Angular adaptivity with spherical harmonics for Boltzmann transport

S. Dargaville[a], A.G. Buchan[b], R.P. Smedley-Stevenson[c,a], P.N. Smith[d,a], C.C. Pain[a]

[a]*Applied Modelling and Computation Group, Imperial College London, SW7 2AZ, UK*
[b]*School of Engineering and Material Sciences, Queen Mary University of London, E14 NS, UK*
[c]*AWE, Aldermaston, Reading, RG7 4PR, UK*
[d]*ANSWERS Software Service, Wood PLC, Kimmeridge House, Dorset Green Technology Park, Dorchester, DT2 8ZB, UK*

Abstract

This paper describes an angular adaptivity algorithm for Boltzmann transport applications which uses P_n and filtered P_n expansions, allowing for different expansion orders across space/energy. Our spatial discretisation is specifically designed to use less memory than competing DG schemes and also gives us direct access to the amount of stabilisation applied at each node. For filtered P_n expansions, we then use our adaptive process in combination with the net amount of stabilisation applied to compute a spatially dependent filter strength that does not depend on *a priori* spatial information. This applies heavy filtering only where discontinuities are present, allowing the filtered Pn expansion to retain high-order convergence where possible. Regular and goal-based error metrics are shown and both the adapted P_n and adapted filtered P_n methods show significant reductions in DOFs and runtime. The adapted filtered P_n with our spatially dependent filter shows close to fixed iteration counts and up to high-order is even competitive with P^0 discretisations in problems with heavy advection.

Keywords: Angular adaptivity, Goal based, Spherical harmonics, Filtered, Boltzmann transport

1. Introduction

The Boltzmann transport equation (BTE) is used to model the transport of neutral particles through an interacting medium and can be difficult to solve because of its mixed hyperbolic/parabolic nature and its (up to) seven-dimensional phase space. The large size of this phase space comes from the three spatial dimensions, two angular, one energy/frequency and one time dimensions. As such, many authors who use deterministic techniques to solve the BTE area turning to adaptive techniques to ensure accuracy in their discretisation while reducing the size of the resulting linear system.

Performing adaptivity in the angular domain of the BTE is becoming more popular, as high angular resolution is required in many problems to resolve the streaming terms (i.e., when a particle propogates a long distance without interaction). Previously, in the AMCG we have investigated using several different angular discretisations with adaptivity across a range of different Boltzmann transport applications [1, 2, 3, 4, 5, 6, 7, 8]. Recently [8] we showed scalable angular adaptivity implemented matrix-free with Haar wavelets that demonstrated $O(n)$ scaling in both runtime and memory usage on some test problems. This Haar wavelet discretisation in angle is equivalent to a P^0 FEM discretisation on the sphere, and like all non-rotationally invariant angular discretisations produces ray-effects in the scalar flux. As such, angular adaptivity becomes key, as resolution must be applied anisotropically on the sphere to smooth out ray-effects and to capture different streaming paths across space/energy.

The only rotationally invariant angular discretisation widely used in Boltzmann transport applications is a spherical harmonics discretisation of the sphere, otherwise known as P_n. This discretisation does not suffer from ray-effects, which would make it ideal for applications with heavy streaming, but its global span on the sphere and high-order nature give increasingly ill-conditioned systems. This is due to Gibbs oscillations when trying to capture the pathology

---

of a hat function in angle due to heavy advection (or other discontinuities in angle, caused by material properties, etc). As such, P_n discretisations in angle are traditionally limited to problems with sufficient smoothness.

Over the last several years however, filtered P_n (we refer to these as FP_n) approximations have been introduced [9, 10, 11, 12, 13, 14], which aim to improve conditioning with discontinuities present. This follows the reasoning behind high-order spectral element methods in fields like aeronautics, where the high-order basis functions are filtered near discontinuities while still retaining high-order convergence in smooth regions. This however creates additional difficulties, as we must now decide what type of filter to apply, how much to filter and where a filter should be applied. If we filter too much in smooth areas for example, we destroy the convergence of the method. This leads naturally to making the filter strength spatially dependent. [12, 13] are the only authors to do this, and showed an example in a 2D duct where they *a priori* determined a region to filter heavily to improve their solutions.

This paper aims to combine both filtered P_n discretisations and angular adaptivity in order to produce ray-effect free solutions to the BTE, with good conditioning in problems with heavy streaming. To begin, we produce a novel matrix-free angular adaptivity scheme for P_n and FP_n that uses both regular and goal-based error metrics. We then turn to the development of a spatially dependent filter strength for our FP_n method. Previously we introduced a sub-grid scale FEM discretisation [15, 16, 17, 18, 8] that uses less memory than equivalent DG schemes. This scheme explicitly separates our solution into "fine" and "coarse" scales, with the fine scale representing the amount of stabilisation that is applied.

Rather than use *a priori* information to determine where to filter heavily, we use the net amount of stabilisation applied in each of our adapt steps to scale our filter strength across space. This automatically applies a heavy filter in areas with discontinuities, no matter the cause (e.g., material properties, heavy streaming, boundary conditions, or sources). This combination of spatially dependent filter and adapted FP_n method is a powerful tool for achieving accurate, ray-effect free solutions in problems with varying degrees of smoothness. We show evidence of the success of our spherical harmonics adaptivity scheme on three example problems, with reductions in DOFs, runtime and memory consumption and highlight the benefit of combining this with filtering.

2. Boltzmann Transport Equation

We begin with the BTE and without loss of generality, we write the mono-energetic steady-state BTE in first-order form as

$$\mathbf{\Omega} \cdot \nabla_{\mathbf{r}} \psi(\mathbf{r}, \mathbf{\Omega}) + \Sigma_t \psi(\mathbf{r}, \mathbf{\Omega}) - S(\psi(\mathbf{r}, \mathbf{\Omega})) = S_e(\mathbf{r}, \mathbf{\Omega}), \tag{1}$$

where $\psi(\mathbf{r}, \mathbf{\Omega})$ is the angular flux in direction $\mathbf{\Omega}$, at spatial position $\mathbf{r}$. The macroscopic cross sections define the material that particles are moving through and Σ_t is the total cross-section. The interaction/source terms have been separated into, S, which is dependent on $\psi(\mathbf{r}, \mathbf{\Omega})$, and those which can be considered purely "external", S_e. In this work the only source term is the scattering from angle $\mathbf{\Omega}'$ into angle $\mathbf{\Omega}$ as particles interact with the medium they are propogating in, or

$$S(\psi(\mathbf{r}, \mathbf{\Omega})) = \int_{\mathbf{\Omega}'} \Sigma_s(\mathbf{r}, \mathbf{\Omega}' \to \mathbf{\Omega}) \psi(\mathbf{r}, \mathbf{\Omega}') d\mathbf{\Omega}', \tag{2}$$

where Σ_s is the macroscopic scatter cross-sections. We now move to discussing spherical harmonics in more detail, along with an overview of the previous work on angular adaptivity.

3. Background

3.1. Spherical harmonics

We begin by writing a spherical harmonics expansion up to order N of a function f, on the unit sphere, S^1, as

$$f(\mathbf{\Omega}) = \sum_{l=0}^{N} \sum_{m=-l}^{l} f_{l,m} Y_{l,m}(\mathbf{\Omega}), \tag{3}$$

where $f_{l,m}$ are the expansion coefficients and $Y_{l,m}$ are the real, orthonormal spherical harmonics. If f is smooth, (3) will converge to to f with spectral accuracy as $N \to \infty$. If f is not smooth, Gibbs oscillations will result and the approximation can result in negative solutions, which can be problematic in many applications.

3.2. Filtered spherical harmonics

One of the main motivations [9] considered in formulating a filtered P_n method was to enforce positivity in a spherical harmonics expansion. We write the FP_n method using the generalised notation of [10], where the P_n expansion (3) with filtering becomes

$$f(\mathbf{\Omega}) = \sum_{l=0}^{N} \sum_{m=-l}^{l} \left[\sigma \left(\frac{l}{N+1} \right) \right]^s f_{l,m} Y_{l,m}(\mathbf{\Omega}), \tag{4}$$

where $\sigma(\eta)$ is a filter function that obeys certain properties [10], including having a value of 1 on the isotropic moment, and that $\sigma(\eta)$ only depends on l, making (4) rotationally invariant like (3). A strength parameter, s is also introduced into (4). We note that as $N \to \infty$, or if $s = 0$, (4) reduces to a normal spherical harmonics expansion.

[10] also showed that this formulation is equivalent to including a forward-peaked scattering operator to (1), making the implementation of a filtered P_n method simple. This form simply adds a diagonal term to the discretised system, namely $-\Sigma_f \log(\sigma(l/m))$, where Σ_f is a free parameter that is independent of the mesh spacing (and time step if applicable), can be made spatially dependent and as such can be considered as similar to a material property, which determines the "filter strength" (Σ_f contains s).

Hence we are left with two choices: what filter function, $\sigma(\eta)$, we chose and the value of Σ_f. We chose to use a Lanczos filter exclusively in this work given by $\sigma(\eta) = \sin(\eta)/\eta$, as [10] shows this filter works well. A number of authors have investigated the effect of different filter functions [10, 13, 12, 11]. [11] in particular prove global convergence properties of filtered P_n expansions, and determined as one may expect, that in smooth problems the order of the filter determines the rate of convergence. For non-smooth problems, it is the underlying smoothness of the transport problem that determines the rate of convergence. Given this work is largely focused on transport problems with a lack of smoothness, we do not examine the use of different filters. Finally we must determine the filter strength, Σ_f. To begin we assume Σ_f is constant across the spatial mesh, but in Section 6 we make Σ_f spatially dependent.

3.3. Adaptivity with P_n

Only a handful of authors have investigated using adaptivity with spherical harmonics in Boltzmann transport applications. We do not discuss other adapted angular discretisations, please see [4, 8] for a detailed overview. [19, 20] first used different P_n expansion orders across space in a 1D test problem, where the expansion order was determined *a priori*. [21, 22] then used the even-parity form of the BTE (which is well suited to P_n simulations as neither work in voids) and built a matrix-free combined space/angle adaptive algorithm that used goal-based metrics and showed reductions in total NDOFs applied. [23] performed regular adaptivity with spherical harmonics in 1D, that showed a reduction in NDOFs. [2, 4] used the first order form of the BTE and combined this with regular and goal-based error metrics to produce an adaptive P_n algorithm, showing reductions in NDOFs across a range of 1D, 2D and 3D problems with energy dependence. The only example problem the authors showed with runtimes however, indicated their adaptive scheme did not improve upon the runtimes of an unadapted spherical harmonics discretisation. Finally [24] formed a combined space/angle adaptivity algorithm using double spherical harmonics expansion and also showed a reduction in NDOFs.

This work is focused on applying angular adaptivity to the first order form of the BTE, which is much more difficult to solve than the even-parity form, given the resulting linear system is not SPD. Naturally the aforementioned authors have not investigated the use of P_n expansions in voids given the conditioning problems of high order expansions, but unfortunately, even in the smooth problems presented, many of the authors do not show the runtimes of their adaptive methods. Our previous work [8] focused heavily on the scalability of an angular adaptivity scheme based on Haar wavelets. This work follows on from [8] and as such we discuss the scalability and runtime/memory consumption of our adaptive spherical harmonics scheme.

The only barriers to scalable P_n solutions in general Boltzmann transport problems are the calculation of half-range integrals/application of BCs across faces and the increasing conditioning problems when capturing discontinuities (e.g., problems with streaming) leading to an increase in iteration count. The sparsity of the P_n angular matrices does not increase significantly with increasing angular order, and the typical expression of scattering kernels as Legendre moments makes the source/interaction terms easy to compute in many nuclear/radiative transfer applications. It is these two "scalability" issues this paper aims to tackle.

At best, the calculation of inflow/outflow across a face (or the application of BCs) with spherical harmonics is $O(n^2)$ in angle size. There are two main approaches to computing these terms, including Riemann decomposition using an eigendecomposition [25], or simply applying a quadrature to compute an approximate half-range integral. Both of these methods can benefit from the rotational invariance of spherical harmonics and the sparsity of the spherical harmonics angular matrix in the z direction to reduce either the number of eigendecompositions calculated, or the number of quadrature points needed (for examples and reviews, please see [26, 21, 27]). Even if these calculations could be computed quickly, simply the application of the resulting dense Riemann decomposition on each face is $O(n^2)$. The only feasible strategy (aside from a breakthrough in spherical harmonics rotation algorithms) to reducing the overall runtime is therefore to reduce the size of the angular expansions on each face if possible, through an angular adapt.

The second "scalability" barrier is in problems with heavy transport, and the aforementioned filtered P_n methods offers the hope of improved conditioning in streaming problems. [12] (Fig. 3.5) for example shows the number of GMRES iterations required to solve a 2D duct problem with FP_n as n is increased (up to FP_{17}). They present an almost constant iteration count for certain filter values as the angular order is increased. This combined with the potential reduction in the size of high-order inflow/outflow calculations required due to angular adaptivity offers the potential for fast spherical harmonics calculations in general problems.

4. Spatial discretisation

We use a sub-grid scale spatial discretisation in this work [15, 16, 17, 18], which provides stabilisation while allowing for low-memory use in general applications. A brief overview of the spatial discretisation is given below, for more details please see [18]. We start by decomposing the solution to (1) as $\psi = \phi + \theta$, where ϕ and θ are the solutions on the "coarse" and "fine" scales, respectively. We then represent each of the solutions with a different finite element representation (like [16]); on the coarse scale we use a continuous representation spanned by η_N basis functions, while on the fine scale we use a discontinuous representation spanned by η_Q basis functions, or

$$\phi(\boldsymbol{r}, \boldsymbol{\Omega}) \approx \sum_{i=1}^{\eta_N} N_i(\boldsymbol{r})\tilde{\phi}_i(\boldsymbol{\Omega}); \qquad \theta(\boldsymbol{r}, \boldsymbol{\Omega}) \approx \sum_{i=1}^{\eta_Q} Q_i(\boldsymbol{r})\tilde{\theta}_i(\boldsymbol{\Omega}), \tag{5}$$

where N_i and Q_i are the basis functions for the continuous and discontinuous spaces respectively, with $\tilde{\phi}_i$ and $\tilde{\theta}_i$ being the associated expansion coefficients.

To discretise in angle, we use the P_n and filtered P_n expansions described in detail in Section 3.1, though here we refer to an arbitrary angular discretisation. We can represent the expansion coefficients $\tilde{\phi}_i$ and $\tilde{\theta}_i$ in (5) with an arbitrary angular discretisation, with the basis functions $G_j(\boldsymbol{\Omega})$ and different numbers of basis functions η_A^i and η_D^i for each expansion coefficient in space, on the coarse and fine scales respectively. Finally if we write the expansion coefficients $\tilde{\phi}_{i,j}$ and $\tilde{\theta}_{i,j}$, we have

$$\tilde{\phi}_i(\boldsymbol{\Omega}) \approx \sum_{j=1}^{\eta_A^i} G_j(\boldsymbol{\Omega})\tilde{\phi}_{i,j}; \qquad \tilde{\theta}_i \approx \sum_{j=1}^{\eta_D^i} G_j(\boldsymbol{\Omega})\tilde{\theta}_{i,j}. \tag{6}$$

If we apply the FEM as usual to (1), by integrating and applying Green's theorem, we can recover the linear system

$$\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix} \begin{bmatrix} \boldsymbol{\Phi} \\ \boldsymbol{\Theta} \end{bmatrix} = \begin{bmatrix} \mathbf{S_\Phi} \\ \mathbf{S_\Theta} \end{bmatrix}, \tag{7}$$

or equivalently

$$(\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})\tilde{\boldsymbol{\Phi}} = \mathbf{S_\Phi} - \mathbf{B}\mathbf{D}^{-1}\mathbf{S_\Theta}. \tag{8}$$

where $\mathbf{S_\Phi}$ and $\mathbf{S_\Theta}$ are the discretised source and $\tilde{\boldsymbol{\Phi}}$ and $\tilde{\boldsymbol{\Theta}}$ are vectors containing the coefficients of the coarse and fine discretised solutions, $\tilde{\phi}_{i,j}$ and $\tilde{\theta}_{i,j}$, respectively. Further details of $\mathbf{A}, \mathbf{B}, \mathbf{C}$ and $\mathbf{D}$ are provided in [18], though we should note that (8) has the same number of DOFs as the continuous problem; indeed $\mathbf{A}$ is simply the linear system

that would result from discretising (1) with continuous finite elements. Once we have solved (8), we can reconstruct the fine solution with

$$\mathbf{\Theta} = \mathbf{D}^{-1}(\mathbf{S_\Theta} - \mathbf{C\Phi}), \tag{9}$$

and then form our discrete solution $\mathbf{\Psi} = \mathbf{\Phi} + \mathbf{\Theta}$ (the coarse solution $\mathbf{\Phi}$ can easily be projected onto the fine space when performing this addition).

When we have an adapted P_n/FP_n simulation, we also follow [2, 4] and remove the coupling in $\mathbf{A}, \mathbf{B}, \mathbf{C}$ and $\mathbf{D}$ in (8) (and (9)) between the angular moments at a node and any higher order moments on neighbouring nodes. Without this modification, non-physical discontinuities appear in the scalar flux; see [4] for a more detailed explanation of why this is necessary.

The $\mathbf{BD}^{-1}\mathbf{C}$ term in (8) can be considered as a stabilisation term. To improve the performance of our discretisation, we have made a number of approximations to $\mathbf{D}$ in order to decrease the cost of inverting it (see [18, 28, 29]). These approximations do not affect the conservation of our scheme, as our conservation statement is on the coarse scale of our discretisation.

We begin (as in [18, 28, 30, 8]) by making $\mathbf{D}$ element local, by replacing the DG jump terms with a vacuum condition on all element boundaries. Furthermore, we enforce a block-diagonal form, but note that the block diagonal form is specific to the angular discretisation used. As such we use a different blocking to [8], instead of grouping over the spatial nodes in an element for each angular coefficient, we block together "shells" of angular coefficients across the element. In three dimensions, these shells correspond to four angular coefficients, in two dimensions either three or four angular coefficients are blocked together. Making this block size constant ensures that computing/storing our block diagonal $\mathbf{D}^{-1}$ remains scalable as the angular order increases. This block diagonal form is more expensive to form/store than that shown in [8], but produces stable results when used with P_n/FP_n. For example, if we have an element with three spatial nodes in two dimensions, with angular orders P_1, P_3 and P_5 on each of the nodes respectively, we have a maximum number of 21 angular basis functions present on each spatial node of the element. We group together the angular basis functions into the shells $[1-3, 4-7, 8-10, 11-14, 15-18, 19-21]$ and remove the off-diagonal coupling in $\mathbf{D}$ between shells. The first shell consisting of basis functions 1 to 3 is present on all spatial nodes, hence that block is of size 9×9. The final shell of functions 19-21 however is only present on a single spatial node, giving a block size of 3×3.

Given the large cost of computing the inflow/outflow across a face discussed in Section 3.3, we further approximate $\mathbf{D}$ on internal faces by grouping together face normals across the spatial mesh that are "close" to each other. We consider a base normal $\boldsymbol{n}_1$ on a face and group other normals, say $\boldsymbol{n}_2$, with this base normal if in two dimensions $\boldsymbol{n}_1 \cdot \boldsymbol{n}_2 < \cos(\pi/500)$, and in three dimensions if $\boldsymbol{n}_1 \cdot \boldsymbol{n}_2 < 1 - \pi/10000$. We never apply this grouping on external faces, to ensure we always accurately represent boundary conditions on the domain of the problem. Importantly, if we have adapted angle, the spatial nodes present on a face may not have the same angular resolution (this is somewhat ameliorated by the resolution smoothing discussed in Section 5.2). This is further exacerbated when we group our internal face normals, as there may be many different angular orders that must be represented on a face across a single grouped normal. To calculate our inflow/outflow, we explicitly compute an eigendecomposition across the face (rather than using a quadrature method). We note of course that once we have computed our eigendecomposition, they are no longer hierarchical (that is the Riemann computation across a face at P_{11} is not a subset of that at P_{13}). For large angular order however, they are close and so we simply calculate the Riemann decomposition for the highest angular order present on a grouped face. Given this is not a good approximation at low order, we also chose to group together normals that only include P_1/FP_1 approximations separately. We could extend this to only grouping together normals that are "close" to each other in angular size as well as normal direction, but we found that this does not give substantial benefit.

Given the block-diagonal form of $\mathbf{D}$ described above, we need only store an block-diagonal form of the internal Riemann decompositions across our grouped normals. This storage (as well as that for the block-diagonal $\mathbf{D}^{-1}$) scales linearly with our adapted angle size, meaning we need only compute these decompositions once per adapt step and store the result, as we need to reuse them when computing our goal-based residuals (see Section 5.1.2). Furthermore, given the rotational invariance of the spherical harmonics expansion, we only need to compute an eigendecomposition once per grouped normal to calculate the Riemann decomposition needed for both the inflow and outflow across a face. We can also reuse both these decompositions in the adjoint problem if we are using goal-based error metrics.

These features and modifications to $\mathbf{D}$ combined with our angular adaptivity significantly decreases both the

number and size of eigendecompositions required by our discretisation, and result in a performant spherical harmonics algorithm. We discuss this further in Section 8.

We also scale our element blocks of $\mathbf{D}^{-1}$ by γ ($0 < \gamma < 1$) defined in [31], to prevent locking in pure scatter regions. The benefit to solving (8) as opposed to standard DG formulation is that the static condensation (given the approximations applied to $\mathbf{D}$, our discretisation can be considered as formed from an approximate Schur-complement) allows us to solve for $\tilde{\mathbf{\Phi}}$ and then reconstruct $\tilde{\mathbf{\Psi}}$. Particularly in 3D, the size of $\tilde{\mathbf{\Phi}}$ on the CG mesh is much smaller than $\tilde{\mathbf{\Psi}}$ on the DG mesh. Furthermore, we use linear basis functions in both the continuous and discontinuous spatial expansions given by (5), hence we can often reuse temporary data during our matrix-free matrix-vector product, making our matvecs less expensive in practice.

5. Angular adaptivity

5.1. Error metrics

We consider two forms of angular adaptivity in this work, regular and goal-based adaptivity. We will refer to (1) as the "forward" problem, with exact solution ψ_{exact} and residual $\mathcal{R}$, hence $\mathcal{R}(\psi_{\text{exact}}) = 0$. In this section, we are trying to compute an approximation, $\mathbf{e}$, to the exact error, $\boldsymbol{\epsilon} = \psi_{\text{exact}} - \psi$, in order to guide our adaptivity.

5.1.1. Regular adaptivity

Regular adaptivity is simple with a hierarchical discretisation, as the coefficients in the expansion can be thresholded, with small coefficient guaranteeing small contribution to the norm of the function we are representing; our only job is picking a thresholding tolerance, τ. We therefore define our regular error metric as $\mathbf{e} = \boldsymbol{\epsilon} \approx |\psi|/\tau$. For convenience, in the results presented below, we also scale $\mathbf{e}$ by the maximum scalar flux across the problem; this simply helps make the choice of τ more problem agnostic.

5.1.2. Goal-based adaptivity

Goal-based adaptivity focuses resolution wherever needed to reduce the error in some arbitrary functional. In this section, we briefly review the formulation of goal-based error metrics through a dual-weighted residual method, described by [3, 4] and used by [8]. We can write the goal of the calculation in terms of a functional, F, of the solution as

$$F(\psi) = \int_P f(\psi) \, \mathrm{d}P,$$

where f is an arbitrary function of the solution and P represents the phase-space. Functionals can be easily defined for quantities such as the average flux over a region, current over given surfaces, reaction rates and even eigenvalues. We can approximate the error in our functional as

$$|F(\psi_{\text{exact}}) - F(\psi)| \approx \boldsymbol{\epsilon}^{\mathrm{T}} \mathbf{R}^* \tag{10}$$

or equivalently

$$|F(\psi_{\text{exact}}) - F(\psi)| \approx \boldsymbol{\epsilon}^{*\mathrm{T}} \mathbf{R} \tag{11}$$

where $\boldsymbol{\epsilon}^{\mathrm{T}}$ and $\boldsymbol{\epsilon}^{*\mathrm{T}}$ are the discrete forward and adjoint solution error, respectively, with $\mathbf{R}$ and $\mathbf{R}^*$ the discrete forward and adjoint residuals computed using ψ^* and ψ^*_{exact}, which are the approximate and exact solutions of the adjoint equation with source term derived from the response function respectively.

In order to avoid (10) and (11) both being zero due to Galerkin orthogonality, we further approximate (10) and (11) by computing "reduced-accuracy" discrete residuals $\hat{\mathbf{R}}$ and $\hat{\mathbf{R}}^*$. We must also pick a target error for our goal-based adaptivity, similar to the thresholding tolerance in Section 5.1.1; we denote this tolerance again as τ. We form our approximate error metric for each angular coefficient by computing the pointwise maximum of both the forward and adjoint pointwise errors, and scaling by the target error in each DOF, namely

$$\mathbf{e} = \frac{\max\{|\boldsymbol{\epsilon} \odot \hat{\mathbf{R}}^*|, |\boldsymbol{\epsilon}^* \odot \hat{\mathbf{R}}|\} N_{\text{DOF}}}{\tau}, \tag{12}$$

where $\odot$ denotes pointwise multiplication. The use of the max operator in (12) ensures that features present in both the forward and adjoint solutions are resolved by the adaptivity (we define the orientation of our adjoint angular space to be the opposite of the forward, so we can easily compute products involving both our forward and adjoint wavelet coefficients). We are now left to define both the solution errors ϵ and ϵ^* and the reduced-accuracy residuals.

Similar to the regular adaptivity, we choose $\epsilon \approx |\psi|$ and $\epsilon^* \approx |\psi^*|$. Rather than simply using a diagonal approximation for the reduced-accuracy residuals like [8], which lead to a pathological effectivity index, we compute our reduced accuracy coarse and fine residuals, $\hat{\mathbf{R}}_\Phi$ and $\hat{\mathbf{R}}_\Theta$, respectively, using

$$\begin{bmatrix} \hat{\mathbf{R}}_\Phi \\ \hat{\mathbf{R}}_\Theta \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{A}} & \tilde{\mathbf{B}} \\ \tilde{\mathbf{C}} & \tilde{\mathbf{D}} \end{bmatrix} \begin{bmatrix} \Phi \\ \Theta \end{bmatrix} - \begin{bmatrix} \mathbf{S}_\Phi \\ \mathbf{S}_\Theta \end{bmatrix}. \tag{13}$$

The modified submatrices, $\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}$ and $\tilde{\mathbf{D}}$ (note we still use the approximations to $\mathbf{D}$ described in Section 4) are formed by coarsening the angular resolution (reducing the angular order by 2) in the diagonal blocks corresponding to the multiplication of each spatial basis function by itself. Note this is not the same as forming a residual by coarsening both Φ and Θ. These coarsened residuals can easily be calculated by a single matrix-free matrix-vector product, and will not be zero. These residuals are then combined like the discrete solution in Section 4 to form our reduced accuracy discrete residual, $\hat{\mathbf{R}}$, as

$$\hat{\mathbf{R}} = \hat{\mathbf{R}}_\Phi + \hat{\mathbf{R}}_\Theta \tag{14}$$

We must also take care when using goal-based adaptivity in streaming regions, to ensure that the coarsest angular resolution used produces a response in our goal. If we do not, then the adaptivity algorithm will not refine. This is a problem faced by all goal-based error metrics in the presence of advection, not just in Boltzmann transport problems. In particular, we discussed this previous in [8], where ray-effects caused by a non-rotationally invariant angular discretisation can cause both the forward/adjoint solutions/residuals to be zero in areas that a coarse discretisation cannot "see". Our P_n/FP_n discretisations are rotationally invariant so there are no ray-effects, but insufficient angular resolution can change how far radiation propogates along a streaming path. Section 8 examines this further, but we note in general problems that P_1 may not be suitable as the coarsest angular discretisation used as part of a goal-based adaptivity simulation.

5.2. Adaptivity algorithm

We now consider our iterative algorithm for the angular adaptivity. We begin the first adapt step by first solving the forward linear system with our coarse angular discretisation, then solve the coarse adjoint linear system if goal-based adaptivity is used. We then compute the regular/goal-based error metric and perform refinement/coarsening. This is then followed by further adapt steps, up to some maximum refinement level.

As mentioned, the direction space Ω in our adjoint problem is explicitly written as the negative of our forward angular domain (i.e., our adjoint angular domain is a reflection about the origin). Our error metric (12) ensures that refinement is triggered in areas important to both the forward and adjoint solutions. We do this as it simplifies our implementation, as we can then apply the "same" angular discretisation to our forward and adjoint problems. A disadvantage of this approach is we may be applying more DOFs in angle than by performing adaptivity separately for the forward and adjoint problems. In practice however, we find this is not significant.

Given our spatial discretisation described in Section 4, the error metrics given in Section 5.1 are all computed using Ψ, which is formed from the sum of our coarse and fine solutions. This solution and hence our error metrics are computed on the fine mesh (i.e., the DG mesh), but we perform our angular adaptivity on the CG mesh. We therefore take the maximum error over the DG nodes that share their position with each CG node, to form an error metric, $\tilde{\mathbf{e}}$, on the CG mesh for each angular coefficient. We then take the maximum of the angular coefficients in $\tilde{\mathbf{e}}$ over each CG node. This gives us a single coefficient per CG spatial node i, which we denote as $\tilde{\mathbf{e}}_i^{max}$, which drives our adaptivity. At each CG spatial node, we trigger refinement and increase the order of our expansion by two if $\tilde{\mathbf{e}}_i^{max} \geq 1.0$. We allow the order to increase by four if $\tilde{\mathbf{e}}_i^{max} \geq 2.0$, and allow coarsening by reducing the order by two if $\tilde{\mathbf{e}}_i^{max} < 0.1$.

We then smooth the resulting angular orders at each CG spatial node by averaging the angular order across the mesh connectivity of that CG spatial node. This ensures we have a smooth transition between areas of low and high angular resolution. The expansion order present on the DG nodes of the spatial mesh are then slaved to their CG counterparts and share the same angular discretisation. The choice to adapt on the CG spatial mesh means that

adjacent faces in our mesh share the same angular discretisation. In intermediate adapt steps, to improve our runtimes we reduce the tolerance of our linear solves, as only the final linear solve with the finest discretisation needs to be solved to a high tolerance. For all adapted P_n/FP_n simulations, the linear solves prior to the final step are to a relative tolerance of 1×10^{-4}, with the final step solved to 1×10^{-10}.

6. Spatially dependent filter

As mentioned in Section 3.2, we can allow the filter strength, Σ_f, to be spatially dependent. [13] suggest two different ways to set the value of Σ_f. The first of these is running a calculation with no filter, and then filtering where the scalar flux in the original problem is negative, with the second using a coarse spatial mesh and a high angular order (N_0) to determine an acceptable unfiltered solution, then computing the filter values as $\Sigma_f = \Sigma_t / f(1, N_0)$. Unfortunately neither of these approaches are suitable for problems with heavy streaming, as we found even (unfiltered) P_1 solutions in many problems are too poorly conditioned to solve. As such, we take a different approach and connect the value of our spatially dependent Σ_f to our adaptivity algorithm discussed in Section 5.2.

We start by picking a constant (large) Σ_f across the entire spatial domain for our first, coarse angular adapt step. This ensures that even the FP_1 solution can be easily computed in the first adapt step. After this first step, we note that we have computed the "coarse" and "fine" scale sub-grid components, $\boldsymbol{\Phi}$ and $\boldsymbol{\Theta}$, (and of course our full angular flux, $\boldsymbol{\Psi}$) due to the decomposition performed as part of our spatial discretisation (see Section 4). The fine solution, $\boldsymbol{\Theta}$, represents the amount of stabilisation applied to each angular coefficient at each of the discontinuous spatial nodes. If the solution is smooth (in space/angle), $\boldsymbol{\Theta}$ is small, and if there is a significant discontinuity, $\boldsymbol{\Theta}$ is larger. The calculation of $\boldsymbol{\Theta}$ is also agnostic to the cause of the discontinuity, which is vital.

Given this, it is then easy to determine a spatially dependent filter value by using $\boldsymbol{\Theta}$ to scale Σ_f so it is large in spatial regions where discontinuities are present. We begin by computing Σ_{stab}, the integral of $\boldsymbol{\Theta}$ over angle at each of the discontinuous nodes,

$$\Sigma_{\text{stab}} = \int_{\boldsymbol{\Omega}} \boldsymbol{\Theta} \, d\boldsymbol{\Omega}, \tag{15}$$

which is trivial to compute given the first isotropic moment of a P_n/FP_n expansion. We then take the maximum (in magnitude) of Σ_{stab} across the DG nodes and form $\tilde{\Sigma}_{\text{stab}}$ on the CG nodes. If we denote the constant filter value we set in the first adapt step as Σ_f^1, subsequent spatially dependent filter values are then calculated as

$$\Sigma_f = \Sigma_f^1 \left(\frac{|\tilde{\Sigma}_{\text{stab}}|}{\max(|\tilde{\Sigma}_{\text{stab}}|)} \right)^{(1/3)}. \tag{16}$$

Much in the same way as our adaptive angular order, the DG values of this spatially dependent Σ_f are then simply taken from their equivalent CG values. Equation 16 has the effect of scaling down Σ_f^1 across the spatial grid. The cube root we use in (16) is not necessary, we simply use it to compress the possible filter values, as the magnitude of $\boldsymbol{\Omega}$ can vary by many orders of magnitude. We can also scale Σ_{stab} by the scalar flux at each spatial node, to ensure the filter value remains consistent in problems where the angular flux varies by many orders of magnitudes (we do this in one problem in Section 8). As for chosing a value of Σ_f^1, experimentation suggests a filter value of $\Sigma_f^1 = 1$ is a good starting value in many problems, though (16) often has the effect of reducing the importance of picking a "good" Σ_f^1 (we discuss this in Section 8).

Equation 16 is trivial to compute and naturally means the filter strength changes as we progress through our adaptive process. This is particular important, as we find that on the boundaries of spatial regions where we have adapted (for both P_n and constant filter strength FP_n), if we move from a first order expansion $N = 1$ to any other angular order, we see visible discontinuities in the scalar flux, in some problems (see the discussion in Section 8.2). [4, 3] also found this (e.g., see Figure 2 in [3]), but did not discuss it. We find that this does not seem to occur in the transition region between higher-order expansions. One simple remedy would be to have our coarsest angular discretisation be larger than P_1/FP_1 with a cost of reducing the effectivity of our angular adaptivity. Instead we chose to rely on our spatially dependent filter. As mentioned, the fine scale solution, $\boldsymbol{\Theta}$, stabilises discontinuities, regardless of the cause. As such, we find that the value of $\boldsymbol{\Theta}$ in subsequent adapt steps around these regions is large and so our spatially dependent filter is also large, smoothing out the discontinuities. This is a powerful feature of our adaptive filtered P_n method. It is also easy to see how (16) could also be scaled by the goal-based error metrics

described in Section 5.1.2, allowing very high filter values in unimportant regions of the problem, improving the overall conditioning and decreasing the iteration count in the linear solver. We have found success in doing so, but discussing this in detail is left for future work.

When performing goal-based adaptivity, we also explicitly allow different spatially dependent filter values for our forward and adjoint problems. We could easily combine the values of Σ_{stab} across both both the forward and adjoint problems, but we found in many problems that the difference in magnitude between the forward and adjoint solutions at the same spatial position meant both ended up overfiltered. Given these separate spatially dependent filter values for the forward/adjoint solutions, care should be taken when forming the error metrics (10)–(12). For example in (10), the discrete forward error and discrete adjoint residuals are convolved, though the angular coefficients at corresponding indices represent different filtered P_n expansions. The solution for this would be to project onto a common space and then form the error metric; we however do not do this. For simplicity, we ignore the different filter values and calculate our error metrics as written in Section 5.1.2. We find this is an acceptable solution given our error metrics are at best an approximation; Section 8.2 discusses the impact this has on our effectivity index.

7. Linear solver

We use the same matrix-free method as [8] to solve our linear system, namely FGMRES(30) preconditioned by a matrix-free multigrid method. Both the spatial tables and $\mathbf{D}^{-1}$ on the lower multigrid levels are built by using an agglomerate-local Galerkin projection on the top-grid element matrices. The lower-grid CG and DG spatial nodes take the adapted angles present on their fine equivalent, and the lower-grid CG and DG spatial nodes also use the filter value, Σ_f, present on their fine equivalent. Importantly, the fact that we do not use jump-terms in our discretisation means we do not need to worry about the non-straight element boundaries on the lower grids, formed from agglomerates of unstructured elements (or having to group these non-straight "normals" as described in Section 4). This allows us to perform matrix-free matrix-vector products on all multigrid levels. Our smoothers differ from [8], as the block-grouping of $\mathbf{D}^{-1}$ for P_n/FP_n makes it difficult to assemble a diagonal scalably. As such, we use GMRES(3) preconditioned by the diagonal up to P_7/FP_7, then switch simply to GMRES(3) for higher order discretisations.

8. Results

We use three examples to test our angular adaptivity algorithm, with varying levels of smoothness. We expect the FP_n algorithms presented to be of most benefit in problems with little smoothness. We start the adaptivity algorithm with a coarse uniform resolution of P_1/FP_1 unless otherwise stated and set a maximum level of refinement. Memory use is profiled using massif (from valgrind) which measures peak heap usage. Very little memory in our simulations is not on the heap, so this gives an accurate measurement of our total peak memory use.

All non-P_n solutions shown use are taken from [8] for comparative purposes (using the same mesh, run on the same machine with the same compiler/optimisation flags, etc). These other angular discretisations, along with the adapted P_n and FP_n discretisations presented in this work have been implemented in FETCH2, the multi-physics, coupled Boltzmann transport code developed at the AMCG. Both the P_n and FP_n use high order P_n calculations on the same spatial meshes as reference solutions, as we know the FP_n converges to P_n as $N \to \infty$.

We should note that for all the adapted simulations shown in this Section, the runtime shown includes all the adapt steps (i.e., all the linear solves performed, computation of error metrics, refinement/coarsening, etc) required to get to that order. For example, if we perform a regular adaptive simulation with 5 adapt steps, the runtime shown includes the time required to perform 5 linear solves. An equivalent goal-based simulation would include 9 linear solves; we do not solve the adjoint problem on the final adapt step. All uniform linear solves were performed in serial to an absolute/relative tolerance of 1×10^{-10}.

8.1. Brunner lattice problem

The first example is the lattice problem from [32]. We discretise this problem in space with the same mesh used by [8], namely an unstructured triangular mesh with 3378 elements (1690 CG nodes and 10,134 DG nodes). This problem has regions of smoothness in angle, but still features discontinuities, particularly in the corners between the scattering and absorbing regions. Previously, we showed [8] that uniform P_n performed well in this problem compared to other

angular discretisations. We use regular adaptivity in this problem, as large regions of the phase-space are important to the final solution and we allow a maximum of 10 adapt steps for both P_n and FP_n. The reference solution used is uniform P_{101} with 5253 DOFs in angle, using 60M DOFs. We compute the relative error in the 2-norm of the scalar flux in this problem.

To begin, we examine the impact of the thresholding tolerance used; too small of a tolerance and the adapt process will add unecessary angles, too large and it will not add enough to reach a desired error. Fig. 1 shows the results from modifying the thresholding tolerance from 1×10^{-3} to 1×10^{-5}. We can see in Fig. 1a that a tolerance of 1×10^{-3} and 1×10^{-4} causes the adaptivity to plateau with the number of CDOFs, with 1×10^{-5} producing a solution that matches the uniform P_n until we reach an order of P_{21}. At this point the adapted P_n produces a solution of equivalent accuracy with less CDOFs than the uniform. Fig. 1b shows that the uniform P_n is quicker than the adapted P_n, until we reach uniform P_{41} and P_{51}, when the adapted is quicker. This is because, for low order, we do not save a substantial number of DOFs by adapting in this problem, and the cumulative cost of solving the linear system at each adapt step outweights the cost of the Riemann decompositions. At higher order however, this balance changes.
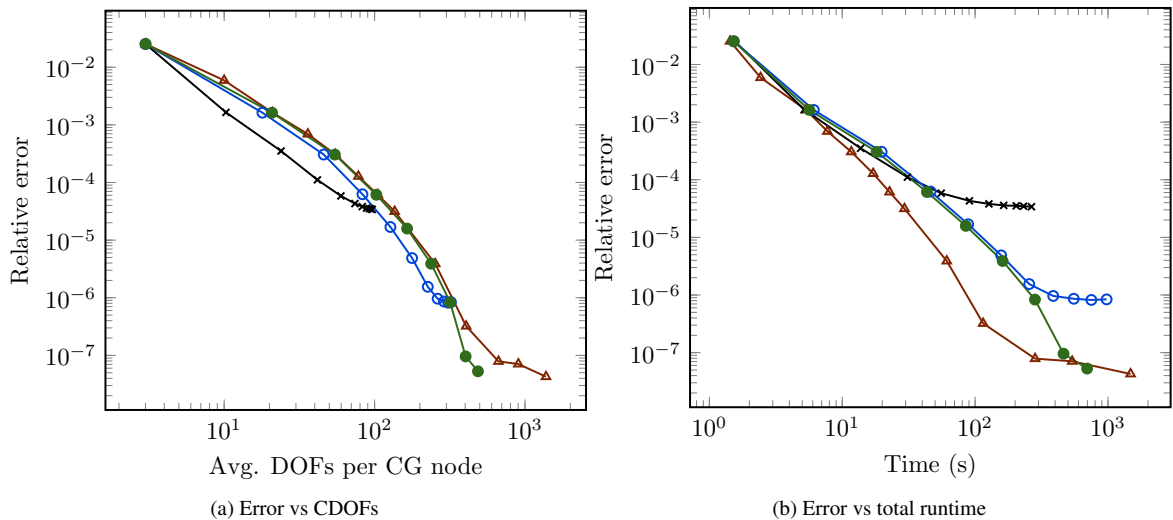

(a) Error vs CDOFs

(b) Error vs total runtime

Figure 1: Performance of the regular angular adaptivity with P_n, in the relative error of the 2-norm of the scalar flux across the domain, for the Brunner problem. The x, o and ● markers use threshold coefficients 1×10^{-3}, 1×10^{-4} and 1×10^{-5}, respectively, with the △ uniform (unadapted).

To begin examining the results from the FP_n method, we fix the filter strength to a constant value and examine the impact of changing the thresholding tolerance, to verify that our FP_n adaptivity still behaves similarly to the P_n. We do not focus heavily on adapted FP_n with a constant filter strength in this paper, for reasons that will become evident below. Fig. 2 shows the impact of setting the filter strength to a constant $\Sigma_f = 1$. We can see in Fig. 2a that even with uniform angular resolution, the filter has significantly degraded the convergence in this problem compared with Fig. 1a, to the point of non-monotonicity. This is to be expected, as this problem has enough smoothness to benefit from the spectral nature of P_n. Examining the impact of the adaptivity and the thresholding tolerance, we see that again, choosing a tolerance that is too large (1×10^{-3}) causes the adaptivity to plateau. We can also see in Fig. 2a that choosing a tolerance that is too small (1×10^{-5}) causes the adaptivity to include too many angles, whereas a tolerance of 1×10^{-4} achieves both a reduction in DOFs for a given error and a decrease in runtime, shown in Fig. 2b.

Fig. 3 visualises the impact of using filtered P_n on the angular flux at a single point in the Brunner problem. Fig. 3a shows that the P_n solution features oscillations across the sphere, with a strong filter of $\Sigma_f = 100$ smoothing the angular flux considerably, and decreasing the magnitude of negativity as expected.

Fig. 4 shows where the adaptivity for both the P_n and FP_n with a strong constant filter has chosen to place angular resolution, and we can see that as expected when using regular adaptivity, the areas of high resolution are focused on areas with large flux. We can see in Fig. 4b that the FP_n has used fewer angles than the P_n shown in Fig. 4a for a given tolerance (this is dependent on the filter strength/type and problem).

(a) Error vs CDOFs
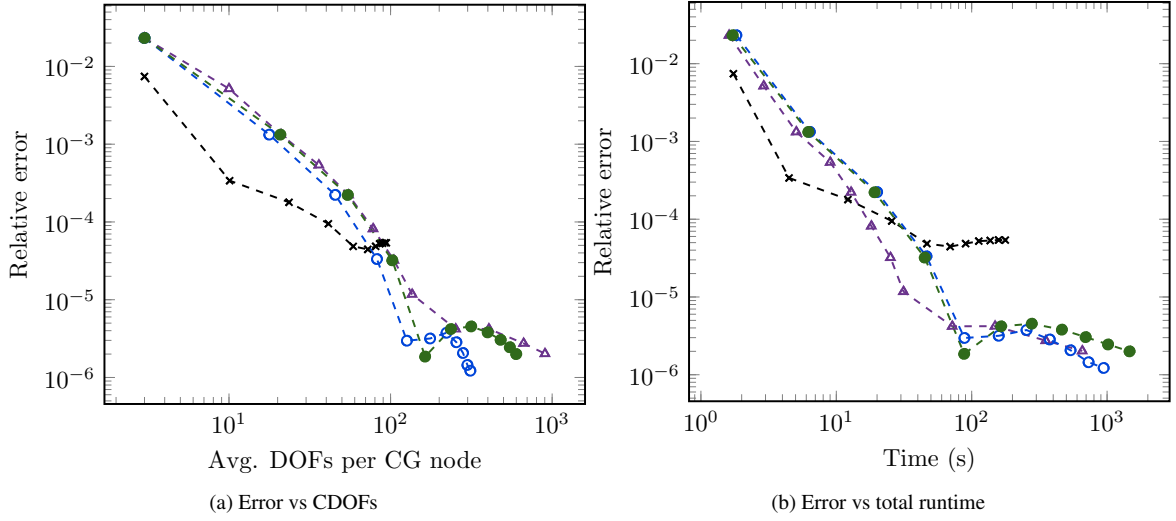
(b) Error vs total runtime

Figure 2: Performance of the regular angular adaptivity with FP$_n$ with $\Sigma_f = 1$, in the relative error of the 2-norm of the scalar flux across the domain, for the Brunner problem. The x, o and ● markers use threshold coefficients 1×10^{-3}, 1×10^{-4} and 1×10^{-5}, respectively, with the △ uniform (unadapted).

We can now turn to examining the behaviour of our FP$_n$ method as we change the filter strength. Fig. 3 showed that a filter degrades convergence in a global norm in this problem, as theoretical results predict [11]. Fig. 5 shows the results from using both uniform FP$_n$ with a fixed constant filter strength, and comparing to an adapted FP$_n$ simulation with spatially dependent filter values. If we examine Fig. 5a we can see that the uniform FP$_n$ all converge worse than the uniform P$_n$ in this problem. As we decrease the size of Σ_f from 100 to 1, the convergence improves, though like in Fig. 3 we see non-monotonic convergence. Fig. 5b also shows starts to show the effect of computing the Riemann decompositions in this problem at high order, as the runtimes of the uniform FP$_n$ method are increasing nonlinearly with the NDOFs.

When we enable our spatially dependent filter strength, we can see in Fig. 5a that for the same initial filter strength Σ_f^1 of between 100 and 1, the convergence is improved significantly compared to when the same constant filter strength is applied. Furthermore, the adapted FP$_n$ method with $\Sigma_f^1 = 1$ is significantly outperforming the uniform P$_n$ method per DOF, reaching an equivalent accuracy at high order with an average of 305 DOFs per CG node, compared to 1378 DOFs. We can see in Fig. 5b that this adapted calculation at high order is also approximately three times faster to compute.

Fig. 6 shows the size of both Σ_{stab} and our spatially dependent Σ_f in this adapted FP$_n$ simulation. We can see in Fig. 6a that our net stabilisation is applied heavily around the material property boundaries in this problem, in greatest magnitude near the source. This is where we would expect discontinuities in space/angle to be large in this problem. Fig. 6b shows the resulting filter values applied in this problem, where the largest value is given by $\Sigma_f^1 = 1$, and scaled down elsewhere in space according to (16). This results in a significantly smaller average filter value applied in this problem, and as noted improves convergence.

Fig. 7 compares the best performing results from the adapted P$_n$ and adapted FP$_n$ with spatially dependent filter from Figures 1 and 5, respectively, with a number of different discretisations taken from [8]. We can see in Fig. 7a that both the adapted P$_n$ and adapted FP$_n$ outperform all other discretisations per DOF. Indeed, in this problem the adapted FP$_n$ even outperforms the adapted P$_n$ method, which is an interesting result. This shows that even in a problem with sufficient smoothness that uniform P$_n$ exhibits good convergence when compared to a first order method like uniform LS P^0 FEM, there can still be a benefit to applying a filter. Both adapted P$_n$ and adapted FP$_n$ achieve similar runtime improvements over the uniform P$_n$ in this problem, as shown in Fig. 7b. This is the impact of the adaptivity reducing the number of large Riemann decompositions that must be performed.

In particular, Table 1 shows that although Fig. 7 suggests that the adapted FP$_n$ with spatially dependent filter is the best performing method in this problem, the cumulative runtime per final DOF is not constant, indicating that the
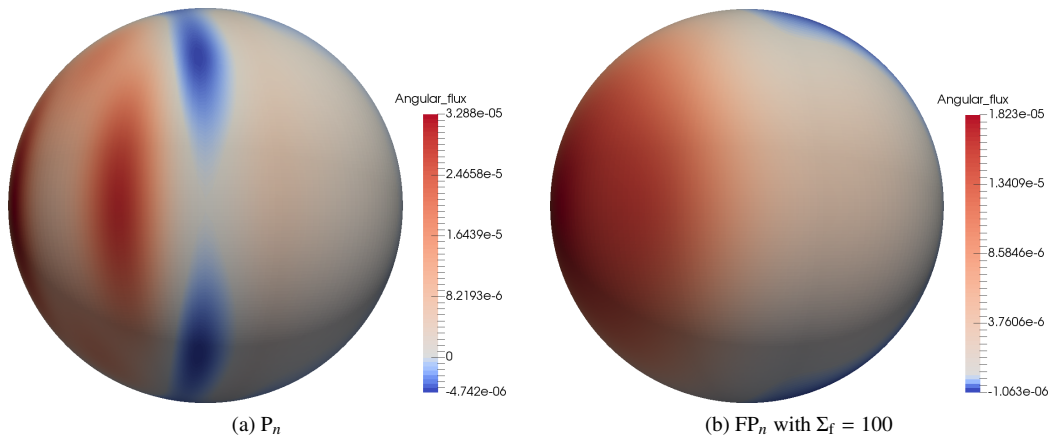
(a) P$_n$ (b) FP$_n$ with $\Sigma_f = 100$

Figure 3: Angular flux in the Brunner problem at $x = 2.5$, $y = 2.5$, with regular angular adaptivity after 5 adapt steps, with threshold coefficient 1×10^{-3}.

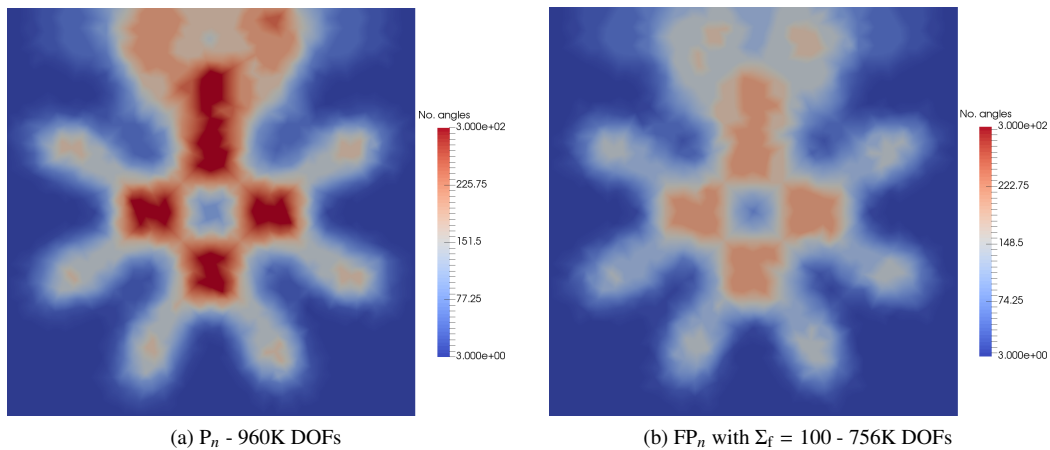(a) P$_n$ - 960K DOFs   (b) FP$_n$ with $\Sigma_f = 100$ - 756K DOFs

Figure 4: Number of basis functions across the spatial domain for the Brunner problem, plotted on the CG mesh, on the 7th step of regular angular adaptivity with threshold coefficient 1×10^{-3}.

overall method cannot be considered scalable (unlike the wavelet adaptivity in [8] in this problem). This is entirely due to the increasing cost of performing/applying the Riemann decompositions, as Table 1 shows that linear solver for our adapted FP$_n$ method actually requires fewer iterations as the maximum angular order is increased. Over half the overall runtime for the 10th adapt step is spent computing Riemann decompositions. The overall time spent computing these is still much less than if we were running a uniform simulation (as shown in Fig. 7b), but it is still the main contributing factor in this problem. Thankfully however, Table 1 shows that the block form of $\mathbf{D}^{-1}$ means our peak memory use remains constant as the angular order is increased, settling down to approximately 37 copies of the angular flux (much like [8] this can be reduced considerably by decreasing our GMRES restart parameter and computing $\mathbf{D}^{-1}$ on the fly with little impact on convergence/runtimes). We now turn to a problem with much larger discontinuities, where we would expect P$_n$ to perform poorly.

8.2. 2D dogleg problem

We now examine the use of spherical harmonics adaptivity in a 2D duct problem [4]. We discretise this problem in space with the same mesh as [8], an unstructured triangular mesh with 2824 elements (1477 CG nodes and 8472 DG nodes). This problem features highly anisotropic flux throughout the duct regions and features heavy streaming region. This problem is highly challenging for high-order methods like P$_n$, with heavy discontinuities in space/angle.
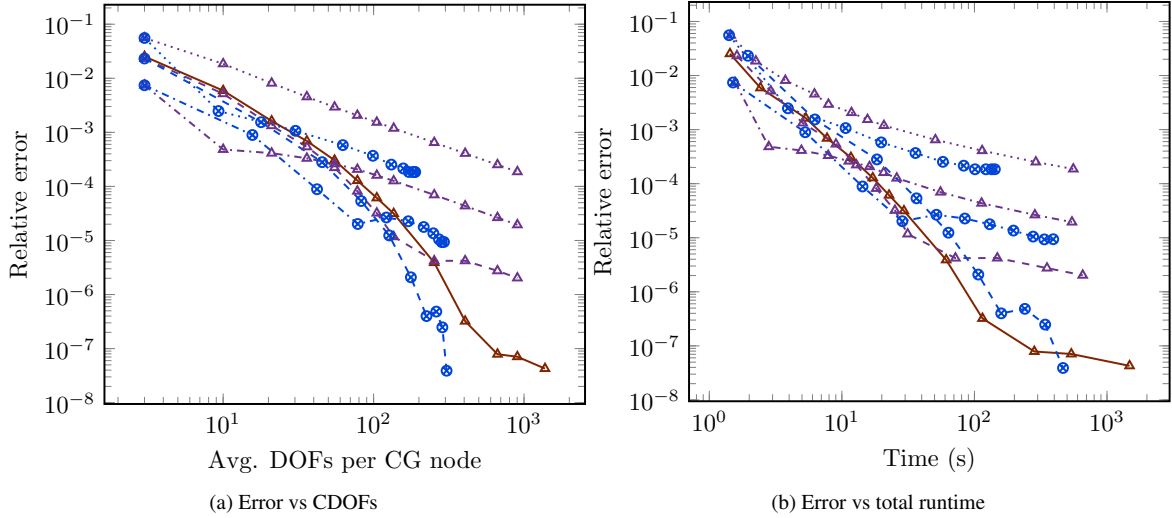
(a) Error vs CDOFs	(b) Error vs total runtime

Figure 5: Investigating the impact of a spatially-dependent filter strength with FP_n, in the relative error of the 2-norm of the scalar flux across the domain, for the Brunner problem. The solid $\triangle$ is uniform P_n, with the dotted $\triangle$ FP_n with $\Sigma_f = 100$, dash-dotted FP_n with $\Sigma_f = 10$ and dashed FP_n with $\Sigma_f = 1$. The $\otimes$ are regular adapted FP_n with threshold tolerance 1×10^{-4}, with spatially dependent filter strength and reduced tolerance solves, with the dotted $\Sigma_f^1 = 100$, the dash-dotted $\Sigma_f^1 = 10$ and dashed $\Sigma_f^1 = 1$

Adapt step (max. order):	1 (1)	2 (5)	3 (9)	4 (13)	5 (17)	6 (21)	7 (25)	8 (29)	9 (33)	10 (37)
Cum. runtime (μs) per final DOFs:	57.2	30.6	35.2	38.5	43.6	52.6	61.7	80.3	104.1	132.9
No. iterations:	12	14	20	20	19	19	19	18	18	17
% runtime Riemann decomp.:	0.1	0.8	1.3	3	6.7	12.7	21.7	33.3	45.3	55.9
Peak memory use:	201.2	61.1	44.2	39.5	37.6	36.7	36.7	37	37.5	37.2

Table 1: Runtime, iteration count, percentage of runtime spent computing Riemann decompositions and peak memory used for the Brunner problem, for the regular FP_n adapt with threshold coefficient 1×10^{-4}, spatially dependent $\Sigma_f^1 = 1$ and reduced tolerance solves. Peak memory use is on the heap (measured by massif) scaled to the size of the angular flux. The runtime is the cumulative runtime of all adapt steps up to that level, scaled by the NDOFs in the final adapt step. The iteration count is from the linear solve at the final adapt step.

We use goal-based adaptivity in this problem, with the goal being the average flux at the end of the duct. The reference solution for the P_n and FP_n is uniform P_{91} with 4278 DOFs in angle, using 42M DOFs.

In this example we move directly to examining the impact of different filter values, the choice of the error tolerance used in this section (1×10^{-1}) was made following a similar process to that shown in Section 8.1 and [8]. Fig. 8a shows that in this problem, uniform FP_n with a constant filter value improves convergence per DOF for low order, but that there is still sufficient smoothness that at high order the uniform P_n method converges better. We can see that overfiltering with $\Sigma_f = 10$ also results in non-monotonic convergence in this problem. One of the key features of Fig. 8b however is that the uniform FP_n with constant filter values significantly outperform the P_n method in runtime until high order. For example, achieving an error of approximately 1×10^{-2} with FP_n with $\Sigma_f = 10$ took only 13 seconds, with the P_n taking around 1000 seconds. This is because the conditioning in this problem is improved considerably by filtering; Table 2 shows that the iteration count for uniform P_{45} is 968, compared with only 41 for uniform FP_{45} with $\Sigma_f = 10$. Table 2 does show that the iteration count for FP_n with $\Sigma_f = 10$ is still growing slowly however as the angular order is increased.

Importantly, Fig. 8 also shows the results from allowing our FP_n method to adapt and use a spatially dependent filter. We can see that this combination reduces both the average NDOFs per CG node and the runtime in this problem by at least two orders of magnitude. As in the previous example problem, allowing the filter to become spatially dependent improves the convergence when compared to the constant filter case. Fig. 9 shows the scalar flux for both the forward and adjoint problem, along with where the goal-based adaptivity has increased the angular order in this
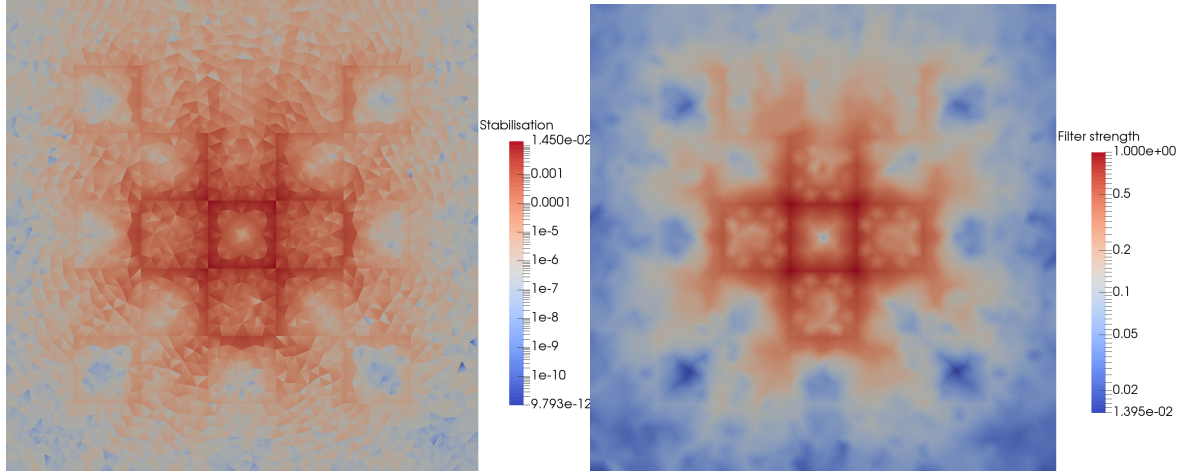
(a) Absolute value of Σ_{stab} across space. This can be considered the net "amount" of stabilisation we apply at each node.

(b) Spatially-dependent Σ_f

Figure 6: Computation of a spatially dependent Σ_f, with $\Sigma_f^1 = 1$, on the the 10th step of FP_n regular adaptivity with threshold coefficient 1×10^{-4} for the Brunner problem.

Order:	1	3	5	7	9	11	13	15	21	27	35	45
Uniform P_n	47	62	95	119	207	244	277	309	420	511	741	968
Uniform FP_n, $\Sigma_f = 0.1$	38	45	63	73	127	141	156	171	225	269	353	482
Uniform FP_n, $\Sigma_f = 10$	9	13	14	16	20	21	22	24	27	31	36	41
Adapt step (max. order):	**1 (1)**	**2 (5)**	**3 (9)**	**4 (13)**	**5 (17)**	**6 (21)**	**7 (25)**	**8 (29)**	**9 (33)**	**10 (37)**	**11 (41)**	**12 (45)**
Goal-based P_n	47	89	165	190	223	274	308	351	400	421	458	-
Goal-based FP_n, $\Sigma_f^1 = 10$	9	32	66	82	82	83	88	85	-	-	-	-

Table 2: Number of iterations for the different uniform/adapted discretisations shown in Fig. 11, for the 2D dogleg problem. For the adapted results, we take the iteration count from the last forward linear solve, which is always performed to the same tolerance as the uniform.

problem. Fig. 9a shows that the highest angular orders have been applied in the streaming path between the source and goal, as would be expected. One important feature to note in the forward/adjoint scalar fluxes shown in Figures 9b and 9c is that the solution is entirely positive, and as mentioned in Section 6, we can see discontinuities where the angular resolution has transitioned from P_1.

Thankfully, as discussed this discontinuity is picked up by Σ_{stab} shown in Figures 10a and 10c in the forward and adjoint solutions respectively. This means our spatially dependent filter is also filtering heavier where this change in angular resolution has caused discontinuities, as shown in Figures 10b and 10d. The other large discontinuities in this problem are present around the sources and in the duct regions, as expected.

We now compare the results from the goal-based adapted filtered P_n shown above with an goal-based adapted P_n method and also with a range of different angular discretisations from [8]. We can see in Fig. 11a adapated P_n in this problem reduces NDOFs by an order of magnitude in this problem when compared to the uniform P_n, and results in a drop in runtime at high order, shown in Fig. 11b. Table 2 shows that the the adapted P_n has resulted in a smaller number of iterations in the linear solver when compared to uniform P_n. The adapted FP_n with spatially dependent filter gives excellent results in this problem, giving a further order of magnitude reduction in DOFs, and due to the improved conditioning, almost two orders of magnitude decrease in runtime when compared to the uniform P_n. Again we can see this reflected in the iteration count in Table 2, and importantly we see a constant iteration count with increasing maximum angular order. This is a powerful result for a filtered spectral method in a streaming problem. Indeed Fig. 11 shows that the goal-based adaptive FP_n with spatially dependent filter is even competitive with the

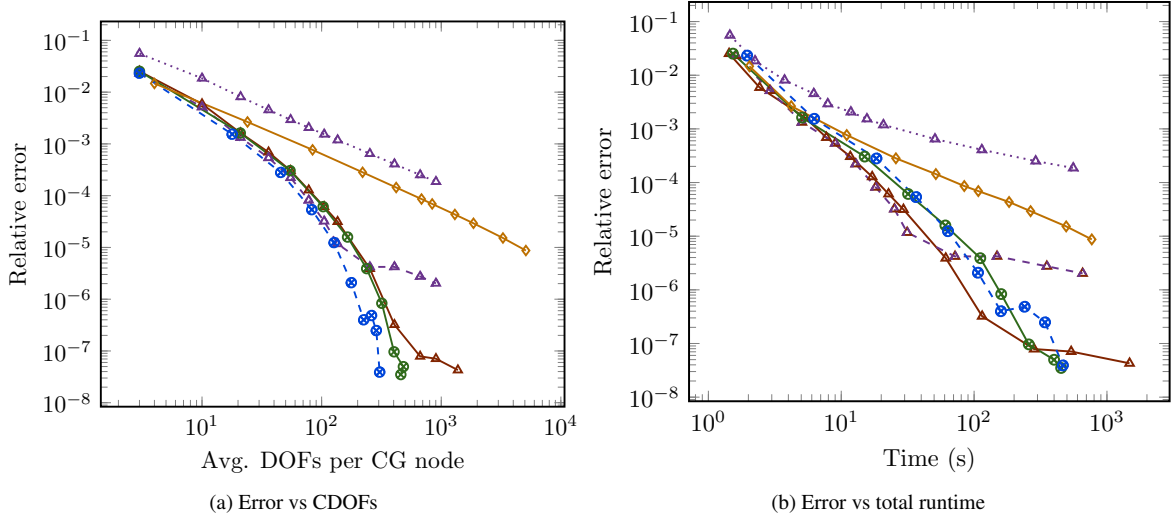| (a) Error vs CDOFs | (b) Error vs total runtime |

Figure 7: Comparison of the relative error of the 2-norm of the scalar flux across the domain, for different angular discretisations, for the Brunner problem. The $\otimes$ are regular P_n adapts with threshold coefficient 1×10^{-5} and reduced tolerance solves, with the dashed $\otimes$ regular FP_n adapts with threshold coefficient 1×10^{-4}, spatially dependent Σ_f, with $\Sigma_f^1 = 1$ and reduced tolerance solves. The solid $\triangle$ is uniform P_n, with the dashed $\triangle$ FP_n with $\Sigma_f = 1$, the dotted $\triangle$ FP_n with $\Sigma_f = 100$ and $\diamond$ uniform LS P^0 FEM.

Adapt step (max. order):	1 (1)	2 (5)	3 (9)	4 (13)	5 (17)	6 (21)	7 (25)	8 (29)
Cum. runtime (μs) per final DOFs:	45.3	113.5	162.6	174	184.6	209.2	228.2	256.6
% runtime Riemann decomp.:	0.002	0.8	1.4	1.8	3.3	4.8	5.5	6.5
Peak memory use:		266.6	239.3	127.9	86.8	71.6	64.8	63.2	63.5

Table 3: Runtime, percentage of runtime spent computing Riemann decompositions and peak memory used for the 2D dogleg problem, for the goal-based FP_n adapt with error target 1×10^{-1}, spatially dependent $\Sigma_f^1 = 10$ and reduced tolerance solves. Peak memory use is on the heap (measured by massif) scaled to the size of the angular flux. The runtime is the cumulative runtime of all adapt steps up to that level, scaled by the NDOFs in the final adapt step.

scalable goal-based P^0 wavelet calculation from [8], both in NDOFs and runtime.

This behaviour will not continue as the maximum adaptive order of our FP_n method increases given the non-scalable computation/application of Riemann conditions discussed above, whereas the wavelets remain scalable given any level of refinement. Indeed we see this in Table 3, where the cumulative runtime per DOF is starting to grow slowly with angular order for the adaptive FP_n. Given we have a fixed iteration count, we know this is due to the Riemann decompositions, and we can see in Table 3 the overall percentage runtime spent computing these decompositions also rising. Even so, the ability to produce a ray-effect free solution in this streaming problem accurate to approximately 1×10^{-3} in only 120 seconds should not be discounted. Thankfully again we see the peak memory consumption for our adaptive method is scalable however, stabilising at around 63 copies of the angular flux (which is slightly less than twice that shown in Table 1, given we solve a forward and adjoint problem and benefit from reuse of some data structures).

Finally in this problem, we examine the effectivity index of our goal-based error metrics. In [8], we showed evidence of a scalable error metric that despite producing excellent results, had a pathological effectivity index. The improved error metric discussed in this work removes this pathology; Table 4 shows that for our goal-based adaptive P_n method our effectivity index varies between 0.04 to 4.52 throughout the adapt process. This is due to both the improved error metric, but also the the absence of ray-effects in our spherical harmonics solution. This is an important result and we will examine this further in future work. Table 4 also shows that the effectivity index for our adapted FP_n with spatially dependent filter is much worse, increasing to around 22 in the 7th adapt step. This is also not pathological and the results from Fig. 11 indicate that even with a worse effectivity index than the adapted P_n, our
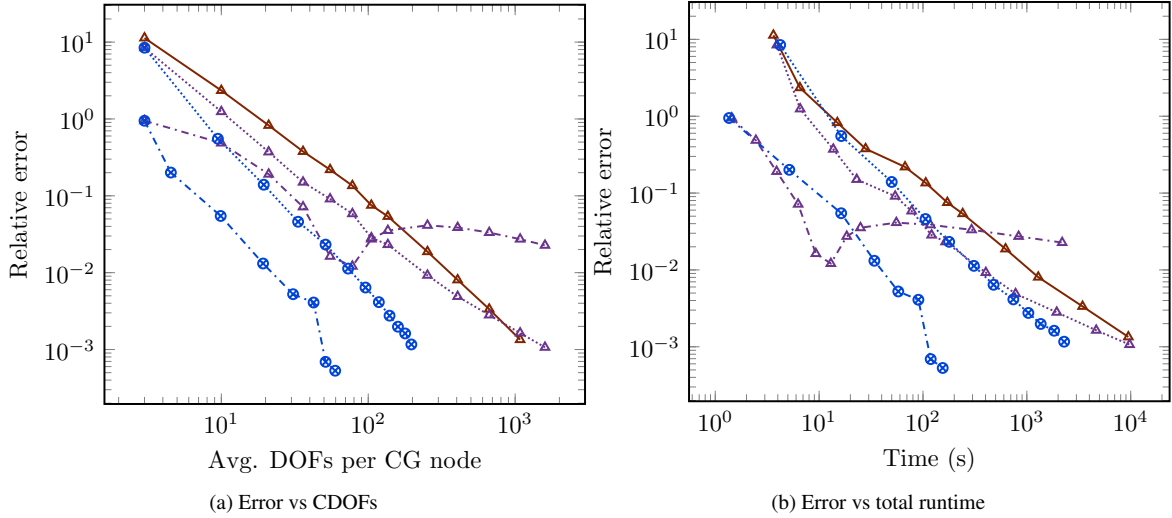
15

(a) Error vs CDOFs	(b) Error vs total runtime

Figure 8: Effect of changing the filter strength of FP_n, in the relative error of the detector response, for the 2D dogleg problem. The solid △ is uniform P_n and the △ are uniform FP_n, with densely dotted $\Sigma_f = 0.1$ and dash dotted $\Sigma_f = 10$. The ⊗ are goal-based FP_n adapts, error target 1×10^{-1} and reduced tolerance solves, spatially dependent Σ_f with densely dotted $\Sigma_f^1 = 0.1$ and dash-dotted $\Sigma_f^1 = 10$

error metric still results in refinement in correct regions, much like in [8]. Given the discussion in Section 5.1.2, we know this is due to ignoring the differing filter values across space in the forward and adjoint problems, and if we wish to improve our effectivity index in this case we simply need to project into a common space; we leave examining this to future work.

Adapt step (max. order):	1 (1)	2 (5)	3 (9)	4 (13)	5 (17)	6 (21)	7 (25)	8 (29)	9 (33)	10 (37)	11 (41)
Goal-based P_n	0.48	4.52	3.76	1.58	1.06	1.32	0.08	0.04	1.33	0.74	1.52
Goal-based FP_n, $\Sigma_f^1 = 10$	0.48	8.75	16.4	10.1	13.3	6.17	22.46	11.79	-	-	-

Table 4: Effectivity index for the goal-based adapted discretisations shown in Fig. 11, for the 2D dogleg problem.

8.3. 3D void problem

Our final example problem is a source/detector problem in 3D, with a pure vacuum in the duct. This is the same as that shown in [8], except we have shortened the length of the duct from 320cm to 100cm. We discretise this problem with an unstructured tetrahedral mesh with 31,542 elements (9,525 CG nodes and 126,168 DG nodes) and (this is a different mesh from [8] as we do not have ray-effects that we would like to keep aligned with the mesh/geometry). This is a challenging problem for a spherical harmonics method, given the pure streaming, while also being trivial given we know the intensity of radiation must fall off from a (point) source like r^{-2}. We could not compute a numerical solution with our iterative method for P_n in this problem, given the poor conditioning. As such, this is an excellent problem for understanding the behaviour of the FP_n method in the streaming limit. We also produce a reference solution in this problem with the non-standard Haar wavelets from [8], with fixed refinement up to 14 levels between $\mu \in [-0.0099, 0.0099]$ and $\omega[1.5611.5807]$.

Fig. 12 shows the results from uniform FP_n simulations in this problem, with two different constant filter strengths. We can see in Fig. 12a that with constant filter of $\Sigma_f = 1$, the FP_9 solution is the first to be both non-zero and positive along the entire length of the duct, with further refinement causing the solution to approach a highly refined P^0 solution (that agrees with the r^{-2} drop-off). If we apply a stronger constant filter of $\Sigma_f = 10$, we can see in Fig. 12b that even refinement up to FP_{21} has failed to produce a non-zero solution at the end of the duct. This is to be expected, as heavier filtering smooths out peaks in the angular flux (by acting like "negative" forward-peaked scatter), making our solution

less anisotropic. Of course this heavier filtering results in improved conditioning of our linear system and hence less iterations; Table 5 shows that the iteration count for FP_{21} with $\Sigma_f = 10$ is almost half that of $\Sigma_f = 1$. Importantly, we can see that even with a filter of $\Sigma_f = 1$, we have an almost constant iteration count with angular refinement in a problem with pure streaming. Although Fig. 12a shows that FP_{21} with $\Sigma_f = 1$ would not be considered high enough angular resolution to resolve this problem, it is perhaps surprising that a filtered spherical harmonics method can produce a non-zero solution (that is naturally free of ray-effects) in a duct with aspect ratio of 1:100 with such low resolution.

Order:	1	5	9	11	15	21
Uniform P_n	-	-	-	-	-	-
Uniform FP_n, $\Sigma_f = 1$	46	61	179	185	196	207
Uniform FP_n, $\Sigma_f = 10$	32	35	92	95	102	112
Adapt step (max. order):	**-**	**-**	**1 (9)**	**2 (11)**	**3 (15)**	**4 (21)**
Goal-based FP_n, $\Sigma_f^1 = 1$	-	-	179	195	207	210

Table 5: Number of iterations for the 3D void problem, solved to a relative tolerance of 1×10^{-14}. The P_n simulations all diverged given the pure vacuum.

If we are to apply our goal-based adaptivity in this problem, we must have a non-zero response in our functional at the coarsest angular resolution, as discussed in Section 5.1.2. These results indicate that for this 3D duct problem our FP_9 discretisation with $\Sigma_f^1 = 1$ would serve as a suitable coarse discretisation. We therefore use this as our coarse angular discretisation with our goal-based adaptivity and a spatially dependent filter value. In order to ensure our spatially dependent filter does not go to zero down the length of the duct (given the 15 orders of magnitude drop in scalar flux), as mentioned in Section 6, we scale Σ_{stab} for both the forward and adjoint problems by the respective scalar flux. Fig. 13a shows that our goal-based adaptivity successfully refines from our coarse FP_9 discretisation, producing a solution that is converging towards the refined P^0 solution (we also plot the refined adjoint solution to show it is similarly non-zero and positive). We don't plot the angular order at each spatial node in this problem, as the the goal-based adaptivity has resulted in angular refinement at every spatial node up to the max. order at each step, giving the same resolution as the uniform, as we might expect from a goal-based error metric in this problem.

Fig. 13b shows the size of the spatially dependent filter down the length of the duct for both the forward and adjoint solutions, and we can see that the filter strength is almost constant. This implies that our net stabilisation drops off in a similar manner to the magnitude of the scalar flux, and that the scaled version of Σ_{stab} is suitable for computing our Σ_f in streaming problems. The magnitude has been reduced slightly from $\Sigma_f^1 = 1$ and this explains the improved solution shown from the FP_{21} solution in Fig. 13a compared with that in Fig. 12a. This slight drop in filter strength does not heavily affect the iteration count with refinement, as shown in Table 5, with the goal-based FP_n method using 210 iterations in the final adapt step. We should also note that the cumulative runtime per final DOF does increase with adapt step in this problem, increasing from 596 μs, 1079 μs, 1692 μs to 1794 μs in the final adapt step. Note at such low order this is not due to computing the Riemann decompositions, which take only 1.5% of the total runtime in the final adapt step (given our grouping, etc), but instead to applying them on each face.

Table 6 also shows the effectivity index for our goal-based FP_n simulation, we can see that the our index is quite poor for early adapt steps, but improves considerably as we refine. This is to be expected in this problem, as Figures 12 and 13a show our coarse angular solution is also poor near the end of the duct. The key point is that our effectivity index is non-zero, and Fig. 13 shows it is causing refinement in the correct spatial regions.

Adapt step (max. order):	1 (9)	2 (11)	3 (15)	4 (21)
Goal-based FP_n, $\Sigma_f^1 = 1$	8.5×10^{-9}	5.3×10^{-4}	3.1×10^{-3}	INSERT WHEN DONE **??**

Table 6: Effectivity index for the goal-based adapted discretisations shown in Fig. 11, for the 3D duct problem.

9. Conclusions

This paper has presented an angular adaptivity algorithm for spherical harmonics that features both regular and goal-based error metrics, in problems with a range of smoothness. In particular, we examined the use of filtered P_n methods in problems with heavy streaming. We used angular adaptivity in conjuction with FP_n and this allowed us to easily build spatially dependent filter strengths that filter heavily near discontinuities in space/angle, while retaining their spectral order in smooth regions. This is in an attempt to produce fast, ray-effect free solutions to problems with heavy streaming and/or voids.

We found that for problems without heavy streaming, (i.e., those with some smoothness), adaptive P_n and adaptive FP_n methods can perform better in NDOFs applied and runtime than uniform discretisations. In particular, care must be taken with FP_n methods to not filter "too much". The introduction of a spatially dependent filter however made this process more robust, and our spatially dependent filter always outperformed a constant filter value. For problems with heavy streaming, we found our adaptive P_n methods at least performed better than uniform P_n, or unsurprisingly were too poorly conditioned to solve. Surprisingly we found our adaptive FP_n methods highly performant in streaming problems, with the spatially dependent filter giving performance comparable to adapted P^0 discretisations up to a certain order. In particular, we found close to fixed iteration counts with angular refinement.

This gives a spherical harmonics method that is very close to "scalable". Of course the last barrier to true scalability in spherical harmonics methods is applying/computing inflow/outflow/BCs. At high-order the (at best) $O(n^2)$ nature of these computations will always dominate, but our adaptive process at least allows the number/size of these computations to be minimised. We also found that using a spatially dependent filter in a goal-based FP_n simulations with a pure vacuum allowed us to get non-zero responses in our functionals at the end of a duct with aspect ratio of 1:100, even with low-order FP_n discretisations. This triggered refinement in our spherical harmonics in a problem where previous work [8] has found significant difficulties in developing robust goal-based metrics.

One fundamental disadvantage of all FP_n methods presented, both in this work and in the literature is picking a filter strength for a given problem. This work showed that allowing an adaptive process to make this filter strength spatially dependent makes the initial choice for a filter value less important, and helps reduce the risk of over-filtering. Indeed, based on this work we would always recommend running with a spatially dependent filter, as it improved the convergence in all the cases we investigated. We believe this work shows that even without $O(n)$ scaling in angle size, adaptive filtered spherical harmonics are a powerful tool for fast, ray-effect free solutions for problems that can include pure streaming.

### References

[1] A. G. Buchan, C. C. Pain, M. D. Eaton, R. P. Smedley-Stevenson, A. J. H. Goddard, Self-adaptive spherical wavelets for angular discretisations of the boltzmann transport equation, Nucl. Sci. Eng. 158 (2008) 244–263.

[2] M. A. Goffin, A. G. Buchan, A. C. Belme, C. C. Pain, M. D. Eaton, P. N. Smith, R. P. Smedley-Stevenson, Goal-based angular adaptivity applied to the spherical harmonics discretisation of the neutral particle transport equation, Ann. Nucl. Energy 71 (2014) 60–80.

[3] M. A. Goffin, A. G. Buchan, S. Dargaville, C. C. Pain, P. N. Smith, R. P. Smedley-Stevenson, Goal-based angular adaptivity applied to a wavelet-based discretisation of the neutral particle transport equation, Journal of Computational Physics 281 (2015) 1032–1062.

[4] M. Goffin, Goal-based adaptive methods applied to the spatial and angular dimensions of the transport equation, Ph.D. thesis, Imperial College London, 2015.

[5] A. Adam, A. G. Buchan, M. D. Piggott, C. C. Pain, J. Hill, M. A. Goffin, Adaptive Haar wavelets for the angular discretisation of spectral wave models, Journal of Computational Physics 305 (2016) 521–538.

[6] A. Adam, Finite element, adaptive spectral wave modelling, Ph.D. thesis, Imperial College London, 2016.

[7] L. Soucasse, S. Dargaville, A. G. Buchan, C. C. Pain, A goal-based angular adaptivity method for thermal radiation modelling in non grey media, Journal of Quantitative Spectroscopy and Radiative Transfer 200 (2017) 215–224.

[8] S. Dargaville, A. G. Buchan, R. P. Smedley-Stevenson, P. N. Smith, C. C. Pain, Scalable angular adaptivity for Boltzmann transport, arXiv:1901.04929 [physics] (2019). ArXiv: 1901.04929.

[9] R. G. McClarren, C. D. Hauck, Robust and accurate filtered spherical harmonics expansions for radiative transfer, Journal of Computational Physics 229 (2010) 5597–5614.

[10] D. Radice, E. Abdikamalov, L. Rezzolla, C. D. Ott, A new spherical harmonics scheme for multi-dimensional radiation transport I. Static matter configurations, Journal of Computational Physics 242 (2013) 648–669.

[11] M. Frank, C. Hauck, K. Kuepper, Convergence of filtered spherical harmonic equations for radiation transport, Commun. Math. Sci 14 (2016) 1443–1465.

[12] V. M. Laboure, Improved Fully-Implicit Spherical Harmonics Methods for First and Second Order Forms of the Transport Equation Using Galerkin Finite Element, Thesis, 2016.

[13] V. M. Laboure, R. G. McClarren, C. D. Hauck, Implicit Filtered PN for High-Energy Density Thermal Radiation Transport using Discontinuous Galerkin Finite Elements, Journal of Computational Physics 321 (2016) 624–643. ArXiv: 1601.08242.

[14] M. P. Laiu, C. D. Hauck, Positivity Limiters for Filtered Spectral Approximations of Linear Kinetic Transport Equations, Journal of Scientific Computing (2018).

[15] T. J. R. Hughes, G. R. Feijo, L. Mazzei, J.-B. Quincy, The variational multiscale methoda paradigm for computational mechanics, Computer Methods in Applied Mechanics and Engineering 166 (1998) 3–24.

[16] T. J. R. Hughes, G. Scovazzi, P. B. Bochev, A. Buffa, A multiscale discontinuous galerkin method with the computational structure of a continuous galerkin method, Computer Methods in Applied Mechanics and Engineering 195 (2006) 2761–2787.

[17] A. S. Candy, Subgrid scale modelling of transport processes., Thesis or dissertation, Imperial College London, 2008.

[18] A. G. Buchan, A. S. Candy, S. R. Merton, C. C. Pain, J. I. Hadi, M. D. Eaton, A. J. H. Goddard, R. P. Smedley-Stevenson, G. J. Pearce, The inner-element subgrid scale finite element method for the boltzmann transport equation, Nuclear science and engineering 164 (2010) 105–121.

[19] R. T. Ackroyd, W. E. Wilson, Discontinuous finite elements for neutron transport analysis, Progress in Nuclear Energy 18 (1986) 39–44.

[20] R. T. Ackroyd, W. E. Wilson, Composite finite element solutions for neutron transport, Ann. Nucl. Energy 15 (1988) 397–419.

[21] H. Park, Coupled Space-Angle Adaptivity and Goal-Oriented Error Control for Radiation Transport Calculations, Ph.D. thesis, Georgia Institute of Technology, 2006.

[22] H. Park, C. R. E. de Oliveira, Coupled Space-Angle Adaptivity for Radiation Transport Calculations, Nuclear Science and Engineering 161 (2009) 216–234.

[23] K. Rupp, T. Grasser, A. Jungel, Adaptive variable-order spherical harmonics expansion of the Boltzmann Transport Equation, in: 2011 International Conference on Simulation of Semiconductor Processes and Devices (SISPAD), pp. 151–154.

[24] O. Safarzadeh, A. S. Shirani, A. Minuchehr, Hybrid spaceangle adaptivity for whole-core particle transport calculations, Annals of Nuclear Energy 80 (2015) 254–260.

[25] A. G. Buchan, S. R. Merton, C. C. Pain, R. P. Smedley-Stevenson, Riemann boundary conditions for the Boltzmann transport equation using arbitrary angular approximations, Annals of Nuclear Energy 38 (2011) 1186–1195.

[26] M. J. Mohlenkamp, A fast transform for spherical harmonics, Journal of Fourier Analysis and Applications 5 (1999) 159–184.

[27] C. Lessig, T. de Witt, E. Fiume, Efficient and accurate rotation of finite spherical harmonics expansions, Journal of Computational Physics 231 (2012) 243–250.

[28] S. Dargaville, M. A. Goffin, A. G. Buchan, C. C. Pain, R. P. Smedley-Stevenson, P. N. Smith, G. Gorman, Solving the boltzmann transport equation with multigrid and adaptive space/angle discretisations, Annals of Nuclear Energy 86 (2015) 99–107.

[29] B. J. Adigun, A. G. Buchan, A. Adam, S. Dargaville, M. A. Goffin, C. C. Pain, A Haar wavelet method for angularly discretising the Boltzmann transport equation, Progress in Nuclear Energy 108 (2018) 295–309.

[30] A. G. Buchan, C. C. Pain, An efficient space-angle subgrid scale discretisation of the neutron transport equation, Annals of Nuclear Energy 94 (2016) 440–450.

[31] J. C. Ragusa, J. L. Guermond, G. Kanschat, A robust SN-DG-approximation for radiation transport in optically thick and diffusive regimes, Journal of Computational Physics 231 (2012) 1947–1962.

[32] T. A. Brunner, Forms of approximate radiation transport, Sandia report (2002).

19
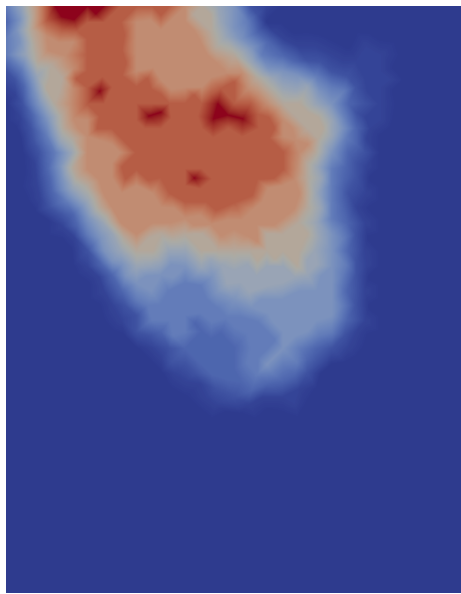
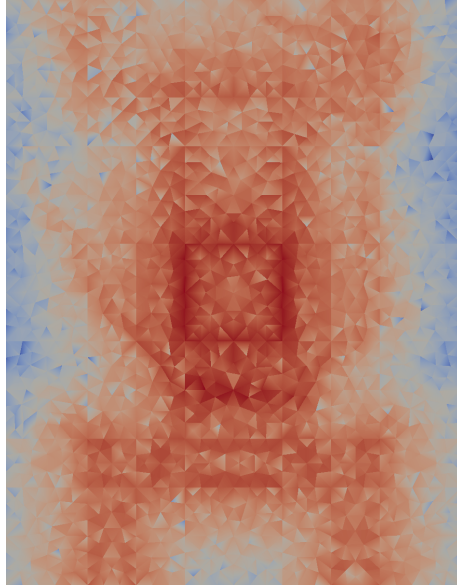(a) Number of angular basis functions across the spatial domain - 603K DOFs
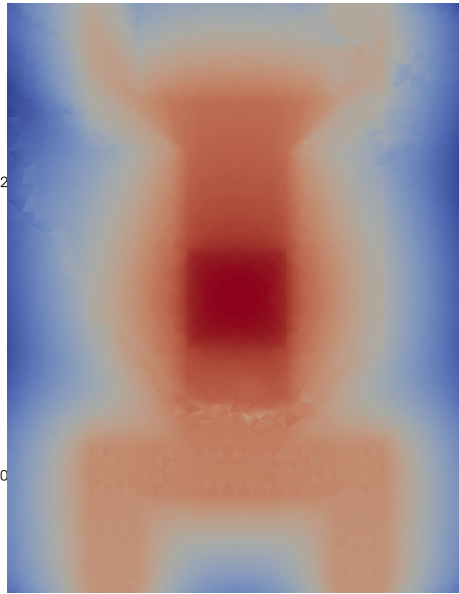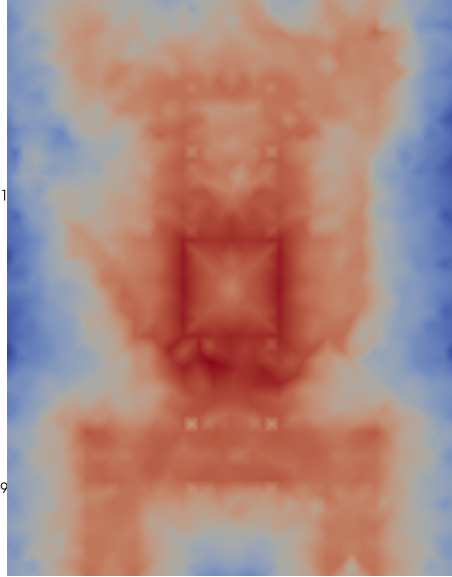

(b) Scalar flux across the domain
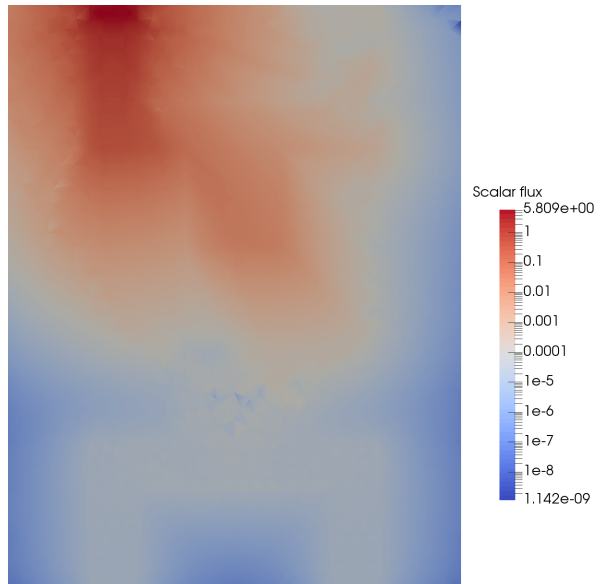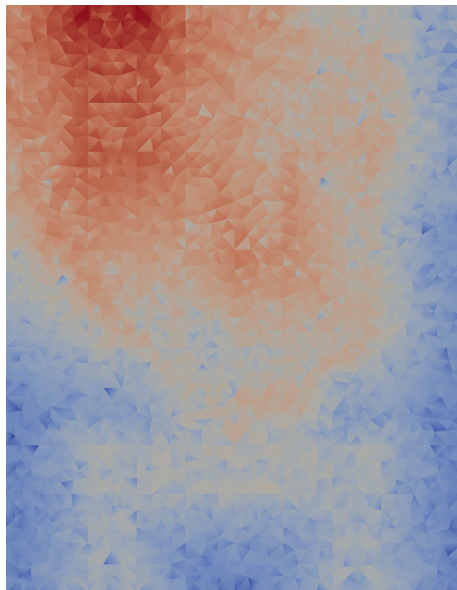

(c) Scalar flux in the adjoint

Figure 9: Results from the goal-based adaptivity with FP_n on the 8th adapt step with error target 1×10^{-1} for the 2D dogleg problem, with spatially dependent Σ_f, with $\Sigma_f^1 = 10$.
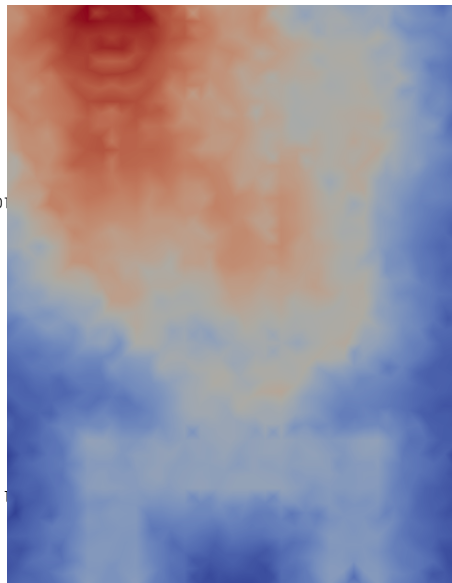
(a) Absolute value of Σ_{stab} across space. This can be considered the net "amount" of stabilisation we apply at each node.

(b) Spatially-dependent Σ_{f}

(c) Absolute value of Σ_{stab} across space for the adjoint. This can be considered the net "amount" of stabilisation we apply at each node.

(d) Spatially-dependent Σ_{f} in the adjoint

Figure 10: Computation of a spatially dependent Σ_{f}, with $\Sigma_{\text{f}}^1 = 10$, on the the 8th step of FP_n goal-based adaptivity with error target 1×10^{-1} for the 2D dogleg problem.
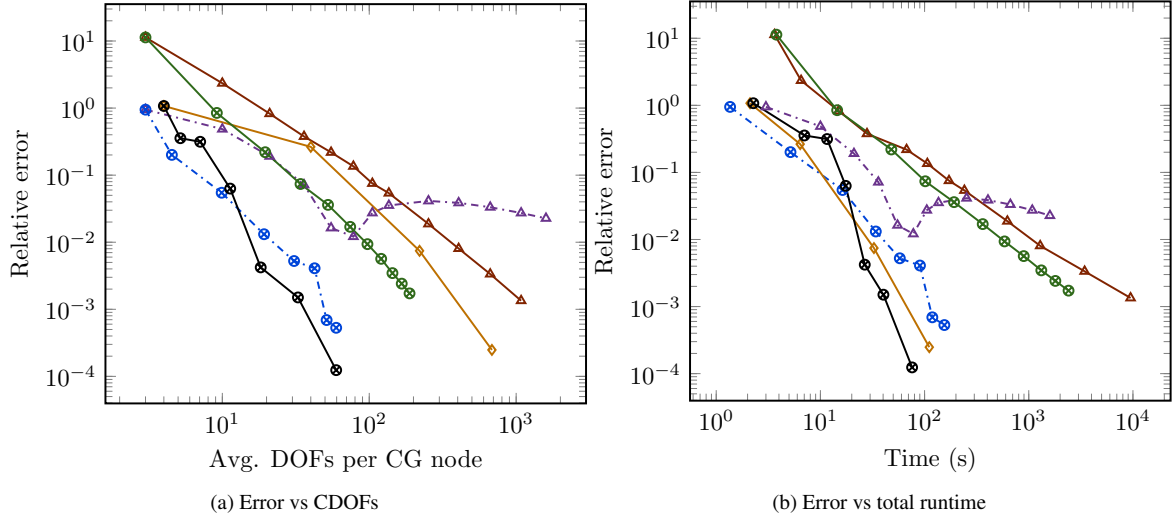
(a) Error vs CDOFs

(b) Error vs total runtime

Figure 11: Comparison of the relative error of the 2-norm of the scalar flux across the domain, for different angular discretisations, for the 2D dogleg problem. The $\otimes$ are goal-based P_n adapts with error target 1×10^{-1} and reduced tolerance solves, with the dashed $\otimes$ goal-based FP_n adapts with error target 1×10^{-1}, spatially dependent Σ_f, with $\Sigma_f^1 = 10$ and reduced tolerance solves. The solid $\triangle$ is uniform P_n, the dash-dotted $\triangle$ is uniform FP_n with $\Sigma_f = 10$ and $\diamond$ uniform LS P^0 FEM. The $\otimes$ are goal-based adapted non-standard Haar wavelets with error target 1×10^{-3} and one extra adapt step (from [8])
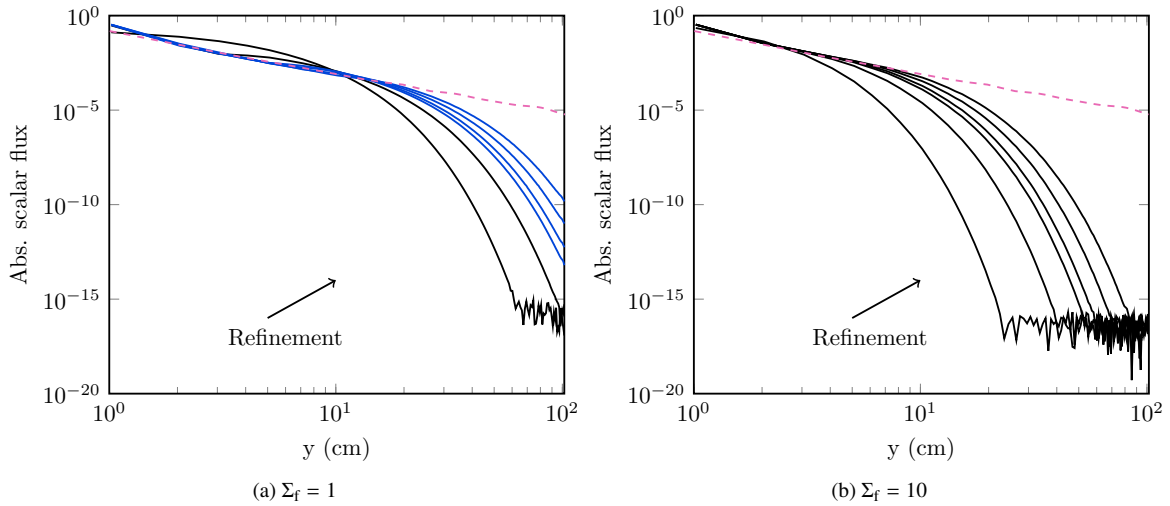

(a) $\Sigma_f = 1$

(b) $\Sigma_f = 10$

Figure 12: Absolute value of the scalar flux in the 3D duct along the length of the duct, from (0.5, 0.0, 0.5) to (0.5, 102, 0,5). The solid lines are FP_n, orders 1, 5, 9, 11, 15 and 21. The FP_n solid lines are shaded blue when they are entirely positive down the length of the tube. The dashed pink line corresponds to the non-standard Haar wavelet discretisation in [8] with 14 levels of refinement.

22
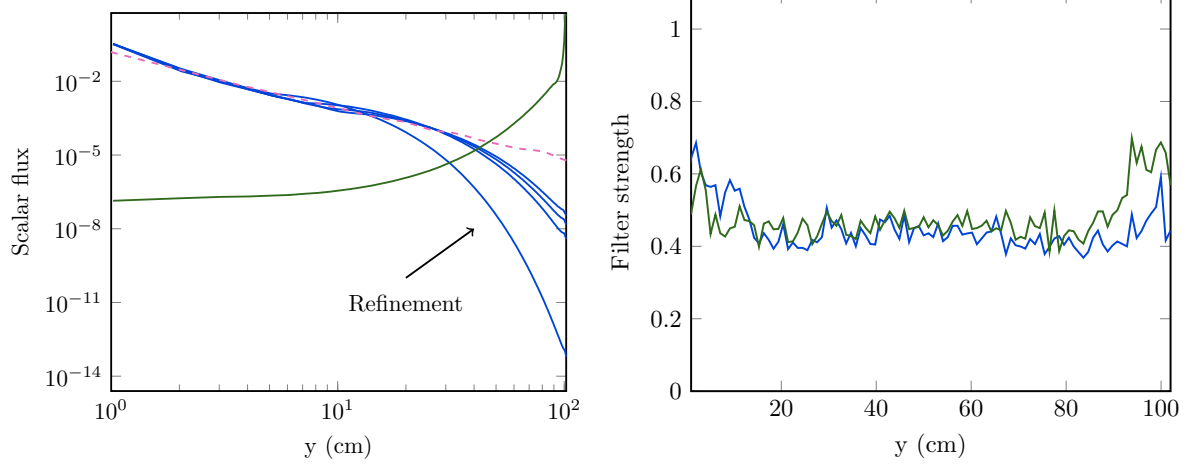
(a) Scalar flux along the length of the duct, from (0.5, 0.0, 0.5) to (0.5, 102, 0,5). The blue lines are the forward problems, the green line is the adjoint problem on the 3rd adapt step. The dashed pink line corresponds to the non-standard Haar wavelet discretisation in [8] with 14 levels of refinement.

(b) Filter strength along the length of the duct, from (0.5, 0.0, 0.5) to (0.5, 102, 0,5) on the 4th adapt step. The blue line is the forward problem, the green line is the adjoint.

Figure 13: Results from goal-based FP_n in the 3D duct problem with $\Sigma_f^1 = 1$, starting from FP_9 with four adapt steps.