


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Journal of Computational Physics

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

An algebraic multigrid solver for transonic flow problems

Shlomy Shitrit^{a,*}, David Sidilkover^a, Alexander Gelfgat^b^a Soreq NRC, Propulsion Physics Laboratory, Yavne 81800, Israel^b School of Mechanical Engineering, Faculty of Engineering, Tel-Aviv University, Israel

ARTICLE INFO

Article history:

Received 29 July 2010

Received in revised form 17 November 2010

Accepted 22 November 2010

Available online 1 December 2010

Keywords:

Algebraic multigrid (AMG)

Transonic flow

Full potential equation

ABSTRACT

This article presents the latest developments of an algebraic multigrid (AMG) based on full potential equation (FPE) solver for transonic flow problems with emphasis on advanced applications. The mathematical difficulties of the problem are associated with the fact that the governing equation changes its type from elliptic (subsonic flow) to hyperbolic (supersonic flow). The flow solver is capable of dealing with flows from subsonic to transonic and supersonic conditions and is based on structured body-fitted grids approach for treating complex geometries. The computational method was demonstrated on a variety of problems to be capable of predicting the shock formation and achieving residual reduction of roughly an order of magnitude per cycle both for elliptic and hyperbolic problems, through the entire range of flow regimes, independent of the problem size (resolution).

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

One of the major early breakthroughs in potential flow computation was the work by Murman and Cole on numerical solution of the small disturbances equation for transonic flow [1]. Their paper laid the ground work for the years that follow. Research on potential flow thrived throughout the 1970s and into the early 1980s. It was a period of rapid development and various improvements of potential flow solvers, by numerous researchers. This direction was abandoned, however, while still being in its infancy in favor of the Euler and Navier–Stokes solvers. Development of such solvers together with advancements in the computer hardware made solving the Navier–Stokes equations practical even for rather complex configurations. It looks tough, for some time (more than a decade) that the further development of this methodology, especially with regard to efficiency improvement reached a certain stagnation. A substantial departure from the existing methodology may be required in order to facilitate further progress. One of such possible directions originates from the recommendation by Brandt [2], that since the system of equations is of the mixed type, it is beneficial in terms of efficiency to address each of the co-factors separately instead of treating the whole system in the same way. This idea was successfully realized in the past for the incompressible high Reynolds number flow equations (see for instance [3]), but the progress towards applying it to the compressible flow was rather slow and the success is very limited. The explanation for this is in the complexity of the issues that need to be resolved. One such difficulty is that the standard discretization schemes in multidimensions introduce non-physical coupling between the different co-factors of the system. This difficulty is addressed by the emerging class of the so-called *factorizable* methods [4]. In this light the task of constructing an efficient FPE solver attains a great importance, since such a solver can be used not only by itself, but becomes an integral part of the overall methodology for solving the flow equations based upon the factorizable discretization. Potential flow analysis still plays an important role in aircraft design and optimization because of its simplicity and efficiency [5].

* Corresponding author. Tel.: +972 8 9434853; fax: +972 8 9434227.

E-mail address: shlomy_shitrit@yahoo.com (S. Shitrit).

One of the first multilevel methods toward solving partial differential equations fast and efficiently was the multigrid method [6,7] which was first introduced by Fedorenko [8] in 1964. Other mathematicians extended Fedorenko's idea to general elliptic boundary value problems with variable coefficients; see e.g., [9]. However, the full efficiency of the multigrid approach was realized in the works of Brandt [10,11]. He also introduced the modification of multigrid methods for nonlinear problems – Full Approximation Scheme (FAS) [11,12]. Another achievement in the formulation of multigrid methods was the full multigrid (FMG) scheme [11,12], based on the combination of nested iteration techniques and multigrid methods.

In geometric “standard” multigrid methods the coarse-grids are uniformly coarsened or semi-coarsened, thus the freedom in the selection of the coarse-grids is limited. The grids hierarchy is constructed based on the grid geometry information rather than properties of the differential operator. In addition, the definition of smoothness of the error involves grid geometry. This geometric dependency imposes certain limits on types of the problems it can be used to solve. For problems with general anisotropies (such anisotropies may occur not only as a property of the operator itself but also as a result of a grid configuration) a construction of an efficient geometric multigrid method may become rather cumbersome task. In the 1980s algebraic multigrid (AMG) methods were developed [13–16] to address these issues by extending the main ideas of geometric multigrid methods to an algebraic setting. AMG is a method for solving algebraic systems based on multigrid principles with no explicit reliance upon the grid geometry. It uses the matrix's properties only to construct the operators involved in the algorithm. The AMG framework usually employs a simple pointwise relaxation method whose role is to smooth (in the algebraic sense) error and then attempts to correct the algebraically smooth error that remains after relaxation by a coarse-level correction.

The purpose of this research is to develop a FPE solver which is based on the AMG method. The previous results were reported in [17] while using the same mathematical basis for solving elementary applications (such as flow around a cylinder, channel with a bump). The emphasis of this paper is on the more advanced applications. The practical goal of this is twofold: First, to develop an important building block for the factorizable methodology. Second, to develop a stand-alone “optimally” efficient FPE solver. The flow solver is to be capable to deal with flow ranging from subsonic to transonic and supersonic regimes. It can be especially useful for engineers during the design process where multiple computations need to be performed as small changes to the geometry are made. The paper is organized as follows: The discretization of the FPE presented in Section 2. The extensions of the AMG method to the transonic flow computations are discussed in Section 3. In Section 4, several applications for body-fitted structured grids are presented and convergence results for various flow speeds are given.

2. The FPE discretization

Transonic flow can be described by the full potential equation (FPE) which is derived from the Euler equations by assuming that the flow is inviscid, isentropic, and irrotational. The FPE in the conservation form reads as follows:

$$\frac{\partial}{\partial x}(\rho u) + \frac{\partial}{\partial y}(\rho v) = 0, \tag{1}$$

where u and v are the velocity components in the Cartesian coordinates x and y , respectively, and ρ is the density. The velocity components are given by the gradient of the potential ϕ ,

$$u = \frac{\partial \phi}{\partial x}, \quad v = \frac{\partial \phi}{\partial y}. \tag{2}$$

The density ρ is computed from the isentropic formula:

$$\frac{\rho}{\rho_\infty} = \left(1 + \frac{\gamma - 1}{2} (V_\infty^2 - \phi_x^2 - \phi_y^2) \right)^{\frac{1}{\gamma - 1}}, \tag{3}$$

where γ is the ratio of specific heats and V_∞ is the free-stream velocity and ρ_∞ is the free-stream density. The relation between the local speed of sound a and the flow speed is defined by Bernoulli's equation:

$$a = \left(a_\infty^2 - \frac{\gamma - 1}{2} (V_\infty^2 - \phi_x^2 - \phi_y^2) \right). \tag{4}$$

The discretization of the FPE in the conservation form is based on the same rationale that was applied to the quasi-linear form of the equation. For this purpose we address the reader to our previous work [18] while in this section we shall only briefly review the general idea.

The strategy of discretizing the FPE in the conservation form is based upon an idea similar to that of the rotated difference approach introduced by Jameson [19] and implemented initially in the quasi-linear form. However, this approach is not made directly. Instead it is accomplished indirectly by following the same rationale.

We review briefly our approach starting with the FPE in the quasi-linear form,

$$\nabla^2 \phi - M_\infty^2 \frac{\partial^2 \phi}{\partial s^2} = 0, \tag{5}$$

where M_∞ is free-stream Mach number ($M_\infty = \frac{V_\infty}{a_\infty}$) and s -axis denotes the flow direction. Let us look at both the terms from a numerical standpoint. Note, that when M_∞ is close to zero (incompressible flow) the second term can be neglected and we are left with Laplacian $\nabla^2\phi$ which is discretized by certain type of central differencing, according to [18]. As the Mach number increases, the second term, which describes the potential second derivative in the streamwise direction, actually determines the “dynamics” of the flow. When the flow is subsonic, a central difference is used for the second derivative in the flow direction (ϕ_{ss}), while pointwise relaxation is applied directly in conjunction with the matrix A (or operator L). If the flow is supersonic, we apply relaxation directly with the product of downwind and upwind operators $\tilde{L}\tilde{L}$ (or matrix $A\tilde{A}$), as is detailed in our previous work [18]. We would like to apply the same rationale to the discretization of the FPE in the conservation form, while the advantages that were obtained in the quasi-linear case discretization, would be implemented. Expanding the FPE and rearranging terms, (5) can be reformulated as,

$$\rho\nabla^2\phi + \left(\phi_x\frac{\partial}{\partial x} + \phi_y\frac{\partial}{\partial y}\right)\rho = 0, \quad (6)$$

where the density ρ is given in (3). Note that (5) and (6) have a similar structure. The density parameter ρ plays two functions. In the first term it serves as a coefficient. In the second term it serves as a variable. As one can see, the second terms in both of the above equations are identical. For the first term $\nabla^2\phi$, the fluxes are computed by a central discretization independent on the flow's speed. The dynamics of the flow is reflected in the second term which is discretized in such a way that the results is a “wide” approximation in the streamwise direction. The flux at given face is a product of the flow velocity vector and the density which are both functions of the potential ϕ . The type of the discretization approach, central fluxes or upwind fluxes, is determined by the local Mach number across the cell's face. For further details we refer the reader to [20].

3. Classical AMG

In this paper we assume the reader to have some basic knowledge of the “traditional” geometric multigrid. He should be familiar with smoothing and coarse-grid correction and with the recursive definition of multigrid cycles. We limit our discussion to the basic principles of the AMG method and shall give here a brief description of the classical AMG in the spirit of Stueben [15] followed by a description of the solution process. Regarding more detailed information on geometric and algebraic multigrid relevant to this work, we also refer to [21,18] and the extensive list of references given therein. We address the reader to [15,22,23] for a detailed description of the AMG algorithm, while in this paper we shall only briefly review the algorithm emphasizing its aspects that require special treatment for the purpose of this work.

Consider a certain boundary-value problem for a scalar PDE in domain Ω . Its discretization will result in a linear algebra problem of the form $Au = f$, where A is an $n \times n$ matrix with entries a_{ij} with $i = 1, \dots, n, j = 1, \dots, n, u = \{u_j\}$ is the vector of unknowns, $b = \{b_j\}$ is the forcing term vector and n is the number of points in the computational grid covering the domain.

In geometric multigrid the definition of smoothness of the error involves grid geometry. The absence of grids in the AMG context renders this definition meaningless. Therefore, the concept of smoothness has to be generalized to some meaningful measurable quantity, which can be computed based on the discrete operator only. An error is algebraically smooth if it is slow to converge with respect to the fine level. Therefore, this smooth error has to be approximated on a coarser level in order to improve the convergence. Consequently, the coarse-grid correction must deal with the remaining slow components. The characterization of such slow components, e , is: $Ae \simeq 0$.

3.1. Smoothing

Usually, the relaxation used as an ingredient of an AMG algorithm is a pointwise one. One of the central contributions of our approach is developing a stable pointwise direction-independent relaxation for the entire range of the flow speed, from low Mach number flow up until transonic and supersonic regimes. This development was a prerequisite for considering application of AMG to the transonic flow problem. For further details we refer the reader to [18].

3.2. Measuring the algorithm efficiency

The computational complexity concept is intended to measure the algorithm's requirements for computer resources: computer storage (memory), and CPU time. There are three types of complexity measure that are commonly considered: convergence rate, grid complexity, and operator complexity.

Convergence factor: the rate in which the residual is decreased between consecutive V-cycles. This parameter gives an indication of how many iterations are needed in order to reduce the residual to a sufficient level.

Grid complexity: is the total number of elements in the coarse-levels divided by the number of elements on the fine-level. The grid complexity provides a direct measure for the storage required for the solution and right hand side vectors and it is a useful tool to compare different coarsening strategies.

Operator complexity: is defined as the sum of the number of nonzero matrix elements in all the coarse-levels, divided by the number of the nonzero matrix elements on the fine-level. The amount of work required by the relaxation and residual computations is directly proportional to the number of the nonzeros on the coarse-levels. Therefore, low value of the operator complexity that increase linearly with the problem's resolution signifies a linear complexity operator.

4. Improved AMG components for transonic flow computations

4.1. Dynamic threshold

In the context of an algebraic multigrid we are going to deal at each level with a linear system of equations

$$A^m u^m = f^m, \tag{7}$$

where m is the level index. The coarsening process is derived based on the strong and weak connections between unknowns, which essentially measures the relative size of the off-diagonal entries. Connections between neighboring variables are considered strong if the size of the corresponding matrix entry exceeds a certain threshold, relative to the maximum entry of the row. This threshold value is very important for constructing a good coarse-grid. According to Ref. [15] a point i is said to be strongly connected to point j , if

$$-a_{ij}^m \geq \varepsilon \max_{k \neq i} (-a_{ik}^m). \tag{8}$$

The threshold value ε is kept fixed for most applications, with a typical value of 0.25. It is well-known that devising a robust coarsening process, which results in an accurate interpolation, is one of the keys for achieving a good AMG convergence. There are special cases, for example, anisotropic elliptic problems (anisotropy can be a result of the computational grid or the equation itself), or problems where the equation changes type from elliptic to hyperbolic (transonic flow problem) that a fixed threshold parameter can result in an inadequate coarsening process.

A problem under consideration (supersonic case, for instance) usually leads to a matrix with significant negative off-diagonal entries. Therefore, it is important to redefine the definition of strong and weak connections, so that it becomes more adequate for our case. We would like to allow a connection with a negative coefficient, whose absolute value is sufficiently large, to be classified as a strong one. Therefore, as it was suggested in [24], we modify the definition (8): a point i is considered to be strongly connected to j if

$$|a_{ij}^m| \geq \varepsilon \max_{k \neq i} |a_{ik}^m|. \tag{9}$$

Our further generalization of this idea is to select strong connections using a different threshold parameter for each row of the matrix A^m , while computing its value by analyzing this row's entries. Each row represents a difference operator at a specific grid point and, possibly, boundary conditions. Intuitively speaking, the same fixed threshold value (as used in the classical AMG) cannot reflect the relative size of the row's entries for various cases (isotropic, strongly anisotropic, or hyperbolic). Therefore, in order to maintain a good AMG performance across the entire variety of flow speeds, we propose to compute the threshold adaptively "on the fly" for each row of the matrix.

Let us define a set of points which are strongly connected to i by S_i . Let $N_i^m = \{j \in \Omega^m : j \neq i, a_{ij}^m \neq 0\}$ denote the neighborhood of a point $i \in \Omega^m$. $D_i = N_i^m - C_i$, $D_i^s = D_i \cap S_i$, $D_i^w = D_i - D_i^s$. The neighboring coarse grid points are denoted by C_i , and the computational domain is denoted by Ω . Assuming the operator A^m is known, we start from the following equation:

$$a_{ii}^m u_i^m + \sum_{j \in N_i^m} a_{ij}^m u_j^m = f_i^m \quad \forall i \in \Omega^m. \tag{10}$$

Let us define the threshold parameter ε_i for each row i as

$$\varepsilon_i = \frac{\sum_{j \neq i} |a_{ij}| \frac{|a_{ij}|}{\max_{k \neq i} |a_{ik}|}}{\sum_{j \neq i} |a_{ij}|}. \tag{11}$$

As one can see, the threshold parameter is simply the average value in a given row. Several observations relating to the constant-coefficient case are in order here. If a periodic boundary condition is applied, the matrix A^m is composed of identical rows and the threshold parameter can be calculated once for each coarse-level. But, while applying other boundary conditions (Dirichlet, Neumann), as for a general problem, it is necessary to calculate the threshold for each row separately.

4.2. The coarse-level and restriction operators

In the classical approach, suggested by Ruge and Stueben, the restriction operator is defined as the transpose of the interpolation, $I_m^{m+1} = [I_{m+1}^m]^T$. Then the coarse-grid operator is defined by the Galerkin-type algorithm, $A^{m+1} = I_m^{m+1} A^m I_{m+1}^m$. We shall refer, following [24], to this approach as Algorithm 1.

Although this is the simplest way to construct the restriction and coarse-grid operators, it leads to poor convergence when the matrix A^m is not an M-matrix.

Alternatively, a second algorithm discussed in [24] is to use direct approximations based on the fine grid operator A^m to construct A^{m+1} and I_m^{m+1} . Assuming the operator A^m is known, we start from the following equation:

$$a_{ii}^m u_i^m + \sum_{j \in C_i^m} a_{ij}^m u_j^m + \sum_{j \in D_i^m} a_{ij}^m u_j^m \approx f_i^m \quad i \in C^m. \tag{12}$$

In order to derive the coarse-grid operator A^{m+1} , the terms associated with $u_j^m, j \in D^m$, in the i th equation, $i \in C^m$, should be approximated. The simplest way to account for these elements is to distribute their values among the entries corresponding to strong connections in this row. The resulting coarse-level operator may not always provide in this case a sufficiently accurate correction to a fine level solution that in its turn, would lead to poor convergence rates. Alternatively the terms $u_j^m, j \in D^m$, in the i th equation can be replaced by the j th equation. This operation does not eliminate all $u_j^m, j \in D^m$, but reduces magnitudes of their coefficients. Then $u_j^m, j \in D^m$ are eliminated by the interpolation formula if $|a_{ij}^m/a_{ii}^m| \leq \varepsilon_c$, or, again, are replaced by the means of the j th equation when $|a_{ij}^m/a_{ii}^m| > \varepsilon_c$, where $\varepsilon_c > 0$ is a small number, i.e.,

$$u_j^m = \begin{cases} \sum_{k \in C_j^m} w_{jk} u_k^m, & \text{if } |a_{ij}^m/a_{ii}^m| \leq \varepsilon_c \\ \frac{f_j^m}{a_{jj}^m} - \frac{1}{a_{jj}^m} \sum_{k \in N_j^m} a_{jk}^m u_k^m, & \text{if } |a_{ij}^m/a_{ii}^m| > \varepsilon_c \end{cases}. \tag{13}$$

The new equations are obtained

$$a_{ii}^{(2)} u_i^m + \sum_{j \in C_i^m} a_{ij}^{(2)} u_j^m + \sum_{j \in D_i^m} a_{ij}^{(2)} u_j^m \approx f_i^m - \sum_{j \in D_i^m} \frac{a_{ij}^m}{a_{jj}^m} f_j^m, \quad i \in C^m, \tag{14}$$

where $D_i^m = \{j : |a_{ij}^m/a_{ii}^m| > \varepsilon_c, j \in D^m\}$. This process is repeated until $u_j^m, j \in D^m$ is eliminated from the i th equation. Therefore,

$$a_{ii}^{(L)} u_i^m + \sum_{j \in C_i^m} a_{ij}^{(L)} u_j^m \approx f_i^m - \sum_{j \in D_i^m} \frac{a_{ij}^m}{a_{jj}^m} f_j^m, \quad i \in C^m, \tag{15}$$

where $a_{ij}^{(1)} = a_{ij}^m, D_i^m = \{j : |a_{ij}^{(l)}/a_{ii}^{(l)}| > \varepsilon_c, j \in D^m\}$. The coarse-grid operator is defined as $A^{m+1} = (a_{ij}^{(L)})$. Note that using this method we get $I_m^{m+1} \neq (I_{m+1}^m)^T$ (see [24] for more details). This approach works well for the problems considered in this paper.

4.3. The interpolation operator

Our formulation of the interpolation operator is based on Stueben's approach [15], which was demonstrated to be efficient for the M-matrices. The construction of the interpolation operator is identical to the standard interpolation except for the definition of weak/strong connections. Namely, the modified criteria are based upon comparing the absolute values of the matrix entries and not on their values. This interpolation formula is more accurate (especially in the supersonic flow regime) than that used in the standard AMG method because of the following: unlike the standard approach, it allows sufficiently large negative off-diagonal entries also to be considered as strongly connected points. Our modified interpolation works well for the entire range of the flow speed, from low Mach number flow up until transonic and supersonic regimes.

4.4. The full multigrid (FMG) method in the context of AMG

When a nonlinear problem is involved, it can happen that many AMG cycles are “wasted” in order to reach a suitable approximation when the initial condition is far from the exact solution. This situation can be avoided by applying the FMG method. The original geometric FMG is based on geometric multigrid cycles in order to improve the initial guess on the fine-level. We suggest to replace the MG V-cycles with an AMG V-cycles.

Fig. 1 shows the schedule of the grids in the order in which they are visited. The coarse “base” grids of the FMG sequence are based on the classical coarsening strategies – doubling the mesh size in each direction, that is, by $h \rightarrow 2h$ coarsening.




Fig. 1. Schedule of grids for the improved FMG scheme on four levels. The bold arrows indicates the AMG V-cycle. The thin arrows present an interpolation of the approximation to the next finer level.

Practically, we solve the problem first on a coarse-grid and then interpolate the solution to the finer level. The interpolated solution serves as an initial approximation for the AMG cycle on the finer level. Formally, such an FMG algorithm, based on four levels, can be described as follows:

1. Initialize $I_n^{2h} f^h \rightarrow f^{2h}$, $I_{2h}^{4h} f^{2h} \rightarrow f^{4h}$, ...
2. Solve the problem on coarsest grid $A^{8h} v^{8h} = f^{8h}$.
3. $I_{8h}^{4h} v^{8h} \rightarrow v^{4h}$.
4. $V_{AMG}^{4h}(v^{4h}, f^{4h}) \rightarrow v^{4h}$, v_0 times.
5. $I_{4h}^{2h} v^{4h} \rightarrow v^{2h}$.
6. $V_{AMG}^{2h}(v^{2h}, f^{2h}) \rightarrow v^{2h}$, v_0 times.
7. $I_{2h}^h v^{2h} \rightarrow v^h$.
8. $V_{AMG}^h(v^h, f^h) \rightarrow v^h$, v_0 times.

Instead of transferring the f^h to the coarse-level by restriction, the original right hand side f is used for the coarsest level. The cycling parameter v_0 sets the number of AMG V-cycles done at each level. Practical experience has shown that v_0 depends on the problem. Usually for elliptic problems $v_0 = 1$ is sufficient to produce good initial guess for the fine-level.

It is important to mention that this simple FMG algorithm seems still far from being optimal. However, it already appears to be very instrumental in solving nonlinear problems. We believe there is still much room for optimization, albeit, the results indicate that the algorithm is very efficient even as is.

5. Applications and performance

Several two dimensional flow calculations have been performed to test the performance of the algebraic multigrid method implemented within the body-fitted structured grid and finite volume context. The tests cases were chosen to address two major requirements: First, the flow model problem has to agree the potential flow limitations. Second, it should allow to examine the capability of the code to deal with irregular structured grids together with an equation that becomes extremely anisotropic near the sonic case and changes type to hyperbolic in the supersonic flow regime. Two dimensional solutions will be given for the following problems: transonic diffuser, rocket chamber, NACA-0012 airfoil, and nonsymmetric NACA-2822 airfoil.

The problems were tested in subsonic and transonic flow regimes in different resolutions. All computational data in the subsonic flow regime is two order accurate in space and the supersonic region is first order accurate. We consider several measures of the effectiveness of the algorithm. Our focus is on solving the FPE in each case to a high accuracy, namely by reducing the residual by at least ten orders of magnitude. Another relevant measure is the asymptotic convergence factor of the resulting cycle. A lower factor requires fewer iterations in the solution phase. For each V-cycle the columns denoted by C_f reports the convergence factor. In the row denoted by C_Ω , the grid complexity is reported, and the operator complexity is reported as C_L . The following default settings were used throughout the calculations, unless explicitly stated otherwise: The second pass process [22] was applied only on the fine-level in order to satisfy the interpolation requirements. Strong connectivity is defined by a fixed threshold $\varepsilon = 0.25$. The dynamic threshold was applied only where the fixed threshold fails. A symmetric Gauss–Seidel, two pre and two post smoothing steps being the default.

5.1. Transonic diffuser

Converging–diverging diffuser is an important and basic component associated with propulsion and the high speed flow of gases. This application often places high demands on CFD codes. For example, one purpose of the diffuser is to decelerate the flow ahead of the engine. Another purpose of the diffuser is to connect the inlet with the engine. In many cases, the inlet and engine axes are offset, and the diffuser must turn the flow [25]. In this exercise, both subsonic and transonic flow through a converging–diverging diffuser are investigated. Mach number variation and shock formation may be examined. This flow problem can be found practically in any gas dynamics textbook, for example [26,25].

5.1.1. Problem definition and boundary condition

The diffuser has a rectangular section. It has a flat bottom wall and a converging–diverging channel with a maximum divergence angle of 10° at the top wall. The ratio of the inlet area to the throat area is $\frac{A_{in}}{A_{throat}} = 1.4114$ and the ratio of the exit area and the throat area is $\frac{A_{out}}{A_{throat}} = 1.5$. This problem was solved on grids with three different levels of resolutions in order to check the AMG's scalability. These grids are as follows: 128×48 , 64×24 , and 32×12 , while the finer mesh is presented in Fig. 2. The grid is clustered close to upper and lower walls vertically and in the throat horizontally, with a stretching factor of 1.1.

The following boundary conditions were specified:

- A subsonic flow at the inlet and the same mass flow rate at the outlet, in the x -direction.
- A solid-wall boundary condition at the top and bottom of the diffuser.




Fig. 2. Mesh used for the flow through the transonic diffuser.

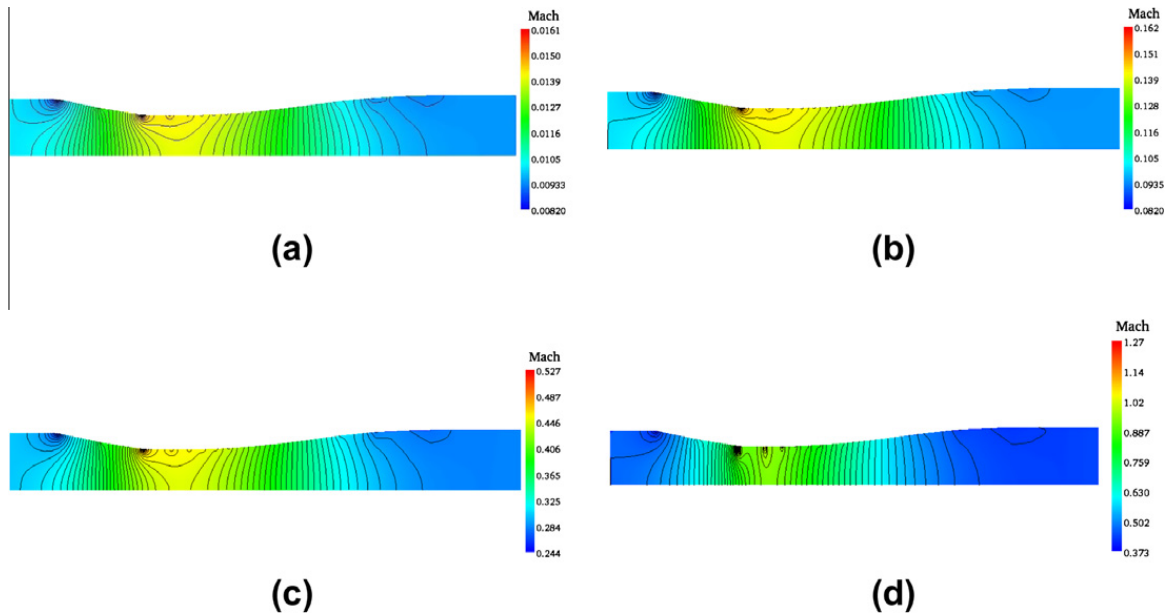


Fig. 3. Mach number isolines as computed on the (128×48) mesh, for an inlet Mach number of: (a) $M_\infty = 0.01$, (b) $M_\infty = 0.1$, (c) $M_\infty = 0.3$, (d) $M_\infty = 0.46$. Fifty contours were sketched between the maximum and minimum presented Mach numbers. Observe the shock wave in the fourth case.

5.1.2. Qualitative results

Fig. 3(a) shows the flow through the diffuser when it is completely subsonic. The flow accelerates out of the chamber through the converging section, reaching its maximum speed at the throat. The flow then decelerates through the diverging section. Fig. 3(d) presents the case where the Mach number at the inlet is increased to $M_\infty = 0.46$. The flow at the throat reaches the sonic speed (shocked throat). Unlike the subsonic flow, the supersonic flow accelerates as the cross-section area is increased. The region of the supersonic acceleration is terminated by a normal shock wave, as can be seen in Fig. 3(d).

5.1.3. AMG performance

Table 1 presents the complexities for each case of Mach number described above. The residual norm decreased rapidly for eight to ten cycles, while in the last few cycles a constant value was reached (asymptotic convergence rate). In the first two cases of $M_\infty = 0.01$, $M_\infty = 0.1$ the setup phase was performed only once, followed by eight cycles. Solving the problem on reduced resolutions, the structure of the fine-scale problem and the convergence factor remain very similar and so does the convergence factor, verifying that the solver's efficiency is independent on the problem size. The case of $M_\infty = 0.3$, where the compressibility effects (together with the nonlinearity) became significant, the results were obtained by repeating the setup phase six times while applying two V-cycles between each update. It results in a total of eight V-cycles until the residual decreased to the desired level of $\sim 10^{-10}$. The transonic case, $M_\infty = 0.46$, provides a more difficult test for the algorithm. In this case the setup phase was repeated six times resulting in 12 V-cycles. The convergence rate in the first three cycles was deteriorated slightly. Although the first three cycles are slower to converge, in the remaining V-cycles the convergence rate is lower than 0.1.

The convergence histories for the four cases are depicted in Fig. 4. The effect of the Mach number on the AMG performance is clearly shown. In the first three subsonic cases, it takes around eight cycles to reduce the L_2 -norm of the residual

Table 1
Convergence factor C_f , grid complexity C_Ω , and operator complexity C_L for four cases of Mach number.

Complexities	$M_\infty = 0.01$	$M_\infty = 0.1$	$M_\infty = 0.3$	$M_\infty = 0.46$
C_f	0.08	0.06	0.08	0.10
C_Ω	1.74	1.74	1.77	2.03
C_L	2.44	2.48	2.54	3.80

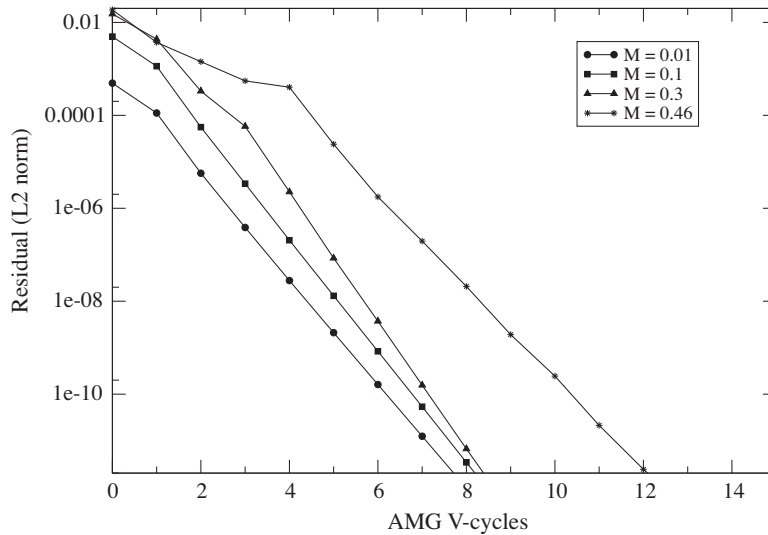


Fig. 4. Convergence histories.

to a level of 10^{-10} . The case of $M_\infty = 0.46$ results in a supersonic flow regime which is terminated by a shock wave. Convergence is somewhat slower here. It requires nearly 12 multigrid iterations to decrease the L_2 -norm of the residual to a level of 10^{-10} .

The grid complexity and operator complexity are presented in Table 1. It is clear that when the flow is subsonic the complexities are bounded. In the case of $M_\infty = 0.46$ an operator complexity of $C_L = 3.80$ is too high, relative to our requirements as stated previously, while the reasons for the increased complexities are similar to those stated in the previous problems. Two possible ways to improve the complexities in the transonic case ($M_\infty = 0.46$) are by more aggressive coarsening or truncation of the coarse-level operator. The aggressive coarsening was implemented but, as expected, convergence became considerably slower (above 0.3) though still bounded for the larger grid resolution. Aggressive coarsening not only causes the smoothing to be less effective but also interpolation is significantly less accurate. Decreasing the truncation coarse-level parameter μ , in its turn, is easy to implement but it is too dangerous since, if not done with great care, a divergence can occur.

This problem was solved for two more meshes consisting of 64×24 and 32×12 points. Both the L_2 -norm of the residual and the convergence factor for two cases of $M_\infty = 0.1$ and $M_\infty = 0.46$ are presented in Table 2. The grid complexity and operator complexity for these cases are also presented in Table 2. As one can see, for both grid sizes the residual norm decreased by a relatively constant factor with each V-cycle. This continues until it levels off after about 12 V-cycles.

5.2. Rocket engine

A rocket in its simplest form is a chamber enclosing a gas under pressure. A small opening at one end of the chamber allows the gas to escape, and this results in a thrust that propels the rocket in the opposite direction. The gas is produced by burning propellants that can be solid or liquid. In the case of solid fuel, the propellant is line-up along the walls of a cavity called the combustion chamber. The propellant is burned leading to build-up of high temperature and pressure. The expanding gas escapes through the nozzle. In the example considered the ratio of the chamber area to the throat area is $\frac{A_m}{A_{throat}} = 5.75$ and the ratio of the exit area and the throat area is $\frac{A_{out}}{A_{throat}} = 1.75$. The computational mesh is shown in Fig. 5. The mesh size is

Table 2

The table shows the results of AMG V-cycles applied to the flow through the transonic diffuser. The discrete second norm of the residual $\|R^m\|_2$ after each V-cycle and the convergence factor C_f are presented for various Mach numbers and two grids of 32×12 and 64×24 .

Grid size →	$M_\infty = 0.01$		$M_\infty = 0.1$		
	32×12	64×24	32×12	64×24	
C_f	0.09	0.08	0.09	0.09	
C_Ω	1.73	1.80	1.78	1.80	
C_L	1.88	2.41	2.01	2.42	
		$M_\infty = 0.3$		$M_\infty = 0.46$	
C_f		0.06	0.08	0.08	0.08
C_Ω		1.67	1.81	2.02	1.98
C_L		1.94	2.46	2.46	3.23

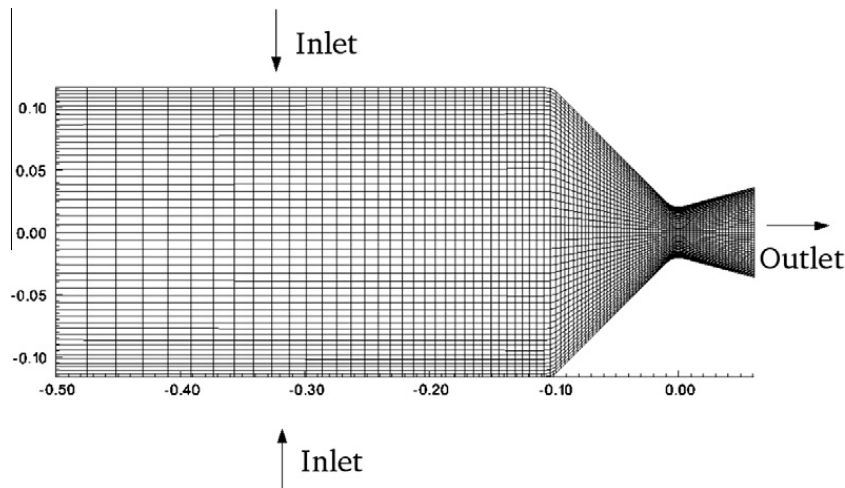


Fig. 5. Mesh used for the 2D rocket test case.

96×48 and it is clustered close to the upper and lower walls vertically and in the throat horizontally, with a stretching factor of 1.1.

5.2.1. Boundary conditions

- The solid propellant is lined-up along the bottom and top walls (see the diagram). The burning surface is simply treated as a wall through which mass injection is applied at a constant rates.
- The mass conservation determines the outflow velocity through the nozzle. In practice, ghost cells were used in order to determine the axial velocity.

Fig. 6 demonstrates the results of computations concerned with subsonic and transonic flow. The Mach contours for the various values of mass injection rate are considered in this investigation. One can readily see that when the Mach number near the upper and lower walls is $V_w = 0.028$ we get a shocked throat. The fifth case of $V_w = 0.03$ results in highly supersonic pockets with strong shock waves at the diverging section of the nozzle. The maximum Mach number at the shock was 2.28. It is clear that in this case the potential flow model is not applicable, since the potential model implies a constant entropy and has therefore no mechanisms to generate entropy variations over the shock. There is a very good agreement between these results and the theory as can be seen in Ref. [25].

5.2.2. AMG performance

The complexities C_f , C_{Ω} , and C_L are presented in Table 3. We see that after 12 V-cycles, the residual L_2 -norm reached 10^{-10} and the process attained an asymptotic convergence factor of 0.05 per V-cycle. When $V_w = 0.001$ the equation is nearly isotropic and the entries of the matrices A , \tilde{A} , and $A^* = \tilde{A}\tilde{A}$ may vary, though also very little, even in the throat area where the velocity gradient is maximal. Consequently, the setup phase needs to be performed only once, at the beginning of the solving phase. This case was solved with 12 V-cycles. As the Mach number increases the effect of the compressibility becomes more significant. In these cases the matrix A^* could not be “frozen” but rather updated while the setup phase has to be performed more frequently. For instance, in the case of $V_w = 0.008$, the setup phase was repeated three times followed by four V-cycles each time. The cases of $V_w = 0.018$ and $V_w = 0.028$ required the setup phase to be performed six times with two V-cycles between each update.

In this context, the effect of the nonlinearity and the presence of shock waves is well observed in the convergence properties for $V_w = 0.028$, where the first three cycles are slow to converge and then the residual reduction is stabilized around 0.1 for the rest V-cycles. It takes around four “waste” cycles to converge close enough to the solution so that the nonlinearity effects stop hampering the further convergence process.

Fig. 7 shows the convergence histories for the algorithm consisting of V-cycles, which was used for solving the problem for various Mach numbers. We can observe that the convergence behavior is comparable to that of the nozzle problem except that the convergence is slightly slower here. Ten cycles are needed in the fastest case, $V_w = 0.001$, to reach convergence.

A grid dependence study has been conducted to verify the independence of the AMG algorithm on the grid resolution; three grid levels have been used with (32×12) , (48×24) and (96×48) grid points in the axial and vertical directions. It is important to remember that our main interest here is to verify the robustness of the code rather than achieving the highest efficiency. Therefore in each resolution the AMG components (for instance, dynamic threshold, coarse-level parameter η , and second pass) were “fixed”, that is, they were not locally adjusted to particular requirements of a given resolution. These components remain exactly the same for all levels of resolution 96×48 .




Fig. 6. Mach number isolines computed on the (96×48) mesh, for the following flow conditions: (a) $V_w = 0.001$, (b) $V_w = 0.008$, (c) $V_w = 0.018$, (d) $V_w = 0.028$, (e) $V_w = 0.03$. Observe the shocks in the diverging section appearing at $V_w = 0.028$ and $V_w = 0.03$.

Table 3

Convergence factor C_f grid complexity C_G and operator complexity C_L for four cases of Mach number. The mesh size is 96×48 .

Complexities	$V_w = 0.001$	$V_w = 0.008$	$V_w = 0.018$	$V_w = 0.028$	$V_w = 0.03$
C_f	0.05	0.11	0.13	0.13	0.14
C_G	1.63	1.63	1.70	1.90	1.90
C_L	2.19	2.22	2.43	2.76	4.76

The L_2 -norm of the residual $\|R^m\|_2$ and the convergence rate C_f are shown in Table 4. For each of the grid resolutions the results demonstrate rapid convergence for all the four cases of Mach numbers. For each case of Mach number there is a minor degradation in the convergence rate with the increase of resolution. For instance, for the case $V_w = 0.008$ the convergence rate of $C_f = 0.1$ was obtained in mesh size (32×12) compared to $C_f = 0.15$ in mesh size (48×24) . This degradation is a result of two possible reasons. First, the interpolation which is much less accurate as the resolution decreased, especially were the irregularity of the grid is dominant. Second, for small resolution the symmetric Gauss–Seidel method serves not only as a smoother (on the fine and coarse-levels as well) but also serves as a solver. In this case a better convergence rate can be achieved.

5.3. Symmetric airfoil – NACA-0012 with $\alpha = 0^\circ$

5.3.1. Problem definition

A grid convergence study was done in order to determine the appropriate grid dimensions in terms of the distance between the airfoil and farfield, grid resolution etc. Three different grids were tested. They are as follows: (128×48) , (64×24) , (32×12) . The first number corresponds to the number of cells that wrap around the airfoil surface, and the second dimension is the number of cells normal to airfoil's surface. An O -topology mesh is generated by applying the algorithm of Hawken [27]. A close-up of the grid near the airfoil is shown in Fig. 8(b). We can see from the rendering of the O -topology




Fig. 7. Convergence histories.

Table 4

The results of AMG V-cycles applied to the flow through a rocket chamber. The second norm of the residual after each V-cycle $\|R^m\|_2$ and the convergence factor C_f are detailed for two different mesh sizes.

Grid	$V_w = 0.001$		$V_w = 0.008$	
	24×12	48×24	24×12	48×24
C_f	0.08	0.13	0.10	0.14
C_Ω	1.87	1.74	1.87	1.75
C_L	2.17	2.28	2.17	2.28
Grid	$V_w = 0.018$		$V_w = 0.028$	
	24×12	48×24	24×12	48×24
C_f	0.05	0.05	0.06	0.15
C_Ω	1.82	1.79	1.79	1.66
C_L	2.17	2.46	2.27	2.80

that the aerodynamic body is enclosed by one family of grid lines ($\eta = const.$). The second family of grid lines ($\xi = const.$) is spanned in the radial direction between the body and the farfield. The complete boundary line $\eta = 1$ represents the contour of the body, from a to b . The coordinate cut is defined by the boundaries $\xi = 1$ (nodes a–c) and $\xi = max$ (nodes b–d) in the computational space.

The farfield boundary ($\eta = max$) is ten chord lengths away from the airfoil, where the inflow and outflow boundary conditions were imposed. Since the potential assumption implies the flow is irrotational, the local vorticity production is zero. When solving a flow problem involving an airfoil which produces lift, there is a finite circulation Γ around the airfoil. In order to obtain lift, this circulation must be imposed. Its magnitude is obtained by applying Kutta–Joukowski condition, which states that the circulation value is such that the stagnation point is located at the trailing edge. The details on how to implement the Kutta–Joukowski condition are described below.

5.3.2. The Kutta–Joukowski condition

Consider an airfoil at an angle of attack, producing a nonzero lift. The following line integral over any contour enclosing such an airfoil will produce a nonzero result, known as the circulation around an airfoil Γ ,

$$\oint_{\partial S} \vec{V} \cdot d\vec{S} = \Gamma \quad \text{around airfoil} \tag{16}$$

In (16), $d\vec{S}$ is an infinitesimal line segment vector, tangential to the contour. The integrand may be written as follows:

$$\vec{V} \cdot d\vec{S} = \vec{\nabla} \phi \cdot d\vec{S} = (\phi_x \hat{i} + \phi_y \hat{j})(dx \hat{i} + dy \hat{j}) = \frac{\partial \phi}{\partial x} dx + \frac{\partial \phi}{\partial y} dy = d\phi. \tag{17}$$

The above result means that somewhere in the x – y plane, the potential function must experience an abrupt jump in its value, by an amount equal to the circulation Γ . In our analysis, we locate the disturbance potential jump along a cut that starts at the airfoil trailing edge, and ends at the downstream infinity. Fig. 9 gives an illustration of the trailing edge in 2D airfoil problem. Consider a typical cell $(i,1)$ above the cut, downstream of airfoil trailing edge. The circulation is first

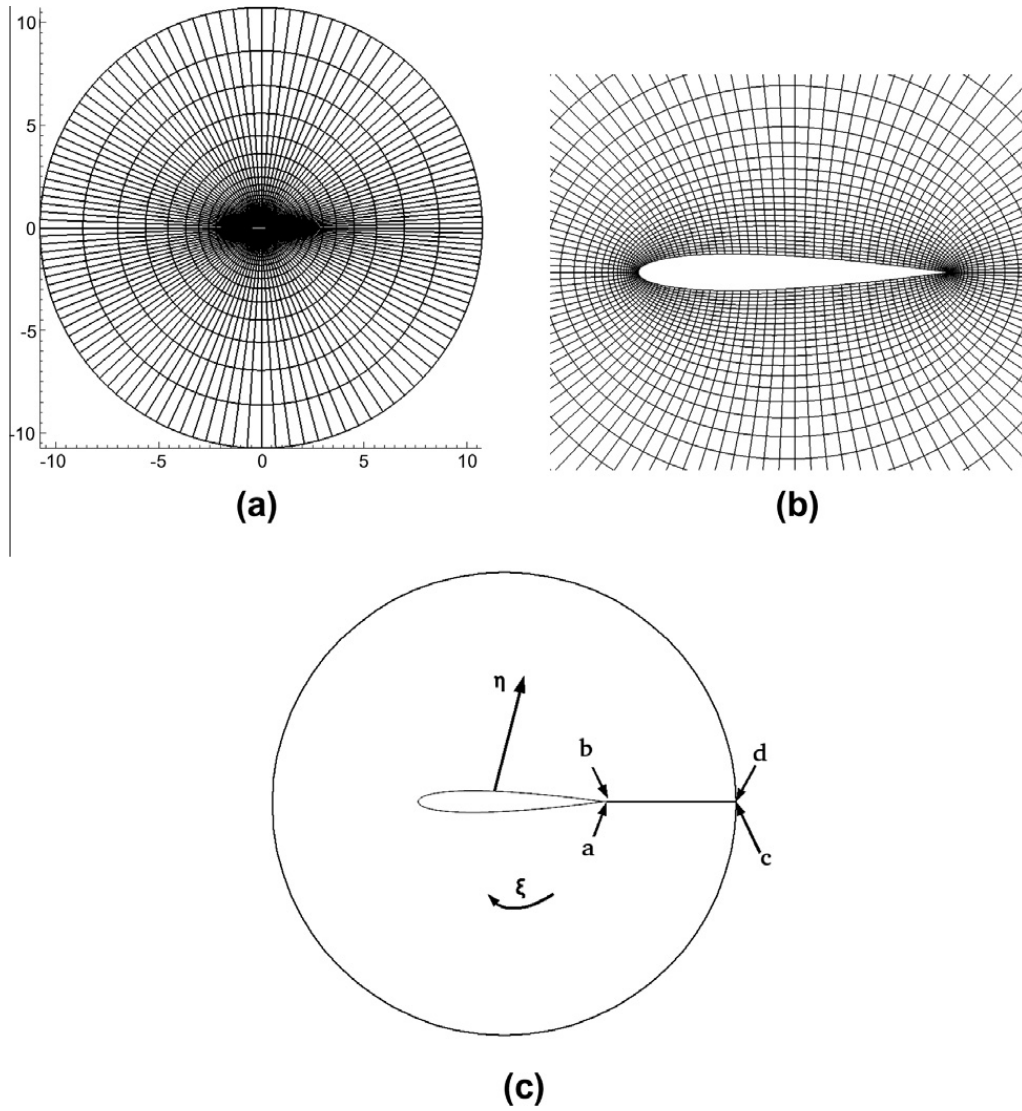


Fig. 8. Structured curvilinear body-fitted grid of the O-type. (a) Mesh used for the NACA-0012 airfoil flow case study. (b) A close-up of the grid near the airfoil. (c) Diagram of the airfoil in O-type topology.




Fig. 9. The Kutta condition in the finite volume method.

computed as, $\Gamma = \phi_{TE,top} - \phi_{TE,bottom}$, where $\phi_{TE,top}$ and $\phi_{TE,bottom}$ are the potential values at the trailing edges for the top and bottom points, respectively. In Fig. 9, the values of the potential $\phi_{TE,top}$ and $\phi_{TE,bottom}$ are known, where the subscripts *TE, top* and *TE, bottom* correspond to the trailing edge upper and lower *i* coordinates, respectively. Since these values are defined at the cell's center, and are not known at the trailing edge, an extrapolation is needed. Once the circulation value is calculated, the next step is to apply the potential jump. All the cells above and below the cut which is drawn from the airfoil surface to the farfield boundary, will have their value of potential modified to satisfy the proper jump condition. It is done as follows:

$$\begin{aligned} \phi_{i,0} &= \phi_{i,jmax} + \Gamma, \\ \phi_{i,-1} &= \phi_{i,jmax-1} + \Gamma, \\ \phi_{i,jmax+1} &= \phi_{i,1} - \Gamma, \\ \phi_{i,jmax+2} &= \phi_{i,2} - \Gamma. \end{aligned} \tag{18}$$

5.3.3. Boundary conditions

For the airfoil surface, a solid wall boundary condition is applied (no penetration), $\rho \vec{V} \cdot \hat{n} = \rho \frac{\partial \phi}{\partial n} = 0$. There are two ways in which the farfield boundary condition are imposed. The first method, which is chosen to be implemented in this research work, is by using a Neumann boundary condition at the farfield. The flow field is assumed to be known at the farfield and is taken to be free-stream condition. Practically, it is implemented by projecting the free-stream velocity V_∞ normal to cell's face.

5.3.4. Qualitative results for $\alpha = 0^\circ$

Fig. 10 demonstrates the results of test cases in which a subsonic and transonic flow are involved. The Mach number isolines are depicted in Fig. 10 for $M_\infty = 0.1-0.76$ with $\alpha = 0^\circ$. The flow-field around the airfoil includes strong local velocity gradient in the leading edge region, where over a 2–3 cells downstream of the stagnation point, the velocity increases from nearly zero to values larger than the incoming free-stream velocity. The following basic features are visible, indicating that qualitatively accurate results are being depicted.

The pressure coefficient $C_p = \frac{p - p_\infty}{\frac{1}{2} \rho_\infty u_\infty^2}$ is the difference between the local static pressure and the free-stream static pressure, nondimensionalized by the free-stream dynamic pressure. Fig. 11 shows the pressure distribution on the upper and lower surfaces computed by the pressure coefficient with negative C_p , toward the top, following the usual convention. Observe the




Fig. 10. Distribution of the velocity as computed on the (128×48) mesh, for different incident Mach numbers. (a) $M_\infty = 0.1$, (b) $M_\infty = 0.3$, (c) $M_\infty = 0.5$, (d) $M_\infty = 0.76$. Observe the shocks appearing at the top and bottom surfaces.




Fig. 11. The distribution of pressure coefficient for an NACA-0012 airfoil as computed using (128×48) mesh, for four Mach numbers.

sharp decrease in C_p across the shock wave. The $M_\infty = 0.1$ and $M_\infty = 0.3$ cases have a slight difference in both the convergence and C_p solution. The flow is nearly incompressible at these speeds, giving very little differences in the results. At $M_\infty = 0.5$ the compressibility of the flow become more significant, and at $M_\infty = 0.76$ the compressibility effects are very evident.

5.3.5. AMG performance

In the subsonic cases, $M_\infty \leq 0.5$ the coarse-levels were obtained while applying a fixed threshold ($\varepsilon = 0.25$). It was sufficient to apply the second-pass process in the first coarsening step and maintain standard coarsening for all subsequent levels. In the transonic case, $M_\infty = 0.76$ a dynamic threshold produced much better results (in terms of convergence properties) but on the expense of somewhat higher memory requirements.

The strength of connection between the matrix variables can be clearly seen from Fig. 12, where the finest and first coarse-levels are plotted. For $M_\infty = 0.1$ the operator is nearly isotropic and due to the setting of $\varepsilon = 0.25$ the AMG algorithm treats most of the connections contained in the corresponding matrix as strong (at least, there is an azimuthal symmetry in the strong connection). Consequently, the first coarse-level corresponds to red–black coarsening in the most of the domain, far from the airfoil surface. There is a significant deviation from this coarsening pattern near the airfoil surface where an isotropic coarsening is observed. The reason is that the grid is highly stretched near the airfoil surface and the cell's aspect ratio becomes nearly unity. In few cells above the airfoil the anisotropy is the largest and this is well reflected in the one-dimensional coarsening pattern, while the AMG coarsened in the direction of the strong connectivity, in the η -direction. This coarsening pattern is essentially the same for the rest of the cases.

Table 5 presents the convergence factor, the grid and operator complexities for the above Mach numbers. When the flow is subsonic, $M_\infty \leq 0.5$, the convergence rate is bounded below an order of magnitude. After 12 V-cycles a residual smaller than 10^{-10} was reached. In the subsonic cases $M_\infty \leq 0.5$ the coefficient matrices of A , \tilde{A} , and A^* were changed very slowly from one cycle to the next, especially above and below the airfoil's surface and at the leading and trailing edges, where the velocity gradient is maximal. Consequently, performing the setup phase only once, before executing the first V-cycle, is sufficient to achieve an efficient convergence, and these cases were solved with 12 V-cycles.

The influence of the compressibility on the overall convergence is well observed in the transonic case $M_\infty = 0.76$, where, as expected, convergence becomes considerably slower. The results were obtained by repeating the setup phase 6 times while applying two V-cycles between each update. It results in a total of 12 V-cycles until the residual decreased to the desired level of $\sim 10^{-10}$. A possible way to improve convergence is by investing more effort by additional smoothing sweeps (~ 1 – 2 local iterations) on the fine-level. This approach was preferred in order to improve the overall convergence properties in a cheap way and to overcome this convergence degradation. In addition, the main purpose of this “extra” work is mainly to “locate” the shock and to establish its profile, i.e., to “clean-up” the area of the strong nonlinearity as the coarse-level correction is distributed on the fine-level. In this way we also significantly reduce the accumulation of errors which were interpolated from the coarse-levels.

The convergence histories are demonstrated in Fig. 13. It is clear that as the Mach number increases so does the number of iterations required to solve the problem. We observe that for subsonic cases, $M_\infty = 0.1$ – 0.5 we need less than 10 V-cycles to reduce the L_2 -norm of the residual by 10 orders of magnitude. As for the transonic case of $M_\infty = 0.76$, the convergence is somewhat slower here while about 12 cycles are required to reduce the residual by 10 orders of magnitude.

This extra increase in the number of iterations is due mainly to the strong anisotropy and the existence of two supersonic pockets which are terminated by shock waves. For $M_\infty = 0.76$, the problem was converged in less than an order of magnitude




Fig. 12. The finest and first coarse-level for mesh size (128×48) . The red color corresponds to the a C-point and the blue color corresponds to an F-point. (a) $M_\infty = 0.1$, (b) $M_\infty = 0.76$. The pictures on the right are magnified views of the airfoil's region. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 5

Convergence factor C_f , grid complexity C_Ω and operator complexity C_L for four cases of Mach number.

Complexities	$M_\infty = 0.1$	$M_\infty = 0.3$	$M_\infty = 0.5$	$M_\infty = 0.76$
C_f	0.04	0.05	0.11	0.18
C_Ω	1.90	1.93	2.13	1.92
C_L	3.01	2.95	3.87	3.89

per V-cycle, and a much higher operator complexity of $C_\Omega = 3.89$ was obtained. This behavior is typical of the more general problems which are characterized by a strong anisotropy and nonlinearity. Combinations of AMG components that produce the best convergence rates are also those with the higher costs in terms of grid and operator complexity. Parameter selection is largely the art of finding a compromise between performance and cost.

This problem was solved for two more grid sizes of 64×24 and 32×12 . The complexities are presented in Table 6. This is a case where applying AMG convergence acceleration is a very attractive possibility since it is precisely a situation for which constructing an efficient geometric multigrid approach would become extremely cumbersome. We observe that for subsonic flow, the AMG exhibits the same type of convergence that was observed in the previous problems. As the Mach number increases the convergence deteriorates but, we find the solver to still be very efficient for this problem, despite of the strong nonlinearity and the anisotropy of the problem. The storage requirement is higher than that for the subsonic case (isotropic case), the reason being that AMG performs one-dimensional coarsening in the azimuthal direction at the farfield since the strong connection in the ξ -direction (azimuthal direction) arises from a large grid spacing in η -direction. In this respect, the complexities deteriorate significantly.




Fig. 13. Convergence histories.

Table 6

The table shows the results of AMG V-cycles applied to the flow around NACA-0012 airfoil. The convergence factor C_f , grid complexity C_Ω , and operator complexity C_L are detailed for two different mesh sizes.

Grid size →	$M_\infty = 0.1$		$M_\infty = 0.76$	
	32×12	64×24	32×12	64×24
C_f	0.05	0.07	0.09	0.09
C_Ω	1.88	1.91	2.00	2.09
C_L	2.62	2.93	3.18	3.54

Table 7

Results of AMG V-cycles applied to the flow around NACA-0012 airfoil.

Level	Number of rows		Number of non-zeros		Average entries per row	
	$M_\infty = 0.1$	$M_\infty = 0.76$	$M_\infty = 0.1$	$M_\infty = 0.76$	$M_\infty = 0.1$	$M_\infty = 0.76$
A^m	6144	6144	108,983	110,289	17.7	17.9
A^{m+1}	3072	3323	104,678	148,812	34	44.7
A^{m+2}	1437	1758	60,275	86,978	41.9	49.4
A^{m+3}	659	968	33,783	45,543	51.26	47
A^{m+4}	266	532	14,869	23,786	55.89	44.7
A^{m+5}	104	278	5805	12,115	55.8	43.5

Let us examine the storage requirements for two extreme cases of $M_\infty = 0.76$ and $M_\infty = 0.1$. For $M_\infty = 0.1$ the problem was solved using six multigrid levels and the solver's characteristics properties are presented in Table 7. The initial coarse-level, A^{m+1} has 3072 points, exactly half the number of points on the fine-level. This can be clearly seen in the coarsening pattern presented in Fig. 12(a). The coarsening in the next coarse-levels is more "aggressive", while each succeeding coarse-level has less than half the number of points as the next finer level. In can be seen in the last column that the average number of non-zero entries per row increases as the grids become coarser. For $M_\infty = 0.76$ the situation is even more dramatic.

We would like to recall here that our main goal in these applications is to demonstrate how different AMG components may influence the overall performance. We have not tried to find a combination that results in an optimal performance, but rather confine ourselves to a set of components or parameters which result in an acceptable performance and in robust multi-purpose algorithm. The above results clearly show that optimized parameter setting (for instance: dynamic threshold, matrix reduction parameter η , second-pass process), depending on the application, may improve the performance substantially further.

5.4. Symmetric airfoil – NACA-0012 with $\alpha = 1.25^\circ$

The following results were obtained using the 128×48 mesh. Fig. 14 presents the results of test cases in which a subsonic and transonic flow are involved. The Mach number isolines for $M_\infty = 0.1$ – 0.71 with $\alpha = 1.25^\circ$ are depicted. It can be observed

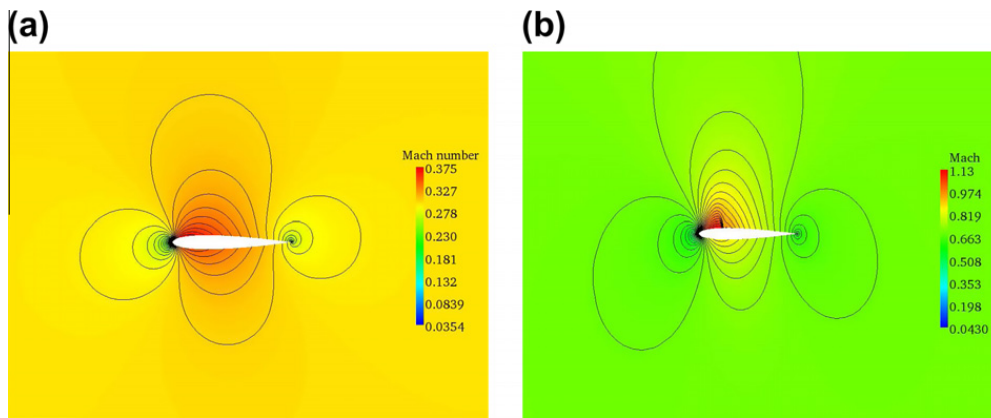


Fig. 14. The velocity distribution as computed on the (128×48) mesh, for different free-stream Mach numbers. (a) $M_\infty = 0.3$, (b) $M_\infty = 0.71$. Observe the shock appears at the bottom surface.

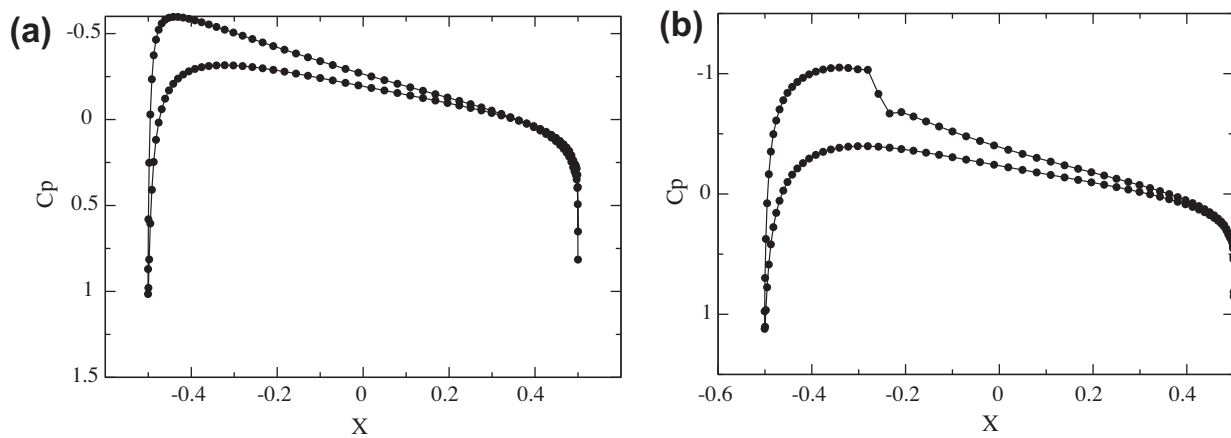


Fig. 15. The distribution of pressure coefficient for an NACA-0012 airfoil at $\alpha = 1.25^\circ$ as computed using (128×48) mesh, for two Mach numbers: (a) $M_\infty = 0.3$, (b) $M_\infty = 0.71$.

that the Kutta condition is satisfied and the stagnation point is located on the trailing edge. Fig. 15 shows the pressure coefficient C_p on the upper and lower surfaces. The presence of shock waves above and below the airfoil are well observed in the sharp decrease of C_p when $M_\infty = 0.71$.

5.4.1. AMG performance

A symmetric airfoil NACA-0012 with $\alpha = 1.25^\circ$ produces lift due to the imposed finite circulation Γ around the airfoil. The value of the circulation is unique for each set of flow conditions. This problem is much more complicated to solve with the AMG method. In fact, this problem becomes a kind of an optimization problem while the goal is to improve the circulation until convergence is reached for both the circulation Γ and the equation itself (the residual has to reach the level of discretization). Since Γ is updated only on the fine-level, the overall AMG performance significantly suffers. The reason is that in order to solve the problem, both the residual and the circulation have to reach convergence. Each circulation update, however, results in a new set of potential values along the cut, which leads to a somewhat different solution and, therefore, has an extreme impact on the overall convergence of the algorithm. The coarse-level correction loses its relevance at the moment the circulation is updated. As a consequence, the convergence rate becomes extremely slow (around 0.85). It is important to note here that it does not matter if the residual is dropped five orders of magnitude (or even more) by successive V-cycles (without updating Γ). Once the circulation value is modified, the residual jumps three or four orders of magnitude and most that we have gained in the AMG V-cycles is lost! Note that as the circulation Γ is “frozen” the AMG performance becomes visible again with convergence rate below an order of magnitude per cycle. However, this situation makes the use of the AMG algorithm very unpractical.

Practical experience has shown that in order to “cure” this problem a possible way is to obtain a good approximation of Γ as an initial condition for the algorithm. A possible way to obtain a good approximation Γ_0 is by solving the problem on a coarser grid (usually a grid coarser twice in each direction than the finer one). This is done by the FMG method in the context of AMG. A properly designed FMG scheme can be much more effective in general than just a sequence of V-cycles. The key principle regarding the FMG approach in general approach is that, before addressing the Ω^m problem, we have to solve

the Ω^{2m} problem to the level of discretization error. In our case here this also result in a good approximation to the circulation value Γ_0 (its value is within $\sim 1\text{--}3\%$ of the value of Γ that is obtained on the fine-level Ω^m) as an initial condition. It is very important to mention that the convergence rate of these nonlinear iterations (especially in the transonic case where the nonlinearity is significant) depends dramatically on a quality of the initial condition for Γ . The better the initial guess Γ_0 used on the fine-level, the more effective the AMG V-cycle will be. In our test-case the FMG scheme relies on four levels, while on the coarsest level (16×6 grid) the problem was solved to the level of machine zero.

The practical implementation of this problem, relies heavily upon the way the treatment of the circulation value is incorporated in each V-cycle. The value of the circulation Γ is updated at the end of each V-cycle, immediately after the correction from the coarse-levels is distributed to the points on the fine-level. A possible way to improve the convergence is to use a simple 2D linear extrapolation of the circulation value. If the circulation value is known in the present and last iterations the next value can be calculated as follows: an “over-relaxation” parameter $1 \leq \omega \leq 2$ can be added in order to accelerate convergence:

$$\Gamma_{new} = \Gamma_{n+1} + \omega(\Gamma_{n+1} - \Gamma_n). \tag{19}$$

The first coarse-level of the AMG algorithm for each case described above is sketched in Fig. 16 as a magnified view of the area around the airfoil. When the flow is incompressible, the operator is nearly isotropic and the points which construct the coarse-level are distributed uniformly, as can be seen in Fig. 16(a). As M_∞ is increased the equation becomes more anisotropic, which is well reflected in the coarsening pattern, where the number of coarse points increases around the sonic line, above and below the airfoil, where the anisotropy is strongest.

The discrete L_2 -norm of the residual and the convergence rate are shown in Table 8 for each V-cycle, in four different cases of Mach number. In Table 8 the grid and operator complexities are presented for the above Mach numbers. When the flow is subsonic, $M_\infty \leq 0.5$, the convergence rate is bounded below an order of magnitude. After 12 V-cycles a residual smaller than 10^{-10} was reached for the case $M_\infty = 0.71$. In the subsonic cases $M_\infty \leq 0.5$ the coefficient matrices A, \tilde{A} , and A^* were changed very slowly from one cycle to the next, even above and below the airfoil surface and at the leading and trailing edges where the velocity gradient is maximal. Consequently, performing the setup phase only once, at the beginning of the solving phase, is sufficient to obtain a stable algorithm and efficient convergence. These problems were solved to the machine zero by applying 12 V-cycles on the finest level.

The influence of the compressibility on the overall performance is well observed in the transonic case $M_\infty = 0.71$. The convergence, as expected, becomes considerably slower. The results were obtained by repeating the setup phase 6 times while applying two V-cycles between each update. It results in a total of 12 V-cycles until the residual decreased to the desired level of $\sim 10^{-10}$. In this particular case more effort was invested by additional local smoothing sweeps (~ 2 local iterations) in the neighborhood of the shocks. This appeared very helpful in order to improve the overall performance.

The convergence histories for the five cases are depicted in Fig. 17. The effect of the Mach number on the AMG performance is clearly visible. The case of $M_\infty = 0.71$ is characterized by two supersonic pockets which are terminated by a shock




Fig. 16. A magnified view of the fine and first coarse-level of the airfoil with $\alpha = 1.25^\circ$ for two cases of Mach number (a) $M_\infty = 0.3$, (b) $M_\infty = 0.71$. Blue point indicates an F-point and red point indicates a C-point. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 8

Grid complexity C_Ω , operator complexity C_L , and convergence rate for four cases of Mach number.

Complexities	$M_\infty = 0.1$	$M_\infty = 0.3$	$M_\infty = 0.5$	$M_\infty = 0.71$
C_f	0.08	0.09	0.09	0.18
C_Ω	3.02	2.94	1.95	2.05
C_L	1.90	1.92	2.99	4.72

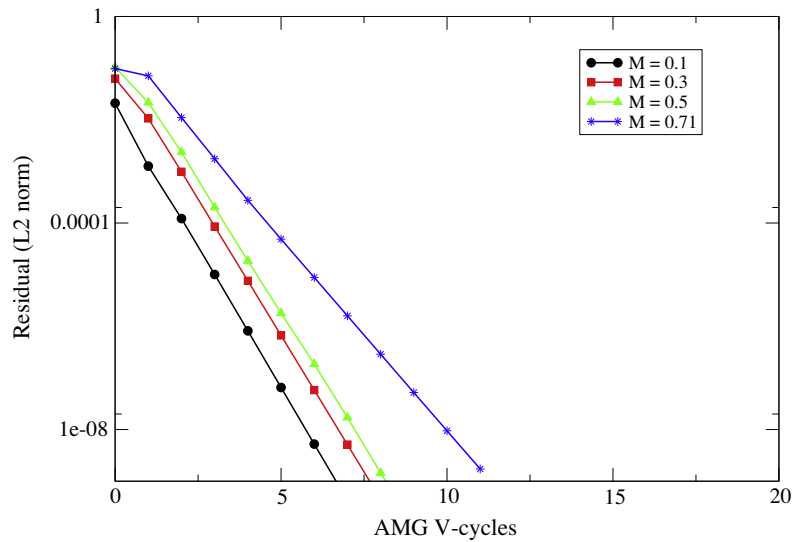


Fig. 17. Convergence histories.

wave, above and below the airfoil. Convergence is somewhat slower here. It requires nearly 11 V-cycles to decrease the L_2 -norm of the residual to a level of 10^{-10} . The case of $M_\infty = 0.01$ is converged to the desired level of 10^{-10} just by 7–8 cycles. Although the convergence rates are very good, the grid and operator complexities are much higher ($C_\Omega = 2.05$ and $C_L = 4.72$) compared to our requirements.

This problem was also solved for two more grid sizes of 64×24 and 32×12 . The complexities are presented in Table 9. For $M_\infty < 0.5$ we observe that for subsonic flow the AMG exhibits the same type of convergence that was observed for this problem solved on the finer grid. The residual L_2 -norm decreases by a nearly constant factor, below an order of magnitude per V-cycle. It takes about 8 V-cycles to reach a residual of 10^{-12} . As the Mach number increases, the convergence degrades, but is still very good for this problem, in spite of the strong anisotropy. The storage requirements are higher than those for the subsonic case (isotropic case) the reason being that AMG performs one-dimensional coarsening in the azimuthal direction at the farfield. This is because the strong connections in the ξ -direction (azimuthal direction) arise from a large grid spacing in η -direction. This fact strongly has a rather severe adverse effect on the complexities.

Table 10 illustrates the FMG performance in resolving the flow around an airfoil NACA-0012 with $\alpha = 1.25^\circ$ for the cases concerning four different free-stream Mach numbers with mesh size of 128×48 . The row labeled “FMG(1,1), number of levels” lists how many coarse-levels (geometrically coarsened $h \rightarrow 2h$) were needed in the FMG scheme. The row labeled “AMG V(1,1) cycles” gives the number of AMG V-cycles used until the residual decreased by 2–3 orders of magnitude in each coarse-level. Also, the average convergence factor is presented for each level.

5.5. Nonsymmetric airfoil – NACA-2822 with $\alpha = 0^\circ$

Here we take one step further and consider a nonsymmetric airfoil. First, the non-symmetry results in an irregular stretched grid near the airfoil surface. Second, such an airfoil generates lift even at a zero angle of attack. Therefore, the circulation must be calculated after each smoothing sweep in order to satisfy the Kutta condition. The grid has 48 cells in the normal direction and 128 cells around the airfoil. The farfield is located ten chord lengths away from the airfoil. The grid and a magnified view of the grid near the airfoil’s surface are presented in Fig. 18.

Table 9

The results of AMG V-cycles applied to the flow around NACA-0012 airfoil with $\alpha = 1.25^\circ$. The convergence factor C_f grid complexity C_Ω and operator complexity C_L are detailed for two different mesh sizes, 32×12 and 64×24 .

Grid size →	$M_\infty = 0.1$		$M_\infty = 0.3$	
	32×12	64×24	32×12	64×24
C_f	0.07	0.08	0.08	0.08
C_Ω	1.85	1.91	1.88	1.91
C_L	2.53	2.89	2.51	2.78
	$M_\infty = 0.5$		$M_\infty = 0.71$	
C_f	0.07	0.11	0.12	0.16
C_Ω	1.95	1.92	1.99	1.99
C_L	2.70	2.80	3.78	4.32

Table 10

The FMG performance applied to the flow around NACA-0012 airfoil with $\alpha = 1.25^\circ$ for four Mach numbers. The FMG (1, 1) scheme one symmetric Gauss–Seidel relaxation on the descent phase and one symmetric Gauss–Seidel relaxation on the ascent phase. The FMG (1, 1) number of levels is presented in the first row. The second row includes the AMG V-cycles performed in the FMG process. The average convergence factor between successive V-cycles is presented in the third row.

Grid size →	$M_\infty = 0.1$			$M_\infty = 0.3$		
	32×12	64×24	128×48	32×12	64×24	128×48
FMG (1, 1), number of levels	3	3	4	3	3	4
AMG V (1, 1) – cycles	1	1	5	1	1	5
Average convergence factor, C_f	–	–	0.12	–	–	0.12
	$M_\infty = 0.5$			$M_\infty = 0.71$		
FMG (1, 1), number of levels	3	3	4	3	3	4
AMG V (1, 1) – cycles	1	1	5	1	10	10
Average convergence factor, C_f	–	–	0.13	–	0.12	0.14




Fig. 18. Structured curvilinear body-fitted grid of the O-type for NACA-2822 airfoil. (a) The mesh used for the NACA-2822 airfoil flow case study. (b) A close-up of the grid near the airfoil.

The same boundary conditions were applied as for NACA-0012 airfoil test case. The Mach contours for the cases of two different free-stream Mach numbers are presented in Fig. 19. Fig. 19(a) depicts a subsonic case ($M_\infty = 0.3$). At $M_\infty = 0.75$, a shock is present in the flow above the upper surface of the airfoil as can be seen in Fig. 19(b). Fig. 20 shows the C_p distribution for the airfoil at $M_\infty = 0.1, 0.3, 0.5, 0.75$ and $\alpha = 0^\circ$.

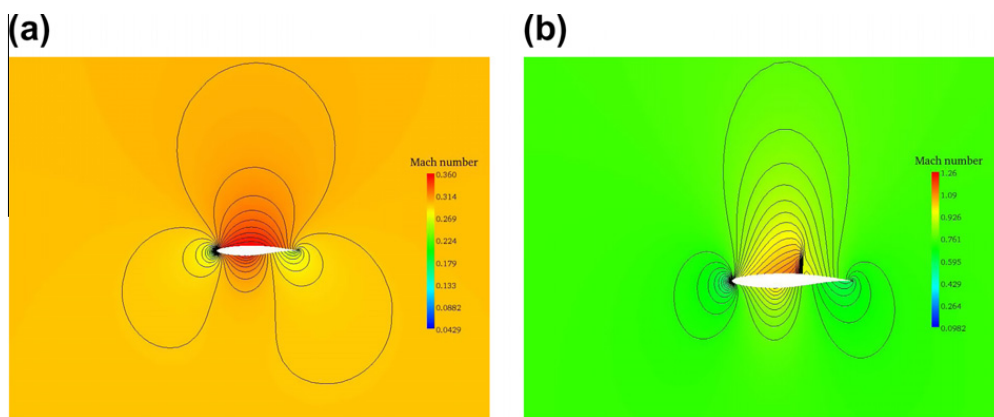


Fig. 19. Distribution of velocity as computed on the (128×48) mesh, for different incident Mach numbers. (a) $M_\infty = 0.3$, (b) $M_\infty = 0.75$. Observe the shocks appearing at top and bottom surfaces.




Fig. 20. The distribution of pressure coefficient as computed using (128×48) mesh, for different Mach numbers at $\alpha = 0^\circ$.




Fig. 21. A Magnified view of the coarsening area around the airfoil. Fine and first coarse-level of the airfoil with $\alpha = 0^\circ$ for four cases of Mach number: (a) $M_\infty = 0.3$, (b) $M_\infty = 0.75$. Blue point indicates an *F*-point and red point indicates a *C*-point. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

5.5.1. AMG performance

The first coarse-level for cases $M_\infty = 0.3$ and $M_\infty = 0.75$ described above is sketched in Fig. 21. When the flow is practically incompressible, the operator is nearly isotropic and the points which comprise the coarse-level are distributed uniformly, as can be seen in Fig. 21(a).

Table 11 presents the complexities for cases of different free-stream Mach numbers. In the first three cases of $M_\infty = 0.1$, $M_\infty = 0.3$ and $M_\infty = 0.5$, the setup phase was implemented only once, followed by 12 V-cycles. For the transonic case of $M_\infty = 0.75$, where the compressibility became significant, the results were obtained by repeating the setup phase 6 times while applying two V-cycles between each update. A total of 12 V-cycles were performed until the residuals decreased to the desired level of $\sim 10^{-10}$. Although the first three cycles are slow to converge, in the remaining V-cycles the convergence rate is lower than 0.17 without any extra relaxation sweeps on the fine-level.

The convergence factor and the complexities are presented in Table 11. It is clear that when the flow is subsonic the complexities are bounded. In the case of $M_\infty = 0.75$ the complexities are somewhat high.

Table 11

Convergence rate C_f , grid complexity C_Ω , and operator complexity C_L for four cases of Mach number.

Complexities	$M_\infty = 0.1$	$M_\infty = 0.3$	$M_\infty = 0.5$	$M_\infty = 0.75$
C_f	0.08	0.09	0.10	0.15
C_Ω	1.90	1.92	1.96	2.08
C_L	3.03	2.94	3.02	4.95




Fig. 22. Convergence histories.

The convergence histories for the four cases are described in Fig. 22. The effect of the Mach number on the AMG performance is clearly visible. The case of $M_\infty = 0.75$ results in a supersonic flow regime (in the upper surface) which is terminated by a shock wave. It requires nearly 12 V-cycles to decrease the L_2 -norm of the residual to a level of 10^{-10} , which is twice the number of AMG iterations required in the case of $M_\infty = 0.1$.

This problem was solved for two more grid sizes of 64×24 and 32×12 . The convergence factor and the complexities are presented in Table 12. We observe that for subsonic flow the AMG exhibits the same type of convergence that was observed on the finer grid. As the Mach number increases the convergence degrades, albeit we find it still very efficient for this problem, in light of the strong irregularity and anisotropy. The storage requirements are higher than that for the subsonic case (isotropic case) the reason being that AMG performs one-dimensional coarsening in the azimuthal direction at the farfield

Table 12

The table shows the results of AMG V-cycles applied to the flow around NACA-2822 airfoil with $\alpha = 0^\circ$. The convergence factor C_f , grid complexity C_Ω , and operator complexity C_L are detailed for two different mesh sizes of 32×12 and 64×24 .

Grid size →	$M_\infty = 0.1$		$M_\infty = 0.3$	
	32×12	64×24	32×12	64×24
C_f	0.07	0.07	0.09	0.07
C_Ω	1.85	1.90	1.89	1.91
C_L	2.54	2.91	2.53	2.82
Grid size →	$M_\infty = 0.5$		$M_\infty = 0.75$	
	32×12	64×24	32×12	64×24
C_f	0.08	0.11	0.11	0.14
C_Ω	1.91	1.92	1.98	2.02
C_L	2.63	2.82	3.94	4.60

Table 13

The table shows the FMG performance applied to the flow around an NACA-2822 airfoil with $\alpha = 0^\circ$ for four Mach numbers. The FMG (1,1) scheme one symmetric Gauss–Seidel relaxation on the descent phase and one symmetric Gauss–Seidel relaxation on the ascent phase. The FMG (1,1) number of levels is presented in the first row. The second row includes the AMG V-cycles performed in the FMG process. The average convergence factor between successive AMG V-cycles is presented in the third row.

Grid size →	$M_\infty = 0.1$			$M_\infty = 0.3$		
	32×12	64×24	128×48	32×12	64×24	128×48
FMG (1,1), number of levels	3	4	4	3	4	4
AMG V (1,1) – cycles	1	5	5	1	5	5
Average convergence factor, C_f	–	0.09	0.10	–	0.09	0.10
Grid size →	$M_\infty = 0.5$			$M_\infty = 0.75$		
	32×12	64×24	128×48	32×12	64×24	128×48
FMG (1,1), number of levels	3	4	4	3	4	4
AMG V (1,1) – cycles	1	5	5	1	10	10
Average convergence factor, C_f	–	0.11	0.12	–	0.12	0.12

since the strong connection in the ξ -direction (azimuthal direction) arises from a large grid spacing in η -direction. This fact has no adverse effect on the complexities.

Table 13 shows the FMG performance for cases of four different free-stream Mach numbers with mesh size of 128×48 cells. In the case of $M_\infty < 0.5$ five AMG V-cycles were applied on each coarse-level as part of the FMG algorithm. Observe that the average convergence rate was around an order of magnitude per cycle. In the transonic case ten V-cycles were applied in each coarse-levels.

6. Conclusions

Transonic flow problem is a rather complex one from a computational point of view. One of the main difficulties is the fact that the differential operator changes its type between elliptic for subsonic flow regime and hyperbolic (with respect to the flow direction) in the supersonic flow regime. Another (sub-) difficulty is that the subsonic flow regime itself presents two extremities (and all the possible cases in between): nearly isotropic operator for the flow speed case and highly anisotropic operator for a nearly sonic flow speed. While the standard AMG algorithm can treat the latter difficulty, it has never been applied yet, to the best of our knowledge, to the supersonic regime. The difficulties here begin with the fact that a simple pointwise relaxation procedure (a desirable component of AMG) appears unstable in the hyperbolic case. One of the main achievements of our work is the development of a pointwise relaxation procedure that is stable (and constitutes a good smoother – in the algebraic sense) for both the subsonic and supersonic flow regimes. Second, we have constructed a variant of an AMG algorithm that employs the new relaxation procedure and allows to achieve very good convergence for both elliptic and hyperbolic cases.

A 2D body-fitted structured grid FPE solver which is based on the AMG method was developed. The flow solver is capable of resolving flows from subsonic to transonic and supersonic regimes. The flow solver can accommodate complex geometries. Several two dimensional flow calculations have been performed to test the performance of the algebraic multigrid method. The computational method was demonstrated to be capable of predicting the shock formation and achieving residual reduction of roughly an order of magnitude per cycle both for elliptic and hyperbolic problems, through the entire range of flow regimes, independent of the problem size (resolution). The purpose of this work was not to design an algorithm which is fine-tuned for a specific problem with the aim to obtain the highest possible efficiency, but a robust algorithm which can address most of the relevant cases regarding the FPE and achieve a good performance without tuning its parameters any further. The AMG algorithm and all its components are constructed on the basis of compromise between the general applicability for the FPE and overall efficiency.

References

- [1] E.M. Murman, J.D. Cole, Calculation of plane steady transonic flows, *AIAA* 9 (1) (1971) 114–121.
- [2] A. Brandt, *Multigrid Techniques: 1984 Guide with Applications to Fluid Dynamics*, The Weizmann Institute of Science, Rehovot, Israel, 1984.
- [3] I. Yavneh, *Multigrid techniques for incompressible flows*, PhD thesis, The Weizmann Institute of Science, Rehovot, Israel, 1991.
- [4] D. Sidilkover, A factorizable scheme for the equations of fluid flow, *Appl. Numer. Math.* 41 (2002) 423–426.
- [5] X.-C. Cai, W.D. Gropp, D.E. Keyes, R.G. Melvin, D.P. Young, Parallel Newton–Krylov–Schwarz algorithm for the transonic full potential equation, *SIAM J. Sci. Comput.* 19 (1998) 246–265.
- [6] U. Trottenberg, A. Schuller, C.W. Oosterlee, *Multigrid*, Elsevier Academic Press, 2001.
- [7] P. Wesseling, *An Introduction to Multigrid Methods*, John Wiley & Sons, New York, 1992.
- [8] R.P. Fedorenko, The speed of convergence of one iterative process, *USSR Comput. Math. Phys.* 4 (1964) 227–235.
- [9] N.S. Bakhvalov, On the convergence of a relaxation method with natural constraints on the elliptic operator, *USSR Comput. Math. Math. Phys.* 6 (1966) 101.
- [10] A. Brandt, Multi-level adaptive technique (MLAT) for fast numerical solutions to boundary value problems, in: *Proceedings of 3rd International Conference on Numerical Methods in Fluids Mechanics*, Lecture Notes in Physics, vol. 18, Springer-Verlag, 1973, pp. 82–89.
- [11] A. Brandt, Multi-level adaptive solutions to boundary-value problems, *Math. Comput.* 31 (1977) 333–390.
- [12] W. Hackbusch, *Multigrid Methods and Applications*, Springer, Heidelberg, 1985.
- [13] A. Brandt, S.F. McCormick, J.W. Ruge, Algebraic multigrid (AMG) for sparse matrix equations, in: *Sparsity and Its Applications*, Cambridge University Press, Cambridge, 1984.
- [14] A. Brandt, Algebraic multigrid theory: the symmetric case, *Appl. Math. Comput.* 19 (1986) 23–56.
- [15] K. Stueben, An introduction to algebraic multigrid, Technical report, German National Research Center for Information Technology (GMD), San Diego, 2001.
- [16] J.W. Ruge, K. Stueben, Algebraic multigrid, in: S.F. McCormick (Ed.), *Multigrid Methods*, vol. 3, SIAM, Philadelphia, 1987, pp. 73–130.
- [17] S. Shitrit, D. Sidilkover, A full potential equation solver based on the algebraic multigrid method: elementary applications. *J. Comput. Appl. Math.*, in press.
- [18] S. Shitrit, D. Sidilkover, Toward applying algebraic multigrid to transonic flow problems, *SIAM J. Sci. Comput.* 32 (July) (2010) 2007–2028.
- [19] A. Jameson, Iterative solution of transonic flows over airfoils and wings, including flows at Mach 1, *Pure Appl. Math.* 27 (1974) 283–309.
- [20] S. Shitrit, D. Sidilkover, The discretization of the full potential equation, *Comput. Fluids*, in press.
- [21] S. Shitrit, *Transonic flow computations by algebraic multigrid*. PhD thesis, The Zandman-Slaner Graduate School of Engineering, Tel-Aviv University, Tel-Aviv University, July 2010.
- [22] W.L. Briggs, V.E. Henson, S.F. McCormick, *A Multigrid Tutorial*, 2 ed., Society for Industrial and Applied Mathematics, Philadelphia, 2000.
- [23] R.D. Falgout, An introduction to algebraic multigrid, in: *Computer in Science Engineering*, Lawrence Livermore National Laboratory, 2006.
- [24] Q. Chang, Y. Shu, Wong H. Fu, On the algebraic multigrid, *J. Comput. Phys.* 125 (1996) 279–292.
- [25] A.H. Shapiro, *The Dynamics and Thermodynamics of Compressible Fluid Flow*, vol. 1, John Wiley & Sons, New York, 1953.
- [26] J.D. Anderson, *Modern Compressible Flow: With Historical Perspective*, McGraw-Hill series in aeronautical and aerospace engineering, 2003.
- [27] D. Hawken, 2009. Available from <www.telusplanet.net>.