Annals of Nuclear Energy 102 (2017) 134–147

Contents lists available at ScienceDirect

Annals of Nuclear Energy

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

Hybrid Multi-level solvers for discontinuous Galerkin finite element discrete ordinate diffusion synthetic acceleration of radiation transport algorithms

B. O'Malley [a,*], J. Kópházi [a], R.P. Smedley-Stevenson [b], M.D. Eaton [a]

[a] Nuclear Engineering Group, Department of Mechanical Engineering, City and Guilds Building, Imperial College London, Exhibition Road, South Kensington, London SW7 2AZ, United Kingdom
[b] AWE PLC, Aldermaston, Reading, Berkshire RG7 4PR, UK

ARTICLE INFO

Article history:
Received 5 July 2016
Received in revised form 8 November 2016
Accepted 29 November 2016
Available online 28 December 2016

Keywords:
Diffusion
Multilevel solver
Multigrid
Discontinuous galerkin
Finite element method
DG-FEM
Interior penalty
Diffusion synthetic acceleration
DSA

ABSTRACT

This paper examines two established preconditioners which were developed to accelerate the solution of discontinuous Galerkin finite element method (DG-FEM) discretisations of the elliptic neutron diffusion equation. They are each presented here as a potential way to accelerate the solution of the Modified Interior Penalty (MIP) form of the discontinuous diffusion equation, for use as a diffusion synthetic acceleration (DSA) of DG-FEM discretisations of the neutron transport equation. The preconditioners are both two-level schemes, differing in the low-level space utilised. Once projected to the low-level space a selection of algebraic multigrid (AMG) preconditioners are utilised to obtain a further correction step, these are therefore "hybrid" preconditioners. The first preconditioning scheme utilises a continuous piece-wise linear finite element method (FEM) space, while the second uses a discontinuous piece-wise constant space. Both projections are used alongside an element-wise block Jacobi smoother in order to create a symmetric preconditioning scheme which may be used alongside a conjugate gradient algorithm. An eigenvalue analysis reveals that both should aid convergence but the piece-wise constant based method struggles with some of the smoother error modes. Both are applied to a range of problems including some which are strongly heterogeneous. In terms of conjugate gradient (CG) iterations needed to reach convergence and computational time required, both methods perform well. However, the piece-wise linear continuous scheme appears to be the more effective of the two. An analysis of computer memory usage found that the discontinuous piece-wise constant method had the lowest memory requirements.

1. Introduction

The diffusion synthetic acceleration (DSA) method for the source iteration of the S_N equations is a long standing area of research (Adams and Larsen, 2002) and now forms an important constituent of many S_N codes. Developing such schemes for discontinuous Galerkin finite element (DG-FEM) transport discretisations proved particularly difficult. While some successful methods are based on a set of carefully adjusted continuous FEM equations (Wareing et al., 2001) there have also been studies in fully discontinuous DSA formulations (Adams and Martin, 1992) in order to improve consistency or to exploit the opportunities offered by the additional freedom of DG-FEM methods, e.g. their applicability for non-conforming meshes allowing more flexibility in adaptivity. MIP, a fully discontinuous DSA method has recently been developed (Wang and Ragusa, 2010) which is a modification of the Symmetric Interior Penalty (SIP) method (DiPietro and Ern, 2012).

The MIP equations are usually solved by direct application of some algebraic multigrid preconditioning (Kraus and Margenov, 2009). Although this procedure is robust, it is also significantly more expensive than the solution of the DSA schemes based on continuous FEM equations. The higher cost mainly originates from the increased number of degrees of freedom in the discontinuous discretisation leading to more non-zero entries in the matrix. For linear hexahedral elements for example the number of non-zero terms increases by a factor of about 5. In this situation, a significant proportion of the computational time is spent solving the DSA equations.

Such problems could be mitigated by turning to a two-level approach for preconditioning the MIP equations. Recently, such novel matrix solution algorithms have been developed (Dobrev

* Corresponding author.
E-mail address: bo712@ic.ac.uk (B. O'Malley).

http://dx.doi.org/10.1016/j.anucene.2016.11.048
0306-4549/

et al., 2006; Van Slingerland and Vuik, 2012) to improve the solution time of the SIP equations. These algorithms are combined with an algebraic multigrid preconditioning step on the low-level problem in order to create a preconditioner which is a hybrid of two different multi-level methods.

This paper investigates the computational efficiency of two such two-level approaches when preconditioning the MIP equations. In the first method the coarse correction is based on a projection of MIP equations onto the problem domain of continuous finite elements (modified by some impact from the penalty terms on the external boundaries of the problem). In the second method it is projected onto a domain with a single unknown per element. Both methods take advantage of the fact that MIP creates a matrix with a predictable block structure. This enables them to create a low-level scheme which has a good reduction in number of non-zero terms while remaining a good approximation of the problem. An element-wise block smoother is used alongside both methods, which is also well suited to the block structure of the matrix and lends itself well to parallelism. For both methods the low-level correction is obtained using an AMG method, hence why they may be considered to be "hybrid" schemes.

Section 2 briefly describes the MIP diffusion equations and then defines both two-level preconditioners which are to be studied. In Section 3 these preconditioners will then be used to solve a series of problems in order to test their effectiveness, using established AMG methods as a benchmark.

2. Method

2.1. Interior penalty diffusion

The neutron diffusion equation is an approximation of the neutron transport equation in which several approximations are made in order to eliminate the terms involving the neutron current J ($\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$). With scalar neutron flux ϕ ($\mathrm{cm}^{-2}\mathrm{s}^{-1}$), macroscopic removal cross-section Σ_r (cm^{-1}), diffusion coefficient D (cm) and neutron source S ($\mathrm{cm}^{-3}\,\mathrm{s}^{-1}$), the steady state monoenergetic form of the neutron diffusion equation at position $\mathbf{r}$ is given as

$$\nabla \cdot D(\mathbf{r})\nabla\phi(\mathbf{r}) - \Sigma_r(\mathbf{r})\phi(\mathbf{r}) + S(\mathbf{r}) = 0 \tag{1}$$

In the case where it is to be discretised within a discontinuous framework using interior penalty methods, the diffusion equation may be defined using discontinuous Galerkin bilinear and linear forms. Let $(\cdot,\cdot)$ represent an inner product on $\mathscr{L}^2(V)$ for spatial domain V with border ∂V and element edges E, where

$$(a,b)_{\mathscr{L}^2(V)} = \int_V ab\,dV$$

then, the bilinear form $a(\phi,\phi^*)$ and linear form $l(\phi^*)$ are combined to form the variational form $a(\phi,\phi^*) = l(\phi^*)$. Since this paper exclusively deals with problems containing only a single neutron energy, Σ_r may be replaced by the macroscopic absorption cross-section Σ_a. The weak form of the Symmetric Interior Penalty (SIP) equations is then given as (Wang and Ragusa, 2010; DiPietro and Ern, 2012)

$$\begin{aligned} a(\phi,\phi^*) \quad &= (\Sigma_a\phi,\phi^*)_V + (D\nabla\phi,\nabla\phi^*)_V \\ &+(\kappa[[\phi]],[[\phi^*]])_E + (\{[[\phi]]\},\{\{D\nabla\phi^*\}\})_E + (\{\{D\nabla\phi\}\},[[\phi^*]])_E \\ &+(\kappa\phi,\phi^*)_{\partial V} - \tfrac{1}{2}(\phi,D\nabla\phi^*)_{\partial V} - \tfrac{1}{2}(D\nabla\phi,\phi^*)_{\partial V} \end{aligned} \tag{2}$$

where κ represents the penalty term at edge E or on the domain boundary ∂V and

$$l(\phi^*) = (Q_0,\phi^*)_V \tag{3}$$

for source Q_0.

At element boundaries let

$$[[\phi]] = \hat{\mathbf{n}}^+\phi^+ + \hat{\mathbf{n}}^-\phi^- \quad \text{and} \quad \{\{\phi\}\} = (\phi^+ + \phi^-)/2 \tag{4}$$

represent the boundary flux jump and average respectively, where $\hat{\mathbf{n}}$ represents the outward pointing normal vector at each face. The $+$ and $-$ are used to represent either side of the face.

Providing that κ is sufficiently large there will be a stable solution to the MIP equations. At element edges, κ is calculated using the diffusion coefficient D and a length scale h, which is equal to the volume of the element divided by the area of the surface at the edge (element area divided by edge length in 2D elements). For each element τ with set of edges ϵ the property ξ_τ is defined as

$$\xi_\tau = \sum_\epsilon \frac{C_\tau D}{2h_\epsilon} \tag{5}$$

where C_τ is a coefficient from the trace inequalities chosen as

$$C_\tau = \begin{cases} 3 & \text{triangular elements} \\ 4 & \text{quadrilateral elements} \\ 5 & \text{all other elements} \end{cases} \tag{6}$$

these values for C_τ are for linear elements only.

For a given edge ϵ let ξ_ϵ^- and ξ_ϵ^+ represent the values of ξ in the two elements connected by that edge. If the edge lies on a domain boundary ξ_ϵ represent the value of ξ for the element which borders that edge. Then for each edge the formula for κ is

$$\kappa_\epsilon^{\mathrm{SIP}} = \begin{cases} \xi_\epsilon^- + \xi_\epsilon^+ & \text{on interior edges} \\ 2\,\xi_\epsilon & \text{on domain boundary} \end{cases} \tag{7}$$

The modified interior penalty (MIP) method is a similar scheme to SIP. It is used here because it is more consistent with the discretised transport equations in the thick domain where DSA acceleration is important. Therefore it leads to a DSA scheme which more effectively reduces the transport iteration number. In the MIP method, κ is given as

$$\kappa_\epsilon^{\mathrm{MIP}} = \max\left(\kappa_\epsilon^{\mathrm{SIP}}, \frac{1}{4}\right) \tag{8}$$

the SIP and MIP equations are otherwise identical (Wang and Ragusa, 2010). The MIP scheme has been shown to be stable and effective at reducing the spectral radius of iterative transport problems (Turcksin and Ragusa, 2014).

2.2. Two-level preconditioners

This paper presents two preconditioners for accelerating the solution of the MIP equations. Both are two-level variations of multi-level methods, well established preconditioning techniques (Brandt, 1977). In order to define a two-level scheme it is necessary to specify both a smoother and a low-level correction. These work together by splitting the problem so that the smoother and the low-level correction both help to reduce the errors by tackling different error modes.

The choice of smoother has a significant impact on the effectiveness of a multi-level method. When solving a discontinuous diffusion problem the matrix A has a block based structure, with each discontinuous element containing n_e nodes being represented by a $(n_e \times n_e)$ block. This problem structure suggests that a block based smoother such as element-wise block Jacobi should work well. Since all diagonal blocks are independent of each other this smoother is well suited to parallel implementation. Additionally, for a cost of relatively little memory the inverses of the diagonal blocks may be calculated and stored in advance, so that instead of inverting the blocks at every smoothing step it is only done once during the setup.

The low-level correction defines a coarse space Ω^c which should provide a good approximation to the high-level fine space Ω. A