An angular multigrid preconditioner for the radiation transport equation with Fokker–Planck scattering

Lathouwers, Danny; Perkó, Zoltán

DOI

[10.1016/j.cam.2018.10.010](https://doi.org/10.1016/j.cam.2018.10.010)

Publication date

2019

Document Version

Final published version

Published in

Journal of Computational and Applied Mathematics

Citation (APA)

Lathouwers, D., & Perkó, Z. (2019). An angular multigrid preconditioner for the radiation transport equation with Fokker–Planck scattering. *Journal of Computational and Applied Mathematics*, 350, 165-177.
<https://doi.org/10.1016/j.cam.2018.10.010>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' – Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

An angular multigrid preconditioner for the radiation transport equation with Fokker–Planck scattering

Danny Lathouwers^{*}, Zoltán Perkó

Department of Radiation Science and Technology, Delft University of Technology, Mekelweg 15, 2629JB Delft, The Netherlands

ARTICLE INFO

Article history:

Received 26 April 2018

Received in revised form 21 September 2018

Keywords:

Discontinuous Galerkin

Fokker–Planck

Particle transport

Radiation transport

Multigrid

Interior penalty

ABSTRACT

In a previous paper (Hennink and Lathouwers, 2017) we developed a finite element discretization for the Boltzmann transport equation with forward peaked scatter modeled by the Fokker–Planck approximation. The discretization was based on the discontinuous Galerkin method in both space and angle. It was expected and found that the regular source iteration algorithm for the Boltzmann equation is not effective in solving the discretized system and becomes excessively expensive for problems with many angular degrees of freedom. The purpose of this paper is to develop a multigrid scheme as preconditioner for the above mentioned discretization. The method exploits the nested nature of the meshes and the natural prolongation/restriction between meshes by Galerkin projection. A set of test problems ranging from pure spherical diffusion to the complete Boltzmann transport problem in 3D are presented to illustrate that the method is very effective, resulting in iteration counts nearly independent of problem size even for highly non-isotropically refined angular meshes.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

Radiative transport as modeled by the Boltzmann transport equation arises in modeling of nuclear reactors, radiotherapy applications, imaging techniques and astrophysics, among others. The solution of the linear Boltzmann equation (LBE) is computationally expensive due to its high-dimensionality: 3D transient cases correspond to 7 independent variables. Studies of numerical techniques for the LBE mostly focus on the steady state, mono-energetic case which is considered a basic building block for transport methods. The basic discretization of the LBE is usually performed by the discrete ordinates (S_N) method in angle, which is a collocation technique utilizing a discrete set of angles and associated weights to perform numerical integration. An advantage of this technique is that it is easily implemented. A drawback however is that anisotropic directional adaptivity is not possible with standard S_N quadrature sets. A similar issue exists with the spherical harmonic expansion alternative. For spatial discretization one nowadays often uses discontinuous Galerkin (DG) finite element techniques which are robust against large jumps in material properties. Whatever the chosen discretization technique, the solution of the LBE is complex, especially in heterogeneous media.

The most well-known solution technique, source iteration, is effective for thin media but becomes increasingly ineffective in optically thick, scattering materials. It is well-known that the slowly converging mode in thick scattering media exhibits little angular dependence. The so-called diffusion synthetic acceleration (DSA) procedure makes use of this by accelerating exactly these error components through the solution of a scalar diffusion equation. Although this procedure has been applied with great success it is well known to be unstable for highly heterogeneous media for general spatial discretizations due to a

^{*} Corresponding author.

E-mail addresses: d.lathouwers@tudelft.nl (D. Lathouwers), z.perko@tudelft.nl (Z. Perkó).

lack of consistency between the high-order transport and the diffusion discretization. However, stability and efficiency are recovered when using the methods as preconditioner within a Krylov context [1].

DSA is also ineffective for highly anisotropic scatter and for thin media where the nearly isotropic assumption no longer holds. A more generally effective procedure can be obtained by use of a multigrid method. Multigrid methods have been met with great success in essentially all fields of science and engineering [2,3]. The DSA can in fact be viewed as a two-level method operating on the fine transport level and the coarse diffusion level. The essential operating procedure of a multigrid method is to effectively damp error components of different wavelengths by smoothing the error on a hierarchy of grids. Here, we focus on multigrid methods operating in the angular domain. These methods are a potential refinement of the two-level method formed by DSA. Some earlier works have appeared in the area of angular multigrids for the LBE which we briefly review here.

In their pioneering work, Morel and Manteuffel [4] constructed an angular multigrid method for the 1D S_N equations. The scheme is based on formulating a hierarchy of transport problems of $S_{N/2}$, down to the coarsest S_4 set which is in turn accelerated by the P_1 equations. The multigrid method is used as a solver rather than a preconditioner. The 1D formulation was later extended to 2D by Pautz et al. [5] by introducing several modifications to the coarse problems and the filtering of high-frequency modes. Without this filtering the method was found to be unstable.

Oliveira et al. [6] used a hierarchical preconditioning of the even-parity transport equation as discretized by spherical harmonics. It builds upon the hierarchic nature of the P_N expansion. The authors show the relative merits of various multigrid cycling schemes. It appears however that the required number of iterations is still considerable.

Lee [7,8] devised a moment-based multigrid method for the S_N discretized equations based on coarsening in angle combined with selective coarsening in space. His work uses a Krylov smoother on each level and special interpolation procedures that do not require the hierarchic meshes to be nested. This is done through Nystrom interpolation. The results show that high efficiency is obtained for both isotropic as well as for forward-peaked scatter. In his work the method is used directly as a solver rather than as a preconditioner.

Turcksin, Ragusa, and Morel [9] presented a multigrid preconditioning scheme for the S_N equations. Their work combines the use of Galerkin quadrature sets [10] with the multigrid method developed earlier by Pautz, Adams and Morel [5], where the latter is recast as a preconditioner. The last step is crucial as the multigrid scheme by itself was found to exhibit instability for sufficiently forward-peaked scatter. When combined with a Krylov method stability and efficiency is restored. The authors applied the method to various problems in a 2D square domain and found considerable speedup compared to standard source iteration accelerated with diffusion synthetic acceleration.

Recently Drumm and Fan [11] have described a two-level scheme as used in the SCEPTRE neutral and charged particle radiation transport code. Two variations are presented: one where the coarse grid problem is formed by the P_N problem and one where the discrete ordinates (S_N) method is used with a lower angular resolution. Applications are given to practical electron transport problems and speedups up to a factor of 10 are reported.

The above works show that angular multigrid holds great promise and we further investigate it in the present paper.

Most of the work in this area was based on the S_N discretization which has the drawback that anisotropic angular refinement is difficult. Previously we have developed an angular discretization based on the use of discontinuous finite elements on the sphere [12]. These discretizations were found to exhibit very high convergence rates for the scalar flux. In later work [13] we extended the Boltzmann equation to include Fokker–Planck scatter as a model for highly-forward peaked scatter as encountered in charged particle transport modeling. Essentially the Fokker–Planck approximation is the asymptotic limit of the Boltzmann equation taking the average scatter angle to zero and the scatter cross section to infinity while keeping the product constant. The discretization was found to be accurate but – as expected – the numerical solution by the use of standard source iteration was ineffective. In the present work we develop a very effective preconditioner for the Fokker–Planck problem based on a multigrid procedure exploiting the nested nature of the angular mesh hierarchy.

This paper is organized as follows. In Section 2 we review the discontinuous Galerkin discretization in space-angle that we introduced [12] for the linear Boltzmann equation with the Fokker–Planck approximation of the forward-peaked scatter for which we seek acceleration. In Section 3, we review the standard source iteration scheme when used as a preconditioner for a Krylov method. Section 4 presents the multigrid procedure with an explanation of the mesh hierarchy, the interpolation operators transferring the residual and solution corrections, and the smoother which is based on the transport sweep. Results for a pure diffusion problem on the sphere and for the complete 3D LBE are given in Section 5 where we compare the various choices for preconditioning. We draw final conclusions in Section 6.

2. Space-angle discretization of the transport equation

Charged particle transport exhibiting highly anisotropic scatter can be described by the linear Boltzmann transport equation which is discretized in energy by a multigroup approach with either no or linear dependence on E within each group. The focus of this paper is the acceleration of a single group iterative method. Any improvement for the single-group algorithm immediately carries over to the multi-group case. The linear Boltzmann equation without energy dependence reads

$$\Omega \cdot \nabla \phi(\mathbf{r}, \Omega) + \Sigma_t(\mathbf{r})\phi(\mathbf{r}, \Omega) = \int_D \Sigma_s(\mathbf{r}, \mu_0)\phi(\mathbf{r}, \Omega')d\Omega' + Q(\mathbf{r}, \Omega) \quad (1)$$

where $\mathbf{r}$ is the spatial coordinate, $\boldsymbol{\Omega}$ is the unit direction vector, ϕ is the angular flux density, Q is the independent volumetric source density, Σ_t is the total macroscopic cross section and Σ_s macroscopic scatter cross section that depends on the deflection cosine μ_0 . The LBE is augmented with appropriate conditions on the domain boundary, Γ_I ,

$$\phi(\mathbf{r}, \boldsymbol{\Omega}) = \phi_{in}(\mathbf{r}, \boldsymbol{\Omega}), \quad \mathbf{r} \in \Gamma_I, \quad \boldsymbol{\Omega} \cdot \hat{\mathbf{n}} < 0 \quad (2)$$

In the case of charged particle transport, particle scatter can be highly anisotropic with very high scatter cross sections and accompanied by very small deflection angles. One can take the asymptotic limit of the Boltzmann scatter operator where the average scatter cosine, $\bar{\mu}_0$, tends to unity (no deflection) and the frequency of interaction tends toward infinity while the product $\alpha = \Sigma_s(1 - \bar{\mu}_0)$ stays finite. The result of this procedure is

$$\boldsymbol{\Omega} \cdot \nabla \phi(\mathbf{r}, \boldsymbol{\Omega}) + \Sigma_a \phi(\mathbf{r}, \boldsymbol{\Omega}) = \frac{\alpha}{2} \nabla_s^2 \phi(\mathbf{r}, \boldsymbol{\Omega}) + Q(\mathbf{r}, \boldsymbol{\Omega}) \quad (3)$$

where the spherical Laplacian operator is

$$\nabla_s^2 \phi = \frac{\partial}{\partial \mu} (1 - \mu^2) \frac{\partial \phi}{\partial \mu} + \frac{1}{1 - \mu^2} \frac{\partial^2 \phi}{\partial \chi^2} \quad (4)$$

Here μ is the directional cosine of the z-component of $\boldsymbol{\Omega}$ and χ is the azimuthal angle. The angular Laplacian term is the Fokker–Planck term describing small angle scatter. Physically it represents that a mono-directional beam diverges during travel by a continuous scatter process, modeled by a diffusion term on the sphere surface.

2.1. Phase space elements

The spatial domain V is split into elements V_k , where k is the spatial element index. We consider structured and unstructured meshes consisting of triangles, quadrilaterals, tetrahedrons and hexahedrons. A discontinuous solution space $S_{h,p}$, is defined containing polynomials of order p at most. This is a standard approach and we focus attention on the angular discretization.

The construction of angular elements is based on hierarchical sectioning of the sphere into patches, D_κ , where κ is the patch index. The coordinate planes divide the sphere into cardinal octants, which are spherical triangles. We also assign a level, l_κ , to a patch. The spherical triangles at the coarsest level are assigned a level of $l_\kappa = 0$. Spherical triangles with $l_\kappa = 1$ are obtained by halving the edges of the $l_\kappa = 0$ patches and subsequently connecting the emerged points with great circles. Every patch is hereby split into four daughters. This procedure can be repeated to arbitrary depth and locally on the sphere (anisotropic refinement).

The present tessellation of the sphere was chosen over others for several reasons: (i) for 2D cases one can choose to tessellate the top hemisphere only (impossible when basing tessellation on the icosahedron or on a projection of a hexahedron) and (ii) the coarsest (level-0) tessellation is relatively coarse.

The angular subdivision of the sphere can be described by determining a set P of patch indices such that $\cup_{p \in P} D_p \equiv D$ with D the unit sphere surface and $D_p \cap D_q \equiv \emptyset$ for $\forall p, q \in P$. The phase space mesh is then obtained by assigning an angular subdivision P_k to each spatial element V_k . In our scheme different angular subdivisions can be assigned to different spatial elements allowing for flexibility in local refinement at the element level. This ability sets the present work apart from that based on standard or Galerkin discrete ordinate.

2.2. Basis functions on the sphere

Two different sets of basis functions are used for $\boldsymbol{\Omega}$ in this paper. Both sets are local to the patch D_κ by setting $\psi_{[\kappa]\alpha}(\boldsymbol{\Omega}) = 0$ if $\boldsymbol{\Omega} \notin D_\kappa$ and can thus be discontinuous at the patch boundary. Here $\psi_{[\kappa]\alpha}(\boldsymbol{\Omega})$ denotes the α -th basis function on the patch with index κ . The property that the functions are local to the patch ensures that the streaming-removal terms will not couple to non-overlapping patches. As will be explained later, this enables sweeping of the equations. Also, refinement can be handled with considerable ease. The basis functions considered are the following:

1. **Lin** A nodal set of three functions obeying $\psi_{[\kappa]\alpha}(\boldsymbol{\Omega}_\beta) = \delta_{\alpha\beta}$ where $\boldsymbol{\Omega}_\beta$ are the vertices of the patch. These basis functions are constructed by projecting the standard Lagrange functions on the flat triangle onto the sphere. The flat triangle is the triangle on the octahedron that is formed by projecting the vertices of the patch to the octahedron.
2. **Lin- Ω** This nodal set consists of four basis functions that span the linear space $\text{span}\{1, \Omega_x, \Omega_y, \Omega_z\}$. The coefficients of the functions are chosen in such a way that $\psi_{[\kappa]\alpha}(\boldsymbol{\Omega}_\beta) = \delta_{\alpha\beta}$. The $\boldsymbol{\Omega}_\beta$ are the vertices of the patch and the center of the patch which is determined as the renormalized average of the vertices. Note that these are similar to those devised in the work of Jarell [14].

2.3. Discretization

In the following we present a shortened version of the complete derivation of the discretization given in [12,13]. Here we prefer the presentation in terms of physical fluxes (both a physical flux view as well as a tensor math view are presented

in [12]). The flux in each energy group is approximated as a product of spatial and angular basis functions

$$\phi(\mathbf{r}, \boldsymbol{\Omega}) = \sum_{k,i} \sum_{p \in P_k, d} \phi_{[k,p]}^{id} \Phi_{[k]i}(\mathbf{r}) \Psi_{[p]d}(\boldsymbol{\Omega}) \quad (5)$$

where $\Phi_{[k]i}(\mathbf{r})$ is the i th spatial basis function of element k and $\Psi_{[p]d}(\boldsymbol{\Omega})$ is the d th angular basis function on patch (angular element) p .

2.3.1. Removal, streaming and source discretization

The Galerkin procedure is applied by substituting the expansion into the transport equation (Equation (3)) and multiplying with spatial and angular test functions $\Phi_{[j]l}(\mathbf{r})$ and $\Psi_{[q]m}(\boldsymbol{\Omega})$ and integrating over the complete phase space. After this we obtain

$$\sum_f \gamma_{[f,j,q]lm} - \sum_{\xi=1}^3 \sum_i \sum_d V_{[j]li\xi} \phi_{[j,q]}^{id} A_{[q,p]md}^{\xi} + \Sigma_a \sum_i \sum_d N_{[j]li} \phi_{[j,q]}^{id} M_{[q]md} = Q_{[j,q]lm}, \quad (6)$$

where the streaming term at the faces reads

$$\gamma_{[f,j,q]lm} = \int_{\partial V_j^f} \Phi_{[j]l}(\mathbf{r}) \sum_{k \in \{j, j'_f\}} \sum_{p \in P_k} \phi_{[k,p]}^{id} \Phi_{[k]i}(\mathbf{r}) A_{[f,q,p]md} d\mathbf{r} \quad (7)$$

and the following shorthand notations have been introduced

$$N_{[j]li} = \int_{V_j} \Phi_{[j]l}(\mathbf{r}) \Phi_{[j]i}(\mathbf{r}) d\mathbf{r} \quad (8)$$

$$V_{[j]li\xi} = \int_{V_j} \nabla_{\xi} \Phi_{[j]l}(\mathbf{r}) \Phi_{[j]i}(\mathbf{r}) d\mathbf{r} \quad (9)$$

$$A_{[q,p]md}^{\xi} = \int_{D_q} \Omega^{\xi} \Psi_{[q]m}(\boldsymbol{\Omega}) \Psi_{[p]d}(\boldsymbol{\Omega}) d\boldsymbol{\Omega} \quad (10)$$

$$A_{[f,q,p]md} = \sum_{\xi=1}^3 \hat{n}_{[f]\xi} A_{[q,p]md}^{\xi} \quad (11)$$

$$A_{[f,q]md} = A_{[f,q,q]md} \quad (12)$$

$$M_{[q,p]md} = \int_{D_q} \Psi_{[q]m}(\boldsymbol{\Omega}) \Psi_{[p]d}(\boldsymbol{\Omega}) d\boldsymbol{\Omega} \quad (13)$$

$$M_{[q]md} = M_{[q,q]md} \quad (14)$$

$$Q_{[j,q]lm} = \int_{V_j} \int_{D_q} \Phi_{[j]l}(\mathbf{r}) \Psi_{[q]m}(\boldsymbol{\Omega}) Q(\mathbf{r}, \boldsymbol{\Omega}) d\boldsymbol{\Omega} d\mathbf{r}. \quad (15)$$

Here, element j has faces indexed by f with its neighbor at face f denoted as j'_f and $\hat{n}_{[f]\xi}$ is the ξ coordinate of the outward normal of face f in element j . One recognizes mass matrices in space (N) and angle (M) where the angular matrices contain mixed levels; volumetric integrals arising from streaming (V), and angular Jacobians arising from streaming (A).

The above discretization is not complete as the surface integrals (Equation (7)) are ambiguous and need further specification by upwinding contributions. Unlike in S_N schemes, the angular components on a patch need to be handled together. Riemann procedures have been introduced by various authors (see e.g. [15]) to separate the surface terms into inward and outward contributions. Outward fluxes are based on flux values at element j while the inward fluxes are based on flux values in element j'_f . Besides having to deal with various basis functions simultaneously requiring Riemann handling, a further complication is that the elements neighboring the face may not have the same angular discretization (because of local refinement), hence $P_j \neq P_{j'_f}$.

In [12] we have given a thorough derivation of the way upwinding needs to be performed for the various cases where the neighbor is either equally refined, coarser or finer. Here we present a brief derivation of the upwinding procedure in terms of numerical fluxes for the equally refined case. We first write the flux as

$$\phi_{[q]}(\mathbf{r}, \boldsymbol{\Omega}) = \underline{\phi}_{[q]}(\mathbf{r}) \cdot \underline{\psi}_{[q]}(\boldsymbol{\Omega}) \quad (16)$$

where $\underline{\phi}_{[q]}(\mathbf{r})$ describes the spatial dependence of the solution's angular coefficients in patch q and $\underline{\psi}_{[q]}$ is the vector containing the angular basis functions on patch q . Multiplying the transport equation with test function $\underline{\psi}_{[q]m}(\boldsymbol{\Omega})$, substitution of the expansion and integrating over patch we obtain (highlighting only the streaming terms)

$$\text{Removal} + \nabla \cdot \underline{\mathbf{A}}_{[q] \rightarrow [q]} \underline{\phi}_{[q]}(\mathbf{r}) = \text{Source} \quad (17)$$

where $\mathbf{A}_{\equiv[q]md} = \int_D \psi_{[q]m}(\Omega) \Omega \psi_{[q]d}(\Omega) d\Omega$ is the angular Jacobian with $p = q$. This is a coupled set of differential equations in space for the angular coefficients $\phi_{\equiv[q]}(\mathbf{r})$. The quantity $\mathbf{A}_{\equiv[q] \equiv[q]}(\mathbf{r})$ is called flux in the CFD literature. To avoid confusion such quantities are called contribution in the present work. Discretizing the system further by multiplying with a spatial test function $\Phi_l(\mathbf{r})$ and performing spatial integration gives the usual volumetric and surface integrals

$$\int_{\partial V} \Phi_l(\mathbf{r}) \hat{\mathbf{n}} \cdot \mathbf{A}_{\equiv[q] \equiv[q]} \phi_{\equiv[q]}(\mathbf{r}) d\mathbf{r} - \int_V \nabla \Phi_l(\mathbf{r}) \cdot \mathbf{A}_{\equiv[q] \equiv[q]} \phi_{\equiv[q]} d\mathbf{r}. \quad (18)$$

Rewriting the term $\hat{\mathbf{n}} \cdot \mathbf{A}_{\equiv[q] \equiv[q]} \phi_{\equiv[q]}(\mathbf{r})$ as the numerical flux $h_{[q]}$ gives

$$\int_{\partial V} \Phi_l(\mathbf{r}) h_{[q]}(\phi_{\equiv[q]}^+(\mathbf{r}), \phi_{\equiv[q]}^-(\mathbf{r})) d\mathbf{r} - \int_V \nabla \Phi_l(\mathbf{r}) \cdot \mathbf{A}_{\equiv[q] \equiv[q]} \phi_{\equiv[q]} d\mathbf{r} \quad (19)$$

The angular Jacobian is projected on the face normal $\mathbf{A}_{\equiv[q]} = \hat{\mathbf{n}} \cdot \mathbf{A}_{\equiv[q]}$ and requires further analysis for upwinding. As stated above, a Riemann procedure is required for separating incoming and outgoing contributions on a face. Riemann procedures are based on eigenvalue decomposition of the angular Jacobian with the signs of the eigenvalues acting to separate incoming from outgoing. Positive eigenvalues reflect outgoing information and negative eigenvalues incoming information. As such the eigenvalue decomposition is a generalization of a scalar equation case to a system. In [12] we have argued based on tensor analysis that because of the non-orthogonality of the basis functions, we in fact need to construct the generalized eigenvalue decomposition that includes the mass matrix as additional metric, hence we have introduced the term generalized Riemann decomposition. If the neighbor element has the same angular discretization (i.e. the patch of interest exists at the neighbor) then this decomposition reads

$$\mathbf{A}_{\equiv[q]} = \mathbf{M}_{\equiv[q]} \mathbf{P}_{\equiv[q]} \mathbf{G}_{\equiv[q]} \mathbf{P}_{\equiv[q]}^{-1}. \quad (20)$$

Here, $\mathbf{M}_{\equiv[q]}$ contains the entries of $M_{[q]md}$, $\mathbf{P}_{\equiv[q]}$ contains the eigenvectors, while $\mathbf{G}_{\equiv[q]}$ contains the eigenvalues. Such decompositions can be conveniently constructed by calling a LAPACK routine. The parts of the angular Jacobian that select the appropriate component of the projection are then defined as

$$\mathbf{A}_{\equiv[q]}^{+/-} = \mathbf{M}_{\equiv[q]} \mathbf{P}_{\equiv[q]} \mathbf{G}_{\equiv[q]}^{+/-} \mathbf{P}_{\equiv[q]}^{-1} \quad (21)$$

where $\mathbf{G}_{\equiv[q]}^+$ and $\mathbf{G}_{\equiv[q]}^-$ are the diagonal matrices formed by the positive and negative eigenvalues respectively by replacing the opposite signed elements with zeros. The numerical flux, $h_{[q]}$ can finally be constructed as

$$h_{[q]}(\phi_{\equiv[q]}^+(r), \phi_{\equiv[q]}^-(r)) = \mathbf{A}_{\equiv[q]}^+ \phi_{\equiv[q]}^+(r) + \mathbf{A}_{\equiv[q]}^- \phi_{\equiv[q]}^-(r) \quad (22)$$

The above formulation carefully selects the in- and outflow contributions on a face through the generalized eigenvalue analysis. In case of an S_N -discretization this formulation leads to standard upwinding. For the formulations where the neighbor patch is either coarser or finer we refer to the original expositions [12,13].

Note that on structured quadrilateral or hexahedral meshes where face normals are oriented along the coordinate axes the decomposition does not need to be performed as it will lead to one-sided upwinding because any patch is always fully contained within an angular octant, making all directions in such a patch to move information uni-directionally. Further note that, the decomposition is only selectively performed, i.e. when bi-directionality is suspected (as indicated by the directionality of the vertices of the patch). On fine angular meshes combined with unstructured spatial meshes, only a very small fraction of all angular elements actually require such a decomposition (see [12] for more details on testing bi-directionality).

2.3.2. Fokker–Planck scatter

Based on our previous work [13], the continuous scatter term $\nabla_S^2 \phi$ is discretized with the symmetric interior penalty method (SIP) [16]. SIP uses the following bilinear form, a^{SIP} , for the angular operator.

$$a^{SIP}(u, v) = \int_D \nabla_S u \cdot \nabla_S v d\Omega - \sum_{f_a} \int_{f_a} [[v]] \{ \nabla_S u \} \cdot \mathbf{n} + [[u]] \{ \nabla_S v \} \cdot \mathbf{n} dl + \sum_{f_a} \int_{f_a} \frac{\eta}{h_{f_a}} [[u]] [[v]] dl \quad (23)$$

where u is the angle-dependent solution and v is the angle dependent test function. The face f_a has two angular element neighbors j_a and j'_a where the normal points from the first to the latter. The jump and averaging operator are defined by

$$[[\cdot]]_{f_a} := \cdot|_{j_a} - \cdot|_{j'_a} \quad \text{and} \quad \{\cdot\}_{f_a} := \frac{1}{2} (\cdot|_{j_a} + \cdot|_{j'_a}) \quad (24)$$

respectively. The final term contains the stabilization that penalizes the solution discontinuity at interfaces and depends on the penalty parameter η and on the characteristic length-scale h_{f_a} . In the present work we set $\eta = 3$ and h_{f_a} is the arc-length of the angular element side. For a discussion of the rationale behind this we refer to [13].

Generalization to the present problem where we include spatial dependence is straightforward as the Fokker–Planck operator has only angular derivatives. The final form is obtained by adding the term

$$\int_{V_k} a^{SIP}(\phi, \Phi_{[ki]}(\mathbf{r})\Psi_{[p]d}(\Omega))d\mathbf{r} \quad (25)$$

to the discrete system (6).

3. Standard solution approach

The discretized transport problem reads

$$L\phi = S\phi + \mathbf{f} \quad (26)$$

where L denotes the streaming matrix operator combined with particle removal, S is the scatter matrix operator and $\mathbf{f}$ the independent source arising from volumetric sources and boundary conditions. We use the non-symmetric preconditioned Krylov method Bicgstab [17] to solve the system. A common choice for the preconditioner is to use L^{-1} . This corresponds to the Krylov-accelerated version of the standard source iteration procedure

$$L\phi^{k+1} = S\phi^k + \mathbf{f} \quad (27)$$

The operator L^{-1} is easily applied in discrete ordinates codes as the ordinates are independent as far as the streaming operator is concerned. With the correct visiting order of the elements a single sweep suffices to invert the system. In the present case of a finite element discretization in angle, such a single sweep is no longer exact due to possible bi-directional dependence between elements (on unstructured meshes). This bi-directionality is caused by the fact that some directions in a patch are incoming with respect to the element of interest whereas others are outgoing for the very same face. We have devised a sweep algorithm to precondition the system as follows: (i) For each octant a sweep direction through the spatial mesh is determined and stored before starting the solution process. These directions coincide with an S_2 ordinate set. (ii) For each octant the sweep direction through the spatial mesh is traversed and all angular elements belonging to the octant are visited and the local discrete system is solved with the most up to date information from neighbors. Hence we perform a block Gauss–Seidel iteration with a very specific ordering of the spatial and angular elements. Splitting L into implicit and explicit parts as $L = L_I + L_E$ the Gauss–Seidel iteration is

$$L_I\phi^{k+1} = (S - L_E)\phi^k + \mathbf{f} \quad (28)$$

and the preconditioner then reads L_I . We have demonstrated in our previous work that the sweep described above through the domain serves as a good preconditioner (see [12] for more details). In cases where scatter is highly anisotropic (as with the Fokker–Planck model) this procedure unfortunately becomes ineffective.

The Krylov accelerated source iteration outlined above is slightly modified for Fokker–Planck scatter to obtain somewhat better convergence characteristics. Just as with the transport operator, we also split the Fokker–Planck scatter operator in an explicit and an implicit part corresponding to a Gauss–Seidel numbering

$$L\phi^{k+1} = S_E\phi^k + S_I\phi^{k+1} + \mathbf{f} \quad (29)$$

This then leads to the following preconditioner: $(L_I - S_I)^{-1}$. This will be referred in the remainder of the paper as the standard sweep preconditioner. Note that this standard sweep preconditioner is just as easy to apply as L_I^{-1} .

4. An angular multigrid preconditioner

As already stated in the introduction the multigrid method operates by smoothing the error on a given mesh and subsequently transferring the residual to a coarser mesh where the problem's lower frequencies error components can be more effectively attenuated. Here we deal with a linear problem. Hence we use the recursive version of the linear multigrid algorithm as shown in Fig. 1. Since the multigrid method has been extensively documented and applied for a wide variety of problems (see e.g. [2]), we only discuss the multigrid components relevant for the present work.

4.1. Hierarchy of angular meshes

The considered angular meshes are obtained from refining a uniformly discretized sphere consisting of 8 spherical triangles (corresponding to octants). The coarsest mesh, T_0 has level 0. Refined meshes are obtained by (locally) refining specific angular elements as explained earlier. The locality of refinement is constrained by restricting angular meshes to possess only up to two-irregularity, i.e. neighboring angular elements are allowed to differ by two levels at most. This constraint is used to bound the cost of a transport sweep (finding all neighbors and include the transport contributions becomes time-consuming at some point). A series of triangulations, $\{T_l\}$ is thus obtained where the maximum angular element level on T_l is $l_k = l$. The coarsest angular mesh only contains angular elements of level $l_k = 0$. The finest angular mesh is denoted as T_L . The maximum occurring element refinement in the problem is also $l_k = L$. An example of a series of angular meshes obtained from local refinement is shown in Fig. 2.


```

Algorithm  $LMG(\phi_l, \mathbf{f}_l, l)$ 
if  $(l = 0)$  then
   $S(\phi_l, \mathbf{f}_l, \nu_{coarse})$ 
else
   $S(\phi_l, \mathbf{f}_l, \nu_{pre})$ 
   $\mathbf{r}_l = \mathbf{f}_l - A_l \phi_l$ 
   $\mathbf{f}_{l-1} = R \mathbf{r}_l$ 
   $\phi_{l-1} = 0$ 
   $LMG(\phi_{l-1}, \mathbf{f}_{l-1}, l-1)$ 
   $\phi_l = \phi_l + P \phi_{l-1}$ 
   $S(\phi_l, \mathbf{f}_l, \nu_{post})$ 
endif
end Algorithm LMG

```

Fig. 1. Recursive linear multigrid (LMG) algorithm using the V-cycle [2].

Fig. 2. Example hierarchy of angular meshes. Mesh T_5 is the original finest mesh. Note that the mesh elements are shown as bisected triangles rather than as spherical triangles. This is due to the plotting.

4.2. Interpolation operators

Prolongation of the solution from the coarse to the fine angular mesh is natural as we are dealing with a discontinuous Galerkin scheme with nested spaces. Prolongation from level l to level $l+1$ is thus carried out by a straightforward Galerkin projection. Due to the nested nature of the spaces, this projection is exact. The restriction operator, R is chosen as the transpose of prolongation, i.e. $R = P^T$.

4.3. Coarse mesh problem

The coarse mesh operator is obtained from direct discretization on the coarser level (DCA). We choose this method as it is most compatible with our matrix-free implementation. See [2] for a detailed discussion of the merits of DCA.

Coarse grid solution is done by performing 10 smoothing steps. Note that the coarsest angular mesh consists of just 4 (in 2D) or 8 (3D) angular elements. It has been verified for the problem of dealing only with diffusion on the sphere (i.e. without spatial transport and particle absorption) that this is sufficient to converge to the exact solution.

4.4. Smoother

The linear system at level l reads

$$L_l \phi_l = S_l \phi_l + \mathbf{f}_l \quad (30)$$

We distinguish two procedures for smoothing (i) the standard block Gauss–Seidel as explained in Section 3 where all elements in the mesh are visited in the specified order and (ii) an alternative method where we relax only elements that have a refinement level l_k that meets a certain criterion. To elaborate the latter method further we write the linear system on T_l as

$$(L_l - S_l) \phi_l = \begin{pmatrix} A_l^l & C_l^{l,k} \\ C_l^{k,l} & A_l^k \end{pmatrix} \begin{pmatrix} \phi_l^l \\ \phi_l^k \end{pmatrix} = \begin{pmatrix} f_l^l \\ f_l^k \end{pmatrix}. \quad (31)$$

Here we have ordered the elements such that all level- $l_k = l$ elements are numbered first. The solution vector ϕ_l is partitioned into a vector ϕ_l^l representing the solution coefficients for elements of level- l on T_l , and a vector ϕ_l^k containing the solution coefficients of level- $l_k = k$ elements with $k < l$. Similarly, the $(L_l - S_l)$ operator is partitioned into the matrices A_l^l and A_l^k representing the bilinear form for the level- l and level- k elements respectively on T_l . The C -matrices contain the coupling between elements having different levels.

Following Kanschat [18,19], the smoother on T_l operates only on unknowns belonging to angular elements of level- $l_k = l$. The solution on the other, coarser, angular elements in T_l is untouched during smoothing. The smoother, H_l on level T_l can thus be written as

$$H_l = \begin{pmatrix} H_l^l & 0 \\ 0 & 0 \end{pmatrix} \quad (32)$$

We effectively perform relaxation on the system

$$A_l^l \phi_l^l = f_l^l - C_l^{l,k} \phi_l^k \quad (33)$$

The smoother in this work is the sweeping procedure highlighted earlier in this paper for the single grid case. During smoothing the order of visiting elements is probably not very important as it is the smoothing property we are interested in. On the coarsest mesh T_0 the same procedure is also followed, but here the order is much more important as the smoother is used as an actual solver.

5. Results

The main reason for deteriorated convergence of the algorithm on single grids is through the diffusion operator on the sphere. We therefore first examine the performance of the multigrid method as preconditioner on a pure spherical diffusion problem (Section 5.1). In the second problem (Section 5.2) we consider various performance aspects for the complete transport equation on a 3D cubic domain, both with a volumetric source as well as with an angular Dirac function on the boundary of the cube combined with an anisotropically refined angular mesh.

5.1. Spherical diffusion (angular Poisson problem)

We use a manufactured solution that depends only on angle to investigate the performance of the multigrid method for the case of spherical diffusion. We implemented this by turning spatial streaming off. Also, to stress the algorithm we use zero absorption ($\Sigma_a = 0$). Note that we still use the non-symmetric Krylov solver Bicgstab for the present test problem even though it is symmetric and could be done more efficiently with a conjugate gradient (CG) solver. We do so in order to be more comparable with the results obtained hereafter including streaming, which is non-symmetric. Furthermore if CG was used, the preconditioner would need to be symmetrized.

We expect that if the algorithm is efficient for this case, it will be also effective for cases that are less dominated by the Fokker–Planck scatter. Our manufactured solution is chosen as

$$\phi_{ms}(\Omega) = 4 + \Omega_x + 2\Omega_y + 3\Omega_x^2 \quad (34)$$

The solution is quadratic in angle, meaning that none of the basis function choices are able to represent this solution exactly. We solve the Fokker–Planck equation with an independent source Q_{ms} , such that

$$\mathcal{L}_{FP} \phi_{ms}(\Omega) = \nabla_\Omega^2 \phi_{ms}(\Omega) = Q_{ms}(\Omega) \quad (35)$$

where the transport cross section, α was taken as 2. Note that the manufactured source, $Q_{ms}(\Omega)$, obeys the compatibility relation

$$\int_D Q_{ms}(\Omega) d\Omega = \mathbf{f}_L = 0 \quad (36)$$

This relation is also enforced on the coarser levels by

$$\mathbf{f}_l = \mathbf{f}_l - \frac{(\mathbf{f}, \mathbf{1})}{(\mathbf{1}, \mathbf{1})} \mathbf{1}, \quad (37)$$

where $\mathbf{1}$ denotes the vector containing 1's as elements (see [3]).

Table 1Iteration counts and timing (between parentheses, in seconds) for isotropically refined angular meshes up to level $l_{\max}$ for the spherical Poisson problem.

$l_{\max}$	Basis	H	$V(1, 1)$	$V(2, 1)$
0	lin	9 (0.007)	–	–
1	lin	14 (0.038)	5 (0.009)	4 (0.008)
2	lin	24 (0.27)	6 (0.03)	5 (0.03)
3	lin	41 (1.86)	7 (0.19)	5 (0.19)
4	lin	82 (14.9)	8 (0.92)	7 (0.96)
5	lin	164 (119)	11 (5.41)	8 (5.09)
6	lin	339 (1119)	15 (27.4)	10 (24.9)
0	lin- Ω	10 (0.009)	–	–
1	lin- Ω	15 (0.047)	6 (0.01)	5 (0.01)
2	lin- Ω	26 (0.31)	7 (0.04)	5 (0.04)
3	lin- Ω	44 (2.23)	8 (0.19)	6 (0.26)
4	lin- Ω	88 (18.1)	10 (1.11)	7 (1.41)
5	lin- Ω	173 (146)	12 (5.96)	9 (6.50)
6	lin- Ω	364 (1358)	15 (29.8)	11 (29.2)

Fig. 3. Angular meshes refined in a circular region corresponding to $l_{\max} = 4$ (left), 5 (middle), 6 (right). The meshes are viewed from the positive Ω_z -direction.

Results for uniformly refined spheres using different choices of preconditioner are shown in Table 1. We compare the standard sweep Gauss–Seidel preconditioner denoted as H to the multigrid $V(\nu_{pre}, \nu_{post})$ -cycle which is shorthand for the LMG($\phi_l, \mathbf{f}_l, l$) algorithm with ν_{pre} and ν_{post} as pre- and postconditioning sweeps, respectively.

In all problems in this paper the iterations are stopped when the L_2 norm of the residual has dropped by 10 orders of magnitude (we start the iteration with the zero solution). It is clear that Gauss–Seidel sweeping is a poor preconditioner as the amount of iterations grows almost linear with the refinement level, $l_{\max}$. The solution time correspondingly grows around a factor 8 per added refinement instead of the optimal factor 4. The V -cycle preconditioners behave much better with very limited increase in the number of iterations with problem size. It is well-known that the block Gauss–Seidel method that is based on clustering of unknowns by element is not optimal in terms of smoothing properties. Hemker [20,21] showed that another choice of blocks leads to superior smoothing for DG methods. These methods are however less natural to implement than the present scheme. The slight growth of the iteration counts is thought to be caused by non-optimal smoothing properties of the present Gauss–Seidel scheme. The timing of the $V(1, 1)$ and $V(2, 1)$ preconditioners shows that sometimes one is more favorable and sometimes the other. This indicates that additional smoothing by changing the pre or post-processing counts would not be beneficial.

To investigate the effect of anisotropic refinement on the efficiency of the preconditioners, we define a region around the pole of the sphere $\Omega_z > 0.95$ that is refined up to a given level. Note that the refinement algorithm restricts the irregularity of the spherical mesh to 2, i.e. refinement level of neighboring elements is allowed to differ by 2 at most. Some meshes with differing maximum levels of refinement are shown in Fig. 3. The number of angular elements obtained for maximum element levels ranging between 0 and 9 is 8, 20, 32, 44, 92, 440, 1316, 4556, 17096 and 66236 respectively.

The same test problem has been run with the single grid and multigrid preconditioners. Here we smooth all angular elements present in the mesh T_l . Results are compiled in Table 2. Again, the classic block Gauss–Seidel method as preconditioner shows strong growth of the iteration counts with the maximum refinement level, $l_{\max}$, and corresponding increases in computational time. The multigrid methods on the other hand show only slight growth, likely caused by non-optimal smoothing associated with element-based blocking. The iteration counts obtained here are similar to those of the work of Kanschat [19] where 2D and 3D planar geometries are investigated.

Following Kanschat [19], in an attempt to reduce costs of a single multigrid cycle we have investigated the effect of restricting the smoother to visit only elements of refinement level l when smoothing on mesh T_l . The results are listed in Table 3 and indicate that convergence rate is deteriorated quite strongly. This is contrary to the results of Kanschat where almost mesh-independent convergence was obtained. Note however that his results were obtained on meshes that were limited to level-1 irregularity and in Cartesian geometry. It is expected that the convergence deteriorates with increasing mesh irregularity due to less effective smoothing.

Table 2

Iteration counts and timing (between parentheses, in seconds) for non-isotropically refined angular meshes up to level $l_{\max}$ for the spherical Poisson problem. For a given level l , smoothing consists of visiting all elements in the mesh T_l .

$l_{\max}$	Basis	H	$V(1, 1)$	$V(2, 1)$
1	lin	13 (0.02)	5 (0.006)	4 (0.006)
2	lin	15 (0.05)	5 (0.01)	4 (0.009)
3	lin	18 (0.07)	5 (0.02)	4 (0.02)
4	lin	28 (0.25)	7 (0.04)	5 (0.03)
5	lin	52 (2.2)	10 (0.31)	7 (0.18)
6	lin	93 (11.2)	13 (1.24)	9 (0.87)
7	lin	174 (72.6)	14 (4.37)	10 (3.70)
8	lin	333 (524)	21 (21.9)	14 (18.7)
9	lin	636 (3917)	25 (101.2)	17 (85.9)
1	lin- Ω	14 (0.03)	5 (0.008)	4 (0.007)
2	lin- Ω	17 (0.06)	6 (0.01)	4 (0.01)
3	lin- Ω	19 (0.09)	6 (0.02)	5 (0.02)
4	lin- Ω	28 (0.29)	7 (0.04)	6 (0.04)
5	lin- Ω	50 (2.3)	10 (0.34)	7 (0.21)
6	lin- Ω	94 (13.1)	12 (1.29)	10 (1.09)
7	lin- Ω	171 (80.5)	17 (5.58)	11 (4.40)
8	lin- Ω	315 (568)	20 (22.5)	16 (23.1)
9	lin- Ω	706 (4970)	27 (121)	17 (95.2)

Table 3

Iteration counts and timing (between parentheses, in seconds) for non-isotropically refined angular meshes up to level $l_{\max}$ for the spherical Poisson problem. For a given level l , smoothing consists of visiting only elements in the mesh T_l that have refinement level l .

$l_{\max}$	Basis	$V(1,1)$	$V(2,1)$
1	lin	7 (0.009)	6 (0.009)
2	lin	7 (0.01)	7 (0.02)
3	lin	8 (0.02)	8 (0.02)
4	lin	11 (0.05)	10 (0.06)
5	lin	13 (0.23)	12 (0.41)
6	lin	17 (1.32)	12 (1.19)
7	lin	20 (5.70)	15 (5.07)
8	lin	26 (26.0)	23 (29.5)
9	lin	32 (123.9)	25 (123.9)
1	lin- Ω	8 (0.01)	7 (0.01)
2	lin- Ω	9 (0.02)	8 (0.02)
3	lin- Ω	10 (0.03)	8 (0.03)
4	lin- Ω	12 (0.06)	12 (0.07)
5	lin- Ω	15 (0.31)	13 (0.50)
6	lin- Ω	17 (1.56)	15 (1.55)
7	lin- Ω	24 (7.38)	18 (6.57)
8	lin- Ω	29 (30.42)	21 (28.5)
9	lin- Ω	36 (173.1)	26 (141.6)

Based on the present results, in the remainder of this paper smoothing will be done on all elements in a mesh T_l . It is concluded that near mesh-independent results are obtained with the multigrid $V(2, 1)$ preconditioner.

5.2. Boltzmann–Fokker–Planck in a cubic domain

The V -cycle preconditioner was successful in the spherical diffusion problem of the previous section and we will extend its use here for the complete Boltzmann–Fokker–Planck problem. The test problem is based on a 2D problem used by Turcksin and Morel [9] where a multigrid scheme for the S_N equation with highly forward peaked scatter was developed and tested. Here we extend the problem to 3D.

The problem consists of a cube of size $5 \times 5 \times 5 \text{ cm}^3$. The geometry is meshed with a structured mesh containing $30 \times 30 \times 30$ hexahedral elements. All meshes in this work have been made using the Gmsh software [22]. Boundary conditions are vacuum on all sides and a uniform, isotropic, source is applied throughout the cube. The transport cross section, α is set to unity, hence the geometry is 5 transport lengths thick. Linear basis functions are used in space (4 spatial degrees of freedom per hexahedron). Angular refinement is isotropic. Note that this problem contains a large number of degrees of freedom. Even at the coarsest angular mesh (8 elements) with a linear basis in angle we have over 2.6×10^6 degrees of freedom increasing by a factor 4 on each subsequent refinement. Results are listed in Table 4. The results show trends similar to the spherical diffusion problem with increasing iteration counts for the block Gauss–Seidel method and near mesh-independent counts for the V -cycle methods. It is clear that the use of multigrid as preconditioner leads to great savings in resources for the complete Boltzmann–Fokker–Planck problem as well. Contrary to our work, in the work of

Table 4

Iteration counts and timing (between parentheses, in minutes) for isotropically refined angular meshes up to level l_{max} for the cubic domain with a volumetric source on a hexahedral mesh.

l_{max}	Basis	H	$V(1, 1)$	$V(2, 1)$
0	lin	17 (15)	–	–
1	lin	42 (104)	6 (55)	5 (55)
2	lin	126 (848)	7 (152)	6 (154)
3	lin	309 (6397)	7 (438)	6 (423)
0	lin- Ω	18 (19)	–	–
1	lin- Ω	43 (126)	7 (72)	6 (70)
2	lin- Ω	129 (1064)	7 (186)	6 (182)
3	lin- Ω	320 (8094)	9 (623)	6 (548)

Table 5

Iteration counts and timing (between parentheses, in minutes) for isotropically refined angular meshes up to level l_{max} for the cubic domain with a volumetric source on a tetrahedral mesh.

l_{max}	Basis	H	$V(1, 1)$	$V(2, 1)$
0	lin	22 (16)	–	–
1	lin	47 (91)	6 (48)	5 (44)
2	lin	126 (672)	7 (127)	6 (122)
3	lin	291 (4702)	8 (402)	6 (376)
0	lin- Ω	23 (19)	–	–
1	lin- Ω	51 (118)	7 (65)	6 (57)
2	lin- Ω	132 (851)	7 (153)	6 (148)
3	lin- Ω	295 (5867)	8 (477)	7 (488)

Turcksin and Morel [9], the iteration counts show significant increase with the size of the mesh. The main difference from our work is that they use the discrete ordinates method with a Galerkin quadrature set. The fact that such a set is not nested is likely to be responsible for the observed differences. Furthermore, in the present DG framework restriction and prolongation are defined naturally, whereas in their work these are affected indirectly by the flux moments.

To investigate the effect of the choice of spatial mesh on these results, the same test case has been done with an unstructured tetrahedral mesh of similar size (approximately 26000 elements). An effect might be expected as the spatial transport in the DG method depends strongly on the orientation of the element faces. For a hexahedral mesh the upwind side of an element face is clearly defined. As a consequence, in case of no scatter a single well chosen sweep through the domain is enough to solve the system exactly. In case of arbitrarily oriented faces such as in an unstructured tetrahedral mesh, so-called bi-directional faces occur where such a sweep direction does not exist. Even when scatter is absent, iteration is still required. Therefore, in a tetrahedral mesh two reasons for iteration occur simultaneously (scatter and bi-directionality). More details of the spatial transport can be found in [12]. The results of iteration counts and timing are listed in Table 5. The results indicate that the iteration counts are very similar to those obtained in the hexahedral mesh. It thus appears that the convergence is dominated by the Fokker–Planck scatter and is not adversely affected further by the bi-directionality.

5.3. A cubic domain with a boundary source

Following the work of Turcksin and Morel [9] we also investigate a problem with a boundary source, consisting of a uni-directional beam in the z -direction in part of the cube surface

$$\phi(\mathbf{r}, \Omega) = \delta(\Omega - \Omega_z), \quad z = 0, 2 < x, y < 3 \quad (38)$$

In their work, to approximate the boundary condition, all ordinate intensities are set to zero for incoming directions except for the one closest to the z -direction. In the present work the actual Dirac delta is used. In the DG finite element approach this requires integrals of the form

$$\int_p \Psi_{[p]\alpha}(\Omega) \delta(\Omega - \Omega_z) d\Omega = \Psi_{[p]\alpha}(\Omega_z). \quad (39)$$

To adequately capture the peaked nature of the boundary condition we use a non-isotropically refined angular mesh that concentrates effort around the z -pole of the sphere surface. More specifically, for a given maximum level $l_{max} = L$ in the angular mesh, elements with $\Omega_z > 0.97$ are assigned this refinement level, elements with $0.9 < \Omega_z < 0.97$ are assigned level $l_{max} - 1$, elements with $0.8 < \Omega_z < 0.9$ are assigned level $l_{max} - 2$, elements with $0.6 < \Omega_z < 0.8$ are assigned level $l_{max} - 3$, and all other elements are assigned level $l_{max} - 4$. Specifically, the meshes have respectively 20, 32, 44, 128, 440, 1388 angular elements for l_{max} ranging from 1 to 6. To give an idea of the magnitude of these calculations: for $l_{max} = 6$ combined with linear-omega basis functions, we have around 5.8×10^8 degrees of freedom. Some example meshes with $l_{max} = 2, 4$ and 6 are shown Fig. 4. The ability to non-isotropically and non-uniformly refine the angular meshes is one of the strengths of the present method and opens up possibilities not offered by standard or Galerkin discrete ordinate quadratures.

Fig. 4. Non-isotropically refined angular meshes corresponding to $l_{\max} = 2$ (left), 4 (middle), 6 (right). The meshes are viewed from the positive Ω_z -direction.

Table 6

Iteration counts and timing (between parentheses, in minutes) for non-isotropically refined angular meshes up to level $l_{\max}$ for the cubic domain with a boundary source on a hexahedral mesh. For a given level l , smoothing consists of visiting all elements in the mesh T_l .

$l_{\max}$	Basis	H	$V(1, 1)$	$V(2, 1)$
1	lin	36 (60)	5 (51)	5 (40)
2	lin	68 (159)	6 (67)	4 (55)
3	lin	95 (285)	6 (98)	5 (90)
4	lin	234 (1420)	8 (240)	6 (224)
5	lin	512 (8491)	10 (796)	8 (761)
6	lin	(–)	13 (2607)	10 (2615)
1	lin- Ω	36 (70)	7 (61)	5 (50)
2	lin- Ω	73 (195)	7 (92)	5 (80)
3	lin- Ω	98 (342)	8 (150)	6 (141)
4	lin- Ω	247 (1815)	10 (337)	8 (325)
5	lin- Ω	537 (10710)	11 (1061)	9 (1092)
6	lin- Ω	(–)	16 (3787)	11 (3467)

The results are shown in Table 6. Even for these highly non-isotropic meshes, the convergence of the Krylov method when combined with a $V(2, 1)$ cycle is nearly independent of the number of refinement levels, whereas for the single grid sweep preconditioner the solution is always more expensive, even on the coarsest levels.

In the Poisson problem on the sphere surface, we concluded that the smoother where only angular elements with $l_k = l$ are included on mesh T_l , led to deteriorated convergence compared to the smoother where all angular elements on T_l are included. In an attempt to overcome this convergence penalty, we tried an intermediate approach where we incorporated elements of levels $l_k = l, l - 1, l - 2$. This way at least all the nearest neighbors of the level $l_k = l$ elements are included. However, we found that the results did not lead to much advantage in terms of CPU time. We have therefore decided to discard this option.

6. Conclusions

Regular transport sweeps are not effective preconditioners for the Boltzmann transport equation with forward scatter modeled by Fokker–Planck angular diffusion. Diffusion synthetic acceleration is also known not be very effective [9]. To improve on this, angular multigrid preconditioning/solver procedures have been proposed for the discrete ordinate equations. In the present work, we have devised an angular multigrid procedure as preconditioner for the transport equation discretized with discontinuous finite elements in angle. An advantage with respect to discrete ordinates procedures is that restriction and prolongation operators are defined naturally. The coarse grid problem is defined on the 8 basic spherical triangles (octants) of the unit sphere. As smoother on each multigrid level we use a block Gauss–Seidel method differing from a standard transport sweep by incorporation of parts of the angular diffusion operator.

Application to purely spherical diffusion problems (excluding spatial dependence and absorption) showed that our multigrid preconditioning scheme is highly effective resulting in near constant iteration counts independent of the angular refinement. These results were obtained for different angular basis sets.

Similar efficiencies were found for the complete transport equation when tested in a 3D geometry with both volumetric source and a boundary source with a highly non-isotropic angular mesh. The angular multigrid preconditioner using a single V-cycle is found to always outperform the standard sweep preconditioner, and by large margin on finer meshes. The present work makes possible the high-performance solution of Fokker–Planck transport, especially as the multigrid scheme presented is inherently parallelizable.

The present work is based on the Fokker–Planck term for anisotropic scatter. Future work will be focused on combining goal-oriented angular adaptivity [23,24] with the present solution method and on accelerating forward peaked scatter based on both the Boltzmann scatter operator and more general scattering kernels.

References

- [1] J.S. Warsa, T.A. Wareing, J.E. Morel, Krylov iterative methods and the degraded effectiveness of diffusion synthetic acceleration for multidimensional SN calculations in problems with material discontinuities, *Nucl. Sci. Eng.* 147 (2004) 218–248.
- [2] P. Wesseling, *An Introduction to Multigrid Methods*, John Wiley and Sons, 1992.
- [3] W.L. Briggs, V.E. Hensen, S.F. McCormick, *A Multigrid Tutorial*, SIAM, 2000.
- [4] J.E. Morel, T.A. Manteuffel, An angular multigrid acceleration technique for SN equations with highly forward-peaked scattering, *Nucl. Sci. Engng.* 107 (4) (1991) 330–342.
- [5] S. Pautz, J.E. Morel, M.L. Adams, An angular multigrid acceleration method for the SN equations with highly forward-peaked scattering, in: *Proceedings of International Conference on Mathematics and Computation, Reactor Physics and Environmental Analyses in Nuclear Applications*, American Nuclear Society, 1999.
- [6] C.R.E. de Oliveira, C.C. Pain, M. Eaton, Hierarchical Angular preconditioning for the finite element spherical harmonics radiation transport method, in: *Proceedings of PHYSOR 2000, International Topical Meeting on Advances in Reactor Physics and Mathematics and Computation into the Next Millennium*, American Nuclear Society, 2000.
- [7] B. Lee, A novel multigrid method for SN discretizations of the mono-energetic Boltzmann transport equation in the optically thick and thin regimes with anisotropic scatter, part I, *SIAM J. Sci. Comput.* 31 (6) (2010) 4744–4773.
- [8] B. Lee, Improved multiple-coarsening methods for SN discretizations of the Boltzmann equation, *SIAM J. Sci. Comput.* 32 (5) (2010) 2497–2522.
- [9] B. Turcksin, J.C. Ragusa, J.E. Morel, Angular multigrid preconditioner for Krylov-Based solution techniques applied to the SN equations with highly forward-peaked scattering, *Transport Theory Statist. Phys.* 41 (1–2) (2012) 1–22.
- [10] R. Sanchez, J.C. Ragusa, On the construction of Galerkin angular quadratures, *Nucl. Sci. Eng.* 169 (2011) 133–154.
- [11] C.R. Drumm, W.C. Fan, Multilevel acceleration of scattering-source iterations with application to electron transport, *Nucl. Eng. Technol.* 49 (2017) 1114–1124.
- [12] J. Kópházi, D. Lathouwers, A Space-Angle DGfEM approach for the Boltzmann radiation transport equation with local angular refinement, *J. Comput. Phys.* 297 (2015) 637–668.
- [13] A. Hennink, D. Lathouwers, A discontinuous Galerkin method for the mono-energetic Fokker-Planck equation based on a spherical interior penalty formulation, *J. Comput. Appl. Math.* 330 (2018) 253–267.
- [14] J.J. Jarrell, M.L. Adams, Discrete-ordinates quadrature sets based on linear discontinuous finite elements, in: *Proceedings of International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, American Nuclear Society, 2011.
- [15] C.C. Pain, M.D. Eaton, R.P. Smedley-Stevenson, A.J.H. Goddard, M.D. Piggott, C.R.E. de Oliveira, Space-Time streamline upwind Petrov-Galerkin methods for the Boltzmann transport equation, *Comput. Methods Appl. Mech. Engrg.* 195 (2006) 4448–4472.
- [16] D. Arnold, An interior penalty finite element method with discontinuous elements, *SIAM J. Numer. Anal.* (1982) 742–760.
- [17] H.A. Van der Vorst, Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 13 (2) (1992) 631–644.
- [18] J. Gopalakrishnan, G. Kanschat, A multilevel discontinuous Galerkin method, *Numer. Math.* 95 (2003) 527–550.
- [19] G. Kanschat, Multilevel methods for discontinuous Galerkin FEM on locally refined meshes, *Comput. Struct.* 82 (2004) 2437–2445.
- [20] P.W. Hemker, W. Hoffmann, M.H. Van Raalte, Two-Level Fourier analysis of a multigrid approach for discontinuous Galerkin discretizations, *SIAM J. Sci. Comput.* 25 (3) (2003) 1018–1041.
- [21] P.W. Hemker, M.H. Van Raalte, Fourier two-level analysis for higher dimensional discontinuous Galerkin discretisation, *Comput. Vis. Sci.* 7 (2004) 159–172.
- [22] C. Geuzaine, J.F. Remacle, Gmsh: A Three-dimensional finite element mesh generator with built-in pre- and post-processing facilities, *Internat. J. Numer. Methods Engrg.* 79 (11) (2009) 1309–1331.
- [23] D. Lathouwers, Spatially adaptive eigenvalue estimation for the SN equations on unstructured triangular meshes, *Ann. Nucl. Energy* 38 (9) (2011) 1867–1876.
- [24] D. Lathouwers, Goal-Oriented spatial adaptivity for the SN equations on unstructured triangular meshes, *Ann. Nucl. Energy* 38 (6) (2011) 1373–1381.