency of large-magnitude ω is reduced on the *pg* grids, compared to the state on the *GLL* grids. This is primarily due to the smoothing effect of integrating the nodal point values over control volumes (H18). The larger ω values are even less frequent on the *pg2* grid due to integrating over control volumes $\frac{9}{4}$ times greater than the *pg3* control volumes.

References

- Barth, T., & Jespersen, D. (1989). The design and application of upwind schemes on unstructured meshes. In *Proc. AIAA 27th Aerospace Sciences Meeting*, Reno.
- Brdar, S., Baldauf, M., Dedner, A., & Klöforn, R. (2013). Comparison of dynamical cores for NWP models: Comparison of COSMO and Dune. *Theoretical and Computational Fluid Dynamics*, 27(3–4), 453–472. <https://doi.org/10.1007/s00162-012-0264-z>
- Computational and Information Systems Laboratory (2017). Cheyenne: HPE/SGI ICE XA System (AMP Model Developments), National Center for Atmospheric Research, Boulder, CO.
- Dennis, J. M., Edwards, J., Evans, K. J., Guba, O., Lauritzen, P. H., Mirin, A. A., et al. (2012). CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model. *International Journal of High Performance Computing*, 26(1), 74–89. <https://doi.org/10.1177/10943420111428142>
- Giraldo, F. X., & Restelli, M. (2008). A study of spectral element and discontinuous Galerkin methods for the Navier Stokes equations in nonhydrostatic mesoscale atmospheric modeling: Equation sets and test cases. *Journal of Computational Physics*, 227, 3849–3877. <https://doi.org/10.1016/j.jcp.2007.12.009>
- Hack, J., Caron, M., Danabasoglu, G., Oleson, K. W., Bitz, C., & Truesdale, J. (2006). CCSM-CAM3 climate simulation sensitivity to changes in horizontal resolution. *Journal Climate*, 19(1), 2269–2289.
- Held, I. M., & Suarez, M. J. (1994). A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. *Bulletin of the American Meteorological Society*, 73, 1825–1830.
- Herrington, A. R., Lauritzen, P. H., Taylor, M. A., Goldhaber, S., Eaton, Brian E., Bacmeister, J. T., et al. (2018). Physics-dynamics coupling with element-based high-order Galerkin methods: Quasi equal-area physics grid. *Monthly Weather Review*, 47, 69–84.
- Herrington, A. R., & Reed, K. A. (2017). An explanation for the sensitivity of the mean state of the community atmosphere model to horizontal resolution on aquaplanets. *Journal Climate*, 30(13), 4781–4797. <https://doi.org/10.1175/jcli-d-16-0069.1>
- Herrington, A. R., & Reed, K. A. (2018). An idealized test of the response of the Community Atmosphere Model to near-grid-scale forcing across hydrostatic resolutions. *Journal of Advances in Modeling Earth Systems*, 10, 560–575. <https://doi.org/10.1002/2017MS001078>

Acknowledgments

The CESM project is supported primarily by the National Science Foundation (NSF). This material is based upon work supported by the National Center for Atmospheric Research (NCAR), which is a major facility sponsored by NSF under cooperative agreement 1852977. Computing and data storage resources, including the Cheyenne supercomputer (<https://doi.org/10.5065/D6RX99HX>), were provided by the Computational and Information Systems Laboratory (CISL) at NCAR. Herrington, Reed, and Lauritzen are indebted to the NCAR Advanced Study Program graduate visitor program for funding Herrington's 12-month visit. Reed was partially supported by U.S. Department of Energy Office of Science grant DE-SC0019459. Goldhaber was partially supported by the U.S. Department of Energy Office of Biological and Environmental Research, Work Package 12-015334 “Multiscale Methods for Accurate, Efficient, and Scale-Aware Models of the Earth System.” Model output used in this study is publicly available at the website (<https://github.com/adamrher/2017-physres/tree/master/data>). The lower resolution physics grid configuration is available in the CESM2.1.1 release version.

- Jablonowski, C., & Williamson, D. L. (2011). The pros and cons of diffusion, filters and fixers in atmospheric general circulation models. In P. H. Lauritzen, R. D. Nair, C. Jablonowski, & M. Taylor (Eds.), *Numerical techniques for global atmospheric models* (Vol. 80, Lecture notes in computational science and engineering. Berlin and Heidelberg: Springer.
- Jeevanjee, N. (2017). Vertical velocity in the gray zone. *Journal of Advances in Modeling Earth Systems*, 9, 2304–2316. <https://doi.org/10.1002/2017MS001059>
- Jeevanjee, N., & Romps, D. M. (2016). Effective buoyancy at the surface and aloft. *Quarterly Journal of the Royal Meteorological Society*, 142(695), 811–820.
- Jöckel, P., von Kuhlmann, R., Lawrence, M. G., Steil, B., Brenninkmeijer, C. A. M., Crutzen, P. J., et al. (2001). On a fundamental problem in implementing flux-form advection schemes for tracer transport in 3-dimensional general circulation and chemistry transport models. *Quarterly Journal of the Royal Meteorological Society*, 127(573), 1035–1052.
- Jones, P. W. (1999). First- and second-order conservative remapping schemes for grids in spherical coordinates. *Monthly Weather Review*, 127, 2204–2210.
- Kessler, E. (1969). On the distribution and continuity of water substance in atmospheric circulations. *Meteorological Monographs*, 10(32), 88.
- Lander, J., & Hoskins, B. J. (1997). Believable scales and parameterizations in a spectral transform model. *Monthly Weather Review*, 125, 292–303.
- Lauritzen, P. H., Bacmeister, J. T., Callaghan, P. F., & Taylor, M. A. (2015). NCAR global model topography generation software for unstructured grids. *Geoscientific Model Development Discussions*, 8(6), 4623–4651. <https://doi.org/10.5194/gmdd-8-4623-2015>
- Lauritzen, P. H., Conley, A. J., Lamarque, J.-F., Vitt, F., & Taylor, M. A. (2015). The terminator “toy” chemistry test: A simple tool to assess errors in transport schemes. *Geoscientific Model Development*, 8(5), 1299–1313.
- Lauritzen, P. H., Nair, R. D., Herrington, A. R., Callaghan, P., Goldhaber, S., Dennis, J. M., et al. (2018). NCAR CESM2.0 release of CAM-SE: A reformulation of the spectral-element dynamical core in dry-mass vertical coordinates with comprehensive treatment of condensates and energy. *Journal of Advances in Modeling Earth Systems*, 10, 1537–1570. <https://doi.org/10.1029/2017MS001257>
- Lauritzen, P. H., Nair, R. D., & Ullrich, P. A. (2010). A conservative semi-Lagrangian multi-tracer transport scheme (CSLAM) on the cubed-sphere grid. *Journal of Computational Physics*, 229, 1401–1424.
- Lauritzen, P. H., Skamarock, W. C., Prather, M. J., & Taylor, M. A. (2012). A standard test case suite for 2D linear transport on the sphere. *Geoscientific Model Development*, 5, 887–901.
- Lauritzen, P. H., Taylor, M. A., Overfelt, J., Ullrich, P. A., Nair, R. D., Goldhaber, S., & Kelly, R. (2017). CAM-SE-CSLAM: Consistent coupling of a conservative semi-Lagrangian finite-volume method with spectral element dynamics. *Monthly Weather Review*, 145(3), 833–855.
- Lauritzen, P. H., & Thuburn, J. (2012). Evaluating advection/transport schemes using interrelated tracers, scatter plots and numerical mixing diagnostics. *Quarterly Journal of the Royal Meteorological Society*, 138(665), 906–918.
- Lauritzen, P. H., Ullrich, P. A., & Nair, R. D. (2011). Atmospheric transport schemes: Desirable properties and a semi-Lagrangian view on finite-volume discretizations. In P. H. Lauritzen, R. D. Nair, C. Jablonowski, & M. Taylor (Eds.), *Numerical techniques for global atmospheric models, Lecture Notes in Computational Science and Engineering* (Vol. 80). Berlin and Heidelberg: Springer.
- Lauritzen, P. H., & Williamson, D. L. (2019). A total energy error analysis of dynamical cores and physics-dynamics coupling in the community atmosphere model (cam). *Journal of Advances in Modeling Earth Systems*, 11, 1309–1328. <https://doi.org/10.1029/2018MS001549>
- Lin, S.-J. (2004). A ‘vertically Lagrangian’ finite-volume dynamical core for global models. *Monthly Weather Review*, 132, 2293–2307.
- Medeiros, B., Williamson, D. L., & Olson, J. G. (2016). Reference aquaplanet climate in the community atmosphere model, version 5. *Journal of Advances in Modeling Earth Systems*, 8, 406–424. <https://doi.org/10.1002/2015MS000593>
- Nair, R., Choi, H.-W., & Tufo, H. (2009). Computational aspects of a scalable high-order discontinuous Galerkin atmospheric dynamical core. *Computers & Fluids*, 38(2), 309–319. <https://doi.org/10.1016/j.compfluid.2008.04.006>
- Nair, R. D., & Lauritzen, P. H. (2010). A class of deformational flow test cases for linear transport problems on the sphere. *Journal of Computational Physics*, 229, 8868–8887.
- Neale, R. B., Chen, C.-C., Gettelman, A., Lauritzen, P. H., Park, S., Williamson, D. L., et al. (2012). The APE atlas. <https://doi.org/10.5065/D6FF3QBR>
- Neale, R. B., & Hoskins, B. J. (2000). A standard test for AGCMS including their physical parametrizations: I: The proposal. *Atmospheric Science Letters*, 1(2), 101–107.
- Pauluis, O., & Garner, S. (2006). Sensitivity of radiative-convective equilibrium simulations to horizontal resolution. *Journal of the Atmospheric Sciences*, 63(7), 1910–1923.
- Randall, D., Khairoutdinov, M., Arakawa, A., & Grabowski, W. (2003). Breaking the cloud parameterization deadlock. *Bulletin of the American Meteorological Society*, 84(11), 1547–1564.
- Shutts, G. (2005). A kinetic energy backscatter algorithm for use in ensemble prediction systems. *Quarterly Journal of the Royal Meteorological Society*, 131, 3079–3102.
- Skamarock, W. C. (2011). Kinetic energy spectra and model filters. In P. H. Lauritzen, R. D. Nair, C. Jablonowski, & M. Taylor (Eds.), *Numerical techniques for global atmospheric models* (vol. 80, Lecture Notes in Computational Science and Engineering: Springer.
- Taylor, M. A., Edwards, J., & St-Cyr, A. (2008). Petascale atmospheric models for the Community Climate System Model: New developments and evaluation of scalable dynamical cores. *Journal of Physics: Conference Series*, 125(1), 12023. <https://doi.org/10.1088/1742-6596/125/1/012023>
- Wan, H., Rasch, P. J., Taylor, M. A., & Jablonowski, C. (2015). Short-term time step convergence in a climate model. *Journal of advances in modeling earth systems*, 7, 215–225. <https://doi.org/10.1002/2014MS000368>
- Wedi, N. P. (2014). Increasing horizontal resolution in numerical weather prediction and climate simulations: Illusion or panacea? *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 372(2018).
- Weisman, M. L., Skamarock, W. C., & Klemp, J. B. (1997). The resolution dependence of explicitly modeled convective systems. *Monthly Weather Review*, 125(4), 527–548.
- Williamson, D. L. (1999). Convergence of atmospheric simulations with increasing horizontal resolution and fixed forcing scales. *Tellus A*, 51, 663–673.
- Williamson, D. L. (2013). The effect of time steps and time-scales on parametrization suites. *Quarterly Journal of the Royal Meteorological Society*, 139(671), 548–560.
- Williamson, D. L., & Olson, J. G. (2003). Dependence of aqua-planet simulations on time step. *Quarterly Journal of the Royal Meteorological Society*, 129(591), 2049–2064.
- Zhang, G. J., & McFarlane, N. A. (1995). Sensitivity of climate simulations to the parameterization of cumulus convection in the Canadian climate center general-circulation model. *Atmosphere-Ocean*, 33(3), 407–446.