Journal of Computational Physics 371 (2018) 651–662

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Bandwidth-based mesh adaptation in multiple dimensions

Elliott S. Wise *, Ben T. Cox, Bradley E. Treeby

Department of Medical Physics and Biomedical Engineering, University College London, Gower Street, London, WC1E 6BT, United Kingdom

ARTICLE INFO

Article history:
Received 13 September 2017
Received in revised form 16 February 2018
Accepted 1 June 2018
Available online xxxx

Keywords:
Adaptive moving mesh method
Pseudospectral method
Bandwidth

ABSTRACT

Spectral methods are becoming increasingly prevalent in solving time-varying partial differential equations due to their fast convergence properties. However, they typically use regular computational meshes that do not account for spatially varying resolution requirements. This can significantly increase the overall grid density when resolution requirements vary sharply over the modelled domain. Moving mesh methods offer a remedy for this, by allowing the position of mesh nodes to adapt to the simulated model solution. In this paper, a mesh specification is presented that is based on a local measure of the spatial bandwidth of the model solution. This addresses the rate of decay of the model solution's frequency components by producing high-sampling rates when this decay is slow. The spatial bandwidth is computed using a combination of the original solution and its Riesz transformed counterparts. It is then integrated into a Fourier spectral moving mesh method, using the parabolic Monge–Ampère equation for mesh control. This method is used to solve a multidimensional version of the viscous Burgers equation, and a heterogeneous advection equation. The performance of bandwidth-based mesh adaptation is compared with arclength- and curvature-based adaptation, and against a static mesh. These numerical experiments show that the bandwidth-based approach produces superior convergence rates, and hence requires fewer mesh nodes for a given level of solution accuracy.

1. Introduction

Spectral methods are increasingly being used for the numerical solution of differential equations [1]. To do this, they use globally-defined basis functions and, when compared with low-order, local methods, often produce high accuracy with relatively coarse spatial discretisations [2]. The sampling rate for spectral methods is related to the rate of decay of the solution's frequency components: functions that vary rapidly require denser sampling than those that vary more smoothly [1]. To meet these sampling requirements, spectral meshes typically used fixed, standardised meshes whose spacing is chosen to ensure that some maximum frequency component is supported. For example, the Fourier collocation method is usually applied using equispaced meshes whose spacing ensures at least two points per minimum wavelength of interest—a choice arising from the Shannon–Nyquist sampling theorem. Any frequency components beyond this will not be supported by the mesh. For many problems of interest, spatial resolution requirements are not uniform throughout the simulated domain. For example, in [3–5] a Fourier collocation method was used to simulate high-intensity focussed ultrasound fields. These contain tightly localised shock fronts that require dense computational meshes, but most of the field only has power at low

* Corresponding author.
E-mail addresses: elliott.wise.14@ucl.ac.uk (E.S. Wise), b.cox@ucl.ac.uk (B.T. Cox), b.treeby@ucl.ac.uk (B.E. Treeby).

https://doi.org/10.1016/j.jcp.2018.06.009
0021-9991/

frequencies. As the meshes used in those works were uniform, the strict shock front sampling requirement was applied everywhere, and large-scale, high-performance computing resources were needed to store and process field variables at each time-step.

To accommodate varying resolution requirements, spectral methods can be implemented within a moving mesh method framework [6]. These dynamically adapt mesh node positions throughout a simulation in a solution-dependent manner. This allows both temporally- and spatially-varying resolution requirements to be met. Monitor functions are used to link the model solution to the mesh, and hence guide mesh adaptation. The choice of monitor function is critical to the performance of moving mesh methods. There have been a number of past multidimensional spectral moving mesh methods, including Fourier, Galerkin, and Chebyshev types [7–11]. These all used arclength-like monitor functions that cluster mesh nodes where the gradient of the model solution is large. In all cases, these methods were applied to problems whose solutions were characterised by steep fronts, and so it is unsurprising that gradient-based mesh adaptation was effective. However, many problems include features for which gradient-based adaptation is not an obvious choice and, additionally, gradient-based mesh adaptation has not been theoretically justified in the context of spectral methods.

A mesh adaptation approach that is tailored to spectral collocation methods was presented in [12] and applied with success to one-dimensional problems. It used a high-pass filter to find regions with large high-frequency solution components, and increased the mesh node density there accordingly. A weakness of this method is that the high-pass filtering step requires parameter choices that are problem- and interpolant-specific. Following this, the (spatially) local bandwidth was presented as a parameterless and robust approach to frequency-based mesh adaptation, and applied to a variety of one-dimensional acoustics problems [13]. In that work, the bandwidth measured the local rate of decay of the solution's frequency components, and the sampling density was chosen to be proportional to its reciprocal. When compared with arclength-based mesh adaptation, the bandwidth-based approach considerably improved the convergence rates of Chebyshev, Fourier, and even finite-difference methods. However, the algorithm presented in that work is limited to one-dimensional problems, as it uses the analytic signal to decouple the spatial phase and amplitude of the model variable.

This paper introduces a multidimensional bandwidth-based mesh adaptation method. It works by first decoupling the spatial phase and amplitude of the model solution using the *monogenic signal* [14,15]. From this, the local bandwidth of the solution is computed and used as a specification for mesh adaptation. This specification is integrated into a Fourier spectral moving mesh method, and assessed against arclength- and curvature-based mesh adaptation. To do so, a multidimensional viscous Burgers equation and a heterogeneous advection equation are used to simulate the formation and propagation of a shock front and a sharp peak.

2. Bandwidth-based mesh adaptation

2.1. Multidimensional local bandwidth

To optimally sample a function using a nonuniform grid, the (spatially-varying) spatial frequency content of that function can be used. Specifically, for a band-limited function the local sampling rate should reflect the maximum spatial frequency present at that point [16]. For functions that are not band-limited, the local sampling rate should similarly reflect the rate at which local spatial frequencies decay. In this regard, the local bandwidth has been shown to be an effective measure of this decay for one-dimensional mesh adaptation [13].

To perform local, multidimensional spatial frequency analysis, it is useful to consider the monogenic signal [14]. This is a multidimensional generalisation of the analytic signal, which contains the original signal as one component and a quadrature signal as the other. By augmenting the original signal, the monogenic signal decouples the local amplitude and phase, making local frequency analysis more straightforward. Given a d-dimensional scalar field u, the monogenic signal can be written as a vector field $\mathbf{v}$ with $d+1$ components consisting of the original scalar field and its Riesz-transformed counterparts

$$\mathbf{v} = \begin{pmatrix} u & \mathcal{R}_1 u & \cdots & \mathcal{R}_d u \end{pmatrix}^T .$$

Here, the Riesz transform $\mathcal{R}_j$ is defined in Fourier-space by

$$\mathcal{R}_j u = \mathcal{F}^{-1}\left\{-\frac{ik_j}{\|\mathbf{k}\|}\mathcal{F}\{u\}\right\}, \tag{1}$$

where i is the imaginary unit, $\mathcal{F}$ is the Fourier transform, $\mathbf{k}$ is a vector-field of wavenumbers corresponding to $\mathbf{x}$, and j indicates the coordinate axis. Now let $\mathbf{V}(\mathbf{k})$ be the Fourier transform of $\mathbf{v}(\mathbf{x})$, and assume without loss of generality that

$$\int_{\mathbb{R}^d} \|\mathbf{V}\| \, d\mathbf{k} = 1.$$

Then, the square of the global spatial bandwidth B_j aligned with coordinate axis j is defined as the expected value of k_j^2. That is,