

Numbering Techniques for Preconditioners in Iterative Solvers for Compressible Flows

Bernhard Pollul ^{*} [†] Arnold Reusken ^{*} [†]

March 2, 2006

Subject Classification: [msc2000] 65F10, 65N22

Key Words: Euler equations, Krylov subspace methods, preconditioning, ordering algorithms

Copyright notice: This is a preprint of an article published in *International Journal for Numerical Methods in Fluids* Volume 55(3), 241–261

Abstract

We consider Newton-Krylov methods for solving discretized compressible Euler equations. A good preconditioner in the Krylov subspace method is crucial for the efficiency of the solver. In this paper we consider a point-block Gauss-Seidel method as preconditioner. We describe and compare renumbering strategies that aim at improving the quality of this preconditioner. A variant of reordering methods known from multigrid for convection-dominated elliptic problems is introduced. This reordering algorithm is essentially black-box and significantly improves the robustness and efficiency of the point-block Gauss-Seidel preconditioner. Results of numerical experiments using the QUADFLOW solver and the PETSc library are given.

1 Introduction

We are interested in efficient numerical techniques for the numerical simulation of two- and three-dimensional compressible flows. One important issue in this field is the solution of large sparse nonlinear systems of equations that arise after spatial discretization combined with an implicit time integration method. Two popular approaches for solving such nonlinear systems of equations are nonlinear multigrid solvers and Newton-Krylov methods. Well-known nonlinear multigrid techniques are the FAS method by Brandt [12], the nonlinear multigrid method by Hackbusch [17] and the algorithm introduced by Jameson [21]. It has been shown that a nonlinear multigrid approach can result in very efficient solvers which can even have optimal complexity for a certain class of problems [22, 26, 27, 29]. Multigrid methods, however, require a coarse-to-fine grid hierarchy, whereas Newton-Krylov algorithms only need the matrix

^{*}This research was supported by the German Research Foundation through the Collaborative Research Centre SFB 401

[†]Institut für Geometrie und Praktische Mathematik, RWTH-Aachen, D-52056 Aachen, Germany pollul@igpm.rwth-aachen.de , reusken@igpm.rwth-aachen.de

of the linearized system. Due to this and to the fact that efficient implementations of many (preconditioned) Krylov subspace algorithms in sparse matrix libraries are available, the Newton-Krylov algorithms are in general much easier to implement than multigrid solvers. Furthermore, concerning efficiency, the Newton-Krylov approach with appropriate preconditioning can be competitive with multigrid. Thus it is not surprising that Newton-Krylov techniques are often used in practice (cf., for instance, [28, 30, 31, 32, 33, 45]).

In this paper we consider the Newton-Krylov approach. A method of this class has been implemented in the QUADFLOW package, which is an adaptive multiscale finite volume solver for stationary and instationary compressible flow computations. Descriptions of this solver are given in [5, 9, 10, 11, 39]. For the linearization we use a standard (approximate) Newton method. The resulting linear systems are solved by a preconditioned BiCGSTAB method using methods implemented in the PETSc-Library [3, 4].

Incomplete LU-factorization and Gauss-Seidel techniques are popular preconditioners that are used in solvers in the numerical simulation of compressible flows [1, 2, 33, 38]. The “point-block”-variants of these preconditioners are obtained by applying the original point versions to the blocks of unknowns corresponding to each cell. Both preconditioners depend on the ordering of the cells (grid-points) [2, 8, 18, 19, 33, 36, 42]. In combination with PBILU the reverse Cuthill-McKee ordering algorithm [13, 14] is often used. This ordering yields a matrix with a “small” bandwidth.

In this paper we focus on ordering algorithms for the PBGS preconditioner. We do not know of any literature that deals with ordering techniques for Gauss-Seidel preconditioners applied to linearized Euler equations. The ordering algorithms consist of three steps. First a weighted directed graph, in which every vertex corresponds to a block unknown, is constructed. Then this graph is reduced by deleted edges with relatively small weights. This graph reduction is very similar to techniques used in algebraic multigrid methods [40]. Finally a renumbering of the vertices in this reduced graph all pictures can also used in black/white is determined. For this we consider three different algorithms. Two of them are known (due to Bey, Wittum [8] and Hackbusch [16, 18]) from the field of robust multigrid solvers for convection-dominated elliptic problems. The third one is new. These methods are implemented in the QUADFLOW solver using the PETSc library. A systematic comparative study shows that for our problem class the new variant yields the best results. The reordering algorithm is essentially black-box. Using this reordering we can improve the robustness of the iterative solver: For large CFL numbers we encounter linear systems for which the BiCGSTAB method with PBGS preconditioner converges only if we first apply the reordering. Using the reordering we can also improve the efficiency of the linear solver significantly. The execution time of the iterative solver part can be reduced by 10% (for complex transonic flows) up to 50% (for supersonic flows).

The remainder of this paper is organized as follows. In the following section we outline the discretization and linearization methods that are used in QUADFLOW for the numerical solution of the compressible Euler equations. In section 3 we describe the point-block-Gauss-Seidel preconditioner. Section 4 gives a detailed description of three renumbering algorithms. In section 5 we apply these algorithms, implemented in the QUADFLOW solver, to some test problems. Finally we summarize some main results of the paper (section 6).

2 Discrete Euler equations

We consider the conservative formulation of the Euler equations. For an arbitrary control volume $V \subset \Omega \subset \mathbb{R}^d$ ($d = 2, 3$) one has equations of the form

$$\int_V \frac{\partial \mathbf{u}}{\partial t} dV + \oint_{\partial V} \mathbf{F}^c(\mathbf{u}) \mathbf{n} dS = 0. \quad (1)$$

Here \mathbf{n} is the outward unit normal on ∂V , $\mathbf{u} = (\rho, \rho \mathbf{v}, \rho e_{tot})^T$ the vector of unknown conserved quantities and the convective flux is given by

$$\mathbf{F}^c(\mathbf{u}) = \begin{pmatrix} \rho \mathbf{v} \\ \rho \mathbf{v} \circ \mathbf{v} + p \mathbf{I} \\ \rho h_{tot} \mathbf{v} \end{pmatrix}. \quad (2)$$

The symbol \circ denotes the dyadic product and h_{tot} is the total enthalpy. The system is closed by the equation of state for a perfect gas and suitable initial and boundary conditions. For the numerical simulation of the compressible Euler equations we use the software package QUADFLOW, which is currently under development at Aachen University, cf. [5, 9, 10, 11, 39]. We briefly describe a few main features of this solver. QUADFLOW contains methods for the numerical simulation of two- and three-dimensional compressible Euler- and Navier-Stokes equations. It is based on block-structured grids. The geometry of these blocks is described using tensor-product B-splines. For discretization finite volume techniques are applied. Several upwind methods, for instance flux-difference splitting (HLLC [41]), flux-vector splitting (van Leer [25], Hänel/Schwane [20]) and AUSDMV(p) [15, 46] have been implemented. A key ingredient in QUADFLOW is the use of local grid refinement in regions of high activity, for example in the neighborhood of shocks. Both explicit and implicit time integration routines are available. The computation of an accurate approximation of a stationary solution is based on a nested iteration approach. One starts with an initial coarse grid and an initial CFL number γ_0 , which determines the size of the timestep. After each timestep in the time integration the CFL number (and thus the timestep) is increased by a constant factor until an a-priori fixed upper bound γ_{max} is reached. Time integration is continued until a tolerance criterion for the residual is satisfied. Then a (local) grid refinement is performed and the procedure starts again with an interpolated initial condition and a starting CFL number equal to γ_0 . The indicator for the local grid refinement is based on a multiscale analysis using wavelets. The nonlinear systems that arise in each timestep of an implicit method are solved using a Newton-Krylov approach. In every timestep one approximate Newton iteration is performed. The resulting linear equations are solved using preconditioned Krylov-subspace methods that are available the PETSc library [3, 4]. For this an interface between QUADFLOW and PETSc has been developed. A first parallel version of QUADFLOW is available now. In this paper, however, we restrict ourselves to the sequential version. To give an impression of the multi-block and adaptivity features of QUADFLOW we show grids that are used in a simulation of an inviscid flow around a BAC 3-11/RES/30/21 airfoil (cf. section 5.2) in fig. 1.

In most simulations an implicit time integration is used. Then the computational work for solving the large sparse systems in the Newton-Krylov method determines to a large extent the total computing time in a simulation run.

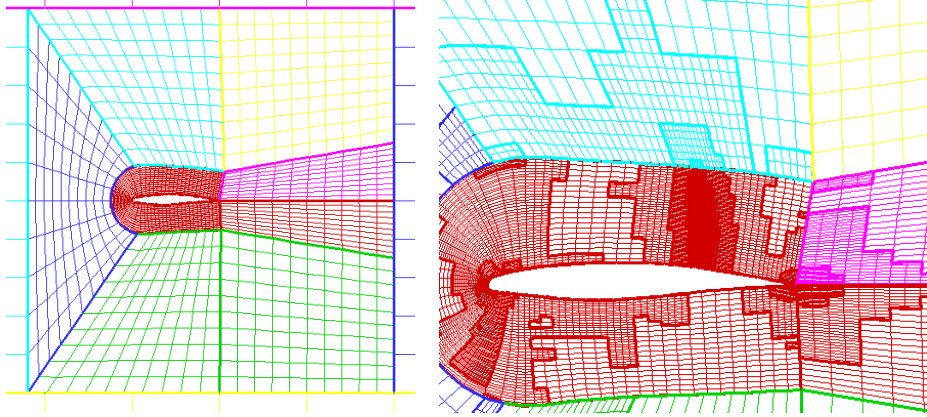


Figure 1: 12-block grid of BAC 3-11/RES/30/21 airfoil: part of the original grid and of the grid after 10 levels of local refinement. Test problem 2, cf. section 5.2

Hence, the efficiency of the iterative solvers for these systems is an important issue. In general for stationary problems this issue plays a bigger role than for nonstationary problems. We therefore focus on stationary problems in this paper.

For the discretization we choose methods that are available in QUADFLOW. For spatial discretization the flux-vector splitting by Hänel and Schwane [20] is applied. A linear reconstruction technique is used to obtain second order accuracy in regions where the solution is smooth. This is combined with the Venkatakrishnan limiter [44]. Although we are interested in stationary solutions the time derivative is not skipped. This time derivative is discretized by a numerical integration method which then results in a numerical method for approximating the stationary solution. To obtain fast convergence towards the stationary solution one wants to use large timesteps and thus an implicit time discretization method is preferred. We use the b2-scheme by Batten et. al. [7]. This approach then results in a nonlinear system of equations in each timestep. Per timestep one inexact Newton iteration is applied. In this inexact Newton method an *approximate* Jacobian is used in which the linear reconstruction technique is neglected and the Jacobian of the first order Hänel-Schwane discretization is approximated by one-sided difference operators (as in [43]). These Jacobian matrices have the structure

$$DF(\mathbf{U}) = \text{diag}\left(\frac{|V_i|}{\Delta t}\right) + \frac{\partial R^{HS}(\mathbf{U})}{\partial \mathbf{U}}, \quad (3)$$

where $|V_i|$ is the volume of cell V_i , Δt the (local) timestep and $R^{HS}(\mathbf{U})$ the residual vector corresponding to the Hänel-Schwane fluxes. Details are given in [11]. Note that in general a smaller timestep will improve the condition number of the approximate Jacobian in (3).

In this paper we introduce and compare several renumbering techniques that aim at improving the efficiency of preconditioned Krylov subspace methods for solving these linear systems in the approximate Newton linearization. We emphasize that for these ordering techniques the particular choice of discretization

components is not essential. The renumbering methods show a similar behavior, if instead of the Hänel-Schwane method, one uses another upwind method (see above), or if, instead of the Batten b2-scheme, one uses another implicit time integration method.

3 Point-block-Gauss-Seidel preconditioner

The approximate Newton method described above leads to large sparse linear systems of equations. For solving these systems we use a standard preconditioned Krylov subspace method, available in PETSc. We choose BiCGSTAB with a point-block-Gauss-Seidel preconditioner. We briefly explain the latter.

If the cells are numbered $i = 1, \dots, N$, then the approximate Jacobian has a point-block structure $DF(\mathbf{U}) = \text{blockmatrix}(A_{i,j})_{0 \leq i,j \leq N}$ with $A_{i,j} \in \mathbb{R}^{d \times d}$ for all i, j and $A_{i,j} \neq 0$ only if $i = j$ or i and j correspond to neighboring cells. Thus we have linear systems of the form

$$\mathbf{A}\mathbf{x} = \mathbf{b} \ , \quad \mathbf{A} = \text{blockmatrix}(A_{i,j})_{1 \leq i,j \leq N} \ , \quad A_{i,j} \in \mathbb{R}^{d \times d} \ . \quad (4)$$

For the right hand side we use a block representation $\mathbf{b} = (b_1, \dots, b_N)^T$, $b_i \in \mathbb{R}^d$, that corresponds to the block structure of \mathbf{A} . The same is done for the iterands \mathbf{x}^k that approximate the solution of the linear system in (4). The point-block-Gauss-Seidel method (PBGs) is the standard block Gauss-Seidel method applied to (4). Let \mathbf{x}^0 be a given starting vector. For $k \geq 0$ the iterand $\mathbf{x}^{k+1} = (x_1^{k+1}, \dots, x_N^{k+1})^T$ should satisfy

$$A_{i,i}x_i^{k+1} = b_i - \sum_{j=1}^{i-1} A_{i,j}x_j^{k+1} - \sum_{j=i+1}^N A_{i,j}x_j^k \ , \quad i = 1, \dots, N \ . \quad (5)$$

This method is well-defined if the $d \times d$ linear systems in (5) are uniquely solvable, i.e., if the diagonal blocks $A_{i,i}$ are nonsingular. In our applications this was always the case. This elementary method is very easy to implement and needs no additional storage. The algorithm is available in the PETSc library [3, 4].

4 Renumbering techniques

Incomplete LU-decomposition and Gauss-Seidel techniques are often used for preconditioning Krylov subspace methods applied to linear systems that arise in numerical simulations of compressible flows (cf. [1, 2, 33, 38]). Both preconditioners depend on the ordering of the cells (points) [8, 18, 19, 33, 36, 42]. This holds for the point-block variants point-block-ILU (PBILU) and PBGS, too. There are many studies available on numbering techniques for ILU preconditioners (cf. [2, 37] and references therein). For PBILU a reverse Cuthill-McKee ordering algorithm [13, 14] often leads to good results. This ordering yields a matrix with a “small” bandwidth which is favorable for PBILU. Such PBILU methods combined with reordering techniques are often used in iterative solvers for compressible flow problems. A PBGS preconditioner is particularly useful in parallel and/or matrix-free iterative solvers. As for PBILU this preconditioner can be improved significantly by reordering techniques. For PBGS the ordering

should be such that one approximately follows the directions in which information is propagated. In this section we introduce three renumbering methods that aim at realizing this. The first two of these algorithms are from the field of robust multigrid methods for convection-dominated problems and are due to Bey, Wittum [8] and Hackbusch [18]. The third one is a new variant, which for our applications turns out to be better.

All three algorithms are completely matrix-based, in the sense that one needs as input only the block-structured matrix from (4). In these algorithms we distinguish the following three steps:

1. *Construct a weighted directed matrix graph* in which every vertex corresponds to a block unknown and each edge to a nonzero off-diagonal block of the given matrix \mathbf{A} .
2. *Construct a reduced weighted directed matrix graph*. The reduction is obtained by deleting edges with relatively small weights.
3. *Determine a renumbering of the vertices*, based on the reduced weighted matrix graph. This provides a point-block-permutation of the given matrix \mathbf{A} .

While for all three algorithms presented below steps 1 and 2 are identical, they differ in the methods used in step 3. We explain these first two steps in sections 4.1 and 4.2.

4.1 Construction of weighted directed matrix graph $\mathcal{G}(\mathbf{A})$

We introduce standard notation related to matrix graphs. Let $\mathcal{V} = \{1, \dots, N\}$ be a vertex set (each vertex corresponds to a discretization cell). The set of edges \mathcal{E} contains all directed edges

$$\mathcal{E} = \{(i, j) \in \mathcal{V} \times \mathcal{V} \mid A_{i,j} \neq 0, i \neq j\} \quad (6)$$

Note that \mathcal{E} does not contain edges (i, i) . The mapping

$$\omega : \mathcal{E} \rightarrow (0, \infty) \quad (7)$$

assigns to every directed edge $(i, j) \in \mathcal{E}$ a weight

$$\omega_{ij} := \omega(i, j) := \|A_{i,j}\|_F \quad (8)$$

We take the Frobenius-norm because it is easy to compute and all entries in a block $A_{i,j}$ are weighted equally. This yields a *weighted directed* matrix graph $\mathcal{G} = \mathcal{G}(\mathbf{A})$

$$\mathcal{G}(\mathbf{A}) := (\mathcal{V}, \mathcal{E}, \omega) \quad (9)$$

Every edge $(i, j) \in \mathcal{E}$ is called an *inflow edge* of vertex $i \in \mathcal{V}$ and an *outflow edge* of vertex $j \in \mathcal{V}$. For $(i, j) \in \mathcal{E}$ we call j a *predecessor* of i and i a *successor* of j . The set of predecessors of vertex $i \in \mathcal{V}$ is denoted by

$$\mathcal{I}_i := \{j \in \mathcal{V} \mid (i, j) \in \mathcal{E}\} \quad (10)$$

In the construction of $\mathcal{G}(\mathbf{A})$ one only has to compute the weights ω_{ij} in (8). For storage of this information we use a sparse matrix format. Note that the size of the sparse matrix corresponding to $\mathcal{G}(\mathbf{A})$ is $N \times N$ (and not $Nd \times Nd$, as for \mathbf{A}). Hence, the costs both for the computation and the storage of $\mathcal{G}(\mathbf{A})$ are low.

4.2 Construction of reduced matrix graph $\hat{\mathcal{G}}$

Based on reduction techniques from algebraic multigrid methods in which *strong couplings* and *weak couplings* are distinguished [40, 35, 24], we separate *strong edges* from *weak edges*. For every vertex $i \in \mathcal{V}$ we neglect all inflow edges $(i, j) \in \mathcal{E}$ with a weight smaller than τ -times the average of the weights of all inflow edges of vertex i . Thus we obtain a reduced set of *strong* edges $\hat{\mathcal{E}}$ and a corresponding reduced (weighted directed) graph $\hat{\mathcal{G}}(\mathbf{A})$:

$$\sigma_i := \frac{1}{|\mathcal{I}_i|} \sum_{j \in \mathcal{I}_i} \omega_{ij} , \quad (11)$$

$$\hat{\mathcal{E}} := \{(i, j) \in \mathcal{E} \mid \omega_{ij} \geq \tau \cdot \sigma_i\} , \quad (12)$$

$$\hat{\mathcal{G}}(\mathbf{A}) := (\mathcal{V}, \hat{\mathcal{E}}, \omega|_{\hat{\mathcal{E}}}) . \quad (13)$$

This simple construction of a reduced matrix graph $\hat{\mathcal{G}}(\mathbf{A})$ can be realized with low computational costs. In the rest of the reordering method we do not need $\mathcal{G}(\mathbf{A})$ anymore, and thus we do not need additional storage because we can overwrite $\mathcal{G}(\mathbf{A})$ with $\hat{\mathcal{G}}(\mathbf{A})$.

In the following sections we present three different methods that are used in step 3, resulting in three different ordering algorithms.

4.3 Downwind numbering based on $(\mathcal{V}, \hat{\mathcal{E}})$ (Bey and Wittum)

A numbering algorithm due to Bey and Wittum (Algorithm 4.3 in [8]) is presented in fig. 2 and denoted by “BW”. It is used in multigrid methods for scalar convection-diffusion problems to construct so-called robust smoothers. To apply this algorithm for our class of problems we need the reduced directed graph $(\mathcal{V}, \hat{\mathcal{E}})$ as input. Note that the weights ω_{ij} are *not* used.

```

for all  $P \in \mathcal{V}$  : Index( $P$ ) := -1 ;
 $n_F := 1$ 
for  $P \in \mathcal{V}$ 
  (if Index( $P$ ) < 0 ) SetF( $P$ );
end  $P$ 

procedure SetF( $P$ )
  (if all predecessors  $B$  of  $P$  have Index( $B$ ) > 0 )
    Index( $P$ ) :=  $n_F$ ;
     $n_F := n_F + 1$ ;
    for  $Q$  successor of  $P$ 
      if (Index( $Q$ ) < 0) SetF( $Q$ );
    end  $Q$ 
  end if
end procedure

```

Figure 2: Downwind numbering algorithm BW

remark 1 In the loop over $P \in \mathcal{V}$ in algorithm BW the ordering of the block-unknowns (cells) corresponding to the input matrix \mathbf{A} is used. In the procedure $\text{SetF}(P)$ a vertex is assigned the next number if all its predecessors have already been numbered. Hence, the first number is assigned to a vertex that has no inflow edges. Note that in the procedure $\text{SetF}(P)$ there is freedom in the order in which the successors Q are processed. In our implementation we again use the ordering induced by the given matrix \mathbf{A} . The BW numbering is applied to the reduced matrix graph. If that graph is cycle-free the algorithm returns a renumbering that is optimal in the sense that this reordering applied to the matrix corresponding to $\hat{\mathcal{G}}(\mathbf{A})$ results in a lower triangular matrix. However, in our problem class the reduced graphs in general contain cycles. In that case, after algorithm BW has finished there still are vertices $P \in \mathcal{V}$ with $\text{Index}(P) = -1$, i.e., there are $N - n_F > 0$ vertices that have no (new) number. The numbers n_F, \dots, N are assigned to these remaining vertices in the order induced by the input matrix ordering. The two variants of BW that are treated below in general have fewer of such “remaining” vertices.

Note that in this algorithm there are logical operations and assignments but no arithmetic operations.

4.4 Down- and upwind numbering based on $(\mathcal{V}, \hat{\mathcal{E}})$ (Hackbusch)

In fig. 3 we present an ordering algorithm, denoted by “HB”, that is due to Hackbusch [18]. As input for this algorithm one needs the reduced directed graph $(\mathcal{V}, \hat{\mathcal{E}})$ (*no* weights required). The presentation of this algorithm is as in section 2.1 in [16]. The Routine “SetF” is the same as in the BW algorithm in figure 2.

```

for all  $P \in \mathcal{V}$  :  $\text{Index}(P) := -1$  ;
 $n_F := 1$ ;  $n_L := N$ ;
for  $P \in \mathcal{V}$ 
  (if  $\text{Index}(P) < 0$  )  $\text{SetF}(P)$ ;
  (if  $\text{Index}(P) < 0$  )  $\text{SetL}(P)$ ;
end  $P$ 

procedure  $\text{SetL}(P)$ 
  (if all successors  $B$  of  $P$  have  $\text{Index}(B) > 0$  )
     $\text{Index}(P) := n_L$ ;
     $n_L := n_L - 1$ ;
    for  $Q$  predecessor of  $P$ 
      if ( $\text{Index}(Q) < 0$ )  $\text{SetL}(Q)$ ;
    end  $Q$ 
  end if
end procedure

```

Figure 3: Down- and upwind numbering algorithm HB

remark 2 In the BW algorithm the vertices are ordered in one direction, namely “downwind” (in the “flow direction”). The algorithm due to Hackbusch uses two directions: “downwind” (setF) and “upwind” (setL). In [18]

and [16] techniques for handling cycles are presented. These techniques are rather complicated and often computationally expensive. In multigrid codes for convection-dominated problems one usually encounters the ordering algorithm HB as in fig. 3 which does not treat cycles. If the reduced matrix graph $(\mathcal{V}, \hat{\mathcal{E}})$ is not cycle-free there are remaining vertices. These are treated as described in remark 1. The computational cost of algorithm HB is comparable to that of BW.

4.5 Weighted reduced graph numbering based on $(\mathcal{V}, \hat{\mathcal{E}}, \omega_{|\hat{\mathcal{E}}})$

In this section we present a modification of the methods of Bey, Wittum and Hackbusch. As input for our method we now need the *weighted* reduced graph $(\mathcal{V}, \hat{\mathcal{E}}, \omega_{|\hat{\mathcal{E}}})$. The algorithm is denoted by “WRG” and is given in figure 4.

```

for all  $P \in \mathcal{V}$  : Index( $P$ ) := -1 ;
 $n_F := 1$ ;  $n_L := N$ ;

/* (i) apply SetF and SetL to starting vertices */
do in an outflow-ordered list: for  $P \in \mathcal{V}$  (14)
  (if Index( $P$ ) < 0 ) SetF( $P$ , 1);
end  $P$ 
do in an inflow-ordered list: for  $P \in \mathcal{V}$  (15)
  (if Index( $P$ ) < 0 ) SetL( $P$ );
end  $P$ 

/* (ii) number remaining vertices */
do in an outflow-ordered list: for  $P \in \mathcal{V}$  (16)
  (if Index( $P$ ) < 0 ) SetF( $P$ , 0);
end  $P$ 

procedure SetF( $P$ ,  $s$ )
  (if all predecessors  $B$  of  $P$  have Index( $B$ ) > 0 ) or ( $s = 0$ )
    Index( $P$ ) :=  $n_F$ ;
     $n_F := n_F + 1$ ;
    do in an outflow-ordered list: for  $Q$  successor of  $P$  (17)
      if (Index( $Q$ ) < 0) SetF( $Q$ , 1);
    end  $Q$ 
  end if

procedure SetL( $P$ )
  (if all successors  $B$  of  $P$  have Index( $B$ ) > 0 )
    Index( $P$ ) :=  $n_L$ ;
     $n_L := n_L - 1$ ;
    do in an inflow-ordered list: for  $Q$  predecessor of  $P$  (18)
      if (Index( $Q$ ) < 0) SetL( $Q$ );
    end  $Q$ 
  end if
end if

```

Figure 4: Weighted reduced graph numbering algorithm WRG

remark 3 There are two important differences to the algorithms HB and BW. The first difference is related to the arbitrariness of the order in which the vertices are handled in the loops in HB and BW, cf. remark 1. If there are different possibilities for which vertex is to be handled next we now use the *weights* ω_{ij} of the reduced graph to make a decision. This decision is guided by the principle that edges with larger weights are declared to be more important than those with relatively small weights. A weight based sorting occurs at several places, namely in (14) - (18). In (14) the vertices with no inflow edges are sorted (“starting” vertices) using the sum of the weights of the outflow edges at each vertex. Similarly, in (15) the vertices with no outflow edges are sorted. The “remaining” vertices are finally sorted based on the sum of the outflow edges at each vertex in (16). In all three cases the number of vertices to be sorted is much smaller than N and thus the time for sorting is acceptable. Sorting is also used in (17) and (18) to determine the order in which successors and predecessors are handled. In $\text{SetF}(\cdot, \cdot)$ the successors Q of the current P are sorted using the sum over the weights of all outflow edges for each Q . This is done similarly in $\text{SetL}(\cdot)$ for all predecessors of the current P .

The second difference is that the loop over the numbering routine SetF is called *two* times. The first call $\text{SetF}(P, 1)$ in part (i) of algorithm WRG is similar to the call of $\text{SetF}(P)$ in the algorithms BW and HB but now with an ordering procedure used in SetF . The second call $\text{SetF}(P, 0)$ (in part (ii) in WRG) is introduced to handle the remaining vertices that still have index value -1 . In this call we do not consider the status of inflow edges and continue numbering in downwind direction ($\text{SetF}(\cdot, 0)$). The inner call $\text{SetF}(Q, 1)$ to number the successors still requires that all predecessors have been numbered. After part (ii) of the algorithm is finished the only possibly not yet numbered vertices are trivial ones, in the sense that these are vertices that have no edges to other vertices.

Due to the additional sorting routines in (14) - (18) the computational costs of the renumbering algorithm WRG are higher than of those BW and HB. However, if we use algorithm WRG in step 3 the total time needed for the execution of the steps 1,2,3 is still acceptable, cf. remark 4.

remark 4 As indicated in our comments above, in all three algorithms the computational time that is needed and the storage requirements are modest compared to other components of the iterative solver. Of course this will not be true for general matrices but it does hold for the class of large sparse point-block-matrices that forms our problem class. In our pseudo-time integration we have a sequence of time steps on every level of adaptation. The time needed for solving the linear systems is typically increasing during the discrete time integration. This is due to the increase of the CFL-number, cf. section 2. Since the Jacobian matrices of consecutive timesteps are in some sense similar we apply reordering not in each iteration but only “now and then” and keep it for the subsequent time steps, cf. section 5. Thus the total execution time for the reordering routines is very small compared to the total time needed for the linear solves with the preconditioned Krylov-subspace method. In our test problems the reordering routines consume at most a few percent of the total execution time of the iterative solver.

Both the computational costs and the quality of the reordering algorithm depend on the parameter τ used in step 2, cf. (12). For large τ -values the

reduced set of edges $\hat{\mathcal{E}}$ contains only few elements and thus the reduced graph $\hat{\mathcal{G}}(\mathbf{A})$ is close to a trivial one. The computational costs for constructing the corresponding renumbering (step 3) are relatively low but the resulting renumbering will in general hardly improve the quality of the PBGS preconditioner. The choice of the value for the parameter τ is discussed in section 5.

5 Numerical experiments

In this section we present results of numerical experiments. We will illustrate the behavior of the different numberings presented above for a few test problems.

In all experiments below we use a left preconditioned BiCGSTAB method. The approximate Jacobian matrices as in (3) are computed in QUADFLOW. For the preconditioned BiCGSTAB method and the PBGS preconditioner we use routines from the PETSc library [3, 4]. As described in section 2, in the time integration on a given discretization level the CFL-number is increased, aiming at fast convergence towards the stationary solution. In QUADFLOW the default strategy for determining this CFL-number γ_k in the k -th timestep is as follows: $\gamma_k = \max\{\gamma_0 \gamma^k, \gamma_{\max}\}$. In all experiments we set $\gamma_0 = 1, \gamma = 1.1$ (default values in QUADFLOW). We continue time integration on every discretization level until the residual of the density has been reduced by a factor 10^2 . On the finest discretization level we require a reduction by a factor 10^4 . The number of discretization levels used depends on the problem and on certain parameters used in the adaptive refinement strategy.

In a typical computation most time is spent on solving the linear equation systems on the grid that corresponds to the finest level of adaptation. Therefore we present the number of iterations of the preconditioned BiCGSTAB method that is needed to reduce the starting residual of the linear (Jacobian) system by a factor 10^4 on the finest grid in order to measure the quality of the renumberings. We compare four different numberings. The BW, HB and WRG methods have been explained above. The fourth numbering is the one induced by the discretization routines in QUADFLOW and is denoted by QN. One central feature of the QUADFLOW solver is the multiscale analysis that is used for error estimation and induces local refinement. This results in a hierarchy of locally refined grids, cf. section 2. In this process the cells are numbered levelwise from the coarsest to the finest level. This leads to a sort of hierarchical block-structure of the matrix. A typical pattern of the Jacobian is shown in figure 5.

After a prolongation to the next finer level in the nested iteration method we perform a renumbering after the first timestep. In each of the following timesteps we have a new Jacobian system to which a renumbering algorithm can be applied. For efficiency reasons we do *not* apply the renumbering method (steps 1,2,3) to every new Jacobian but use the known renumbering as computed in the first time step. We determine a new renumbering only after every k_r timesteps. Typical values for k_r are $k_r = 10, k_r = \infty$. All three numbering techniques are sensitive with respect to the choice of the value for the parameter τ . In our sub- and supersonic problems $\tau = 1.25$ turned out to be a good default value. In transonic problems the performance can often be improved by taking a somewhat large τ -value (e.g. $\tau = 2.00$).

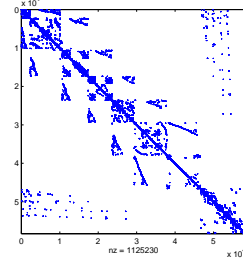


Figure 5: Test problem 2: Nonzero pattern of the matrix on the finest level (BAC 3-11/RES/30/21)

5.1 Test problem 1: Stationary flow around NACA0012 airfoil

The first problem is a standard test case for inviscid compressible flow solvers. We consider the inviscid, transonic stationary flow around the NACA0012 airfoil (cf. [23]). In this section we present some results for the following three test cases:

	M_∞	α
test case A	0.80	1.25°
test case B	0.95	0.00°
test case C	1.20	0.00°

Table 1: Test problem 1, cases A,B,C: parameters for NACA0012 airfoil

Results of a numerical simulation for case B are shown in figure 6. Renumbering is applied only once after every prolongation to the next finer discretization level ($k_r = \infty$). The maximum CFL-number was set to $\gamma_{\max} = 1000$. Computations are done as in [11]: We allow 8 maximum levels of refinement. In the cases A and C 10 cycles of adaptations are performed, 13 levels are used in case B.

Tables 5.1 - 5.1 show the average iteration count on the finest level for the different orderings. The average is taken over all timesteps used on the finest discretization level. The savings compared to the original QUADFLOW numbering QN are displayed in the last row. In all three cases the savings were not improved significantly when using smaller k_r values.

Numbering	QN	BW	HB	WRG
Average iteration count	32.0	30.6	28.6	23.0
Saving	0%	4.4%	10.6%	28.1%

Table 2: Case A, average iteration count on finest level (10th discretization level)

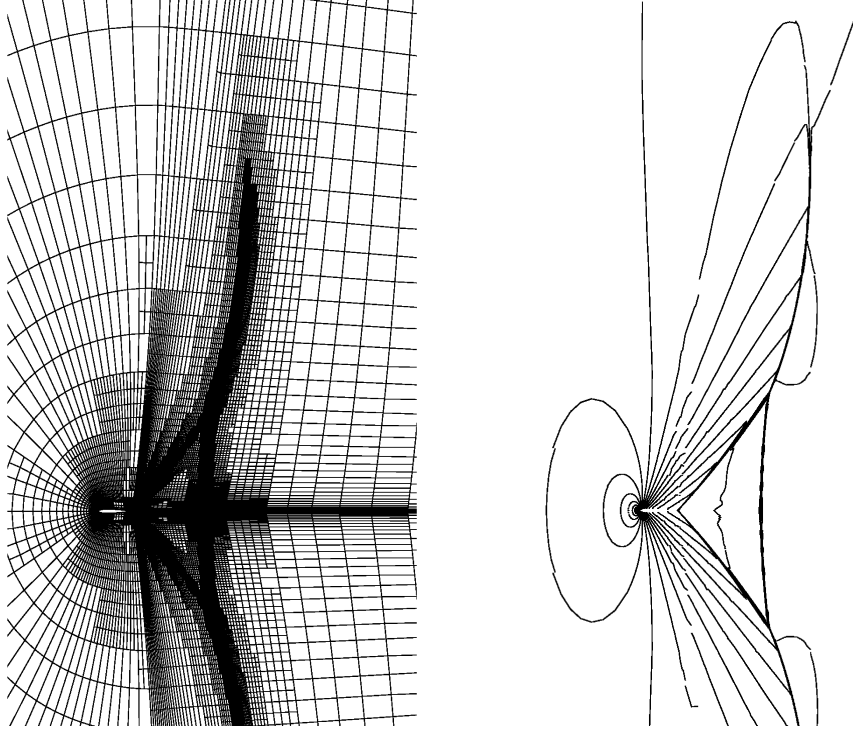


Figure 6: Case B: computational grid (left) and Mach distribution (right), $M_{min} = 0.0$, $M_{max} = 1.45$

Numbering	QN	BW	HB	WRG
Average iteration count	20.2	20.1	18.2	18.4
Saving	0%	0.5%	9.9%	8.9%

Table 3: Case B, average iteration count on finest level (13th discretization level)

Numbering	QN	BW	HB	WRG
Average iteration count	20.6	12.1	12.2	10.0
Saving	0%	41.1%	40.1%	51.5%

Table 4: Case C, average iteration count on finest level (10th discretization level)

In all cases the reduced matrix graph was constructed with $\tau = 1.25$. With the WRG renumbering method we save between 9% and 52% of PBGS-preconditioned BiCGSTAB iterations on the finest level compared to the original numbering QN. Since the renumbering has to be computed only once ($k_r = \infty$) the additional computational costs for WRG are negligible. The improvement is

strongest for case C, which is due to the fact that in this case the flow is almost supersonic and thus there is a main stream in which information is transported.

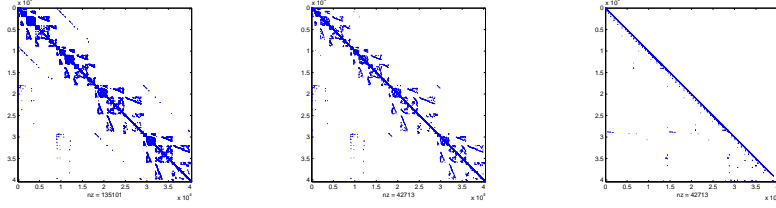


Figure 7: Test problem 1, case C: graph $\mathcal{G}(\mathbf{A})$, reduced graph $\hat{\mathcal{G}}(\mathbf{A})$ and renumbered reduced graph of Jacobian matrix on finest grid.

In case B the results for WRG numbering can be improved by a stronger reduction of the graph. With $\tau = 2.00$ the saving with WRG is about 21%. In this transonic case the pattern of directions in which information is propagated has a more complex structure than in the other cases. Therefore the savings are less than in the other examples. We want to point out that the ordering QN induced by the QUADFLOWS discretization routines is already quite good. If namely a (point-block) *random* numbering is used, then the PBGS preconditioned BiCGSTAB method turns out to diverge in most cases, even when computing supersonic flow.

In cases with higher CFL-numbers γ_{\max} the linear systems are in general harder to solve and the importance of an improvement due to a better numbering increases.

5.2 Test problem 2: Stationary flow around BAC 3-11/RES/30/21 airfoil

This test case is a standard cruise configuration [6] of the Collaborative Research Center SFB 401 [39] with $M_\infty=0.77$ and $\alpha=0.00^\circ$, see also [34]. In figure 1 we give a typical grid that is used in the simulation. We take parameter values $\tau = 1.25$, $\gamma_{\max} = 200$ and $k_r = 10$. For a typical Jacobian \mathbf{A} we show graph $\mathcal{G}(\mathbf{A})$, reduced graph $\hat{\mathcal{G}}(\mathbf{A})$ and the effect of the WRG renumbering in figure 8.

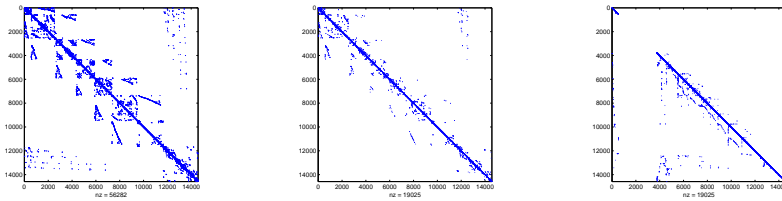


Figure 8: Test problem 2: graph $\mathcal{G}(\mathbf{A})$, reduced graph $\hat{\mathcal{G}}(\mathbf{A})$ and renumbered reduced graph of Jacobian matrix \mathbf{A} from figure 5.

The behavior of the preconditioned BiCGSTAB method is illustrated in figure 9. In this figure we give the number of iterations that the PBGS-preconditioned BiCGSTAB method needs to satisfy the stopping criterion for the linear solver in every timestep. We only give results for the timesteps after the last (10th) adaptation.

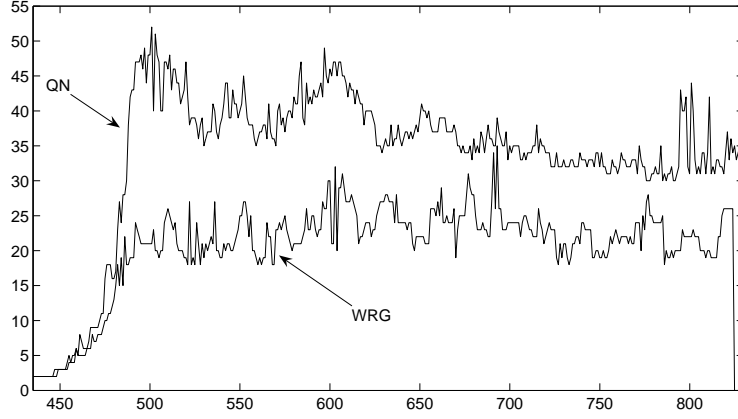


Figure 9: Test problem 2: number of PBGS-preconditioned BiCGSTAB-iterations in every timestep, timesteps on finest level.

There is a clear systematic improvement when using the WRG renumbering. The savings are about 38%. A comparison to the BW and HB renumbering methods is shown in table 5.2.

Numbering	QN	BW	HB	WRG
Average iteration count	33.5	22.2	22.2	20.6
Saving	0%	33.7%	33.6%	38.4%

Table 5: Test problem 2: average iteration count on finest level (10th discretization level)

5.3 Test problem 3: Stationary flow in oblique 3D-channel

In this problem we consider a flow through a 3D-channel with a bump at the bottom. Cross-sections of this channel with the x - y and x - z -plane are given in figure 10. The non-rectangular form is used to obtain a truly three-dimensional flow. Inflow and outflow conditions are prescribed at both ends of the channel. At inflow we take $M_\infty = 1.3$ and $\alpha = 0.00^\circ$.

The parameters in this test case are $\gamma_{\max} = 200$, $\tau = 1.25$ and $k_r = \infty$. Some results are presented in table 5.3. If instead of $\gamma_{\max} = 200$ we take $\gamma_{\max} = 1000$ then with the orderings resulting from QN, BW and HB the PBGS-preconditioned BiCGSTAB solver diverged in at least one timestep during the time integration on the finest discretization level. With WRG renumbering,

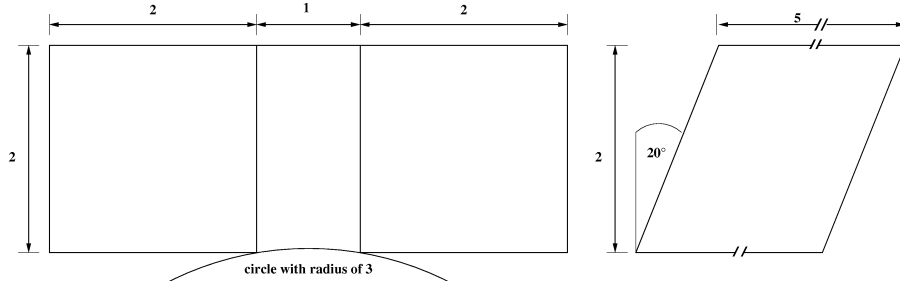


Figure 10: Test problem 3: Oblique channel with a bump. Left: $x-y$ plane. Right: $x-z$ plane

however, this was not the case. Thanks to the higher value $\gamma_{\max} = 1000$ we need about 16.1% of timesteps less than with $\gamma_{\max} = 200$. The average iteration count then is 13.9 for WRG. When summing up all Krylov-Iterations on the finest level, the total amount of iterations is 22.3% less than with QN numbering and $\gamma_{\max} = 200$.

Hence, this illustrates a further important advantage of the WRG renumbering, namely that it improves the *robustness* of the linear solver.

Numbering	QN	BW	HB	WRG
Average iteration count	15.0	14.3	14.1	12.9
Saving	0%	4.2%	6.0%	13.9%

Table 6: Test problem 3, average iteration count on finest level (after 4th adaptation)

6 Summary

Both the PBILU and PBGS methods are useful preconditioners in Newton-Krylov methods for compressible flow simulations. The behavior of these preconditioners depends on the ordering of the block-unknowns (cells). In this paper we present *ordering techniques for the PBGS method* that use ideas from algebraic multigrid methods. First a reduced weighted directed graph is constructed and then a renumbering of the vertices in this graph is determined. For this renumbering we use methods from the field of multigrid solvers for convection-dominated problems (BW and HB) and a modification of these (WRG). All three methods are implemented in QUADFLOW using the PETSc library. The reordering algorithm is *black-box*, except for the (critical) graph reduction parameter τ in (12). In most test cases a good choice for this grid-reduction parameter turns out to be $\tau = 1.25$. A systematic comparative study shows that for our problem class the WRG reordering yields the best results. Using this reordering we improve the *robustness* of the iterative solver. Even with large CFL-numbers (e.g. 200, 1000, 5000) the linear solver always converges if we use PBGS with WRG reordering, whereas with other orderings the solver

sometimes diverges. This implies that with WRG reordering it is possible to use larger CFL-numbers in order to reduce the total number of time steps. Using the reordering one can improve the *efficiency* of the linear solver significantly. The execution time of the iterative solver part can be reduced by 10% (for complex transonic flows) up to 50% (for supersonic flows). For efficiency reasons the reordering is not computed for each new Jacobian but kept fixed in a number of time steps.

The reordering algorithm can also be applied in the setting of (linear or nonlinear) multigrid solvers with block-Gauss-Seidel type smoothers for compressible flow problems.

Acknowledgment

The experiments in section 5 are done using the QUADFLOW solver developed in the Collaborative Research Center SFB 401. The authors acknowledge the fruitful collaboration with several members of the QUADFLOW research group.

References

- [1] K. Ajmani, W.-F. Ng and M. Liou, Preconditioned Conjugate Gradient Methods for the Navier-Stokes Equations, *Journal of Computational Physics*, 1994; **110**: 68–81.
- [2] E. F. D’Azevedo, P.A. Forsyth and W.-P. Tang, Ordering Methods for Preconditioned Conjugate Gradient Methods Applied to Unstructured Grid Problems, *SIAM Journal on Matrix Analysis and Applications*, 1992; **13(3)**: 944-961.
- [3] S. Balay, K. Buschelman, W. D. Gropp, D. Kaushik, M. Knepley, L. C. McInnes, B. F. Smith and H. Zhang, PETSc, <http://www-fp.mcs.anl.gov/petsc/>, 1992
- [4] S. Balay, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith and H. Zhang, PETSc Users Manual, ANL-95/11 - Revision 2.1.5, Argonne National Laboratory, 2004.
- [5] J. Ballmann (editor), Flow Modulation and Fluid-Structure-Interaction at Airplane Wings, *Numerical Notes on Fluid Mechanics*, Springer 2003; **84**.
- [6] J. Ballmann, Flow Modulation and Fluid-Structure-Interaction at Airplane Wings - Survey and Results of the Collaborative Research Center SFB 401. *DGLR 2002-009*, 2002.
- [7] P. Batten, M. A. Leschziner and U.C. Goldberg, Average-State Jacobians and Implicit Methods for Compressible Viscous and Turbulent Flows, *Journal of Computational Physics*, 1997; **137**: 38–78.
- [8] Bey, J. Wittum, Downwind numbering: robust multigrid for convection-diffusion problems, *Applied Numerical Mathematics*, 1997; **23**: 177-192.

- [9] K.H. Brakhage and S. Müller, Algebraic-hyperbolic Grid Generation with Precise Control of Intersection of Angles, *International Journal for Numerical Methods in Fluids*, 2000; **33**: 89–123.
- [10] F. Bramkamp, B. Gottschlich-Müller, M. Hesse, Ph. Lamby, S. Müller, J. Ballmann, K.-H. Brakhage, W. Dahmen, H-adaptive Multiscale Schemes for the Compressible Navier-Stokes Equations: Polyhedral Discretization, Data Compression and Mesh Generation, 2001; in [5], 125–204.
- [11] F. Bramkamp, Ph. Lamby and S. Müller, An adaptive multiscale finite volume solver for unsteady and steady state flow computations, *Journal of Computational Physics*, 2004; **197/2** 460–490.
- [12] A. Brandt, Multi-level Adaptive Solutions to Boundary Value Problems, *Mathematics of Computation*, 1997; **31**: 333–390.
- [13] E. Cuthill, Several strategies for reducing the band width of matrices. *Sparse Matrices and Their Applications*, D. J. Rose and R. A. Willoughby, eds., New York, 1997; 157–166.
- [14] E. Cuthill, J. McKee, Reducing the bandwidth of sparse symmetric matrices. in: Proc. ACM Nat. Conf., New York, 1969; 157–172.
- [15] J. Edwards and M.S.Liou, Low-Diffusion Flux-Splitting Methods for Flows at All Speeds. *AIAA Journal*, 1993; **36(9)**: 457–497.
- [16] S. Gutsch., T. Probst, Cyclic and feedback vertex set ordering for the 2D convection-diffusion equation. *Technical Report*, Universit Kiel, 1997; **97-22**.
- [17] W. Hackbusch, Multi-grid Methods and Applications, Springer 1985.
- [18] W. Hackbusch, On the Feedback Vertex Set for a Planar Graph, *Computing*, 1997; **58**: 129–155.
- [19] W. Hackbusch and T.Probst, Downwind Gauß-Seidel Smoothing for Convection Dominated Problems, *Numerical Linear Algebra With Applications*, 1997; **4(2)**: 85–102.
- [20] D. Hänel and F. Schwane, An Implicit Flux-Vector Splitting Scheme for the Computation of Viscous Hypersonic Flow, *AIAA paper*, 1989; **0274**.
- [21] A. Jameson, Solution of the Euler Equations for Two-Dimensional Transonic Flow by a Multigrid Method, *Applied Mathematics and Computation*, 1983; **13**: 327–356.
- [22] A. Jameson, D.A. Caughey, How Many Steps are Required to Solve the Euler Equations of Steady, Compressible Flow: in search of a Fast Solution Algorithm, *AIAA paper*, 2001; **2673**.
- [23] D. J. Jones, Reference Test Cases and Contributors, Test Cases For Inviscid Flow Field Methods. *AGARD Advisory Report*, 1986; **211(5)**.

- [24] F. Kicking, Algebraic Multigrid for Discrete Elliptic Second-Order Problems, Multigrid Methods V. Proceedings Of The 5th European Multigrid Conference (W.Hackbusch ed.), *Springer Lecture Notes in Computational Science and Engineering*, 1998; **3**: 157-172.
- [25] B. van Leer, Flux Vector Splitting for the Euler Equations. In: Proceedings of the 8th International Conference on Numerical Methods in Fluid Dynamics (E. Krause, ed.). *Lecture Notes in Physics*, Springer, Berlin, 1982; **170**: 507-512.
- [26] B. van Leer and D. Darmofal, Steady Euler Solutions in $O(N)$ Operations, Multigrid Methods (E. Dick, K. Rienslagh and J. Vierendeels, editors), 1999; **VI**: 24-33.
- [27] I. Lepot, P. Geuzaine, F. Meers, J.-A. Essers and J.-M. Vaassen, Analysis of Several Multigrid Implicit Algorithms for the Solution of the Euler Equations on Unstructured Meshes, Multigrid Methods (E. Dick, K. Rienslagh and J. Vierendeels, editors), 1999; **VI**: 157-163.
- [28] H. Luo, J. Baum and R. Löhner, A Fast, Matrix-free Implicit Method for Compressible Flows on Unstructured Grids, *Journal of Computational Physics*, 1998; **146**: 664-690.
- [29] D.J. Mavriplis and V. Venkatakrishnan, Implicit method for the Computation of Unsteady Flows on Unstructured Grids, *Journal of Computational Physics*, 1996; **127**: 380-397.
- [30] P.R. McHugh and D.A. Knoll, Comparison of Standard and Matrix-Free Implementations of Several Newton-Krylov Solvers, *AIAA Journal*, 1994; **32**: 394-400.
- [31] A. Meister, Comparison of different Krylov Subspace Methods Embedded in an Implicit Finite Volume Scheme for the Computation of Viscous and Inviscid Flow Fields on Unstructured Grids, *Journal of Computational Physics*, 1998; **140**: 311-345.
- [32] A. Meister, Th. Sonar, Finite-Volume Schemes for Compressible Fluid Flow. *Surveys on Mathematics for Industry*, 1998; **8**: 1-36.
- [33] A. Meister, C. Vömel, Efficient Preconditioning of Linear Systems Arising from the Discretization of Hyperbolic Conservation Laws, *Advances in Computational Mathematics*, 2001; **14**: 49-73.
- [34] I.R.M. Moir, Measurements on a Two-Dimensional Aerofoil with High-Lift-Devices, volume 1 and 2 *AGARD Advisory Report 303*: A Selection of Experimental Test Cases for the Validation of CFD Codes, 1994.
- [35] A. Reusken, On the Approximate Cyclic Reduction preconditioner, *SIAM J. Scientific Comp.*, 2000; **21**: 565-590.
- [36] H. Rentz-Reichert and G. Wittum, A comparison of smoothers and numbering strategies for laminar flow around cylinder, in E. Hirschel, ed., Flow Simulation with High-Performance Computers II, *Notes on Numerical Fluid Mechanics*, Vieweg 1996; **52**: 134-149.

- [37] Y. Saad, Iterative methods for sparse linear systems, PWS Publishing Company, Boston (1996).
- [38] Y. Saad, Preconditioned Krylov Subspace Methods for CFD Applications, in: *Solution techniques for Large-Scale CFD-Problems*, ed. W.G. Habashi, Wiley 1995; 139–158.
- [39] SFB 401, Collaborative Research Center, Modulation of flow and fluid-structure interaction at airplane wings, RWTH Aachen University of Technology, <http://www.lufmech.rwth-aachen.de/sfb401/kufa-e.html>
- [40] K. Stüben, An Introduction in Algebraic Multigrid, Appendix A in: U. Trottenberg, C. W. Oosterlee, A. Schller, Multigrid, Academic Press, GMD Birlinghoven, St.Augustin, 2001.
- [41] E.F. Toro, M.Spruce and W.Speares, Restoration of the Contact Surface in the HLL Riemann Solver. *Shock Waves*, 1994; **4**: 25–34.
- [42] S. Turek, On ordering strategies in a multigrid algorithm. In *Notes on Numerical Fluid Mechanics* **41**, *Proceedings 8th GAMM-Seminar*, Vieweg 1992.
- [43] K.J. Vanden and P. D. Orkwis, Comparison of Numerical and Analytical Jacobians, *AIAA Journal*, 1996; **34(6)**: 1125–1129.
- [44] V. Venkatakrishnan, V., Convergence to Steady State Solutions of the Euler Equations on Unstructured Grids with Limiters, *Journal of Computational Physics*, 1995; **118**: 120–130.
- [45] V. Venkatakrishnan, Implicit schemes and Parallel Computing in Unstructured Grid CFD, *ICASE-report*, 1995; **28**.
- [46] Y. Wada and M.S.Liou A Flux Splitting Scheme with High-Resolution and Robustness for Discontinuities, *AIAA Paper*, 1994; **94-0083**.