

Adaptive-Anisotropic Wavelet Collocation Method on general curvilinear coordinate systems



Eric Brown-Dymkoski^a, Oleg V. Vasilyev^{b,c,a,*}

^a Department of Mechanical Engineering, University of Colorado Boulder, Boulder, CO 80309, USA

^b Skolkovo Institute of Science and Technology, Skolkovo Innovation Center, 3 Nobel St., Moscow 143026, Russia

^c NorthWest Research Associates, 3380 Mitchell Ln., Boulder, CO 80301, USA

ARTICLE INFO

Article history:

Received 7 May 2016

Received in revised form 14 October 2016

Accepted 20 December 2016

Available online 3 January 2017

Keywords:

CFD

Adaptive meshing

Wavelet collocation

Curvilinear mesh

ABSTRACT

A new general framework for an Adaptive-Anisotropic Wavelet Collocation Method (A-AWCM) for the solution of partial differential equations is developed. This proposed framework addresses two major shortcomings of existing wavelet-based adaptive numerical methodologies, namely the reliance on a rectangular domain and the “curse of anisotropy”, i.e. drastic over-resolution of sheet- and filament-like features arising from the inability of the wavelet refinement mechanism to distinguish highly correlated directional information in the solution. The A-AWCM addresses both of these challenges by incorporating coordinate transforms into the Adaptive Wavelet Collocation Method for the solution of PDEs. The resulting integrated framework leverages the advantages of both the curvilinear anisotropic meshes and wavelet-based adaptive refinement in a complimentary fashion, resulting in greatly reduced cost of resolution for anisotropic features. The proposed Adaptive-Anisotropic Wavelet Collocation Method retains the *a priori* error control of the solution and fully automated mesh refinement, while offering new abilities through the flexible mesh geometry, including body-fitting. The new A-AWCM is demonstrated for a variety of cases, including parabolic diffusion, acoustic scattering, and unsteady external flow.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

The adaptive wavelet collocation method (AWCM) [1] is a powerful tool for computational solutions of parabolic [2], elliptic [3], and hyperbolic [4] partial differential equations, as it allows highly efficient use of memory resources and rigorous control of discretization error. It dynamically adapts the grid using a wavelet basis decomposition to track localized structures in the solution, exploiting spatio-temporal intermittency to optimize the total grid point count. As a result, AWCM has been used to effectively analyze many physical problems, including turbulence [5–7], reacting flows [8,9], and geophysics [10].

There are two major deficiencies inherent to AWCM, however, that have limited its application: 1) the “curse of anisotropy,” whereby isotropic wavelet refinement is unable to introduce a spatially varying aspect ratio of the mesh elements and 2) the reliance on a rectangular domain [11]. The first limitation is related to the adaptive mechanism used by AWCM to analyze and refine the mesh, which is isotropic in nature. It is unable to distinguish and exploit highly correlated

* Corresponding author at: Skolkovo Institute of Science and Technology, Skolkovo Innovation Center, 3 Nobel St., Moscow 143026, Russia.

E-mail addresses: O.Vasilyev@skoltech.ru, Oleg.Vasilyev@nwra.com, Oleg.Vasilyev@colorado.edu (O.V. Vasilyev).

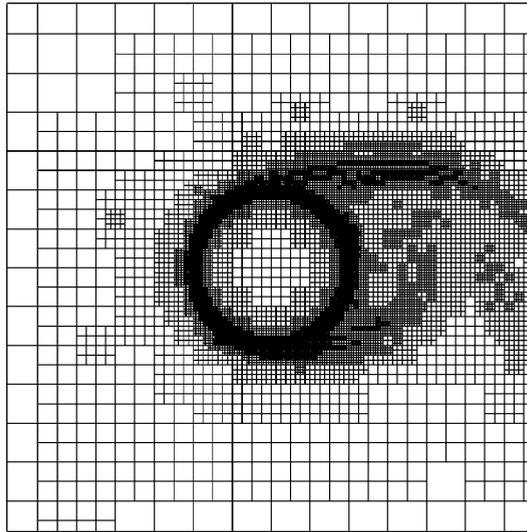


Fig. 1. The adaptive mesh for the unsteady flow around a cylinder imposed by volume penalization. Isotropic refinement via AWCM has created a very large number of points around the immersed surface and boundary layer.

directions in a variable field where the solution is changing less rapidly than in other directions. Such features, like flame sheets or interfaces, have anisotropic resolution requirements which permit coarser tangential discretization. Wavelet-based adaptation is unable to efficiently exploit these highly correlated and oriented structures. With the associated isotropic adaptation mechanism, increased refinement for gradients in a single direction triggers refinement in all directions irrespective of the need. These filament- and sheet-like structures consequentially become over-resolved, severely degrading the benefits of dynamic adaptivity.

The second limitation is that the use of efficient tensorial wavelets, such as a second generation wavelet basis [12], imposes the requirement for a rectilinear grid. This impedes the efficient application of AWCM towards problems involving complex geometries. In these cases, volume penalization methods have been employed to impose arbitrarily-shaped boundaries [13–16]. They are simulated through additional forcing terms on the continuous governing equations in nonphysical regions of the computational domain. While this approach circumvents the limitation of the grid topology, it comes with severe drawbacks, among which the most serious are a) approximate nature of boundary conditions, even with guaranteed a priori prescribed accuracy convergence, b) higher surface-normal resolution requirements, c) increased computational stiffness related to penalization terms, and d) additional computational cost associated with the penalization points within the nonphysical volume.

While alternative immersed boundary methods might bypass some of the difficulties associated with volume penalization, all of these approaches ultimately introduce surfaces into the AWCM domain, and therefore suffer from the previously mentioned costs associated with isotropic refinement. Fig. 1 shows a cross-section from the unsteady flow around a cylinder modeled with the characteristic-based volume penalization method [15]. Due to the solid interface and the high velocity gradient in the boundary layer, the mesh has adapted to the highest permitted level. While the resolution is appropriate in the wall-normal direction, it is excessive in the tangential. Refinement in this manner means that the number of points required to capture surface effects increases with the inverse square of the characteristic normal lengthscale, δ_N (the viscous boundary layer thickness) and the surface area S , resulting in the computational cost scaling as $O(S/\delta_N^2)$ in three dimensions. This quickly becomes prohibitive in many applications, such as high Reynolds number wall-resolved turbulent regimes or geometries with a large surface area.

Structured curvilinear meshes reconcile this need for anisotropic refinement with the topological limitations of adaptive wavelet grids. A flexible physical space mesh can be mapped to a regular rectilinear geometry in a computational space where AWCM can be implemented. For the purposes of this paper, the *mesh* is defined as the geometry of the physical-space discretization, while the *grid* is the structured, rectilinear assembly of collocation points defining the discrete topology and the computational space.

The use of curvilinear coordinate systems to leverage more optimal discretization is a well established approach for solving PDEs [e.g., 17,18] on nonadaptive grids. Features that are stationary and spatially straightforward to predict, such as boundary layers, benefit from the flexibility of curvilinear mappings to locally and anisotropically control resolution. Additionally, enabling body fitted meshes obviates the need for immersed boundary approaches in many cases.

While it may seem that the use of static anisotropic structured meshes is antithetical to the automated refinement of AWCM, their hybrid implementation is complimentary. Curvilinear meshes provide a continuous mapping that defines local anisotropy and yields a more optimal discretization. On the other hand, a rectangular computational space permits the native implementation of AWCM, which locally refines the grid based on structures in the transformed solution, guaranteeing

the accuracy of the solution and quality of the computational grid automatically and *a priori*. This property of solution assurance, irrespective of mesh geometry, is a distinguishing feature of A-AWCM. The coordinate transform then benefits the discretization by reducing the computational cost. Well constructed curvilinear mappings normalize anisotropy in the transformed field variables, permitting more efficient wavelet-based compression into a sparse data representation. The mapping fields themselves are maintained on the grid as additional variables so that AWCM ensures sufficient support to maintain transform operations at a consistent accuracy and, thus, the mesh quality.

In this paper, the general framework for an integrated Adaptive-Anisotropic Wavelet Collocation Method (A-AWCM) is described using curvilinear coordinates. Advantages over isotropic and static mesh optimization are demonstrated by implementing structured curvilinear meshes based on explicit analytic mappings. This approach, however, is presented as a fully general framework that can be implemented with a variety of meshing approaches, including the use of automated elliptic mesh generation [e.g., 17,19–22] as well as spatio-temporally evolving geometries [e.g., 23]. Assurance of consistent global accuracy through the dynamic local refinement of AWCM permits especially flexible use of redistribution and mapping algorithms. Even though the proposed framework is general and highly flexible, the development of automated meshing and anisotropic refinement based on the solution, statistical quantities of interest, and/or for geometries with moving and deformable boundaries is outside of the scope of this paper and left as the subject of future research.

2. Adaptive wavelet collocation mesh refinement

The adaptive wavelet collocation method is a generalized approach for compressing discrete data fields into sparse grid representations. It uses multiresolution analysis to dynamically adapt and refine local resolution to form grids that are suitable for the time-accurate solutions of PDEs.

A spatially dependent field, $u(\mathbf{x})$, is discretized onto a set of dyadic, nested grids and decomposed as

$$u(\mathbf{x}) = \sum_{\mathbf{l} \in \mathcal{L}^0} c_{\mathbf{l}}^0 \phi_{\mathbf{l}}^0(\mathbf{x}) + \sum_{j=1}^{j_{\max} 2^n - 1} \sum_{\mu=1} \sum_{\mathbf{k} \in \mathcal{K}} d_{\mathbf{k}}^{\mu,j} \psi_{\mathbf{k}}^{\mu,j}(\mathbf{x}), \tag{1}$$

where $c_{\mathbf{l}}^0$ and $d_{\mathbf{k}}^{\mu,j}$ are the respective coefficients on scaling functions $\phi_{\mathbf{l}}^0(\mathbf{x})$ and wavelet basis function, $\psi_{\mathbf{k}}^{\mu,j}(\mathbf{x})$. Multidimensional wavelet families, μ , arise from the tensorial product of one-dimensional wavelet and scaling functions, as defined on alternating grid points in each axis-major direction [1]. The resolution levels are identified by level j , where the higher are more refined. Each wavelet has localized wavenumber and spatial support. It is uniquely identified and corresponds one-to-one with the grid points given by \mathbf{k} and \mathbf{l} .

The sparse grid is created by filtering out collocation points whose wavelet coefficient falls below a prescribed threshold. The coefficients are effectively set to zero, and the corresponding points are discarded. The resulting approximation, denoted $u(\mathbf{x})_{\geq}$, is given by a subset of the original wavelet basis,

$$u(\mathbf{x})_{\geq} = \sum_{\mathbf{l} \in \mathcal{L}^0} c_{\mathbf{l}}^0 \phi_{\mathbf{l}}^0(\mathbf{x}) + \sum_{j=1}^{j_{\max} 2^n - 1} \sum_{\mu=1} \sum_{\substack{\mathbf{k} \in \mathcal{K} \\ |d_{\mathbf{k}}^{\mu,j}| \geq \varepsilon \|u\|}} d_{\mathbf{k}}^{\mu,j} \psi_{\mathbf{k}}^{\mu,j}(\mathbf{x}), \tag{2}$$

consisting of the dominant information containing modes. Typically, a relative filter threshold ε is used together with the absolute thresholding scale $\|u\|$. This characteristic amplitude scale is often taken as either the L_2 - or L_{∞} -norm of the field u or a derived quantity of interest, such as the Reynolds averaged fluctuating component, $u - \bar{u}$, for turbulent flows. The reconstruction error incurred through the wavelet filtering operation can be shown to be [24]

$$\|u_{\geq} - u\| \leq O(\varepsilon \|u\|). \tag{3}$$

For a sparse dataset, often a weighted norm of $u(\mathbf{x})$ is taken as the characteristic variational scale, since a strict L_2 norm across discrete sample points can become locally biased by local grid refinement and therefore unrepresentative. In cases where multiple variable fields are adapted upon, $\mathbf{u}(\mathbf{x}) = \{u_1(\mathbf{x}), u_2(\mathbf{x}), \dots\}$, the grid is taken as the union of the grids filtered for each component where the wavelet thresholds, $\boldsymbol{\varepsilon} = \{\varepsilon_1(\mathbf{x}), \varepsilon_2(\mathbf{x}), \dots\}$, can be prescribed independently for each variable of interest.

As the solution evolves, it is necessary to account for wavelets that might become significant in the next iteration or time step. Therefore, the filtered grid is augmented by nearest neighbors to every active point on the current, next higher and next coarser resolution levels. This is justified by the finite propagation rate of information across the grid, which is determined by the type of PDE and iterative or time integration method. Successively filtering and augmenting steps adaptively refine upon and dynamically track local structures in the signal.

The interpolating-lifted wavelet transform at the core of AWCM is recursively computed for one dimension by

$$d_{\mathbf{k}}^j = \frac{1}{2} (c_{2\mathbf{k}+1}^{j+1} - P_{\mathbf{k}}^j), \tag{4}$$

$$c_{\mathbf{k}}^j = c_{2\mathbf{k}}^{j+1} + \tilde{P}_{\mathbf{k}}^{j+1}, \tag{5}$$

where $c_{2\mathbf{k}+1}^{j+1}$ is a discrete sample of $u(\mathbf{x})$ at the $j + 1$ (next finer) resolution level, $d_{\mathbf{k}}^j$ are the details at j level of resolution, $P_{\mathbf{k}}^j$ is the corresponding interpolant taken from coarser nodes on the j level, and $\tilde{P}_{\mathbf{k}}^j$ is the interpolant from the details on the finer $j + 1$ level of resolution. Lagrange polynomial interpolants with finite stencils are used to calculate $P_{\mathbf{k}}^j$ and $\tilde{P}_{\mathbf{k}}^j$. For many implementations it is common to omit the update step (5), especially for distributed memory solvers [25]. The inverse transform is computed simply by applying the algebraic inverse of (4)–(5), starting from the coarsest resolution level. As new nodes are added during mesh adaptation, the solution is interpolated to them by explicitly including those collocation points with coefficient $d_{\mathbf{k}}^{\mu,j} = 0$ and performing the inverse transform operation. Finite differencing is used to calculate derivative approximations, for which wavelet interpolation provides stencil support for collocation points not included in the sparse, filtered grid in (2).

Multidimensional wavelets are the tensorial product of their 1D counterparts, therefore the multidimensional transform is taken by consecutive application of (4)–(5) across every direction, and recursively for each resolution level. While the wavelet basis functions for some multidimensional families, μ , are anisotropic in a strict sense, they cannot be exploited for anisotropic refinement as they result from the lifting transform itself. The families are a direct consequence of the discrete location on the collocation grid; furthermore, the shapes align with the rectilinear axes. Since AWCM tightly integrates mesh adaptation, interpolation, and the computational environment, attempting to impose an artificial anisotropic refinement method externally would introduce excessive artifacts by violating the consistency of the wavelet-space representation. Restricting refinement is equivalent to discarding significant information-carrying modes from the perspective of the wavelet filter, (2). This discrete damping is unlike the action of the filter thresholding in that its effects are unbounded. In such a case, summation across the wavelet basis would no longer form a valid approximation for the original signal $u(\mathbf{x})$ at the *a priori* threshold level. Even if such additional processing of the grid were tenable, the geometry only permits anisotropic control in the major axis directions.

Despite this shortcoming, tensorial lifted-interpolated wavelets are highly desirable for the solution of PDEs due to their simplicity and computational efficiency, as the transform (4)–(5) can be accomplished in $O(N)$ operations.

3. Adaptive-Anisotropic Wavelet Collocation Method framework for curvilinear meshes

The Adaptive-Anisotropic Wavelet Collocation Method circumvents the shortcomings of the wavelet transform by employing the general framework for a harmonic mesh map [e.g., 26], which maps the physical domain $\Omega_p(\mathbf{x})$ to the computational domain $\Omega_c(\xi)$ (both assumed to be Riemannian manifolds):

$$\xi(\mathbf{x}) : \Omega_p \rightarrow \Omega_c, \xi = (\xi_1, \xi_2, \xi_3), \mathbf{x} = (x_1, x_2, x_3). \tag{6}$$

This allows for flexibility of the mesh geometry and more optimal point distribution by partially decoupling the physical mesh from the natural coordinate system implied by the rectilinear computational grid. Spatially varying cell aspect ratios and orientation, including body fitting, can then be prescribed within the physical domain, Ω_p , based on the problem being solved. This mesh is mapped back onto a rectilinear grid, which is suitable for the discrete operators, including the wavelet transform and derivative approximations.

Assuming that the final mesh is well conditioned and of sufficient quality to support the discrete PDE solution, this approach is agnostic to the methods that are used to construct it. Explicit, analytic curvilinear mappings are sufficient for the theoretical consideration of this general framework, though their use is limited for applied problems of engineering interest. In such cases, more advanced mappings and automated generation methods, either continuous or discrete, would prove more useful.

Transform between the Cartesian physical space basis, $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$, and the computational curvilinear basis, $\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$, proceeds as

$$\mathcal{J}_{ij}\mathbf{e}_i = \mathbf{b}_j, \tag{7}$$

where \mathcal{J} is the Jacobian matrix,

$$\mathcal{J}_{ij} = \frac{\partial x_i}{\partial \xi_j}, \tag{8}$$

and summation is implied across indices.

Given that the computational grid is rectilinear, the unmodified AWCM approach is directly implemented, where it adapts and refines the grid upon the transformed solution fields $\mathbf{u}(\xi)$. The inverse coordinate mappings $\mathbf{x}(\xi)$, which define the physical mesh geometry and support the discrete version of (7), are likewise tracked. The Jacobian matrix \mathcal{J}_{ij} is therefore readily computed from these mappings using the integrated finite difference framework. For the construction of PDE terms, which are naturally prescribed in physical space, computational derivative approximations are calculated as

$$\frac{\partial}{\partial x_j} = \mathcal{J}_{ji}^{-1} \frac{\partial}{\partial \xi_i}, \tag{9}$$

where \mathcal{J}_{ji}^{-1} is the inverse Jacobian matrix.

Adapting upon the coordinate mappings, $\mathbf{x}(\xi)$, with a low, independent wavelet threshold, $\varepsilon_{\mathbf{x}}$, ensures consistent accuracy of transformation operations, and automatically controls the quality of the sparse representation of the mesh. Spurious artifacts can otherwise arise if the local accuracy of the discrete Jacobian matrix approximation (8) deteriorates as a result of grid coarsening.

On the transformed solution, $u(\xi)$, the wavelet filter maintains *a priori* accuracy of the L_2 - and L_∞ -norm in the usual manner following (3). Given that the transform of u is purely spatial, the corresponding L_∞ -norms are equal,

$$\|u(\xi) - u(\xi)_{\geq}\|_{\infty} = \|u(\mathbf{x}) - u(\mathbf{x})_{\geq}\|_{\infty}. \quad (10)$$

Furthermore, the L_2 -norms can be shown to be equivalent, as $\int_{\Omega_p} f(\mathbf{x})d\mathbf{x} = \int_{\Omega_c} f(\mathbf{x}(\xi))\mathcal{J}(\xi)d\xi$. Therefore, error convergence of the physical space solution, $u(\mathbf{x})$, which is of primary interest, is likewise guaranteed by (3).

Local anisotropic refinement and unidirectional stretching of the physical space mesh is introduced by adjusting the location and distribution of points through the definition of the coordinate mappings, $\mathbf{x}(\xi)$. This compliments wavelet-based adaptation, which adjusts the local mesh point density but preserves the underlying geometry and local aspect ratio of the physical mesh. With properly constructed meshes, highly correlated gradients in physical space become normalized once transformed to computational space. This results in more efficient isotropic refinement through AWCM, a smaller memory footprint, and a more optimal use of computational resources. However, since accuracy convergence through AWCM is unified in both spaces, Ω_p and Ω_c , the solution *fidelity* is independent of the anisotropic refinement. With this general decoupling of the accuracy from the mesh geometry, A-AWCM is more robust and flexible than nonadaptive structured meshes. Note that this property is contingent upon a well conditioned geometry that is capable of numerically supporting the transform operations. High skewness, degenerate Jacobian matrix ($\mathcal{J} \rightarrow 0$), and non-smooth mesh lines can cause a degradation of the transform operations (7) that can not be compensated for by wavelet grid refinement. One consequence of solution/geometry independence is that it essentially eliminates the need for traditional *a posteriori* mesh convergence studies to guarantee the solution quality. Furthermore, it provides enormous robustness and flexibility for moving mesh [e.g., 23] and mesh redistribution [e.g., 27] methods for time-marching integration schemes, as AWCM will preserve the solution if the mesh geometry becomes temporarily sub-optimal.

In general, well constructed mapping fields, $\mathbf{x}(\xi)$, can be provided by a variety of means. In the simplest configurations and with *a priori* understanding of the solution field, explicit algebraic functions can impose body fitting and provide some anisotropic benefit to the computational cost. More complex problem geometry and more optimal mesh distribution are typically achieved through automated methods, including the solution of elliptic or hyperbolic PDEs [17,21,22]. Many approaches involve dynamic redistribution and moving meshes [23], either as a cost optimization or to capture moving/deformable geometry. Development of automatic meshing, mesh redistribution, and anisotropic refinement methods that exploit the particular computational environment and take advantage of A-AWCM are a topic of future research.

Each time the grid is dynamically refined through the wavelet mechanism, the mesh must be updated with the physical space positions of the new points in the discrete mapping. Added nodes are treated in one of two ways, depending on how the mesh is explicitly defined. For meshes that are exactly defined at every potential grid point, either through the deterministic solution of an equation (e.g. algebraic or PDE) or some other procedural method, the mapping must be consistently prescribed for every point that is introduced. Fully discrete maps must have some mechanism for updating grid points, such as a multiresolution nested definition or the ability to be refined at will. For some methods, this can involve a potentially expensive global/iterative solution that can add substantial cost to the refinement. Alternatively, the mesh can be strictly defined by the continuous, static wavelet representation $\mathbf{x}_{0,\geq}$ of any arbitrary mapping, \mathbf{x}_0 , filtered to a fixed subset of nonzero wavelet basis functions $\mathcal{K}_0 \subset \mathcal{K}$. From this, the value of the map (and thus the physical coordinate) of a new point at any location in Ω_p is interpolated from $\mathbf{x}_{0,\geq}$ using the wavelet transform. The subset, \mathcal{K}_0 , of wavelet functions at grid points and the coefficient values are determined by an initial mesh generation step, and retained throughout the course of solution. For static geometries, this has the ability to amortize the cost of highly sophisticated generation methods and to ensure that grid refinement will maintain a piecewise-smooth mesh suitable for the polynomial-based numerical framework. It should be noted that the adaptive wavelet filter is not positive preserving, so care must be taken to ensure that the continuous field \mathbf{x}_{\geq} is monotonic to within a tolerance specified by the finest mesh resolution. This can be enforced by AWCM in the usual manner.

4. Numerical results

The A-AWCM framework is fully general for the solution of elliptic, parabolic and hyperbolic PDEs. In order to demonstrate the central principles of the method, it has been applied and examined here for parabolic and hyperbolically dominant cases. Corresponding mesh geometry includes body fitting and the introduction of local anisotropy to match general trends in the solution. This is shown to optimize dynamically adaptive meshes while retaining the fidelity control and error convergence inherent to AWCM.

While sophisticated meshing techniques can be employed, analytic mappings are used in these cases for simplicity and to highlight the general framework and demonstrate the concepts of A-AWCM, namely the implementation of spatially inhomogeneous anisotropic meshes, *h*-refinement in arbitrary geometries, and rigorous error control. Development of automated meshing techniques that exploit this framework is left as a topic for future research.

The cases considered here are parabolic diffusion, acoustic scattering, and unsteady external compressible viscous flow. While AWCM, and by extension A-AWCM, can be implemented with arbitrary order wavelets and finite-differences [3], all of the results shown here use 4-th order for both.

4.1. Parabolic diffusion

Time-accurate parabolic diffusion is examined on a 2D rectangular domain, defined by the physical bounds $\mathbf{x} \in [0, \pi] \times [0, \pi]$. A nonhomogeneous profile of the diffusive scalar, u , is imposed along one edge, generating a steep initial gradient that drives diffusion into the domain.

The parabolic diffusion equation is defined in physical space as

$$\frac{\partial u}{\partial t} = \frac{\partial}{\partial x_j} \left(\kappa \frac{\partial u}{\partial x_j} \right), \tag{11}$$

where the diffusion coefficient used here is constant, $\kappa = 1$. In order to be solved, it is transformed into the computational coordinates as

$$\frac{\partial u}{\partial t} = \mathcal{J}_{ji}^{-1} \frac{\partial}{\partial \xi_i} \left(\kappa \mathcal{J}_{ji}^{-1} \frac{\partial u}{\partial \xi_i} \right). \tag{12}$$

It should be noted that this transform is fully procedural and can be easily abstracted, minimizing the cost of implementation in existing codebases.

The boundary conditions are given as

$$\left. \frac{\partial u}{\partial x} \right|_{x=0, \pi} = 0, \tag{13}$$

$$u|_{y=0} = 0, \tag{14}$$

$$u|_{y=\pi} = f(x), \tag{15}$$

where the Cartesian coordinates $\mathbf{x} = (x, y)$ will be used interchangeably with $\mathbf{x} = (x_1, x_2)$ for simplicity of the discussion. This convention is used hereafter.

A fifth order polynomial is specified for $f(x)$. The initial condition is defined by the series

$$u|_{t=0} = \sum_m \left[A_m \frac{\sinh(my)}{\sinh(m\pi)} + \sum_n C_{m,n} \sin(ny) \right] \cos(mx), \tag{16}$$

where the expansion coefficients A_m , and $C_{m,n}$ are computed to match the boundary conditions, (13)–(15), and an initial steep exponential decay in the y direction. This provides a smooth approximation of an impulsive perturbation at the $y = \pi$ edge, generating a steep gradient in the transported scalar.

The grid is given by a base mesh of $N = [5 \times 5]$ points, with $j_{\max} = 10$ levels of resolution, yielding equivalent, non-adaptive resolution of $N_{\max} = [2049 \times 2049]$. The physical mesh is intentionally non-tensorial and introduces skewness to demonstrate the generality of the method. An algebraic curvilinear, two-dimensional coordinate mapping, $\mathbf{x} = \{x, y\}$, define the mesh across the computational domain $\xi = \{\xi, \eta\} \in [0, \pi] \times [0, \pi]$ as

$$x = \phi \left[\frac{-1.9}{\pi^2} \xi^3 + \frac{2.85}{\pi} \xi^2 - 0.95\xi \right] + (1 - \phi) \xi, \tag{17}$$

$$y = \frac{-0.65}{\pi^2} \eta^3 + \frac{0.45}{\pi} \eta^2 + 1.1\eta, \tag{18}$$

where a blending function controls the local expression of x -direction anisotropic stretching,

$$\phi = \frac{-2}{\pi^3} \eta^3 + \frac{3}{\pi^2} \eta^2. \tag{19}$$

The full grid at the $j = 3$ resolution level is illustrated in Fig. 2. This mesh has been stretched in the y -direction to more efficiently resolve the steep initial gradient that is formed on the $y = \pi$ edge at early timesteps. The mild stretching in the x -direction near the perturbed edge captures some of the variability in the boundary profile, $f(x)$. In comparison, a standard isotropic wavelet grid in Fig. 3 shows increased refinement in the corners resulting from the initial temperature field and the one-sided interpolating stencils near the grid edges.

It should be noted that this particular mesh has been generated to approximate general trends in the solution, but has not been intentionally or optimally constructed to capture the functional form of the initial conditions, as wavelet refinement provides any additional refinement and automatically filters out over-resolved regions.

The error convergence at $t = 0.05$ is shown in Fig. 4. As with the standard rectilinear implementation, both the physical space L_2 - and L_∞ -norm are well controlled and scale with the adaptive wavelet threshold parameter as $O(\varepsilon)$. Furthermore, for standard rectilinear AWCM the total number of grid points, \mathcal{N} , scales with the wavelet threshold as $O(\mathcal{N}^{-p/n})$, where p is the order of the wavelets and spatial derivative operator and n is dimensionality of the problem [3].

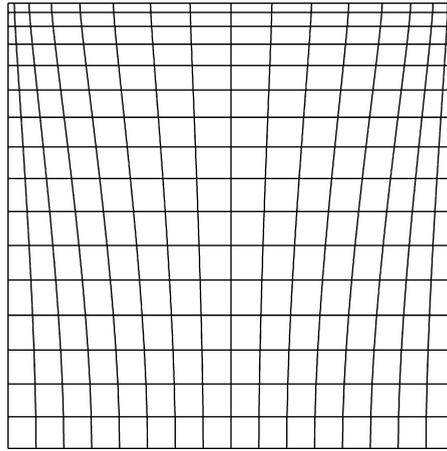


Fig. 2. The non-adapted mesh at the $j = 3$ resolution level used to solve a parabolic diffusion equation across a square plate. Anisotropic refinement is prescribed on the top boundary to account for steep gradients at early timesteps.

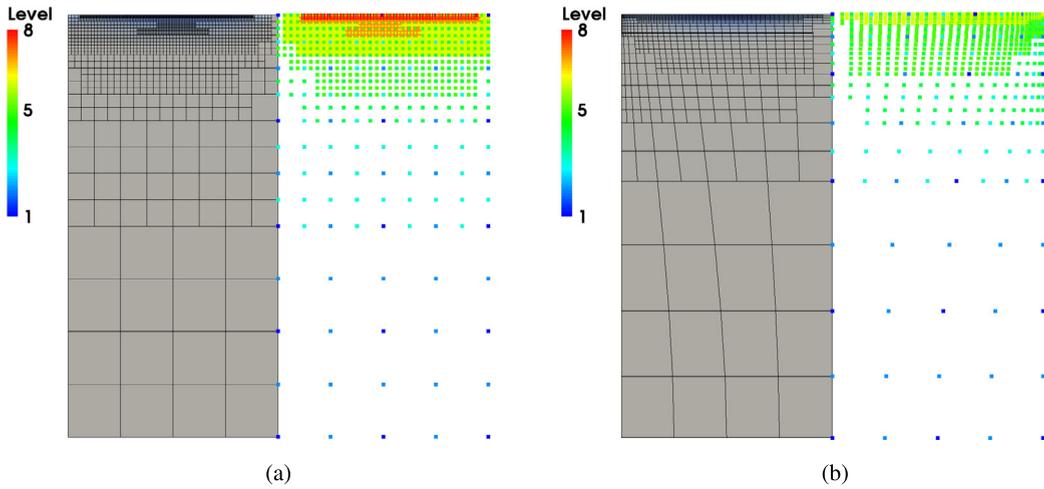


Fig. 3. The mesh as adapted upon the initial conditions of u with a steep decay from a quintic polynomial BC. Using the same wavelet threshold, $\varepsilon = 10^{-3}$, the fully rectilinear mesh (a) uses both a greater number of total points and a higher resolution level than the more optimal curvilinear mesh geometry (b). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

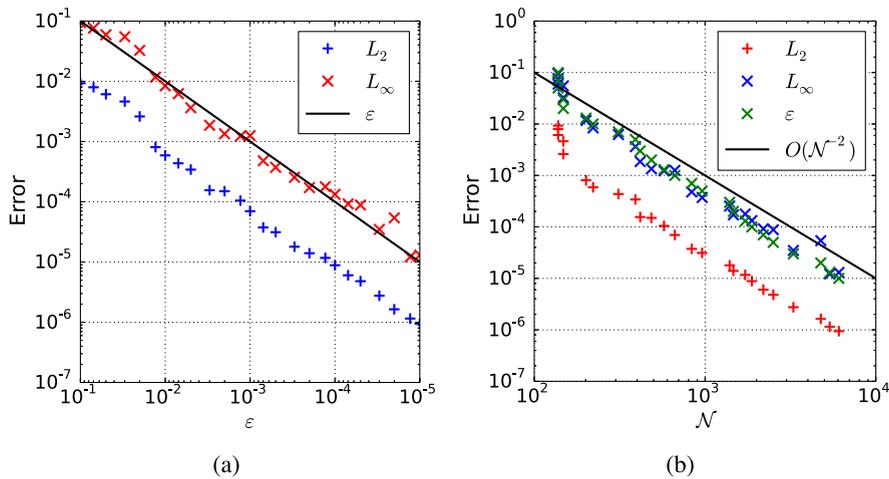


Fig. 4. Both L_2 - and L_∞ -norm errors of the physical-space numerical solution converge linearly with the wavelet threshold parameter, ε . As shown by Vasilyev and Kevlahan [3] for 4-th order wavelets and finite differences, the error and threshold ε scale with the total number of grid points as $O(\mathcal{N}^{-2})$.

4.2. Acoustic scattering

For an examination of hyperbolically-dominant PDEs, consider the acoustic scattering problem of Tam and Hardin [28]. A low amplitude pulse reflects off of the surface of a cylinder, and is scattered into the far field. Here the cylinder is modeled using a body-fitted O-type mesh with appropriate BCs, obviating an immersed boundary approach.

The fully compressible Navier–Stokes equations

$$\frac{\partial \rho}{\partial t} = - \frac{\partial \rho u_j}{\partial x_j}, \tag{20}$$

$$\frac{\partial \rho u_i}{\partial t} = - \frac{\partial}{\partial x_j} (\rho u_i u_j + P \delta_{ij} - \tau_{ij}), \tag{21}$$

$$\frac{\partial \rho e}{\partial t} = - \frac{\partial}{\partial x_j} (\rho u_j (e + P) - u_i \tau_{ij} + q_j) \tag{22}$$

are transformed and solved in the curvilinear coordinates in a rectangular computational space, Ω_c , as before.

The energy for the acoustic signal is provided by a small, Gaussian perturbation in the initial condition,

$$u' = p' = 10^{-3} \times 2^{-\frac{(x-c)^2 + y^2}{0.04}}, \tag{23}$$

superimposed on a quiescent field, centered at $c = 4$. Nonreflective boundary conditions in the far field are simulated by a characteristic-based buffer region [29], where convective forcing is aligned with the radial direction for computational efficiency.

A high acoustic Reynolds number of $Re_a = 10^5$, based on the cylinder diameter and far field fluid properties, is used to approximate the Euler equations and match the analytic solution given for an inviscid fluid [28]. Given the finite Reynolds number, the BCs on the cylinder are prescribed to be no-slip and isothermal, resulting in a minute boundary layer forming during the transient interaction between the acoustic pulse and the solid surface. Correspondingly, the mesh is anisotropically refined in the radial direction by the function

$$r = 1 + \frac{\tanh\left(\frac{3}{2} \left(\frac{\xi - \xi_{\min}}{\xi_{\max} - \xi_{\min}} - 1\right)\right)}{\tanh\left(\frac{3}{2}\right)} \tag{24}$$

in order to efficiently resolve the steep gradient that will be generated in the wall normal direction. At the cylinder surface, the mesh takes upon a high aspect ratio of $A \approx 4 : 1$ in physical space, with the first gridpoint at a distance of $r/D \approx 0.00014$ at the highest resolution level. Mesh lines are evenly spaced in the azimuthal direction. The computational grid at maximum resolution is equivalent to $N_{\max} = [3073 \times 5632]$, and utilizes a limit of 10 levels of adaptive refinement, sufficient to fully resolve all simulations without artificial truncation. Fig. 5 shows the mesh at the $j = 3$ level of resolution. As with the previous problem, only very general and conservative considerations have been taken in construction of this mesh. The use of AWCM ensures that regions of over- or under-resolution will be refined accordingly to optimize the grid for the underlying geometry.

The reflection of the acoustic pulse is shown in Fig. 6 for a threshold of $\varepsilon = 5 \times 10^{-3}$. In the locations of high strain rate at the cylinder surface, the mesh has refined to a high level, while remaining coarser at other surface regions. As an explicit Runge–Kutta method is used for time integration, the small resolved spatial scales severely, though temporarily, restrict the computational timestep, Δt , in order to maintain stability in accordance with the CFL condition. This is limited to a relatively concise window in time when the acoustic pulse interacts with the cylinder.

The structure of the acoustic wave spans a large portion of the domain. In the far field, a greater number of levels of refinement are used to resolve the wave structure. Typically, structured O-type meshes yield non-uniform fidelity, as the mesh becomes stretched and enlarged as the radius increases, degrading the resolution power and spatial accuracy of derivative approximations. With A-AWCM, however, a uniform accuracy specified by the wavelet threshold, ε , is maintained at all radial locations, provided that there are sufficient adaptive levels, j_{\max} , to support the necessary refinement. Fig. 6b shows an approximation of the local element size across the physical domain. Consistent physical resolution is provided for the wave structure that spans the entire radial distance. The physical mesh geometry, therefore, does not degrade the solution quality. This ability to maintain consistent spatial representation, and therefore globally constant fidelity, is unique to A-AWCM.

Furthermore, A-AWCM is able to exploit transients in the solution. Before and after the period where the main acoustic waves are actively interacting with the solid surface, the strain rate at the surface is negligible, and the mesh is refined only to a coarse level as shown in Fig. 7. The resulting timestep is then much larger, and only limited by the mesh spacing required to resolve the pulse in the far field.

The pressure at point $\mathbf{x} = (\sqrt{2}, \sqrt{2})$ is shown in Fig. 8. Both the initial wave and its reflections from the cylinder are visible. As the wavelet threshold is decreased, the error and spurious artifacts in the transient solution can be clearly seen to converge towards the analytic inviscid solution.

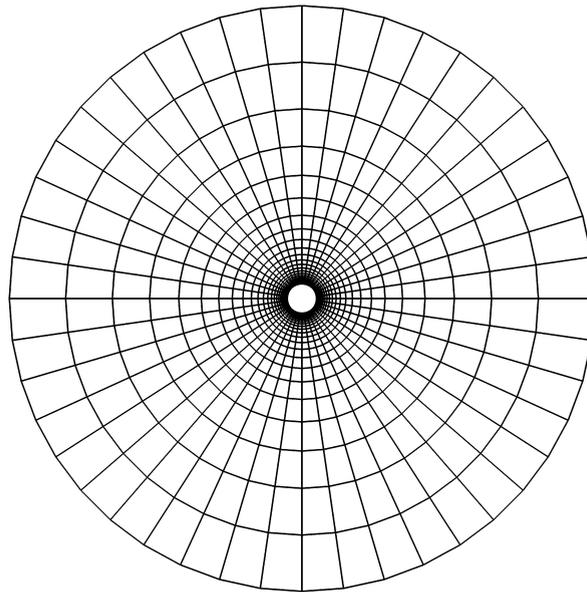


Fig. 5. An O-type mesh is used to simulate acoustic scattering off of a solid cylinder. The mesh is shown with all gridlines at the $j = 3$ resolution level displayed. Stretching in the radial direction introduces anisotropy in the mesh at the surface of the cylinder in order to more efficiently resolve the boundary layer.

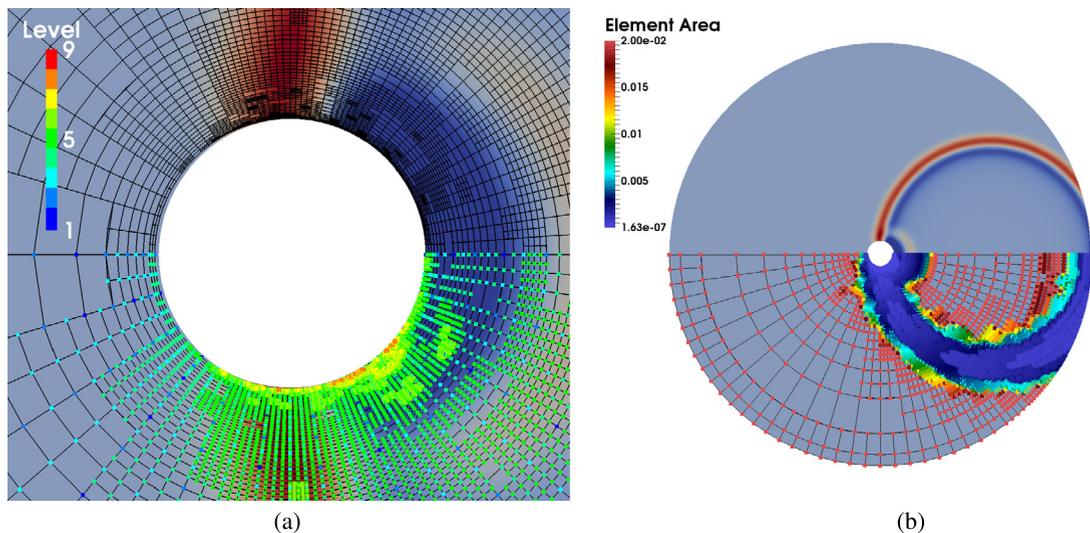


Fig. 6. The pressure field at $ta/D = 4$, when the acoustic wave is passing and reflecting off of the surface of the cylinder. The minuscule boundary layer is causing the grid to adapt to the $j = 9$ level with $\varepsilon = 5 \times 10^{-3}$. Examination of the approximate element area shows that uniform resolution is maintained through wavelet-based adaptation across the acoustic pulse, which spans from the cylinder to the far field, regardless of local stretching in the mesh geometry. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.3. External flow

For unsteady, 2-D flow around a cylinder, the compressible Navier–Stokes equations (20)–(22) are solved for $Re = 250$ at subsonic Mach number $Ma = 0.1$, where the Reynolds number is based on the inflow velocity, cylinder diameter, and far-field fluid properties. In this configuration, the governing system of PDEs demonstrates both strong parabolic, hyperbolic, and convective effects. While the flow is still laminar at this Reynolds number, a well structured von-Kármán vortex street forms in the cylinder wake.

The computational grid has an equivalent finest resolution of $N_{\max} = [640 \times 384]$, employing 7 adaptive levels. In physical space, an O-type mesh is used, with azimuthal and radial distribution of mesh points following that of Beaudan and Moin [30]. A more conservative estimation of the wake region is prescribed to ensure that the vortex street is fully resolved. Wavelet-based adaptation compensates for the overly dense mesh geometry. The boundary layer is anisotropically refined

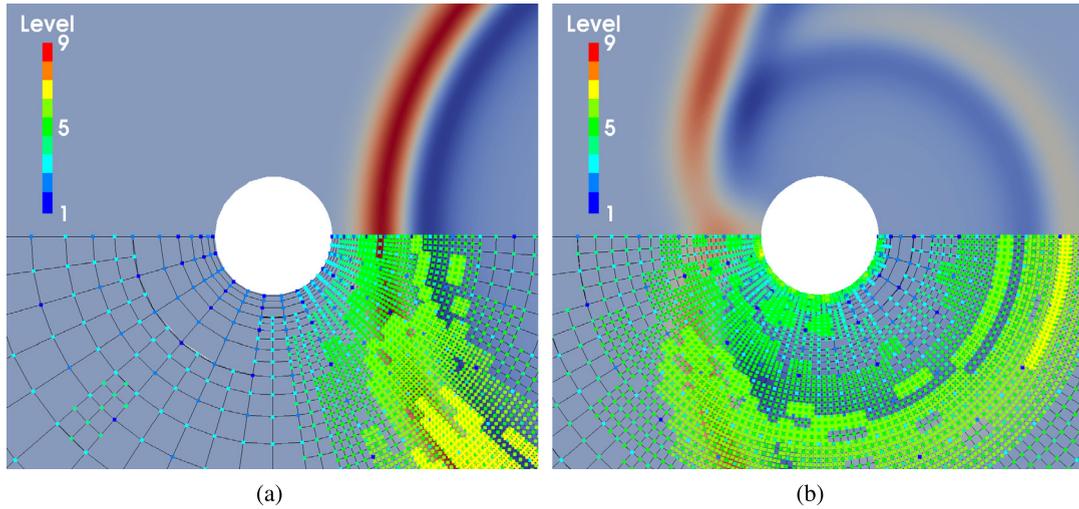


Fig. 7. The pressure field and adaptive mesh at $ta/D = 3$ (a) and $ta/D = 5$ (b). When the main pulse is not directly interacting with the cylinder, the grid locally coarsens, thereby relaxing the explicit time step through the CFL condition. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

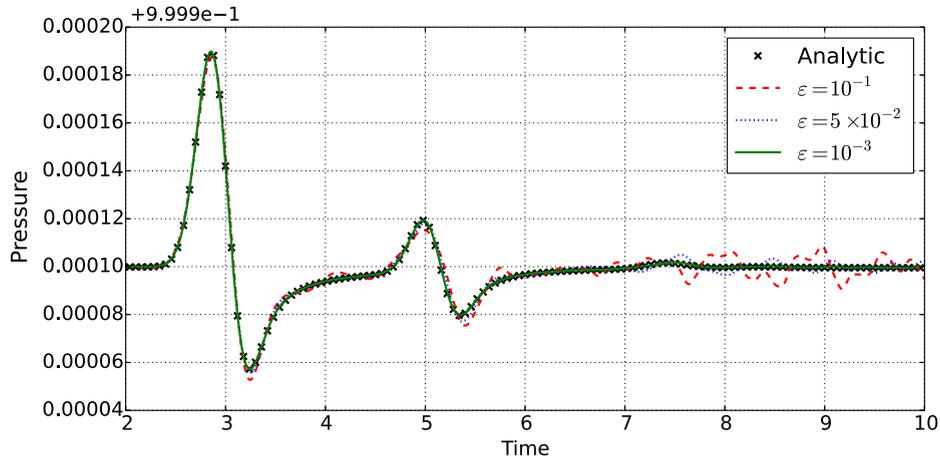


Fig. 8. Acoustic scattering measured at point $\mathbf{x} = (\sqrt{2}, \sqrt{2})$. The error and numerical artifacts in the reflected wave can clearly be seen to converge as ϵ decreases. Correspondingly, the percentage of the total nodes retained in the adaptive grid at $t = 4.0$ is $\mathcal{R} = 0.9\%$ and 1.3% for $\epsilon = 10^{-1}$, and 5×10^{-2} , respectively, which were resolved using 8 resolution levels. Since the case where $\epsilon = 10^{-3}$, required 10 levels to be fully resolved, a direct comparison of the compression ratios is not possible.

using geometric cell growth starting from an aspect ratio of approximately $2.5 : 1$ in the radial direction, with the first point from the surface placed at $r/D \approx 0.0032$. The non-adaptive mesh at the $j = 4$ level is shown in Fig. 9. In order to capture important physical behavior in the solution, the wavelet filter is set to adapt upon momentum, vorticity and the magnitude of the strain rate tensor. Three simulations were run with the normalized wavelet threshold filter set to ranges across $10^{-3} \leq \epsilon \leq 10^{-1}$.

Time-accurate lift and drag coefficients are shown in Fig. 10, and the statistical results are summarized in Table 1. For the most well refined case, where $\epsilon = 10^{-3}$, the average drag coefficient (\bar{C}_D), RMS of the lift coefficient ($C_{L,RMS}$), and Strouhal number (St) for the shedding frequency agree well with established results [31,32]. At the higher threshold value of $\epsilon = 10^{-2}$, indicating discretization error of $O(1\%)$, these low-order statistics are reasonably approximated using only 5.1% of the total grid points. At the highest threshold value, the solution is only very roughly approximated on the nearly coarse grid.

From Fig. 11b, the advantage of a body fitted mesh with A-AWCM is clearly visible. The resolution is not uniform along the surface, but rather tracks unsteadiness within the boundary layer and optimizing the number of points to resolve separately the stagnation and separation points, and near-wake bubble. There are only very limited regions where the highest resolution level has been reached, indicating that the solution is indeed captured at a numerical accuracy $O(10^{-3})$ without effects from the finite available resolution levels. Furthermore, AWCM exploits the near-surface intermittency of

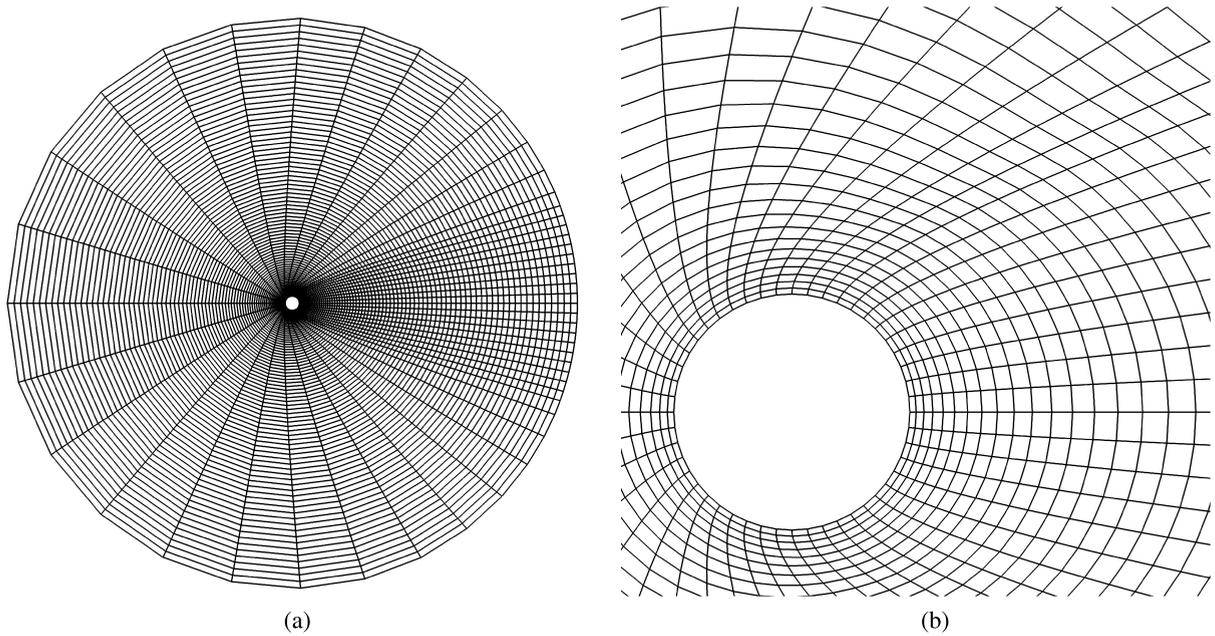


Fig. 9. An O-type mesh following the distribution method of Beaudan and Moin [30] is used to resolve the unsteady wake and boundary layer for external flow past a cylinder. It is shown here with all gridlines at the $j = 4$ resolution level.

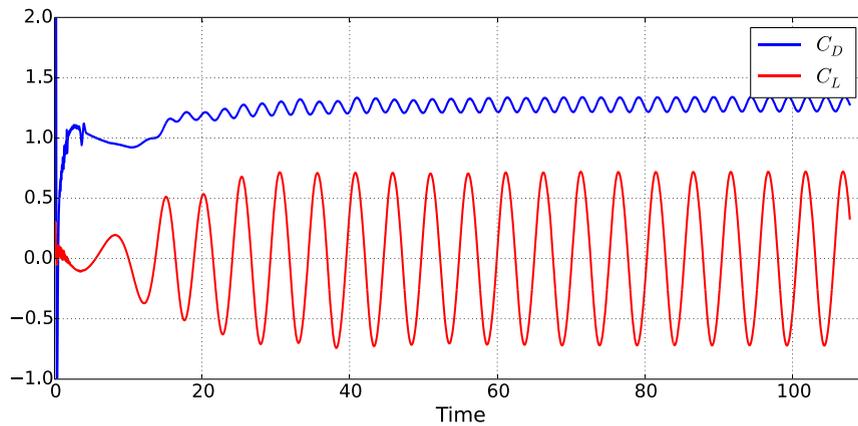


Fig. 10. Lift and drag forces calculated on the cylinder ($\varepsilon = 10^{-3}$).

Table 1

Computed drag (\bar{C}_D) and lift ($C_{L,RMS}$) statistics, Strouhal number (St) and percentage of adaptive points retained (\mathcal{R}) for pseudo-incompressible flow around a cylinder at $Re = 250$ at various values of the wavelet threshold. The values were computed for 15 shedding cycles once the solution had reached statistically steady state.

ε	\bar{C}_D	$C_{L,RMS}$	St	\mathcal{R}
10^{-1}	1.026	0.312	0.195	3.6%
10^{-2}	1.265	0.495	0.201	5.1%
10^{-3}	1.273	0.507	0.204	19.5%

the oscillating shear layer resulting from boundary layer separation. The anisotropically stretched elements on the surface reduce the number of refinement levels that are used to resolve much of the boundary layer.

As vortical structures are convected into the far-wake, increasingly higher resolution levels again compensate for the larger cells inherent to O-meshes, maintaining highly accurate spatial resolution. While the smallest physical lengthscales are typically found at the surface and in the near-wake, the curvilinear transformation effectively regularizes the regional scales in addition to providing anisotropy for the boundary layer. This allows AWCM to compensate for the underlying stretching and coarsening that occurs in the far-field of O-type meshes without the compounding costs of over-resolving the near region.

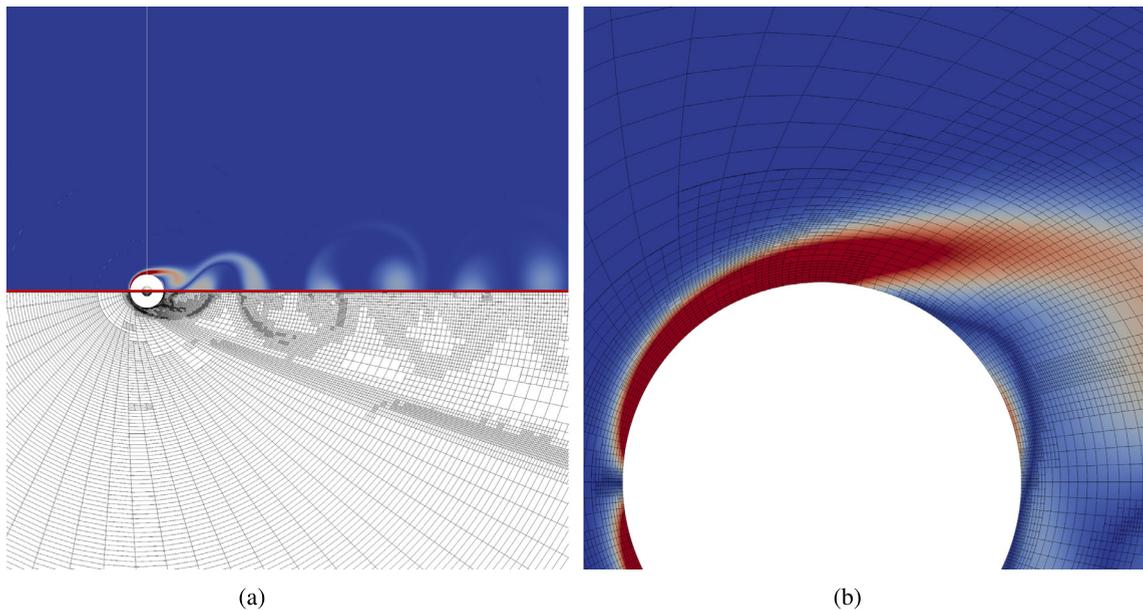


Fig. 11. Vorticity field and the adaptive mesh for external flow. Dissipation of vortices in the wake causes decreased grid density farther from the cylinder. To maintain the accuracy of the transform operations, adaptation occurs upon the mesh mappings $\mathbf{x}(\xi)$, and is clearly seen in (a).

5. Conclusions

The general framework for A-AWCM leverages the advantages of both curvilinear mesh geometries and wavelet-based adaptive refinement. The use of a transformed coordinate system allows the simultaneous use of complex geometry meshes as well as efficient wavelet algorithms for the solution of parabolic, elliptic, and hyperbolic PDEs. Spatially varying anisotropic and complex geometries can therefore be introduced to more optimally discretize the domain for a given problem, thereby circumventing the deficiencies that accompany the isotropic refinement mechanism of AWCM. This permits greatly reduced computational resources, especially with solutions that contain filament- and sheet-like structures that are often prohibitively expensive for isotropic refinement approaches. Furthermore, this enables AWCM to be used in a body-fitted context for solving applied problems.

Within this framework, many of the most attractive properties of AWCM have been preserved. The grid adapts to track structures in the solution and thereby exploits intermittency to greatly reduce the memory footprint. Fidelity of the physical space solution is rigorously controlled *a priori* through the wavelet threshold, ε , and error convergence is demonstrated for both the L_2 and L_∞ norms. Furthermore, the accuracy is maintained globally. The adaptive wavelet transform yields uniform spatial discretization in physical space regardless of the underlying mesh geometry, provided that the local mesh structure is well conditioned for transformation operations. This is a unique property of A-AWCM, and largely decouples the mesh geometry from the quality of the solution. The mesh therefore generally implicates only the computational cost and efficiency of the discretization.

While this method has been demonstrated on a series of problems with analytically created meshes, the framework itself is agnostic to the particular generation method. It allows for the use of more sophisticated and automated meshing and mesh redistribution approaches. A high degree of flexibility can be exercised with these methods since a consistent solution fidelity can be guaranteed through the adaptive wavelet mechanism. This implies an optimization of the grid for a given geometry by removing excess grid points in over-resolved regions. Advanced meshing methods that exploit the particularities of the integrated A-AWCM environment are topics of ongoing research.

Acknowledgements

Partial support for this work was provided by the Office of Naval Research under Grant No. N00014-11-1-069. OVV was also partially supported by Russian Science Foundation (Project 16-11-10350) and the Polish National Science Centre (NCN) project DEC-2013/10/M/ST6/00531. This support is gratefully acknowledged.

References

- [1] O.V. Vasilyev, C. Bowman, Second generation wavelet collocation method for the solution of partial differential equations, *J. Comput. Phys.* 165 (2000) 660–693.
- [2] O.V. Vasilyev, Solving multi-dimensional evolution problems with localized structures using second generation wavelets, in: *Special Issue on High-Resolution Methods in Computational Fluid Dynamics*, *Int. J. Comput. Fluid Dyn.* 17 (2003) 151–168.

- [3] O.V. Vasilyev, N.K.R. Kevlahan, An adaptive multilevel wavelet collocation method for elliptic problems, *J. Comput. Phys.* 206 (2005) 412–431.
- [4] J.D. Regele, O.V. Vasilyev, An adaptive wavelet-collocation method for shock computations, *Int. J. Comput. Fluid Dyn.* 23 (2009) 503–518.
- [5] D.E. Goldstein, O.V. Vasilyev, Stochastic coherent adaptive large eddy simulation method, *Phys. Fluids* (1994-present) 16 (2004) 2497–2513.
- [6] A. Nejadmalayeri, A. Vezolainen, O.V. Vasilyev, Reynolds number scaling of CVS and SCALES, *Phys. Fluids* 25 (2013).
- [7] G. De Stefano, O.V. Vasilyev, Wavelet-based adaptive simulations of three-dimensional flow past a square cylinder, *J. Fluid Mech.* 748 (2014) 433–456.
- [8] J.D. Regele, D.R. Kassoy, O.V. Vasilyev, Effects of high activation energies on acoustic timescale detonation initiation, *Combust. Theory Model.* 16 (2012) 650–678.
- [9] S. Paolucci, Z.J. Zikoski, D. Wirasaet, WAMR: an adaptive wavelet method for the simulation of compressible reacting flow. Part I. Accuracy and efficiency of algorithm, *J. Comput. Phys.* 272 (2014) 814–841.
- [10] N.K. Kevlahan, T. Dubos, M. Aechtner, Adaptive wavelet simulation of global ocean dynamics, *Geosci. Model Dev. Discuss.* 8 (2015).
- [11] A. Cohen, Adaptive methods for PDE's wavelets or mesh refinement?, in: *Proceedings of the ICM*, vol. 1, 2002, pp. 607–620, arXiv:math/0212414.
- [12] W. Sweldens, The lifting scheme: a construction of second generation wavelets, *SIAM J. Math. Anal.* 29 (1998) 511–546.
- [13] O.V. Vasilyev, N.K.R. Kevlahan, Hybrid wavelet collocation-Brinkman penalization method for complex geometry flows, *Int. J. Numer. Methods Fluids* 40 (2002) 531–538.
- [14] N.K.R. Kevlahan, O.V. Vasilyev, An adaptive wavelet collocation method for fluid-structure interaction, *SIAM J. Sci. Comput.* 6 (2005) 1894–1915.
- [15] E. Brown-Dymkoski, N. Kasimov, O.V. Vasilyev, A characteristic based volume penalization method for general evolution problems applied to compressible viscous flows, *J. Comput. Phys.* 262 (2015) 344–357.
- [16] Q. Liu, O.V. Vasilyev, Brinkman penalization method for compressible flows in complex geometries, *J. Comput. Phys.* 227 (2007) 946–966.
- [17] J.F. Thompson, F.C. Thames, C.W. Mastin, Automatic numerical generation of body-fitted curvilinear coordinate system for field containing any number of arbitrary two-dimensional bodies, *J. Comput. Phys.* 15 (1974) 299–319.
- [18] R.H. Pletcher, J.C. Tannehill, D. Anderson, *Computational Fluid Mechanics and Heat Transfer*, CRC Press, 2012.
- [19] A.M. Winslow, Numerical solution of the quasilinear Poisson equation in a nonuniform triangle mesh, *J. Comput. Phys.* 1 (1966) 149–172.
- [20] J.F. Thompson, F.C. Thames, C.W. Mastin, Tomcat—a code for numerical generation of boundary-fitted curvilinear coordinate systems on fields containing any number of arbitrary two-dimensional bodies, *J. Comput. Phys.* 24 (1977) 274–302.
- [21] P.D. Thomas, J.F. Middlecoff, Direct control of the grid point distribution in meshes generated by elliptic equations, *AIAA J.* 18 (1980) 652–656.
- [22] U.K. Kaul, Three-dimensional elliptic grid generation with fully automatic boundary constraints, *J. Comput. Phys.* 229 (2010) 5966–5979.
- [23] T. Tang, Moving mesh methods for computational fluid dynamics, *Contemp. Math.* 383 (2005) 141–174.
- [24] D.L. Donoho, *Interpolating Wavelet Transforms*, Technical Report, Department of Statistics, Stanford University, 1992.
- [25] A. Nejadmalayeri, A. Vezolainen, E. Brown-Dymkoski, O.V. Vasilyev, Parallel adaptive wavelet collocation method for PDEs, *J. Comput. Phys.* 298 (2015) 237–253.
- [26] A.S. Dvinsky, Adaptive grid generation from harmonic maps on Riemannian-manifolds, *J. Comput. Phys.* 95 (1991) 450–476.
- [27] H. Tang, T. Tang, P. Zhang, An adaptive mesh redistribution method for nonlinear Hamilton–Jacobi equations in two-and three-dimensions, *J. Comput. Phys.* 188 (2003) 543–572.
- [28] C.K.W. Tam, J.C. Hardin (Eds.), *Second Computational Aeroacoustics (CAA) Workshop on Benchmark Problems*, 1997.
- [29] J.B. Freund, Proposed inflow/outflow boundary condition for direct computation of aerodynamic sound, *AIAA J.* 35 (1997) 740–742.
- [30] P. Beaudan, P. Moin, Numerical Experiments on the Flow Past a Circular Cylinder at Sub-Critical Reynolds Number, Technical Report, DTIC Document, 1994.
- [31] H.M. Blackburn, R.D. Henderson, A study of two-dimensional flow past an oscillating cylinder, *J. Fluid Mech.* 385 (1999) 255–286.
- [32] C. Norberg, Fluctuating lift on a circular cylinder: review and new measurements, *J. Fluids Struct.* 17 (2003) 57–96.