A Fully Implicit Jacobian-Free High-Order Discontinuous Galerkin Mesoscale Flow Solver

Amik St-Cyr[1,*] and David Neckels[2]

[1] National Center for Atmospheric Research (NCAR), Boulder, USA
amik@ucar.edu
[2] Previously NCAR now Beckman Coulter Inc., Fullerton, California, USA

Abstract. In this work it is shown how to discretize the compressible Euler equations around a vertically stratified base state using the discontinuous Galerkin approach on collocated Gauss type grids. A stiffly stable Rosenbrock W-method is combined with an approximate evaluation of the Jacobian to integrate in time the resulting system of ODEs. Simulations with fully compressible equations for a rising thermal bubble are performed. Also included are simulations of an inertia gravity wave in a periodic channel. The proposed time-stepping method accelerates the simulation times with respect to explicit Runge-Kutta time stepping procedures having the same number of stages.

1   Introduction

With modern climate models currently able to reach nonhydrostatic resolutions, it is generally perceived that the next generation of general circulation models will be able to run globally for both weather and climate prediction. The Reynolds numbers involved in global and mesoscale atmospheric modeling explains the choice of numerical weather prediction (NWP) centers for the compressible Euler equations. The eventual availability of petascale calculators forces the research stream into highly-scalable numerical methods. A very popular approach nowadays for applications is the so called discontinuous Galerkin method which enjoys some of the properties of finite-volumes and finite-elements methods [1]. It benefits from a variable polynomial degree within each element, it is highly scalable since the cost of parallel communications can be almost completely hidden [2]. For global atmospheric computations such comparable methods are performing very well [3]. However, at increasingly finer resolutions, time-stepping the compressible Euler equations is problematic. Most models utilize an explicit time-stepping procedure and are thus strongly dependent on the Courant-Friedrich-Lewy (CFL) condition which dictates the maximum allowable time-step to be proportional to the minimal spatial mesh size. The consequences of such techniques are such that, in order to solve a global climate problem at a

---

$1km$ resolution at scientifically relevant integration rates, a 2^{20} folds increase in computing power will be required. Following Moore's Law such a supercomputer will be available in 30 years. Therefore developments in new algorithms are of utmost importance.

In the present work a fully implicit framework for solving the compressible Euler equations at low Mach number is proposed. First, the governing equations are presented with their non-dimensionalisation. The space discretization then follows. Next, the time-discretization is discussed and the Rosenbrock W-method is introduced with the low Mach treatment of the numerical flux. Finally, various standard numerical experiments proposed to test mesoscale flow solvers are presented.

2   Governing Equations

The model system of equations that will be solved in this work is a system of m-conservation laws with source term, in two spatial dimensions, written generically as

$$\mathrm{U}_t + \nabla \cdot \mathbf{F}(\mathrm{U}) = S(\mathrm{U}). \tag{1}$$

where $\mathbf{F}(\mathrm{U}) \equiv (F, G) : \mathbb{R}^m \to \mathbb{R}^m \times \mathbb{R}^m$ and $\mathrm{U} = \mathrm{U}(\mathbf{x}, t) : \mathbb{R}^2 \times \mathbb{R}^+ \to \mathbb{R}^m$. For convenience we define the hyperbolicity matrix as

$$A(\mathrm{U}; \hat{\mathbf{n}}) \equiv \sum_{i=1}^{2} A_i(\mathrm{U}) \hat{n}_i \tag{2}$$

with $\mathbf{A} = (A_1, A_2)$ where $A_1 = F'(\mathrm{U})$ and $A_2 = G'(\mathrm{U})$. The above generic equation permits the study of various sets of equations. In [4], the conservative form of the compressible Euler equations involves the conservation of potential temperature instead of the total energy as traditionally favored in computational fluid dynamics. The state vector and the fluxes associated with this formulation are

$$\mathrm{U} \equiv \begin{pmatrix} \rho \\ \rho u \\ \rho w \\ \rho \theta \end{pmatrix} \equiv \begin{pmatrix} \rho \\ U \\ W \\ \Theta \end{pmatrix}, \quad F \equiv \begin{pmatrix} U \\ \frac{UU}{\rho} + p \\ \frac{UW}{\rho} \\ \frac{U\Theta}{\rho} \end{pmatrix}, \quad G \equiv \begin{pmatrix} W \\ \frac{WU}{\rho} \\ \frac{WW}{\rho} + p \\ \frac{W\Theta}{\rho} \end{pmatrix}$$

and $S = (0, 0, -\rho g, 0)^T$. In the above system ρ represents density, u and w are the horizontal and vertical velocities and θ is the potential temperature. To close the system, an expression for the pressure is required:

$$p = p_{\mathrm{ref}} \left(\frac{R\Theta}{p_{\mathrm{ref}}} \right)^{\gamma} \tag{3}$$

with $\gamma = c_p / c_v = 1004.6/717.6$ and where $p_{\mathrm{ref}} = 1013.25$ kpa. For certain tests, the Eady model is required and consists into adding the $\hat{y}$ momentum equation with an f plane Coriolis force. The entire approach is still 2D and all derivatives in the $\hat{y}$ direction are zero.

The system of conservation laws is re-written around a hydrostatically balanced base state. This step is necessary in order to later avoid any kind of roundoff issues. The hydrostatic assumption reads

$$\frac{\partial \bar{p}}{\partial z} = -\bar{\rho}g.$$

Thus the following assumption is made on the dependent variables:

$$p = \bar{p}(z) + p' \tag{4}$$
$$\rho = \bar{\rho} + \rho' \tag{5}$$
$$U = \rho'u + \bar{\rho}u = U' + U \tag{6}$$
$$W = \rho'w + \bar{\rho}w = W' + W \tag{7}$$
$$\Theta = \bar{\rho}(z)\bar{\theta}(z) + \Theta'. \tag{8}$$

To minimize numerical cancellation errors associated with the original set of equations, the equation are written in their non-dimensional form. Thus a length scale x_0, a time scale t_0 along with a reference velocities u_0, pressure γp_0, ρ_0 and θ_0 are introduced into the equations using the simple replacement $\rho \rightarrow \rho\rho_0$. This leads to the apparition of the Strouhal number $\frac{x_0}{t_0 u_0}$ in front of all time derivatives and to the inverse Mach number in front of the pressure gradient. The buoyancy term has a factor $\frac{gx_0}{u_0^2}$ which can be made equal to one if $u_0 = \sqrt{gx_0}$. If a reference length x_0 is chosen, the time scale is determined by setting the Strouhal number to one. Therefore, the final form of the conserved variables and fluxes used in the discretization are

$$\mathbf{U} \equiv \begin{pmatrix} \rho' \\ (\bar{\rho} + \rho')u \\ (\bar{\rho} + \rho')w \\ (\rho\theta)' \end{pmatrix} \equiv \begin{pmatrix} \rho' \\ U \\ W \\ \Theta' \end{pmatrix}, \quad F \equiv \begin{pmatrix} U \\ \frac{UU}{(\bar{\rho}+\rho')} + \frac{1}{M^2}p' \\ \frac{UW}{(\bar{\rho}+\rho')} \\ \frac{U\Theta}{(\bar{\rho}+\rho')} \end{pmatrix}, \quad G \equiv \begin{pmatrix} W \\ \frac{WU}{(\bar{\rho}+\rho')} \\ \frac{WW}{(\bar{\rho}+\rho')} + \frac{1}{M^2}p' \\ \frac{W\Theta}{(\bar{\rho}+\rho')} \end{pmatrix}$$

The boundary conditions employed in the numerical experiments are either of the periodic or slip type. This condition is expressed simply as $\mathbf{u} \cdot \hat{\mathbf{n}} = 0$ and will be imposed weakly in the spatial discretization.

3 Spatial Discretization

The discontinuous Galerkin method can be formulated in two different ways. The first approach is one popularized by Cockburn and various authors, see [5], and is known as the *weak formulation*. The second one, the *strong formulation*, was first considered in [6] and employed herein. Also, instead of considering modal orthogonal basis functions we consider only nodal ones. Thus, in a nodal discontinuous Galerkin discretization, the computational domain Ω is partitioned into J quadrilateral elements Ω_j in which the dependent and independent variables are approximated by N–th order tensor-product polynomial expansions.

A function $U : \mathbb{R}^2 \times \mathbb{R}^+ \to \mathbb{R}$ is spatially expanded, on element Ω_j, in terms of the N–th degree Lagrangian interpolants h_i as defined in [7],

$$U_h(\mathbf{x}(r_1^j, r_2^j), t)|_{\Omega_j} \equiv \sum_{k=0}^{N} \sum_{l=0}^{N} (U_{jkl})(t)\, h_k(r_1^j(\mathbf{x}))\, h_l(r_2^j(\mathbf{x})), \tag{9}$$

where $\mathbf{x} \to (r_1^j(\mathbf{x}), r_2^j(\mathbf{x}))$ is an affine transformation, T_j, from the quadrilateral element Ω_j, on the original domain, to the reference element $[-1, 1] \times [-1, 1]$ and the $(\mathbf{u}_j)_{kl}$ are the nodal basis coefficients defined at the Gauss-Lobatto Legendre (GLL) quadrature points $\{\hat{\xi}_i\}_{i=0}^{N}$. Introducing the finite dimensional space $W_h = \{v_h \in L^2(\Omega) \mid \forall \Omega_j \in \Omega, \mathbf{v}_h \circ T_j \in \mathbb{P}_k\}$, the discrete weak formulation is obtained by integration by parts of the differential equation (1), using the representation (9), on an element Ω_j:

$$\int_{\Omega_j} \mathbf{v}_h \frac{dU_h}{dt} dA + \int_{\Omega_j} \mathbf{v}_h \nabla \cdot \mathbf{F}(U_h) dA =$$

$$M \int_{\Omega_j} \mathbf{v}_h U_h dA + \int_{\partial\Omega_j \backslash \partial\Omega} \mathbf{v}_h (\mathbf{F}(U_h) - \mathbf{F}_{num}(U_h)) \cdot \hat{\mathbf{n}} d\sigma$$

$$+ \int_{\bar{\Omega}_j \cap \partial\Omega} \mathbf{v}_h (\mathbf{F}(U_h) - \mathbf{F}_{bnd}(U_h)) \cdot \hat{\mathbf{n}} d\sigma \tag{10}$$

where $\mathbf{v}_h$ and U_h are in $[W_h]^4$. The flux is not uniquely defined at the boundary $\partial\Omega_j$ and a numerical flux was introduced in the above expression. The Russanov flux, or local Lax-Friedrich, is

$$\mathbf{F}_{num}(U_h) \cdot \hat{\mathbf{n}} = \frac{1}{2}(\mathbf{F}(U_h^R) + \mathbf{F}(U_h^L)) \cdot \hat{\mathbf{n}} - \frac{1}{2}\mathcal{D}(U_h^R, U_h^L; \hat{\mathbf{n}})(U_h^R - U_h^L) \tag{11}$$

at the interface $\partial\Omega_j$ with U_h^R defined as the discrete counterpart of $U^+(\mathbf{x}) \equiv \lim_{\epsilon \to 0^+} U((1-\epsilon)\mathbf{x} + \epsilon\hat{\mathbf{n}})$ where $\mathbf{x} \in \partial\Omega_j$ and $\hat{\mathbf{n}}$ points in the outward direction of element Ω_j. Notice that with $\epsilon \to 0^+$, U_h^L is obtained. The matrix $\mathcal{D}$, defined in more details in a forthcoming section, is a function of $(U_h^{R,L}$ and of the interface points $\mathbf{x} \in \partial\Omega_j$. Notice that the numerical flux is the only function connecting the element Ω_j to its neighboring elements. The integrals are directly evaluated using Gauss-Lobatto quadratures. The latter, for surfaces, are obtained as follows:

$$(f, g)_j \equiv \sum_{k,l=0}^{N} f(\mathbf{x}^j(\xi_k, \xi_l))|_{\Omega_k} \cdot g(\mathbf{x}^j(\xi_k, \xi_l))|_{\Omega_j} |\mathcal{J}^j(\xi_k, \xi_l)| \rho_k \rho_l, \tag{12}$$

where $|\mathcal{J}^j(\xi_k, \xi_l)|$ is the Jacobian of the transformation $\mathbf{x}^j(\mathbf{r})$ and $\{\rho_i\}_{i=0}^{N}$ are the Gauss-Legendre weights associated with the quadrature points $\{\xi_i\}_{i=0}^{N}$. The boundary integrals are also performed using the GLL points. Once the spatial discretization is performed, the semi-discrete problem in time is

$$\frac{dU_h}{dt} = R_h(U_h) \tag{13}$$

and a method of lines can be applied to the above system of ordinary differential equations.

4   Time Discretization

A general Rosenbrock method for solving the ODE (13) with suitable initial conditions, with an error estimator $\hat{U}_h^{n+1}$, can be written as

$$r_i = \sum_{j=1}^{i} \gamma_{ij} k_j, \quad k_i = \frac{1}{\gamma_{ii}} r_i - \sum_{j=1}^{i-1} c_{ij} r_j$$

$$(\frac{1}{\Delta t \gamma_{ii}} - J_h) r_i = R_h(U_h^n + \sum_{j=1}^{i-1} a_{ij} r_j) + \sum_{j=1}^{i-1} (\frac{c_{ij}}{\Delta t}) r_j \tag{14}$$

$$U_h^{n+1} = U_h^n + \sum_{j=1}^{s} m_j r_j, \quad \hat{U}_h^{n+1} = U_h^n + \sum_{j=1}^{s} \hat{m}_j r_j$$

where $J_h \equiv \left. \frac{\partial R_h}{\partial u} \right|_{u^n}$ is simply the Jacobian. When the Jacobian J_h is exact, the above Rosenbrock method has the same stability domain as the equivalent diagonally implicit Runge-Kutta (DIRK) scheme [8]. They can be L-stable and stiffly accurate provided that the number of stages exceeds the order of the method. For an explicit Runge-Kutta, having the same number of stages s, a cost function for the average number of Krylov iterations, per stage, necessary to solve the linear systems for a Rosenbrock method can be established as $\overline{iter} \leq \frac{\Delta t_i}{\Delta t_e \, accel}$ were $\Delta t_i / \Delta t_e$ is the ratio between the implicit time-step size and the CFL constrained explicit time-step and *accel* is the desired acceleration factor with respect to explicit. The latter clearly establishes that outperforming an explicit scheme is very difficult unless very large $\Delta t_i / \Delta t_e$ ratios are employed. It should be noted that in the case of a Jacobian free Newton-Krylov approach, for DIRK, the upper bound would still be valid but divided by the number of Newton iterations. Notice that the above cost formula is idealized in the sense that the costs for setting up or applying a non-trivial preconditioning technique were not included. Because the linear system at each stage can *at most* be solved approximatively, due to the use of iterative procedures, traditional Rosenbrock methods cannot be used. In turn, methods able to support arbitrary matrices instead of the true Jacobian are available. These methods are called Rosenbrock W-methods and were first investigated in [9]. Rang and Angermann have derived L-stable Rosenbrock methods of order 3 with 4 stages and this is the method used here [10]. The coefficients of the method in table 1 are for use with the Rosenbrock method in the form of (14) which avoids multiplication by the Jacobian in the rhs. The coefficients were computed using arbitrary precision arithmetics and truncated at the 15^{th} digit. Finally, the linear problem is inverted using a restarted generalized conjugate residual (GCR) Krylov accelerator. The latter is equivalent to GMRES.

4.1   Construction of the Jacobian for Euler

In order to have a practical method, it is necessary to have an efficient evaluation of the Jacobian matrix for the compressible Euler equations. One approach is the

Table 1. L-stable Rosenbrock W-method coefficients with $\gamma = 0.4358665215084590$

$a_{21} = 2.0000000000000000$	$c_{21} = -4.588560720558083$
$a_{31} = 1.4192173174557647$	$c_{31} = -4.184760482319161$
$a_{32} = -0.2592322116729697$	$c_{32} = 0.285192017355496$
$a_{41} = 4.1847604823191607$	$c_{41} = -6.368179200128358$
$a_{42} = -0.2851920173554959$	$c_{42} = -6.795620944466836$
$a_{43} = 2.2942803602790417$	$c_{43} = 2.870098604331056$
$m_1 = 0.242123807060954$	$\hat{m}_1 = 3.907010534671192$
$m_2 = -1.223250583904515$	$\hat{m}_2 = 1.118047877820503$
$m_3 = 1.545260255335102$	$\hat{m}_3 = 0.521650232611491$
$m_4 = 0.435866521508459$	$\hat{m}_4 = 0.500000000000000$

Jacobian free technique where the action of the Jacobian on a vector is *approximated* using a Gâteau derivative. Since a Rosenbrock-W method is employed a slight error in the Jacobian has no effects in the consistency of the method but could adversely affect the stability. The multiplication of the vector $\underline{V}_h$ by the Jacobian, frozen at $\underline{U}_h^n$, of the discretized right hand side of the Euler equations is defined as

$$J\underline{V}_h = \frac{\partial R_h}{\partial u}\Big|_{\underline{U}_h^n}\underline{V}_h = \frac{R_h(\underline{U}_h^n + \epsilon\underline{V}_h^n) - R_h(\underline{U}_h^n)}{\epsilon} + O(\epsilon). \tag{15}$$

5 Low Mach Number Preconditioning

Current compressible flow solvers are not suitable to simulate, without any modifications, flow fields transitioning from incompressible to compressible regimes. The literature on the subject demonstrates that the extension of a compressible solver to the incompressible regime leads to asymptotically wrong solutions at the low Mach limit. By adding time derivatives in an incompressible system it is possible to obtain a clustering of the hyperbolic eigenvalues such that they are all of the same magnitude when compared to one another: see [11] and references therein. However these techniques were developed in order to march the equations towards a steady state. Viozat and collaborators [12,13] have shown how to taylor such an approach to unsteady low Mach number simulations in the finite-volumes case. In what follows we adapt the Russanov flux for the equation proposed in [4] for the low Mach regime. The first step consist into rewriting in entropy variables $W = (p, u, v, s)^T$ the system using the following *passage* matrices:

$$\frac{\partial W}{\partial U} = \begin{pmatrix} c^2 & 0 & 0 & 0 \\ -u/\rho & 1/\rho & 0 & 0 \\ -w/\rho & 0 & 1/\rho & 0 \\ \gamma/\rho & 0 & 0 & \frac{\gamma\,R\,p_{\mathrm{ref}}^{\frac{1-\gamma}{\gamma}}}{p^{1/\gamma}} \end{pmatrix} \quad \text{and} \quad \frac{\partial U}{\partial W} = \begin{pmatrix} 1/c^2 & 0 & 0 & 0 \\ u/c^2 & \rho & 0 & 0 \\ w/c^2 & 0 & \rho & 0 \\ \frac{1}{\gamma R}\frac{p}{p_{\mathrm{ref}}}^{\frac{1-\gamma}{\gamma}} & 0 & 0 & \frac{p_{\mathrm{ref}}}{\gamma R}\left(\frac{p}{p_{\mathrm{ref}}}\right)^{1/\gamma} \end{pmatrix}$$

where $U = (\rho, \rho u, \rho w, \rho \theta)^T$ and $s \equiv \log p/\rho^\gamma$. The preconditioning employed is the same as in [13] and consist, in the entropy variables, in a diagonal matrix denoted by $P^{-1}(W) = \text{diag}\{\beta^2, 1, 1, 1\}$ with $\beta = O(M)$.

$$P(W)W_t + P(W)\mathbf{G}(W) \cdot \nabla W = P(W)\tilde{M}W \tag{16}$$

where the quasi-linear form of the compressible Euler equations was employed. Transforming back to the conserved variables yields the preconditioning matrix

$$P(U) \equiv \frac{\partial U}{\partial W}P(W)\frac{\partial W}{\partial U}$$

and the corresponding quasi-linear system of conservation laws

$$P(U)U_t + P(U)\mathbf{A}(U) \cdot \nabla U = P(U) \; MU. \tag{17}$$

The Russanov flux is modified next only in the dissipative term in analogy to what was proposed by [11]. More precisely the maximum eigenvalues are computed with respect to the hyperbolicity matrix $P(U)\mathbf{A}(U) \cdot \hat{n}$ and multiplied by $P^{-1}(U)$:

$$\mathbf{F}(U^R, U^L) \cdot \hat{n} = \frac{1}{2}(\mathbf{F}(U^R)+\mathbf{F}(U^L)) \cdot \hat{n} - \frac{1}{2}P^{-1}(\frac{U^R + U^L}{2})D(U^R, U^L; \hat{n}) \tag{18}$$

with

$$D(U^R, U^L; \hat{n}) = \text{diag}(\max\{|\text{eigen}\{P(U)\mathbf{A}(U) \cdot \hat{n}\}|\})(U^R - U^L).$$

Thus, as remarked in [14] the only modification required to precondition the system are performed in the dissipation matrix. It will later be seen that these modifications lead to lower iteration counts in the iterative solution process when the considered domains are non-hydrostatic.

6   Numerical Experiments

6.1   Rising Smooth Bubble Experiment

We carry an integration using as initial condition a continuous perturbation of the potential temperature in an initially hydrostatically balanced atmosphere. The complete details of the initial condition can be found in [15]. The computational domain consist in a rectangle of dimension $20km$ in the horizontal by $10km$ in the vertical with slip boundary conditions on the top and bottom of the atmosphere and periodic conditions on both ends. The initial perturbation is integrated for a 1000 seconds. The same test is performed with two different values of β. In table 2 the solver is used with the low Mach preconditioning (LM) set with $\beta = 0.4$ while the modification is turned off by using $\beta = 1.0$. The acceleration nearly doubles at large Courant numbers.

Table 2. Effects of low Mach (LM) number preconditioning for the raising bubble experiment. Relative residual tolerance set to Solver 1×10^{-6}, with 7^{th} degree polynomials per element and 16 elements in the horizontal direction and 8 in the vertical.

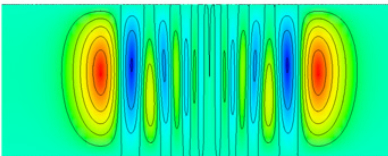| Time step	Iterations with LM	acceleration	Iterations **without** LM	acceleration
$1.0s$	30	3.2	33	2.8
$2.0s$	36	5.1	45	4.1
$10.0s$	69	13.5	103	9.1
$50.0s$	207	22.7	493	10.2

6.2   Inertia Gravity Waves in a Periodic Channel

An initial perturbation in a periodic channel triggers internal gravity waves that are transported by a background mean flow of $20m/s$ to the right of the domain. The original full definition of the initial condition can be found in [16]. The domain is again a rectangle of dimensions $[300, 10]km$. The translating perturbation is integrated for 3000 seconds. The domain is covered using $[90, 3]$ elements of polynomial of degree 8 leading to a relative mesh spacing of $\Delta x = \Delta z = 500m$ for all simulations. In Fig. 1 left panel, a time step of 12 seconds is employed which is twice the time step used in [16] for a comparable resolution. For a large unphysical time-step, the solution is significantly degraded but stable: Fig. 1 right panel. For time-steps up to 50 seconds which is twice the advective scale $(500/20 = 25)$ the approach produces reasonable solutions.

6.3   Eady Model

The same test as in the previous section is modified by considering an *Eady model*. The forcing term f is set to 0.0001 and the periodic channel is transformed to a very thin shell of dimensions $[6000, 10]km$. The latter is tiled using 600 elements of dimension $1km^2$ each of 7^{th} degree. The model is integrated for 60000 seconds with a resolution of approximatively $1.5km$ in both spatial directions. In this configuration $\beta = 1.0$ lead to faster compute times. The time steps employed in this simulations were ranging from $500s$ to $2000s$ at which an acceleration of

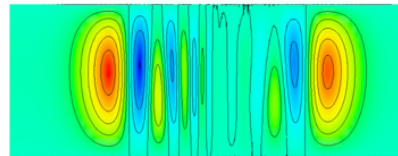a: b:

Fig. 1. Inertia gravity wave in a periodic channel using the equations in potential temperature density form. The various time-steps (in alphabetical order) for the third order Rosenbrock W-method are: 12 secs and 100 secs.
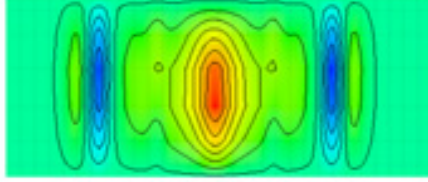
Fig. 2. Inertia gravity wave in a periodic channel in hydrostatic configuration using the Eady model. The various time for the third order Rosenbrock W-method is 500 secs.

45 with respect to explicit was observed. The extreme gains are explained by the exact vertical preconditioning in this case: a consequence of using one element in the vertical direction.

7   Conclusions and Future Directions

In this work it was shown how to discretize the stratified compressible Euler equations for low Mach, nearly incompressible atmospheric simulations consisting in slight perturbation of a hydrostatically balanced state. A low Mach preconditioning of the numerical flux used for yielding asymptotically valid results leads to lower iteration counts for the Krylov accelerator. All simulations result in lower simulation times for simple block Jacobi preconditioning[1]. However, the flux preconditioning needs to take into account the aspect ratio of the domain in order to differentiate between a nonhydrostatic and hydrostatic regimes. Moreover, the numerical results presented show that the nonlinear cycle present in Jacobian-Free Newton Krylov method can be *avoided* and provides a robust framework for unsteady computations. For very thin domains, where the hydrostatic assumption could be employed, the solver shows very good accelerations. If these results are extrapolated to global climate or weather modeling, the corresponding time-steps sizes at 10km resolution would be close to 30 minutes: the time scale employed by most physical parameterizations. Future work will concern the coupling to idealized physics, inclusion of moisture transport, adaptive hybrid non-conforming grids, extension to three spatial dimensions, higher-order Rosenbrock methods and, most importantly, improved preconditioning.

References

1. Cockburn, B., Karniadakis, G.E., Shu, C.W.: Discontinuous Galerkin Methods: Theory, Computation, and Applications. Lecture Notes in Computational Science and Engineering, vol. 11. Springer, New York (2000)

---

[1] The quality of the non-linear solutions can be readily compared with more traditional approaches at `http://www.mmm.ucar.edu/projects/srnwp_tests/`

2. St-Cyr, A., Thomas, S.J.: Parallel atmospheric modeling with high-order continuous and discontinuous galerkin methods. In: Deane, A., Periaux, J., Ecer, A., Satofuka, N., McDonough, J. (eds.) Parallel Computational Fluid Dynamics 2005: Theory and Applications: Proceedings of the Parallel CFD 2005 Conference, pp. 485–492. Elsevier Science, Amsterdam (2006)
3. Bhanot, G., Dennis, J.M., Edwards, J., Grabowski, W., Gupta, M., Jordan, K., Loft, R.D., Sexton, J., St-Cyr, A., Thomas, S.J., Tufo, H.M., Voran, T., Walkup, R., Wyszogrodzki, A.A.: Early experiences with the 360tf ibm bluegene/l platform. Int. J. Comput. Meth. 5(2), 237–253 (2008)
4. Skamarock, W.C., Klemp, J.B., Dudhia, J., Gill, D.O., Barker, D.M., Wang, W., Powers, J.G.: A description of the advanced research WRF version 2. NCAR Tech. Note TN-468+STR, National Center for Atmospheric Research, 1850 Table Mesa Drive, Boulder, Colorado, 80305, USA (2005) (revised January 2007)
5. Bassi, F., Rebay, S.: High-order accurate discontinuous finite element solution of the 2d euler equations. Journal of Computational Physics 138(2), 251–285 (1997)
6. Giraldo, F.X., Hesthaven, J.S., Warburton, T.: Nodal high-order discontinuous galerkin methods for the spherical shallow water equations. Journal of Computational Physics 181(2), 499–525 (2002)
7. Ronquist, E.M.: Optimal Spectral Element Methods for the Unsteady Three-Dimensional Incom- pressible Navier-Stokes Equations. Ph.D thesis. MIT, Cambridge, MA, USA (1988)
8. Hairer, E., Wanner, G.: Solving Ordinary Differential Equations II, 2nd edn. Springer Series in Computational Mathematics, vol. 14. Springer, Heidelberg (1996)
9. Steihaug, T., Wolfbrandt, A.: An attempt to avoid exact jacobian and nonlinear equations in the numerical solution of stiff differential equations. Math. Comp. 33, 521–534 (1979)
10. Rang, J., Angermann, L.: New Rosenbrock W-methods of order 3 for partial differential algebraic equations of index 1. BIT Numerical Mathematics 45, 761–787 (2005)
11. Meister, A.: Asymptotic based preconditioning technique for low Mach number flows. ZAMM 83(1), 3–25 (2003)
12. Viozat, C.: Calcul d'écoulements stationnaires et instationnaires à petit nombre de Mach, et en maillages étirés. Ph.D thesis, Université de Nice-Sophia Antipolis, Nice, France (October 1998)
13. Guillard, H., Viozat, C.: On the behaviour of upwind schemes in the low mach number limit. Comput. Fluids 28, 63–86 (1999)
14. Mavriplis, D.J.: Multigrid strategies for viscous flow solvers on anisotropic unstructured meshes. Journal of Computational Physics 145, 141–165 (1998)
15. Wicker, L.J., Skamarock, W.C.: Time splitting methods for elastic models using forward time schemes. Mon. Wea. Rev. 130, 2088–2097 (2002)
16. Skamarock, W.C., Klemp, J.B.: Efficiency and accuracy of the Klemp-Wilhelmson time-splitting technique. Mon. Wea. Rev. 122, 2623–2630 (1994)