Journal of Computational Physics 228 (2009) 3917–3935

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

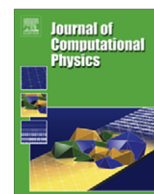

Preconditioning methods for discontinuous Galerkin solutions of the Navier–Stokes equations

Laslo T. Diosady *, David L. Darmofal

Aerospace Computational Design Laboratory, Massachusetts Institute of Technology, 77 Massachusetts Ave. 37-401, Cambridge MA 02139, United States

A R T I C L E   I N F O

Article history:
Received 2 July 2008
Received in revised form 1 December 2008
Accepted 3 February 2009
Available online 9 March 2009

Keywords:
Discontinuous Galerkin
Implicit solvers
GMRES
ILU factorization
Multigrid
In-place factorization

A B S T R A C T

A Newton–Krylov method is developed for the solution of the steady compressible Navier–Stokes equations using a discontinuous Galerkin (DG) discretization on unstructured meshes. Steady-state solutions are obtained using a Newton–Krylov approach where the linear system at each iteration is solved using a restarted GMRES algorithm. Several different preconditioners are examined to achieve fast convergence of the GMRES algorithm. An element Line-Jacobi preconditioner is presented which solves a block-tridiagonal system along lines of maximum coupling in the flow. An incomplete block-LU factorization (Block-ILU(0)) is also presented as a preconditioner, where the factorization is performed using a reordering of elements based upon the lines of maximum coupling. This reordering is shown to be superior to standard reordering techniques (Nested Dissection, One-way Dissection, Quotient Minimum Degree, Reverse Cuthill–Mckee) especially for viscous test cases. The Block-ILU(0) factorization is performed in-place and an algorithm is presented for the application of the linearization which reduces both the memory and CPU time over the traditional dual matrix storage format. Additionally, a linear *p*-multigrid preconditioner is also considered, where Block-Jacobi, Line-Jacobi and Block-ILU(0) are used as smoothers. The linear multigrid preconditioner is shown to significantly improve convergence in term of number of iterations and CPU time compared to a single-level Block-Jacobi or Line-Jacobi preconditioner. Similarly the linear multigrid preconditioner with Block-ILU smoothing is shown to reduce the number of linear iterations to achieve convergence over a single-level Block-ILU(0) preconditioner, though no appreciable improvement in CPU time is shown.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Discontinuous Galerkin (DG) discretizations have become increasingly popular for achieving accurate solutions of conservation laws. Specifically, DG discretizations have been widely used to solve the Euler and Navier–Stokes equations for convection-dominated problems [6–8,13,14,5]. DG methods are attractive since the elementwise discontinuous representation of the solution provides a natural way of achieving higher-order accuracy on arbitrary, unstructured meshes. A detailed overview of DG methods for the discretization of the Euler and Navier–Stokes equations is provided by Cockburn and Shu [14]. They, among others [21,30], have noted that while DG discretizations have been extensively studied, development of solution methods ideally suited for solving these discretizations have lagged behind.

The use of *p*-multigrid for the solution of a DG discretization of a two-dimensional convection problem was presented in [22]. Fidkowski [19] and Fidkowski et al. [21] first used a multigrid strategy to solve DG discretizations of compressible flows. They used a *p*-multigrid scheme with an element-line smoother to solve the non-linear system of equations. Recently,

* Corresponding author.
E-mail addresses: diosady@mit.edu (L.T. Diosady), darmofal@mit.edu (D.L. Darmofal).

0021-9991/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2009.02.035

several other authors have used p-multigrid methods to solve DG discretizations of the Euler or Navier–Stokes equations [23,30,29,26]. Nastase and Mavriplis [30,29] used both p-multigrid (where coarse solutions are formed by taking lower order approximations within each element), and hp-multigrid, where an h-multigrid scheme was used to provide a solution update for the $p = 0$ approximation. Nastase and Mavriplis used this hp-multigrid scheme with an element Block-Jacobi smoother to solve the non-linear system as well as to solve the linear system arising from a Newton scheme for the compressible Euler equations.

The Newton-GMRES approach has been widely used for finite volume discretizations of the Euler and Navier–Stokes equations [1,12,11,39,27,25,31]. In the context of DG discretizations, GMRES was first used to solve the steady 2D compressible Navier–Stokes equations by Bassi and Rebay [8,9]. GMRES has also been used for the solution of the linear system arising at each iteration of an implicit time stepping scheme for the DG discretization of the time dependent Euler or Navier–Stokes equations [40,17,36,38]. Persson and Peraire [36,38] developed a two level scheme as a preconditioner to GMRES to solve the linear system at each step of an implicit time stepping scheme. They used an ILU(0) smoother for the desired p and solved a coarse grid problem ($p = 0$ or $p = 1$) exactly.

Much of the work in the development of solvers for DG discretizations has built upon ideas developed for finite difference or finite volume discretizations. While solution methods developed for finite difference or finite volume discretizations may be adapted to solve DG discretizations, Persson and Peraire [38] noted that the matrix structure arising from DG discretizations has a block structure which may be exploited to develop a more efficient solver. This work examines several preconditioners which take advantage of the block structure of the Jacobian matrix for the solution of the steady-state Euler and Navier–Stokes equations. While results presented here are used to solve steady-state problems, the methods are also suitable for solving time dependent problems.

This paper is a completion of work originally presented in [15]. Section 2 provides an overview of the DG discretization and the Newton–Krylov approach for solving systems of non-linear conservation laws. Section 3 presents the Block-Jacobi, Line-Jacobi and Block-ILU(0) stationary iterative methods that are used as single-level preconditioners or as smoothers on each level of the linear multigrid preconditioner. By considering the Block-ILU preconditioner as a stationary iterative method, a memory efficient implementation is developed which requires no additional storage for the incomplete factorization, while reducing the total time required per linear iteration compared to the traditional dual matrix storage format. Section 4 presents a new matrix reordering algorithm for the Block-ILU factorization based upon lines of maximum coupling between elements in the flow. This line reordering algorithm is shown to significantly improve the convergence behaviour, especially for viscous problems. Section 5 presents the linear multigrid algorithm and discusses memory considerations involved in the development of a memory efficient preconditioner. Finally, Section 6 presents numerical results comparing the convergence of the different preconditioning algorithms.

2. Solution method

2.1. DG discretization

The time dependent, compressible Navier–Stokes equations using index notation are given by:

$$\partial_t u_k + \partial_i F_{ki}(\mathbf{u}) - \partial_i F^v_{ki}(\mathbf{u}) = 0, \quad k \in [1, n_s], \tag{1}$$

where u_k is the kth component of the conservative state vector $\mathbf{u} = [\rho, \rho v_i, \rho E]$, ρ is the density, v_i are the components of the velocity, and E is the total energy. The size of the conservative state vector n_s, is 4 and 5, for two- and three-dimensional flows, respectively (assuming turbulence modeling or other equations are not included). $F_{ki}(\mathbf{u})$ and $F^v_{ki}(\mathbf{u})$ are inviscid and viscous flux components, respectively, such that Eq. (1) is a compact notation for the conservation of mass, momentum, and energy.

The DG discretization of the Navier–Stokes equations is obtained by choosing a triangulation T_h of the computational domain Ω composed of triangular elements κ, and obtaining a solution in $\mathcal{V}^p_h$, the space of piecewise polynomials of order p, which satisfies the weak form of the equation. We define $\mathbf{u}_h$ to be the approximate solution in $(\mathcal{V}^p_h)^{n_s}$, while $\mathbf{v}_h \in (\mathcal{V}^p_h)^{n_s}$ is an arbitrary test function. The weak form is obtained by multiplying Eq. (1) by the test functions and integrating over all elements. The weak form is given by

$$\sum_{\kappa \in T_h} \int_\kappa v_k \partial_t u_k dx + \mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_h) = 0, \tag{2}$$

where

$$\mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_h) = \sum_{\kappa \in T_h} [\mathbb{E}_\kappa(\mathbf{u}_h, \mathbf{v}_h) + \mathbb{V}_\kappa(\mathbf{u}_h, \mathbf{v}_h)], \tag{3}$$

$$\mathbb{E}_\kappa(\mathbf{u}_h, \mathbf{v}_h) = -\int_\kappa \partial_i v_k F_{ki} dx + \int_{\partial\kappa} v^+_k \widehat{F}_{ki}(\mathbf{u}^+_h, \mathbf{u}^-_h)\hat{n}_i ds \tag{4}$$

and $\mathbb{V}_\kappa(\mathbf{u}_h, \mathbf{v}_h)$ is the discretization of the viscous terms. In Eq. (4), $()^+$ and $()^-$ denote values taken from the inside and outside faces of an element, while $\hat{n}$ is the outward-pointing unit normal. $\widehat{F}_{ki}(\mathbf{u}^+_h, \mathbf{u}^-_h)\hat{n}_i$ is the Roe numerical flux function

approximating $F_{ki}\hat{n}_i$ on the element boundary faces [41]. The viscous terms, $\mathbb{V}_\kappa(\mathbf{u}_h, \mathbf{v}_h)$ are discretized using the BR2 scheme of Bassi and Rebay [8]. The BR2 scheme is used because it achieves optimal order of accuracy while maintaining a compact stencil with only nearest neighbour coupling. Further details of the discretization of the viscous terms may be found in Fidkowski et al. [21].

The discrete form of the equations is obtained by choosing a basis for the space $\mathcal{V}_h^p$. The solution vector $\mathbf{u}_h(x,t)$ may then be expressed as a linear combination of basis functions $\mathbf{v}_{h_i}(x)$ where the coefficients of expansion are given by the discrete solution vector $\mathbf{U}_h(t)$, such that:

$$\mathbf{u}_h(x,t) = \sum_i \mathbf{U}_{h_i}(t)\mathbf{v}_{h_i}(x). \tag{5}$$

Two sets of basis functions are used in the context of this work: a nodal Lagrange basis and a hierarchical basis. Further details of the bases may be found in Fidkowski et al. [21].

Given a basis for the space $\mathcal{V}_h^p$, the weak form of the Navier–Stokes equations given in Eq. (2) can be written in semi-discrete form as:

$$\mathcal{M}_h \frac{d\mathbf{U}_h}{dt} + R_h(\mathbf{U}_h(t)) = 0, \tag{6}$$

where R_h is the discrete non-linear residual such that $R_h(\mathbf{U}_h)_i = \mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_{h_i})$, while $\mathcal{M}_h$ is the mass matrix given by

$$\mathcal{M}_{h_{ij}} = \int_\kappa \mathbf{v}_{h_i}\mathbf{v}_{h_j} dx. \tag{7}$$

Since the basis functions are piecewise polynomials which are non-zero only within a single element, the mass matrix is block-diagonal.

To discretize Eq. (6) in time, we introduce a time integration scheme given by:

$$\mathbf{U}_h^{m+1} = \mathbf{U}_h^m - \left(\frac{1}{\Delta t}\mathcal{M}_h + \frac{\partial R_h}{\partial \mathbf{U}_h}\right)^{-1} R_h(\mathbf{U}_h^m). \tag{8}$$

A steady-state solution of the Navier–Stokes equations is given by $\mathbf{U}_h$ satisfying:

$$R_h(\mathbf{U}_h) = 0. \tag{9}$$

The steady-state solution is obtained by using the time integration scheme given in Eq. (8) and increasing the time step Δt, such that $\Delta t \to \infty$. Directly setting $\Delta t = \infty$ is the equivalent of using Newton's method to solve Eq. (9), however convergence is unlikely if the initial guess is far from the solution. On the other hand, if the solution is updated using Eq. (8), then the intermediate solutions approximate physical states in the time evolution of the flow, and convergence is more likely.

2.2. Linear system

The time integration scheme given by Eq. (8) requires the solution of a large system of linear equations of the form $A\mathbf{x} = \mathbf{b}$ at each time step, where

$$A = \frac{1}{\Delta t}\mathcal{M}_h + \frac{\partial R_h}{\partial \mathbf{U}_h}, \quad \mathbf{x} = \Delta \mathbf{U}_h^m, \quad \mathbf{b} = -R_h(\mathbf{U}_h^m). \tag{10}$$

The matrix A is commonly refered to as the Jacobian matrix. Since the Jacobian matrix is derived from the DG discretization, the Jacobian matrix has a block-sparse structure with N_e block rows of size n_b, where N_e is the number of elements in the triangulation T_h, while n_b is the number of unknowns for each element. Here $n_b = n_s \times n_m$, where n_m is the number of modes per state. n_m is a function of the solution order p and the spatial dimension, as summarized in Table 1. Each block row of the Jacobian matrix has a non-zero diagonal block, corresponding to the coupling of states within each element, and n_f off-diagonal non-zero blocks corresponding to the coupling of states between neighbouring elements, where n_f is the number of faces per element (3 and 4 for 2D triangular and 3D tetrahedral elements, respectively). When the time step, Δt, is small, the Jacobian matrix is block-diagonally dominant and the linear system is relatively easy to solve iteratively. On the other hand as the time step increases the coupling between neighbouring elements becomes increasingly important and the linear system generally becomes more difficult to solve.

Table 1
Number of modes per element, n_m, as a function of solution order, p.

p	0	1	2	3	4	p
n_m, 2D	1	3	6	10	15	$\frac{(p+1)(p+2)}{2}$
n_m, 3D	1	4	10	20	35	$\frac{(p+1)(p+2)(p+3)}{6}$

2.3. Linear solution method

The block-sparse structure of the Jacobian matrix and the large number of unknowns suggest the use of an iterative method, more specifically a Krylov-subspace method, to solve the linear system. Since the Jacobian matrix is non-symmetric (though structurally symmetric), the method of choice is the restarted GMRES [43,42] algorithm which finds an approximate solution, $\tilde{\mathbf{x}}$, in the Krylov subspace, $\mathcal{K} = \{\mathbf{b}, A\mathbf{b}, A^2\mathbf{b}, \ldots, A^{n-1}\mathbf{b}\}$, that minimizes the $L-2$ norm of the linear residual $\mathbf{r} = \mathbf{b} - A\tilde{\mathbf{x}}$.

The convergence of the GMRES algorithm has been shown to be strongly dependent upon eigenvalues of the Jacobian matrix, A [43,42,44]. To improve the convergence properties of GMRES, a preconditioner is used which transforms the linear system $A\mathbf{x} = \mathbf{b}$ into a related system with better convergence properties. In this work, only left preconditioning is used, where the linear system is multiplied on the left by a preconditioner P^{-1}, resulting in the linear system: $P^{-1}A\mathbf{x} = P^{-1}\mathbf{b}$. Though the preconditioner, P, is presented as a matrix, any iterative method may be used as a preconditioner.

2.4. Residual tolerance criterion

When solving the DG discretization of the steady-state Navier–Stokes equations using the time stepping scheme presented in Eq. (8), it is often unnecessary to solve the linear system of equations exactly at each iteration. When the time step is small, or the solution estimate is far from the exact solution, the linear system only needs to be solved to a limited tolerance, which depends upon the non-linear residual. Kelley and Keyes [24] considered three phases of a time stepping scheme to solve the steady-state Euler equations: the initial, midrange, and terminal phases. Kelley and Keyes proved super-linear convergence of the non-linear residual in the terminal phase of an inexact Newton iteration given sufficient reduction of the linear residual in each iteration. In this section, an exit criterion is developed for the solution of the linear system to realize the super-linear convergence during the terminal phase. To develop this exit criterion, we consider the convergence of Newton's method to solve Eq. (9), such that the solution update is given by:

$$\mathbf{U}_h^{m+1} = \mathbf{U}_h^m - \left(\frac{\partial R_h}{\partial \mathbf{U}_h}\right)^{-1} R_h(\mathbf{U}_h^m), \tag{11}$$

where $\mathbf{U}_h^m$ is the approximate solution at iteration m of Newton's method. Defining $\epsilon_h^m = \mathbf{U}_h - \mathbf{U}_h^m$ to be the solution error at iteration m, quadratic convergence of the error can be proven as $\epsilon_h^m \to 0$. Namely,

$$\|\epsilon_h^{m+1}\| = C_1 \|\epsilon_h^m\|^2 \tag{12}$$

for some constant C_1 [24]. Similarly quadratic convergence of the solution residual is observed,

$$\left\|R_h(\mathbf{U}_h^{m+1})\right\| = C_2 \|\mathcal{R}_h(\mathbf{U}_h^m)\|^2 \tag{13}$$

for some different constant C_2. Based on this observation, an estimate of the reduction in the solution residual may be given by:

$$\frac{\left\|R_h(\mathbf{U}_h^{m+1})\right\|}{\|R_h(\mathbf{U}_h^m)\|} \sim \left(\frac{\|R_h(\mathbf{U}_h^m)\|}{\left\|R_h(\mathbf{U}_h^{m-1})\right\|}\right)^2 = (d^m)^2, \tag{14}$$

where $d^m = \frac{\|R_h(\mathbf{U}_h^m)\|}{\left\|R_h(\mathbf{U}_h^{m-1})\right\|}$, is the decrease factor of the non-linear residual at iteration m. When the expected decrease of the non-linear residual is small, it may not be necessary to solve the linear system at each Newton step exactly to get an adequate solution update. It is proposed that the linear system given by $A_h\mathbf{x}_h = \mathbf{b}_h$ should have a reduction in linear residual proportional to the expected decrease in the non-linear residual. Defining the linear residual at linear iteration k to be $\mathbf{r}_h^k = \mathbf{b}_h - A_h\mathbf{x}_h^k$, the linear system is solved to a tolerance of:

$$\frac{\|\mathbf{r}_h^n\|}{\|\mathbf{r}_h^0\|} \leqslant K(d^m)^2, \tag{15}$$

where K is a user defined constant, typically chosen in the range $K = [10^{-3}, 10^{-2}]$. Since left preconditioning is used, the linear residual is not available at each GMRES iteration and computing this linear residual can be computationally expensive. As a result, the preconditioned linear residual norm, $\left\|P^{-1}(\mathbf{b}_h - A_h\mathbf{x}_h^k)\right\|$, is used, which can be computed essentially for free at each GMRES iteration. The reduction in the preconditioned residual also provides an estimate of the reduction of the norm of the linear solution error, $\left\|A_h^{-1}b - \mathbf{x}_h^k\right\|$, since

$$\frac{\left\|(A_h^{-1}\mathbf{b}_h - \mathbf{x}_h^k)\right\|}{\left\|(A_h^{-1}\mathbf{b}_h - \mathbf{x}_h^0)\right\|} = \frac{\left\|\left(A_h^{-1}P\right)\left(P^{-1}(\mathbf{b}_h - A_h\mathbf{x}_h^k)\right)\right\|}{\left\|\left(A_h^{-1}P\right)\left(P^{-1}(\mathbf{b}_h - A_h\mathbf{x}_h^0)\right)\right\|} \leqslant \kappa\left(P^{-1}A_h\right)\frac{\left\|P^{-1}(\mathbf{b}_h - A_h\mathbf{x}_h^k)\right\|}{\left\|P^{-1}(\mathbf{b}_h - A_h\mathbf{x}_h^0)\right\|}, \tag{16}$$

where $\kappa(P^{-1}A_h)$ is the condition number of $P^{-1}A_h$. With increasingly effective preconditioning, $P^{-1}A_h$ approaches the identity matrix and the reduction in the preconditioner residual norm more closely approximates the reduction in the linear solution error.

Since the non-linear residual may increase at some iteration m, the tolerance for the linear system presented in Eq. (15) is modified to be:

$$\frac{\left\|P^{-1}\mathbf{r}_h^n\right\|}{\left\|P^{-1}\mathbf{r}_h^0\right\|} \leqslant K(\min\{1, d^m\})^2. \tag{17}$$

This criterion for the reduction of the linear residual is then used to determine n, the number of GMRES iterations to perform each Newton step.

3. In-place preconditioning

3.1. Stationary iterative methods

Stationary iterative methods used to solve the system of linear equations $A\mathbf{x} = \mathbf{b}$ involve splitting the matrix A into two parts such that $A = M + N$, where M in some sense approximates the matrix A and is relatively easy to invert. Since an iterative scheme is typically used directly as a preconditioner to GMRES, M is commonly referred to as the preconditioning matrix. Applying a stationary iterative method, $\mathbf{x}$ is updated using

$$\mathbf{x}^{k+1} = (1-\omega)\mathbf{x}^k + \omega M^{-1}(\mathbf{b} - N\mathbf{x}^k), \tag{18}$$

where ω is the under relaxation factor. An equivalent form of Eq. (18) is

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \omega M^{-1}\mathbf{r}^k, \tag{19}$$

where $\mathbf{r}^k$ is the linear residual given by

$$\mathbf{r}^k = \mathbf{b} - A\mathbf{x}^k. \tag{20}$$

In practice, stationary iterative methods involve a preprocessing stage and an iterative stage. The iterative stage involves repeated solution updates according to Eq. (18) or Eq. (19), where Eq. (18) is used if the application of N is computationally less expensive than the application of A, otherwise Eq. (19) is used. In addition, if the stationary iterative method is used as a smoother for linear multigrid, then the iterative stage will involve repeated calculation of the linear residual, $\mathbf{r}$, using Eq. (20). In the preprocessing stage the matrix A is factorized such that the application of M^{-1}, M, N and A in Eqs. ()()()(18)–(20) may be evaluated at a fraction of the computational cost of the preprocessing stage. In our implementation, the preprocessing stage is performed in place such that the original matrix A is rewritten with a factorization F. As a result the iterative method uses only the memory required to store the original matrix A, with no additional memory storage required for M, M^{-1} or N.

3.2. Block-Jacobi solver

The first and most basic stationary iterative method used in this work is a Block-Jacobi solver. The Block-Jacobi solver is given by choosing M to be the block-diagonal of the matrix A. In the preprocessing stage each diagonal block is LU factorized and the factorization, F, is stored, where

$$F = \begin{bmatrix} LU(A_{11}) & A_{12} & A_{13} \\ A_{21} & LU(A_{22}) & A_{23} \\ A_{31} & A_{32} & LU(A_{33}) \end{bmatrix}. \tag{21}$$

This factorization allows for the easy application of both M and M^{-1} during the iterative stage. N is given by the off-diagonal blocks of A which are not modified in the preprocessing stage. Table 2 gives the asymptotic operation counts per element for forming F (given A), as well as the application of M^{-1}, M, N and A. The operation counts presented in Table 2 are

Table 2
Block-Jacobi solver asymptotic operation count per element.

Operation	Operation count	2D	3D
Form F	$\frac{2}{3}n_b^3$	$\frac{2}{3}n_b^3$	$\frac{2}{3}n_b^3$
$x = M^{-1}x$	$2n_b^2$	$2n_b^2$	$2n_b^2$
$y = Mx$	$2n_b^2$	$2n_b^2$	$2n_b^2$
$y = Nx$	$2n_f n_b^2$	$6n_b^2$	$8n_b^2$
$y = Ax$	$2(n_f+1)n_b^2$	$8n_b^2$	$10n_b^2$

asymptotic estimates, in that lower order terms in n_b have been ignored. The application of A is computed as the sum of the applications of M and N. Thus, the Block-Jacobi iterative step uses Eq. (18), since the application of A is computationally more expensive than the application of N.

3.3. Line-Jacobi solver

The second stationary iterative method presented in this work is a Line-Jacobi solver. The Line-Jacobi solver is given by forming lines of maximum coupling between elements and solving a block-tridiagonal system along each line. The coupling between elements is determined by using a $p = 0$ discretization of the scalar transport equation:

$$\nabla \cdot (\rho u \phi) - \nabla \cdot (\mu \nabla \phi) = 0. \tag{22}$$

The lines are formed by connecting neighbouring elements with maximum coupling. For purely convective flows, the lines are in the direction of streamlines in the flow. For viscous flows solved using anisotropic grids, the lines within the boundary layer are often in non-streamline directions. Further details of the line formation algorithm are presented in the theses of Fidkowski [19] and Oliver [33].

For the Line-Jacobi solver, M is given by the block-tridiagonal systems corresponding to the lines of maximum coupling, while N is given by the blocks associated with the coupling between elements across different lines. In the preprocessing stage, M is factorized using a block-variant of the Thomas algorithm given by:

$$F = \begin{bmatrix} LU(A_{11}) & A_{12} & A_{13} \\ A_{21} & LU(A'_{22}) & A_{23} \\ A_{31} & A_{32} & LU(A'_{33}) \end{bmatrix}, \tag{23}$$

where $A'_{22} = A_{22} - A_{21}A_{11}^{-1}A_{12}$ and $A'_{33} = A_{33} - A_{32}A'^{-1}_{22}A_{23}$. The corresponding LU factorization of M is given by:

$$M = \begin{bmatrix} A_{11} & A_{12} & \\ A_{21} & A_{22} & A_{23} \\ & A_{32} & A_{33} \end{bmatrix} = \begin{bmatrix} I & & \\ A_{21}A_{11}^{-1} & I & \\ & A_{32}A'^{-1}_{22} & I \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} & \\ & A'_{22} & A_{23} \\ & & A'_{33} \end{bmatrix}. \tag{24}$$

The factorization given by Eq. (23) is stored as opposed to the LU factorization given by Eq. (24) to reduce the computational cost of the preprocessing stage. The reduction in computational cost of storing the factorization given by Eq. (23) is offset by an increase in the computational cost of applying M and M^{-1} during the iterative stage. The total computational cost for both the preprocessing and iterative stages using the factorization given by Eq. (23) is lower than the LU factorization given by Eq. (24), as long as the total number of linear iterations is less than the block size, n_b.

Table 3 gives the asymptotic operation counts per element for the preprocessing stage as well as the application of M^{-1}, M, N and A. The application of A is once again computed as a sum of the applications of M and N. As with the Block-Jacobi solver, the solution update for the Line-Jacobi solver is given by Eq. (18), since the application of N is computationally less expensive than the application of A.

3.4. Block-ILU solver

The final iterative method presented in this work is a block incomplete-LU factorization (Block-ILU). ILU factorizations have been successfully used as preconditioners for a variety of aerodynamic problems [1,11,39,27,25,36,31]. Typically the LU factorization of a sparse matrix will have a sparsity pattern with significantly more non-zeros, or fill, than the original matrix. The principle of an incomplete-LU factorization is to produce an approximation of the LU factorization of A, which requires significantly less fill than the exact LU factorization. The incomplete LU factorization, $\tilde{L}\tilde{U}$, is computed by performing Gaussian elimination on A but ignoring values which would result in additional fill. The fill level, k, indicates the distance in the sparsity graph of the neighbours in which coupling may be introduced in the ILU(k) factorization. In the context of this work ILU(0) is used, hence no additional fill outside the sparsity pattern of A is permitted. To simplify the notation, for the remainder of this work we use ILU to denote an ILU(0) factorization unless explicitly stated otherwise.

Table 3
Line-Jacobi solver asymptotic operation count per element.

Operation	Operation count	2D	3D
Form F	$\frac{14}{3}n_b^3$	$\frac{14}{3}n_b^3$	$\frac{14}{3}n_b^3$
$x = M^{-1}x$	$8n_b^2$	$8n_b^2$	$8n_b^2$
$y = Mx$	$8n_b^2$	$8n_b^2$	$8n_b^2$
$y = Nx$	$2(n_f - 2)n_b^2$	$2n_b^2$	$4n_b^2$
$y = Ax$	$2(n_f + 2)n_b^2$	$10n_b^2$	$12n_b^2$

Though incomplete-LU factorizations are widely used, most implementations store both the linearization A and the incomplete factorization $\tilde{L}\tilde{U}$. Since in most aerodynamic applications the majority of the memory is used for the storage of the linearization and its factorization, such duplicate memory storage may limit the size of the problems which may be solved on a given machine [11,28,36]. In this section, an algorithm is developed that performs the incomplete-LU factorization in-place, such that no additional memory is required for the storage of the factorization. This in-place storage format is an enabling feature which allows for the solution of larger and more complex problems on a given machine. Assuming the majority of the memory is used for the storage of the Jacobian matrix and the Krylov vectors, the increase in the size of the problem which may be solved on a given machine is given by $\frac{2+\eta}{1+\eta}$, where η is the ratio of the memory required to store the Krylov vectors to the memory required to store the Jacobian matrix. For a typical range $\eta \in [0.1, 1.0]$, this represents an increase of 50–90% in the size of problem which may be solved.

To develop an ILU implementation where the memory usage is no greater than that required for the Jacobian, we consider the ILU factorization as a stationary iterative method. In the context of stationary iterative methods, M is given by the product $\tilde{L}\tilde{U}$. It can be easily shown that A differs from M only where fill is dropped in the incomplete LU factorization. Correspondingly, N is given by a matrix containing all fill which was ignored in the ILU factorization. Namely, defining the sparsity of the matrix A by:

$$\mathcal{S}(A) \equiv \{(i,j) : A_{ij} \neq 0\}.$$

It is easily shown that:

$$A = M + N = \tilde{L}\tilde{U} + N \quad \mathcal{S}(A) \neq \mathcal{S}(M) \neq \mathcal{S}(N),$$

where

$$A_{ij} = M_{ij} = (\tilde{L}\tilde{U})_{ij} \quad \forall (i,j) \in \mathcal{S}(A),$$
$$M_{ij} + N_{ij} = (\tilde{L}\tilde{U})_{ij} + N_{ij} = 0 \quad \forall (i,j) \notin \mathcal{S}(A).$$

To construct an in-place storage for ILU, note that both A and N may be reconstructed from $\tilde{L}$ and $\tilde{U}$ given the original sparsity pattern of A. Namely, A may be computed by taking the product $\tilde{L}\tilde{U}$ and ignoring those values not within the original sparsity pattern. Similarly N may be computed by taking the values of $-\tilde{L}\tilde{U}$ outside the sparsity pattern of A. Though recomputing A and N in this manner is possible, it is impractical since the computational cost is of the same order as the original ILU factorization and requires additional memory storage. Fortunately, only the application of A or N is required, and these products can be computed efficiently using $\tilde{L}$ and $\tilde{U}$.

The remainder of this section describes the implementation and computational efficiency of the in-place Block-ILU solver. The operation count estimates for the Block-ILU solver is based on the assumption that neighbours of an element do not neighbour one another. This assumption leads to the fact that the upper triangular part of A and $\tilde{U}$ are identical. Persson and Peraire [38] took advantage of this property by developing a solver which stored $\tilde{L}, A$ and the LU factors of the block diagonal of $\tilde{U}$. Where the assumption that neighbours of an element do not neighbour one another failed they simply ignored the connection between those neighbouring element, noting that it is only an incomplete factorization. The Block-ILU(0) solver presented in this work essentially takes advantage of this same property, but only $\tilde{L}$ and $\tilde{U}$ are stored. Additionally, the assumption that neighbours of an element do not neighbour one another is only used for operational count analysis while the actual implementation does not make this assumption.

In the preprocessing stage, the block incomplete-LU factorization of A is performed in-place where A is replaced by the factorization F. An example of one step of the factorization is given below:

$$\begin{bmatrix} A_{11} & & A_{13} & & A_{15} & A_{16} \\ & A_{22} & & & & \\ A_{31} & & A_{33} & & & \\ & & & A_{44} & & \\ A_{51} & & & & A_{55} & \\ A_{61} & & & & & A_{66} \end{bmatrix} \Rightarrow \begin{bmatrix} LU(A_{11}) & & A_{13} & & A_{15} & A_{16} \\ & A_{22} & & & & \\ (A_{31}A_{11}^{-1}) & & A'_{33} & & & \\ & & & A_{44} & & \\ (A_{51}A_{11}^{-1}) & & & & A'_{55} & \\ (A_{61}A_{11}^{-1}) & & & & & A'_{66} \end{bmatrix},$$

where $A'_{33} = A_{33} - A_{31}A_{11}^{-1}A_{13}, A'_{55} = A_{55} - A_{51}A_{11}^{-1}A_{15}$, and $A'_{66} = A_{66} - A_{61}A_{11}^{-1}A_{16}$. Based on the assumption that neighbours of an element do not neighbour one another, only two of the blocks A_{ij}, A_{ik}, and A_{jk} may be non-zero for any $i \neq j \neq k$. This implies that when eliminating row i only elements A_{ji} and $A_{jj}, j \geqslant i$ are modified. In addition, fill is ignored at A_{jk} and A_{kj}, if elements $j, k > i$ both neighbour element i. In the general case where the assumption is violated, A_{jk} and A_{kj} are non-zero, and these terms are modified in the Block-ILU factorization such that: $A'_{jk} = A_{jk} - A_{ji}A_{ii}^{-1}A_{ik}$ and $A'_{kj} = A_{kj} - A_{ki}A_{ii}^{-1}A_{ij}$. The number of non-zero blocks in the matrix N is given by $\sum_{i=1}^{N_e} \tilde{n}_{f_i}(\tilde{n}_{f_i} - 1)$ where, $\tilde{n}_{f_i}$ is the number of larger ordered neighbours of element i. While the number of non-zero blocks is dependent upon the ordering of the elements in the ILU factorization, it is possible to obtain an estimate by assuming an ordering exists where, $\tilde{n}_{f_i} = \left\lceil \frac{i}{N_e} n_f \right\rceil$. The corresponding estimate for the number of non-zero blocks in N is $N_e(n_f^2 - 1)/3$.

In the iterative stage, the application of M^{-1} is performed using backward and forward substitution of $\widetilde{L}$ and $\widetilde{U}$. The application of A is performed by multiplying by those components of $\widetilde{L}$ and $\widetilde{U}$ which would not introduce fill outside the original sparsity pattern of A. Similarly, the application of N may be performed by multiplying by the components of $\widetilde{L}$ and $\widetilde{U}$ which introduce fill outside the original sparsity pattern of A.

The application of A and N is best illustrated with a simple example. Consider the 3×3 matrix A below, and the corresponding ILU factorization, $\widetilde{L}\widetilde{U}$:

$$A = \begin{bmatrix} 4 & 5 & -6 \\ 8 & 3 & 0 \\ -12 & 0 & 26 \end{bmatrix}, \quad \widetilde{L} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ -3 & 0 & 1 \end{bmatrix}, \quad \widetilde{U} = \begin{bmatrix} 4 & 5 & -6 \\ 0 & -7 & 0 \\ 0 & 0 & 8 \end{bmatrix}.$$

The corresponding matrices M, N and F are given by:

$$M = \begin{bmatrix} 4 & 5 & -6 \\ 8 & 3 & -12 \\ -12 & -15 & 26 \end{bmatrix}, \quad N = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 12 \\ 0 & 15 & 0 \end{bmatrix}, \quad F = \begin{bmatrix} 4 & 5 & -6 \\ 2 & -7 & 0 \\ -3 & 0 & 8 \end{bmatrix}$$

The application of A to a vector x, may be performed by multiplying x by those components of $\tilde{L}$ and $\tilde{U}$ which would not introduce fill outside the original sparsity pattern of A. For the sample matrix, fill was ignored in the ILU factorization at (2,3) and (3,2) when eliminating row 1. Hence, for the sample matrix the application of A may be performed as follows:

$$\begin{aligned} y_1 &= \tilde{U}_{11}x_1 + \tilde{U}_{12}x_2 + \tilde{U}_{13}x_3 = 4x_1 + 5x_2 - 6x_3 \\ y_2 &= \tilde{L}_{21}\tilde{U}_{11}x_1 + \tilde{L}_{21}\tilde{U}_{12}x_2 + \cancel{\tilde{L}_{21}\tilde{U}_{13}x_3} + \tilde{U}_{22}x_2 = 2(4x_1) + 2(5x_2) - 7x_2 \\ y_3 &= \tilde{L}_{31}\tilde{U}_{11}x_1 + \cancel{\tilde{L}_{31}\tilde{U}_{12}x_2} + \tilde{L}_{31}\tilde{U}_{13}x_3 + \tilde{U}_{33}x_3 = -3(4x_1) - 3(-6x_3) + 8x_3 \end{aligned}$$

Clearly, the operation count for computing the application of A in this manner is more expensive than simply applying A in the original form. However, it is important to recognize that in the case of block matrices, each of the terms $\widetilde{L}_{ij}$ and $\widetilde{U}_{ij}$ are matrices and x_i's are vectors, and hence the (matrix–vector) multiplications become significantly more expensive than the (vector) additions. Hence, to leading order, the computational cost is given by the number of matrix–vector multiplications. The total number of multiplications may be reduced by recognizing that certain products $(\widetilde{U}_{11}x_1, \widetilde{U}_{12}x_2, \widetilde{U}_{13}x_3)$ are repeated. Taking advantage of the structure of the matrix A, based on the assumption that neighbours of an element do not neighbour one another, it is possible to show that the application of A using $\widetilde{L}\widetilde{U}$ may be performed at a computational cost of $2(\frac{3}{2}n_f + 1)n_b^2 N_e$.

The application of N is performed by multiplying those components of $\widetilde{L}$ and $\widetilde{U}$ which would introduce fill outside the original sparsity pattern of A. For the sample matrix, fill was ignored at (2,3) and (3,2) when eliminating row 1. Hence, the application of N to a vector x may be performed as follows:

$$\begin{aligned} y_1 &= 0, \\ y_2 &= -\widetilde{L}_{21}\widetilde{U}_{13}x_3 = -2(-6x_3) = 12x_3, \\ y_3 &= -\widetilde{L}_{31}\widetilde{U}_{12}x_2 = 3(5x_2) = 15x_2. \end{aligned}$$

Once again, the computational cost is dominated by (matrix–vector) multiplications, and additional efficiency may be attained by recognizing that some products may be repeated. The operation count for the application of N is a function of $\tilde{n}_{f_i}$, the number of larger ordered faces of each element. While the operation count for the application of N is dependent upon the ordering of the elements in the ILU factorization, it is possible to obtain an estimate by assuming an ordering exists where, $\tilde{n}_{f_i} = \left\lceil \frac{i}{N_e} n_f \right\rceil$. The corresponding estimate for the operation count for applying N is given by $2/3(n_f + 4)(n_f - 1)n_b^2 N_e$.

This estimate of the operation count for the application of N tends to overestimate actual operation counts for practical computational grids. A revised estimate for the application of N may be obtained by considering a particular reordering algorithm based on lines of maximum coupling which is presented in Section 4. Using the ordering of the elements based upon lines effectively reduces the number of free faces for all but the first element in each line since at least one of the faces corresponds to a lower ordered neighbour. The revised estimate for the operation count for the application of N may then be obtained by replacing n_f by $n_f - 1$ in the initial estimate given above. Namely, the revised estimate for the operation count is given by: $\frac{2}{3}(n_f + 3)(n_f - 2)n_b^2 N_e$.

Table 4 shows this revised estimate of the operation count for the application of N normalized by the operation count for the application of A using the traditional dual matrix storage format, for both 2D and 3D problems. Table 4 also shows timing results from several sample 2D and 3D problems. For each grid, timing results are presented for $p = 1$ as well as the largest value of p for which the Jacobian matrix could fit into memory on a single machine. For the $p = 1$ cases the actual timing results exceed the revised estimate. However, for large p the actual timing results closely match the revised estimate in 2D, and are bounded by the revised estimate in 3D. The poorer performance for the $p = 1$ cases may be attributed to the effects of lower order terms in n_b, which become significant since the block size for the $p = 1$ solution is relatively small.

Table 5 shows the asymptotic operation count per element for the preprocessing stage and components of the iterative stage for the Block-ILU solver using the in-place storage format. Note that if the Block-ILU factorization $\widetilde{L}\widetilde{U}$ is stored as a sep-

Table 4
Revised timing estimate for application of N for in-place Block-ILU(0) normalized by a Jacobian vector product.

Dim	Type	# Elements	p	Timing
2D	Estimate			**0.50**
	Structured	2432	1	0.78
	Unstructured	7344	1	0.84
	Cut cell	1250	1	0.69
	Structured	2432	4	0.51
	Unstructured	7344	4	0.52
	Cut cell	1250	4	0.46
3D	Estimate			**0.93**
	Structured	1920	1	0.86
	Unstructured	45,417	1	1.02
	Cut cell	2883	1	0.98
	Structured	1920	3	0.77
	Cut cell	2883	3	0.85

Table 5
Block-ILU solver asymptotic operation count per element.

Operation	Operation count	2D	3D
Form F	$2(n_f+1)n_b^3$	$8n_b^3$	$10n_b^3$
$x = M^{-1}x$	$2(n_f+1)n_b^2$	$8n_b^2$	$10n_b^2$
$y = Mx$	$2(n_f+1)n_b^2$	$8n_b^2$	$10n_b^2$
$y = Nx$ (initial estimate)	$\frac{2}{3}(n_f+4)(n_f-1)n_b^2$	$9\frac{1}{3}n_b^2$	$16n_b^2$
$y = Nx$ (revised estimate)	$\frac{2}{3}(n_f+3)(n_f-2)n_b^2$	$4n_b^2$	$9\frac{1}{3}n_b^2$
$y = Ax$	$2(\frac{3}{2}n_f+1)n_b^2$	$11n_b^2$	$14n_b^2$
$y = Ax$ (full storage)	$2(n_f+1)n_b^2$	$8n_b^2$	$10n_b^2$

arate matrix such that the original matrix A is still available, the cost of computing $y = Ax$ is $2(n_f+1)N_e n_b^2$. Based on the operation counts presented in Table 5, a linear iteration in 2D should be performed using Eq. (18) since the application of A is more expensive than the application of N. Based on the initial estimate for the application of N, in 3D it appears as though the cost of applying A is less than applying N and hence a linear iteration should be performed using Eq. (19). However, in practice a linear iteration in 3D is also performed using Eq. (18) since the revised timing estimate for the application of N is less than the application of A.

3.5. Timing performance

In the previous sections, timing estimates were presented in terms of the operations counts for the different components of each solver. In addition, actual timing results were presented to validate the resived estimate for the ILU application of N. Here all three preconditioners are compared using actual timing results obtained based on a sample 2D test grid with 2432 elements using a $p = 4$ discretization. The actual and estimated timing results are presented in Table 6 where the time has been normalized by the cost of a single matrix vector product of the Jacobian matrix. As shown in Table 6 the actual timing results closely match the estimates based on operation counts.

Table 7 gives the asymptotic operation counts for the different solvers presented in this work. As shown in Table 7, the operation count of performing a linear iteration using the in-place storage format is 25% and 5% less than that using the traditional dual matrix storage format for 2D and 3D, respectively. The in-place matrix storage format is superior to the traditional dual matrix storage format since the application of N is computationally less expensive than the application of A. In this case, the dual storage format could be modified to store M and N as opposed to M and A, so that a linear iteration may be performed according to Eq. (18). A linear iteration could then be performed faster using the modified dual matrix storage

Table 6
Solver asymptotic operation count per element normalized by a Jacobian vector product for $p = 4$, 2432 element mesh.

Operation	Block-Jacobi		Line-Jacobi		Block-ILU	
	Estimate	Actual	Estimate	Actual	Estimate	Actual
$x = M^{-1}x$	0.25	0.39	1.00	1.24	1.00	1.16
$y = Nx$	0.75	0.76	0.25	0.28	0.50	0.51
$y = Ax$	1.00	1.14	1.25	1.34	1.38	1.43

Table 7
Linear iteration asymptotic operation count per element (in multiples of n_b^2).

Preconditioner	2D	3D
Block-Jacobi	8	10
Line-Jacobi	10	12
Block-ILU in-place	12	$19\frac{1}{3}$
Block-ILU dual storage	16	20

format than the in-place matrix storage format. However, the modified dual matrix storage format would require computing N in the preprocessing stage, such that the total computational time for both the preprocessing and iterative stages would still be faster using the in-place storage format if fewer than approximately $3n_b$ linear iterations are performed.

3.6. In-place ILU factorization of general matrices

The in-place ILU algorithm developed in this section has been tailored for DG discretizations and may not be generally applicable to sparse matrices arising from other types of discretizations. While the application of A and N may be computed using the ILU factorization for any sparse matrix, the use of an in-place factorization may be unfeasible due to the number of operations required. The number of non-zero blocks in N and correspondingly, the computational cost for the application of N scales with the square of the number of off-diagonal blocks in the stencil of A. Similarly, if the assumption that neighbours of an element do not neighbour one another is removed, the operation count for the application of A using the ILU factorization also scales with the square of the number of off-diagonal blocks in the stencil. The in-place ILU algorithm is feasible for DG discretizations since there is only nearest neighbour coupling, resulting in a stencil with few off-diagonal blocks. On the other hand, discretizations such as high-order finite volume discretizations have much wider stencils, involving 2nd and 3rd order neighbours [5,31], making the in-place ILU factorization algorithm unfeasible.

4. ILU reordering

In the development of an efficient Block-ILU(0) preconditioner for DG discretizations, the ordering of the equations and unknowns in the linear system is critical. Matrix reordering techniques have been widely used to reduce fill in the LU factorization for direct methods used to solve large sparse linear systems [42]. These reordering techniques have also been used with ILU preconditioners of Krylov methods [11,39,10,31]. Benzi et al. [10] performed numerical experiments comparing the effect of different reordering techniques on the convergence of three Krylov-subspace methods used to solve a finite difference discretization of a linear convection–diffusion problem. They showed that reordering the system of equations can both reduce fill for the incomplete factorization, and improve the convergence properties of the iterative method [10]. Blanco and Zingg [11] compared Reverse Cuthill–Mckee, Nested Dissection, and Quotient Minimum Degree reorderings for ILU(k) factorizations of a finite volume discretization of the Euler Equations. They showed that the Reverse Cuthill–Mckee reordering reduced the fill and resulted in faster convergence for ILU(2). Similarly, Pueyo and Zingg [39] used Reverse Cuthill–Mckee reordering to reduce fill and achieve faster convergence for the finite volume discretization of the Navier–Stokes equations. In the context of ILU(0) factorizations, no additional fill is introduced, hence reordering the system of equations effects only the convergence properties of the iterative method. However, Benzi et al. [10] showed that even for ILU(0), reordering the systems of equations can significantly reduce the number of GMRES iterations required to reach convergence. In the context of ILU factorizations for DG discretizations, Persson and Peraire developed a reordering algorithm for the Navier–Stokes equations that performed well over a wide range of Mach and Reynolds numbers [38]. This reordering algorithm was based on minimizing the magnitude of the discarded fill in the ILU(0) factorization.

In this section, we present a new matrix reordering algorithm for the DG discretization of the Navier–Stokes equations based upon lines of maximum coupling within the flow. This ordering algorithm is motivated by the success of line solvers for both finite volume and DG discretizations [27,20]. We note that the lines of maximum coupling will produce an ILU(0) preconditioner in which the magnitude of the dropped fill will often be small because of the weaker coupling in the off-line directions. For Persson and Peraire's minimum discarded fill algorithm, the magnitude of the fill for each block was quantified using a Frobenius norm [38]. In our approach, we first reduce to a scalar $p = 0$ convection–diffusion system and then measure the coupling directly from the resulting matrix. This new reordering algorithm is compared with several standard reordering techniques; Reverse Cuthill–Mckee, Nested-Dissection, One-Way Dissection, Quotient Minimum Degree and the natural ordering produced by the grid generation. The numerical results for the standard matrix reordering algorithms were determined using the PETSc package for numerical linear algebra [2,4,3].

4.1. Line reordering

The lines of maximum coupling described in Section 3.3 may be used to order the elements for ILU preconditioning. Specifically, the elements may be ordered as they are traversed along each line. Such an ordering of elements ensures that the

coupling between elements within a line, captured by the Line-Jacobi preconditioner, is maintained. A line-ordered Block-ILU preconditioner also captures some additional coupling between lines which is ignored by the Line-Jacobi preconditioner. We note that the lines do not produce a unique reordering, since each line may be traversed in either the forward or backward directions or the lines themselves may also be reordered. While a systematic approach may be developed to choose an optimal permutation for the lines, the natural ordering produced by the line creation algorithm is used for the test cases presented. For these test cases, reordering the lines according to the standard reordering techniques (Reverse Cuthill–Mckee, Nested-Dissection, One-Way Dissection and Quotient Minimum Degree) or reversing the direction of the lines from the natural ordering did not significantly impact the convergence rate.

4.2. Numerical results

To investigate the effectiveness of a reordering based upon lines, numerical results are presented for two representative test cases: an inviscid transonic flow and a subsonic viscous flow. The convergence plots are presented in terms of the number of linear iterations since the computational cost of performing the ILU(0) factorization or a single linear iteration is independent of the matrix reordering when using the traditional dual matrix storage format.

The first test case is an Euler solution of the transonic flow over the NACA 0012 airfoil at a freestream Mach number of $M = 0.75$ and angle of attack of $\alpha = 2.0^{\circ}$. The flow is solved using a $p = 4$ discretization on an unstructured mesh with 7344 elements. Fig. 1 shows the convergence plot of the non-linear residual starting from a converged $p = 3$ solution. The fastest convergence is achieved using the reordering based on lines, which requires only 946 linear iterations for a 10 order drop in residual. One-Way Dissection and Reverse Cuthill–Mckee algorithms also perform well requiring only 1418 and 1611 iterations to converge, respectively. On the other hand, Quotient Minimum Degree and Nested Dissection reorderings result in convergence rates which are worse than the natural ordering of the elements. The second test case is a Navier–Stokes solution of the subsonic flow over the NACA0012 airfoil at zero angle of attack with a freestream Mach number of $M = 0.5$ and a Reynolds number of $Re = 1000$. A $p = 4$ solution is obtained on a computational mesh with 2432 elements, where the solution procedure is restarted from a converged $p = 3$ solution. Fig. 2 presents the convergence plot of the non-linear residual versus linear iterations. The reordering based upon lines is superior to all other reorderings; requiring only 341 iterations to converge. The second best method for this test case is the natural ordering of elements which requires 1350 iterations. The natural reordering performs well for this test case since a structured mesh is used (though the solution procedure does not take advantage of the structure), and hence the natural ordering of the elements involves some inherent structure. Among the other reordering algorithms, Reverse Cuthill–Mckee performs best, requiring 1675 iterations, followed by One-Way Dissection, Quotient Minimum Degree and finally Nested Dissection.

Clearly, reordering the elements according to the lines of maximum coupling results in superior convergence for both inviscid and viscous test cases. The advantages of the line reordering algorithm is especially obvious in the viscous case where reordering according to lines results in a convergence rate nearly 5 times faster than the standard matrix reordering

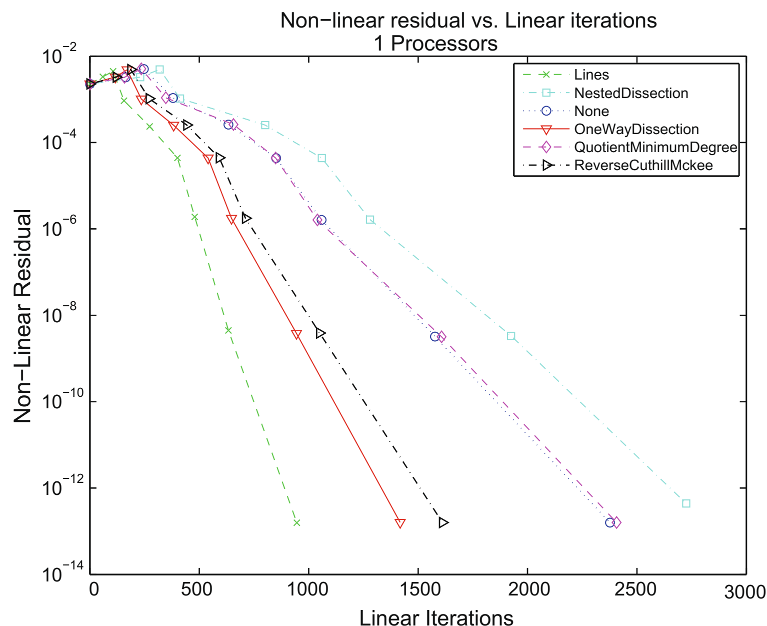

Fig. 1. Non-linear residual vs linear iterations using the Block-ILU(0) preconditioner with different reordering techniques for a transonic Euler solution of the flow about the NACA0012 airfoil (7344 elements, $p = 4$).

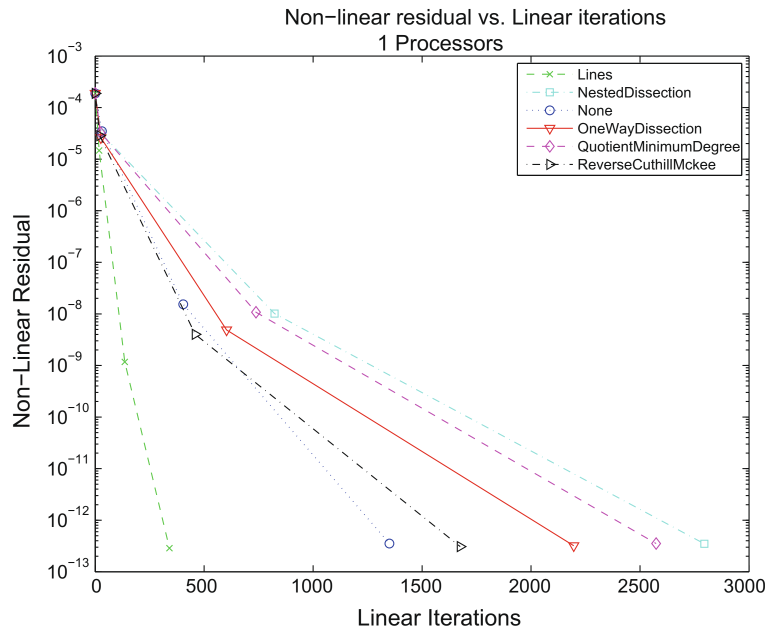

Fig. 2. Non-linear residual vs linear iterations using the Block-ILU(0) preconditioner with different reordering techniques for a Navier–Stokes solution of the flow about the NACA0012 airfoil (2432 elements, $p = 3$).

algorithms available in the PETSc package. Due to the clear success of the line reordering algorithm for these two sample problems, the line reordering method is used for the remainder of the work presented here.

5. Linear multigrid

Multigrid algorithms are used to accelerate the solution of systems of equations arising from the discretization of a PDE-based problem by applying corrections based on a coarser discretization with fewer degrees of freedom. The coarse discretization may involve a computational mesh with fewer elements (h-multigrid) or a lower order solution space (p-multigrid). The DG discretization naturally lends itself to a p-multigrid formulation as a coarser solution space may be easily created by using a lower order polynomial interpolation within each element. Multigrid algorithms may be used to directly solve a non-linear system of equations (non-linear multigrid), or to solve the system of linear equations arising at each step of Newton's method (linear multigrid). This section presents a linear p-multigrid algorithm which is used as a preconditioner to GMRES and makes use of the stationary iterative methods presented in Section 3 as linear smoothers on each multigrid level.

5.1. Linear multigrid algorithm

The basic two-level linear-multigrid algorithm is presented below. While only a two-level system is presented here, in general the multigrid formulation involves multiple solution levels.

- Perform pre-smoothing: $\tilde{\mathbf{x}}_h^k = (1-\omega)\mathbf{x}_h^k + \omega M_h^{-1}(\mathbf{b}_h - N_h\mathbf{x}_h^k)$;
- Compute linear residual: $\tilde{\mathbf{r}}_h^{\tilde{k}} = \mathbf{b}_h - A_h\tilde{\mathbf{x}}_h^k$;
- Restrict linear residual: $\mathbf{b}_H = I_H^h\tilde{\mathbf{r}}_h^k$, where I_H^h is the restriction operator;
- Define coarse level correction: $\mathbf{x}_H^0 = 0$;
- Perform coarse level smoothing: $\mathbf{x}_H^{j+1} = (1-\omega)\mathbf{x}_H^j + \omega M_H^{-1}\left(\mathbf{b}_H - N_H\mathbf{x}_H^j\right)$;
- Prolongate coarse level correction: $\hat{\mathbf{x}}_h^k = \tilde{\mathbf{x}}_h^k + I_h^H\mathbf{x}_H$, where I_h^H is the prolongation operator;
- Perform post-smoothing: $\mathbf{x}_h^{k+1} = (1-\omega)\hat{\mathbf{x}}_h^k + \omega M_h^{-1}(\mathbf{b}_h - N_h\hat{\mathbf{x}}_h^k)$.

As presented in Section 2.1, the solution space for the DG discretization is given by $\mathcal{V}_h^p$, the space of piecewise polynomials of order p spanned by the basis functions $\mathbf{v}_{h_i}$. The corresponding coarse solution space is given by $\mathcal{V}_h^{p-1}$, the space of piecewise polynomials of order $p-1$ spanned by the basis functions $\mathbf{v}_{H_k}$. Since $\mathcal{V}_h^{p-1} \in \mathcal{V}_h^p$, the coarse level basis functions may be expressed as a linear combination of the fine level basis functions:

$$\mathbf{v}_{H_k} = \sum_i \alpha_{ik} \mathbf{v}_{h_i}. \tag{25}$$

The matrix of coefficients α_{ik} form the prolongation operator I_h^H. The coefficients of expansion may also be used to define the restriction operator by considering the restriction of a component of the residual:

$$\mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_{H_k}) = \mathcal{R}_h(\mathbf{u}_h, \sum_i \alpha_{ik} \mathbf{v}_{h_i}) = \sum_i \alpha_{ik} \mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_{h_i}). \tag{26}$$

Hence the restriction operator is given by $I_H^h = (I_h^H)^T$. In our implementation of the linear multigrid algorithm, the coarse grid Jacobian A_H is given by a simple Galerkin projection of the fine grid Jacobian:

$$A_H = I_H^h A_h I_h^H. \tag{27}$$

In this work the linear p-multigrid scheme is used as a preconditioner to GMRES. Multigrid levels are given by each p from the solution order down to $p = 0$. The multigrid preconditioner involves a single V-cycle where one pre- and post- smoothing iteration is used on each multigrid level. On the coarsest multigrid level ($p = 0$) a fixed number (5–10) smoothing iterations are performed. Hence, in general the coarse problem is never solved exactly, however the preconditioner remains fixed at each GMRES iteration.

5.2. Memory considerations

For a linear multigrid preconditioner significant additional memory is required for the storage of the lower order Jacobians on each multigrid level. Table 8 shows the additional memory required for all lower order Jacobians in terms of the fine grid Jacobian for $p = 1 \rightarrow 5$.

Several authors [28,19] have argued that a linear multigrid preconditioner may be unfeasible for large problems due to the additional memory cost of storing these lower order Jacobians. Alternatively, others have advocated for skipping multigrid levels to reduce memory usage. For example, Persson and Peraire [36,38] employed a multi-level scheme where only $p = 0$ and $p = 1$ corrections were applied. Though the linear multigrid method may require significant additional memory for the storage of the lower order Jacobians, faster convergence of the GMRES method is expected and hence fewer Krylov vectors may be required to obtain a converged solution. Hence, to provide a memory equivalent comparison between a single- and multi-level preconditioner, the total memory usage for the Jacobians and Krylov vectors must be considered. In the context of a restarted GMRES algorithm this is equivalent to increasing the GMRES restart value for the single-level preconditioner so that the total memory used by the single and multi-level preconditioners is the same. Table 8 also gives the additional memory for the storage of all lower order Jacobians for the linear multigrid solver in terms of the number of solution vectors on the fine grid. These values may also be viewed as the additional number of GMRES vectors allocated for the single-level preconditioner to provide a memory equivalent comparison with the multigrid preconditioner.

6. Numerical results

The performance of the three preconditioners presented in Section 3, as well as the linear multigrid preconditioner presented in Section 5 are evaluated using three representative test cases: an inviscid transonic flow, a subsonic laminar viscous flow, and a subsonic turbulent viscous flow.

6.1. Inviscid transonic flow over NACA0012 airfoil, $M = 0.75, \alpha = 2^\circ$

The first test case is an Euler solution of the transonic flow over the NACA0012 airfoil at an angle of attack of $\alpha = 2^\circ$ with a free-stream Mach number of $M = 0.75$. This flow involves a weak shock over the upper surface of the airfoil which is captured using an artificial viscosity approach similar to that presented by Persson and Peraire [37]. This approach involves adding artificial viscosity of order h/p near the shock based on a non-linear shock indicator. This flow is solved using a hierarchical basis on a set of three grids with 276, 1836 and 7344 element, respectively. Fig. 3 shows a portion of the coarse grid and the corresponding $p = 4$ solution of density on this grid.

Table 8
Additional memory usage for lower order Jacobians for linear multigrid as a percent of the fine grid Jacobian and number of fine grid solution vectors.

Solution order	% Fine Jacobian		Solution vectors	
	2D	3D	2D	3D
$p = 1$	11.1	6.25	5	6
$p = 2$	27.7	17.0	27	43
$p = 3$	46.0	29.3	74	146
$p = 4$	64.9	42.2	156	369
$p = 5$	84.1	55.5	283	778

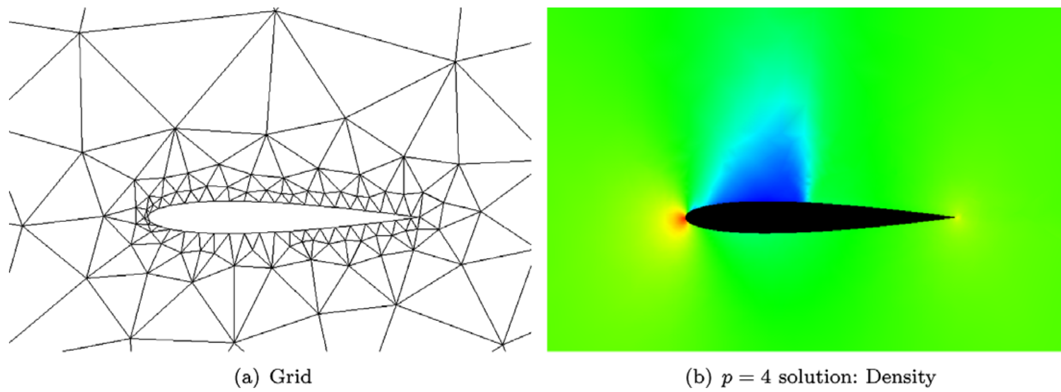

(a) Grid (b) $p = 4$ solution: Density

Fig. 3. Computational grid and solution on coarse grid (276 elements, $p = 4$) for NACA0012 transonic test case.

Solutions are obtains for $p = 0, 1, 2, 3, 4$, where each solution procedure, is initialized with the previously converged flow solution at lower p except for $p = 0$ which is initialized using free-stream values. The solution procedure is converged to a non-linear residual value of 10^{-12}. A GMRES restart value of 200 is used for the single-level preconditioners while memory equivalent GMRES restart values of 195, 170, 125 and 40 are used for p = 1, 2, 3, and 4, respectively. The linear multigrid preconditioner involves a single V-cycle, where one pre- and post- smoothing iteration is used on each multigrid level, while five coarse level ($p = 0$) smoothing iterations are used. The number of linear iterations taken in each Newton step is determined by the tolerance criterion specified in Eq. (17) up to a maximum of 10 GMRES outer iterations.

Tables 9–11 show the convergence results for the different preconditioners in terms of the number of non-linear Newton iterations, linear GMRES iteration and CPU time. The residual tolerance criterion developed in Section 2.3 ensures sufficient convergence of the linear system in each Newton step so that super-linear convergence of the non-linear residual is observed. Additionally, the residual tolerance criterion developed in Section 2.3 ensures that the convergence history of the non-linear residual is the nearly the same for these preconditioners. Hence the number of non-linear iterations for each preconditioner is the same for $p = 0$–4 on the coarsest grid. While, on the medium grid the number of non-linear iterations is nearly constant for each preconditioner except for $p = 3$ where the Block-Jacobi preconditioner is unable to converge due to stalling of the restarted GMRES algorithm. Similarly, for the finest grid, stalling of the restarted GMRES algorithm prevents the convergence of the Block-Jacobi preconditioner for all but $p = 1$, and the linear multigrid preconditioner with Block-Jacobi smoothing for $p = 4$.

Using the single-level Block-ILU preconditioner significantly reduces the number of linear iterations required to converge compared to the single-level Line-Jacobi and Block-Jacobi preconditioners. This improved convergence using the Block-ILU

Table 9
Convergence results of the inviscid transonic NACA0012 coarse grid test case (276 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (s).

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	36	15,797	2.9	36	6151	2.1	36	3703	1.9	–	–	–	–	–	–	–	–	–
$p = 1$	28	23,271	10.2	28	9412	4.7	28	4474	3.2	28	5842	5.4	28	3168	4.8	28	1905	3.8
$p = 2$	28	32,487	37.1	28	13,071	15.0	28	4606	8.2	28	6004	15.8	28	3453	13.3	28	1607	9.3
$p = 3$	30	33,853	75.4	30	13,041	28.6	30	5211	20.1	30	6343	37.7	30	3886	34.2	30	1999	25.7
$p = 4$	31	33,038	149.9	31	13,108	66.2	31	4938	47.2	31	5338	76.2	31	3142	67.5	31	1816	60.6

Table 10
Convergence results of the inviscid transonic NACA0012 medium grid test case (1836 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min). '*' denotes cases which did not converge due to stalling of restarted GMRES algorithm.

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	50	63,737	1.0	50	20,590	0.4	50	11,092	0.3	–	–	–	–	–	–	–	–	–
$p = 1$	41	98,865	5.3	41	30,032	1.8	41	11,818	0.7	42	19,400	1.6	42	7447	1.0	41	4400	0.7
$p = 2$	33	80,081	10.2	32	25,232	3.4	32	10,314	1.5	33	14,817	3.5	35	6151	2.2	32	2720	1.4
$p = 3$	*	*	*	38	34,096	10.7	38	12,918	4.2	38	17182	9.4	38	6081	5.2	38	3305	4.0
$p = 4$	34	11,4381	64.5	32	24,854	16.2	32	9187	7.2	32	11779	17.2	34	4827	10.3	32	2247	8.0

Table 11
Convergence results of the inviscid transonic NACA0012 fine grid test case (7344 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min). '*' denotes cases which did not converge due to stalling of restarted GMRES algorithm.

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	*	*	*	90	157,936	5.9	90	68,733	3.0	–	–	–	–	–	–	–	–	–
$p = 1$	52	204,664	58.2	56	71,766	24.8	55	25,445	7.1	50	31681	13.5	55	13,552	7.3	52	7661	4.4
$p = 2$	*	*	*	51	134,543	71.4	53	53,242	18.4	54	74,045	45.0	54	22,177	19.8	52	10,230	11.9
$p = 3$	*	*	*	37	53,768	108.7	37	15,879	25.1	37	20,900	65.1	37	6381	26.7	37	3489	18.1
$p = 4$	*	*	*	35	29,285	110.1	37	13,169	43.6	*	*	*	36	5476	53.7	36	3053	39.6

preconditioner ensures that the GMRES restart value is reached less often. On the other hand, the GMRES restart value is reached in nearly all Newton iterations for the Block-Jacobi preconditioner and most Newton iterations for the Line-Jacobi preconditioner. The repeated restarting of the GMRES algorithm degrades the convergence rate and leads to the stalling of the GMRES algorithm using the Block-Jacobi preconditioner. While both the preprocessing and the iterative stages of the Block-ILU preconditioner are more expensive than the corresponding stages of the Line-Jacobi or Block-Jacobi preconditioners, the significant reduction in the number of linear iterations ensures that the Block-ILU preconditioner achieves fastest convergence in terms of CPU time.

The linear multigrid preconditioners with Block-Jacobi, Line-Jacobi and Block-ILU smoothing significantly reduce the number of linear iterations required to achieve convergence compared to the corresponding single-level preconditioners. The improved convergence rate in terms of the number of linear iterations ensure that the GMRES restart value is not reached as often for the multi-level preconditioners despite the memory equivalent GMRES restart value being smaller than for the single-level preconditioners. Note that this is the case even for $p = 4$ where the GMRES restart value for the single-level preconditioner is five times larger than for the corresponding multigrid preconditioner. That the GMRES restart value is not reached as often for the multigrid preconditioner ensures that GMRES stall is not seen as often with the linear multigrid preconditioner using Block-Jacobi smoothing.

Though the linear multigrid preconditioner significantly reduces the number of linear iterations required to converge this problem, the cost of each application of the linear multigrid preconditioner is more expensive than the single-level preconditioner. For the coarsest grid, fastest convergence in the range $p = 1$—4 is achieved by the Block-ILU preconditioner. On the medium grid both Block-ILU and linear multigrid using Block-ILU smoothing perform equally well. While on the finest grid fastest convergence is achieved using linear multigrid preconditioner with Block-ILU smoothing.

6.2. Viscous subsonic flow over NACA0005 airfoil, $M = 0.4, \alpha = 0^\circ, Re = 50000$

The second test case is a Navier–Stokes solution of a subsonic, $M = 0.4$ flow over the NACA0005 airfoil at zero angle of attack with Reynolds number $Re = 50,000$. A steady, laminar solution of this flow is obtained using an output based adaptation scheme using simplex cut-cell meshes [18]. Convergence studies are performed on grids 2, 4, 6 and 8 from the adaptation procedure, where solutions are obtains for p = 0, 1, 2, and 3 using a Lagrange basis on each grid. The four meshes for which convergence results are presented have 3030, 3433, 4694 and 6020 elements, respectively. Fig. 4 shows a portion of the grid # 2 and the corresponding $p = 3$ solution of the Mach number on this grid.

The solution procedure is initialized with the previously converged flow solution at lower p except for $p = 0$ which is initialized using free-stream values. A GMRES restart value of 120 is used for the single-level preconditioners, while a memory equivalent 115, 90, and 40 GMRES iterations are used for the linear multigrid preconditioners for p = 1, 2 and 3, respectively. The linear multigrid preconditioner involves a single V-cycle, where one pre- and post- smoothing iteration is used on each multigrid level, while 5 coarse level ($p = 0$) smoothing iterations are used. The non-linear residual is converged to a tolerance of 10^{-10}, while the linear system at each Newton iteration is converged based on the criterion described in Section 2.3. The convergence data for the four grids are summarized in Tables 12–15.

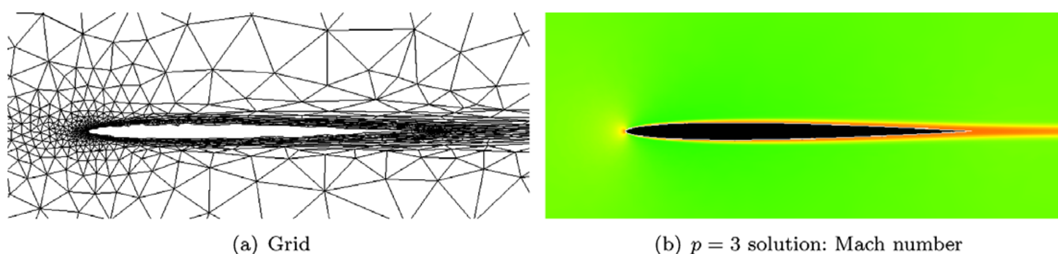

Fig. 4. Computational grid and solution on grid # 2 for NACA0005 viscous test case (3030 elements, $p = 3$).

Table 12
Convergence results of the viscous NACA0005 test case with adapted grid # 2 (3030 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min).

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	31	45,543	1.5	31	14,672	0.9	31	7131	0.8	–	–	–	–	–	–	–	–	–
$p = 1$	25	69,468	4.9	24	19146	2.2	24	7658	1.1	24	13,228	3.3	24	4231	1.6	24	2408	1.1
$p = 2$	27	110,001	16.1	24	25,845	6.9	24	9034	2.9	24	17,642	12.3	24	4970	4.7	24	2901	2.9
$p = 3$	25	89,138	36.8	22	21576	16.2	22	6817	5.9	22	12,234	20.3	22	3997	10.3	22	2098	6.0

Table 13
Convergence results of the viscous NACA0005 test case with adapted grid # 4 (3433 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min).

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	32	47,949	1.7	32	15,437	1.0	32	7755	0.9	–	–	–	–	–	–	–	–	–
$p = 1$	31	172,474	8.9	24	23,862	3.3	24	9065	1.5	24	16876	5.3	24	4682	1.9	24	2567	1.2
$p = 2$	30	164,594	28.2	24	29,954	11.6	24	10,048	4.2	24	21368	18.1	24	5658	6.4	24	3118	3.6
$p = 3$	23	60,482	34.4	22	25,424	22.6	22	7673	7.9	22	12763	23.9	22	5004	16.6	22	2169	7.3

Table 14
Convergence results of the viscous NACA0005 test case with adapted grid # 6 (4694 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min).

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	32	56,198	2.6	32	16,827	1.5	32	8629	1.2	–	–	–	–	–	–	–	–	–
$p = 1$	37	327,309	17.6	25	34,588	6.0	25	12,976	2.6	25	26,244	10.4	25	6375	3.5	25	3446	2.1
$p = 2$	31	186,272	45.0	24	34,978	19.0	24	12,741	7.8	24	27,634	31.0	24	7042	11.7	24	3483	5.8
$p = 3$	24	64,508	50.5	22	20,308	27.7	22	7883	13.0	22	11,607	31.5	22	4891	26.5	22	2296	11.5

Table 15
Convergence results of the viscous NACA0005 test case with adapted grid # 8 (6020 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min).

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	31	80,053	5.1	30	20,356	2.8	30	9986	2.0	–	–	–	–	–	–	–	–	–
$p = 1$	45	624,315	60.5	26	48,556	16.6	25	17,031	7.2	25	32,911	23.2	25	8328	8.3	25	4315	4.3
$p = 2$	31	181,094	73.2	26	33,924	28.2	24	9585	9.3	25	27,661	48.2	24	5968	17.5	24	3256	8.6
$p = 3$	30	109,457	106.2	24	14,203	30.0	25	7225	20.2	29	37,082	97.3	25	5979	34.6	25	3370	24.3

To achieve fast convergence for this viscous test case, it is necessary that the preconditioner sufficiently resolves the coupling between elements in the boundary layer. Since the Block-Jacobi preconditioner ignores all inter-element coupling, the restarted GMRES algorithm stalls and the linear system is not sufficiently solved such that several additional Newton iterations are required to converge the non-linear residual. On the other hand, the Line-Jacobi and Block-ILU preconditioners which make use of the lines of maximum coupling within the flow are able to sufficiently converge the linear system at each Newton step. Hence, the same super-linear convergence of the non-linear residual is observed for both Line-Jacobi and Block-ILU preconditioners.

As with the previous test cases, the use of the linear multigrid preconditioner significantly reduces the number of linear iterations required to converge the linear system at each Newton step. The GMRES restart value is reached less often in the case of the linear multigrid preconditioners despite the GMRES restart value being larger for the single-level preconditioners. This ensures that the linear multigrid preconditioner with Block-Jacobi smoothing is able to solve the linear system sufficiently to have essentially the same convergence of the non-linear residual as the Line-Jacobi and Block-ILU preconditioners. On average fastest convergence in terms of CPU time is achieved using the linear multigrid preconditioner with Block-ILU smoothing, which performs on average about 5% faster than the single-level Block-ILU preconditioner.

6.3. Turbulent viscous subsonic flow over NACA0012 airfoil, $M = 0.25, \alpha = 0^{\circ}, Re = 10^6$

The final test case is a Reynolds-Averaged Navier–Stokes (RANS) solution of a subsonic, $M = 0.25$ flow over the NACA0012 airfoil at a Reynolds number of $Re = 10^6$. The single equation Spalart–Allmaras turbulence model is used, where the source terms are discretized using a dual-consistent formulation [35,34]. The flow solution is obtained on a sequence of higher-order meshes using an output based adaptation scheme [32,34]. Convergence studies are performed on grids 2, 4, and 6 from the adaptation procedure, where solutions are obtains for p = 0, 1, 2, and 3 using a hierarchical basis on each grid. The three meshes for which convergence results are presented have 1209, 1522, and 3113 elements, respectively. Fig. 5 shows a portion of the grid # 2 and the corresponding $p = 3$ solution of the Mach number on this grid.

The solution procedure is initialized with the previously converged flow solution at lower p except for $p = 0$ which is initialized using free-stream values. The GMRES restart values and convergence criteria are the same as for the previous test case. The convergence data for the three grids are summarized in Tables 16–18.

For this RANS test case the non-linear residual history for $p = 1$ differs significantly from $p = 2$ and $p = 3$, typically requiring a larger number of non-linear iterations in order to obtain a converged solution. In addition the solution procedure fails using the Block-Jacobi and Block-ILU preconditioners for grid #4 due to divergence of the non-linear solution algorithm. This

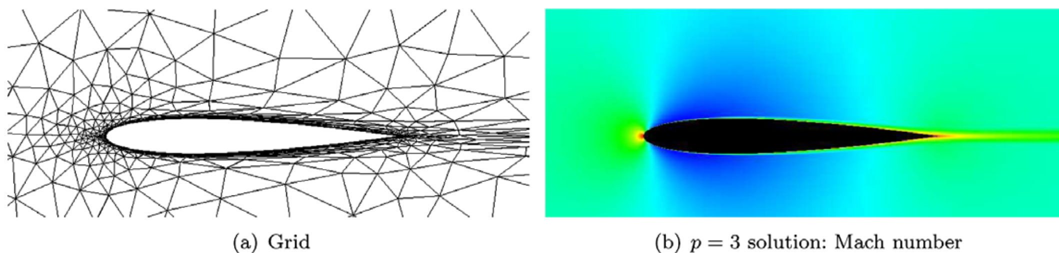

(a) Grid (b) $p = 3$ solution: Mach number

Fig. 5. Computational grid and solution on grid # 2 for NACA0012 RANS test case (1209 elements, $p = 3$).

Table 16
Convergence results of the NACA0012 RANS test case with adapted grid # 2 (1209 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min).

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	51	68,596	1.9	51	24,872	1.7	51	13,490	1.6	–	–	–	–	–	–	–	–	–
$p = 1$	95	812,446	10.1	88	207,363	5.6	90	76,455	4.1	106	261,297	7.7	67	34,995	4.2	103	49,716	5.3
$p = 2$	59	230,845	14.7	53	56,324	7.2	53	16,859	4.1	53	27,918	7.9	53	15,960	6.5	53	8466	4.8
$p = 3$	56	167,661	27.1	59	167,066	29.7	52	14,068	7.3	53	36,097	24.0	52	14,662	15.1	52	7761	9.5

Table 17
Convergence results of the NACA0012 RANS test case with adapted grid # 4 (1522 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min). '*' denotes cases which did not converge.

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	53	93,452	2.6	53	31,453	2.3	53	17,023	2.1	–	–	–	–	–	–	–	–	–
$p = 1$	*	*	*	104	352,881	9.6	*	*	*	100	321,138	11.6	100	147,480	9.5	87	29,364	5.5
$p = 2$	58	197,944	19.3	56	77,980	11.7	55	19,065	5.7	55	36,148	12.1	55	18,053	9.2	55	9400	6.5
$p = 3$	78	878,595	89.1	54	64,958	26.9	54	15,156	10.2	55	42,180	35.5	54	19,725	25.9	54	7712	12.4

Table 18
Convergence results of the NACA0012 RANS test case with adapted grid # 6 (3113 elements). Iter: total non-linear iterations, GMRES: total number of linear GMRES iterations, Time: total run time (min). '*' denotes cases which did not converge.

	Block			Line			ILU			MG-Block			MG-Line			MG-ILU		
	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time	Iter	GMRES	Time
$p = 0$	67	411,896	9.4	66	93,632	6.2	66	44,721	5.6	–	–	–	–	–	–	–	–	–
$p = 1$	104	1,777,595	44.1	101	754,669	27.9	105	118,884	13.2	101	305,972	24.4	106	92,106	17.2	105	47,713	14.4
$p = 2$	73	594,225	76.8	70	302,661	54.5	66	42,604	15.9	66	150,312	66.9	66	42,511	30.8	65	18,379	18.1
$p = 3$	93	1,620,620	250.5	*	*	*	58	21,007	25.1	69	148,697	139.9	61	38,265	73.4	62	14,052	34.6

behaviour at $p = 1$ may be due to the poor starting condition provided by the $p = 0$ solution. In practice this problem may be avoided during an adaptation procedure by starting with an initial solution interpolated from a converged solution on a previous grid.

The Block-ILU preconditioner performs significantly better than the other single-level preconditioners for this test case. The Block-Jacobi solver takes many more non-linear iterations in order to converge the $p = 3$ solution for grids #4 and #6 due to stalling of the restarted GMRES algorithm. Additionally the Line-Jacobi solver fails to converge the grid #6 case for $p = 3$. On the other and for the Block-ILU preconditioner, the GMRES restart value is reached in only half of the Newton iterations and stalling does not occur when converging $p = 3$ on grid #6.

As in the previous test cases, the use of the linear multigrid preconditioner significantly reduces the number of linear iterations compared to the single-level preconditioners. For linear multigrid with Block-Jacobi and Line-Jacobi smoothing, this ensures that a better solution update is obtained prior to restarting the GMRES algorithm. Hence, linear multigrid with both Block-Jacobi smoothing and Line-Jacobi smoothing generally require the same number of non-linear iteration as the Block-ILU preconditioner. Though the linear multigrid preconditioner with Block-ILU smoothing significantly reduces the number linear iterations compared to the single-level Block-ILU preconditioner, fastest convergence in terms of CPU time is generally seen by the single-level Block-ILU preconditioner.

7. Conclusions and discussion

An in-place Block-ILU(0) factorization algorithm has been developed, which has been shown to reduce both the memory and computational cost over the traditional dual matrix storage format. A reordering technique for the Block-ILU(0) factorization, based upon lines of maximum coupling in the flow, has also been developed. The results presented show that this reordering technique significantly reduces the number of linear iterations required to converge compared to standard reordering techniques, especially for viscous test cases.

A linear p-multigrid algorithm has been developed as a preconditioner to GMRES. The linear multigrid preconditioner is shown to significantly reduce the number of linear iterations and CPU times required to obtain a converged solution compared to a single-level Block-Jacobi or element Line-Jacobi preconditioner. The linear p-multigrid preconditioner with Block-ILU(0) smoothing also reduces the number of linear iterations relative to the single-level Block-ILU(0) preconditioner though not necessarily the total CPU time.

The solution of complex 3D problems necessitates the use of parallel computing. The development of an efficient solver for DG discretizations must therefore necessarily consider the implications of parallel computing. Except for the Block-Jacobi preconditioners, the preconditioners presented have some inherent serialism as they require elements to be traversed sequentially. Thus, while the Block-Jacobi preconditioners can be trivially parallelized, the Line-Jacobi and Block-ILU methods are more difficult. In this paper only a serial implementation is presented, while a basic parallel implementation has been discussed in [16]. While an efficient parallel implementation has yet to be developed, the preconditioners presented in this work may serve as local solvers for a more sophisticated parallel solver based on domain decomposition methods.

Acknowledgments

The authors would like to thank the anonymous reviewers for their suggestions, which significantly improved this paper. This work was partially supported by funding from The Boeing Company with technical monitor Dr. Mori Mani.

References

[1] W. Anderson, R. Rausch, D. Bonhaus, Implicit multigrid algorithms for incompressible turbulent flows on unstructured grids, No. 95-1740-CP, in: Proceedings of the 12th AIAA CFD Conference, San Diego CA, 1995.
[2] S. Balay, K. Buschelman, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L.C. McInnes, B.F. Smith, H. Zhang, Petsc users manual, Tech. Rep. ANL-95/11 – Revision 2.1.5, Argonne National Laboratory, 2004.
[3] S. Balay, K. Buschelman, W.D. Gropp, D. Kaushik, M.G. Knepley, L.C. McInnes, B.F. Smith, H. Zhang, PETSc Web page, 2007. <http://www.mcs.anl.gov/petsc>.
[4] S. Balay, W.D. Gropp, L.C. McInnes, B.F. Smith, Efficient management of parallelism in object oriented numerical software libraries, in: E. Arge, A.M. Bruaset, H.P. Langtangen (Eds.), Modern Software Tools in Scientific Computing, Birkhäuser Press, 1997.
[5] T. Barth, Numerical methods for conservation laws on structured and unstructured meshes, VKI March 2003 Lecture Series, 2003.
[6] F. Bassi, S. Rebay, High-order accurate discontinuous finite element solution of the 2d Euler equations, Journal of Computational Physics 138 (2) (1997) 251–285.
[7] F. Bassi, S. Rebay, A high-order discontinuous finite element method for the numerical solution of the compressible Navier–Stokes equations, Journal of Computational Physics 131 (1997) 267–279.
[8] F. Bassi, S. Rebay, GMRES discontinuous Galerkin solution of the compressible Navier–Stokes equations, in: K. Cockburn, Shu (Eds.), Discontinuous Galerkin Methods: Theory, Computation and Applications, Springer, Berlin, 2000, pp. 197–208.
[9] F. Bassi, S. Rebay, Numerical evaluation of two discontinuous Galerkin methods for the compressible Navier–Stokes equations, International Journal for Numerical Methods in Fluids 40 (2002) 197–207.
[10] M. Benzi, D.B. Szyld, A. van Duin, Orderings for incomplete factorization preconditioning of nonsymmetric problems, SIAM Journal on Scientific Computing 20 (5) (1999) 1652–1670.
[11] M. Blanco, D.W. Zingg, A fast solver for the Euler equations on unstructured grids using a Newton-GMRES method, AIAA Paper 1997-0331, January, 1997.

[12] X.-C. Cai, W.D. Gropp, D.E. Keyes, M.D. Tidriri, Newton–Krylov–Schwarz methods in CFD, in: Proceedings of the International Workshop on Numerical Methods for the Navier–Stokes Equations, 1995.
[13] B. Cockburn, G. Karniadakis, C. Shu, The development of discontinuous Galerkin methods, Lecture Notes in Computational Science and Engineering, vol. 11, Springer, 2000.
[14] B. Cockburn, C.-W. Shu, Runge–Kutta discontinuous Galerkin methods for convection-dominated problems, Journal of Scientific Computing (2001) 173–261.
[15] L. Diosady, D. Darmofal, Discontinuous Galerkin solutions of the Navier–Stokes equations using linear multigrid preconditioning, AIAA Paper 2007-3942, 2007.
[16] L.T. Diosady, A linear multigrid preconditioner for the solution of the Navier–Stokes equations using a discontinuous Galerkin discretization, Masters thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, May 2007.
[17] V. Dolejší, M. Feistauer, A semi-implicit discontinuous Galerkin finite element method for the numerical solution of inviscid compressible flow, Journal of Computational Physics 198 (1) (2004) 727–746.
[18] K. Fidkowski, D. Darmofal, An adaptive simplex cut-cell method for discontinuous AIAA Paper 2007-3941, Massachusetts Institute of Technology, 2007.
[19] K.J. Fidkowski, A high-order discontinuous Galerkin multigrid solver for aerodynamic applications, Masters Thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2004.
[20] K.J. Fidkowski, D.L. Darmofal, Development of a higher-order solver for aerodynamic applications, AIAA Paper 2004-0436, January 2004.
[21] K.J. Fidkowski, T.A. Oliver, J. Lu, D.L. Darmofal, *p*-multigrid solution of high-order discontinuous Galerkin discretizations of the compressible Navier–Stokes equations, Journal of Computational Physics 207 (1) (2005) 92–113.
[22] B.T. Helenbrook, D. Mavriplis, H.L. Atkins, Analysis of *p*-multigrid for continuous and discontinuous finite element discretizations, AIAA Paper 2003-3989, 2003.
[23] K. Hillewaert, N. Chevaugeon, P. Geuzaine, J.-F. Remacle, Hierarchic multigrid iteration strategy for the discontinuous Galerkin solution of the steady Euler equations, International Journal for Numerical Methods in Fluids 51 (9) (2005) 1157–1176.
[24] C.T. Kelley, D.E. Keyes, Convergence analysis of pseudo-transient continuation, SIAM Journal of Numerical Analysis 35 (2) (1998) 508–523.
[25] D.A. Knoll, D.E. Keyes, Jacobian-free Newton–Krylov methods: a survey of approaches and applications, Journal of Computational Physics 193 (1) (2004) 357–397.
[26] H. Luo, J.D. Baum, R. Löhner, A *p*-multigrid discontinuous Galerkin method for the Euler equations on unstructured grids, Journal of Computational Physics 211 (1) (2006) 767–783.
[27] J. Mavriplis, On convergence acceleration techniques for unstructured meshes, AIAA Paper 1998-2966, 1998.
[28] D.J. Mavriplis, An assessment of linear versus nonlinear multigrid methods for unstructured mesh solvers, Journal of Computational Physics 175 (1) (2002) 302–325.
[29] C.R. Nastase, D.J. Mavriplis, High-order discontinuous Galerkin methods using a spectral multigrid approach, AIAA Paper 2005-1268, January 2005.
[30] C.R. Nastase, D.J. Mavriplis, High-order discontinuous Galerkin methods using an *hp*-multigrid approach, Journal of Computational Physics 213 (1) (2006) 330–357.
[31] A. Nejat, C. Ollivier-Gooch, Effect of discretization order on preconditioning and convergence of a higher-order unstructured Newton–Krylov solver for inviscid compressible flows, AIAA Paper 2007-0719, January 2007.
[32] T. Oliver, D. Darmofal, An unsteady adaptation algorithm for discontinuous Galerkin discretizations of the RANS equations, AIAA Paper 2007-3940, 2007.
[33] T.A. Oliver, Multigrid solution for high-order discontinuous Galerkin discretizations of the compressible Navier–Stokes equations, Masters Thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2004.
[34] T.A. Oliver, A higher-order, adaptive, discontinuous Galerkin finite element method for the Reynolds-averaged Navier–Stokes equations, Ph.D. Thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2008.
[35] T.A. Oliver, D.L. Darmofal, An analysis of dual consistency for discontinuous Galerkin discretization of source terms, ACDL report, Massachusetts Institute of Technology (2007).
[36] P.-O. Persson, J. Peraire, An efficient low memory implicit DG algorithm for time dependent problems, AIAA Paper 2006-0113, 2006.
[37] P.-O. Persson, J. Peraire, Sub-cell shock capturing for discontinuous Galerkin methods, AIAA Paper 2006-0112, 2006.
[38] P.-O. Persson, J. Peraire, Newton-GMRES preconditioning for discontinuous Galerkin discretizations of the Navier–Stokes equations, SIAM Journal for Scientific Computing 30 (6) (2008) 2709–2722.
[39] A. Pueyo, D.W. Zingg, An efficient Newton-GMRES solver for aerodynamic computations, AIAA Paper 1997–1955, June 1997.
[40] P. Rasetarinera, M.Y. Hussaini, An efficient implicit discontinuous spectral Galerkin method, Journal of Computational Physics 172 (1) (2001) 718–738.
[41] P.L. Roe, Approximate Riemann solvers, parameter vectors, and difference schemes, Journal of Computational Physics 43 (2) (1981) 357–372.
[42] Y. Saad, Iterative methods for sparse linear systems, Society for Industrial and Applied Mathematics (1996).
[43] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM Journal on Scientific and Statistical Computing 7 (3) (1986) 856–869.
[44] L.N. Trefethen, D. Bau, Numerical linear algebra, Society for Industrial and Applied Mathematics (1997).