

Continuous and Discontinuous Galerkin Methods for a Scalable Three-Dimensional Nonhydrostatic Atmospheric Model: Limited-Area Mode

James F. Kelly and Francis X. Giraldo

Department of Applied Mathematics, Naval Postgraduate School, Monterey, CA

Abstract

This paper describes a unified, element based Galerkin (EBG) framework for a three-dimensional, nonhydrostatic model for the atmosphere. In general, EBG methods possess high-order accuracy, geometrical flexibility, excellent dispersion properties and good scalability. Our nonhydrostatic model, based on the compressible Euler equations, is appropriate for both limited-area and global atmospheric simulations. Both a Continuous Galerkin (CG), or spectral element, and discontinuous Galerkin (DG) model are considered using hexahedral elements. The formulation is suitable for both global and limited-area atmospheric modeling, although we restrict our attention to 3D limited-area phenomena in this study; global atmospheric simulations will be presented in a follow-up paper. Domain decomposition and communication algorithms utilized by both our CG and DG models are presented. The communication volume and exchange algorithms for CG and DG are compared and contrasted. Numerical verification of the model is performed using two test cases: flow past a 3D mountain and buoyant convection of a bubble in a neutral atmosphere; these tests indicate that both CG and DG can simulate the necessary physics of dry atmospheric dynamics. Scalability of both methods is shown up to 8192 CPU cores, with near ideal scaling for DG up to 32768 cores.

Keywords: compressible flow, Euler, Lagrange, Legendre, Navier-Stokes, nonhydrostatic, parallelization.

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 09 MAR 2012		2. REPORT TYPE		3. DATES COVERED 00-00-2012 to 00-00-2012	
4. TITLE AND SUBTITLE Continuous and Discontinuous Galerkin Methods for a Scalable Three-Dimensional Nonhydrostatic Atmospheric Model: Limited-Area Mode			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School, Department of Applied Mathematics, 833 Dyer Road, Monterey, CA, 93943			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This paper describes a unified, element based Galerkin (EBG) framework for a three-dimensional, nonhydrostatic model for the atmosphere. In general EBG methods possess high-order accuracy, geometrical flexibility, excellent dispersion properties and good scalability. Our nonhydrostatic model, based on the compressible Euler equations, is appropriate for both limited-area and global atmospheric simulations. Both a Continuous Galerkin (CG), or spectral element, and discontinuous Galerkin (DG) model are considered using hexahedral elements. The formulation is suitable for both global and limited-area atmospheric modeling, although we restrict our attention to 3D limited-area phenomena in this study; global atmospheric simulations will be presented in a follow-up paper. Domain decomposition and communication algorithms utilized by both our CG and DG models are presented. The communication volume and exchange algorithms for CG and DG are compared and contrasted. Numerical verification of the model is performed using two test cases: flow past a 3D mountain and buoyant convection of a bubble in a neutral atmosphere; these tests indicate that both CG and DG can simulate the necessary physics of dry atmospheric dynamics. Scalability of both methods is shown up to 8192 CPU cores, with near ideal scaling for DG up to 32768 cores.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 42	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1. Introduction

As the resolution of numerical weather prediction (NWP) models increase, nonhydrostatic effects become relevant. Almost all current limited-area models utilize a nonhydrostatic core, while global models are currently transitioning from the hydrostatic to the nonhydrostatic regime. A host of challenging, non-trivial numerical problems arise when one enters the nonhydrostatic regime: these include 1) choosing the appropriate equation set to ensure efficiency, accuracy, and conservation, 2) effectively resolving multi-scale flow features, that may require adaptive mesh refinement (AMR), 3) developing efficient time-integrators and/or “soundproof” [23, 30] equation sets to confront the fast acoustic and gravity waves, and 4) developing scalable parallel codes for shared, distributed, and hybrid architectures based on items 1) -3).

Virtually all current nonhydrostatic NWP models are based on a combination of finite-difference discretization in space and either split-explicit or semi-implicit (i.e., implicit-explicit or IMEX) discretization in time. Examples include WRF [24] (NCAR), Lokal Modell [32] (DWD), COAMPS [19] (US Navy), and UM [3] (UK Met Office). Although finite difference schemes are very efficient, they may suffer from several problems, including: dispersion error (due to low-order approximations), geometrical inflexibility, and lack of scalability to large (e.g. tens of thousands) of processors. To overcome some of the limitations of finite-difference approximations, several emerging NWP models utilize finite-volume spatial discretization, such as MPAS [39] and MCORE [40], which utilize polynomial reconstruction of the inter-element fluxes. In order to achieve higher-order accuracy, these finite-volume based methods utilize a larger halo, which may impede scalability at high processor counts. Other approaches within a finite volume framework include the Flux-Based Characteristic Semi-Lagrangian (FBCSL) method [29], which uses the method of characteristics to enable larger time-steps that may present an obstacle to scalability.

Alternatively, element-based Galerkin (EBG) methods have recently been proposed for several next generation non-hydrostatic NWP models [14, 15, 31], as well as for hydrostatic models using both continuous Galerkin (CG) [7, 6] and discontinuous Galerkin (DG) [28] formulations. We note that the hydrostatic models are only valid for horizontal resolutions coarser than 10 km and for this reason are not viable options for mesoscale (regional) modeling. EBG methods possess several desirable attributes for future NWP nonhydrostatic models, such as high-order accuracy, geometrical flexibility,

whereby the solver is completely independent of the grid, excellent dispersion properties [26] and minimal communication overhead within a parallel implementation. High-order accuracy, which allows fine-scale atmospheric flow features to be resolved, is achieved by representing the prognostic variables via a polynomial basis function expansion within each element. Geometrical flexibility, which is inherent to all EBGs (both low and high order), is advantageous since any terrain-following coordinate may be used within the same solver; also, both static and dynamic adaptivity may be retro-fitted to the existing solver. Finally, in a distributed memory environment (e.g. cluster), low communication overhead is critical to maintain linear scalability to hundreds of thousands of processor cores. In our previous work, high-order accuracy and geometrical flexibility were addressed within explicit [14] and IMEX [15] frameworks for 2D (x - z slices) problems using a serial implementation. However, parallel implementation was not explicitly addressed. The purpose of the present work is to extend the work begun in [14], [31], and [15] to realistic three-dimensional (3D) domains using a parallel, MPI-based implementation.

The present work is guided by a need for highly scalable models in distributed memory environments. In the past five years, clock speeds of processors have remained stagnant; to achieve increased floating-point performance, chip manufacturers have developed multiple core processors that allow many threads to execute in parallel. In tandem, high performance computing (HPC) has evolved towards clusters with processors counts exceeding 100,000. As we approach the exaflop era, core counts are expected to approach 1,000,000. Therefore, next-generation NWP models must be based upon scalable numerical methods that allow arbitrarily large processor counts with minimal communication overhead. EBG methods, both CG and DG, have proved effective in this respect in modeling massive biological flow [17], modeling the hydrostatic atmosphere [11, 16], incompressible flows using low-order finite elements [20], and geodynamical problems [41].

This paper presents a scalable, 3D nonhydrostatic atmospheric model based on a unified element based Galerkin (EBG) method targeted toward distributed memory architectures; to our knowledge, these are the first 3D continuous Galerkin (spectral element) and discontinuous Galerkin models for a nonhydrostatic atmosphere. We are developing a unified dynamical core appropriate for both mesoscale and global simulations; this nonhydrostatic unified model of the atmosphere we call NUMA. The current implementation utilizes tensor products of Lagrange polynomials in a hexahedral grid for

maximum computational efficiency; however, more flexible grids based on either tetrahedra or triangular prisms may be incorporated into future versions. The remainder of this paper is structured as follows. In Section 2, we formulate the nonhydrostatic compressible Euler equations (Set 2NC and 2C from [15]), which constitute the governing equations of our dynamical core. To ensure numerical stability, these nonhydrostatic equations are solved about a hydrostatic base state. In Section 3, we present both the continuous and discontinuous Galerkin discretization, along with the explicit time-integrator, boundary conditions, and artificial diffusion. Section 4, which forms the core of the paper, outlines the parallelization algorithm for both the CG and DG methods. The communication volume and memory access patterns of both CG and DG are compared in this section. Numerical results for a 3D linear hydrostatic mountain and a 3D rising thermal bubble are shown in Section 5 along with the results of the scalability experiments. Scalability is demonstrated to 8192 processor cores for both the CG and DG methods, with near ideal scaling for DG up to 32768 cores.

2. Governing Equations

We consider the fully compressible, nonhydrostatic Euler equations in both conservative and non-conservative form. These equations (Sets 2NC and 2C), which are valid for spatial resolutions finer than 10 km, have previously been considered in [14] for 2D limited-area atmospheric flows. Set 2NC is considered within a continuous Galerkin (CG) framework, whereas Set 2C is considered within a discontinuous Galerkin (DG) framework.

2.1. Non-Conservative Form (Set 2NC)

Of the five equation sets considered in [15], set 2NC proved to be both computationally efficient and provides acceptable mass and energy conservation properties; in addition, replacing the advection operator in the momentum equation allows for formal conservation of energy up to time truncation error. Both diabatic forcing and the effects of moisture are neglected; in other words, we consider a *dry* dynamical core without sub-grid scale turbulence closure. In the present study, we consider three-dimensional flow in Cartesian coordinates (x - y - z) subject to gravitational and Coriolis forces, yielding

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

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \frac{1}{\rho} \nabla P + g \hat{\mathbf{k}} + \mathbf{f} \times \mathbf{u} = 0 \quad (1b)$$

$$\frac{\partial \theta}{\partial t} + \mathbf{u} \cdot \nabla \theta = 0 \quad (1c)$$

with the prognostic variables defined as $(\rho, \mathbf{u}^T, \theta)$, where ρ is density, $\mathbf{u} = (u, v, w)^T$ is velocity, and θ is potential temperature; the superscript T denotes the transpose operator. In addition, P is pressure, g is the gravitational constant, $\mathbf{f} = 2\Omega \hat{\mathbf{k}}$ is the Coriolis parameter (with Ω the angular frequency of the earth), and $\hat{\mathbf{k}}$ is the unit vector in the z direction. Eq. (1a) enforces mass conservation, Eq. (1b) enforces conservation of momentum, and Eq. (1c) enforces conservation of entropy. To close the system of conservation laws given by Eq. (1), a thermodynamic equation of state is required. We utilize the ideal gas law given by

$$P = P_A \left(\frac{\rho R \theta}{P_A} \right)^\gamma \quad (2)$$

where P_A is the atmospheric pressure at ground, $R = c_p - c_v$ is the ideal gas constant, and $\gamma \equiv \frac{c_p}{c_v} \approx 1.4$ is the ratio of specific heats. In three dimensions, Eqs. (1) and (2) constitute a closed system of nonlinear partial differential equations in five unknowns.

To facilitate the solution of the compressible Euler equations and maintain numerical stability, we split the density, pressure, and potential temperatures about their mean hydrostatic values:

$$\rho(x, y, z, t) = \rho_0(z) + \rho'(x, y, z, t) \quad (3a)$$

$$\theta(x, y, z, t) = \theta_0(z) + \theta'(x, y, z, t) \quad (3b)$$

$$P(x, y, z, t) = P_0(z) + P'(x, y, z, t) \quad (3c)$$

where ρ_0 , θ_0 , and P_0 are the hydrostatic reference states. For all the test cases considered in this paper, the reference states are functions of the vertical coordinate z ; however, more sophisticated test cases require reference states that are functions of x , y , and z . Inserting Eq. (3) into Eq. (1) and applying hydrostatic balance

$$\frac{dP_0}{dz} = -\rho_0 g \quad (4)$$

yields the system

$$\frac{\partial \rho'}{\partial t} + \mathbf{u} \cdot \nabla \rho' + \mathbf{u} \cdot \nabla \rho_0 + (\rho' + \rho_0) \nabla \cdot \mathbf{u} = 0 \quad (5a)$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \frac{1}{\rho' + \rho_0} \nabla P' + \frac{\rho'}{\rho' + \rho_0} g \hat{\mathbf{k}} + \mathbf{f} \times \mathbf{u} = 0 \quad (5b)$$

$$\frac{\partial \theta'}{\partial t} + \mathbf{u} \cdot \nabla \theta' + \mathbf{u} \cdot \nabla \theta_0 = 0. \quad (5c)$$

Defining a solution vector $\mathbf{q} = (\rho', \mathbf{u}^T, \theta')^T$, Eq. (5) is written in condensed form as

$$\frac{\partial \mathbf{q}}{\partial t} = S^{2NC}(\mathbf{q}) \quad (6)$$

where $S(\mathbf{q})$ is a nonlinear, first-order differential operator.

2.2. Conservative Form (Set 2C)

Eq. (1) may be written in flux, or conservative form

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{U} = 0 \quad (7a)$$

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \left(\frac{\mathbf{U} \otimes \mathbf{U}}{\rho} + P \mathbf{I}_3 \right) + \rho g \hat{\mathbf{k}} + \mathbf{f} \times \mathbf{U} = 0 \quad (7b)$$

$$\frac{\partial \Theta}{\partial t} + \nabla \cdot \left(\frac{\Theta \mathbf{U}}{\rho} \right) = 0 \quad (7c)$$

complemented by the equation of state

$$P = P_A \left(\frac{R\Theta}{P_A} \right)^\gamma \quad (8)$$

with the prognostic variables defined as $(\rho, \mathbf{U}^T, \Theta)^T$, where ρ is density, $\mathbf{U} = \rho \mathbf{u}$ is momentum, $\mathbf{u} = (u, v, w)^T$ is velocity, and $\Theta = \rho \theta$ is density potential temperature. In addition, the operator $\otimes$ denotes the tensor product and $\mathbf{I}_3$ is the rank-3 identity matrix. The variables in these equations are split in a similar fashion to Eq. (1) except that we now split the density potential temperature as follows

$$\Theta(x, y, z, t) = \Theta_0(z) + \Theta'(x, y, z, t). \quad (9)$$

Applying the hydrostatic decomposition to Eq. (7) yields

$$\frac{\partial \rho'}{\partial t} + \nabla \cdot \mathbf{U} = 0 \quad (10a)$$

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \left(\frac{\mathbf{U} \otimes \mathbf{U}}{\rho} + P' \mathbf{I}_3 \right) + \rho' g \hat{\mathbf{k}} + \mathbf{f} \times \mathbf{U} = 0 \quad (10b)$$

$$\frac{\partial \Theta'}{\partial t} + \nabla \cdot \left(\frac{\Theta \mathbf{U}}{\rho} \right) = 0 \quad (10c)$$

which is written in vector form as

$$\frac{\partial \mathbf{q}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{q}) = S(\mathbf{q}) \quad (11)$$

where $\mathbf{F}(\mathbf{q})$ is the flux tensor and the gravitational and Coriolis terms are incorporated into $S(\mathbf{q})$. Finally, we can write this equation in the condensed form as

$$\frac{\partial \mathbf{q}}{\partial t} = S^{2C}(\mathbf{q}). \quad (12)$$

3. Numerical Methods

In this section, we briefly discuss the numerical discretization used by NUMA which includes: the spatial discretization of Eq. (1) using a continuous Galerkin (CG) method and Eq. (7) using a discontinuous Galerkin (DG) method (Sec. 3.1); the explicit time-integrator used to evolve the solution forward in time (Sec. 3.2); the necessary boundary conditions required for limited-area problems (Sec. 3.3); and the artificial diffusion used to control overshoots (Sec. 3.4).

3.1. Spatial Discretization

Let us now describe the approximation of the continuous spatial operators using both CG and DG. We begin with a description of the decomposition of the global spatial domain into elements, the basis functions defined within these elements, and element-wise integrals.

3.1.1. Basis Functions, Elements, and Integrals

Element-based Galerkin methods such as the continuous Galerkin and discontinuous Galerkin methods require the decomposition of the global domain $\Omega \subset \mathcal{R}^3$ into N_e non-overlapping elements Ω_e via

$$\Omega = \bigcup_{e=1}^{N_e} \Omega_e. \quad (13)$$

In the current formulation of NUMA, we let Ω_e be hexahedra, which provides simple grid generation and efficient (fast) evaluation of the necessary differentiation and integration operators. We note, however, that Ω_e may be replaced by tetrahedra, pyramidal elements, or triangular prisms, in a future version of NUMA; while this can be done for both CG and DG, we envision doing this only for the DG version since DG does not require global matrices to be stored. CG would require the inversion of a sparse global mass matrix which would be prohibitive in three-dimensions. We also note that tetrahedral and/or pyramidal elements would produce an unstructured grid in the vertical, which would impede the incorporation of column-based microphysics and physical parameterizations. Finally, sum-factorization may be used with hexahedral elements, reducing the complexity of a 3D method based on N -th order polynomials from $\mathcal{O}(N^6)$ to $\mathcal{O}(N^4)$, making EBG methods one to three orders of magnitude more efficient for typical values of $3 \leq N \leq 10$.

Letting the unit cube $(\xi, \eta, \zeta) \in E = [-1, 1]^3$ be the reference hexahedral element, a transformation $\mathcal{F}_e : \Omega_e \rightarrow E$ mapping physical space to computational space is defined for each element, yielding $(x, y, z) = \mathcal{F}_e(\xi, \eta, \zeta)$ where the associated Jacobian of $\mathcal{F}_e$ is denoted by J_e .

Within each element Ω_e , a finite-dimensional approximation $\mathbf{q}_N$ is formed by expanding $\mathbf{q}(\mathbf{x}, t)$ in basis functions $\psi_j(\mathbf{x})$ such that

$$\mathbf{q}_N^{(e)}(\mathbf{x}, t) = \sum_{j=1}^{M_N} \psi_j(\mathbf{x}) \mathbf{q}_j^{(e)}(t) \quad (14)$$

where $M_N = (N + 1)^3$ is the number of nodes per element, N is the order of the basis functions, and the superscript (e) denotes element-wise (local) values. For basis functions, we construct tensor products of Lagrange polynomials given by

$$\psi_i(\mathbf{x}) = h_\alpha(\xi) \otimes h_\beta(\eta) \otimes h_\gamma(\zeta) \quad (15)$$

where $h_\alpha(\xi)$ is the Lagrange polynomial associated with the Legendre-Gauss-Lobatto (LGL) points ξ_i and (ξ, η, ζ) are functions of the physical variable $\mathbf{x}$. These LGL points satisfy

$$(1 - \xi^2)P'_N(\xi) = 0 \quad (16)$$

where $P_N(\xi)$ is the N -th order Legendre polynomial. We can now use all of this machinery defined to construct discrete approximations of the continuous

spatial derivatives. For example, the derivative of $\mathbf{q}$ can be obtained by differentiating Eq. (14) as follows

$$\nabla \mathbf{q}_N^{(e)} x(\mathbf{x}, t) = \sum_{j=1}^{M_N} \nabla \psi_j x(\mathbf{x}) \mathbf{q}_j^{(e)}(t) \quad (17)$$

which, after multiplying by the test (basis) function ψ and integrating within an element yields

$$\int_{\Omega_e} \psi_i \sum_{j=1}^{M_N} \nabla \psi_j x(\mathbf{x}) \mathbf{q}_j^{(e)}(t) d\Omega_e \quad (18)$$

where we shall use Nth-degree LGL integration points. We will make use of these definitions in order to define the CG and DG approximations.

3.1.2. Continuous Galerkin

For CG we shall use Eq. (1) as our governing equations. Using Eq. (14) to approximate $\mathbf{q}_N$, multiplying by a test function and integrating as in Eq. (18) yields the element-wise definition of the problem

$$\int_{\Omega_e} \psi_i \frac{\partial \mathbf{q}_N}{\partial t} d\Omega_e = \int_{\Omega_e} \psi_i S^{2NC}(\mathbf{q}_N) d\Omega_e \quad (19)$$

and applying the global assembly, or direct stiffness summation (DSS) operator required by all CG methods yields the weak formulation for CG: find $\mathbf{q}_N \in \mathcal{V}_N^{CG}$ such that

$$\int_{\Omega} \psi_I \frac{\partial \mathbf{q}_N}{\partial t} d\Omega = \int_{\Omega} \psi_I S^{2NC}(\mathbf{q}_N) d\Omega \quad \forall \psi \in \mathcal{V}_N^{CG} \quad (20)$$

where

$$\mathcal{V}_N^{CG} = \{\psi \in H^1(\Omega) : \psi \in P^N(I), e = 1, \dots, N_e\} \quad (21)$$

and P^N denotes the space of all polynomials of degree N . Notice that the requirement $\psi \in H^1(\Omega)$ implies $\mathcal{V}_N^{CG} \subset C^0(\Omega)$. The DSS operator $\bigwedge_{e=1}^{N_e}$ forms global matrices from the element (local) matrices. For example, for the local mass matrix defined as

$$M_{ij}^{(e)} = \int_{\Omega_e} \psi_i \psi_j d\Omega_e \quad (22)$$

the DSS operator produces the global mass matrix

$$M_{IJ} = \bigwedge_{e=1}^{N_e} M_{ij}^{(e)} \equiv \bigwedge_{e=1}^{N_e} \int_{\Omega_e} \psi_i \psi_j d\Omega_e \quad (23)$$

where the DSS sums the contribution of all the elements $e = 1, \dots, N_e$ and $i = 1, \dots, M_N$ and stores them in the global grid points $I = 1, \dots, N_P$ as follows $(i, e) \rightarrow I$. Extracting the global mass matrix from Eq. (19) and writing the operator S as a vector, results in the following compact matrix-vector form

$$\frac{\partial \mathbf{q}_I}{\partial t} = M_{IJ}^{-1} S_J^{CG}(\mathbf{q}) \quad (24)$$

where the mass matrix $M_{IJ} = \int_{\Omega} \psi_I \psi_J d\Omega$ is diagonal because the interpolation and integration points have been carefully chosen to be co-located; this approximation is valid for $N \geq 4$ while incurring a small error in integration [13]. Denoting the right-hand side (RHS) of Eq. (24) by $R_I(\mathbf{q})$, Eq. (24) is expressed as

$$\frac{\partial \mathbf{q}_I}{\partial t} = \bigwedge_{e=1}^{N_e} R_i^{(e)}(\mathbf{q}_i^{(e)}) . \quad (25)$$

Note that the DSS operator maps local, element-wise coordinates (i, e) to global coordinates I . Eq. (25) forms the core of the spectral element method, allowing local, element-wise information $\mathbf{q}_i^{(e)}$ to propagate to adjacent elements via the DSS operator.

3.1.3. Discontinuous Galerkin

For DG we use Eq. (7) because the DG method requires the equations to be in conservation form. We have also developed a CG code with Set 2C but we use Set 2NC for CG because it represents the optimal equation set for our needs as described in [15].

Beginning with Eq. (11), using the basis function expansion, Eq. (14), multiplying by a test function, and integrating element-wise yields

$$\int_{\Omega_e} \psi_i \frac{\partial \mathbf{q}_N^{(e)}}{\partial t} d\Omega_e + \int_{\Omega_e} \psi_i \nabla \cdot \mathbf{F}^{(e)}(\mathbf{q}_N) d\Omega_e = \int_{\Omega_e} \psi_i S^{(e)}(\mathbf{q}_N) d\Omega_e. \quad (26)$$

Integrating the divergence of the flux by parts yields the weak formulation

for DG: find $\mathbf{q}_N^{(e)} \in \mathcal{V}_N^{DG}$

$$\begin{aligned} \int_{\Omega_e} \psi_i \frac{\partial \mathbf{q}_N^{(e)}}{\partial t} d\Omega_e + \sum_{k=1}^{N_{faces}} \int_{\Gamma_e} \psi_i \hat{\mathbf{n}}^{(e,k)} \cdot \mathbf{F}^{(e,k)}(\mathbf{q}_N) d\Gamma_e - \int_{\Omega_e} \nabla \psi_i \cdot \mathbf{F}^{(e)}(\mathbf{q}_N) d\Omega_e \\ = \int_{\Omega_e} \psi_i S^{(e)}(\mathbf{q}_N) d\Omega_e \quad \forall \psi \in \mathcal{V}_N^{DG} \end{aligned} \quad (27)$$

where the DG finite-dimensional space is defined as

$$\mathcal{V}_N^{DG} = \{ \psi \in L^2(\Omega) : \psi \in P^N(I), e = 1, \dots, N_e \}. \quad (28)$$

Notice that, contrary to Eq. (21), there is no global continuity requirement, so that $\mathcal{V}_N^{DG} \not\subset C^0(\Omega)$. This is possible because in Eq. (27) differentiability is required separately within each element, and not within the entire domain Ω . The coupling between neighboring elements is then recovered through the numerical flux (or Riemann solver) $\mathbf{F}^{(e,k)}$, which is required to be a single valued function on the inter-element boundaries and the precise definition of which is given below.

For simplicity we use the Rusanov flux defined as

$$\mathbf{F}^{(e,k)} = \frac{1}{2} [\mathbf{F}^{(e)} + \mathbf{F}^{(k)} - \hat{\mathbf{n}}^{(e,k)} c_{max} (\mathbf{q}^{(k)} - \mathbf{q}^{(e)})] \quad (29)$$

where c_{max} is the maximum wave speed of the Euler equations (i.e., maximum eigenvalue of the Jacobian matrix) which turns out to be $c_{max} = |\hat{\mathbf{n}} \cdot \mathbf{u}| + a$ where a denotes the speed of sound. Note that other types of numerical fluxes are possible including multi-dimensional Riemann solvers (e.g., [9]). We have initially chosen this simple Rusanov flux because it vastly simplifies the parallel strategy. Using a one-dimensional flux, as we have, reduces the element communication to edge neighbors but restricts the maximum allowable time-step. Conversely, using a multi-dimensional Riemann solver will increase the maximum allowable time-step but will also increase the size of the communication stencil. The vices and virtues of this approach need to be studied in detail. Such a study is beyond the scope of the present work; we leave this topic for future work.

3.2. Explicit RK methods

We implemented a strong stability preserving (SSP) Runge-Kutta third-order, five stage time-integrator (RK35) proposed in [37]. This time-integrator

is stable for (acoustic) Courant numbers of 1.3 or less when using a continuous Galerkin discretization, whereas for DG, the time-integrator is stable for (acoustic) Courant numbers of 0.85 or less.

We emphasize that this explicit RK method will not be used in an operational setting due to the stringent CFL requirement; rather, we are developing IMEX methods which allow much larger time-steps. Nevertheless, this particular time integrator was chosen for the following reasons: 1) to provide high-order accuracy in time to complement the high-order spatial discretization, 2) to minimize the amount of numerical dissipation, 3) to allow a relatively larger time-step (with respect to explicit methods), and 4) to provide stable solutions for both CG and DG methods. Criteria 1) and 2) are necessary for resolving fine-scale flow features present in many atmospheric phenomena, whereas criterion 3) is necessary for model efficiency, and criterion 4) is required in order to facilitate the comparison of CG and DG. In a previous study involving nonhydrostatic modeling [5], all available third-order multi-stage methods in addition to an assortment of multi-step methods were compared and RK35 was found to be the most efficient. In addition, RK35, unlike leapfrog, does not have a computational mode and therefore does not require a time filter (DG also would not work with leapfrog due to the eigenvalues for DG not residing along the imaginary axis as they do for CG). Note, however, that within a parallel setting, each RK stage requires parallel communication, thus making this choice of time-integrator relatively expensive. However, this extra expense is tolerated in order to maintain high-order accuracy in time. It should also be mentioned that we have chosen an explicit method in order to simplify and focus the discussion of CG and DG with respect to scalability. The optimal scalability will be achieved by an explicit method. A study on the scalability of IMEX methods requires a discussion of a number of topics including: implicit time-integrators, the choice of iterative solver, and specific preconditioning strategies; we leave this for future work.

3.3. Boundary Conditions

Limited-area atmospheric models, such as mesoscale codes, typically require two types of boundary conditions: no-flux boundary conditions (NFBCs), that mimic impenetrable objects (e.g., the ground) and non-reflecting boundary conditions (NRBCs), that mimic an infinite domain (e.g. the top of the atmosphere) by allowing waves to propagate out of the computational domain without generating spurious reflections. In this section, we outline

how NFBCs and NRBCs are imposed for both CG and DG. For additional details, see [14, 15].

3.3.1. No-Flux Boundary Conditions

All of our test cases utilize NFBCs on the bottom boundary, while some test cases (such as the rising thermal bubble) utilize NFBCs on other boundaries as well. For CG, we enforce

$$\hat{\mathbf{n}} \cdot \mathbf{u} = 0 \quad (30)$$

for all points on the boundary Γ , where $\hat{\mathbf{n}}$ is the outward pointing unit normal vector on Γ . In order to apply the NFBC to the prognostic vector $\mathbf{q}$, we augment $\hat{\mathbf{n}}$ to R^5 via $\hat{\mathbf{n}} = (0, \hat{\mathbf{n}}^T, 0)^T$, yielding $\hat{\mathbf{n}} \cdot \mathbf{q} = 0$. To apply these boundary conditions in the *strong* sense, we construct a 3 by 3 projection matrix P via

$$P = \begin{pmatrix} 1 - n_x^2 & -n_x n_y & -n_x n_z \\ -n_y n_x & 1 - n_y^2 & -n_y n_z \\ -n_z n_x & -n_z n_y & 1 - n_z^2 \end{pmatrix}. \quad (31)$$

This matrix is constructed during the initialization phase and applied to the RHS operator after each time-step.

On the other hand, for DG the NFBCs are imposed via the numerical flux as follows. On boundary edges where NFBCs are to be imposed, a ghost-cell is constructed on the other side of the wall that has a velocity vector the negative of the interior element. In other words,

$$\mathbf{U}^{(k)} = -\mathbf{U}^{(e)}$$

where k in this case denotes the edge that is positioned on a no-flux boundary; note that this ensures that the no-flux condition $\hat{\mathbf{n}} \cdot \mathbf{U}$ is satisfied. The scalar variables are copied directly such that $\rho^{(k)} = \rho^{(e)}$, etc.

3.3.2. Non-Reflecting Boundary Conditions

In an operational mesoscale NWP model, the four lateral and the top boundaries should mimic an open domain. That is, waves should smoothly exit the domain without reflection; in addition, information from outside the domain should be allowed to enter the domain of interest. Mathematically modeling this behavior is non-trivial and has attracted the attention of researchers in many disciplines [4, 18, 27]. In our model, we utilize a simple, albeit, effective absorbing sponge layer method. The computational domain

is surrounded by a layer with Newtonian relaxation coefficients $\alpha(\mathbf{x})$ and $\beta(\mathbf{x})$ such that $\alpha = 1$ and $\beta = 0$ in the domain of interest, while $\alpha \rightarrow 0$ and $\beta \rightarrow 1$ as the boundary is approached. Specifically, for the top boundary, we choose

$$\beta = \left(\frac{z - z_s}{z_t - z_s} \right)^4 \quad (32)$$

where z_t is the vertical height of the domain and z_s is the bottom of the sponge layer. Similar functions are used for the lateral boundaries of the domain. Once the sponge layer is constructed, the numerical solution $\tilde{\mathbf{q}}$ given by the RHS operator is relaxed to some known solution at the boundary $\mathbf{q}_b$ via

$$\mathbf{q} = \alpha(\mathbf{x})\tilde{\mathbf{q}} + \beta(\mathbf{x})\mathbf{q}_b. \quad (33)$$

For problems under consideration in this paper, $\mathbf{q}_b$ is a far-field condition (e.g. known wind velocity and zero perturbation of scalars).

3.4. Artificial Diffusion

To add a certain level of numerical dissipation we have implemented second order artificial diffusion operators, although hyper-viscosity is also possible. For Eq. (6) we add the following right-hand-side operator

$$\begin{pmatrix} 0 \\ \nu \nabla \cdot (\nabla \mathbf{u}) \\ \nu \nabla \cdot (\nabla \theta') \end{pmatrix}$$

and for Eq. (12) we add the operator

$$\begin{pmatrix} 0 \\ \nu \nabla \cdot (\rho \nabla \mathbf{u}) \\ \nu \nabla \cdot (\rho \nabla \theta') \end{pmatrix}.$$

Note that the artificial diffusion operator applied to the potential temperature equation must act only on the potential temperature perturbation, θ' , because if it is applied to the full potential temperature variable then it will smoothen the background reference field which must remain intact in order to maintain hydrostatic balance.

4. Parallel Implementations of EBG Methods

NUMA is an MPI-based code targeted toward distributed memory architectures (e.g. clusters). In this section, we discuss the parallel implementation of NUMA, including: a description of the domain-decomposition (Sec. 4.1), communication strategies for both CG and DG (Secs. 4.2 and 4.3), a comparison of these communication volumes (Sec. 4.4), and a discussion of memory access of these EBG methods (Sec. 4.5).

4.1. Domain Decomposition

We decompose Ω into N_p processor elements (PE) Ω_p that consist of local elements $\Omega_{e'}^{(p)}$. Mathematically, we rewrite Eq. (13) as

$$\Omega = \bigcup_{p=1}^{N_p} \bigcup_{e'=1}^{N_e^{(p)}} \Omega_{e'}^{(p)} \quad (34)$$

where $N_e^{(p)}$ is the number of local elements residing on PE p . Since the DSS operator for CG requires global elements e and the flux integrals for DG require global faces f , we must also construct local to global mappings for both elements $e = LGE^{(p)}(e')$ and faces $f = LGF^{(p)}(f')$ that map local elements e' and faces f' on processor p to global elements e and faces f residing on the global domain Ω .

A guiding principle in the construction of NUMA is to maintain independence between grid generation and the CG/DG solver; hence, domain decomposition should be as general as possible and not be constrained by the choice of grid. Therefore, we have implemented a domain decomposition strategy based on the widely used METIS graph partitioning library [21]. METIS requires an adjacency graph where the N_e vertices are elements Ω_e and the “edges” denote the connectivity between elements. Since the DSS operator requires information from elements that share nodes, the “edges” include all forms of geometric connectivity (faces, edges, and vertices). For maximum flexibility, we construct a weighted adjacency graph $G' = (V, E)$ with adjacency matrix A' of size N_e by N_e defined via

$$a'_{ij} = \begin{cases} v & \text{if } i \text{ and } j \text{ are vertex neighbors} \\ e & \text{if } i \text{ and } j \text{ are edge neighbors} \\ 1 & \text{if } i \text{ and } j \text{ are faces neighbors} \end{cases} \quad (35)$$

where v and e are arbitrary weights. Although not explored in this paper, the values of v and e may be chosen to construct a machine-optimal weighted adjacency matrix. Once A' is constructed the adjacency matrix A for the graph $G = (V, F)$, where F are geometrical faces, is simply $a_{ij} = 1$ if $a'_{ij} = 1$ and $a_{ij} = 0$ otherwise. Two example connectivity graphs for a 2D grid are shown in Figure 1, showing both edge and vertex connectivity. The associated 9 by 9 adjacency matrix has nodes with degree ranging from 3 (for the corner nodes) to 8 (for the central node). In contrast, the flux integrals only require face adjacency information; hence, our DG code only utilizes the standard adjacency matrix A . Also shown in Fig. 1 is the adjacency matrix associated with DG, which only shows edge (face) adjacency. These adjacency matrices, along with the number of processors N_p are then passed to METIS, which returns a partition $\mathcal{P} : V \rightarrow \{1, 2, \dots, N_p\}$, that maps global elements to processors. This mapping is then used to construct the local to global mappings LG necessary for global assembly in Eq. (25).

In addition to the element adjacency graph G , a processor-element adjacency graph $G_P = (V_P, E_P)$ is constructed, where the vertices V_P are processor elements and the edges E_P are the connectivity between processor elements. For a CG method, E_P includes vertex, edge, and face connections between processor elements. The N_P by N_P adjacency matrix $A^{(P)}$ associated with G_P is derived from A' as follows: the element $a_{ij}^{(P)} = 1$ if the intersection of all rows i' and columns j' of A' such that $i = \mathcal{P}(i')$ and $j = \mathcal{P}(j')$ has at least one nonzero element; otherwise, $a_{ij}^{(P)} = 0$. From the inter-processor adjacency matrix $A^{(P)}$, the necessary communication data structures are constructed. Specifically, the neighbors of processor element i are simply the non-zero columns of row i , while the number of neighbors for element i is given by $\sum_{j=1}^{N_P} a_{ij}^{(P)}$. Thus, a low-communication partition will have an adjacency matrix $A^{(P)}$ that is as sparse as possible.

4.2. Parallelization: CG

We first consider the parallelization of CG. The global DSS operator in Eq. (25) requires inter-processor communication due to the C^0 requirement of elements at processor element boundaries. A global DSS is required in two parts of the code: construction of the mass matrix and construction of the right-hand side (RHS). To perform this communication, we first construct the boundary nodes of each processor (excluding the physical boundary where BCs are applied). Denote the boundary of processor element i by $\partial\Omega_i$. In




Figure 1: Example 2D grid (left), the CG associated adjacency graph (center), and DG adjacency graph (right). Since CG methods utilize nodal communication, the generalized adjacency graph includes both edge and vertex connectivity, with a maximum degree of eight. For 3D, structured, Cartesian grids, the maximum degree is 26. In contrast, typical DG methods utilize face-based communication and hence only require the adjacency graph shown in the right panel. For 3D Cartesian grids the maximum degree for DG is only 6.

order to communicate between processor elements, we construct two data structures *send* and *recv*. Consider a particular processor i and neighboring processor j . *send* contains the local nodes on processor i that are sent to each neighboring processor j , while *recv* contains the local nodes on processor j that must be sent to i . In order to construct *send*, we utilize the method shown in Algorithm 1. The corresponding *recv* structure contains the same grid points as *send*, although they may be ordered differently.

Algorithm 1 Construction of MPI send/receive communication data structures for both CG and DG.

```

for all NBHs  $j$  of  $i$  do
    MPISEND  $\partial\Omega_i^G \leftarrow LGE^{(i)}(\partial\Omega_i)$  to proc  $j$ 
    MPIRECV  $\partial\Omega_j^G \leftarrow LGE^{(j)}(\partial\Omega_j)$  to proc  $i$ 
     $B \leftarrow \partial\Omega_i^G \cap \partial\Omega_j^G$ 
     $send(j) \leftarrow [LGE^{(i)}]^{-1}(B)$ 
end for

```

Once these data structures have been constructed, the global mass matrix and RHS operators may be constructed via a two step process. The global

mass matrix M_{IJ} and RHS operator in Eq. (25) are decomposed as

$$M_{IJ} = \bigwedge_{e=1}^{N_e} M_{ij}^{(e)} = \bigwedge_{n_p=1}^{N_p} \bigwedge_{e'=1}^{N_e^{(p)}} M_{ij}^{(e')} \text{ and} \quad (36a)$$

$$R_I = \bigwedge_{e=1}^{N_e} R_i^{(e)} = \bigwedge_{n_p=1}^{N_p} \bigwedge_{e'=1}^{N_e^{(p)}} R_i^{(e')}. \quad (36b)$$

Hence, the global DSS is decomposed into a local, on-processor DSS and a global DSS, that requires inter-processor communication. In this way, global continuity between elements is preserved during the construction of each RHS. Note that the communication stencil is simply the boundary of each processor element $\partial\Omega^{(i)}$ and is independent of the polynomial order N ; that is, unlike higher-order finite difference or finite volume methods, the spectral-element method is *halo-free*. We note however, that the message size grows with the surface area of the processor element, or N^2 . In summary, the global DSS procedure may be summarized as follows: The global DSS operation is

Algorithm 2 The Direct Stiffness Summation (DSS) Operation in CG

1. Perform a local DSS on processor i .
 2. Exchange boundary points between processors i and all processor neighbors j using (isynchronous) *isend* and *irecv*.
 3. Perform a global DSS using the boundary data received from neighbors j .
-

represented schematically in the left panel of Figure 2 in a 2D setting. To simplify the discussion, each processor is assumed to own one element. In order to construct the RHS operator on the boundary (red dots), the element E (green) requires nodal information for its 8 nodal neighbors. These neighbors include both edge neighbors (2, 4, 6, and 8) and vertex neighbors (1, 3, 5, and 7). In a 3D setting, a hexahedral element may have (a maximum) of 6 face neighbors, 12 edge neighbors, and 8 vertex neighbors, for a total of 26 nodal neighbors. Note that for tetrahedral grids, the number of vertex neighbors may be much greater. To reduce this communication cost, METIS is used to reduce the total number of vertex neighbors; in a typical 3D partition, a processor element lying away from a physical boundary has 16 or fewer nodal neighbors; hence, the METIS-based decomposition further




Figure 2: Computational stencils for a) CG and b) DG in a 2D setting, where each processor element owns one element. In the left panel, the element E (green) requires nodal information for its 8 nodal neighbors in order to construct the DSS operator on the boundary (red dots). In the right panel, the same element E only requires nodal information for its 4 face neighbors in order to construct flux integrals. In both methods, the interior nodes (yellow dots) do not need to be communicated to adjacent processors, resulting in a *halo-free* algorithm.

reduces communication cost relative to a naive geometric domain decomposition. In particular, we use the routine `METIS_PartGraphRecursive`, which reduces the edge-cut of the graph G .

EBG methods possess purely local communication stencils. Referring to the left panel of Fig. 2, the interior nodes (yellow dots) do not need to be communicated to adjacent processors, resulting in a *halo-free* algorithm. Thus unlike high-order finite difference and finite volume schemes, spectral elements may achieve spectral accuracy without sacrificing their local character.

4.3. Parallelization: DG

Parallelization of DG is much more straightforward than CG. Since there are no global mass matrices, all that is required is a way to communicate boundary data between processor elements. In addition to a local/global element-mapping, an additional local/global face structure is also required. Faces on each processor element (excluding boundary faces) are classified into two types: intra-processor and inter-processor. The necessary send and

receive data structures are then constructed by calculating the intersection of face boundaries on each PE as described in a manner similar to Algorithm 1. Unlike CG, however, the DG algorithm requires only the set of faces f which need to be sent to adjacent processors and not the actual grid points. Once these maps and data structures are constructed, the flux operators are evaluated using the following algorithm, which employs interleaved computation/communication:

Algorithm 3 Interleaved Computation of the Flux Operator in DG

1. Send/Receive boundary data via non-blocking MPI Send/Receive.
While waiting for boundary data do:
 2. Compute volume integrals.
 3. Compute intra-processor flux integrals.
 4. `MPI_Wait()`
 5. Compute inter-processor flux integrals.
 6. Multiply by element-wise inverse mass matrix.
-

The communication stencil is illustrated geometrically in the right panel of Figure 2. Like CG, DG algorithms do not require internal nodes to be communicated to adjacent processors and is hence *halo-less*. Moreover, since all communication is based on fluxes, only elements which share adjacent faces need to communicate. This compact stencil has a crucial impact on both the communication volume and memory access patterns of DG when compared to CG. These issues are explored in the next two sections.

We note that the focus of this study is a comparison of CG and DG from an algorithmic point of view and that the parallelization algorithms outlined in this and the previous section are not optimal. In particular, we have not implemented sophisticated scheduling between the individual MPI processes using `MPI_Waitany` or `Mpi_Waitsome`, nor have we exploited topology aware communication. We are currently addressing these implementation details in our codes. Therefore, it must be understood that both methods can indeed be further optimized; however, the results we present represent the simplest approaches for constructing parallel strategies.

4.4. Communication Volume

A simple model [8] for the communication time (or cost) of a single MPI process is

$$t_c = t_a(\alpha + \beta m) \quad (37)$$

where t_a is the time to do a floating-point operation, α is proportional to the latency (i.e., message startup cost), and β is the asymptotic transfer rate (i.e., inverse bandwidth). Typically, the latency cost α is larger than bandwidth cost β by one to three orders of magnitude, depending on the architecture of the machine. To characterize the trade-off between bandwidth and latency, the ratio β/α provides a useful metric; a typical value is $\beta/\alpha = 0.01$. Using the model given by Eq. (37) the communication overhead of both DG and CG may be quantitatively analyzed.

Consider a hexahedral processor element (PE) with N_e elements discretized with order N polynomials. Geometrically, we state that the length of a face, edge, and vertex message is

$$m_f = n_{var} N_e^{2/3} (N + 1)^2 \quad (38a)$$

$$m_e = n_{var} N_e^{1/3} (N + 1) \quad (38b)$$

$$m_v = n_{var} \quad (38c)$$

where $n_{var} = 5$ is the number of state variables.

For DG, since the inter-processor fluxes only require face communication, and since each PE has an average of six face neighbors, the total communication cost at each stage of the RK time-integrator is

$$C_{DG} = t_a \alpha \left(6 + 6 \frac{\beta}{\alpha} n_{var} N_e^{2/3} (N + 1)^2 \right). \quad (39)$$

Asymptotically, we see that the communication cost scales as $\mathcal{O}(N_e^{2/3} N^2)$. In contrast, the computation volume, which is dominated by the construction of volume integrals, scales as $\mathcal{O}(N_e N^4)$; hence, for large N and/or large N_e , the ratio of communication to computation $\mathcal{O}(N_e^{-1/3} N^{-2})$ tends to zero, thereby ensuring weak scaling.

In contrast, for CG the construction of the DSS operator requires edge and vertex communications in addition to face communication. Since for a hexahedral PE, there are 12 edges and 8 vertices, the communication cost is

$$\begin{aligned} C_{CG} &= 6t_a (\alpha + \beta m_f) + 12t_a (\alpha + \beta m_e) + 8t_a (\alpha + \beta m_v) \\ &= t_a \alpha \left(26 + 6 \frac{\beta}{\alpha} n_{var} N_e^{2/3} (N + 1)^2 + 8 \frac{\beta}{\alpha} n_{var} N_e^{1/3} (N + 1) + 8 \frac{\beta}{\alpha} n_{var} \right). \end{aligned} \quad (40)$$

As with DG, the communication cost scales as $\mathcal{O}\left(N_e^{2/3}N^2\right)$ with a similar estimate for computation; hence, in an asymptotic sense, both CG and DG have the same communication footprints. Geometrically, this can be understood since the surface area of the faces of a PE grow faster than the sizes of edges and vertices. These asymptotic estimates are not significant for practical choices of N and N_e , e.g., for $4 \leq N \leq 16$ and small values of N_e the intermediate behavior of Eqs. (39) and (40) become quite important. We note briefly that the communication footprint of CG methods may be reduced to a footprint similar to DG using the d -directional exchange algorithm proposed in [10]; however, implementation of this algorithm is non-trivial and limited to tensor product meshes, and is therefore not explored in this study.

A typical ratio of bandwidth to latency on modern computers is $\beta/\alpha = 0.01$. Consider the case of one element per MPI process, $N_e = 1$, and $N = 4$. Then $C_{CG}/C_{DG} \approx 2.6$. Hence, CG has a significant communication overhead relative to DG in practice. To understand this communication overhead more extensively, the ratio of Eqs. (40) and (39) is plotted in Figure 3 for three different values of β/α and a range of polynomial orders N . Several trends are evident from this graph. Firstly, for all bandwidth/latency ratios, the communication overhead is greatest for linear ($N = 1$) polynomials and decreases monotonically to one as $N \rightarrow \infty$. This behavior is expected from the asymptotic estimates made above. However, from a practical point of view, the machine architecture and MPI implementation, which is characterized by the ratio β/α , determines the relative communication footprint of CG to DG. As the latency costs increase relative to bandwidth costs (β/α decreases), CG incurs greater communication relative to DG. Since the general trend in HPC is toward greater bandwidth ($1/\beta$ is large therefore β is small) and greater latency (large α), then DG is expected to significantly outperform CG in terms of communication footprint as β/α decreases.

Note that for both CG and DG, the communication volume is $\mathcal{O}\left(N_e^{2/3}N^2\right)$ while the computation volume is $\mathcal{O}\left(N_eN^4\right)$. Hence, for large number of elements and (relatively) large orders N , the computation volume overwhelms the communication volume. Therefore, in the DG algorithm described in Algorithm 3, steps 2 and 3 typically require much more time than the communication in step 1. For this reason, DG is expected to scale to massive numbers of cores in a distributed memory architecture.




Figure 3: The Ratio of communication costs C_{CG}/C_{DG} plotted for transfer rate/latency ratios $\beta/\alpha = 0.1, 0.01$ and 0.001 for polynomial orders $N = 1$ to 16 .

4.5. Memory Access

Fetching data from memory poses a serious bottleneck in any HPC application. The latency associated with reading data from main memory can adversely affect the efficiency and scaling of any HPC application. In general, EBG methods exhibit a high level of data locality since the computational work is organized into discrete elements. In this section, we explore the impact of data layout and access in CG and DG and how this data access affects performance.

Figure 4 shows the grid numbering used by a) CG and b) DG for $N_e = 4$ and $N = 3$. Let us assume for the moment that CG uses a global indexing while DG must use a local indexing. This data layout has an impact on the amount of data that must be fetched from memory and stored in fast cache for an EBG. In DG, the construction of volume integrals is element-local. Since the data for each element lies in a contiguous region of memory, all of the data for each element may be fetched from slow memory loaded into fast cache at the beginning of an element-wise operation. The only operation that requires access to a neighboring element's data is the construction of flux integrals. Hence, during the construction of a RHS, the number of cache misses is small.

For CG, the state vector is typically stored in a continuous block of memory, rather than in an element-wise fashion (it is possible to use DG storage

in CG but here we discuss only the standard finite element numbering). Obviously, this data layout uses less memory than the element-wise DG layout; however, a penalty in performance may be incurred. To construct a RHS, first data must be fetched from the global state vector and stored in a local array. Cache misses may occur at all points on the boundary of the element, especially for large grids, where there is a large stride in indexing. Element-wise operations are then performed, followed by DSS, which requires updating the global state vector, thus resulting in additional cache misses. These cache misses are exacerbated in 3D, where a typical element requires data from 26 adjacent elements. Although these cache misses may be partially alleviated by grid renumbering, a CG method can never achieve the *natural* data locality of a DG method. Even if CG uses DG storage, the cache misses associated with the DSS operator cannot be circumvented; the DSS operator requires indirect memory access that requires striding through the memory in order to enforce C^0 continuity.

For optimal performance, the entire on-processor problem should fit in cache. In this situation, which only occurs for small test problems and/or large processor counts, no cache misses will happen. A rough estimate for the size of the on-processor problem is made for both DG and CG. To construct a RHS requires the following data structures: 1) a state vector, 2) reference vector, 3) RHS vector, 4) the inverse mass “matrix”, and 5) nine metric terms and two Jacobian matrices (for both volume and flux integrals). The data structures in items 1)-3) each have size $n_{var}N_e(N+1)^3$ with $n_{var} = 5$, whereas structures 4) and 5) have size $N_e(N+1)^3$, thus yielding $26N_e(N+1)^3$ real numbers. A typical 64-bit machine uses 8 bytes to store a real, yielding a problem size of $208N_e(N+1)^3$. The memory requirements for CG is slightly smaller, since data is stored using a global, rather than element-based index; however the magnitude is similar. On typical clusters (e.g., TACC’s Ranger), each core (and hence each MPI process) has a cache with size 512 KB, implying that a single element can fit into cache provided $N \leq 12$. Using this estimate, approximately 20 continuous elements fit into cache for $N = 4$, while about 3 elements fit into cache for $N = 8$.

From these estimates, the number of cache misses for both CG and DG may be estimated. For DG, the number of cache misses may vary from 0 to 6. For a very large problem or a small number of MPI processes, a cache miss will occur during each flux construction, when data from an adjacent processor is accessed, yielding 6 misses. For intermediate sizes, a cache miss does not occur when fetching data from an element with continuous numbering,




Figure 4: Grid numbering for a) CG with global indexing and b) DG with local indexing for $N_e = 4$ and $N = 3$.

yielding 2 or 4 misses; finally, for very small problems, all the data is in cache, yielding no misses. For CG, the number of misses ranges from 0 to 52. For large problems, reading and writing any data from an adjacent element will incur a cache miss, yielding $2 \times 26 = 52$ misses. For smaller problems, this count may be reduced to 48 or fewer misses, depending on the configuration of the sub-domain. Excluding sub-domains small enough to fit into cache, the number of cache misses associated with CG will be many times larger than the misses associated with DG.

5. Results

5.1. Test Cases

Although a standard set of 2D mesoscale test cases has been proposed [33] and later used within an element-based Galerkin framework [14], an analogous suite of 3D test cases has not yet been developed. Fortunately, a 3D dynamical core can be run in 2D mode by imposing symmetry in the y -dimension and periodic boundary conditions in the y -lateral boundaries. For initial verification of NUMA, we used the test cases proposed in [33]; later, we ran full 3D test cases with no y -symmetry. In the following analysis, we

consider two test cases: a 3D linear hydrostatic mountain and a 3D rising thermal bubble.

5.1.1. 3D Linear Hydrostatic Mountain

To test the 3D capabilities of NUMA and the NRBCs, we consider stratified flow past an isolated mountain as outlined in [34]. An initial horizontal flow $U = 20$ m/s goes past a mountain with orography given by

$$h(x, y) = \frac{h_0}{(x^2/a^2 + y^2/a^2 + 1)^{3/2}} \quad (41)$$

with mountain half-width $a = 10$ km and height $h_0 = 1$ m. The hydrostatic background is specified by a constant Brunt-Väisälä frequency $N_{bv} = g/\sqrt{c_p T_0}$ with ground temperature $T_0 = 250$ K. In other words, we consider an isothermal atmosphere. No flux boundary conditions are imposed on the bottom of the domain, while NRBCs are imposed on the four lateral boundaries and the top boundary. In order to verify our numerical results, several analytical results from [34] and [35] are used. First, a contour integral solution for the density perturbations ρ' valid under a linear Boussinesq approximation is utilized. Also, the velocity perturbations parallel and perpendicular to the flow (u' and v') are known for observation points near the ground ($z = 0$) (see Eqs. (39) and (41) in [34]):

$$u'(x, y, 0) = hN_{bv} \frac{x/a}{1 + x^2/a^2 + y^2/a^2} \quad (42a)$$

$$v'(x, y, 0) = hN_{bv} \frac{y/a}{1 + x^2/a^2 + y^2/a^2} \quad (42b)$$

under the same approximation.

5.1.2. 3D Rising Thermal Bubble

We consider a 3D buoyant thermal bubble rising in a neutrally stratified atmosphere [38], which is the 3D extension of the 2D thermal bubble originally considered in [36]. The hydrostatic potential temperature $\theta_0(z) = 300$ K (neutral atmosphere) is perturbed with a sphere of radius $r_c = 250$ m centered at $(x_c, y_c, z_c) = (500, 500, 260)$ m by a cosine taper given by

$$\theta' = A \left[1 + \cos \left(\frac{\pi r}{r_c} \right) \right] \quad (43)$$




Figure 5: Comparison of the density perturbations ρ' for the isolated linear hydrostatic mountain using 1) NUMA (solid line) and 2) a contour integral solution (dashed line) [35]. In panel a), the ρ' contours in the plane $y = 120000$ m are shown (the analytical model are solid while the CG model are dashed - the DG results are similar to those for CG and are not shown). In panel b), ρ' is shown as a function of z using $x = y = 120000$ m using numerical results produced by the CG and DG models, and Smith's analytical model.

where $r = \|\mathbf{x} - \mathbf{x}_c\|_2$ where $\|\cdot\|_2$ denotes the 2-norm, and $A = \frac{1}{2}$ is a constant. Unlike the test case used in [38], our problem has a C^1 initial condition, thus mitigating unphysical oscillations at the interface of the bubble. The domain is defined as $(x, y, z) \in [0, 1000]^3$ m. No-flux boundary conditions are applied on all six boundaries.

5.2. Numerical Verification

The numerical verification proceeds in three steps. In phase one, we ran the code in pseudo-2D mode using 1 element and periodic boundary conditions in the y -direction. The numerical results for the standard mesoscale suite [33] are directly compared to the results of our existing 2D model [14]. In phase two, we considered the linear isolated mountain problem, which possesses an approximate analytical solution in the form of a contour integral. In phase three, we consider a three-dimensional buoyant convection problem. Although this problem does not have an analytic solution, its qualitative behavior is well understood.




Figure 6: Comparison of the a) down-stream velocity perturbation u' and b) cross-stream velocity perturbation v' for the isolated mountain at ground level. Agreement between the CG and DG models and the analytical formulas given by Eq. (42) is satisfactory, especially for the cross-stream velocity perturbation.

5.2.1. 3D Linear Hydrostatic Mountain

In order to test the 3D operators, orography, and non-reflecting boundary conditions, flow over a 3D isolated linear hydrostatic mountain was considered. This problem may be solved under the linear Boussinesq approximation via a contour integral technique as described in [35]. In addition, closed form expressions for both the down-stream and cross-stream velocity components may be used to judge the numerical solution but only in a qualitative sense. We used 20 elements in the x and y dimensions and 10 in the z direction with eighth-order polynomials yielding effective resolutions of $\overline{\Delta x} = \overline{\Delta y} = 1.5$ km and $\overline{\Delta z} = 300$ m. The explicit time-integrator was run at the maximum allowable time-step for each method.

Figure 5 compares the density perturbations ρ' of the analytic and numerical solutions. In panel a), contours for ρ' (analytic are solid while NUMA-CG are dashed) are shown after 5 hours of simulation time and compared to the contour integral solution; after 5 hours the models have converged to steady-state. The results of NUMA-DG are similar to the dashed contours in Fig. 5a. Panel b) of this figure compares the numerical results produced by NUMA-CG, NUMA-DG, and Smith's analytical model along the line $x = y = 120$ 000 m. Agreement is very good near the mountain, while the two solutions

begin to deviate as z increases due to the influence of the sponge. Agreement between NUMA and Smith’s results may be brought into closer agreement by increasing the resolution of the model; at least this is the case in the 2D analog of the problem (see [14] and [15]); recall that the “analytic” solution used here is for a Boussinesq model and therefore one should not expect to get exact agreement. Additional verification was performed by comparing the velocity on the surface of the mountain. The down-stream velocity perturbation u' and cross-stream velocity perturbation v' at the ground for both the CG and DG models are compared with the analytical formulas in Eq. (42). Figure 6 displays the results of this comparison after $t = 5$ hours of integration. Agreement between the CG and DG models and the analytical formulas given by Eq. (42) is satisfactory, especially for the cross-stream velocity perturbation. We attribute the slight deviation of the down-stream velocity to the sponge layer in the model, which drives the total down-stream velocity to a constant velocity of 20 m/s at the boundary of the computational domain. This additional forcing term affects the quality of the solution near the sponge layer. In fact, for this particular case, the NRBCs dominate the solution; this exact same behavior was seen in the 2D mountain cases in the NUMA2D codes presented in [14].

5.2.2. 3D Rising Thermal Bubble

Finally, the results of the buoyant convection experiment are shown in Figures 7 and 8. Numerical results using both NUMA-CG and NUMA-DG are shown. A total of 10^3 elements with $N = 8$ polynomials were used, yielding an effective resolution of 12.5m. Due to the coarse resolution of this run, a small amount of artificial diffusion $\nu = 0.5 \text{ m}^2/\text{s}$ was used to suppress grid noise. In the CG code, a time-step of $\Delta t = 0.01$ s was used, whereas the DG code required $\Delta t = 0.005$. A smooth potential temperature perturbation in an initially neutral atmosphere generates vertical updrafts and shear velocities, which cause the bubble to rise, deform, and transition to turbulence. Fig. 7 displays x - z cross-sections of the potential temperature perturbation for $t = 0, 200$, and 400 seconds, while Fig. 8 displays a 1D cross-section along the line $x = y = 500$ m for both CG and DG. Note that the CG result displays significant Gibbs oscillations, while the Gibbs oscillations are less pronounced in the DG code; this behavior is expected, since DG is known to capture sharp gradients more effectively. The rising thermal bubble problem tests the models’ ability to capture the effects of turbulence and turbulent convection. Although no analytical solution exists




Figure 7: Evolution of a 3D rising thermal bubble problem ($x - z$ -slices of the potential temperature perturbation θ') in the $y = 500$ m plane for $t = 0, 200$, and 400 seconds using both NUMA-CG and NUMA-DG. A grid consisting of $N_e = 10^3$ elements with $N = 8$ polynomials was used.




Figure 8: Numerical results for rising thermal bubble problem along the line $x = y = 500$ m using both NUMA-CG and NUMA-DG.

for this problem, the numerical results are physically plausible and resemble previous 3D bubble experiments in [38] and [1]. Similar results are seen for both the NUMA-CG and NUMA-DG codes.

For this test case, the relative conservation of both mass (M) and energy (E) were calculated for both the CG and DG codes. In both codes $M \approx 10^{-13}$ and $E \approx 10^{-7}$. In other words, both CG and DG conserve mass to a level approaching machine precision, whereas energy, which is not a conservation variable in either model, is not conserved. Hence, there is no particular disadvantages to using a non-conservative form of the equations within the CG model.

5.3. Parallel Performance

Optimized versions of both NUMA-CG and NUMA-DG have been deployed on TACC’s Ranger Sun Constellation cluster [2]. It is important to understand that there may exist more optimal parallel implementations and so our results are illustrative of two very specific choices for constructing parallel algorithms (as described in Sec. 4). However, the specific choices we have made in our parallel implementation are either standard or simple in

that they represent the most obvious ways of optimizing both the CG and DG methods.

Two test problems using the rising thermal bubble were executed using a fixed grid comprised of $32^3 = 32768$ elements with both $N = 4$ and $N = 8$ polynomials. The $N = 4$ simulation has an effective resolution of 7.8 km in both the horizontal and vertical, while for $N = 8$ the effective resolution is 3.9 km. The maximum allowable stable time step was used for both the CG and DG models (note that the CG model uses a time-step twice as large as DG and the $N = 4$ simulations allow a time-step twice as large as those for $N = 8$). Due to limitations on the available computational resources, only one run at each processor count was feasible. Also, due to the size of this problem, running the code in serial was not feasible; therefore, the wallclock time for each code at 16 cores was assumed to be 16 times the serial wallclock time.

5.3.1. Results for 4th and 8th Order Polynomials

Figure 9 displays both a) the wallclock time and b) speedup for the CG and DG codes using $N = 4$. In panel b), the ideal speedup (perfect scalability) is also shown. Both codes exhibit nearly perfect linear scaling to 2048 cores. Beyond 2048 cores, however, the scaling of the CG code begins to degrade while the DG code continues to scale ideally up to 8192 cores. At 16384 cores, the DG code scalability plateaus. Let us now see what happens when we increase the polynomial order.

Figure 10 displays the same data for $N = 8$. Both the CG and DG codes exhibit nearly perfect linear scaling to 512 cores. Between 512 to 2048 the CG code loses perfect scaling but recovers beyond 2048 cores. Note that the DG code exhibits perfect linear scaling all the way through to 32768 processors which is the maximum number of processors that a 32^3 element simulation can use (at this point, each processor contains only one element).

5.3.2. Communication Cost

The difference in the scaling behavior between the two codes can be explained by examining the relative communication footprints of both CG and DG. In 3D, each CG processor element (PE) may have up to 26 neighbors, whereas DG has only 6; hence, NUMA-DG's communication footprint is more than 4 times as compact as NUMA-CG. In addition, DG has a more local memory access, since DG stores data on an element-by-element basis.

Hence, only the flux computations require non-local memory fetches; in contrast, CG utilizes a global index, and requires many more non-local memory accesses although this overhead can be eliminated if DG-type local indexing is used.

Effect of Polynomial Order. Comparing Figs. 9 and 10, we also see the dependence of scalability on the order of the method. The number of grid points for a high-order EBG method is $\mathcal{O}(N_e N^3)$, whereas the computation volume is $\mathcal{O}(N_e N^4)$; hence, the computation density is $\mathcal{O}(N)$. In contrast, the communication volume, as described in Sec. 4.4, grows only as $\mathcal{O}(N_e^{2/3} N^2)$. Hence, as we increase the order N , the computation density grows as N whereas the communication density (ratio of communication volume to the number of grid points) decreases as $1/N$. From a computer architecture viewpoint, the *data locality* for both CG and DG increases with order, which is advantageous for both distributed memory implementation as well as shared-memory implementations using graphical processor units (GPUs) [25]. A higher-order method, while more expensive at low processor counts, becomes more efficient as the processor count increases.

For example, at 16 MPI processes, the DG code requires 650 seconds of wallclock time at $N = 4$ and 15000 seconds at $N = 8$ (both using 32^3 elements), yielding a ratio of 23.7 for the wallclock times, i.e.,

$$R_{16} \equiv \left(\frac{WT_{N=8}}{WT_{N=4}} \right)_{NPROC=16} = 23.7$$

where WT is the wallclock time. In contrast, at 8192 MPI processes, the same code requires 1.5 seconds at $N = 4$ and 25 seconds at $N = 8$, for a ratio of $R_{8192} = 16.6$; therefore going from 16 to 8192 has decreased the computational cost ratio between the 4th and 8th order simulations. We attribute this phenomena to the low communication volume and increased data locality (e.g., greater on-processor work per grid point) which is achieved by DG as order N increases. Of course, the same arguments holds for CG although the curves are not as straight as they are for DG.

Note that for a fixed number of elements (in this case 32^3 elements) and increasing the polynomial order from $N = 4$ to $N = 8$ increases the size of the problem by a factor of 32: the number of grid points has increased by a factor of 8 and the number of time-steps by a factor of 4. Although the $N = 8$ simulation is 32 times larger than the $N = 4$ simulation, the

additional cost of increasing the order N decreases as we move to larger core counts. For example, the minimum wallclock time in Fig. 9 for $N = 4$ is 1.5 seconds (corresponding to 8192 processors) while the minimum wallclock time in Fig. 10 for $N = 8$ is 7 seconds (corresponding to 32768 processors). In this case, we see that the cost ratio between 8th and 4th order is now 4.7.

Complexity Analysis: A Geometrical View. One other way of interpreting the results in Figs. 9 and 10 involves quantifying the operation count that takes place on-processor versus the size of the message that has to be communicated across processors. Geometrically, this can be viewed as the ratio between the computational volume (volume of the PE) and communication footprint (surface area of the PE). To simplify the following discussion let us first define the computational volume on a specific hexahedral processor element (PE) as

$$V_{PE} = \mathcal{O}(N_e n_{var} (N + 1)^4)$$

while the surface area of this HPE that defines the communication cost is

$$S_{PE} = \mathcal{O}(N_e^{2/3} n_{var} (N + 1)^2)$$

where N_e denotes the number of elements per PE, n_{var} are the number of variables in our equations ($n_{var} = 5$ in our case), and N denotes the order of the polynomial; these relations come from Eq. (38). Note that the term $N_e^{2/3}$ represents the number of faces of a PE (i.e., the convex hull of the PE). Using these two relations, we can now define the ratio

$$\mathcal{R}_{PE}^N \equiv \frac{V_{PE}}{S_{PE}} = \mathcal{O}(N_e [PE]^{1/3} (N + 1)^2)$$


where $N_e[PE]$ denotes that (for a fixed grid) the number of elements per processor N_e is a function of the number of processors used (PE); this ratio we now use to compare Figs. 9 and 10.

Figure 9 shows that for $N = 4$ both CG and DG lose perfect scalability between 8192 and 16384 processors. Using a grid with 32^3 total elements we can determine that 8192 processors gives $N_e = 4$ and for 16384 processors yields $N_e = 2$. From these values we determine that $\mathcal{R}_{8192}^4 = 39$ and $\mathcal{R}_{16384}^4 = 31$. In contrast for $N = 8$, $\mathcal{R}_{8192}^8 = 127$ and $\mathcal{R}_{16384}^8 = 101$ and scalability is maintained (Fig. 10). This tells us that for the entire range of possible number of processors, for $N = 8$, the ratio of the on-processor


work to inter-processor communication is ≥ 100 which means that the communication is overwhelmed by the computation and we therefore see perfect scalability. The $N = 4$ results tells us that a factor of 40 or less is not sufficient to maintain scalability. Looking again at Fig. 9 we see that for $N = 4$ we maintain perfect linear scalability for both CG and DG at 2048 or fewer processors. For these values we determine that $\mathcal{R}_{2048}^4 = 62$ so we can conjecture that the factor of on-processor work to inter-processor communication can be as low as 62 while still maintaining scalability.

5.3.3. Memory Access

The previous discussion has focused on the effects of load-balancing and communication on scaling. However, memory access must also be considered in order to achieve maximum parallel efficiency. In particular, we wish to fit the local problem into the fast cache on each compute core in order to minimize the number of fetches from *random access memory* (RAM), since each fetch may consume hundreds of clock cycles. To determine the core count at which the local problem fits into cache, consider that the approximate size of each state-vector for the global problem (in bytes) for DG, which uses an element-local storage scheme, is approximately $8n_{var}N_e^g(N+1)^3$, where the coefficient 8 is the size of a real number and N_e^g is the global number of elements. Since we need to store both the state-vector and a RHS, the total size of the global problem for $N = 8$, $N_e^g = 32^3$ and $n_{var} = 5$ is approximately 2 GB. The size of the global problem for CG is slightly smaller due to a global storage scheme, although the difference is small for high polynomial orders. Each quad-core processor on Ranger has 6 MB of shared L3 cache and each core has 512 KB of dedicated L2 cache. Hence the total cache per core on Ranger is 2 MB. Assuming good load balancing, the local problem will fit into cache using 1000 cores. Performing the same calculation for $N = 4$, we see the local problem will fit into cache using approximately 160 cores. Therefore, we expect to see a super-linear speedup at these core counts, which is evident in Figs. 9 and 10 (see panels b) where the CG and DG simulations exceed the ideal scaling curve).




(a) wallclock time




(b) speedup

Figure 9: Scaling study performed using the NUMA-CG and NUMA-DG codes for the RTB problem using 32^3 elements and $N = 4$ polynomials. The wallclock time refers to the total simulation time.



(a) wallclock time



(b) speedup

Figure 10: Scaling study performed using the NUMA-CG and NUMA-DG codes for the RTB problem using 32^3 elements and $N = 8$.

6. Discussion and Conclusion

6.1. Future Work

6.1.1. Microphysical Parameterizations

In conjunction with the Naval Research Laboratory-Monterey, we are incorporating microphysical parameterizations into NUMA. Preliminary experiments using the Kessler scheme [22] have been conducted within a 2D, serial implementation [12]. Since physical parameterizations operate on columns of data independently of adjacent data, the problem is embarrassingly parallel *provided* the domain is decomposed in the horizontal only such that all z values reside on-processor. To facilitate scaling on hybrid shared-distributed memory architectures (such as TACC’s Ranger), hierarchical domain decomposition is desirable, whereby MPI communication based on the algorithm developed in the present paper is utilized in the horizontal and either OpenMP parallelization, appropriate for shared memory, or graphical processor units (GPUs) are employed for fine-grained parallelism in the vertical. For such problems where the domain decomposition is only performed in the xy plane such that all z values are on-processor both the CG and DG methods will always have a much larger computational volume versus communication footprint that will allow both methods to scale perfectly for the maximum number of processors accommodated by a 2D domain decomposition.

6.2. Conclusion

In this paper, we have developed a nonhydrostatic model of the atmosphere (NUMA) based on both CG and DG discretizations in space and explicit discretization in time. We note that purely explicit time-discretization is not feasible due to the stringent CFL restriction; therefore, we will present the results of our implicit-explicit time-integrators in a future work. This paper has subjected NUMA to a couple of limited-area simulations, including orographic flow and buoyant convection problems. The results of these test problems are in agreement with either previous simulations, analytical results, or physical intuition.

The parallel performance study described in Sec. 5 suggests several directions for the future of EBG methods (in particular, DG methods). Firstly, good scalability was exhibited for both the CG and DG codes for $\mathcal{O}(10^4)$ cores, with near-linear strong scaling for the DG code beyond $\mathcal{O}(10^4)$. As described in Secs. 4.4 and 4.5, we attribute these scalability results to a large

computational volume relative to communication volume and high data locality, which yields a low-cache miss percentage. In particular, DG exhibits very low communication volume and outstanding data locality, indicating that DG will be able to achieve ideal scaling up to $\mathcal{O}(10^5)$ cores with no modifications. Secondly, these experiments indicate that high order EBGs (e.g. $N = 8$ relative to $N = 4$) become more efficient at higher core counts. Hence, as we move towards the exascale epoch in which core counts of $\mathcal{O}(10^5)$ and $\mathcal{O}(10^6)$ are rapidly become available, these high order EBGs may become competitive with their established low-order counterparts (finite element and finite volume methods). For these reasons, high-order EBG methods are excellent candidates for next-generation NWP models.

Acknowledgment

The authors acknowledge Shiva Gopalakrishnan for his assistance in analyzing the bottlenecks of the MPI codes as well as running some of the simulations. In addition we thank both Shiva Gopalakrishnan and Michal Kopera for reading over the drafts of the paper. The authors also acknowledge TeraGrid for providing resources on TACC’s Ranger Sun Constellation cluster. We would also like to thank the people who run Ranger for their assistance. This work was funded by ONR Grant PE-0602435N.

References

- [1] N. Ahmad and J. Lindeman. A Godunov-type finite volume scheme for meso- and micro-scale flows in three dimensions. *Pure and Applied Geophysics*, 165(9):1929–1939, 2008.
- [2] Texas Advanced Computing Center. *Ranger user guide*, 2012. <http://www.tacc.utexas.edu/user-services/user-guides/ranger-user-guide>.
- [3] T. Davies, M. J. P. Cullen, A. J. Malcolm, M. H. Mawson, A. Staniforth, A. A. White, and N. Wood. A new dynamical core for the Met Office’s global and regional modelling of the atmosphere. *Q. J. R. Meteorol. Soc.*, 131:1759–1782, 205.
- [4] J. R. Dea, F. X. Giraldo, and B. Neta. High-order non-reflecting boundary conditions for the linearized 2-D Euler equations: No mean flow case. *Wave Motion*, 46:210–220, 2009.

- [5] T. J. DeLuca. Performance of hybrid eulerian-lagrangian semi-implicit time-integrators for nonhydrostatic mesoscale atmospheric modeling. Master’s thesis, Naval Postgraduate School, 2008.
- [6] J. M. Dennis, J. Edwards, K. J. Evans, O. N. Guba, P. H. Lauritzen, A. Mirin, A. St. Cyr, M. Taylor, and P.H. Worley. CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model. *Int. J. of High Perf. Comput. Appl.*, (In Press), 2011.
- [7] J. M. Dennis, M. Levy, R. D. Nair, H. M. Tufo, and T. Voran. Towards and efficient and scalable discontinuous Galerkin atmospheric model. *Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium*, pages 1–7, 2005.
- [8] M. O. Deville, P. F. Fischer, and E. H. Mund. *High-Order Methods for Incompressible Fluid Flow*. Cambridge University Press, 2002.
- [9] M Fey. Multidimensional upwinding. Part I. The method of transport for solving the Euler equations. *J. Comp. Phys.*, 143(1):159–180, JUN 10 1998.
- [10] P. F. Fischer and A. T. Patera. Parallel spectral element solution of the Stokes problem. *J. Comp. Phys.*, 92:380–421, 1991.
- [11] Aimé Fournier, Mark A. Taylor, and Joseph J. Tribbia. The spectral element atmosphere model (SEAM): High-resolution parallel computation and localized resolution of regional dynamics. *Mon. Wea. Rev.*, 132(3):726–748, 2004.
- [12] S. Gabersek, F. X. Giraldo, and J. D. Doyle. Simple microphysics experiments with a spectral element model. *Mon. Wea. Rev. (in review)*, 2010.
- [13] F.X. Giraldo. The Lagrange-Galerkin spectral element method on unstructured quadrilateral grids. *J. Comp. Phys.*, 147(1):114–146, NOV 20 1998.
- [14] F.X. Giraldo and M. Restelli. A study of spectral element and discontinuous Galerkin methods for the Navier–Stokes equations in nonhydrostatic mesoscale atmospheric modeling: Equation sets and test cases. *J. Comp. Phys.*, 227(8):3849–3877, April 2008.

- [15] F.X. Giraldo, M. Restelli, and M. L. J. G. S. Semi-implicit formulations of the Navier-Stokes equations: Applications to non-hydrostatic atmospheric modeling. *SIAM J. Sci. Comp.*, 32:3394–3425, 2010.
- [16] F.X. Giraldo and Thomas E. Rosmond. A scalable spectral element Eulerian atmospheric model (SEE-AM) for NWP: Dynamical core tests. *Mon. Wea. Rev.*, 132(1):133–153, JAN 2004.
- [17] L. Grinberg and G. E. Karniadakis. A new domain decomposition method with overlapping patches for ultrascale simulations: Application to biological flows. *J. Comput. Phys.*, 229:5541–5563, 2010.
- [18] R. L. Higdon. Radiation boundary conditions for dispersive waves. *SIAM J. Numer. Anal.*, 31(1):64–100, 1994.
- [19] RM Hodur. The Naval Research Laboratory’s coupled ocean/atmosphere mesoscale prediction system (COAMPS). *Mon. Wea. Rev.*, 125(7):1414–1430, JUL 1997.
- [20] G. Houzeaux, M. Vázquez, R. Aubry, and J. M. Cela. A massively parallel fractional step solver for incompressible flows. *J. Comp. Phys.*, 228:6316–6332, 2009.
- [21] G. Karypis and V. Kuman. A fast and highly quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comp.*, 20(1):359–392, 1998.
- [22] E. Kessler. *On the Distribution and Continuity of Water Substance in Atmospheric Circulation*. AMS, 1969.
- [23] R. Klein, U. Achatz, D. Bresch, O. M. Knio, and P. K. Smolarkiewicz. Regime of validity of soundproof atmospheric flow models. *J. Atmos. Sci.*, 67(10):3226–3237, 2010.
- [24] J.B. Klemp, W.C. Skamarock, and J. Dudhia. Conservative split-explicit time integration methods for the compressible nonhydrostatic equations. *Mon. Wea. Rev.*, 135(8):2897–2913, 2007.
- [25] A Klockner, T. Warburton, J. Bridge, and J. S. Hesthaven. Nodal discontinuous galerkin methods on graphics processors. *J. Comp. Phys.*, 228(21):7863–7882, 2009.

- [26] Daniel Y. Le Roux. Dispersion relation analysis of the $p^{NC_1} - p^1$ finite-element pair in shallow-water models. *SIAM J. Sci. Comput.*, 27(2):394–414, 2005.
- [27] Joseph M. Lindquist, Beny Neta, and Francis X. Giraldo. A spectral element solution of the Klein-Gordon equation with high-order treatment of time and non-reflecting boundary. *Wave Motion*, 47(5):289 – 298, 2010.
- [28] R. D. Nair, H. W. Choi, and H. M. Tufo. Computational aspects of a scalable high-order discontinuous Galerkin atmospheric dynamical core. *Computers and Fluids*, 30:309–319, 2009.
- [29] M.R. Norman, R. D. Nair, and F. H. M. Semazzi. A low communication and large time step explicit finite-volume solver for non-hydrostatic atmospheric dynamics. *J. Comp. Phys.*, 230:1567–1584, 2011.
- [30] J.M. Prusa, P.K. Smolarkiewicz, and A.A. Wyszogrodzki. EULAG, a computational model for multiscale flows. *Comput. Fluids*, 37:1193–1207, 2008.
- [31] M. Restelli and F. X. Giraldo. A conservative discontinuous galerkin semi-implicit formulation for the navier-stokes equations in nonhydrostatic mesoscale modeling. *SIAM J. Sci. Comp.*, 31:2231–2257, 2009.
- [32] U. Schattler, G. Doms, and J. Steppeler. Requirements and problems in parallel model development at DWD. *Scientific Programming*, 8(1):13–22, 2000 2000.
- [33] W. C. Skamarock, J. D. Doyle, P. Clark, and N. Wood. A standard test set for nonhydrostatic dynamical cores of NWP models . *AMS NWP-WAF Conference Poster*, page P2.17, 2004.
- [34] R. B. Smith. Linear theory of stratified hydrostatic flow past an isolated mountain. *Tellus*, 32:348–364, 1980.
- [35] R. B. Smith. Linear theory of stratified flow past an isolated mountain in isotheric coordinates. *J. Atmos. Sci.*, 45(42):3889–3896, 1988.
- [36] P. K. Smolarkiewicz and J. A. Pudykiewicz. A class of semi-Lagrangian approximations for fluids. *J. Atmos. Sci.*, 49(22):2082–2096, 1992.

- [37] Raymond J. Spiteri and Steven J. Ruuth. A new class of optimal high-order strong-stability-preserving time discretization methods. *SIAM J. Numer. Anal.*, 40(2):469–491, 2002.
- [38] S. J. Thomas, J. P. Hacker, and P. K. Smolarkiewicz. Spectral preconditioners for nonhydrostatic atmospheric models. *Mon. Wea. Rev.*, 131:2464–2478, 2003.
- [39] J. Thuburn, Todd D. Ringler, William C. Skamarock, and Joseph B. Klemp. Numerical representation of geostrophic modes on arbitrarily structured c-grids. *J. Comp. Phys.*, 228:8321–8335, 2009.
- [40] Paul A. Ullrich, C. Jablonowski, and B. van Leer. High-order finite-volume methods for the shallow-water equations on the sphere. *J. Comp. Phys.*, 229(17):6104–6134, 2010.
- [41] L. C. Wilcox, G. Stadler, C. Burstedde, and O. Ghattas. A high-order discontinuous Galerkin method for wave propagation through coupled elastic-acoustic media. *J. Comp. Phys.*, 229(24):9373–9396, 2010.