Computers and Fluids 159 (2017) 327–337

Contents lists available at ScienceDirect

Computers and Fluids

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

A conservative overlap method for multi-block parallelization of compact finite-volume schemes

F. Capuano[a,*], A. Mastellone[b], E.M. De Angelis[a]

[a] Dipartimento di Ingegneria Industriale (DII), Università di Napoli "Federico II", Napoli, 80125, Italy
[b] Centro Italiano Ricerche Aerospaziali (CIRA), Capua, 81043, Italy

ARTICLE INFO

Article history:
Received 8 April 2017
Revised 18 August 2017
Accepted 11 October 2017

Keywords:
Compact schemes
Finite-volume method
Local conservation
MPI Parallel computing
Turbulent flows

ABSTRACT

A conservative approach for MPI-based parallelization of tridiagonal compact schemes is developed in the context of multi-block finite-volume methods. For each block, an enlarged linear system is solved by *overlapping* a certain number of neighbour cells from adjacent sub-domains. The values at block-to-block boundary faces are evaluated by a high-order centered approximation formula. Unlike previous methods, conservation is retained by properly re-computing the common interface value between two neighbouring blocks. Numerical tests show that parallelization artifacts decrease significantly as the number of overlapping cells is increased, at some expense of parallel efficiency. A reasonable trade-off between accuracy and performances is discussed in the paper with reference to both the spectral properties of the method and the results of fully turbulent numerical simulations.

© 2017 Elsevier Ltd. All rights reserved.

1. Introduction

Compact schemes are widely employed in numerical simulations of turbulent flows, because of their beneficial resolution properties and high overall accuracy [1]. In contrast to classical *explicit* schemes, compact methods are global and require the solution of a tri- or penta-diagonal linear system to provide the desired derivatives or interpolated values [2]. The implicit nature of such methods leads to a difficult implementation in a parallel framework, especially when a MPI-based domain decomposition technique is employed. This issue may preclude the potential benefits of compact methods for direct (DNS) or large-eddy simulations (LES) of turbulence, for which massively parallel computations are mandatory. Several algorithms have been developed over the years to tackle this problem efficiently, each with pros and cons; the approaches fall mainly into three categories, which are briefly reviewed as follows.

The first category is typically referred to as *transpose* methods, and is popular in the pseudo-spectral community. In this case, the computational domain is partitioned along one (or two) dimensions at a time (*slabs* or *pencils* respectively), so that the application of the global scheme can be performed exactly along the remaining direction(s). Then, the computational space is transposed by means of *all-to-all* communications to allow completion of the algorithm [3–5]. Although in this case the parallelization is free of errors (i.e., the solution is identical to the serial one), the resulting approach is very communication intensive, as the collective communications require to exchange a volume of data which increases in size as N^3, being N the number of unknowns of each one-dimensional computation - in contrast with the N^2 scaling of conventional techniques. On the other hand, upon use of optimized libraries (such as MPI_Alltoall and FFTW3) and careful processor mapping, some authors were able to obtain excellent scalability on up to $\mathcal{O}(10^5)$ processors [3,6]. By construction, the transpose method is especially suited for simple cartesian domains, while an appropriate generalization to complex geometries may require involved computer programming to properly take into account the domain topology. Also, the number of processors is limited to no more than N (or N^2) when slabs (or pencils) are used.

Alternatives to the transpose method are represented by algorithmic approaches, i.e., methods that aim to parallelize the solution of the banded linear system. Notable examples include the pipelined Thomas [7], the parallel diagonal dominant [8] (which are limited to tridiagonal systems), and the SPIKE algorithms [9]. These methods are powerful and provide exact (pipelined Thomas, SPIKE) or bounded (parallel diagonal dominant) parallelizations errors, but are usually susceptible to penalties in efficiency due to idle times. Also, the computational complexity of the computer code is generally highly increased. To the authors' knowledge, the use of such methods in computational fluid dynamics (CFD) applications is relatively limited. For a comparison of the pipelined Thomas and parallel diagonal dominant algorithms for flow simu-

* Corresponding author.
E-mail address: francesco.capuano@unina.it (F. Capuano).

https://doi.org/10.1016/j.compfluid.2017.10.017
0045-7930/© 2017 Elsevier Ltd. All rights reserved.

lations, see [10]. The SPIKE algorithm has been recently applied to large-eddy simulation in [11].

The third family of methods is constituted by the so-called boundary approximation approach (BAA). In this case, the original linear system is split into disjoint matrix systems that can be solved independently for each sub-domain by exchanging a (small) number of halo cells. The method is thus naturally suited to the widespread MPI-based multi-block partitioning technique, which is at the base of many finite-difference and finite-volume CFD codes. The major drawback of this procedure is that the global dependence of the compact scheme is broken at block-to-block interfaces, leading to a certain degree of deterioration similar to the one occurring at regular boundaries. These effects are mainly attributed to the altered dissipation and dispersion properties due to the boundary closure [12]. However, it is well known that the spectral characteristics of the discretization are crucial for an accurate numerical simulation of multi-scale and acoustics phenomena, such as compressible turbulence [13]. Several boundary approximations have been developed in recent years to reduce parallelization artifacts. These mainly rely on overlapping grids [14,15] or halo points [16] with proper boundary closures at adjacent sub-domain interfaces. In some cases, the resulting schemes are optimized in wavenumber space for accurate acoustic computations, or used in conjunction with suitable filtering operators to remove high-frequency errors [17]. Particularly Gaitonde & Visbal [15] suggested to use a region of overlap between adjacent subdomains aimed to reduce parallelization artifacts in their finite-difference method; boundary closure was achieved by high-order one-sided formulas. However, their approach is not locally conservative, which might be troublesome for flows with shocks [18].

The present work falls within the BAA category and focuses on tridiagonal compact schemes. An overlapping strategy is developed in order to preserve accuracy on the interior points of each sub-domain; however, boundary closure is achieved by explicit centered formulas. Unlike previous studies, which have been concerned with the finite-difference method, a finite-volume (FV) discretization is employed here. A straightforward implementation of the overlapping method in a FV framework leads to the formal loss of local conservation at block-to-block interfaces. In this work, a method aimed to overcome this issue is developed so to retain the inherent conservation properties of the finite-volume method. The proposed algorithm is thus suitable for high-fidelity, parallel computations of compressible shock-free turbulent flows, and could serve as a building block for flows with discontinuities.

The paper is organized as follows. In Section 2, the relevant governing equations are introduced; then, the employed finite-volume discretization and serial numerical method are described. Section 3 presents the parallel method. A modified wavenumber analysis is reported in Section 4, while Section 5 reports a series of numerical tests aimed to characterize the accuracy of the proposed approach. Section 6 focuses on efficiency and parallel performances. Concluding remarks are given in Section 7.

2. Governing equations and serial numerical method

The fully compressible Navier-Stokes equations are considered in this work. In a three-dimensional Cartesian coordinate frame (x, y, z), the motion of a gas with density ρ, velocity $\mathbf{u} = (u, v, w)$, pressure p, temperature T and total energy E is governed by the system:

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}_x}{\partial x} + \frac{\partial \mathbf{F}_y}{\partial y} + \frac{\partial \mathbf{F}_z}{\partial z} = 0, \tag{1}$$

where $\mathbf{U}$ is the vector of conservative variables, $\mathbf{U} = (\rho, \rho u, \rho v, \rho w, E)$. The total energy E is defined as the sum of internal and kinetic energy, $E = \rho e + \rho |\mathbf{u}|^2/2$. The vectors $\mathbf{F}_x = \mathbf{F}_x^c - \mathbf{F}_x^d$, $\mathbf{F}_y = \mathbf{F}_y^c - \mathbf{F}_y^d$ and $\mathbf{F}_z = \mathbf{F}_z^c - \mathbf{F}_z^d$ represent the fluxes along the three components. The inviscid (convective) fluxes are defined as:

$$\mathbf{F}_x^c = \begin{pmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho uw \\ u(E+p) \end{pmatrix}, \quad \mathbf{F}_y^c = \begin{pmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ \rho vw \\ v(E+p) \end{pmatrix},$$

$$\mathbf{F}_z^c = \begin{pmatrix} \rho w \\ \rho uw \\ \rho vw \\ \rho w^2 + p \\ w(E+p) \end{pmatrix}, \tag{2}$$

whereas the diffusive fluxes are

$$\mathbf{F}_x^d = \begin{pmatrix} 0 \\ \tau_{11} \\ \tau_{12} \\ \tau_{13} \\ (\tau \mathbf{u})_1 - q_1 \end{pmatrix}, \quad \mathbf{F}_y^d = \begin{pmatrix} 0 \\ \tau_{21} \\ \tau_{22} \\ \tau_{23} \\ (\tau \mathbf{u})_2 - q_2 \end{pmatrix},$$

$$\mathbf{F}_z^d = \begin{pmatrix} 0 \\ \tau_{31} \\ \tau_{32} \\ \tau_{33} \\ (\tau \mathbf{u})_3 - q_3 \end{pmatrix}. \tag{3}$$

The stress tensor τ_{ij} and the conductive heat flux q_i are expressed by the usual Newton's and Fourier's laws, respectively $\tau_{ij} = \mu\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \frac{2}{3}\delta_{ij}\frac{\partial u_k}{\partial x_k}\right)$ and $q_i = -\lambda \frac{\partial T}{\partial x_i}$, where μ is the molecular viscosity and λ the thermal conductivity of the fluid. Closure is achieved by means of the ideal-gas equation of state, $p = \rho RT$.

2.1. Finite-volume discretization

The computational domain is partitioned into a structured grid of hexahedrons indexed by (i, j, k). The finite-volume method is based upon integration of Eq. (1) over a generic control volume Ω, yielding

$$V_{ijk}\frac{\partial \tilde{\mathbf{U}}_{ijk}}{\partial t} + \int_{\partial\Omega} (\mathbf{F}_x n_x + \mathbf{F}_y n_y + \mathbf{F}_z n_z) d\sigma = 0, \tag{4}$$

where $V_{ijk} = |\Omega|$ is the volume of the region and

$$\tilde{\mathbf{U}}_{ijk} = \frac{1}{V_{ijk}} \int_\Omega \mathbf{U} d\Omega. \tag{5}$$

The integral in Eq. (4) applies to each control-volume; as a consequence, surface integrals over inner cell faces cancel out and discrete global conservation of primary unknowns is guaranteed through the telescopic property. The built-in global conservation feature is indeed one of the major advantages of finite-volume methods.

The meaning of the cell integral in Eq. (5) deserves further discussion. In a so-called *pointwise* approach, nodal values $\mathbf{U}_P$ are usually supposed to be known at the cell-center, resulting in a second-order accurate average, i.e., $\tilde{\mathbf{U}} = \mathbf{U}_P + \mathcal{O}(\Delta^2)$, with Δ being a relevant grid spacing. Therefore, if high-order accuracy is seeked, correspondent high-order formulas are needed for the evaluation of

Download English Version:

https://daneshyari.com/en/article/7156559

Download Persian Version:

https://daneshyari.com/article/7156559

Daneshyari.com