51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition
07 - 10 January 2013, Grapevine (Dallas/Ft. Worth Region), Texas

AIAA 2013-0383

Time-Accurate Flow Simulations Using an Efficient Newton-Krylov-Schur Approach with High-Order Temporal and Spatial Discretization

Pieter D. Boom,* and David. W. Zingg †

Institute for Aerospace Studies, University of Toronto, Toronto, Ontario, M3H 5T6, Canada

In order to demonstrate the potential advantages of high-order spatial and temporal discretizations, implicit large-eddy simulations of the Taylor-Green vortex flow and transitional flow over an SD7003 wing are computed using a variable-order finite-difference code on multi-block structured meshes. The spatial operators satisfy the summation-by-parts property, with block interface coupling and boundary conditions enforced through simultaneous approximation terms. The solution is integrated in time with explicit-first-stage, singly-diagonally-implicit Runge-Kutta methods. Simulations of the Taylor-Green vortex show the clear advantage of high-order spatial discretizations in terms of accuracy and efficiency. The higher-order methods are better able to delay excessive dissipation on coarser grids and are better able to capture the details of the flow on finer grids. Similar dissipation and enstrophy profiles are obtained with a second-order spatial discretization, and a fourth-order spatial discretization with half the number of grid points in each direction, half the number of time steps, and approximately 85% less CPU time. Temporal convergence studies demonstrate the relatively high efficiency of the fourth-order explicit-first-stage, singly-diagonally-implicit Runge-Kutta method, except for simulations requiring only a minimum level of accuracy. Results of the simulation of transitional flow over the SD7003 wing show good agreement with experiment and other computations, despite a relatively coarse grid. The use of high-order discretizations is shown to be essential in obtaining this accuracy efficiently. These results give a clear picture of the benefits of high-order discretizations, along with the advantages of the novel parallel Newton-Krylov-Schur algorithm presented, for high-accuracy unsteady flow simulation.

I. Introduction

The simulation of large-scale, complex unsteady flows is becoming more prevalent due to advances in computer architecture, parallel computing, and numerical methods. However, these simulations remain very expensive, in terms of both computational resources and time. High-order methods present one means of reducing the cost of these simulations. Despite being more computationally expensive per node or per time step, significantly coarser simulations can be used to obtain the same level of accuracy. As the required accuracy becomes more stringent, the reduction in mesh and time step requirements outweighs the increased cost of the methods, thus providing greater efficiency.

Summation-by-parts (SBP) spatial operators are a robust and efficient way of extending finite-difference schemes to higher-order. The SBP property imposes specific constraints on block boundaries to ensure time stability using the energy method. As a result, each block is independently time stable, provided appropriate boundary values are available. An efficient way of providing these values is through simultaneous approximation terms (SATs). Only a single halo node is required for the SBP-SAT approach, regardless of order, keeping inter-block communication to a minimum. Mesh generation is also simplified by the minimal requirement of C^0 continuity at block boundaries. This requirement can be further relaxed, provided a suitable interpolant can be found on each block face.[34] The merits of the SBP-SAT approach have been

*PhD Candidate, AIAA Student Member.

†Professor and Director, Tier 1 Canada Research Chair in Computational Aerodynamics, J. Armand Bombardier Foundation Chair in Aerospace Flight, Associate Fellow AIAA.

demonstrated with a variety of steady[19,40] and unsteady[25,26,40,43,44] simulations, as well as having being applied in the context of aerodynamic shape optimization.[20,21,31]

Implicit Runge-Kutta methods have excellent high-order stability characteristics which are essential for simulating stiff differential equations with stringent accuracy requirements. Particularly, Explicit-first-stage Singly-Diagonally-Implicit Runge-Kutta (ESDIRK) methods have been shown to perform very well.[3,7,24,40,48] Stiff-accuracy and stage-order two, resulting from the explicit first stage, help minimize the effects of order reduction. A balance between accuracy, stability and computational cost are obtained by the semi-implicit form and the constant diagonal coefficient. All the methods of this class are unconditionally stable (A-stable) and furthermore have perfect damping at infinity (L-stability). These characteristics make methods of this class a robust and efficient choice for complex unsteady flow simulations.

The use of high-order spatial and temporal discretizations is of particular interest in high-resolution simulations such as direct numerical simulation (DNS), and large-eddy simulation (LES). The objective is to characterize the performance of the methods used, namely high-order SBP-SAT finite-difference schemes, high-order semi-implicit Runge Kutta schemes, and a Newton-Krylov-Schur algorithm, in the computation of complex unsteady flows. Prediction of transition and turbulence in the Taylor-Green vortex flow and low-Reynolds number flow over the SD7003 wing are chosen for this investigation.

The Taylor-Green vortex flow was originally developed to study the dynamics of turbulence numerically.[49] The initial conditions are smooth, but the flow quickly transitions to turbulence. In the inviscid limit the flow is thought to become singular very rapidly as the vortices stretch and smaller structures are created which are not dissipated.[5] In the viscous case, however, energy is naturally dissipated. Therefore, once the flow has transitioned to turbulence, it immediately begins to decay, mimicking homogeneous non-isotropic turbulence. The simplicity of this case and its wide range of scales and dynamics have made it an attractive case to compare LES and ILES techniques against direct numerical simulation (DNS).[14]

A more practical test case is the simulation of transition in laminar separation bubbles (LSBs) on low-Reynolds-number wings. An LSB can form in a laminar boundary layer when it is subject to an adverse pressure gradient. Once the boundary layer separates, Kelvin-Helmholtz (KH) instabilities become dominant and vortex shedding begins. As the vortices grow downstream, they begin to break down into smaller scale structures. This may be exacerbated by small turbulent structures moving upstream in the recirculation bubble or by acoustic waves generated at the trailing edge. The breakdown continues until the flow is fully turbulent. This increases momentum transport and allows the boundary layer to reattach as a turbulent boundary layer. Evidence of the KH-vortices remains in the turbulent boundary layer as slow moving vortical packets which propagate downstream much more slowly than the bulk flow.[56] LSBs are often classified as long or short. It is common to refer to bubbles having only a small effect on the potential flow as short and those that noticeably alter the pressure distribution as long.[1] Short bubbles do not dramatically alter the performance characteristics of an airfoil, but long bubbles can, often causing rapid changes in lift, drag, and pitching moment.

II. Numerical methods

The Navier-Stokes equations are discretized in space by high-order Summation-By-Parts (SBP) operators and solved on structured multi-block grids. Simultaneous-Approximation-Terms (SATs) are used to enforce block-interface coupling and boundary conditions, while matrix artificial dissipation is used to maintain numerical stability. The resulting system of ordinary differential equations is advanced in time with Explicit-first-stage, Singly-Diagonally-Implicit Runge-Kutta (ESDIRK) methods. A Newton-Krylov algorithm is used to drive the non-linear residual equations to zero. Finally, the linear system is solved with FGMRES with a parallel approximate-Schur preconditioner.

II.A. Summation-By-Parts operators

Summation-by-parts operators are centered finite-difference operators which are constructed to mimic integration-by-parts. Using an energy method, this allows statements to be made about the time-stability of the discretization. For example, the linearized Navier-Stokes equations have been shown to be time-stable,[36] conditional on the use of diagonal norms in curvilinear coordinate systems. In this section we provide a brief description of the operators used in this work, without derivation. For more information on the derivation and analysis of SBP-SAT schemes, see Refs. 12, 13, 29, 30, 33–37, 45–47.

The operator D_1 is a first derivative SBP operator if it approximates the first derivative and has the form $D_1 = H^{-1}\theta$, where H, called the norm, is a positive-definite diagonal matrix, and $\theta + \theta^T = \text{diag}(-1, 0, \ldots, 0, 1)$. As an example, the second-order accurate first derivative SBP operator defined as:

$$D_1 = H^{-1}\theta \tag{1}$$

where

$$H = h\begin{bmatrix} \frac{1}{2} & & & & \\ & 1 & & & \\ & & \ddots & & \\ & & & 1 & \\ & & & & \frac{1}{2} \end{bmatrix}, \quad \text{and} \quad \theta = \frac{1}{2}\begin{bmatrix} -1 & 1 & & & \\ -1 & 0 & 1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 0 & 1 \\ & & & -1 & 1 \end{bmatrix} \tag{2}$$

and $h = \Delta\xi, \Delta\eta,$ or $\Delta\zeta$ for the appropriate coordinate direction.

In the Navier-Stokes equations, these are used to approximate the inviscid flux derivatives and the viscous cross and double-derivatives, for example:

$$\partial_\xi E \approx D_{1,\xi} E, \tag{3}$$

and,

$$\partial_\xi(\phi \partial_\eta q) \approx D_{1,\xi}\phi D_{1,\eta} q, \tag{4}$$

where E is the inviscid flux, ϕ is a spatially varying coefficient and q is some flow quantity. Double-derivatives obtained with the application of first derivative twice do not produce operators with the minimum stencil width or minimum local error.

Compact-stencil second-derivative SBP operators have the form $D_2(\phi) = H^{-1}\{-M + EBD_b\}$.[12] Here, H is the norm, which must be consistent with the norm of first derivative to guarantee time stability, $M = D_1^T HBD_1 + R$ is symmetric positive definite, as is R, the correction term which reduces the stencil width and lowers the truncation error coefficient, $E = \text{diag}(-1, 0, \ldots, 0, 1)$, $B = \text{diag}(\phi_0, \phi_1, \ldots, \phi_N)$, where $\phi_i > 0$, and D_b is an approximation to the first derivative at the boundaries. These operators also mimic integration-by-parts, but for second derivatives.

As an example, the second-order accurate compact-stencil SBP operator defined as:[12]

$$D_2(\phi) = H^{-1}\left\{ -\underbrace{D_1^T HBD_1}_{\text{Term1}} - \underbrace{\frac{1}{4h}\tilde{D}_2^T CB\tilde{D}_2}_{\text{Term2}} + \underbrace{EBD_b}_{\text{Term3}} \right\} \tag{5}$$

where Term1 is the application of the first derivative twice, Term2 = R is the correction term where $\tilde{D}_2$ is a centered undivided difference operator:

$$\tilde{D}_2 = \begin{bmatrix} 1 & -2 & 1 & & \\ 1 & -2 & 1 & & \\ & \ddots & \ddots & \ddots & \\ & & 1 & -2 & 1 \\ & & 1 & -2 & 1 \end{bmatrix}, \quad C = \begin{bmatrix} 0 & & & & \\ & 1 & & & \\ & & \ddots & & \\ & & & 1 & \\ & & & & 0 \end{bmatrix}, \tag{6}$$

and $h = \Delta\xi, \Delta\eta,$ or $\Delta\zeta$ for the appropriate coordinate direction. Finally, Term3 modifies the boundary closure:

$$EBD_b = \frac{1}{h}\begin{bmatrix} \frac{3\phi_1}{2} & -2\phi_1 & \frac{\phi_1}{2} & & \\ 0 & 0 & 0 & & \\ & \ddots & \ddots & \ddots & \\ & & 0 & 0 & 0 \\ & & \frac{\phi_{N+1}}{2} & -2\phi_{N+1} & \frac{3\phi_{N+1}}{2} \end{bmatrix}. \tag{7}$$

II.B. Simultaneous Approximation Terms

Simultaneous approximation terms are penalty terms which provide a weak imposition of block-interface coupling and boundary conditions in conjunction with the SBP operators.

II.B.1. Inviscid terms

The SATs corresponding to the inviscid fluxes terms, E for example in the ξ-direction, have the form:[19,40]

$$\mathrm{SAT}_{\mathrm{inv}} = \mp\frac{1}{J}A_\xi^\pm(Q - Q_{\mathrm{external}}), \tag{8}$$

where Q is the vector of conserved variables, Q_{external} takes on the values of the coincident node on the adjoining block or boundary target values, J is the metric Jacobian of the curvilinear coordinate transformation, the $\pm$ is to account for the difference the high and low sides of the blocks respectively, and $A_\xi^\pm$ is the modified flow Jacobian defined by:

$$A_\xi^\pm = \frac{A_\xi \pm |A_\xi|}{2}, \quad A_\xi = \frac{\partial E}{\partial Q}, \quad |A_\xi| = X_\xi^{-1}|\Lambda_\xi|X_\xi, \tag{9}$$

where X_ξ is the right eigenmatrix of A_ξ, and Λ_ξ is a matrix with the eigenvalues of A_ξ as its diagonal. A small modification to this SAT is required at the outflow in viscous simulations:[40]

$$\mathrm{SAT}_{\mathrm{inv}} = \mp\frac{1}{J}A_\xi^\pm(Q_b - Q_{b-1}) \tag{10}$$

where Q_b is the boundary node and Q_{b-1} is the node one in from the boundary. This modification is not applied when the zonal acoustic boundary treatment is used, described below.

II.B.2. Viscous terms

Block-interface SATs associated with the viscous flux terms have a very similar form to the inviscid SATs:

$$\mathrm{SAT}_{\mathrm{vis}_1} = \mp\frac{\sigma_\nu}{ReJ}B_{\nu,\xi}(Q - Q_{\mathrm{external}}), \tag{11}$$

where $0 \le \sigma_\nu \le 0.5$, and $B_{\nu,\xi}$ is related to the viscous flux Jacobian. In addition, a second SAT is required to account for the increased derivative order:

$$\mathrm{SAT}_{\mathrm{vis}_2} = \mp\frac{1}{Re}(E_\nu - E_{\nu,\mathrm{external}}), \tag{12}$$

where $E_{\nu,\mathrm{external}}$ is the viscous flux of the coincident node on the adjoining block or the target boundary flux value. At the farfield boundary, this is set to zero. This SAT is also modified at solid boundaries to enforce adiabatic or isothermal no-slip conditions. The alternate form at the wall is:

$$\mathrm{SAT}_{\mathrm{viswall}} = \mp\frac{\sigma_{\nu,\mathrm{wall}}}{Re}(Q - Q_{\mathrm{wall}}), \tag{13}$$

where

$$\sigma_{\nu,\mathrm{wall}} \le \frac{\xi_x^2 + \xi_y^2 + \xi_z^2}{J}\frac{\mu}{2\rho}\max\left(\frac{\gamma}{Pr}, \frac{5}{3}\right), \quad Q_{\mathrm{wall}} = \left[\rho_1, 0, 0, 0, \frac{\rho_1 T_*}{\gamma(\gamma-1)}\right], \tag{14}$$

ρ_1 is the density at the boundary node, and T_* is the temperature one node away from the wall for adiabatic conditions or the wall temperature for isothermal conditions.

II.C. Characteristic boundary zones

Characteristic boundary conditions are often used to minimize spurious reflections. The amplitude of incoming waves, $\mathcal{L}_i$, is set to zero at the boundary. However, only linear waves propagating normal to the boundary are completely eliminated. Numerical dissipation often overwhelms the influence of the waves which do get partially reflected, but in simulations which require stringent-accuracy, especially aeroacoustic simulations, these errors can have a significant impact.

A relatively simple technique has been developed by Sandberg and Sandham,[43] who employ characteristic boundary conditions, but non-locally in a buffer zone. The characteristics, λ_i, of the hyperbolic terms are determined along with their amplitudes, $\mathcal{L}_i$. Within the buffer zone, the amplitudes of the incoming waves are ramped to zero,

$$\tilde{\mathcal{L}}_i = g(x)\mathcal{L}_i, \qquad g(x) = 0.5\left[1 + \cos\left(\frac{\pi(x - x_s)}{x_{out} - x_s}\right)\right], \tag{15}$$

where x_s is the onset location of the buffer zone, and x_{out} is the location of the outflow boundary. Notice that the full local characteristic conditions are recovered at the boundary. This non-local boundary scheme has been shown to reduce the reflection of large non-linear structures and of linear waves propagating at incidence to the boundary.[43] It has also been applied with success to a number of aeroacoustic DNS simulations.[25,26,43] The present implementation follows many of the ideas of Kim and Lee.[28]

II.D. Explicit-first-stage, Singly-Diagonally-Implicit Runge-Kutta methods

General s-stage Runge-Kutta methods are described by:

$$\begin{aligned} Y_i &= y^{(n-1)} + h\sum_{j=1}^{s} A_{ij} F(Y_j, t^{(n-1)} + c_i h) \quad \text{for } i = 1, \ldots, s, \\ y^{(n)} &= y^{(n-1)} + h\sum_{j=1}^{s} b_j F(Y_j, t^{(n-1)} + c_i h), \end{aligned} \tag{16}$$

where Y_i are the individual stage values, $y^{(n)}$ the solution at time step n, $h = t^{(n+1)} - t^{(n)}$ is the step size, and A_{ij}, b_j and c_i are the coefficients of the given method, often presented in a Butcher tableau:

$$\begin{array}{c|c} c_i & A_{ij} \\ \hline & b_j \end{array}.$$

II.D.1. Explicit-first-stage and stiff-accuracy

Often the order of the internal stages in a Runge-Kutta method is lower than the global order predicted by classical order theory. Asymptotically, global order convergence is always guaranteed and is also practically realized for relatively non-stiff problems; however, when implicit Runge-Kutta methods are applied to very stiff or differential algebraic problems with finite step sizes, the local error of these internal stages can dominate. This is known as order reduction.

In CFD applications, order reduction is not observed in inviscid or laminar problems; however, order reduction can manifest in URANS simulations.[3,7] Forcing the first stage of ESDIRK methods to be explicit, allows the internal stages to have order two. The local error can be further reduced by enforcing stiff-accuracy, namely $c_s = 1$, and therefore $b_j = A_{sj}$. These conditions imply that the minimum convergence of an ESDIRK method for a stiff ODE will be at least third order. The conditions for stiff-accuracy also mean that the explicit stage needs only to be computed once.

II.D.2. Singly-Diagonally-Implicit methods and stability

Fully implicit Runge-Kutta methods can be generated which have order $2s$, where s again is the number of stages. This is very attractive for lowering the local truncation error and for increasing convergence with step size. However, fully implicit RK methods require an implicit solution to a system of size $(s \times n)^2$, where n is the number of unknowns. For large systems of equations, this can be very expensive in terms of both CPU time and memory. In contrast, diagonally-implicit methods only require the solution to s systems of size n^2. Despite requiring more implicit stages to achieve the same order, diagonally-implicit methods are often more efficient than fully implicit methods.

Letting the diagonal entries of A_{ij} be constant, except A_{11}, which is zero, means that the temporal component of the Jacobian (20) is constant. This can be exploited during the solution process to reduce computational costs. More information can be found in Section II.E.3.

Methods in the ESDIRK class of time-integrators are unconditionally stable (A-stable) and provide complete damping of modes at infinity (L-stable). This is particularly advantageous when the governing equations are stiff. The size of the time steps are, therefore, only limited by accuracy and not stability.

II.D.3. Runge-Kutta methods and order of accuracy

It is well known that A-stable implicit Linear Multistep Methods (LMMs) are limited to second-order.[11] However, this restriction does not apply to implicit Runge-Kutta methods; arbitrarily high-order methods can be generated which are A-stable. High-order methods are desired since they have the potential be be more efficient, especially for simulations which require stringent accuracy.

II.D.4. ESDIRK methods and local truncation error

Incorporating these ideas, the Butcher tableau of an ESDIRK method is of the form:

$$\begin{array}{c|ccccc} 0 & 0 & 0 & 0 & \dots & 0 \\ 2\tilde{\lambda} & \tilde{\lambda} & \tilde{\lambda} & 0 & \dots & 0 \\ c_3 & a_{31} & a_{32} & \tilde{\lambda} & \dots & 0 \\ \vdots & & \vdots & & \ddots & \vdots \\ 1 & a_{s1} & a_{s2} & a_{s3} & \dots & \tilde{\lambda} \\ \hline & a_{s1} & a_{s2} & a_{s3} & \dots & \tilde{\lambda} \end{array}.$$

A final advantage of ESDIRK methods is the relatively small local truncation error coefficients, as seen in Table 1, which compares some common time integration methods. Taking into account the increased number of implicit stages, if the cost of one implicit solve is assumed to be approximately constant, it is easy to see that ESDIRK methods of a given order are very efficient.

II.E. Solution methodology

The semi-discrete form of the Navier-Stokes equations is,

$$\frac{dQ}{dt} = -\tilde{R}(Q), \tag{17}$$

where $\tilde{R}(Q)$ represents the spatially discretized equations, including the numerical dissipation model, and Q represents the solution vector at all points. Using an ESDIRK scheme with s stages, the fully discrete system of non-linear equations becomes,

$$R_k^{(n)}(Q_k^{(n)}, \dots, Q_1^{(n)}, Q^{(n-1)}) = \frac{Q_k^{(n)} - Q^{(n-1)}}{\tilde{\lambda}\Delta t} + \frac{1}{\tilde{\lambda}} \sum_{j=1}^{s} A_{kj} \tilde{R}(Q_j^{(n)}) = 0, \quad k = 2, \dots, s, \tag{18}$$

where $R(Q)$ is the residual. The residual equations are solved by applying Newton's method with iteration counter p:

$$\mathcal{A}_k^{(p)} \Delta Q_k^{(p)} = -\hat{R}^{(n)}(Q_k^{(p)}, Q_{k-1}^{(n)} \dots, Q_1^{(n)}, Q^{(n-1)}), \quad k = 2, \dots, s, \tag{19}$$

where $\Delta Q_k^{(p)} = Q_k - Q_k^{(p-1)}$ and $\mathcal{A}$ is the Jacobian:

$$\mathcal{A}_k^{(p)} = \frac{\mathcal{I}}{\tilde{\lambda}\Delta t} + \frac{\partial R(Q_k^{(p)})}{\partial Q_k^{(p)}}, \quad k = 2, \dots, s. \tag{20}$$

Method	Global Order	External Steps	Total Stages	Implicit Stages	Stability	$\|LTE\|$
BDF2	2	2	1	1	L-Stable	$\approx 0.33z^3$
BDF2OPT(0.5)[52]	2	2	1	1	L-Stable	$\approx 0.16z^3$
ESDIRK2/TRBDF2	2	1	3	2	L-Stable	$\approx 0.04z^3$
BDF3	3	3	1	1	L(86.03°)-Stable	$0.25z^4$
ESDIRK3	3	1	4	3	L-Stable	$\approx 0.0259z^4$
RK4	4	1	4	0	Conditional	$\approx 0.0083z^5$
BDF4	4	4	1	1	L(73.35°)-Stable	$0.2z^5$
MEBDF4[10]	4	3	3	3	L-Stable	$\approx 0.0892z^5$
SDIRK4	4	1	3	3	L-Stable	$\approx 0.1644z^5$
ESDIRK4	4	1	6	5	L-Stable	$\approx 0.0008z^5$

Table 1. List of time marching methods and associated characteristics, $z = \lambda h$, where λ comes from the linear test equation $y' = \lambda y$.

II.E.1. Inexact Newton's Method

Newton's method requires a Jacobian; however, the use of a Krylov method means that only the Jacobian-vector products need to be computed. These are evaluated with first-order forward differences:

$$\mathcal{A}^{(p)}\mathbf{v} = \frac{R(Q^{(p)} + \epsilon\mathbf{v}) - R(Q^{(p)})}{\epsilon}, \tag{21}$$

where

$$\epsilon = \sqrt{\frac{N_u\delta}{\mathbf{v}^T\mathbf{v}}}, \tag{22}$$

N_u is the number of unknowns, and $\delta = 10^{-12}$.

An approximate first-order Jacobian is constructed for the approximate-Schur preconditioner described by Saad and Sosonkina[42] and Hicken and Zingg.[19] The inviscid fluxes are evaluated to second-order along with a second-difference dissipation model. The viscous terms are also modified to reduce their stencil width. They are similarly evaluated to second order, but in addition, the cross derivate terms are dropped and the viscous SATs are modified to neglect tangential derivatives. The final modification is to the viscosity term, which is treated as constant. This reduces the order of the Jacobian to first-order, and more significantly, reduces the number of matrix entries, and therefore, memory requirements.

Finally, the linear system is not solved exactly; rather it is only solved so far as to satisfy the inequality:

$$||R^{(p)} + \mathcal{A}^{(p)}\Delta Q^{(p)}||_2 \leq \eta_n||R^{(p)}||_2 \tag{23}$$

where η_n is a specified forcing parameter. A larger value can reduce simulation time by preventing the linear-system from being over-solved; however, if it is too large, the non-linear convergence will be hindered. A typical value is 0.01.

II.E.2. Polynomial extrapolation

The performance of Newton's method can be improved by providing better initial iterates. Previous solution information can be used to generate low-order inexpensive approximations of the next time step.

Consider a sequence of solution values $u^{(n-1)}, \ldots, u^{(n-k)}$ at times $t^{(n-1)}, \ldots, t^{(n-k)}$. These times do not have to be equally spaced or monotonic. The solution $u^{(n)}$ at time $t^{(n)}$ is then approximated by:

$$u^{(n)} = \sum_{i=1}^{k} l_{n-i}(t^{(n-i)})u^{(n-i)}, \tag{24}$$

where

$$l_{n-i}(t^{(n-i)}) = \prod_{j=1, j\neq i}^{k+1} \left(\frac{t^{(n)} - t^{(n-j)}}{t^{(n-i)} - t^{(n-j)}} \right). \tag{25}$$

Increasing the number of past solutions increases the accuracy of the approximation. In this work, three past solutions are used as a balance between accuracy and memory usage.

II.E.3. Delayed preconditioner updates

The temporal component of the Jacobian (20) is constant and is often significantly larger than the change in the spatial Jacobian over a stage or an entire time step. Therefore, it is possible to freeze the preconditioner over a stage or time step without a significant impact on the convergence of the system. This reduces CPU time and thus increases the efficiency of the solution algorithm. Current results were obtained by freezing the preconditioner over each time step.

II.E.4. Termination of non-linear iterations

The temporal integration has a certain level of truncation error associated with it. The convergence of the residual equations can, therefore, be terminated when the residual is less than this error. This reduces computational cost and is done without any loss in global accuracy. In this work, termination is based on a preset reduction from the initial residual value. The necessary relative tolerance is fairly step size independent since a reduction in step size will result in a better initial iterate from polynomial extrapolation and therefore a lower initial residual. A typical relative tolerance is 10^{-6},[48] requiring between 5 and 15 Newton iterations.

III. Results and Discussion

All results in this section, unless explicitly stated otherwise, were obtained with the fourth-order non-compact-stencil spatial discretization and the ESDIRK4 time-marching method.

III.A. Taylor-Green vortex flow

The first case is the Taylor-Green vortex flow. It was originally developed to study vortex stretching, the creation of small eddies from larger ones, believed to be an important mechanism in turbulent flows.[49] The flow is initialized with a smooth two-dimensional uni-modal velocity field. As the solution develops, higher frequency modes are generated, eventually mimicking homogeneous non-isotropic turbulence. Finally, the turbulence decays as the smallest modes are dissipated due to viscous effects. The initial conditions are:

$$\begin{gathered} u = M_\circ \sin(x)\cos(y)\cos(z), \\ v = -M_\circ \cos(x)\sin(y)\cos(z), \\ w = 0, \\ p = p_\circ + \tfrac{\rho_\circ M_\circ^2}{16}\left(\cos(2x) + \cos(2y)\right), \\ \rho = p/p_\circ, \end{gathered}$$

where $\rho_\circ = 1$ and $p_\circ = \frac{1}{\gamma}$. To minimize the effects of compressibility and to be consistent with the AIAA's 1st International Workshop on High-Order CFD Methods, the free-stream Mach number is set to 0.1. The Reynolds number is 1600, which corresponds to a peak Taylor microscale Reynolds number of about 22, and the Prandtl number is 0.71. The convective time unit is defined as $t_c = \frac{1}{M_\circ}$ and the simulation is advanced to $t_{\text{final}} = 20t_c$. The simulation domain is a periodic box, $-\pi \leq x, y, z \leq \pi$. The reference for this study is the incompressible dealiased spectral DNS of van Rees et al.[51] which contained 512^3 modes. This reference is denoted as 'Spectral' in the Figures.

III.A.1. Basic Definitions

In this study, kinetic energy is defined as:

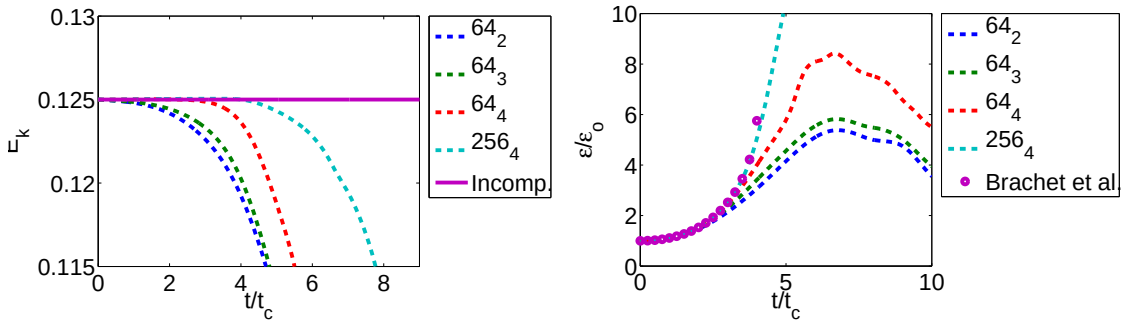

Figure 1. Taylor-Green flow: Inviscid solutions compared with exact incompressible solution (kinetic energy) and semi-analytic solution of Brachet et al.[5] (enstrophy).

$$E_k = \frac{1}{2V}\int_V \rho \mathbf{v}\cdot\mathbf{v} dV,$$

where V is the volume, $\mathbf{v}$ is the velocity vector and ρ is the density. The dissipation rate can then be computed as $\epsilon = -\frac{dE_k}{dt}$. Enstrophy is defined as,

$$\varepsilon = \frac{1}{2V}\int_V \rho\boldsymbol{\omega}\cdot\boldsymbol{\omega} dV,$$

where $\boldsymbol{\omega}$ is the vorticity vector. Finally, the Courant (CFL) numbers quoted are global maximums defined based on the initial conditions. The CFL number at any node is:

$$\text{CFL} = \kappa^*_{\text{inv,max}}\Delta t\left(|U^{(\alpha)}| + a|\nabla\alpha|\right),$$

where a is the speed of sound, U is the contravariant velocity, and $\alpha = \{\xi, \eta, \zeta\}$, and $\kappa^*_{\text{inv,max}}$ is the maximum value of the Fourier symbol corresponding to the derivative operator of the inviscid flux; in this case it is also the maximum value of the modified wave number.

III.B. Inviscid Simulations

Initially the Taylor-Green flow is computed in the inviscid limit. Figure 1 shows the evolution of kinetic energy on a 64^3 grid for different spatial orders of accuracy, along with a fourth-order fine grid solution. These simulations are advanced with a time step of $\Delta t = 0.2$. Also shown is the incompressible solution, for which kinetic energy is preserved exactly. The present simulations are slightly compressible and have the addition of numerical dissipation to maintain stability, leading to the deviation viewed in the figure. The figure shows a greater improvement from third to fourth-order than from second to third, a result of the increased order of the dissipation model required to maintain global order four.

Figure 1 shows a similar result for the evolution of normalized enstrophy ($\epsilon/\epsilon_\circ$). Brachet et al.[5] suggest that a singularity exists in the inviscid limit around $t_c \sim 5.2$, which means that the normalized enstrophy should go to infinity. The fine grid solution matches the semi-analytic solution of Brachet et al.[5] up to $t_c \sim 4$; however, with finite resolution and the inclusion of numerical dissipation, the enstrophy eventually peaks, and then decays.

III.C. High-Order SBP Integration

Preliminary high-order inviscid results were stable and converged; however, integration of the total kinetic energy initially increased in time before decaying as shown in Figure 2. Also, integration of the conserved quantities showed large variations in time, phenomena not present in second-order solutions. After the code and initial conditions were verified, attention turned to the quadrature used to obtain the results, the classical second-order trapezoidal rule. High-order quadratures were sought to eliminate the influence of the spatial discretization order on the results. Hicken and Zingg[22] showed that the norm of an SBP operator, H, is related to the trapezoidal rule, but with end corrections to match the order of the operator. Using the diagonal norm associated with the spatial discretization to perform the integration, alleviated the

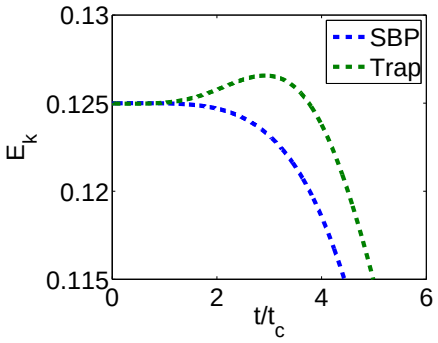

Figure 2. Taylor-Green flow: Fourth-order inviscid solutions comparing second-order trapezoidal and fourth-order SBP quadratures.

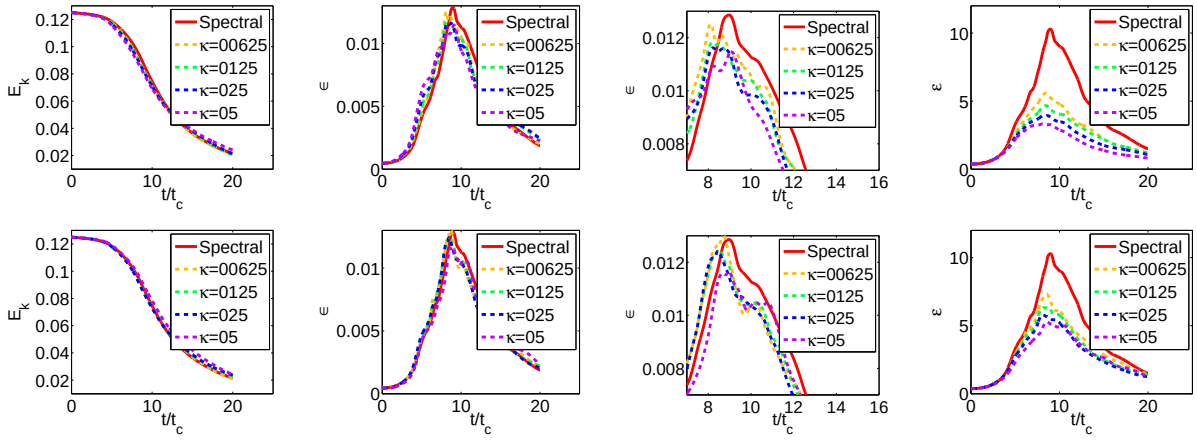

Figure 3. Taylor-Green flow: effect of the numerical dissipation coefficient, κ. Second-order (upper), and fourth-order (lower) spatial discretization with ESDIRK4 time marching and $\Delta t = 0.25$.

above problems. The variation of conserved quantities over time becomes negligible and the kinetic energy monotonically decreases in time, as seen in Figure 2. The original second-order solutions did not show these deviations because the diagonal norm associated with the second-order spatial discretization, is the classical trapezoidal rule.

III.C.1. Numerical Dissipation

The amount of numerical dissipation is a critical factor in high-resolution simulations, especially in ILES simulations. Too much numerical dissipation will damp important information being computed, too little will admit spurious modes. The amount of numerical dissipation applied in the present simulations is controlled by a coefficient, κ, and the grid density. Larger values of κ correspond to more dissipation, of both high-frequency and low-frequency modes.

A study investigating the optimal κ value is undertaken using the 128^3 grid for both second and fourth-order spatial discretizations. The order of accuracy of the dissipation model is set equal to or greater than the global order of the spatial discretization, and solutions are advanced with a non-dimensional time step of $\Delta t = 0.25$. The evolution of the spatially integrated quantities is displayed in Figure 3.

Enstrophy is taken as an indication of the resolving power of the discretization. As the flow transitions to turbulence, smaller scales are generated; increasing the small-scale vorticity content of the closed system leads to an increase in enstrophy. As mentioned in Section III.B, the Taylor-Green vortex flow is thought to become singular in the inviscid limit, driving enstrophy to infinity.[5] Viscous effects in Navier-Stokes simulations dissipate the smallest scales giving the enstrophy profile presented by van Rees et al.[51] In addition to viscous effects, the present simulations introduce numerical dissipation as a function of grid density, damping the small scale content of the flow to maintain numerical stability. This also limits the maximum enstrophy when the numerical dissipation is active above the viscous scale.

Order	Grid size	Avg Newton iters./ ESDIRK stage	Avg GMRES iters./ ESDIRK stage	Walltime (sec.)	Walltime × CPUs (sec.)
2	32^3	12.0	13.1	6822	6822
2	64^3	11.5	12.5	16707	133654
2	128^3	11.3	12.5	36389	2328895
2	256^3	10.5	12.1	82875	42431820
4	32^3	10.9	19.0	15930	15930
4	64^3	10.7	16.2	43326	346607
4	128^3	10.6	16.6	92404	5913834
4	256^3	10.8	16.8	219496	112381866

Table 2. Taylor-Green flow: Computational details of grid convergence study.

As the dissipation coefficient is lowered, the maximum enstrophy increases, indicating more modes are being well-resolved. Lowering the dissipation coefficient also results in better capture of the location and rate of energy dissipation. All of this is expected. However, the lowest dissipation value also results in slight oscillations in the dissipation rate and an over prediction of the peak dissipation rate in the fourth-order simulation. This may be an indication that spurious modes are being admitted to the solution. Therefore, a value of $\kappa = 0.0125$ is selected for the rest of the Taylor-Green simulations.

III.C.2. Grid convergence studies

A grid convergence study is conducted for second and fourth-order spatial discretizations on four successively finer grids: 32^3, 64^3, 128^3 and 256^3 nodes. Each grid is decomposed into blocks of 32^3 nodes with a one-to-one distribution of blocks to processors, keeping the workload per processor constant. The time step is chosen to be constant across simulation of equal grid density and to maintain a constant maximum CFL number across simulations of equal spatial order, ~ 31 for second-order and ~ 50 for fourth-order. The discrepancy is due to the difference in the maximum value of the modified wave number between the second and fourth-order discretizations.

The computational details of the simulations are displayed in Table 2, and simulation results are presented in Figure 4. In all cases, the coarsest grids were not able to accurately capture the decay of kinetic energy. The higher-frequency modes cannot be represented on these grids and are damped by the numerical dissipation in order to maintain stability. As a consequence, less energy is transferred to the higher frequency modes, and this is believed to be the cause of the lower dissipation rate and the higher kinetic energy later in the simulation. The fourth-order simulations do not dissipate as early as the second-order simulations; however, there is still a noticeable deviation from the DNS results.

The finer-mesh simulations, both second and fourth-order, more accurately capture the decay of kinetic energy. These simulations are isolated in Figure 5. The second-order solution simulated on the 128^3 grid still dissipates too early, and only the finest fourth-order simulation lies on top of van Rees' DNS results. However, the accuracy of the evolution of kinetic energy is somewhat surprising considering the large deviation in enstrophy from the DNS results, which suggests that even the finest simulation is under resolved. The second-order 256^3 and fourth-order 128^3 results are similar, accurately capturing the initial dissipation, but under predicting the final dissipation rate later in the simulation. The difference is the CPU time; the fourth-order simulation required approximately 85% less CPU time than the finer second-order simulation.

Contours of the vorticity norm at one of the periodic faces, $x = \pi$, are shown in Figure 6 for the fourth-order result obtained on the finest grid. The structures presented by van Rees et al.[51] are recovered; however, extra structures are visible in the present simulations. These structures are fairly large, but are formed by the lowest vorticity contour lines. Analytically, the y and z-components of vorticity should be identically zero at this face. Removing these components and showing only the x-component of the vorticity, the structures then match very well. The deviation is likely due to the difference in spatial resolution or the effects of compressibility in the present simulation. Figure 7 shows contours of the x-component of vorticity of the fine grid simulations for both second and fourth-order. The high-vorticity component of the second-order 128^3 result is weak and very concentrated. There is indication of an annular structure, but it is only formed by the lowest contour lines. Increasing the order of the simulation, increases maximum vorticity, which now becomes

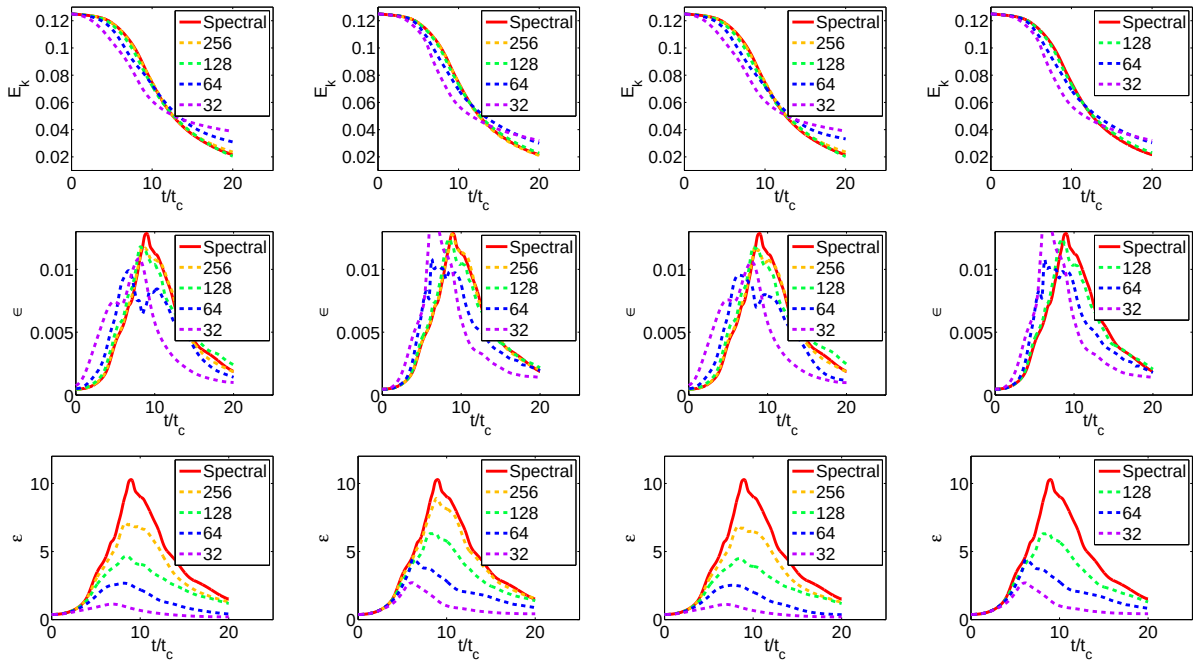

Figure 4. Taylor-Green flow: grid convergence. Legend entries denote spatial resolution in each coordinate direction, i.e. the number of nodes.

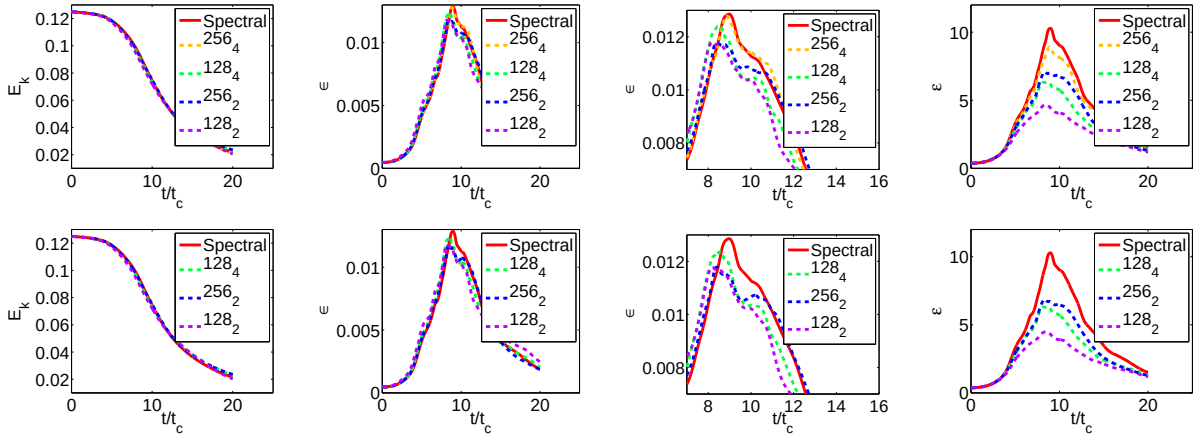

Figure 5. Taylor-Green flow: fine mesh simulation results. (upper) non-compact-stencil discretization, and (lower) compact-stencil discretization. Legend entries denote spatial resolution in each coordinate direction, i.e. the number of nodes, and subscripts denote spatial order of accuracy.

part of the annular structure. It remains less defined than the DNS result, and there are some erroneous artifacts. Increasing the spatial resolution to 256^3, the second-order result becomes more well-defined than the fourth-order result on the coarser grid, and is able to removes the erroneous artifacts. The former may be a result of the interpolation used to create the contour being computed from a finer grid. Regardless, the location and strength of the vorticity is only marginally closer to the DNS results. The structure now has a good likeness to the DNS result, though it remains slightly compressed. Finally, the fourth-order result on the finest grid has the high maximum vorticity, is even more decompressed, and is very similar to the DNS result.

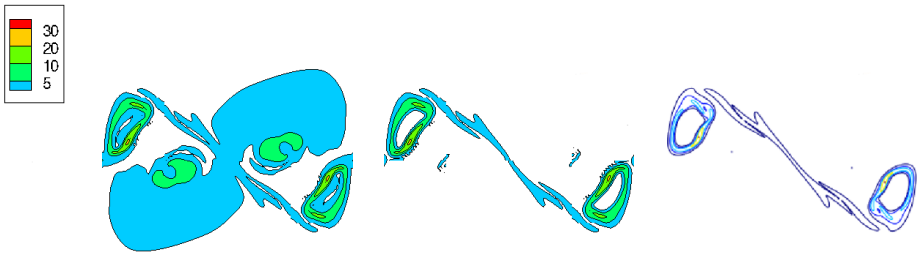

$|\omega|$ $|\omega_x|$ $|\omega|$ - van Rees et al.[51]

Figure 6. Taylor-Green flow (256^3 grid): vorticity norm and x-component of vorticity of the present study compared with vorticity norm from van Rees et al.[51] ($\frac{1}{M_\circ}|\omega|$, $\frac{1}{M_\circ}|\omega_x| = 1, 5, 10, 20, 30$)

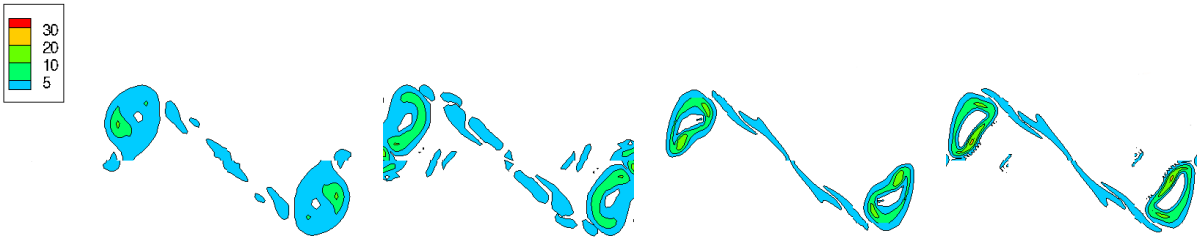

128^3 grid - Second-order 128^3 grid - Fourth-order 256^3 grid - Second-order 256^3 grid - Fourth-order

Figure 7. Taylor-Green flow: x-component of vorticity ($\frac{1}{M_\circ}|\omega_x| = 1, 5, 10, 20, 30$).

III.C.3. Temporal convergence studies

The temporal accuracy and efficiency of a few common time integration methods are evaluated in a temporal convergence study with time steps $\Delta = 0.003125, \ldots, 0.2$; corresponding to CFL $\approx 3, \ldots, 195$. These simulations, along with a reference solution obtained with the classical fourth-order Runge-Kutta (RK4) method and a time step of 0.00078125, CFL ≈ 0.76, were computed on a 128^3 grid. The use of a reference case, computed on the same grid, eliminates the influence of the spatial discretization error, thus isolating the temporal error. The error is computed as the root-mean-square of the difference in kinetic energy:

$$e_{\text{rms}} = \sqrt{\frac{\sum_{i=1}^{m}(E_{k,i} - E_{k,i,ref})^2}{m}}$$

where m is the number of time steps.

The temporal convergence and efficiency of the methods are shown in Figure 8. Not shown in the figure is BDF3, which was not able to produce a stable result. The design order of every other method is recovered; however, the main result is the efficiency of the ESDIRK methods: the CPU time required to obtain a preset level of error. While the trapezoidal method is more efficient than ESDIRK2, the trapezoidal method is not L-stable and its stability contour lies on the imaginary axis which could lead to stability issues in other cases. Regardless, the second-order methods are efficient only for simulations requiring a minimum level of temporal accuracy. As the required accuracy becomes more stringent, ESDIRK4 quickly and decidedly becomes more efficient.

III.D. SD7003 Wing

A major concern with low-Reynolds number flows over wings is the formation of laminar separation bubbles. At moderate angles of attack, the laminar boundary layer becomes susceptible to adverse pressure gradients, causing the formation of the Laminar Separation Bubbles (LSBs). The focus for this paper is to examine the ability of the numerical methods to efficiently predict separation, transition in the separated shear layer, and turbulent reattachment. In future studies we hope to further examine the resulting turbulent boundary layer, the acoustic field, and different passive flow control devices.

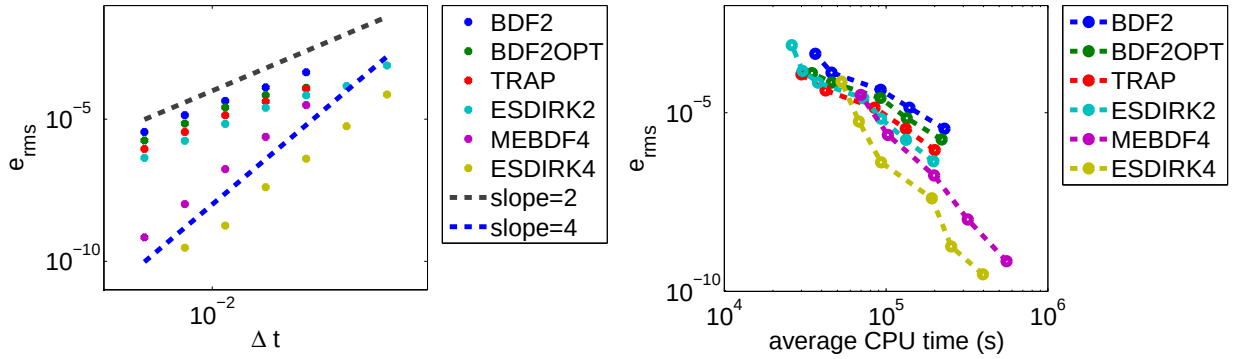

Figure 8. Taylor-Green flow: temporal convergence and efficiency

The SD7003 airfoil is chosen for this study and extruded to a rectangular unswept wing. The flow is simulated with periodic boundary conditions in the spanwise direction to emulate an infinite wing. The parameters for these studies are a Reynolds number of 60,000, a Mach number of 0.2, and angles of attack 4° and 8°. Each simulation is initialized with a fully developed two-dimensional flow field. The solution is extruded to three-dimensions and allowed to develop for $10 - 20$ convective time units before averaged data is collected. The data is averaged over 4 convective time units and over the span.

III.D.1. Definitions

The point of separation is defined as the time-averaged location along the chord line that the friction coefficient becomes negative. Similarly, the reattachment point is the time-averaged location along the chord where the friction coefficient becomes positive again. Defining a point for transition is more difficult since transition occurs over a whole region; however, for the sake of comparison, the time-averaged location at which the normalized Reynolds stress, $\tau_{xy} = -\overline{u'v'}/u_\infty^2$, reaches the critical value of 0.001, is defined as the transition point.[56]

III.D.2. Grids

The SD7003 geometry is modified with a rounded trailing edge to facilitate the use of O-grids. This is not necessary, but is done to be consistent with the 1^{st} AIAA International Workshop on High-Order CFD methods. The grid is generated in two dimensions and then extruded to three dimensions, with the farfield placed at 15 chords, and the spanwise extent set to 0.2 chords. 126 grid points are placed in the off-wall direction, 251 on the upper surface, and 101 on the lower surface and in the spanwise direction. The acoustic boundary zone comprises the 6 outermost nodes in the farfield. The off-wall spacing is 10^{-6}, corresponding to an average $y^+ \sim 5$, and the farfield spacing is set to 1, all measured in chord units. In total, the two and three-dimensional grids have approximately 47,000 and 4.9×10^6 grid points respectively. Figure 9 shows every 5^{th} node of the complete SD7003 grid, and Figure 10 shows a close-up of the near-wall region (denoted as "Baseline").

The performance of the SD7003 grid is evaluated with a series of three finer grids. The first grid has a refined the near-wall region (shown in Figure 11), increasing the number of streamwise nodes on the upper surface and the number of nodes in the off-wall direction by 50 and 25 nodes respectively. The second grid extends the refinement to the rest of the grid, further increasing the number of nodes by 50 and 25 in the circumferential and off-wall directions. The final grid contains 701, 301, and 201 grid points in the circumferential, off-wall, and spanwise directions respectively. This provides a very fine reference against which to evaluate the performance of the other grids. Close-ups of the grids are shown in Figure 10.

This study is carried out at 4° angle of attack with a time step of $\Delta t = 10^{-2}$. Table 3 compares the time-averaged lift and drag coefficients, and Figure 12 presents the pressure and skin friction coefficient distributions for the grids. The points of separation and transition are all nearly identical across the grids. Up until this point the flow is considered to be laminar, so it is not surprising that the results are consistent. There is small variation in the reattachment point; however it does not significantly affect the time-averaged lift and drag values.

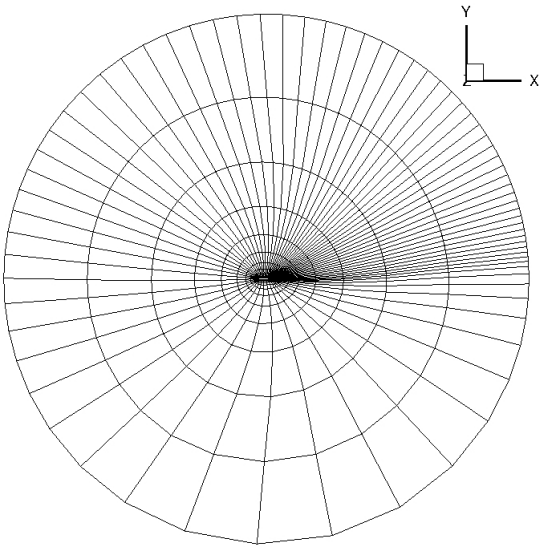

Figure 9. SD7003: Full grid with every 5^{th} node shown.

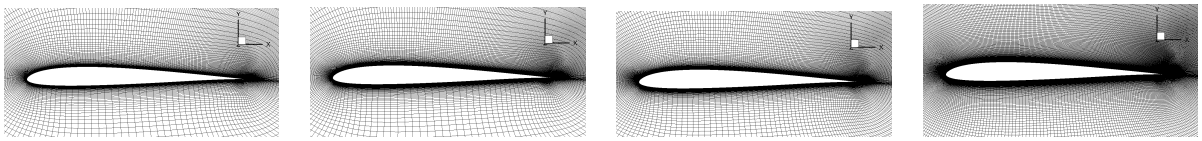

Baseline Grid 1 Grid 2 Grid 3

Figure 10. SD7003: Grids for grid refinement study

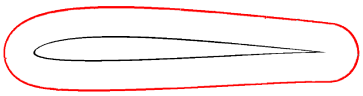

Figure 11. SD7003: Near-wing boundary region definition.

	Separation	Transition	Reattachment	c_l	c_d
Baseline	21%	51%	64%	0.6009	0.0222
Grid 1	21%	48%	65%	0.5970	0.0224
Grid 2	21%	49%	66%	0.5991	0.0228
Grid 3	22%	-	64%	0.5958	0.0217

Table 3. SD7003: Time-averaged locations of separation, transition, and reattachment along with time-average lift and drag coefficients for the grid refinement study.

III.D.3. Order

A comparison of results obtained with second and fourth-order spatial discretizations are found in Figure 13, overlaid with results from Galbraith and Visbal[15] and Zhou and Wang.[56] The results are obtained with the baseline grid at an angle of attack of 4°, and a time step of 10^{-2}. At this grid resolution, the second-order solution deviates significantly from the published results. The pressure recovery after the pressure plateau occurs noticeable sooner and is much shallower. The flow separates slightly further downstream, and reattaches noticeable upstream. The fourth-order results are significantly closer to the published results shown and to those presented in Table 5. Furthermore, the fourth-order simulation only requires approximately 2.5 times more CPU time. If we assume perfect scaling, this would only allow approximately 35% more

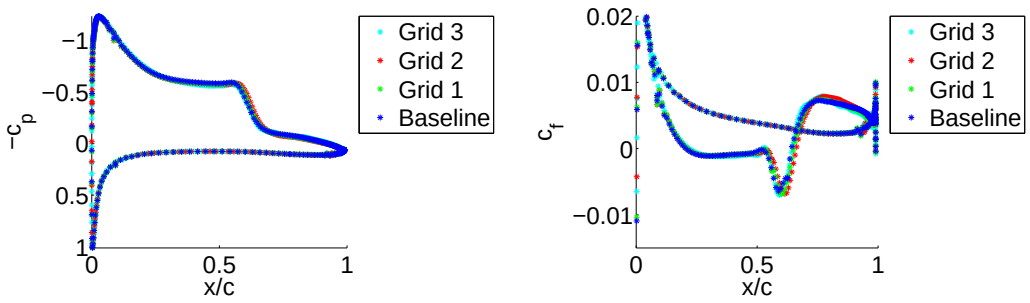

Figure 12. SD7003: time averaged pressure and friction coefficients for grid study.

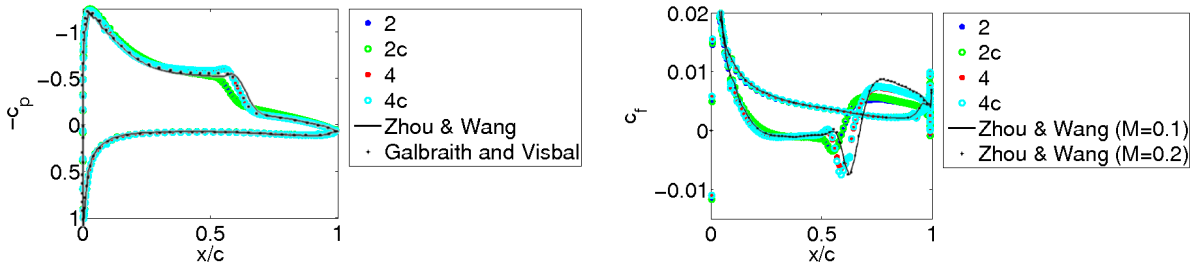

Figure 13. SD7003: order comparison. Legend entries denote spatial order of accuracy with 'c' marking results obtained with the compact stencil spatial discretization.

	c_l	c_d
$\Delta t = 0.08$	0.6008	0.0222
$\Delta t = 0.04$	0.6010	0.0223
$\Delta t = 0.02$	0.6007	0.0222
$\Delta t = 0.01$	0.6009	0.0222
$\Delta t = 0.005$	0.6005	0.0222
$\Delta t = 0.0025$	0.6005	0.0222

Table 4. SD7003: Time averaged lift and drag for time step study.

nodes to be used in each direction of the second-order simulation to have the same computational cost. It is unlikely that with this increased resolution, the accuracy of the second-order simulation competitive with the fourth-order result.

III.D.4. Temporal Integration

The temporal discretization plays a large role in the efficiency of the overall solution process. To determine the most efficient time step size, an investigation using the baseline grid is undertaken, evaluated by the accuracy of time-averaged lift and drag coefficients. The results of the study are shown in Table 4. Within the accuracy of the approximations, the lift and drag coefficients appear to be step size independent. The Courant number of the simulation with largest step size is on the order of 10^4, and while the simulation converges and is stable, the solution of the linear system becomes difficult. The parameters required to ensure convergence of the simulations with the largest step sizes, render them inefficient. Therefore, the step size chosen for the rest of the simulations is $\Delta t = 0.01$.

III.D.5. Angle of Attack

Angles of attack of 4° and 8° are common test conditions for the simulation of the SD7003. These were the subject of investigation at the 1^{st} AIAA International Workshop on High-Order CFD Methods. Using the baseline grid, Table 5 compares the time-averaged results with previously published values, and Figure 14 shows the distribution of the force coefficients. The results for the 4° case match very well to other

Study	Type	Separation	Transition	Reattachment	c_l	c_d
4° Angle of Attack						
TU-BS[41]	Exper.	30%	53%	62%	-	-
HFWT[38]	Exper.	18%	47%	58%	-	-
Carton de Wiart and Hillewaert[8]	DNS	21%	-	65%	0.60	0.022
Uranga et al.[50]	XFoil	26%	57%	60%	0.58	0.018
Uranga et al.[50]	ILES	24%	51%	60%	0.60	0.022
Castonguay et al.[9]	ILES	23%	52%	65%	-	-
Zhou and Wang[56]	ILES	22.7%	52.1%	68.5%	-	-
Galbraith and Visbal[15]	ILES	23%	55%	65%	-	-
Galbraith[16]	ILES	23%	55%	65%	0.59	0.021
Present (2^{nd}-order)	ILES	23%	42%	60%	0.590	0.019
Present (4^{th}-order)	ILES	21%	51%	64%	0.601	0.022
8° Angle of Attack						
Garmann and Visbal[17]	ILES	2.3%	-	26%	0.970	0.039
Zimmerman and Wang[57]	ILES	4%	-	31%	0.916	0.049
Galbraith[16]	ILES	4%	18%	28%	0.92	0.043
Present	ILES	3.7%	10.5%	20%	0.968	0.034

Table 5. SD7003: Experimental and Computational results for transitional flow over the SD7003 wing.

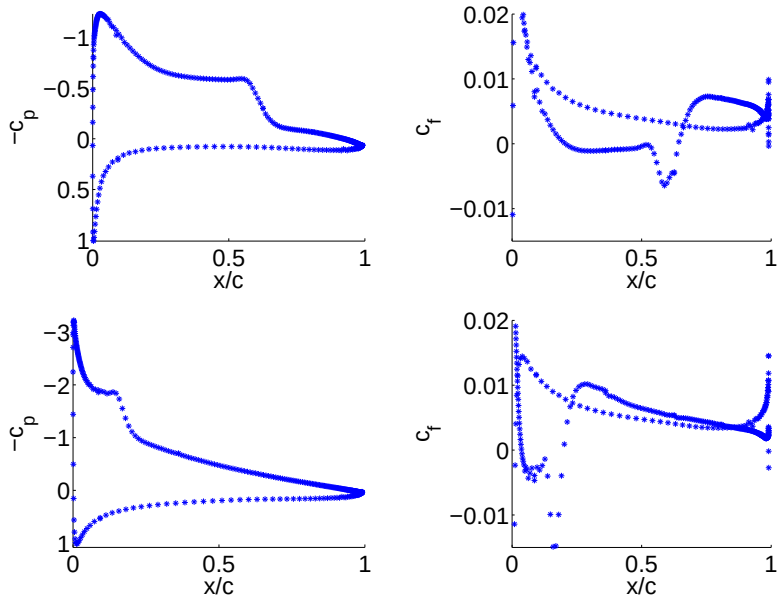

Figure 14. SD7003: Averaged pressure and friction coefficient distributions. Angle of attack 4° (upper), and 8° (lower).

numerical simulations, and discrepancies with experimental results are thought to be caused by low intensity free-stream turbulence not present in the numerical simulations. At angle of attack 8°, the general form of the force distributions match very well with the literature, as well as the lift coefficient. The location of separation is also very good; however, transition and reattachment occur upstream of the locations seen in the other computations.

IV. Conclusions

Investigation of high-order methods, using the Taylor-Green vortex flow, shows the clear advantage in efficiency of high-order methods when stringent accuracy is required. The fourth-order spatial discretization obtains a similar evolution of dissipation rate and enstrophy to the second-order discretization, but with half as many grid points and time steps. The reduction in mesh and time step requirements translates to an 85% reduction in CPU time. The second-order solution on the finer mesh is slightly better at capturing the strength and structure of the vorticity, but does not warrant the significant increase in computational cost. The evaluation of the temporal discretization shows a clear advantage to the higher-order ESDIRK4 method for simulations requiring more than a basic level of accuracy.

Results from the simulation of transitional flow over the SD7003 are in good agreement with experimental and computational results, even with fairly coarse spatial and temporal resolution. The high-order spatial discretization accurately captures the separation, transition in the separated shear layer, and finally turbulent reattachment. This is not the case with the second-order discretization when the same simulation parameters are used; the force distribution profiles significantly deviate from published results.

These results give a clear picture of the benefits of high-order discretizations along with the advantages of the novel parallel Newton-Krylov-Schur algorithm for high-accuracy unsteady fluid simulation.

Acknowledgments

Computations were performed on the GPC supercomputer at the SciNet HPC Consortium. SciNet is funded by the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund -Research Excellence; and the University of Toronto.[32]

References

[1]Alam, M. and Sandham, N., "Direct numerical simulation of 'short' laminar separation bubbles with turbulent reattachment", *J. Fluid Mech.*, Vol. 410, 2000, pp. 1-28.

[2]Alexander, R., "Diagonally Implicit Runge-Kutta Methods for Stiff O.D.E.'s", *SIAM J. Num. An.*, Vol. 14, No. 6, 1977, pp. 1006-1021.

[3]Bijl, H., Carpenter, M.H., Vatsa, V.N. and Kennedy, C.A., "Implicit Time Integration Schemes for the Unsteady Compressible Navier-Stokes Equations: Laminar Flow", *Journal of Computational Physics*, Vol. 179, No. 1, 2002, pp. 313-329.

[4]Boris, J., "On Large Eddy Simulation Using Subgrid Turbulence Models", *Whither Turbulence? Turbulence at the Crossroads*, Vol. 357, 1990, pp. 344-353.

[5]Brachet, M.E. et al., "Small-Scale Structure of the Taylor-Green Vortex", *J. of Fluid Mech.*, Vol. 130, 1983, pp. 411-452.

[6]Burgmann, S., Brucker, S. and Schroder, W., "Scanning PIV Measurements of a Laminar Separation Bubble", *J. Exp. Fluids*, Vol. 41, 2006, pp. 319-326.

[7]Carpenter, M.H., Kennedy, C.A., Bijl, H., Viken, S.A. and Vatsa, V.N., "Fourth-Order Runge-Kutta Schemes for Fluid Mechanics Applications", *Journal of Scientific Computing*, Vol. 25, No. 1-2, 2005, pp. 157-194.

[8]Carton de Wiart, C., and Hillewaert, K., "DNS and ILES of transitional flows around a SD7003 using a high order Discontinuous Galerkin Method", *7th Int. Conf. on Comp. Fluid Dyn.*, ICCFD7-3604, 2012, July, 2012, Big Island, HI.

[9]Castonguay, P., Liang, C., and Jameson, A., "Simulation of Transitional Flow over Airfoils using the Spectral Difference Method", *40th Fluid Dynamics Conference*, 2010-4626, June, 2010, Chicago, IL.

[10]Cash, J., "The integration of Stiff initial value problems in ODEs using modified extended backward differentiation formulae", *Computers and Mathematics with Applications*, Vol. 9, No. 5, 1983, pp. 645-657.

[11]Dahlquist, D., "A special stability problem for linear multistep methods", *BIT*, Vol. 3, 1963, pp. 27-43.

[12]Del Rey Fernández, D.C., and Zingg, D.W., "High-Order Compact-Stencil Summation-By-Parts Operators for the Second Derivative with Variable Coefficients", *7th Int. Conf. on Comp. Fluid Dyn.*, ICCFD7-2803, 2012, July, 2012, Big Island, HI.

[13]Diener, P., Dorband, E.N., Schnetter, E. and Tiglio, M., "Optimized high-order derivative and dissipation operators satisfying summation by parts, and application in three-dimensional multi- block evolutions", *Journal of Scientific Computing*, Vol. 32, No. 1, 2007, pp. 109-145.

[14]Drikakis, D., Fureby, C., Grinstein, F.F. and Youngs, D., "Simulation of transition and turbulence decay in the Taylor-Green vortex", *Journal of Turbulence*, Vol. 8, No. 20, 2007, pp.1-12.

[15]Galbraith, M. and Visbal, M., "Implicit Large Eddy Simulation of Low Reynolds Number Flow Past the SD7003 Airfoil", *46th AIAA Aero. Sci. Meet.*, 2008-225, Jan, 2008, Reno, NV.

[16]Galbraith, M., "Implicit Large Eddy Simulation of Low Reynolds Number Flow Past the SD7003 Airfoil", M.Sc. Thesis, Department of Aerospace Engineering and Engineering Mechanics, University of Cincinnati, Cincinnati, OH, 2009.

[17]Garmann, D.J., and Visbal, M.R., "High-order solutions of transitional flow over the SD7003 airfoil using compact finite-differencing and filtering", *AIAA 1st International Workshop on High-Order CFD Methods*, 2012, `http://zjw.public.iastate.edu/hiocfd/abstracts/C33_GarmannVisbal.pdf` [retrieved 19 December 2012].

[18]Gordnier, R., Visbal, M. and Galbraith, M., "High-Fidelity Computations of Moving and Flexible Wing Sections with Application to Micro Air Vehicles", *DoD HPCMP User Group Conference*, DOI 10.1109, 2008.

[19]Hicken, J.E. and Zingg, D.W., "A parallel Newton-Krylov solver for the Euler equations discretized using simultaneous approximation terms", *AIAA Journal*, Vol. 46, No. 11, 2008, pp. 2773-2786.

[20]Hicken, J.E. and Zingg, D.W., "Aerodynamic Optimization Algorithm with Integrated Geometry Parameterization and Mesh Movement", *AIAA Journal*, Vol. 48, No. 2, 2010, pp. 400-413.

[21]Hicken, J.E. and Zingg, D.W., "Induced-Drag Minimization of Nonplanar Geometries Based on the Euler Equations", *AIAA Journal*, Vol. 48, No. 11, 2010, pp. 2564-2575.

[22]Hicken, J.E. and Zingg, D.W., "Summation-By-Parts Operators and High-Order Quadrature", *Journal of Computational and Applied Mathematics*, Vol. 237, 2013, pp. 111-125.

[23]Hu, H. and Yang, Z., "An Experimental Study of the Laminar Flow Separation on a Low-Reynolds-Number Airfoil", *J. Fluids Eng.*, 051101, Vol. 130, 2008.

[24]Isono, S. and Zingg, D.W., "A Runge-Kutta-Newton-Krylov Algorithm for Fourth-Order Implicit Time Marching Applied to Unsteady Flows", *42nd AIAA Aerospace Sciences Meeting*, 2004-0433, Aug, 2004, Reno, NV.

[25]Jones, L., Sandberg, R. and Sandham, N., "Direct numerical simulations of force and unforced separation bubbles on an airfoil at incidence", *J. Fluid Mech.*, Vol. 602, 2008, pp. 175-207.

[26]Jones, L., Sandberg, R. and Sandham, N., "Stability and receptivity characteristics of a laminar separation bubble on an aerofoil", *J. Fluid Mech.*, Vol. 602, 2008, pp. 175-207.

[27]Kennedy, C. and Carpenter, M., "Additive Runge-Kutta Schemes for Convection-Diffusion-Reaction Equations", *J. App. Num. Math.*, Vol. 44, No. 1-2, 2003, pp. 139-181.

[28]Kim, J. and Lee, D., "Generalized Characteristic Boundary Conditions for Computational Aeroacoustics", *AIAA Journal*, Vol. 38, No. 11, 2000, pp. 2040-2049.

[29]Kamakoti, R. and Pantano, C., "High-order narrow stencil finite-difference approximations of second-derivatives involving variable coefficients", *SIAM Journal on Scientific Computing*, Vol. 31, No. 6, 2009, pp. 4222-4243.

[30]Kreiss, H.-O. and Scherer, G., "Finite element and finite difference methods for hyperbolic partial differen- tial equations", In Carl de Boor, editor, *Mathematical Aspects of Finite Elements in Partial Differential Equations*, Mathematics Research Center, the University of Wisconsin, Academic Press, 1974.

[31]Leung, T.M. and Zingg, D.W., "Aerodynamic Shape Optimization of Wings Using a Parallel Newton-Krylov Approach", *AIAA Journal*, Vol. 50, No. 3, 2012, pp. 540-550.

[32]Loken, C., Gruner, D., Groer, L., Peltier, R., Bunn, N., Craig, M., Henriques, T., Dempsey, J., Yu, C.-H., Chen, J., Dursi, L.J., Chong, J., Northrup, S., Pinto, J., Knecht, N. and Van Zon, R., "SciNet: Lessons Learned from Building a Power-efficient Top-20 System and Data Centre", *Journal of Physics - Conference Series*, Vol. 256, No. 1, 2010, 012026.

[33]Mattsson, K., "Summation by parts operators for finite-difference approximations of second derivatives with variable coefficients", *Journal of Scientific Computing*, Vol. 51, No. 3, 2012, pp 650-682.

[34]Mattsson, K. and Carpenter, M.H., "Stable and Accurate Interpolation Operators for High-Order Multiblock Finite Difference Methods", *SIAM Journal on Scientific Computing*, Vol. 32, No. 4, 2010, pp. 2298-2320.

[35]Mattsson, K. and Nordström, J., "Summation by parts operators for finite-difference approximations of second derivatives", *Journal of Computational Physics*, Vol. 199, No. 2, 2004, pp. 503-540.

[36]Mattsson, K., Svärd, M. and Shoeybi, M., "Stable and accurate schemes for the compressible Navier-Stokes equations", *Journal of Computational Physics*, Vol. 227, No. 4, 2008, pp 2293-2316.

[37]Nordström, J., Gong, J., van der Weide, E. and Svärd, M., "A stable and conservative high order multi-block method for the compressible Navier-Stokes equations", *Journal of Computational Physics*, Vol. 228, No. 24, 2009, pp. 9020-9035.

[38]Ol, M., McAuliffe, B., Hanff, E., Scholz, U., and Káhler, C., "Comparison of Laminar Separation Bubble Measurements on a Low Reynolds Number Airfoil in Three Facilities", *AIAA Paper*, 2005-5149, June, 2005.

[39]Osusky, M., Hicken, J. and Zingg, D.W., "A parallel Newton-Krylov-Schur flow solver for the Navier-Stokes equations using the SBP-SAT approach", *48th AIAA Aerospace Sciences Meeting*, 2010-116, Jan, 2010, Orlando, FL.

[40]Osusky, M., Boom, P.D., Del Rey Fernández, D.C., and Zingg, D.W., "An efficient Newton-Krylov-Schur parallel solution algorithm for the steady and unsteady Navier-Stokes equations", *7th Int. Conf. on Comp. Fluid Dyn.*, ICCFD7-1801, 2012, July, 2012, Big Island, HI.

[41]Radespiel, R., Windte, J. and Scholz, U., "Numerical and experimental flow analysis of moving airfoils with laminar separation bubbles", *44th AIAA Aerospace Sciences Meeting*, 2006-0501, Jan, 2006, Reno, NV.

[42]Saad, Y. and Sosonkina, M., "Distributed Schur complement techniques for general sparse linear systems", *SIAM Journal of Scientific Computing*, Vol. 21, No. 4, 1999, pp. 1337-1357.

[43]Sandberg, R. and Sandham, N., "Nonreflecting Zonal Characteristic Boundary Condition for Direct Numerical Simulation of Aerodynamic Sound", *AIAA Journal*, Vol. 44, No. 2, 2006, pp. 402-405.

[44]Sandberg, R.D., Jones, L.E., Sandham, N.D. and Joseph, P.F., "Direct Numerical Simulation of Tonal Noise Generated by Laminar Flow Past Airfoils", *J. of Sound and Vib.*, Vol. 320, 2009, pp. 838-858.

[45]Strand, B., "Summation by parts for finite difference approximations for d/dx", *Journal of Computational Physics*, Vol. 110, No. 1, 1994, pp. 47-67.

[46]Svärd, M., Carpenter, M.H., and Nordström, J., "A stable high-order finite difference scheme for the compressible Navier-Stokes equations, far-field boundary conditions", *Journal of Computational Physics*, Vol. 225, No. 1, 2007, pp. 1020-1038.

[47]Svärd, M. and Nordström, J., "A stable high-order finite difference scheme for the compressible Navier-Stokes equations, no-slip wall boundary conditions", *Journal of Computational Physics*, Vol. 227, No. 10, 2008, pp 4805-4824.

[48]Tabesh, M. and Zingg, D.W., "Efficient Implicit Time-Marching Methods Using a Newton-Krylov Algorithm", *47th AIAA Aerospace Sciences Meeting*, 2009-164, Jan, 2009, Orlando, FL.

[49]Taylor, G.I. and Green, A.E., "Mechanism of the Production of Small Eddies from Large Ones", *Proc. of the Royal Soc. A*, Vol. 158, No. 3, 1937, pp. 499-521.

[50]Uranga, A., Persson, P.-O., Drela, M. and Peraire, J., "Implicit Large Eddy Simulation of Transitional Flow Over Airfoils and Wings", *19th AIAA Comp Fluid Dyn. Conf.*, June, 2009, San Antonio, TX.

[51]van Rees, W.M., Leonard, A., Pullin, D.I. and Koumoutsakos, P., "A Comparison of Vortex and Pseudo-Spectral Methods for the Simulation of Periodic Vortical Flows at High Reynolds Number", *J. of Comp. Phys.*, Vol. 230, No. 8, 2011, pp. 2794-2805.

[52]Vatsa, V., Carpenter, M. and Lockard, D., "Re-evaluation of an Optimized Second Order Backward Difference (BDF2OPT) Scheme for Unsteady Flow Applications", *48th AIAA Aerospace Sciences Meeting*, 2010-0122, Jan, 2010, Orlando, FL.

[53]Visbal, M., "High-Fidelity Simulation of Transitional Flows past a Plunging Airfoil", *AIAA Journal*, Vol. 47, No. 11, 2009, pp. 2685-2697.

[54]Visbal, M., Gordnier, R. and Galbraith, M., "High-fidelity simulation of moving and flexible airfoils at low Reynolds numbers", *Animal Locomotion*, 2010, pp. 341-360.

[55]Yee, H., Sandham, N. and Djomehri, M., "Low dissipative high order shock-capturing methods using characteristic-based filters", *RIACS NASA Ames Research Center*, NASA/CR-1998, 208350, 1998.

[56]Zhou, Y. and Wang, Z.J., "Implicit Large Eddy Simulation of Transitional Flow over a SD7003 Wing Using High-order Spectral Difference Method", *40th AIAA Fluid Dyn. Conf. and Ex.*, June, 2010, Chicago, IL.

[57]Zimmerman, B., and Wang, Z.J., "Transitional Flow Over an SD7003 Wing", *AIAA 1st International Workshop on High-Order CFD Methods*, 2012, `http://zjw.public.iastate.edu/hiocfd/abstracts/C33_ISU.pdf` [retrieved 19 December 2012].