

The Use of Domain Decomposition in Accelerating the Convergence of Quasi-Hyperbolic Systems

Bernard Parent* and Jean P. Sislian†

This paper proposes an alternate form of the active-domain method [K. Nakahashi and E. Saitoh, Space-Marching Method on Unstructured Grid for Supersonic Flows with Embedded Subsonic Regions, *AIAA J.* **35**, 1280 (1997)] that is applicable to streamwise separated flows. Named the “marching-window”, the algorithm consists of performing pseudo-time iterations on a minimal width subdomain composed of a sequence of cross-stream planes of nodes. The upstream boundary of the subdomain is positioned such that all nodes upstream exhibit a residual smaller than the user-specified convergence threshold. The advancement of the downstream boundary follows the advancement of the upstream boundary, except in zones of significant streamwise ellipticity where a streamwise ellipticity sensor ensures its continuous progress. Compared to the standard pseudo-time marching approach, the marching-window decreases the work required for convergence by up to 24 times for flows with little streamwise ellipticity and by up to 8 times for flows with large streamwise separated regions. Storage is reduced by up to six times by not allocating memory to the nodes not included in the computational subdomain. The marching-window satisfies the same convergence criterion as the standard pseudo-time stepping methods, hence resulting in the same converged solution within the tolerance of the user-specified convergence threshold. The algorithm is not restricted to a discretization stencil and pseudo-time stepping scheme in particular, and is used here with the Yee-Roe scheme and block-implicit approximate-factorization solving the Favre-averaged Navier-Stokes (FANS) equations closed by the Wilcox $k\omega$ turbulence model. The eigenstructure of the FANS equations is also presented.

1. Introduction

THERE is little doubt that the most efficient way to solve supersonic or hypersonic flow with no streamwise ellipticity is through a space-marching method, as numerous extremely efficient marching methods developed over the years can attest (see for example Refs. [1, 2, 3, 4, 5]). The Navier-Stokes equations at supersonic speeds do, however, exhibit some ellipticity in the marching direction through the streamwise viscous terms and the subsonic layer of the boundary layer, and it is necessary for a space-marching method to ignore these mechanisms by solving a reduced set of the original equations of motion, such as the parabolized Navier-Stokes equations (PNS). The PNS are defined here as the equation set obtained from the Navier-Stokes equations by neglecting all viscous terms in the streamwise direction and by modifying the streamwise momentum equation to prevent any pressure disturbance to travel upstream, using characteristics splitting or pressure splitting as suggested by Vigneron *et al.* [1]. The applicability of the space-marching methods is limited to flows with negligible streamwise ellipticity, hence preventing their deployment to many practical flowfields.

The need to tackle streamwise ellipticity prompted the development of the “global iteration” space-marching methods in which a sweep is performed several times on the entire computational domain to permit the upstream propagation of information (see Ref. [6] for a detailed review). Such are characterized, compared to the pseudo-time marching schemes, by a smaller memory requirement due to the storage of temporary variables in one marching plane only and by enhanced wave propagation mechanism in the streamwise direction. The reduced Navier-Stokes (RNS) equations, which are derived from the Navier-Stokes equations by ignoring all streamwise diffusion terms but not altering the momentum convection terms, are usually solved in this manner leading to fast convergence of subsonic/supersonic streamwise unseparated flows [7, 8, 9] and even of viscous/inviscid

*Graduate Student. University of Toronto Institute for Aerospace Studies, 4925 Dufferin St., Downsview ON Canada, M3H 5T6. URL: <http://bernardparent.ca>.

†Professor. University of Toronto Institute for Aerospace Studies, 4925 Dufferin St., Downsview ON Canada, M3H 5T6

interactions creating streamwise separation [10, 11]. In a similar vein, Bardina [12] shows that significant reduction in work is achievable by the use of global marching sweeps to solve the full Navier-Stokes equations for high-speed flows. However, if not limited to some predetermined zones of the computational domain, the global iteration approach loses some performance when solving large reverse flow regions, as the number of sweeps can become excessive due to its dependence on the size of the separation bubble. Further, some computing might be inefficiently allocated to the nodes downstream of the separation bubble, *prior* to its convergence. These deficiencies can be remedied by using a space-marching scheme solving the PNS equations until an elliptic/reverse flow region is encountered, then switching to a global iteration RNS method for the length of the elliptic region, iterating until convergence is reached, and pursuing with the marching PNS scheme (see Miller *et al.* [6] for instance). However, such a strategy forces the solution of the PNS equations in certain regions of the flowfield for which the PNS assumption might induce appreciable errors. The accuracy of the final solution is hence strongly dependent on the ability of the method at predicting correctly which regions of the flowfield can be accurately predicted with the PNS equations, and which regions require the use of the RNS equations.

Recently, a novel approach at solving inviscid supersonic flow with embedded subsonic regions has been proposed [13]. The method, named "active-domain", consists of performing pseudo-time iterations on a small band-like computational domain that advances in the streamwise direction every time the residual of the active-domain near the upstream boundary falls below a user-defined threshold. Using sensors based on the streamwise component of the Mach number, the active-domain boundaries automatically surround any locally subsonic region, on which sufficient iterations are performed to reach steady-state. When the residual inside the subsonic region decreases below the user-defined threshold, the active-domain advances past the subsonic region further downstream. By marching in the streamwise direction, the active-domain results in a decrease in work of up to 10 times compared to standard pseudo-time marching methods for several inviscid problems. However, the ability of the active-domain at solving accurately a streamwise elliptic region is limited by the accuracy of the sensor responsible for the upstream movement of the upstream boundary of the active-domain. Extension of the active-domain method to viscous flow is hampered by the difficulty of formulating a streamwise ellipticity sensor that captures all significant upstream propagating waves while restricting the size of the active-domain to a minimum. Success has been reported in solving viscous flow without streamwise separation by maintaining the active-domain width equal to the height of the boundary layer [14]. However, to the authors' knowledge, the active-domain method has not yet been extended to streamwise separated flows.

This paper proposes an alternate form of the active-domain method, named the "marching-window" algorithm, that is applicable to streamwise separated flows. Similarly to the active-domain, the marching-window performs localized pseudo-time stepping on a subdomain composed of a sequence of cross-stream planes of nodes. The width of the marching-window decreases to only a few planes in regions of quasi-hyperbolic flow and increases to the size of the streamwise-elliptic region when encountered. However, in contrast to the active-domain algorithm, the marching-window is strictly a convergence acceleration technique as it guarantees the residual of all nodes to be below the user-defined threshold when convergence is attained, and hence does not modify the final solution as obtained with the "standard" pseudo-time stepping schemes. This is accomplished by keeping the residual upstream of the marching window subdomain updated at all times, and by positioning the upstream boundary such that the residual of all nodes upstream is below the user-defined threshold. This results in an algorithm that captures all upstream propagating waves affecting the residual significantly. The upstream propagating waves can originate from (but are not necessarily limited to) large subsonic pockets, streamwise separation, streamwise viscous fluxes, or the flux limiters in the streamwise convection flux derivative, for instance. Further, to enhance the performance of the algorithm, a sensor based on the Vigneron splitting [1] is developed to advance the downstream boundary when significant streamwise ellipticity is detected.

Several numerical experiments are presented ranging from the inviscid solution of a supersonic inlet with a blunt leading edge to turbulent shock boundary layer interactions with considerable streamwise flow separation solved with the Favre-averaged Navier-Stokes (FANS) equations closed by the $k\omega$ turbulence model of Wilcox [15]. A time accurate turbulent flow field using dual time stepping is also investigated. A comparison between the marching-window cycle, the active-domain cycle (for the inviscid case only), and the standard pseudo-time marching cycle is made on the basis of CPU time, effective iterations, and storage.

2. Governing Equations

The residual of the full Navier-Stokes equations can be expressed in generalized coordinates in tensor form for any number of dimensions as

$$R = \sum_{i=1}^d \left[\frac{\partial F_i}{\partial X_i} - \sum_{j=1}^d \frac{\partial}{\partial X_i} \left(K_{ij} \frac{\partial G}{\partial X_j} \right) \right] - S, \quad (1)$$

where the minimization of R is sought and where d refers to the number of dimensions. Due to the non-linearity of the system of equations, a fictitious unsteady term $\partial Q / \partial \tau$ is necessary to obtain the right physical root from a given set of initial conditions, *i.e.*

$$\frac{\partial Q}{\partial \tau} = -R. \quad (2)$$

Even though the marching-window method presented in this paper along with the discretization and pseudo-time stepping schemes are not linked to some governing equations in particular, this study focuses on the Favre averaged Navier-Stokes equations (FANS) closed by the Wilcox two-equation $k\omega$ model [15]. The $k\omega$ turbulence model is chosen due to its capability at solving accurately a wide range of realistic flowfields, while maintaining a close resemblance in form to the classical NS equations and avoiding the use of extra low Reynolds number terms (present in the $k\epsilon$ family of turbulence models for instance). Aside their inelegance, additional low Reynolds number terms are detrimental to the performance of the algorithm by increasing the complexity of the source terms which results in bigger meshes necessary to attain grid convergence, and by increasing the stiffness of the governing equations which often translates into increased convergence time. This is avoided by the $k\omega$ modeling which induces a conservative variable Q , a convective flux F_i , and a diffusion term G of

$$Q = \frac{1}{J} \begin{bmatrix} \rho \\ \rho v_1 \\ \vdots \\ \rho v_d \\ \rho E \\ \rho k \\ \rho \omega \end{bmatrix}, \quad F_i = \frac{1}{J} \begin{bmatrix} \rho V_i \\ \rho V_i v_1 + X_{i,1} P^* \\ \vdots \\ \rho V_i v_d + X_{i,d} P^* \\ \rho V_i H \\ \rho V_i k \\ \rho V_i \omega \end{bmatrix}, \quad \text{and} \quad G = \begin{bmatrix} 1 \\ v_1 \\ \vdots \\ v_d \\ T \\ k \\ \omega \end{bmatrix}, \quad (3)$$

where the notation $X_{i,j}$ stands for $\partial X_i / \partial x_j$ and J is the metric Jacobian, both of which obtainable in any dimension following the approach of Viviand [16] and Vinokur [17]. The total energy, enthalpy and effective pressure include molecular and turbulent properties,

$$E = e + k + \frac{1}{2} \sum_{i=1}^d v_i^2, \quad P^* = P + \frac{2}{3} \rho k, \quad \text{and} \quad H = E + \frac{P^*}{\rho}, \quad (4)$$

with e the internal energy of the gas, v_i the velocity component in the Cartesian x_i direction and k the turbulence kinetic energy. Calorically perfect gas assumptions are used to determine the internal energy from the temperature while an ideal gas law is assumed in finding P from ρ and T .

Starting from the tensorial form of the governing equations in Cartesian coordinates, it can be shown that the conductivity matrix K_{ij} corresponds in curvilinear coordinates to

$$K_{ij} = \frac{1}{J} \begin{bmatrix} 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & \mu^* \beta_{ij}^{11} & \dots & \mu^* \beta_{ij}^{1d} & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & \mu^* \beta_{ij}^{d1} & \dots & \mu^* \beta_{ij}^{dd} & 0 & 0 & 0 \\ 0 & \mu^* \sum_k \beta_{ij}^{k1} v_k & \dots & \mu^* \sum_k \beta_{ij}^{kd} v_k & \kappa^* \alpha_{ij} & \mu_k^* \alpha_{ij} & 0 \\ 0 & 0 & \dots & 0 & 0 & \mu_k^* \alpha_{ij} & 0 \\ 0 & 0 & \dots & 0 & 0 & 0 & \mu_\omega^* \alpha_{ij} \end{bmatrix}, \quad (5)$$

where the diffusion coefficients include a molecular and turbulent contribution function of the eddy viscosity $\mu_t = 0.09\rho k/\omega$, molecular viscosity μ and molecular thermal conductivity κ ; that is,

$$\mu^* = \mu + \mu_t, \quad \kappa^* = \kappa + C_P \frac{\mu_t}{0.9}, \quad \mu_k^* = \mu + \frac{\mu_t}{2}, \quad \text{and} \quad \mu_\omega^* = \mu + \frac{\mu_t}{2}, \quad (6)$$

while α and β are function of the metrics only,

$$\alpha_{ij} = \sum_{k=1}^d X_{i,k} X_{j,k} \quad \text{and} \quad \beta_{ij}^{mn} = \alpha_{ij} \delta_{mn}^{Kr} + X_{j,m} X_{i,n} - \frac{2}{3} X_{j,n} X_{i,m}, \quad (7)$$

where δ_{mn}^{Kr} stands for the Kronecker delta, which vanishes should m and n differ, but assumes a value of unity otherwise. The source term S is composed of the sum of the baseline Wilcox $k\omega$ source terms [15] and an unsteady source term,

$$S = \frac{1}{J} \begin{bmatrix} \vdots \\ 0 \\ P_k - \rho k \omega \\ \frac{\omega}{\tilde{k}} \left(\frac{5}{9} P_k - \frac{5}{6} \rho k \omega \right) \end{bmatrix} - \frac{\partial Q}{\partial t}, \quad (8)$$

where the turbulence kinetic energy production term P_k can be shown to correspond to

$$P_k = \sum_{i=1}^d \sum_{j=1}^d \left(-\frac{2}{3} \rho k X_{i,j} \frac{\partial v_j}{\partial X_i} + \sum_{m=1}^d \sum_{n=1}^d \mu^* \beta_{ij}^{mn} \frac{\partial v_m}{\partial X_i} \frac{\partial v_n}{\partial X_j} \right). \quad (9)$$

It is noted that in the Wilcox $k\omega$ model, \tilde{k} is set simply to k which in the freestream is set to a small value to prevent a division by zero. We prefer, however, to specify $k = 0$ in the freestream and, in order to prevent a division by zero in the dissipation rate source term, to define \tilde{k} as

$$\tilde{k} = \max \left[k, \min \left(k_{\text{div}}, \frac{\omega \mu}{\rho} \right) \right], \quad (10)$$

with k_{div} a user-specified constant which is generally set lower than one tenth of the maximum value of k throughout the boundary layer. This is verified numerically not to affect the laminar sublayer but to improve the robustness and efficiency of the integration significantly. The minimum between k_{div} and $\omega \mu / \rho$ is taken so that a clipping occurs *only* in non-turbulent flow regions in which an accurate representation of ω does not affect the accuracy of the flowfield.

3. Discretization

The discretization of all terms in the residual is now presented. The use of tensor form in writing the governing equations in curvilinear coordinates shown in Eq. (1), along with the unique compact K_{ij} matrix described in Eq. (5), simplify greatly the discretization and practical implementation of the viscous terms: less than 100 lines of code are needed to implement the viscous contribution of the residual for all $d \in [1, 2, 3]$. By referring to the discretized form of the derivatives along the X_i coordinate by δ_{X_i} , the discretized residual R_Δ can be written as

$$R_\Delta = \sum_{i=1}^d \left[\delta_{X_i} F_i - \sum_{j=1}^d \delta_{X_i} (K_{ij} \delta_{X_j} G) \right] - S, \quad (11)$$

where the diffusion terms are discretized using second-order accurate centered finite difference stencils which, should $i = j$, give,

$$[\delta_{X_i} (K_{ii} \delta_{X_i} G)]^{X_i} = K_{ii}^{X_i+1/2} (G^{X_i+1} - G^{X_i}) - K_{ii}^{X_i-1/2} (G^{X_i} - G^{X_i-1}), \quad (12)$$

and, should $i \neq j$, give,

$$\begin{aligned} [\delta_{X_i} (K_{ij} \delta_{X_j} G)]^{X_i, X_j} = & \\ & \frac{1}{4} K_{ij}^{X_i+1/2, X_j} (G^{X_i, X_j+1} + G^{X_i+1, X_j+1} - G^{X_i, X_j-1} - G^{X_i+1, X_j-1}) \\ & - \frac{1}{4} K_{ij}^{X_i-1/2, X_j} (G^{X_i-1, X_j+1} + G^{X_i, X_j+1} - G^{X_i-1, X_j-1} - G^{X_i, X_j-1}), \end{aligned} \quad (13)$$

where $K^{X_i+1/2}$ midway between nodes is taken as half the value of K^{X_i+1} and K^{X_i} for example. The convection terms are discretized using a conservative form of the Roe scheme [18], turned second order accurate through a symmetric minmod limiter by Yee *et al.* [19]:

$$[\delta_{X_i} F_i]^{X_i} = \frac{1}{2} \left[F_i^{X_i+1} - F_i^{X_i-1} - \left[\frac{L_i^{-1} |\lambda_i| N_i}{J} \right]^{X_i+\frac{1}{2}} + \left[\frac{L_i^{-1} |\lambda_i| N_i}{J} \right]^{X_i-\frac{1}{2}} \right], \quad (14)$$

where $|\lambda_i|$ stands for the absolute value of the eigenvalues of the convective flux Jacobian $A_i \equiv \partial F_i / \partial Q$, and where

$$N_i^{X_i} = M_i^{X_i} - \text{minmod} \left(M_i^{X_i-1}, M_i^{X_i}, M_i^{X_i+1} \right), \quad (15)$$

$$\text{with } M_i^{X_i} = L_i^{X_i} \left((JQ)^{X_i+1/2} - (JQ)^{X_i-1/2} \right). \quad (16)$$

In the above, the properties at the interface are determined from Roe averaging [18]. For the minmod limiter to be in pseudo control volume form, it is necessary to ensure that all metric terms needed to construct the N matrix at one cell interface are measured at that particular interface.

A small positive value can be added to the eigenvalues to fix the aphysical carbuncle phenomenon originating from the Roe scheme when tackling certain problems, especially blunt bodies. This is usually referred to as “entropy correction”, but is just a convenient way of adding artificial dissipation to the flux discretization [20] and can significantly deteriorate the accuracy of the scheme in turbulent boundary layers (see for example results obtained by Parent and Sislian [21]). Unless otherwise specified, no entropy correction term is used in this paper.

Finally, all partial derivatives of the source terms are discretized using second order accurate three-point stencils, except for the stencil of the time derivative term, which is limited and is set to,

$$\begin{aligned} \delta_t Q = \frac{1}{\Delta t} & \left[Q^t - Q^{t-\Delta t} + \frac{1}{2} \text{minmod} (Q^t - Q^{t-\Delta t}, Q^{t-\Delta t} - Q^{t-2\Delta t}) \right. \\ & \left. - \frac{1}{2} \text{minmod} (Q^t - Q^{t-\Delta t}, Q^{t-\Delta t} - Q^{t-2\Delta t}, Q^{t-2\Delta t} - Q^{t-3\Delta t}) \right], \end{aligned} \quad (17)$$

where the minmod function returns the minimum of its arguments if the arguments are all positive, the maximum if the arguments are all negative, and zero if the arguments are of mixed signs. It is noted that the second order contribution in Eq. (17) is in non-conservative form, but, to the authors’ knowledge, such is unavoidable if no future value of Q (at a time $t + \Delta t$) is included in the stencil and a second-order accurate stencil that results in no spurious oscillations is desired. The non-conservation of the stencil is weak and numerical tests indicate that Eq. (17) performs well.

3.1. Eigenstructure of the Convective Flux Jacobian

Since the Roe scheme is used to discretize the convection derivatives, the determination of the eigenstructure of the convective flux Jacobian $A_i \equiv \partial F_i / \partial Q$ is needed. It can be checked by substitution that the following definition of λ_i ,

$$\lambda_i = \left[V_i, V_i, \rightarrow, V_i + a \hat{X}_i, V_i - a \hat{X}_i, V_i, V_i \right]^D, \quad (18)$$

satisfies the necessary relationship $\det(A_i - w_i I) = 0$ (with w_i any element on the diagonal of λ_i) and is hence a valid eigenvalue matrix. Denoting the flow speed by q , the non-metric effective speed of sound is found to be equal to

$$a = \left(P_\rho + \frac{2}{3}k + P_{\rho E} (H - q^2 - k) \right)^{\frac{1}{2}}, \quad (19)$$

which as might be expected is a function of the kinetic energy of turbulence, contrarily to the “non-turbulent” speed of sound (here denoted by $a_{k=0}$) encountered in the eigenstructure of the convection terms of the molecular Navier-Stokes equations. Dividing both sides of Eq. (19) by $a_{k=0}$ and after some reformatting, the normalized speed of sound can be shown to be equal to

$$\bar{a} \equiv \frac{a}{a_{k=0}} = \left(1 + \frac{P_{\rho E} + 1}{3} M_t^2 \right)^{\frac{1}{2}}, \quad (20)$$




FIGURE 1. Normalized speed of sound $\bar{a} = a/a_{k=0}$ versus the turbulent Mach number $M_t = \sqrt{2k}/a_{k=0}$, for a calorically perfect gas according to Eq. (20).

where for a perfect gas, $P_{\rho E} + 1$ is equal to the ratio of the specific heats, γ . Figure 1 shows the relationship between \bar{a} and the turbulent Mach number for $\gamma = 7/5$ and $\gamma = 5/3$. As the turbulent Mach number increases, its influence on \bar{a} becomes more predominant due to the relative speed of the turbulent vortices, with respect to the average vortex speed of displacement, gradually overtaking the thermodynamic sound speed as the information propagation mechanism. For a perfect or real gas, the derivatives of pressure with respect to the mass-weighted total energy and density are equal to

$$P_{\rho E} = \frac{1}{\rho \frac{\partial}{\partial P} e(P, \rho)} \quad \text{and} \quad P_{\rho} = -P_{\rho E} \left(E - q^2 - k + \rho \frac{\partial}{\partial \rho} e(P, \rho) \right), \quad (21)$$

with e the internal energy of the gas and where it is assumed that any thermodynamic property can be obtained from only two others. The right eigenvectors are not unique, and each column of the matrix can be multiplied by a constant other than 0; here, we choose the multiplying constant of each column such as to keep the same units along each row, except for the last column which is further multiplied by a^2/ω which is found to result in faster convergence:

$$L_i^{-1} = \begin{bmatrix} 1 & 0 & \rightarrow & 1 & 1 & 0 & 0 \\ v_1 & l_i^{1,1} a & \rightarrow & v_1 + \frac{a X_{i,1}}{\widehat{X}_i} & v_1 - \frac{a X_{i,1}}{\widehat{X}_i} & 0 & 0 \\ \vdots & \vdots & \searrow & \vdots & \vdots & \vdots & \vdots \\ v_d & l_i^{d,1} a & \rightarrow & v_d + \frac{a X_{i,d}}{\widehat{X}_i} & v_d - \frac{a X_{i,d}}{\widehat{X}_i} & 0 & 0 \\ H - \frac{a^2}{P_{\rho E}} & \sum_j l_i^{j,1} a v_j & \rightarrow & H + \frac{a V_i}{\widehat{X}_i} & H - \frac{a V_i}{\widehat{X}_i} & a^2 - \frac{2a^2}{3P_{\rho E}} & 0 \\ k & 0 & \rightarrow & k & k & a^2 & 0 \\ \omega & 0 & \rightarrow & \omega & \omega & 0 & a^2 \end{bmatrix} \quad (22)$$

The exactness of the right eigenvectors can be readily verified by the relation $\lambda_i = L_i A_i L_i^{-1}$. Note that the columns of the right eigenvectors containing $l_i^{m,n}$ are not needed in one dimension, while in two dimensions, $l_i^{m,n}$ takes on the form

$$l_i^{m,1} = (-1)^{m+1} X_{i,m+1} / \widehat{X}_i, \quad (23)$$

and in three dimensions, becomes,

$$l_i^{m,1} = \frac{X_{i,m+2} - X_{i,m+1}}{\left[\sum_{j=1}^3 (X_{i,j+1} - X_{i,j})^2\right]^{1/2}}, \quad l_i^{m,2} = \frac{X_{i,m+1}l_i^{m+1,1} - X_{i,m+2}l_i^{m+2,1}}{\widehat{X}_i}. \quad (24)$$

In the above, \widehat{X}_i corresponds to the magnitude of all derivatives of X_i , that is,

$$\widehat{X}_i = \left(\sum_{j=1}^d X_{i,j}^2\right)^{\frac{1}{2}}. \quad (25)$$

4. Pseudo-Time Integration

Using implicit Euler pseudo-time marching the delta form of the discretized equations can be shown to correspond to

$$\frac{\Delta^n Q}{\Delta \tau} + \sum_{i=1}^d \left[\delta_{X_i} \Delta^n F_i - \sum_{j=1}^d \delta_{X_i} (K_{i,j} \delta_{X_j} \Delta^n G) \right] - \Delta^n S = -R_\Delta, \quad (26)$$

where to minimize storage requirements and the inversion effort the LHS is approximated using a multiplication of one-dimensional operators based on a block-implicit approximate factorization algorithm [22, 23] and a linearization strategy of the viscous terms by Chang and Merkle [8]:

$$\left[\prod_{i=1}^d \left(I + \Delta \tau \overline{\delta_{X_i} A_i} - \Delta \tau \sum_{j=1}^d \delta_{X_i} (K_{i,j} \delta_{X_j} B) - \Delta \tau \delta_{1i}^{\text{Kr}} C^- \right) \right] \Delta^n Q = -\Delta \tau R_\Delta, \quad (27)$$

where δ_{1i}^{Kr} is the Kronecker delta, B the linearization Jacobian of the viscous terms ($B \equiv \partial G / \partial Q$), and C_i^- the linearization Jacobian of the negative source terms ($\partial S^- / \partial Q$) for the $i = 1$ sweep but ignored for the other sweeps. Only the negative source terms are linearized to ensure the stability of the implicit algorithm [24] and are set to

$$S^- = \frac{1}{J} \begin{bmatrix} \vdots \\ -\rho k \omega \\ -\frac{5}{6} \rho k \omega^2 / \widetilde{k} \end{bmatrix} - \frac{\partial Q}{\partial t}. \quad (28)$$

The term $\overline{\delta_{X_i} A_i}$ is symbolic and stands for the linearization of the first-order Roe scheme with the Roe Jacobian locally frozen. The use of a fully linearized Roe scheme is shown in Batten *et al.* [20] not to decrease the number of iterations needed for convergence for several test problems (in some cases it is even detrimental) while requiring more work per iteration than the frozen Jacobian approach. Hence, the equation to solve at each node for the i th sweep can be written as

$$\begin{aligned} & \left[-K_{ii}^{X_i - \frac{1}{2}} B^{X_i - 1} - \frac{(J^{-1} L^{-1} |\lambda| L)_i^{X_i - \frac{1}{2}}}{2(J^{-1})^{X_i - 1}} - \frac{A_i^{X_i - 1}}{2} \right] \Delta \widetilde{Q}_i^{X_i - 1} + \left[\frac{I}{\Delta \tau^{X_i}} - \delta_{1i} C^{-X_i} \right. \\ & + \left(K_{ii}^{X_i - \frac{1}{2}} + K_{ii}^{X_i + \frac{1}{2}} \right) B^{X_i} + \frac{(J^{-1} L^{-1} |\lambda| L)_i^{X_i - \frac{1}{2}} + (J^{-1} L^{-1} |\lambda| L)_i^{X_i + \frac{1}{2}}}{2(J^{-1})^{X_i}} \left. \right] \Delta \widetilde{Q}_i^{X_i} + \\ & \left[-K_{ii}^{X_i + \frac{1}{2}} B^{X_i + 1} - \frac{(J^{-1} L^{-1} |\lambda| L)_i^{X_i + \frac{1}{2}}}{2(J^{-1})^{X_i + 1}} + \frac{A_i^{X_i + 1}}{2} \right] \Delta \widetilde{Q}_i^{X_i + 1} = \frac{I}{\Delta \tau^{X_i}} \Delta \widetilde{Q}_{i-1}^{X_i} \end{aligned} \quad (29)$$

where $\Delta \widetilde{Q}_0^{X_i} = -\Delta \tau^{X_i} R_\Delta^{X_i}$ and the total flux increment ΔQ^{X_i} is set to $\Delta \widetilde{Q}_d^{X_i}$.

It is emphasized that the success of approximate factorization relies on the degree of invariance of the linearization matrices, deterring the inclusion of a linearized form of the minmod limiter on the implicit side. Numerical experiments show that a "switch" type of algorithm on the implicit side might induce erratic patterns in the convergence history sometimes preventing a converged solution altogether. For similar reasons, the implicit treatment of the cross diffusion terms is not recommended as their linearization necessarily involves spatial derivatives which are subject to change from iteration to iteration.




FIGURE 2. Examples of the distribution of the node types for a backward facing step (left) and a two-element airfoil (right); the number “3” represents a wall condition, “0” inflow, “1” outflow, while “+” represents inner nodes and “.” is used for unactivated nodes.

Block-implicit approximate-factorization is chosen here as it is still one of the most used techniques in solving the compressible Navier-Stokes equations in the hypersonic range. However, it is noted that the use of a different pseudo-time marching algorithm (such as DDADI [12], MAF(k) [25], or LUSGS [26]) has been observed by the authors not to affect the performance gains obtained with the marching-window algorithm presented herein.

4.1. Local Pseudo-Time Step

One commonly used acceleration technique is local pseudo-time stepping based on the CFL condition which results in a wave traveling speed of one node per iteration for convection dominated flows. However, in multiple dimensions, each dimension assumes a different CFL condition and one faces the dilemma of specifying a wave traveling speed proportional to the dimension exhibiting the lowest CFL condition, commonly referred to as a minimum CFL based local time step, or to the dimension exhibiting the highest CFL condition which is referred to as a maximum CFL based local time step. A formulation including both the minimum and maximum CFL based approaches can take the form

$$\Delta\tau = \text{CFL} \max_{i=1}^d \left(\frac{1}{|V_i| + a\hat{X}_i} \right)^\sigma \min_{i=1}^d \left(\frac{1}{|V_i| + a\hat{X}_i} \right)^{1-\sigma}, \quad (30)$$

where a σ varying between 0 and 1 induces a time step of a magnitude situated respectively between a minimum and a maximum CFL based time step. While it is acknowledged that for viscous dominated regions, a local time step based on the Von Neumann number (VNN) would result in a more equitable wave propagation which might translate into faster convergence, for the purposes of this paper Eq. (30) is used exclusively.

5. Boundary Conditions

A multiblock stratagem is generally required when tackling complex geometries with a structured mesh, but it can significantly complicate the implementation of the domain decomposition algorithms presented herein for reasons that shall become apparent shortly. As a substitute to using multiple blocks connected to the geometry and to one another through their outer edges (or planes in 3D), any node that is part of the computational domain is allowed to be either a boundary, inner or inactive node. Although not as multipurpose as the multiblock, such an approach can be used to solve a wide variety of flowfields while retaining all the simplicity of a single block. Figure 2 shows, for example, how the node types would be distributed for a backward facing step and a two-element airfoil.

Zeroth order extrapolation polynomials are used to obtain the properties from the adjacent inner node at the supersonic outflow boundary (hereafter referred to simply as outflow boundary), while the properties at the supersonic inflow (hereafter referred to as inflow), are unaltered in pseudo-time. At the symmetry boundary node, a first order extrapolation polynomial of the form

$$\psi^X = \frac{4}{3}\psi^{X+1} - \frac{1}{3}\psi^{X+2}, \quad (31)$$




FIGURE 3. Example of computational domain and subdomain notation in two dimensions; the computational domain limits are denoted by the superscripts E and S.

is employed to obtain P^* , k , ω , ρ and the velocity components tangent to the surface, while the perpendicular velocity component is set to zero. At the wall, the turbulence kinetic energy and the velocity are fixed to zero, while the effective pressure and temperature (in the case of an adiabatic wall) are extrapolated as in Eq. (31). Also, following Wilcox [27], the dissipation rate at the wall is specified to

$$\omega_w = \frac{36}{5} \frac{\mu}{\rho d_w^2}, \quad (32)$$

with d_w the distance between the wall node and its nearest neighbor.

It is well known that an implicit treatment of the boundary nodes results in less prohibitive restrictions on the pseudo-time step size for some problems, but when solving strong shock waves or other highly non-linear phenomena it is not uncommon for the time step size to be limited in any case by the flow physics, even if the time stepping scheme can be shown to be Von Neumann unconditionally stable (see the chapter on non linear stability in Laney [28]). Moreover, experience shows that treating the boundary conditions presented herein in an explicit manner does not restrict the size of the local time step more than implicit boundary conditions would. For all numerical experiments presented, an explicit treatment of the boundary nodes is chosen.

6. Domain Decomposition Algorithms

While domain decomposition is generally used for parallel computing purposes or used to enable the implementation of different discretization/integration methods in different subdomains, it is utilized here as a means to accelerate the convergence of quasi-hyperbolic systems. We define as quasi-hyperbolic a system of equations 1) which is elliptic, 2) where some of the terms, but not all, can be regrouped to form a hyperbolic set of equations, and 3) whose solution is very close to the solution of the hyperbolic set of terms. For instance, the steady-state Navier-Stokes equations in the hypersonic regime away from the surfaces would exhibit a weak influence of the diffusion terms (responsible for the ellipticity of the system) on the solution compared to the convection terms (the hyperbolic set) and would hence be classified as quasi-hyperbolic. Similarly, a quasi-parabolic system is defined as a system of equations 1) which is elliptic, 2) where some of the terms, but not all, can be regrouped to form a set of parabolic equations, and 3) whose solution is very close to the solution of the parabolic set of terms. The Favre averaged Navier-Stokes equations closed by the $k\omega$ model solved at steady-state over a turbulent flat plate would be termed quasi-parabolic, as the streamwise diffusion terms and the upstream component of the convection terms play a negligible role compared to the other terms.

The acceleration techniques presented in this paper are aimed at reducing the work needed to solve quasi-hyperbolic or quasi-parabolic systems through the use of domain decomposition. Nonetheless, the effectiveness of the methods is not limited to entirely quasi-hyperbolic/parabolic systems and extends to systems where some regions are quasi-hyperbolic/parabolic and others strictly elliptic. It is emphasized that domain decomposition is used here solely as a convergence acceleration technique, and *does not* modify the discretized residual, the time stepping schemes, and the convergence criterion, except in the case of the active-domain method. Convergence is attained independently of the acceleration technique when

$$\xi \leq \xi_{\text{verge}} \quad \forall \text{ inner nodes}, \quad (33)$$

with ξ a convergence criterion based on the maximum between the discretized continuity and energy residuals,

$$\xi \equiv \max \left(\frac{|R_{\Delta}^{\text{continuity}}|}{J^{-1}\rho}, \frac{|R_{\Delta}^{\text{energy}}|}{J^{-1}\rho E} \right), \quad (34)$$

which is divided by Q to obtain units involving only pseudo-time (*i.e.* $\frac{1}{s}$). The user-defined convergence threshold ξ_{verge} is typically given a value of $100 \frac{1}{s}$, yet this value is not universal and is dependent on the flowfield at hand. Based on dimensional analysis arguments, ξ_{verge} can be thought of as the inverse of a time scale common for all nodes, which we formulate in terms of the free stream flow speed and a characteristic length, *i.e.*

$$\xi_{\text{verge}} \sim \frac{1}{10} \frac{q_{\infty}}{L_c}, \quad (35)$$

where an analogy can be made to the time needed to obtain steady-state flow using an experimental setup. The characteristic length L_c can be taken as the length of the domain for instance. It is noted, however, that the efficiency of the domain decomposition methods presented herein is dependent on the precision of ξ as a convergence criterion, and since this varies from one flow problem to the next, it might not always be possible to achieve at first the proper compromise between optimal convergence rate and acceptable accuracy by using Eq. (35).

Identifying the limits of the computational domain by X_i^s and X_i^e with $i \in [1, \dots, d]$ and the limits of a subdomain by X_i^s and X_i^e , with $i \in [1, \dots, d]$ the region spanned by the subdomain is referred to by the notation $\|X_i^s \Leftrightarrow X_i^e\|_{\forall i}$, as shown in Fig. 3. For a subdomain with limits different from the computational domain limits in only one dimension the notation $\|X_n^s \Leftrightarrow X_n^e\|_n$ is employed, where it is implied that the limits in the dimensions other than the n th do not differ from those of the computational domain (see Fig. 3). Also, $\|X_n^s\|_n$ is a shortcut that stands for the subdomain $\|X_n^s \Leftrightarrow X_n^e\|_n$. A property that is used in conjunction with the domain decomposition algorithms is the number of nodes of dependence of the discretized residual, b_r , which is defined as

$$b_r \equiv \begin{cases} \text{the maximum number of nodes on which the discretized} \\ \text{residual depends on each side of the center node,} \\ \text{or,} \\ \text{half the maximum discretization stencil point minus one} \\ \text{if the stencil is symmetric.} \end{cases} \quad (36)$$

For example, the minmod TVD discretization stencil (which is the longest of all stencils contained in the residual) would give $b_r = 2$ but should a first order Roe scheme be employed instead, then b_r would be set to one. Similarly, the number of nodes of dependence of the boundary nodes is defined as

$$b_b \equiv \begin{cases} \text{the maximum number of nodes any boundary node depends} \\ \text{on along one direction,} \end{cases} \quad (37)$$

which is set to 2, since the properties at the boundary nodes are extrapolated from at most 2 inner nodes using a blend of zeroth and first order extrapolation polynomials.

When the nodes comprised in the subdomain $\|X_i^s \Leftrightarrow X_i^e\|_{\forall i}$ are updated in pseudo-time, then it follows from the definition of b_b that the boundary nodes situated inside $\|X_i^s - b_b \Leftrightarrow X_i^e + b_b\|_{\forall i}$ must be updated. The residual, which depends on both inner and boundary nodes must then be updated between $\|X_i^s - b_b - b_r \Leftrightarrow X_i^e + b_b + b_r\|_{\forall i}$. In many cases where there are no boundary nodes situated in the region $\|X_i^s - b_b \Leftrightarrow X_i^e + b_b\|_{\forall i}$, it is sufficient to update the residual in $\|X_i^s - b_r \Leftrightarrow X_i^e + b_r\|_{\forall i}$. For all methods presented in this paper, however, this shortcut is not implemented.

6.1. Standard Cycle

The “standard cycle” here implies the usual way of updating the solution in pseudo-time, by first finding the residual for all nodes and then updating the solution. The algorithm can be written in the following steps:

1. update the boundary nodes in the domain $\| X_i^s \Leftrightarrow X_i^E \|_{v_i}$,
2. update the residual in the domain $\| X_i^s \Leftrightarrow X_i^E \|_{v_i}$,
3. update Q (by pseudo-time stepping) in the domain $\| X_i^s \Leftrightarrow X_i^E \|_{v_i}$, and,
4. convergence is attained when $\xi \leq \xi_{\text{verge}}$ in the domain $\| X_1^s \Leftrightarrow X_1^E \|_{v_i}$.

6.2. Multizone Cycle

One strategy towards improving the standard cycle is to divide the computational domain into a number of non-overlapping zones of approximately equal size and to update in pseudo-time only the zones in which $\xi > \xi_{\text{verge}}$. This stratagem has been previously employed by Sawley *et al.* [29] as a convergence acceleration technique for supersonic flows, but where the computational domain is split into several blocks, instead of several zones. Note that a “zone” is defined as a computational domain region that can be bounded by boundary and/or inner nodes (see for example Ref. [30]), while a “block” is defined as a region delimited by boundary nodes only. The zone length in each dimension is set to at most ϕ_1 , a user specified constant usually given a value of 20. At each iteration, should the maximum ξ inside each zone be greater than the user-specified threshold ξ_{verge} , the inner nodes up to the zone boundaries are updated in pseudo-time, followed by the update of the boundary nodes up to the zone boundaries expanded by b_b , and the update of the residual up to the zone boundaries expanded by $b_b + b_r$. The residual and properties of all other nodes of the computational domain are not altered. Prior to the first iteration, the computational domain is divided into a number of non-overlapping zones of length in each dimension no greater than ϕ_1 with each zone z defined by the subdomain $\| X_i^{z,s} \Leftrightarrow X_i^{z,e} \|_{v_i}$. Then, at each iteration, the following steps are performed:

1. for each zone z , update Q (by pseudo-time stepping) in the subdomain $\| X_i^{z,s} \Leftrightarrow X_i^{z,e} \|_{v_i}$ if $\xi > \xi_{\text{verge}}$ in the subdomain $\| X_i^{z,s} \Leftrightarrow X_i^{z,e} \|_{v_i}$,
2. for each zone z , update the boundary nodes in the subdomain $\| X_i^{z,s} - b_b \Leftrightarrow X_i^{z,e} + b_b \|_{v_i}$ if $\xi > \xi_{\text{verge}}$ in the subdomain $\| X_i^{z,s} \Leftrightarrow X_i^{z,e} \|_{v_i}$,
3. for each zone z , update the residual in the subdomain $\| X_i^{z,s} - b_b - b_r \Leftrightarrow X_i^{z,e} + b_b + b_r \|_{v_i}$ if $\xi > \xi_{\text{verge}}$ in the subdomain $\| X_i^{z,s} \Leftrightarrow X_i^{z,e} \|_{v_i}$, and,
4. convergence is attained when $\xi \leq \xi_{\text{verge}}$ in the domain $\| X_1^s \Leftrightarrow X_1^E \|_{v_i}$.

It is noted that the multizone cycle ensures the residual on all nodes to be up to date after each iteration but, due to the non self-starting property of this cycle, it is necessary to compute the residual on the entire domain before the first iteration is performed.

6.3. Active-Domain Cycle

The active-domain is an algorithm aimed at decreasing the work needed for convergence of supersonic inviscid flow [13] and refers to a band-like computational domain marching in the flow direction in which localized pseudo-time stepping is performed. The domain width automatically adjusts to the size of subsonic regions when encountered by monitoring the streamwise component of the Mach number, as shown in Fig. 4. In our notation, the active-domain algorithm can be written as follows, denoting the left boundary of the computational window by X_1^s and the right boundary by X_1^e :

1. update the boundary nodes in the subdomain $\| X_1^s \Leftrightarrow X_1^e \|_1$,
2. update the residual in the subdomain $\| X_1^s \Leftrightarrow X_1^e \|_1$,
3. update Q (by pseudo-time stepping) in the subdomain $\| X_1^s \Leftrightarrow X_1^e \|_1$,
4. redefine the active-domain boundaries:

- (a) if $M_1 < 1.001$ for any node in the subdomain $\| X_1^s \Leftrightarrow X_1^e \|_1$ then decrease X_1^e by one,




FIGURE 4. Schematic of the active-domain cycle with the upstream and downstream boundaries in equilibrium surrounding an embedded subsonic region; ϕ_0 is the minimum width of the residual monitor subdomain and ϕ_3 is the minimum width of the active-domain (when no subsonic region is present).

- (b) if $M_1 < 1.001$ for any node in the subdomain $\|X_1^e - (\phi_3 - \phi_0) \Leftrightarrow X_1^e\|_1$ then increment X_1^e by one,
- (c) if $\xi \leq \xi_{\text{verge}}$ for all nodes in the subdomain $\|X_1^s \Leftrightarrow X_1^e - (\phi_3 - \phi_0)\|_1$ then increment X_1^e by ϕ_0 and set $X_1^s = X_1^e - \phi_3$, and,

5. convergence is attained when $\xi \leq \xi_{\text{verge}}$ for all nodes in the subdomain $\|X_1^s \Leftrightarrow X_1^e\|_1$ and when $X_1^e = X_1^E$.

The size of the residual-monitor region ϕ_0 and the size of the active-domain ϕ_3 are user-specified constants typically given values of 4 and 9 respectively. It is emphasized that the active-domain is restricted to inviscid flow due to the “ellipticity sensors” in Steps 4a and 4b being based on the streamwise component of the Mach number. For viscous flows, this would effectively enlarge the active-domain to the size of any object due to the vanishing value of the Mach number in the vicinity of a wall. Aside from being restricted to inviscid flow, the active-domain algorithm does not guarantee that $\xi \leq \xi_{\text{verge}}$ for all nodes of the computational domain when convergence is attained. This is due to the assumption in Step 4a that streamwise ellipticity is present locally only when the streamwise component of the Mach number is less than 1. For inviscid flow, this is exactly true if the discretization stencil of the streamwise convection derivative is upwinded (such as the first-order accurate Roe scheme for instance), but is not true for the Yee-Roe scheme used herein due to the Yee flux limiter being a function of downstream nodes, even when the flow is locally supersonic. Therefore, when used in conjunction with a flux-limiter inducing streamwise ellipticity in supersonic flow, the active-domain does not meet the necessary convergence criterion for a well-posed acceleration technique [as stated previously in Eq. (33)].

6.4. Marching-Window Cycle

An alternate form of the active-domain cycle that permits the solution of viscous streamwise separated flows and satisfies the convergence criterion of Eq. (33) is here presented. Named the marching-window, the algorithm differs from the active-domain on three points, namely 1) a dynamic outflow boundary is forced at the downstream boundary of the marching-window (see Fig. 5), 2) the ellipticity sensor responsible for a shift downstream of the downstream boundary of the marching-window is based on a Vigneron splitting of the streamwise pressure derivative instead of the streamwise component of the Mach number, and 3) the upstream boundary of the marching-window is positioned such that $\xi \leq \xi_{\text{verge}}$ for all nodes upstream, instead of being a function of a residual monitor region and a streamwise ellipticity sensor based on the streamwise component of the Mach number. At the first iteration, the upstream boundary of the marching-window is set to the upstream boundary of the

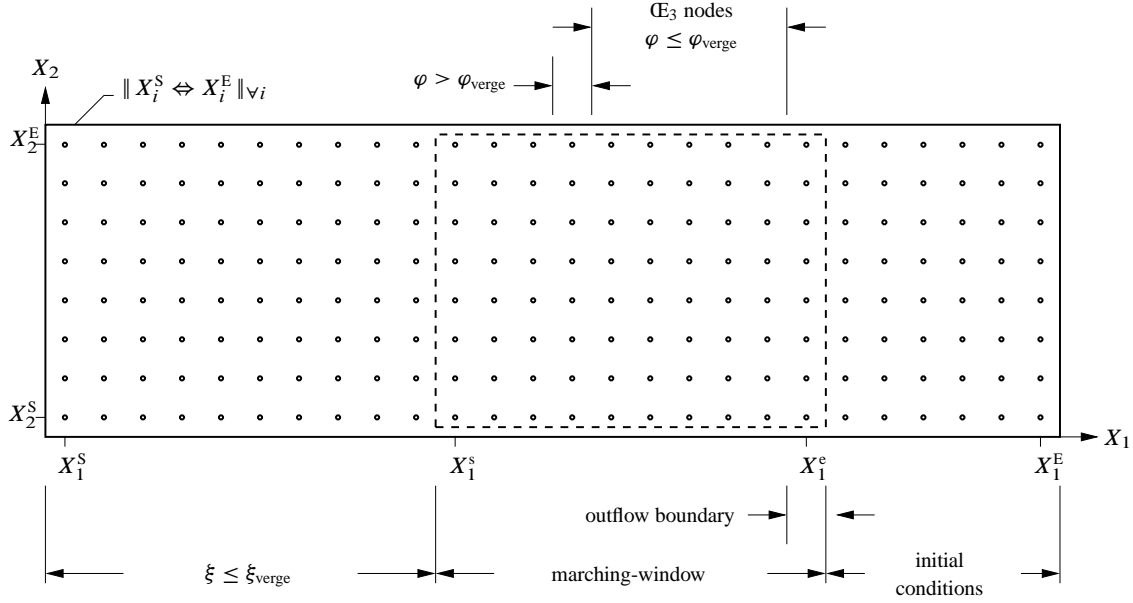


FIGURE 5. Schematic of the marching-window cycle with the upstream and downstream boundaries in equilibrium surrounding an embedded streamwise elliptic region which is bounded upstream by the condition $\xi \leq \xi_{\text{verge}}$ and downstream by the condition $\phi \leq \phi_{\text{verge}}$; a dynamic outflow boundary condition is forced on all inner nodes in $\|X_1^e\|_1$.

computational domain with the downstream boundary of the marching-window separated from the upstream boundary by b_b nodes. Denoting the upstream boundary of the marching-window by X_1^s and the downstream boundary by X_1^e , the marching-window cycle can be written as:

1. update Q (by pseudo-time stepping) in the subdomain $\|X_1^s \Leftrightarrow X_1^e\|_1$,
2. update the boundary nodes in the subdomain $\|X_1^s - b_b \Leftrightarrow X_1^e\|_1$,
3. update the residual (hence, ξ) in the subdomain $\|X_1^s - b_b - b_r \Leftrightarrow X_1^e\|_1$,
4. redefine the marching-window boundaries:
 - (a) find the maximum value for X_1^s such that $\xi \leq \xi_{\text{verge}}$ for all nodes in the subdomain $\|X_1^s \Leftrightarrow X_1^s - 1\|_1$,
 - (b) every ϕ_2 iterations, if $\phi > \phi_{\text{verge}}$ for any node in the subdomain $\|X_1^e - \phi_3 \Leftrightarrow X_1^e\|_1$ or if $X_1^s > X_1^e - \phi_3$ then 1) increment X_1^e by one, 2) update the boundary nodes in the subdomain $\|X_1^e - 1 - b_b \Leftrightarrow X_1^e\|_1$ and 3) update the residual in the subdomain $\|X_1^e - 1 - b_b - b_r \Leftrightarrow X_1^e\|_1$, and,
5. convergence is attained when $\xi \leq \xi_{\text{verge}}$ for all nodes in the subdomain $\|X_1^s \Leftrightarrow X_1^e\|_1$ and when $X_1^e = X_1^E$.

The marching-window cycle is not self-starting and it must be ensured that the residual is updated for all nodes part of the computational window before the first iteration.

The ability of the marching-window at satisfying the convergence criterion of Eq. (33) lies in Steps 1-3, where Q is updated in pseudo-time in Step 1 *before* determining the residual in Step 3. Once Q is updated in Step 1 on the subdomain $\|X_1^s \Leftrightarrow X_1^e\|_1$, since a boundary node depends on at most b_b neighbors, it is sufficient to update the boundary nodes in the subdomain $\|X_1^s - b_b \Leftrightarrow X_1^e\|_1$ to guarantee that all boundary nodes upstream of X_1^e are up to date after Step 2. Once the boundary nodes have been updated in the subdomain $\|X_1^s - b_b \Leftrightarrow X_1^e\|_1$, since the discretized residual depends on at most b_r neighbors, it is sufficient to update the residual in the subdomain $\|X_1^s - b_b - b_r \Leftrightarrow X_1^e\|_1$ to guarantee that the residual of all nodes upstream of X_1^e are up to date after Step 3. Since ξ is a function of the residual, ξ upstream of X_1^e is up to date, and the upstream boundary of the marching-window X_1^s can be positioned correctly in Step 4a by ensuring for all nodes upstream of X_1^e that $\xi \leq \xi_{\text{verge}}$. This serves two purposes: 1) the convergence criterion of Eq. (33) is satisfied if convergence is attained in Step 5, and 2) the upstream boundary of the marching-window moves upstream for any upstream propagating wave that affects the residual significantly and raises ξ above the user-defined convergence threshold ξ_{verge} . Contrarily to the active-domain, the upstream propagating wave

is not limited to locally subsonic flow but includes all *significant* streamwise elliptic phenomena, such as streamwise separated flow, streamwise viscous derivatives, or flux limiters in the streamwise convection derivative, for instance.

Step 4b advances the marching-window downstream boundary when the width of the window is smaller than a user-specified constant ϕ_3 , or when the streamwise ellipticity sensor φ is greater than the user-specified constant φ_{verge} for any node part of the subdomain $\|X_1^e - \phi_3 \Leftrightarrow X_1^e\|_1$. The streamwise ellipticity sensor φ is here chosen as the component of the streamwise convection derivative inducing a streamwise ellipticity. This is derived, following an approach by Vigneron *et al.* [1], by multiplying by ζ the effective pressure terms in the momentum fluxes part of the streamwise convection flux F_1 . The eigenvalues of the streamwise convective flux Jacobian with ζ frozen can then be shown to correspond to:

$$\tilde{\lambda}_1 = \left[V_1, V_1, \rightarrow, V_1 + \frac{1}{2} V_1 P_{\rho E} (1 - \zeta) + \tilde{a}_1 \hat{X}_1, V_1 + \frac{1}{2} V_1 P_{\rho E} (1 - \zeta) - \tilde{a}_1 \hat{X}_1, V_1, V_1 \right]^D,$$

$$\text{with } \tilde{a}_1 = \left\{ \frac{1}{4} [V_1 P_{\rho E} (1 - \zeta)]^2 + \hat{X}_1^2 \zeta \left[P_\rho + \frac{2}{3} k + P_{\rho E} (H - k - q^2) \right] \right\}^{\frac{1}{2}}.$$

Then, for all the eigenvalues to share the same sign (a necessary condition for an hyperbolic system), it is required that

$$\zeta = \min \left(1, V_1^2 \frac{1 + P_{\rho E}}{\hat{X}_1^2 a^2 + V_1^2 P_{\rho E}} \right) = \min \left(1, \frac{M_1^2 (1 + P_{\rho E})}{1 + M_1^2 P_{\rho E}} \right) \quad (38)$$

where the streamwise Mach number M_1 corresponds to $V_1/a\hat{X}_1$. If multiplying by ζ the pressure derivative terms part of the momentum components of $\partial F_1/\partial X_1$ results in a hyperbolic system, it follows that the component of the streamwise convection derivative which induces a streamwise ellipticity is $(1 - \zeta)$ times the pressure derivative terms part of the momentum components of $\partial F_1/\partial X_1$. The product is then normalized with ρa to obtain units of inverse pseudo-time:

$$\varphi \equiv \frac{1}{\rho a} \left\{ \sum_{j=1}^d \left[(1 - \zeta) X_{1,j} \frac{\partial P^*}{\partial X_1} \right]^2 \right\}^{\frac{1}{2}} = \frac{\hat{X}_1}{\rho a} \max \left(0, \frac{1 - M_1^2}{1 + P_{\rho E} M_1^2} \right) \left| \frac{\partial P^*}{\partial X_1} \right|. \quad (39)$$

The ellipticity sensor φ makes two important assumptions: 1) the streamwise ellipticity originating from the streamwise viscous derivative terms and the flux limiter part of the streamwise convection derivative is assumed negligible, and 2) at the point where φ is evaluated, the solution is assumed to be converged to steady-state. The first assumption is remedied by introducing a minimum width of the marching-window, ϕ_3 , which is typically given a value ranging from 9 to 15. The second assumption can lead to some performance degradation of the marching-window when the flow near the downstream boundary is far from convergence. For this reason, the user-adjustable parameter ϕ_2 is introduced in Step 4a, with the consequence of evaluating φ every ϕ_2 iterations only. Therefore, a high value given to ϕ_2 helps in ensuring a more converged solution near the downstream boundary, and reduces the error in the ellipticity sensor φ due to temporarily non steady-state flow. It is suggested to give the ellipticity sensor threshold, φ_{verge} , a value of about 100 times the one given to ξ_{verge} , that is,

$$\varphi_{\text{verge}} \sim 10 \frac{q_\infty}{L_c}, \quad (40)$$

with L_c a characteristic length of the system. In Step 4b, after the downstream boundary of the marching-window is advanced by one station, the update of the boundary nodes in the subdomain $\|X_1^e - 1 - b_b \Leftrightarrow X_1^e\|_1$ and of the residual in the subdomain $\|X_1^e - 1 - b_b - b_r \Leftrightarrow X_1^e - 1\|_1$ is necessary to ensure that the residual is properly updated in the marching-window, which is necessary for Step 1 to be performed correctly at the following iteration.

While the user-definable constants ϕ_2 , ϕ_3 , and φ_{verge} affect the performance of the marching-window cycle as a convergence acceleration technique, they do *not* affect the accuracy of the solution when convergence is attained due to the convergence criterion of Eq. (33) being satisfied.

6.5. Marching-Window / Multizone Cycle

The performance of the marching-window algorithm can be enhanced by introducing multizone decomposition inside the marching-window. Before each iteration, the marching-window subdomain $\|X_1^e \Leftrightarrow X_1^e\|_1$ is decomposed into several zones of length no more than ϕ_1 nodes in each dimension. Then, Steps 1-3 of the marching-window cycle (see Section 6.4.) are replaced by Steps 1-3 of the multizone cycle (see Section 6.2.).

TABLE 1.

Effective iteration count, work, and storage comparison at a CFL number of unity, for the blunt leading edge inviscid supersonic inlet case.

cycle	128 × 64 nodes			512 × 256 nodes		
	iter.	work	stor.	iter.	work	stor.
marching-window / multizone	34.4	1.0	0.11	44.9	22.2	0.91
marching-window	40.4	1.1	0.11	86.8	38.9	0.91
active-domain	55.6	1.3	0.10	102.2	39.2	0.81
multizone	158.5	3.8	1.0	318.6	131.6	16.0
standard cycle	293.0	6.5	1.0	1391.0	524.3	16.0

6.6. Sweeping-Window / Multizone Cycle

Intended for time-accurate simulations using dual-time stepping, the sweeping-window algorithm is identical to the marching-window algorithm, with the exception of no outflow boundary condition forced on the downstream boundary of the sweeping-window. When the converged solution of the previous time level is used as initial conditions for the current time level, not forcing an outflow condition at the downstream boundary helps in attaining faster convergence due to the initial conditions providing a better "guess" at the downstream boundary. For the same reasons, the sweeping-window cycle can also be used to gain extra orders of magnitude of convergence on the solution obtained by the marching-window.

7. Numerical Experiments

Three steady state supersonic flowfields and one unsteady flowfield are solved using the different types of cycles mentioned in the last section, and the performance of each is assessed on the basis of 1) the number of effective iterations, 2) CPU time, and 3) maximum storage required. To enable a fair comparison between the different cycle strategies, the number of effective iterations is defined as

$$\text{effective iterations} \equiv \frac{\text{number of times an inner node is updated}}{\text{total number of inner nodes}}, \quad (41)$$

which is a good measure of the cycle performance as long as most of the computing effort is spent on the pseudo-time stepping instead of the residual, due to the overlap of the residual determination when a multizone decomposition is used. The implicit scheme used herein spends three quarters of its computing effort on the time stepping side, therefore reducing the residual overlap overhead work and justifying the use of Eq. (41) as a performance parameter. In spite of being accurately measured, the number of CPU seconds is not regarded as a more meaningful performance parameter due to the unavoidable bias that might occur in the programming of the cycles and the high dependence of the work on the architecture of the computer. Certain enhancements to the multizone cycle, such as unifying adjacent zones, could be implemented which would result in a non-negligible decrease in work, while the use of a vector computer (of CRAY type) would advantage the longer loops present in the standard cycle. Therefore, both the number of effective iterations and CPU time are monitored for all test cases.

7.1. Inviscid Supersonic Inlet with a Blunt Leading Edge

A first comparison between the different cycles is performed for a steady-state inviscid flow over a 1 m long supersonic inlet. Air enters the channel at a Mach number of 5, a pressure of 4 kPa, and a temperature of 240 K. The grid size is varied between 128 × 64 nodes and 512 × 256 nodes. The user-defined parameters of interest are set to (when applicable)

$$\sigma = 0.5, \quad \xi_{\text{verge}} = 100 \frac{1}{s}, \quad \varphi_{\text{verge}} = 5000 \frac{1}{s}, \quad \phi_0 = 4, \quad \phi_1 = 20, \quad \phi_2 = 3, \quad \text{and} \quad \phi_3 = 9,$$

where the value of 0.5 given to σ translates into a geometric average between the minimum CFL condition based pseudo-time step and the maximum CFL condition based pseudo-time step. The convergence threshold ξ_{verge} is low enough that a decrease in ξ_{verge} would not result in any noticeable difference of the pressure contours in Fig. 6. It is noted that the use of the entropy correction by Yee *et al.* [19] with $\tilde{\zeta} = 0.2$ is here used to avoid a carbuncle phenomenon near the blunt leading edge.

Table 1 shows the CPU time and effective iterations needed to reach convergence for the marching-window, marching-window/multizone, active-domain, multizone, and standard cycles. Due to the CFL = 1 restriction on the traveling speed of the waves in the flowfield to approximately one grid line per iteration, the standard cycle requires a number of iterations




FIGURE 6. Pressure contours for the blunt leading edge inviscid supersonic inlet case obtained using a 512×256 grid; the inflow conditions correspond to $M = 5$, $P = 4$ kPa and $T = 240$ K; no difference is noticeable between the pressure contours obtained with the different cycles.

proportional to the number of grid lines along the streamwise direction, *i.e.* 293 iterations for a 128×64 mesh to 1391 iterations for a 512×256 mesh. The multizone cycle suffers the same symptoms but has the extra advantage of *not* allocating work to the zones where all nodes exhibit a ξ smaller than the user specified threshold, therefore reducing the computing to a smaller and smaller domain as the iteration count progresses and the non-converged flow region moves towards the domain exit. This results in impressive savings in iteration count of 1.8 times for the coarse mesh and of 4.4 times for the fine mesh. Both the active-domain cycle and the marching-window cycle decrease further the iteration count by allowing a computational window to travel in space following the propagation of the waves. This results in a decrease in effective iterations, compared to the standard cycle using the 512×256 mesh, of 14 and 16 times for the active-domain and marching-window respectively. Furthermore, the use of multizone decomposition inside the marching-window focuses the pseudo-time stepping effort to the regions requiring more iterations to reach convergence, such as the region of subsonic flow upstream of the inlet blunt leading edge, hence resulting in only 45 effective iterations to reach convergence and an overall reduction in effective iterations of 31 times compared to the standard cycle.

It is reminded that the standard cycle, the multizone cycle, and the marching-window cycle (with and without multizone decomposition) all guarantee that

$$\xi \leq \xi_{\text{verge}} \quad \forall \text{ inner nodes ,}$$

once convergence is reached [as previously stated in Eq. (33)], which is a necessary condition for a well-posed acceleration technique. The latter is *not* a property of the active-domain method when the discretization stencil for the streamwise convection derivative depends on downstream nodes in locally supersonic flow. The Yee TVD limiter used here has this property, and the active-domain algorithm induces a converged residual that does not satisfy the convergence criterion of Eq. (33).

It could be argued that raising the CFL number would improve the standard cycle over the others for this particular case. Investigation on a change of CFL number is not performed, but is addressed in the subsequent test problems. It is noted that for many realistic problems dominated by non linear phenomena, nonlinear stability conditions restrict the use of high CFL numbers until the waves have started to settle down considerably and for which the use of a fine mesh results in very poor

TABLE 2.
 Effective iteration count, work, and storage comparison for the backward facing step case for a 256×128 grid.

cycle	CFL = 1			$1 \leq \text{CFL} \leq 10$		
	iter.	work	stor.	iter.	work	stor.
marching-window / multizone	459	2.0	1.0	229	1.0	1.0
marching-window	1166	3.9	1.0	355	1.2	1.0
multizone	1259	5.9	4.3	698	3.2	4.3
standard cycle	3164	10.4	4.3	1054	3.5	4.3




FIGURE 7. Mach number contours of the backward facing step case, obtained using a 256×128 node mesh; the inflow conditions are set to a Mach number of 2, Reynolds number per meter of 5 million and temperature of 300 K.

performance of the standard cycle.

7.2. Backward Facing Step

Although the gains in computing efficiency obtained through the marching-window cycle might be expected for an entirely supersonic problem without any reverse flow regions, we proceed to show in this subsection that the marching-window can reduce convergence time considerably even when a substantial portion of the flowfield is separated.

Air enters the computational domain at a Mach number of 2, Reynolds number per meter of 5 million and temperature of 300 K. Symmetry conditions are applied at $x_2 = -0.015$ m and at $x_2 = 0.03$ m, inflow at $x_1 = -0.067$ m, and outflow at $x_1 = 0.2$ m, and an adiabatic wall boundary condition is in effect elsewhere. The Mach number contours shown in Fig. 7 show the limits of the recirculation region ($0 \leq x_1 \leq 0.04$) in which 25% of the grid lines along X_1 and 50% of the grid lines along X_2 are placed. The investigation is performed using a 256×128 mesh, clustered at the surfaces; it is noted that no significant difference in the trends is observed using a coarser mesh of 128×64 nodes.

A minimum/maximum CFL averaged local time step as specified in Eq. (30) is used for all cycles, and the convergence threshold along with the other user-defined constants are specified to

$$\sigma = 0.5, \quad \xi_{\text{verge}} = 100 \frac{1}{s}, \quad \varphi_{\text{verge}} = 5000 \frac{1}{s}, \quad \phi_1 = 20, \quad \phi_2 = 3, \quad \text{and} \quad \phi_3 = 9.$$

TABLE 3.

Effective iteration count, work, and storage comparison for the concatenated channels test case; the mesh size is of 256×128 nodes.

cycle	$0.1 \leq \text{CFL} \leq 1$			$0.1 \leq \text{CFL} \leq 10$		
	iter.	work	stor.	iter.	work	stor.
marching-window / multizone	431	1.9	1.0	219	1.0	1.0
marching-window	1111	4.0	1.0	444	1.7	1.0
multizone	1571	7.8	6.2	1215	5.9	6.2
standard cycle	4547	15.5	6.2	2342	8.0	6.2

TABLE 4.

Sensitivity of the effective iteration count and work to the user-defined constants for the concatenated channels test case; the marching-window/multizone cycle is used with a mesh size of 256×128 nodes, and a CFL range $0.1 \leq \text{CFL} \leq 10$.

ϕ_1	ϕ_2	ϕ_3	φ_{verge}	work	iter.
20	3	9	5×10^3	1.00	1.0
10	3	9	5×10^3	1.20	1.12
40	3	9	5×10^3	1.00	1.02
20	15	9	5×10^3	1.15	1.12
20	1	9	5×10^3	0.96	0.98
20	3	5	5×10^3	1.46	1.41
20	3	18	5×10^3	1.19	1.24
20	3	36	5×10^3	1.89	1.99
20	3	9	5×10^2	1.33	1.40
20	3	9	5×10^4	1.57	1.54

When a variable CFL number is used, it is set to a function of ξ_{max} as opposed to the iteration count to enable a more adequate comparison between the different cycles, since the convergence history has a different dependence on the iteration count for each cycle. Note that ξ_{max} stands for the maximum value of ξ in the computational window, which corresponds to the entire domain for the multizone and standard cycles. In this case, a ξ_{max} varying between 10^5 and 10^3 is made to induce a CFL number varying between 1 and 10. The variation in CFL number is necessary due to nonlinear stability restrictions on the time step size, and it is assumed that an inverse relationship exists between ξ_{max} and the maximum allowable CFL number for stable and predictable convergence.

The influence of a change in CFL number on the efficiency of the different cycles can be seen in Table 2. The marching-window/multizone cycle at a CFL of unity requires 6.9 times fewer iterations to reach convergence than the standard cycle. Yet, if the CFL is varied between 1 and 10 the speed-up is reduced to 4.6 times. As expected, a rise in the CFL number greatly helps the propagation of the waves along the streamwise direction for the standard cycle, while the transmission of the streamwise information is already adequate at a CFL of unity for the marching-window and marching-window/multizone cycles. Nevertheless, the recirculation region is the iteration bottleneck of this problem and the number of steps necessary to solve it is similar for all approaches. For the standard cycle, since the entire computational domain is computed at every step, a region of slow convergence somewhere in the flowfield translates in a very high number of effective iterations, whether the region of slow convergence is very small or not. On the other hand, the marching-window algorithm focuses the effort on the region of slow convergence, consequently resulting in much improved algorithm efficiency.

7.3. Concatenated Channels: Shock / Boundary Layer Interactions

The ability of the marching-window algorithm to solve shock / boundary layer interactions at hypersonic flow conditions is now tested. The geometry involves the concatenation of a 1.0×0.5 m channel to a 0.69×0.38 m channel through a 37° compression ramp. Air enters the first channel at uniform conditions of $M = 5$, $P = 1000$ Pa, and $T = 450$ K. Fixed temperature ($T_{\text{wall}} = 450$ K) wall boundary conditions are applied on bottom and top boundaries, with a grid clustered at both walls. As for the backward facing step case, a geometric averaged local time step is utilized to enhance wave propagation through high




FIGURE 8. Effective pressure contours of the concatenated channels case obtained using a 512×256 mesh; air enters the first channel at a Mach number of 5, a pressure of 1000 Pa and a temperature of 450 K.

aspect ratio cells, while the other user-adjustable parameters are set to

$$\sigma = 0.5, \quad \xi_{\text{verge}} = 100 \frac{1}{s}, \quad \varphi_{\text{verge}} = 5000 \frac{1}{s}, \quad \phi_1 = 20, \quad \phi_2 = 3, \quad \text{and} \quad \phi_3 = 9.$$

From the effective pressure contours of Fig. 8, two recirculation regions are visible: one at the start of the shock formed by the 37° wedge, and one at the point where the shock impinges on the top wall boundary layer. Both recirculation zones are of appreciable size due to the very low Reynolds number of the flow which helps generate thick incoming boundary layers. The major obstacle in converging this flowfield efficiently comes from the high difference in time scales between the convection dominated flow in the middle of the channels and the viscous dominated recirculation zones. Time accurate simulations of a similar problem indicate that the amount of time required for the separated flow regions to reach steady state is typically one order of magnitude more than the time needed for the shock structure to establish itself. Consequently, one would prefer high pseudo-time steps to be used in the recirculation zones for fast convergence, but unfortunately the step size is limited by nonlinear stability restrictions which are of importance especially near the non-converged shock waves. For these reasons, it is not surprising that so many iterations are needed for the standard cycle to reach convergence, as Table 3 shows: 4547 iterations for a CFL number varying between 0.1 and 1 and 2342 iterations for the range $0.1 \leq \text{CFL} \leq 10$. Similarly to the backward facing step, the CFL number is linked to ξ_{max} such that at $\xi_{\text{max}} = 10^4$, the CFL number is 10 and at $\xi_{\text{max}} = 10^6$, the CFL number is 0.1.

The marching-window/multizone cycle performs particularly well as the work is focused on the reverse flow regions, while the rest of the domain is quasi-hyperbolic/parabolic and needs only a small amount of work to reach convergence (see Fig. 9). The use of the marching-window coupled with a multizone strategy makes possible a decrease in effective iterations of 8 times compared to the standard cycle, independently of the CFL number used, as shown in Table 3.

There might be doubts as to the adequacy of a varying CFL number function of ξ_{max} as a means to compare different cycles. For this reason, additional simulations involving the marching-window/multizone and standard cycles are performed in which the CFL number is made a function of the iteration count and raised to 10 as rapidly as the stability conditions permit. 1930 iterations for the standard cycle are needed for convergence while 197 iterations are needed for the marching-window/multizone cycle. Again, while a slight increase in efficiency for the standard cycle is apparent, approximately the same amount is noticeable for the marching-window/multizone cycle.




FIGURE 9. Location of the marching-window upstream and downstream boundaries for the concatenated channels case using the marching-window/multizone cycle, with a variable CFL number, *i.e.* $0.1 \leq \text{CFL} \leq 10$; notice the high amount of work spent on the recirculation zones in the vicinity of $X_1 \sim 128$ and $X_1 \sim 218$, while very few steps are needed to converge the quasi-hyperbolic/parabolic regions.

As in the previous test cases, a convergence criterion of $\xi_{\text{verge}} = 100 \frac{1}{s}$ is found necessary to obtain reasonable accuracy, and no discernible difference is observed between the contours of properties obtained with the different cycles. Even if both the marching-window and the standard cycle guarantee the convergence criterion of Eq. (33) to be satisfied once convergence is attained, the governing equations have multiple roots due to their non-linearity, and a different flow solution could be obtained by the different cycle strategies. For all test cases presented here, however, it is verified that the same root is obtained independently of the acceleration technique.

The sensitivity of the user-adjustable parameters for the marching-window cycle is assessed for this test case in Table 4. It is seen that the performance of the marching-window is not affected considerably by a change of the average zone length ϕ_1 or by a change in ϕ_2 , the latter being the number of iterations before a reading of the streamwise ellipticity sensor ξ is taken. For ϕ_1 varied from 10 to 40, the number of effective iterations is observed to change by only 12%, and for ϕ_2 varied from 1 to 15, the number of effective iterations increases by 14%. On the other hand, the parameters ϕ_3 and ϕ_{verge} are seen to affect the performance of the algorithm considerably. Raising ϕ_3 from 9 to 36 increases twofold the number of effective iterations, and increasing ϕ_{verge} tenfold results in an increase of 54% in the effective iterations count. The high sensitivity of the effective iterations on either ϕ_3 or ϕ_{verge} is due to the high dependence of the width of the marching-window on these parameters. When the marching-window encloses too tightly a zone of streamwise ellipticity, the solution needs to be converged locally several times, hence increasing the work. When the marching-window overestimates the size of a streamwise elliptic region, the high number of iterations needed locally to converge a streamwise-elliptic region is spent on a larger portion of the computational domain, hence resulting in decreased performance.

7.4. Time Accurate Simulation of an Exploding Cavity in a Supersonic Stream

The performance of the different cycles on a time accurate simulation of a stagnant high pressure flow pocket exploding into a Mach 2 air stream is investigated in this subsection. The computational domain has dimensions as shown in Fig. 10, and is spanned by a grid composed of 256×128 nodes of which 110×38 are allocated to the cavity. Inflow, outflow and symmetry conditions are applied to the left, right and top boundaries respectively, while an adiabatic wall condition is in effect elsewhere. The mesh is clustered at the inflow and at all surfaces to capture the turbulent boundary layer correctly and is not altered in

TABLE 5.

Effective iteration count, work and storage comparison for the time accurate simulation of an exploding cavity; the CFL number is unity.

cycle	iterations	work	storage
multizone	1505	1.0	5.0
sweeping-window / multizone	2034	1.5	1.0
sweeping-window	2980	1.9	1.0
standard cycle	4418	2.0	5.0

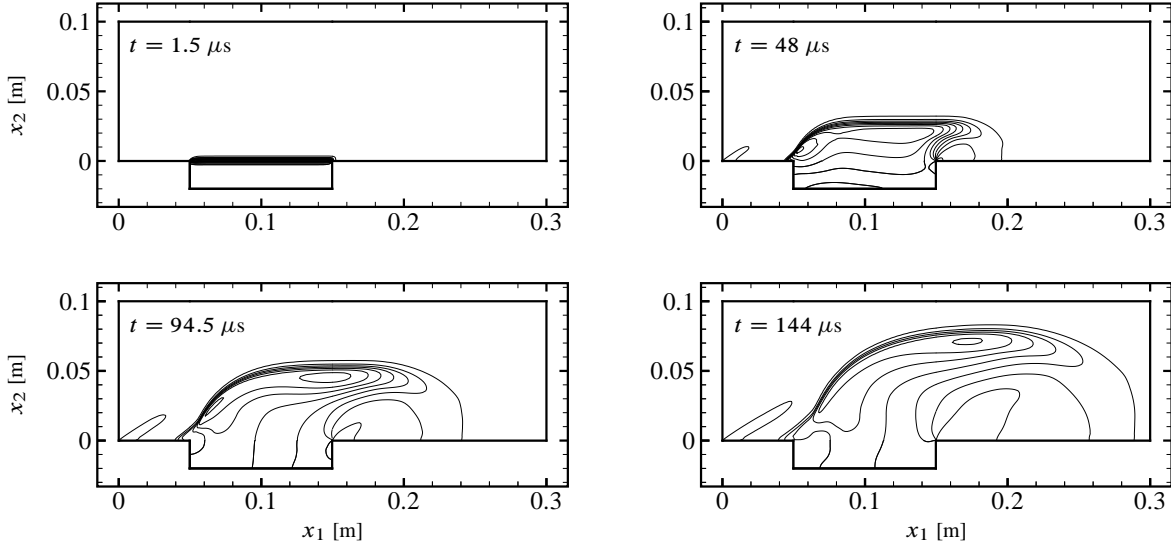


FIGURE 10. Effective pressure contours of the exploding cavity case using a mesh composed of 256×128 nodes, and a time step of 1.5 microsecond; at $t = 0$, the flow outside the cavity is air uniformly distributed at $P = 10$ kPa, $T = 300$ K, and $M = 2$, while the air inside the cavity is set to $P = 100$ kPa, $T = 2000$ K, and $M = 0$.

time. The flowfield at $t = 0$ for $x_2 \geq 0$ is set to a pressure of 10 kPa, a temperature of 300 K and a flow Mach number of 2, while for $x_2 < 0$, the pressure, temperature and Mach number are set to 100 kPa, 2000 K and 0 respectively. The solution is iterated in pseudo-time starting from the converged solution of the previous physical time step until the maximum value ξ in the computational domain falls below ξ_{verge} . The physical time step, Δt , is fixed to 1.5 microsecond, while the following user defined parameters are in use:

$$\sigma = 0.3, \quad \xi_{\text{verge}} = 100 \frac{1}{s}, \quad \varphi_{\text{verge}} = 5000 \frac{1}{s}, \quad \phi_1 = 20, \quad \phi_2 = 3, \quad \text{and} \quad \phi_3 = 9.$$

Due to the strict convergence criterion utilized and the use of the same residual, all acceleration techniques result in the same answer at all time steps despite noticeable differences in CPU work, as Table 5 attests: a twofold decrease in work is achieved through the use of the multizone cycle, while the sweeping-window/multizone cycle decreases the work by one quarter. The performance of the sweeping-window is not particularly good for this problem as the wave propagation direction is more towards time than in the streamwise coordinate due to the relatively small physical time step. Furthermore, since only 20 effective iterations are needed on average per time level, the overhead work induced by sweeping becomes more important, as the greater discrepancy observed for this case between the reduction in effective iterations and CPU work shows. There is, nonetheless, a non-negligible fivefold reduction in storage when using the sweeping-window.

8. Summary and Conclusions

A novel acceleration technique is presented which is aimed at accelerating the convergence of the Favre-averaged Navier-Stokes equations in the supersonic/hypersonic regime for flowfields with large streamwise separated flow regions. Similarly to the active-domain method [13], the marching-window iterates in pseudo-time a band-like computational domain of minimal

width which adjusts to the size of the streamwise elliptic regions when encountered. However, contrarily to the active-domain method, it is shown that the marching-window guarantees the residual on all nodes to be below a user-defined threshold when convergence is reached, and hence results in the same converged solution (within the tolerance of the convergence criterion) as the one obtained by standard pseudo-time marching methods. Further, a streamwise ellipticity sensor based on the Vigneron splitting [1] is developed which ensures the downstream boundary of the marching-window to advance sufficiently such that regions of significant streamwise ellipticity are contained within the marching-window subdomain. It is noted that while the Vigneron splitting sensor does not capture all possible streamwise elliptic phenomena, this does not affect the accuracy of the final solution and only affects the performance of the marching-window as an acceleration technique. Also, a multizone decomposition is implemented inside the marching-window to restrict the computing to the zones where the residual is above the user-defined convergence threshold. This is shown to further decrease the work needed for convergence by close to 2 times for the problems shown herein.

The use of the marching-window with multizone decomposition on a backward facing step and a shock boundary layer interaction flowfield (where one or several large streamwise separated region is present) reveals a 4 to 6 times decrease in storage and a 4 to 8 times decrease in work compared to the standard cycle. The proposed algorithm is also shown to work well at a low CFL number in regions of quasi-hyperbolicity/parabolicity and is recommended for stiff problems with high non-linear stability restrictions on the time step size. A variant of the marching-window designed for time-accurate simulations is observed to result in a fivefold reduction in storage and 25% reduction in work for the time-accurate exploding cavity case investigated herein. The reduction in computational work through the use of the marching-window is made possible by focusing the high number of iterations needed to converge the streamwise separated regions to the region in question. The amount of storage needed is also significantly reduced if no memory is allocated to the nodes outside of the marching-window subdomain.

The marching-window does not impose any restriction on the discretization stencils part of the residual or on the pseudo-time stepping method. While not implemented here, the numerous acceleration techniques available for pseudo-time stepping (such as multigrid, block relaxation [31], preconditioning, Newton-Krylov, *etc.*) can be used in conjunction with the marching-window. Furthermore, the marching-window is not limited to the Favre-averaged Navier-Stokes equations and its extension to other governing equations would only require a redefinition of the ellipticity sensor shown in Eq. (39).

The performance of the algorithm is seen to be sensitive to the user-defined ellipticity threshold constant φ_{verge} and the marching-window minimal width ϕ_3 . It is unclear at this stage by how much these parameters would need to be altered for very different flow properties and physical domain size. The dependency on the problem setup seems not too severe as the same values for the user-specified constants are used for all cases shown in this paper.

References

- [1] VIGNERON, Y. C., RAKICH, J. V., AND TANNEHILL, J. C., "Calculations of Supersonic Viscous Flow Over Delta Wings with Sharp Subsonic Leading Edges," 1978, AIAA Paper 78-1137.
- [2] LAWRENCE, S. L. AND TANNEHILL, J. C., "An Upwind Algorithm for the Parabolized Navier-Stokes Equations," *AIAA/ASME 4th Fluid Mechanics, Plasma Dynamics and Lasers Conference*, 1986, AIAA Paper 86-1117.
- [3] TANNEHILL, J., IEVALTS, J. O., AND LAWRENCE, S., "An Upwind Parabolized Navier-Stokes Code for Real Gas Flows," *26th Aerospace Sciences Meeting*, Jan. 1988, AIAA Paper 89-0713.
- [4] D'AMBROSIO, D. AND MARSILIO, R., "A Numerical Method for Solving the Parabolized Three-Dimensional Navier-Stokes Equations," *Computers & Fluids*, Vol. 26, No. 6, 1997, pp. 587–611.
- [5] MILLER, J. H., TANNEHILL, J. C., LAWRENCE, S. L., AND EDWARDS, T. A., "Parabolized Navier-Stokes Code for Hypersonic Flows in Thermo-Chemical Equilibrium or Nonequilibrium," *Computers & Fluids*, Vol. 27, No. 2, 1998, pp. 199–215.
- [6] MILLER, J. H., TANNEHILL, J. C., AND LAWRENCE, S. L., "Parabolized Navier-Stokes Algorithm for Solving Supersonic Flows with Upstream Influences," *AIAA Journal*, Vol. 38, No. 10, 2000, pp. 1837–1845.
- [7] POWER, G. D. AND BARBER, T. J., "Analysis of Complex Hypersonic Flows with Strong Viscous/Inviscid Interaction," *AIAA Journal*, Vol. 26, No. 7, 1988, pp. 832–840.
- [8] CHANG, C.-L. AND MERKLE, C. L., "The Relation between Flux Vector Splitting and Parabolized Schemes," *Journal of Computational Physics*, Vol. 80, 1989, pp. 344–361.
- [9] YAMALEEV, N. K. AND BALLMANN, J., "Iterative Space-Marching Method for Compressible Sub-, Trans-, and Supersonic Flows," *AIAA Journal*, Vol. 38, No. 2, 2000, pp. 225–233.
- [10] RUBIN, S. G. AND REDDY, D. R., "Analysis of Global Pressure Relaxation for Flows with Strong Interaction and Separation," *Computers & Fluids*, Vol. 11, No. 4, 1983, pp. 281–306.
- [11] BARNETT, M. AND DAVIS, R. T., "Calculation of Supersonic Flows with Strong Viscous-Inviscid Interaction," *AIAA Journal*, Vol. 24, No. 12, 1986, pp. 1949–1955.
- [12] BARDINA, J. E., "Three-Dimensional Navier-Stokes Method with Two-Equation Turbulence Models for Efficient Numerical Simulation of Hypersonic Flows," *Proceedings of the 30th Joint Propulsion Conference and Exhibit, Indianapolis, IN*, 1994, AIAA Paper 94-2950.

- [13] NAKAHASHI, K. AND SAITOH, E., "Space-Marching Method on Unstructured Grid for Supersonic Flows with Embedded Subsonic Regions," *AIAA Journal*, Vol. 35, No. 8, 1997, pp. 1280–1285.
- [14] MORINO, H. AND NAKAHASHI, K., "Space-Marching Method on Unstructured Hybrid Grid for Supersonic Viscous Flows," 1999, AIAA Paper 99-0661.
- [15] WILCOX, D. C., "Reassessment of the Scale Determining Equation for Advanced Turbulence Models," *AIAA Journal*, Vol. 26, No. 11, 1988, pp. 1299–1310.
- [16] VIVIAND, H., "Conservative Forms of Gas Dynamics Equations," *La Recherche Aéronautique*.
- [17] VINKUR, M., "Conservative Equations of Gas-Dynamics in Curvilinear Coordinate Systems," *Journal of Computational Physics*, Vol. 14, 1974, pp. 105–125.
- [18] ROE, P. L., "Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes," *Journal of Computational Physics*, Vol. 43, 1981, pp. 357–372.
- [19] YEE, H. C., KLOPFER, G. H., AND MONTAGNÉ, J.-L., "High-Resolution Shock-Capturing Schemes for Inviscid and Viscous Hypersonic Flows," *Journal of Computational Physics*, Vol. 88, 1990, pp. 31–61.
- [20] BATTEN, P., LESCHZNER, M. A., AND GOLDBERG, U. C., "Average-State Jacobians and Implicit Methods for Compressible Viscous and Turbulent Flows," *Journal of Computational Physics*, Vol. 137, 1997, pp. 38–78.
- [21] PARENT, B. AND SISLIAN, J. P., "Turbulent Hypervelocity Fuel/Air Mixing by Cantilevered Ramp Injectors," *Proceedings of the 10th AIAA International Aerospace Planes Hypersonics Technologies Conference, Kyoto, Japan*, April 2001, AIAA Paper 2001-1888.
- [22] BRILEY, W. R. AND McDONALD, H., "Solution of the Multidimensional Compressible Navier-Stokes Equations by a Generalized Implicit Method," *Journal of Computational Physics*, Vol. 24, 1977, pp. 372–397.
- [23] BEAM, R. AND WARMING, R. F., "An Implicit Factored Scheme for the Compressible Navier-Stokes Equations," *AIAA Journal*, Vol. 16, No. 4, 1978, pp. 393–402.
- [24] PATANKAR, S. V., *Numerical Heat Transfer and Fluid Flow*, Taylor and Francis, 1980.
- [25] MACCORMACK, R. W., "Iterative Modified Approximate Factorization," *Computers & Fluids*, Vol. 30, No. 8, 2001, pp. 917–925.
- [26] YOON, S. AND JAMESON, A., "Lower-Upper Implicit Scheme for High-Speed Inlet Analysis," *AIAA Journal*, Vol. 25, 1987, pp. 1052–1053.
- [27] WILCOX, D. C., *Turbulence Modeling for CFD*, DCW Industries, 1994.
- [28] LANEY, C. B., *Computational Gasdynamics*, Cambridge University Press, Cambridge, New-York, NY, 1998.
- [29] SAWLEY, M. L. AND TEGNER, J. K., "A Data Parallel Approach to Multi-Block Flow Computations," *International Journal for Numerical Methods in Fluids*, Vol. 19, No. 8, 1994, pp. 707–721.
- [30] ROSENFELD, M., "The Alternating Direction Multi-zone Implicit Method," *Journal of Computational Physics*, Vol. 110, 1994, pp. 212–220.
- [31] DE NICOLA, C., TOGNACCINI, R., AND PUOTI, V., "Local Block Relaxation Method for the Solution of Equations of Gasdynamics," *AIAA Journal*, Vol. 38, No. 8, 2000, pp. 1377–1384.