NUCLEAR SCIENCE AND ENGINEERING: **165**, 305–319 (2010)

Goal-Oriented *h*-Adaptivity for the Multigroup SP_N Equations

Bruno Turcksin and Jean C. Ragusa

Texas A&M University, Department of Nuclear Engineering
College Station, Texas 77843

and

Wolfgang Bangerth

Texas A&M University, Department of Mathematics
College Station, Texas 77843

Received June 7, 2009
Accepted November 13, 2009

Abstract—*We investigate application of goal-oriented mesh adaptivity to the SP_N multigroup equations. This technique utilizes knowledge of the computational goal and combines it with mesh adaptivity to accurately and rapidly compute quantities of interest. Specifically, the local error is weighted by the importance of a given cell toward the computational goal, resulting in appropriate goal-oriented error estimates. Even though this approach requires the solution of an adjoint (dual) problem, driven by a specific source term for a given quantity of interest, the work reported here clearly shows the benefits of such a method.*

We demonstrate the level of accuracy this method can achieve using two-dimensional and three-dimensional numerical test cases for one-group and two-group models and compare results with more traditional mesh refinement and uniformly refined meshes. The test cases consider situations in which the radiative flux of a source is shielded and are designed to prototypically explore the range of conditions under which our methods improve on other refinement algorithms. In particular, they model strong contrasts in material properties, a situation ubiquitous in nuclear engineering.

I. INTRODUCTION

Mesh adaptation techniques aim at locally refining or coarsening mesh cells in a computational domain in order to obtain an accurate approximate solution of a partial differential equation with lower memory requirements and in less CPU time compared to a uniform mesh approach. In the past two decades, these techniques have received significant attention in the mathematical[1] and engineering sciences communities[2] and, more recently, in nuclear science and engineering; see, e.g., Refs. 3 and 4 for one-group two-dimensional (2-D) discrete ordinates transport, Refs. 5 and 6 for one-group 2-D diffusion, Ref. 7 for multigroup one-dimensional diffusion, Refs. 6 and 8 for 2-D and three-dimensional (3-D) multigroup diffusion, and Ref. 9 for multigroup 2-D SP_N.

E-mail: ragusa@ne.tamu.edu

Standard adaptive mesh refinement (AMR) techniques are based on the availability of a quantity that estimates or approximates the error locally, i.e., the difference between the exact and numerical solutions, for each cell; one then aims at equi-distributing the error throughout the mesh by refining cells that have a large error and possibly coarsening cells with small errors until the overall error is below a prescribed threshold.[10–13] While such methods have been proven to be highly efficient and accurate compared to uniformly refined meshes, this "traditional" approach to adaptivity is not always optimal for the following reasons and can be further improved:

1. Obtaining a highly accurate solution in every single mesh cell of the computational domain may not be necessary from a practical engineering point of view. Rather, one frequently runs a simulation with a particular goal in mind. Consequently, the ultimate answer sought

may be a functional of the solution rather than the solution itself. In nuclear engineering, examples of such functionals may include reaction rates, fluxes, and currents, either at a single point or integrated over a subdomain, for all or a portion of the energy range.

2. In some cases, where the solution can vary by orders of magnitude over the whole domain, the absolute error in regions of low values may be negligible, whereas the relative error in these same regions may still be unacceptably high. Examples of such situations include shielding problems, where standard AMR techniques would tend to refine the mesh where the solution is large and experiences significant spatial variations, e.g., in the source region and at the source/shield interface, whereas the zone past the shield (often the region of interest) would not be refined and can present large relative errors, even when standard AMR is used.

In order to address these points, goal-oriented adaptivity has been developed over the course of the last decade.[10,11,14–18] This method derives error estimates for the goal of the computation rather than for the global error in the solution. Cellwise refinement indicators are then composed of both the local error (i.e., how well we approximate the exact solution on this cell) as well as a quantity (the "dual weight") that describes how important the accuracy of the solution on a given cell is toward the goal of the computation. For example, it may be that a large error can be tolerated on a cell far away from the region of interest, and a small dual weight will then indicate that this cell does not need to be refined despite its large absolute error. Conversely, a cell with a relatively small error may still require further refinement if this cell is important for our goal as indicated by a large dual weight.

The quantities of interest that drive the adaptation in the goal-oriented calculations are user-defined (see the examples given in Sec. III.D). In the formalism that we outline in Sec. III, it becomes clear that the goal functional defines the source terms of an adjoint equation whose solution (the dual or adjoint solution) will give rise to the dual weights for the goal-oriented error indicators. This process requires that not only the original primal (or forward) problem but also the adjoint problem be solved at each mesh adaptivity cycle. We show that this combined process is still significantly more efficient than using uniform meshes or standard AMR because it can generate meshes that are tailored toward a particular goal rather than resolving the solution everywhere. As a consequence, one is able to achieve greater precision with fewer unknowns and smaller run time. As will be shown through our numerical results, the reduction in CPU time is at least one order of magnitude and sometimes more. Finally, we note that the use of a deterministic adjoint solution is widely spread in the Monte Carlo community to devise automatic variance-reduction algorithms; see, e.g., Refs. 19 and 20.

In this paper, we apply goal-oriented adaptivity to the multigroup SP_N equations. The SP_N equations have been observed to yield improved results when compared to the diffusion approximation of transport phenomena and are less computationally demanding than other higher-order transport approximations such as the spherical harmonics (P_N) and discrete ordinates (S_N) methods.[21] In addition, the SP_N equations can be recast into a set of coupled diffusion-like equations, with Laplace operators, and standard linear algebra techniques such as preconditioned conjugate gradient can easily be used to solve such systems of equations. Moreover, the error estimates used in this work have been shown to perform well for elliptic problems, and the results of this paper demonstrate that this is also the case for the problems considered here. In fact, using the techniques described below, we believe that we are able to drive the numerical error far below the modeling error that results from replacing the transport equations by the SP_N approximation. As a consequence, one can argue that this level of accuracy may not be necessary in the current context. Nevertheless, we believe that the current study serves as an important proof of concept that duality-based error estimation techniques are a useful tool for high-fidelity numerical simulations in nuclear engineering applications. We expect that such techniques can also be extended to higher-fidelity transport approximations, using the methods described in Ref. 22 for the demonstration of AMR applied to the S_N approximation and Ref. 23 for AMR applied to the P_N form of the transport equation. Finally, Refs. 24 and 25 describe techniques to extend goal-oriented techniques to eigenvalue problems; such methods could be used to compute the critical eigenvalue k_{eff} in reactor calculations.

The outline of this paper is as follows. In Sec. II, we briefly review the SP_N equations and their finite element discretization. Section III reviews the principles of mesh adaptation and provides a derivation of the goal-oriented error estimator used herein. Refinement strategies and quantities of interest are also discussed in Sec. III. We present 2-D and 3-D results for one-group and multigroup problems in Sec. IV and give concluding remarks in Sec. V.

II. MATHEMATICAL FORMULATION FOR THE SP_N EQUATIONS

The aim of this paper is to discuss the feasibility of goal-oriented simulations in nuclear science and engineering applications. It is meant to be a proof of principle, and given the high phase-space dimensionality of the transport equation, only the simplified P_N (SP_N) equations[21,26–28] are solved here instead of the full transport equations. Since the focus of this paper is on advanced numerical techniques, the reader is referred to the cited references for a more complete discussion of

the approximations and the validity of the SP_N equations. However, we note that the numerical examples employed here have been shown to be adequately treated with the SP_N approximation,[29] even though in general the SP_N answers do not converge to the exact transport answers.

II.A. Variational Form

The multigroup SP_N approximation is employed. Let g denote the energy group index ($1 \le g \le G$) and consider N flux moments $\Phi_{g,n}(\mathbf{r})$ for a given group g, with $0 \le n \le N$ (N is odd) denoting the moment index. In the SP_N approximation, one can recast the system of $N+1$ coupled first-order differential equations as a system of $M = \frac{1}{2}(N+1)$ coupled second-order differential equations, with the M even modes $\Phi_{g,n}(\mathbf{r})$ ($n = 0, 2, \ldots, N-1$) as unknowns. Using appropriate linear combinations of these even fluxes into a vector of composite even fluxes $\mathsf{F}_g(\mathbf{r}) = (\mathsf{F}_{g,1}, \mathsf{F}_{g,2}, \ldots, \mathsf{F}_{g,M})^T$ of size M, the SP_N approximation then yields the following set of coupled diffusion-like equations:

$$-\boldsymbol{\nabla}\cdot(\boldsymbol{\sigma}^o)^{-1}\boldsymbol{\nabla}\mathsf{F} + C^{-T}\boldsymbol{\sigma}^e C^{-1}\mathsf{F} = C^{-T}S^e - \boldsymbol{\nabla}\cdot(\boldsymbol{\sigma}^o)^{-1}\mathbf{S}^o \ . \tag{1}$$

where

$\boldsymbol{\sigma}^o, \boldsymbol{\sigma}^e$ = $G \times G$ matrices of $M \times M$ matrices encoding the absorption and scattering coefficients and their angular moments

C = $M \times M$ matrix coupling different angular modes

$S^e, \mathbf{S}^o$ = mode-dependent even and odd source terms.

Summation over appropriate indices is implied; for example,

$$(C^{-T}\boldsymbol{\sigma}^e C^{-1}\mathsf{F})_{g,m} = \sum_{g'=1}^{G}\sum_{m',m'',m'''=1}^{M} (C^{-T})_{mm'} \times \sigma^e_{gg',m'm''}(C^{-1})_{m''m'''}\mathsf{F}_{g',m'''} \ .$$

The above equations are complemented by Marshak boundary conditions

$$\mathbf{n}\cdot\nabla\mathsf{F} = \mathbf{n}\cdot\mathbf{S}^o + \boldsymbol{\sigma}^o A_o^{-1} B C^{-1}\mathsf{F} - \boldsymbol{\sigma}^o A_o^{-1} J^{\text{inc}} \ , \tag{2}$$

where

A_o, B = $M \times M$ matrices coupling individual modes

$J^{inc}_{g,m}$ = incoming boundary sources for energy group g and mode m.

We have used a particular Legendre polynomial scaling that makes the matrices A_o, B, and C less cumbersome, resulting in a simpler code implementation. For additional details, we refer the reader to Refs. 9 and 30.

The weak or variational formulation of the SP_N equations is obtained by multiplying Eq. (1) with a test function $b(\mathbf{r})$ (also composed of group and mode components), integrating over the domain V, and integrating second derivatives by parts. This yields the weak form

$$\int_V [\boldsymbol{\nabla}b]^T(\boldsymbol{\sigma}^o)^{-1}[\boldsymbol{\nabla}\mathsf{F}] + \int_V b^T C^{-T}\boldsymbol{\sigma}^e C^{-1}\mathsf{F} - \int_{\partial V} b^T\mathbf{n}\cdot(\boldsymbol{\sigma}^o)^{-1}\boldsymbol{\nabla}\mathsf{F} = \int_V b^T[C^{-T}S^e - \boldsymbol{\nabla}\cdot(\boldsymbol{\sigma}^o)^{-1}\mathbf{S}^o] \ , \tag{3}$$

where contraction over mode and group indices is again implied. In the third term of Eq. (3), we use the boundary condition Eq. (2). Furthermore, we can integrate by parts to obtain $-\int_V b^T\boldsymbol{\nabla}\cdot(\boldsymbol{\sigma}^o)^{-1}\mathbf{S}^o = \int_V[\boldsymbol{\nabla}b]^T(\boldsymbol{\sigma}^o)^{-1}\mathbf{S}^o - \int_{\partial V} b^T\mathbf{n}\cdot(\boldsymbol{\sigma}^o)^{-1}\mathbf{S}^o$. Inserting this relation and Eq. (2) into the weak form, Eq. (3), yields

$$\int_V [\boldsymbol{\nabla}b]^T(\boldsymbol{\sigma}^o)^{-1}[\boldsymbol{\nabla}\mathsf{F}] + \int_V [C^{-1}b]^T\boldsymbol{\sigma}^e[C^{-1}\mathsf{F}] + \int_{\partial V} b^T A_o^{-1} B C^{-1}\mathsf{F} = \int_V b^T C^{-T}S^e + \int_V [\boldsymbol{\nabla}b]^T(\boldsymbol{\sigma}^o)^{-1}\mathbf{S}^o + \int_{\partial V} b^T A_o^{-1} J^{inc} \ . \tag{4}$$

Using abbreviations for the left and right sides, the problem to solve is then to find composite flux moments F such that

$$a(b,\mathsf{F}) = l(b) \quad \text{for all test functions } b \ , \tag{5}$$

where $a(\cdot,\cdot)$ is the bilinear form and $l(\cdot)$ is the linear form associated with the left and right side terms, respectively. Note that the bilinear form is nonsymmetric because of the boundary terms, whereas the domain terms are symmetric in the one-group case ($\boldsymbol{\sigma}^o$ is then a diagonal matrix). In the multigroup case, the diagonal blocks of $\boldsymbol{\sigma}^o$ and $\boldsymbol{\sigma}^e$ contain the intergroup scattering matrices, which are nonsymmetric.

II.B. Discrete Problem

Galerkin methods approximate the solution F of Eq. (5) by considering linear combinations of basis functions $b_j(\mathbf{r})$,

$$\mathsf{F}_h(\mathbf{r}) = \sum_{j=1}^{P} \gamma_j b_j(\mathbf{r}) \ , \tag{6}$$

and determining the P coefficients γ_j by requiring that the following P equations hold:

$$a(b_i, F_h) = l(b_i) \quad \text{for } 1 \le i \le P \ . \tag{7}$$

Expanding F_h in this way, Eq. (7) leads to the following linear problem:

$$\mathbf{AG} = \mathbf{L} \ , \tag{8}$$

where

$$\mathbf{A}_{ij} = a(b_i, b_j)$$

$$\mathbf{G}_i = \gamma_i$$

$$\mathbf{L}_i = l(b_i).$$

For future use, let us subdivide the P coefficients γ_i and the shape functions b_i by energy groups g and angular modes m. For each component g, m, we will then have $P_{g,m}$ shape functions $b_{g,m,i}, i = 1, \ldots, P_{g,m}$; each of these shape functions is still a vector of GM components, but only the g, m'th of these components is nonzero. Correspondingly, $\gamma_{g,m,i}, i = 1, \ldots, P_{g,m}$ are the unknown coefficients. There is a total of $P = \sum_{g=1}^{G} \sum_{m=1}^{M} P_{g,m}$ such shape functions and coefficients. Using this subdivision, we can conceptually think of $\mathbf{A}$ as a matrix with $GM \times GM$ blocks, and $\mathbf{G}$ and $\mathbf{L}$ as vectors with GM blocks. We then solve the linear system in Eq. (8) using a block Gauss-Seidel method, where each block represents one energy and mode. The diagonal blocks are symmetric and positive definite, and we invert them using a standard conjugate gradient method, preconditioned by successive overrelaxation.

In finite element methods, the basis functions are typically defined on a grid composed of triangles/quadrilaterals (in two dimensions) or tetrahedra/hexahedra (in three dimensions), and using functions that are polynomials on each cell and continuous across cell faces. Theoretically, we could use shape functions $b_{g,m,i}$ defined independently on GM different meshes $\mathbb{T}_{g,m}$. Some or all of these meshes may, but need not, coincide. Allowing meshes to differ allows a proper distribution of mesh cells to achieve high accuracy with a reduced number of cells but at a possibly increased algorithmic cost. For traditional AMR methods applied to the multigroup diffusion approximation, this idea is investigated in Ref. 8. In particular, note that using different meshes per group and moment also gives us the freedom to approximate certain energy groups or moments with higher accuracy than others, which is a feature that we use for the goal-oriented AMR techniques discussed below.

Whether such a level of flexibility in choosing different meshes is necessary in practice to achieve efficient solution methods is a question that we investigate in Sec. IV. It will turn out that at least for the cases tested, there is in fact little additional benefit from using different spatial meshes for different angular modes, though this question is far from obvious. On the other hand, as determined in Ref. 8, using different meshes for different energy groups can make computations less costly, and therefore, we always employ different spatial meshes for each energy group. We would like to point out in this context that the computational overhead of working with different meshes compared to only a single mesh is not very large if we can make a number of basic assumptions on the meshes; for instance, efficient algorithms and data structures for this case have been proposed and investigated in the context of the multigroup diffusion approximation in Ref. 8.

III. ERROR ESTIMATION AND *h*-ADAPTIVITY

As mentioned in Sec. II, we would like to assume, at least conceptually, that we have GM different meshes $\mathbb{T}_{g,m}$ on which we discretize the variables $F_{g,m}$ for the g'th energy group and the m'th angular moment. In order to use this flexibility, we need to define a practical way in which we can construct these different meshes. In finite element approximation, mesh adaptivity is typically driven through the use of error estimates. In its most basic form,[31] adaptivity requires that we solve the problem on a coarse mesh and then evaluate an error estimator or error indicator that provides, for any given cell, an estimation of the cell contribution to the global error. Then, for mesh refinement, we select the cells that have the largest errors, and for mesh coarsening, we select those with the smallest errors. We thereby arrive at an adapted mesh on which this procedure is repeated until the overall error is satisfactorily small.

The key to applying such an approach is error indicators. For model problems, a vast body of literature exists (see, for example, Refs. 1, 10, 12, and 13 and the references therein). In this work, however, we have to deal with three nonstandard complications: (a) a fairly complex mathematical model that lacks many of the properties often used in the derivation of error estimates, such as coercivity or symmetry; (b) a multitude of meshes, all of which could be adapted independently; and (c) the fact that the quantities of interest are not natural norms of the solution but are often localized quantities.

III.A. Duality-Based Error Estimates

We approach this problem by basing our refinement indicators on the idea of duality-based or goal-oriented error estimates (see, for example, Refs. 10, 11, and 14 through 18). To this end, let us assume that we are not interested in the solution $F(\mathbf{r})$ itself but in a functional $J(F)$. One particularly simple functional could be the average over area $\omega \subset V$ of the zeroth moment of the composite flux in the third energy group, i.e., $J(F) = (1/|\omega|) \int_\omega F_{3,0}(\mathbf{r})\, \mathrm{d}r$; other examples are given later in this section. For simplicity, we assume that $J(\cdot)$ is a linear

functional, though the formalism can easily be extended to nonlinear functionals as well.[11]

Our goal is then to accurately compute $J(F)$, but all we have available is $J(F_h)$, computed from the numerical approximation F_h on our current mesh. To assess the accuracy of our answer, we would, therefore, need to estimate the error in the quantity of interest $|J(F) - J(F_h)| = |J(e)|$, where the error is defined as $e = F - F_h$. In order to estimate this error in J, let us assume that we have access to the solution $\hat{F}$ of the dual or adjoint problem defined by

$$a(\hat{F}, b) = \hat{a}(b, \hat{F}) = J(b) \quad \text{for all test functions } b \ . \tag{9}$$

Note that compared to Eq. (5), the order of test functions and solution on the left has been inverted, resulting in the use of the adjoint operator, obtained by taking the transpose of matrices appearing in the bilinear form of Eq. (5). In a similar vein to Eq. (7), let us define the numerical approximation $\hat{F}_h$ to $\hat{F}$:

$$a(\hat{F}_h, b_i) = J(b_i) \quad \text{for } 1 \le i \le P \ . \tag{10}$$

Since Eq. (9) has to hold for all test functions b, it is also true for $b = e$, yielding the equation

$$J(F) - J(F_h) = J(e) = a(\hat{F}, e) \ .$$

Next, we use the Galerkin orthogonality property of Galerkin methods: From Eq. (7) and Eq. (5), it follows that $a(b_h, F_h) = l(b_h)$ and $a(b_h, F) = l(b_h)$ for all finite element functions b_h. Consequently, $a(b_h, e) = 0$. Since $\hat{F}_h$ is a finite element function as well, we conclude that also $a(\hat{F}_h, e) = 0$. As a consequence, we obtain

$$J(F) - J(F_h) = a(\hat{F}, e) = a(\hat{F}, e) - a(\hat{F}_h, e) = a(\hat{e}, e) \ ,$$

where the dual error is defined as $\hat{e} = \hat{F} - \hat{F}_h$. Using the definition of the bilinear form in Eq. (4), we can write this equation as a sum over all cells K that appear on the various meshes $\mathbb{T}_{g,m}$:

$$\begin{aligned} J(F) - J(F_h) = & \sum_{g,g'=1}^{G} \sum_{m,m'=1}^{M} \sum_{K \in \mathbb{T}_{g,m}} \int_K [\boldsymbol{\nabla} \hat{e}_{g',m'}]^T (\sigma^o)^{-1}_{g'g,m'm} [\boldsymbol{\nabla} e_{g,m}] \\ & + \sum_{g,g'=1}^{G} \sum_{m,m',m'',m'''=1}^{M} \sum_{K \in \mathbb{T}_{g,m}} \int_K [C^{-1}_{m''m'''} \hat{e}_{g',m'''}]^T \sigma^e_{g'g,m''m'} [C^{-1}_{m'm} e_{g,m}] \\ & + \sum_{g=1}^{G} \sum_{m,m',m'',m'''=1}^{M} \sum_{K \in \mathbb{T}_{g,m}} \int_{\partial K \cap \partial V} \hat{e}^T_{g,m'''} [A_o^{-1}]_{m'''m''} B_{m''m'} C^{-1}_{m'm} e_{g,m} \ . \end{aligned}$$

This equality can be used to define a cellwise error quantity $\eta_{K,g,m}$ for each cell K of each of the meshes $\mathbb{T}_{g,m}$:

$$\begin{aligned} \eta_{K,g,m} = & \sum_{g'=1}^{G} \sum_{m'=1}^{M} \int_K [\boldsymbol{\nabla} \hat{e}_{g',m'}]^T (\sigma^o)^{-1}_{g'g,m'm} [\boldsymbol{\nabla} e_{g,m}] \\ & + \sum_{g'=1}^{G} \sum_{m',m'',m'''=1}^{M} \int_K [C^{-1}_{m''m'''} \hat{e}_{g',m'''}]^T \sigma^e_{g'g,m''m'} [C^{-1}_{m'm} e_{g,m}] \\ & + \sum_{m',m'',m'''=1}^{M} \int_{\partial K \cap \partial V} \hat{e}^T_{g,m'''} [A_o^{-1}]_{m'''m''} B_{m''m'} C^{-1}_{m'm} e_{g,m} \ . \end{aligned} \tag{11}$$

These error quantities could then be used to determine whether a cell K on a mesh $\mathbb{T}_{g,m}$ should be refined, coarsened, or left unchanged. Note that we have the equality $J(F) - J(F_h) = \sum_{g=1}^{G} \sum_{m=1}^{M} \sum_{K \in \mathbb{T}_{g,m}} \eta_{K,g,m}$.

III.B. Duality-Based Refinement Indicators

While accurate, the error quantities defined in Eq. (11) have two significant drawbacks: (a) their computation requires that we know the errors $e, \hat{e}$ for which one would have to know the exact primal and dual solutions $F, \hat{F}$; this problem can be addressed by approximating the errors in various different ways (see, for example, Ref. 11) and (b) the formula is long and would be awkward to implement.

Instead of using the full definition of $\eta_{K,g,m}$ and approximating e and $\hat{e}$, we use here a simpler approach. Rather than implementing an error estimator that evaluates approximations of $\eta_{K,g,m}$ as best as possible, let us define a quantity $\tilde{\eta}_{K,g,m}$ that may not be a mathematically exact approximation but that captures the essential features of $\eta_{K,g,m}$. Here, it is important to remember that a local criterion to determine whether a cell should be refined, coarsened, or left unchanged is all that is required. Hence, cells with large error contributions $\eta_{K,g,m}$ must also have large refinement indicators $\tilde{\eta}_{K,g,m}$ and vice versa. To this end, note that each term in $\eta_{K,g,m}$ contains both the primal

and dual errors $e, \hat{e}$. Intuitively, this means that in order to estimate $J(\mathsf{F}) - J(\mathsf{F}_h)$, we have to consider the error e in the primal solution on each mesh cell and also weigh it with a factor that contains the error in the dual solution $\hat{e} = \hat{\mathsf{F}} - \hat{\mathsf{F}}_h$. The equation that $\hat{\mathsf{F}}$ has to satisfy involves $J(\cdot)$ as its right side source term and transports information into the opposite (adjoint) direction from the primal problem. Thus, the weighting factor indicates how important a cell's contribution is toward computing $J(\mathsf{F})$. For simplicity, goal-oriented adaptivity often employs the same adapted meshes for the primal and dual problems; our implementation follows this prescription. We note that the price one pays for abandoning $\eta_{K,g,m}$ in favor of $\tilde{\eta}_{K,g,m}$ is that we will no longer be able to accurately estimate the actual level of error in our computations. On the other hand, if our goal is only the computation of refinement indicators, our results below show that this simplification is certainly still useful.

One way to define refinement indicators is, therefore, to use the product of primal and dual errors, $\tilde{\eta}_{K,g,m} = \|\nabla e_{g,m}\|_K \|\nabla \hat{e}_{g,m}\|_K$. While this still requires knowledge of the exact primal and dual solutions, we can use the following widely used approximation for elliptic problems[8,32]:

$$\tilde{\eta}_{K,g,m} = h_K \|[\![\partial_n \mathsf{F}_{h,g,m}]\!]\|_{\partial K} \|[\![\partial_n \hat{\mathsf{F}}_{h,g,m}]\!]\|_{\partial K} \ , \qquad (12)$$

where h_K is the diameter of cell K on mesh $\mathbb{T}_{g,m}$ and $[\![\partial_n \psi]\!]$ is the jump of the gradient of quantity ψ across the boundary ∂K between cell K and its neighbors.

It is clear that a sharp mathematical bound of the error may have been lost in the derivation of this indicator. As a consequence, we note that refining meshes based on $\tilde{\eta}_{K,g,m}$ could be less efficient than if we had used $\eta_{K,g,m}$; i.e., we expect that more cells will be necessary to achieve a certain accuracy. Nevertheless, as we demonstrate in our numerical results in Sec. IV, our heuristic derivation for a goal-oriented error indicator still leads to more accurate answers than standard AMR techniques. In particular, we show that our approach is more efficient than global (uniform) mesh refinement and more efficient than using a refinement indicator that only takes into account the accuracy in the primal solution and disregards the importance factor derived from the dual solution. The latter constitutes the standard AMR approach, and as a representative of this class of algorithms, we compare our goal-oriented refinement indicator against the one commonly referred to as the Kelly refinement indicator[32] and defined by

$$\tilde{\eta}^{\text{Kelly}}_{K,g,m} = h_K^{1/2} \|[\![\partial_n \mathsf{F}_{h,g,m}]\!]\|_{\partial K} \ . \qquad (13)$$

III.C. Refinement Strategies

Once refinement indicators $\tilde{\eta}_{K,g,m}$ have been obtained using, for example, Eq. (12) (goal-oriented AMR) or Eq. (13) (standard AMR), we have to decide which cells to refine. There are several strategies that are commonly used in the case that the partial differential equation is discretized using only a single mesh. However, the use of different meshes per group and per moment presents some additional complications. In this work, we consider the following two strategies:

1. *Refining meshes individually:* In this strategy, we consider each mesh individually. For a given energy group g and moment m, we then consider the cells $K \in \mathbb{T}_{g,m}$. Let $N_{g,m}$ be the number of cells in $\mathbb{T}_{g,m}$, and let $0 \le f_r, f_c \le 1$ be the refinement and coarsening fractions (with $f_c + f_r \le 1$). Then we will refine those $f_r N_{g,m}$ cells of $\mathbb{T}_{g,m}$ that have the largest indicators $\eta_{K,g,m}$ and coarsen those $f_c N_{g,m}$ cells with the smallest indicators.

2. *Refining meshes jointly:* In this alternate strategy, we consider all cells on all meshes at the same time. Let the total number of cells be $N = \sum_{g=1}^{G} \sum_{m=1}^{M} N_{g,m}$. Then we refine those $f_r N$ cells among all cells on all meshes that have the largest indicators $\eta_{K,g,m}$, and coarsen those $f_c N$ of all cells with the smallest indicators.

Note that when using the first strategy, all meshes will grow by the same amount in each refinement step (because on each mesh a fraction f_r of all cells is refined and a fraction of f_c is coarsened). When employing the second strategy, some group-dependent or moment-dependent meshes may only be refined whereas some others may only be coarsened. In the latter meshes, the solution in these groups and moments is not as important to achieve accuracy in our goal functional as in the former meshes.

The latter strategy has the advantage that it can tell us which meshes (and energy groups/modes) are important by simply keeping the other meshes much coarser and therefore comparatively inexpensive. However, the choice of the simplified estimator Eq. (12) over the exact form Eq. (11) carries the risk that we miss important contributions to the actual error if meshes diverge too far, for example, if some meshes remain unrefined entirely. In particular, a closer mathematical analysis reveals that Eq. (11) contains terms that indicate the error that stems from projecting solutions $\mathsf{F}_{g,m}$ onto meshes $\mathbb{T}_{g',m'}$ when computing the right side terms for the partial differential equation for $\mathsf{F}_{g',m'}$. If meshes are too different, these projection errors may dominate the errors introduced by numerically approximating the $\mathsf{F}_{g,m}$ on $\mathbb{T}_{g,m}$. Unfortunately, the simplified indicator Eq. (12) has no representation for these projection errors, and it may be better to use the first refinement strategy above to ensure that all meshes have roughly the same degree of mesh refinement. We investigate the choice of refinement strategy in Sec. IV below. In particular, one outcome will be that in the cases considered there, the choice is actually not important.

Finally, we point out that for ease of implementation, it may also be desirable to limit the number of adapted meshes. In Ref. 8, it is found that utilizing

group-dependent meshes is clearly beneficial. Here, we have kept this recommendation but investigate whether adapting each moment separately yields important memory/CPU gains.

III.D. Examples of Quantities of Interest

In the final part of this section, we discuss the choice of the goal functional $J(\cdot)$ that constitutes the right side of the adjoint problem Eq. (9). The functional $J(\cdot)$ extracts from a solution F the quantity that for a given simulation we are most interested in computing accurately. For example, we may be interested in averages or point values of the solution; more generally, any linear functional of the solution could be chosen.

By construction, the functional $J(\cdot)$ acts on the solution F. We are typically not interested in the composite moments F but in the scalar flux Φ_0. However, to translate between the two representations, we use the following relation:

$$\Phi_{g,0} = \sum_{m=1}^{M} [C^{-1}]_{0m} F_{g,m} = \sum_{m=1}^{M} (-1)^{m+1} F_{g,m} = F_{g,1} - F_{g,2} + \cdots + (-1)^{M+1} F_{g,M} \ . \quad (14)$$

By the Riesz representation theorem, we can then write every linear functional of the scalar flux in which we may be interested in the form

$$J(F) = \sum_{g=1}^{G} \int_V \hat{Q}_g(\mathbf{r}) \Phi_0(\mathbf{r})\, dr = \sum_{g=1}^{G} \sum_{m=1}^{M} \int_V \hat{Q}_g(\mathbf{r}) (-1)^{m+1} F_{g,m}\, dr \ , \quad (15)$$

where $\hat{Q}$ is a weight function that may also include distributions (delta functions). By recalling that the dual problem has the form $a(b,F) = J(b)\ \forall b$, we can conclude that the source term of the adjoint equation for group g and mode m equals $(-1)^{m+1}\hat{Q}_g(\mathbf{r})$.

Let us now give three prototypical cases of such goal functionals. First, a classical example in nuclear engineering is the detector reading, given by the absorption reaction rate in the detector volume V_{det}:

$$J(F) = \sum_{g=1}^{G} \int_{V_{\text{det}}} \Sigma_a^g(\mathbf{r}) \Phi_{g,0}(\mathbf{r})\, dr = \sum_{g=1}^{G} \sum_{m=1}^{M} \int_{V_{det}} (-1)^{m+1} \Sigma_a^g(\mathbf{r}) F_{g,m}(\mathbf{r})\, dr \ . \quad (16)$$

For this functional, the source term in the adjoint equation for $\hat{F}_{g,m}$ then equals $(-1)^{1+m}\Sigma_a^g(\mathbf{r})\chi_{V_{det}}(\mathbf{r})$, where the characteristic function is defined as $\chi_{V_{det}}(\mathbf{r}) = 1$ if $\mathbf{r} \in V_{det}$ and zero outside the detector volume.

As a second example, let the goal of the calculation be the average scalar flux in group γ and in subdomain ω; i.e., the goal functional is

$$J(F) = \frac{1}{|\omega|} \int_\omega \Phi_{\gamma,0}(\mathbf{r})\, dr = \sum_{g=1}^{G} \sum_{m=1}^{M} \int_\omega \frac{1}{|\omega|} \delta_{g\gamma} (-1)^{m+1} F_{g,m}(\mathbf{r})\, dr \ , \quad (17)$$

where $\delta_{g\gamma} = 1$ if $g = \gamma$ and zero otherwise. Consequently, the sources for the adjoint equation are now $(1/|\omega|)\delta_{g\gamma}(-1)^{m+1}\chi_\omega(\mathbf{r})$; in particular, this implies that only the adjoint equation for group γ has a nonzero source.

Third, if we are interested in pointwise estimates, such as the pointwise scalar flux at $\mathbf{r}_0$ in group γ, the goal functional is

$$J(F) = \Phi_{\gamma,0}(\mathbf{r}_0) = \sum_{g=1}^{G} \sum_{m=1}^{M} \int_V \delta(\mathbf{r} - \mathbf{r}_0) \delta_{g\gamma} (-1)^{m+1} F_{g,m}(\mathbf{r})\, dr \ . \quad (18)$$

Consequently, the adjoint source terms are $(-1)^{m+1}\delta_{g\gamma}\delta(\mathbf{r} - \mathbf{r}_0)$.

Finally, let us note that the formalism also allows for goal functionals defined on odd moments. For example, the net current along the x-direction at position $\mathbf{r}_0$ in group γ can be written as

$$J(F) = \mathbf{\Phi}_{\gamma,1}(\mathbf{r}_0)\cdot\mathbf{e}_x = -(\sigma^1)^{-1}\nabla F_{\gamma,1}(\mathbf{r}_0)\cdot\mathbf{e}_x = -\sum_{g=1}^{G} \sum_{m=1}^{M} \delta_{g\gamma}\delta_{m1} \int_V \delta(\mathbf{r} - \mathbf{r}_0)[\sigma^1]^{-1}\partial_x F_{g,m}(\mathbf{r})\, dr \ . \quad (19)$$

After integration by parts, it is immediately clear that the adjoint source is $(\sigma^1)^{-1}\delta_{g\gamma}\delta_{m1}\nabla\delta(\mathbf{r} - \mathbf{r}_0)\cdot\mathbf{e}_x$. Note that here, as in the point flux evaluation above, no particular difficulties are associated with the fact that adjoint sources may contain Dirac functions or their derivatives. In fact, in the finite element setting, evaluating $J(b)$ for any test function $b = [b_{g,m}(\mathbf{r})]_{g=1\ldots G, m=1\ldots M}$ when the source terms contain a Dirac function (or its derivative) simply translates into the evaluation of the test function (or its derivative) at point $\mathbf{r}_0$.

IV. NUMERICAL RESULTS

In the following, we present a series of numerical experiments that illustrate the adaptive refinement strategies outlined above. With these experiments, we (a) show that the goal-oriented refinement indicator Eq. (12) is superior to the more traditional Kelly refinement indicator

Eq. (13) as well as uniform global refinement, (b) compare the different refinement strategies introduced in Sec. III.C, and (c) investigate whether the use of different meshes for the individual moments $F_{g,m}$ of the same energy group g leads to more efficient solution schemes (it was already demonstrated in Ref. 8 that using separate meshes for different energy groups is beneficial). We investigate these three questions using two 2-D examples in Sec. IV.A., retaining the full generality of using different meshes for different modes. Two further examples illustrate our claims in three dimensions as well.

All results shown here were obtained with a program based on the Open Source finite element library deal.II (Refs. 33 and 34). Unless otherwise mentioned below, we compute solutions using the SP_3 approximation (i.e., $M = 2$, using two different modes) for one-group problems, and we use bi-quadratic (in two dimensions) and tri-quadratic (in three dimensions) finite elements throughout our experiments. However, in example 1 we will also use the SP_5 approximation, and example 4 employs two energy groups.

IV.A. Two-Dimensional Test Cases

For the 2-D test cases, consider the situation depicted in Fig. 1. A source at the left is partially shielded by an absorbing region in the upper center. This could be considered a model for the shielded source problem. In cases like these, we may be interested in radiation levels behind the shielding to determine exposure to shielded components or persons (example one), or just inside the shielding to determine radiation damage to the shielding

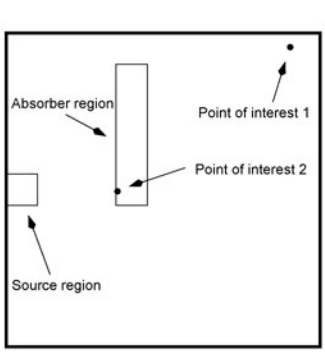

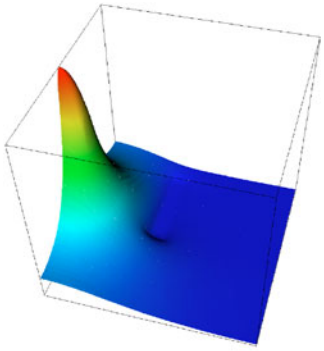

Domain:	$0\text{cm} \le x \le 20\text{cm}$	Points of interest:	$\mathbf{r}_1 = (19.1\text{cm}, 19.1\text{cm})$
	$0\text{cm} \le y \le 20\text{cm}$		$\mathbf{r}_2 = (7.1\text{cm}, 10.1\text{cm})$
Source region:	$0\text{cm} \le x \le 2\text{cm}$	Absorber region:	$7\text{cm} \le x \le 9\text{cm}$
	$9\text{cm} \le y \le 11\text{cm}$		$9\text{cm} \le y \le 18\text{cm}$
	$Q = 1(\text{cm}^2\text{s})^{-1}$		$\Sigma_t = 1\text{cm}^{-1}$
	$\Sigma_t = 1\text{cm}^{-1}$		$\Sigma_s = 0.1\text{cm}^{-1}$
	$\Sigma_s = 0.99\text{cm}^{-1}$	Bulk medium:	$\Sigma_t = 1\text{cm}^{-1}$
			$\Sigma_s = 0.9\text{cm}^{-1}$

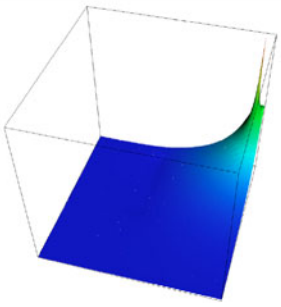

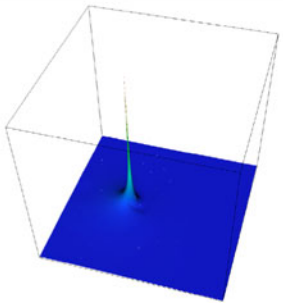

Fig. 1. Two-dimensional test cases, examples 1 and 2. Geometry layout (top left); graph of the solution (top right); material parameters and locations of points of interests (middle row). Adjoint solutions for the two points of interest (bottom row).

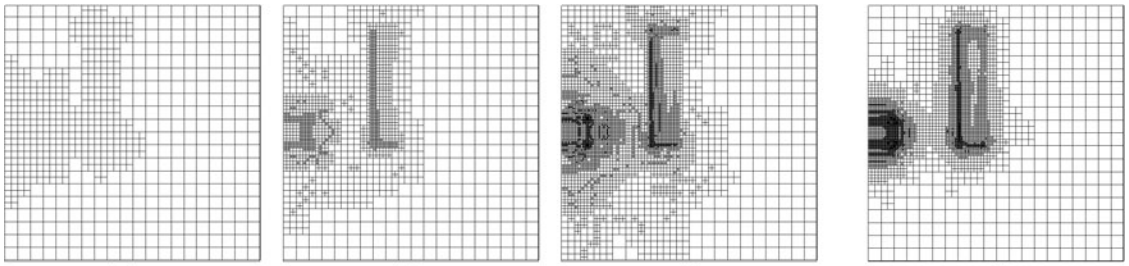

Fig. 2. Example 1: (a), (b), (c) Meshes for F_1 after AMR steps 1, 3, and 5 using the "traditional" Kelly indicator Eq. (13) and (d) mesh for F_2 after five refinement steps.

material (example two). We are, therefore, interested in point values of the flux $\Phi_0(\mathbf{r}_i)$ at one of the two points of interest $\mathbf{r}_1, \mathbf{r}_2$ and employ Eq. (18) as the target functional. The adjoint solutions corresponding to this target functional at the two points of interest are shown in the bottom row of Fig. 1. In all cases, mesh refinement starts from a coarse mesh of 20×20 cells.

Example 1: First Evaluation Point ($\mathbf{r}_1$)

Point $\mathbf{r}_1$ lies in the upper right corner behind the shielding absorber. At this point, the particle flux is much smaller than in the vicinity to the source; in fact, the maximal value at the source is approximately 4.66 whereas $\Phi(\mathbf{r}_1) \approx 0.022$, i.e., smaller by a factor of $\sim$200. It is therefore not surprising that attempts at reducing the error globally will have limited effect; for example, using standard AMR with the Kelly refinement indicator Eq. (13) will mostly refine around the source region and the front part of the absorber, as shown in Fig. 2.[a] Had our goal been to obtain a numerical solution that is accurate everywhere, this would be sensible since these are the areas where the magnitude of the error is in fact larger. However, for our goal of computing $\Phi(\mathbf{r}_1)$, an error that is small everywhere is not required. On the contrary, we need to refine cells where the error is both large and where it is important for our goal. In this case, this will yield a mesh refinement close to the point of interest, as well as on the backside of the absorber, as shown in Fig. 3.

While these meshes make intuitive sense, the success of the method proposed here hinges on whether we can show that it is actually more efficient. To this end, Fig. 4 shows the accuracy in $J(F)$ that can be achieved for a given number of degrees of freedom (or a given amount of compute time) for global uniform refinement, refinement based on the Kelly refinement indicator Eq. (13), and refinement based on the goal-oriented indicator Eq. (12). As can be seen from Fig. 4, the method proposed here can achieve the same accuracy with far fewer cells (by a factor of $>$10) than for the other two refinement methods. In particular, standard refinement by the Kelly indicator proceeds somewhat unpredictably in steps because no refinement is happening at all in the vicinity of the point $\mathbf{r}_1$ where we want to know the solution accurately. For the relatively predictable goal-oriented adaptive and uniform refinement strategies, the error $|J(F) - J(F_h)|$ as a function of N, the overall number of unknowns in the problem, can be computed to be $\mathcal{O}(N^{-1.5})$ and $\mathcal{O}(N^{-1.1})$, respectively, confirming the better convergence order of the goal-oriented adaptive approach.

Figure 4b shows that the savings in the number of cells immediately translate into corresponding savings in compute time despite the fact that in the goal-oriented approach, we now have to solve both a primal problem and a dual problem (a fact that can be seen for the first data point where our method requires twice as much time as the others).

Finally, we have investigated which of the two strategies outlined in Sec. III.C for refining the two meshes used in this problem is more efficient. We found that the differences are marginal and that the curves would be indistinguishable if plotted in Fig. 4. The reason is that even for the strategy that globally determines which cells to refine (comparing cells from different meshes), the number of cells on the two meshes differs by $<$20%. One can conclude that the magnitude of error indicators for both meshes is approximately equal, indicating that both moments are important to accurately compute $J(F)$.

In order to determine whether this is because we have only two modes, we reran the same test case but with the SP_5 approximation (three angular modes). There, the second refinement strategy generated meshes that, on the finest level, had 1 635 627, 1 357 471, and 1 265 128 degrees of freedom for F_1, F_2, F_3, respectively. Yet, this

[a]The meshes shown here contain refined islands and other artifacts that are due to the cutoff strategy employed (fractions f_r and f_c) and do not contribute significantly to a more accurate solution. However, their presence is also not harmful, and they are mostly a visual nuisance. If desired, their generation can easily be suppressed in deal.II, leading to "smoother" adapted meshes.

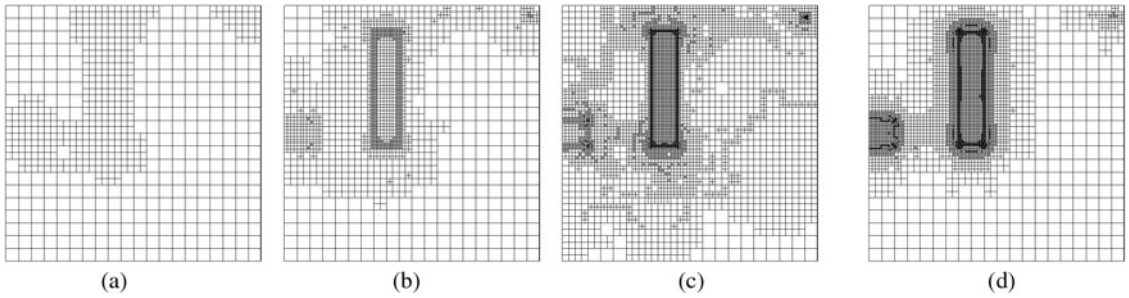

Fig. 3. Example 1: (a), (b), (c) Meshes for F_1 after AMR steps 1, 3, and 5 using the goal-oriented indicator Eq. (12) and (d) mesh for F_2 after five refinement steps.

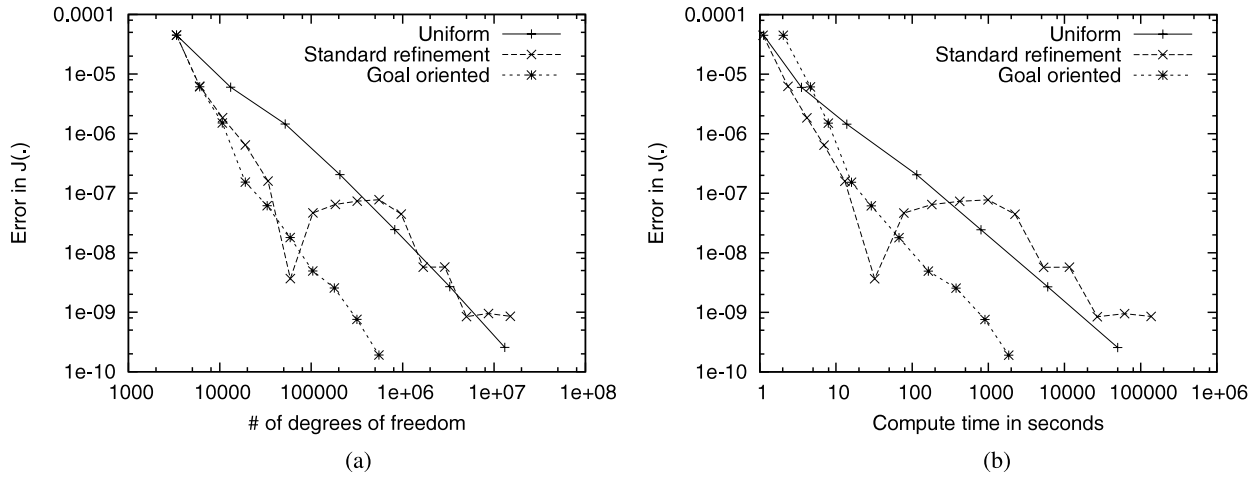

Fig. 4. Example 1: Absolute error $|J(F) - J(F_h)|$ as a function of (a) the number of degrees of freedom and (b) CPU time. For comparison, note that $J(F) \approx 0.02189817737$.

strategy performed only marginally better than the first one outlined in Sec. III.C (the error was some 10 to 20% smaller for the same number of degrees of freedom). It is again apparent that the choice of mesh adaptation strategy is not crucial for the moments.

Example 2: Second Evaluation Point ($\mathbf{r}_2$)

While the test case of the previous section may seem to be made to benefit the goal-oriented method proposed in this work, this second example shows that it is universally better than more traditional AMR approaches. To this end, let us consider the same setup as in example 1 but with an evaluation point that lies just inside the absorbing region, i.e., in a region where the standard AMR approach would preferentially refine mesh cells. Figure 5 shows the meshes produced by the goal-oriented refinement indicator, while the meshes generated by the Kelly indicator of course remain unchanged from Fig. 2 since this refinement criterion does not depend on the quantity of interest. We see that again the goal-oriented indicator emphasizes refinement around the point of interest and neglects refinement of most of the rest of the boundary layer around the absorber region as that is not important for our current goal.

Even though one could expect that the Kelly indicator should perform well for this example, Fig. 6 shows that the proposed method again does significantly better after the first few refinement iterations, producing results that are orders of magnitude better than the ones obtained using the two other methods for the same number of unknowns, or order of magnitude faster for the same error. In fact, goal-oriented adaptive refinement performs even better than for the first test case, with a convergence order $|J(F) - J(F_h)| \approx \mathcal{O}(N^{-2.5})$ as compared to $\mathcal{O}(N^{-1.1})$ for global uniform refinement. Standard AMR and goal-oriented AMR produce equivalent error reduction as long as standard AMR refines in the regions of importance toward the computational goal; after several adaptations, standard AMR refines in regions that bear no importance toward the goal, hence leading to a stagnation in the error, as seen in Fig. 6.

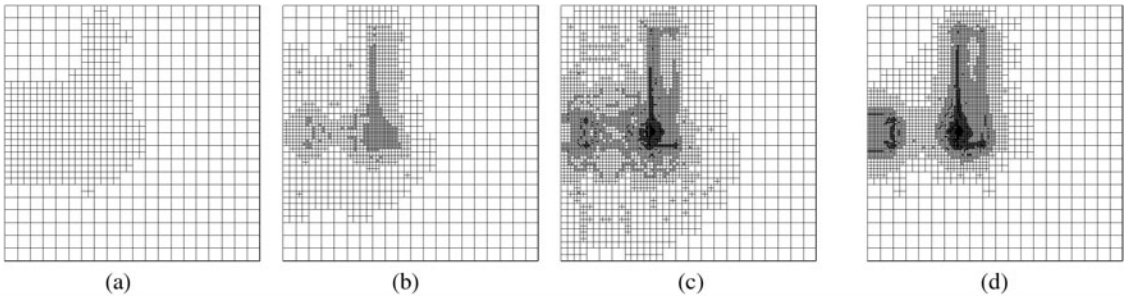

Fig. 5. Example 2: (a), (b), (c) Meshes for F_1 (bottom row) after AMR steps 1, 3, and 5 using the goal-oriented indicator Eq. (12) and (d) for F_2 after five refinement steps.

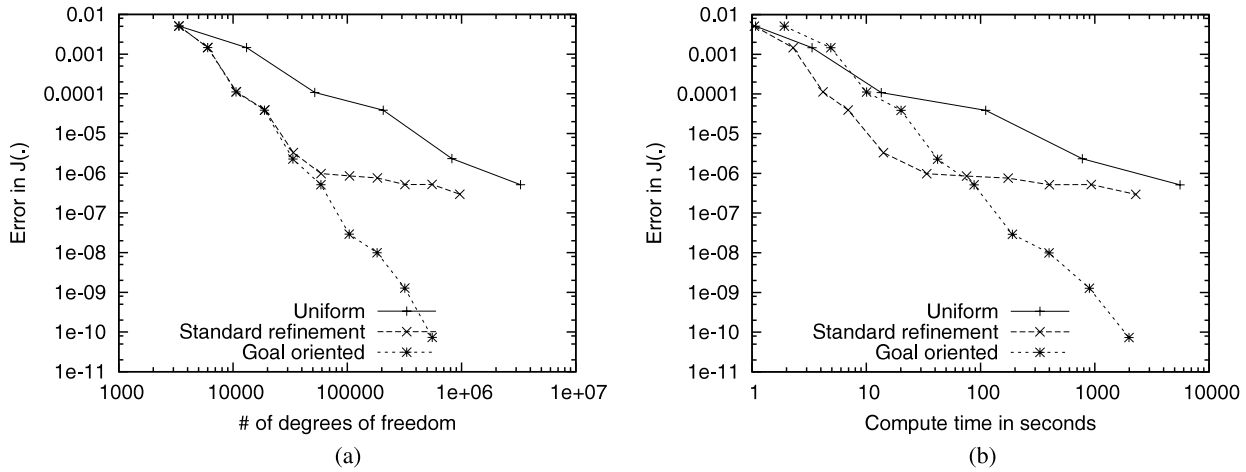

Fig. 6. Example 2: Absolute error $|J(F) - J(F_h)|$ as a function of (a) the number of degrees of freedom and (b) CPU time. For comparison, note that $J(F) \approx 0.19345929975$.

As for the first example, we have compared the two refinement strategies of Sec. III.C and have found no significant difference. The number of degrees of freedom on the two meshes again differed by <20% for this test case.

IV.B. Three-Dimensional Test Cases

For our 3-D test case, we generalize the setting of Fig. 1 by adding a third spatial dimension. The source is located in the area 0 cm $\le x \le$ 2 cm, 9 cm $\le y \le$ 11 cm, and 9 cm $\le z \le$ 11 cm at the center of a face of the domain $[0 \text{ cm}, 20 \text{ cm}]^3$. The absorber is a hexahedron in the area 7 cm $\le x, z \le$ 9 cm, and 9 cm $\le y \le$ 18 cm. In the multigroup case, the source produces particles only in the first energy group (fast particles).

Figure 7a shows isocontours of the solution of example 3 below, along with the locations of the source and absorber. To present as wide a range of examples as possible, for the 3-D examples we consider the situation where we are interested not in the flux, but in the three components of the current, i.e., $\mathbf{\Phi}_{\gamma,1}(\mathbf{r}_0) = -(\sigma^1)^{-1}\nabla F_{\gamma,1}(\mathbf{r}_0)$ with a point of interest at $\mathbf{r}_1 =$ (19.1 cm, 19.1 cm, 19.1 cm); see Eq. (19). For the one-group computation of example 3, $\gamma = 1$ is obviously the only choice. For the two-group calculation of example 4, we choose $\gamma = 2$ as the energy group of interest (whereas the source emits only in group 1).

In this situation where the three components of the current are quantities of interest, we have multiple goal functionals. However, in the goal-oriented approach, we need only a single functional as the right side of the dual problem. To this end, we here use $J(F) = J_x(F) + J_y(F) + J_z(F)$, where we let $J_u(F)$ be the goal functional for the current in direction u given by Eq. (19). Strictly speaking, using this functional, meshes will then be adjusted to compute $J(\cdot)$ with maximal accuracy, not any of its components. In the examples below, we nevertheless show that the meshes are well suited to compute each of its components, in particular, $J_x(\cdot)$, for which we show convergence graphs. We note that in general, computing currents accurately is more difficult than computing fluxes

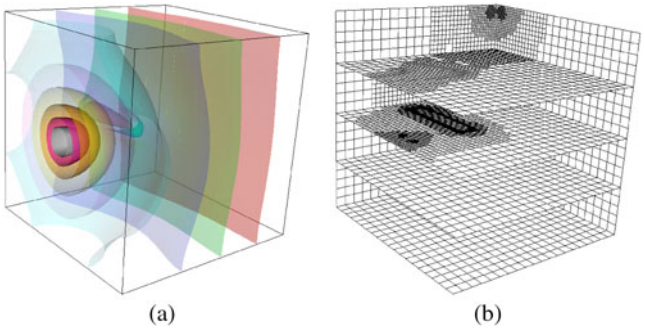

Fig. 7. Example 3: (a) Iso-surfaces of the solution separated by factors in flux of approximately 1.6. (b) Mesh when using the same mesh for all moments and after three cycles of using Eq. (12) for refinement.

since it requires evaluating derivatives of the solution. This can also be seen in the equation for the dual solution, whose source is now the derivative of a Dirac delta function, leading to a highly singular dual solution. Our results show that we can efficiently handle this situation as well.

Example 3: One-Group SP_3

For this test case, we use a single energy group with material properties chosen as in the 2-D examples. For the chosen source and region of interest, we expect the goal-oriented mesh to resolve only that part of the radiation that reaches the point of interest coming from the source. Figure 7b confirms this, showing for example that the goal-oriented AMR mesh does not resolve any radiation going downward at all (in the lower half the mesh is essentially the coarse $20 \times 20 \times 20$ initial mesh).

In our 2-D experiments, we have seen (see Figs. 3c, 3d, 5c, and 5d) that the adapted meshes for the first and second composite moments F_1, F_2 are not significantly different. This can be understood by realizing that the two variables both satisfy diffusion equations with the similar diffusion and absorption coefficients; the regularity of solutions is, therefore, quite similar, unlike the case of different energy groups (see Ref. 8). Likewise, because all composite moments enter in the definition of the scalar flux Φ_0 (and consequently in the goal functional) with equal absolute weight [see, for example, Eq. (18)], no moment is necessarily more important than others. Hence, it is justified that using the same mesh for all moments does not make the method notably less efficient than employing moment-dependent meshes, though it would simplify the implementation significantly.

To this end, Fig. 8 compares the reduction in error in the functional $J_x(\cdot)$ for uniform refinement, goal-oriented refinement using independent meshes for different angular modes, and goal-oriented refinement where we use a single mesh for all modes. First, we note that goal-oriented refinement is again much more efficient. Second, we can confirm our conjecture that using a single adapted mesh is sufficient for all moments. This is also in line with our observation in the 2-D case.

Example 4: Multigroup SP_3

In this final example, we consider a two-group SP_3 case. We choose the same general setup as in example 3, with material properties given in Table I (no upscattering and only isotropic scattering for simplicity).

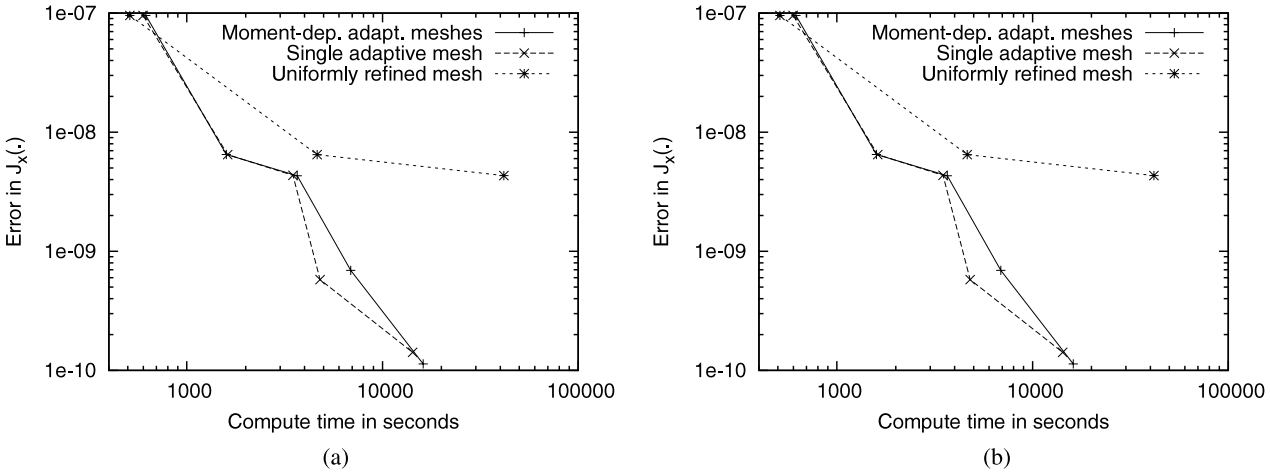

Fig. 8. Example 3: Absolute error $|J_x(F) - J_x(F_h)|$ as a function of (a) the number of degrees of freedom and (b) CPU time. Mesh refinement proceeds for a sequence of uniformly refined meshes; adaptively refined meshes that are used for all modes simultaneously; and adaptively refined meshes where each mode has its own, separately refined mesh. For comparison, note that $J_x(F) \approx 0.0001008117$.

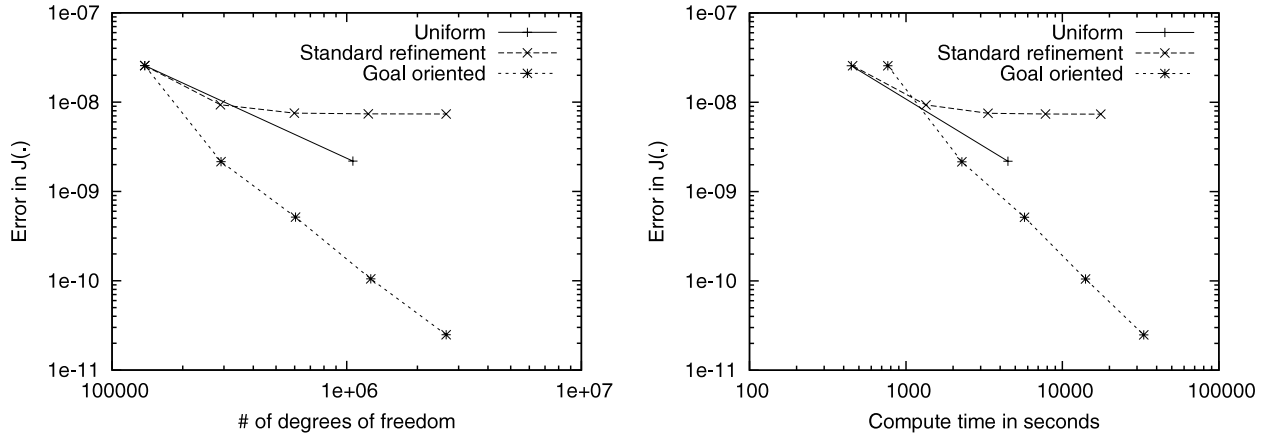

Fig. 9. Example 4: Absolute error $|J_x(F) - J_x(F_h)|$ as a function of (a) the number of degrees of freedom and (b) CPU time. For comparison, note that $J_x(F) \approx 0.00002840484$.

TABLE I

Two-Group Material Properties for Example 4

Region	Σ_t^1	Σ_t^2	$\Sigma_{s,0}^{1\to1}$	$\Sigma_{s,0}^{2\to2}$	$\Sigma_{s,0}^{1\to2}$
Source	0.35	1	0.3	0.9	0.01
Absorber	0.4	1	0.1	0.1	0.1
Bulk medium	0.4	1	0.34	0.99	0.01

*Data in cm^{-1}

Following our discussion above, we here employ the same mesh for all angular moments of the solution, although we use different meshes for different energy groups. For this case, Fig. 9 shows the error in $J_x(\cdot)$ as a function of degrees of freedom and CPU time for the refinement methods already used in previous examples. It is again apparent that goal-oriented adaptive meshes are much more efficient than global or traditional AMR. In particular, it is even more obvious that traditional adaptivity is actually counterproductive in this case, as has already been observed in example 2 (see Fig. 6).

V. CONCLUSIONS

As a proof of concept, we have applied the paradigm of goal-oriented mesh adaptivity to the multigroup SP_N equations to obtain accurate local quantities of interest. Goal-oriented adaptivity combines the benefits of standard AMR with knowledge about the end goal of the computation. Hence, the adaptivity proceeds by refining mesh cells whose errors are large and whose importance toward the computational goal is high. As we have shown, the resulting calculations are much more efficient than either global uniform mesh refinement or traditional mesh adaptivity approaches. In particular, in our experiments, we have observed improvements in the number of degrees of freedom (and consequently in memory consumption) as well as in CPU time of factors of 10 and more, opening the door to simulations that were previously impossible. The error indicators used to drive the goal-oriented adaptivity are, in essence, similar to the ones used for standard adaptivity, except that the product of the error indicators for both a primal and a dual (adjoint) problem is employed. Appropriate adjoint source terms have been defined for various quantities of interest. The cost of solving an adjoint problem is easily absorbed as the adaptivity proceeds.

As part of our investigation, we have also tested different strategies for mesh refinement. In particular, we have noted, at least for the test cases considered, that it is not important whether we refine meshes individually or jointly, in the latter case comparing refinement indicators of cells from different meshes. Likewise, for one-group calculations, it was sufficient to use only a single adapted mesh for all moments rather than using a different adapted mesh per moment. Owing to these observations—and to simplify the algorithmic implementation—multigroup calculations have been carried out using one adapted mesh for all moments belonging to the same energy group; as previously noted in Ref. 8, it is nonetheless important for efficiency that different meshes be used for different energy groups.

ACKNOWLEDGMENTS

Part of this research was funded through U.S. Department of Energy grants DE-FG07-07ID14767 and DE-FG07-05ID14692. Part of this research is based upon work supported

by the U.S. Department of Homeland Security under grant award No. 2008-DN-077-ARI001-02. This publication is also partly based on work supported by award No. KUS-C1-016-04, made by the King Abdullah University of Science and Technology. W.B. is supported by an Alfred P. Sloan Research Fellowship. The authors gratefully acknowledge these sources of support.

REFERENCES

1. R. VERFÜRTH, *A Review of A Posteriori Error Estimation and Adaptive Mesh Refinement Techniques*. Wiley/Teubner, New York, Stuttgart (1996).

2. E. RAMM, E. RANK, R. RANNACHER, K. SCHWEIZERHOF, E. STEIN, W. WENDLAND, G. WITTUM, PETER WRIGGERS, WALTER WUNDERLICH, and ERWIN STEIN, *Error-Controlled Adaptive Finite Elements in Solid Mechanics*, John Wiley and Sons, New York (2001).

3. J. P. JESSEE, W. A. FIVELAND, L. H. HOWELL, P. COLELLA, and R. B. PEMBER, "An Adaptive Mesh Refinement Algorithm for the Radiative Transport Equation." *J. Comput. Phys.*, **139**, *2*, 380 (1998).

4. C. AUSSOURD, "Styx: A Multidimensional AMR S_N Scheme," *Nucl. Sci. Eng.*, **143**, 281 (2003).

5. H. ZHANG and E. E. LEWIS, "An Adaptive Approach to Variational Nodal Diffusion Problems," *Nucl. Sci. Eng.*, **137**, 14 (2001).

6. J. RAGUSA, "A Simple Hessian-Based 3D Mesh Adaptation Technique with Applications to the Multigroup Diffusion Equations," *Ann. Nucl. Energy*, **35**, *11*, 2006 (2008).

7. Y. WANG and J. RAGUSA, "Application of hp Adaptivity to the Multigroup Diffusion Equations," *Nucl. Sci. Eng.*, **161**, 22 (2009).

8. Y. WANG, W. BANGERTH, and J. RAGUSA, "Three-Dimensional *h*-Adaptivity for the Multigroup Neutron Diffusion Equations," *Prog. Nucl. Energy*, **51**, 543 (2009).

9. Y. WANG and J. RAGUSA, "Multi-Dimensional Automatic Mesh Refinement for the Multigroup SP_N Equations," *Proc. Joint Int. Topl. Mtg. Mathematics & Computations and Supercomputing in Nuclear Applications (M&C + SNA 2007)*, Monterey, California, April 15–19, 2007, American Nuclear Society (2006) (CD-ROM).

10. M. AINSWORTH and J. T. ODEN, *A Posteriori Error Estimation in Finite Element Analysis*, John Wiley and Sons, New York (2000).

11. W. BANGERTH and R. RANNACHER, *Adaptive Finite Element Methods for Differential Equations*, Birkhäuser Verlag, Basel, Switzerland (2003).

12. I. BABUŠKA and T. STROUBOULIS, *The Finite Element Method and Its Reliability*, Clarendon Press, New York (2001).

13. D. BRAESS and R. VERFÜRTH, "A Posteriori Error Estimators for the Raviart-Thomas Element," *SIAM J. Numer. Anal.*, **33**, 2431 (1996).

14. K. ERIKSSON, D. ESTEP, P. HANSBO, and C. JOHNSON, "Introduction to Adaptive Methods for Differential Equations," *Acta Numerica*, 105 (1995).

15. R. BECKER and R. RANNACHER, "Weighted A Posteriori Error Control in FE Methods," *Proc. 1st European Conf. Numerical Mathematics and Advanced Applications (ENUMATH 95)*, Paris, France, September 18–22, 1995, p. 621, World Scientific Publishing; published in *Proc. 2nd European Conf. Numerical Mathematics and Advanced Applications (ENUMATH 97)*, Heidelberg, Germany, September 28–October 3, 1997, World Scientific Publishing (1998).

16. J. PERAIRE and A. T. PATERA, "Bounds for Linear-Functional Outputs of Coercive Partial Differential Equations: Local Indicators and Adaptive Refinement," *Proc. Workshop Advances in Adaptive Computational Methods in Mechanics*, Cachan, France, September 19, 1997, p. 199, Elsevier (1998).

17. R. BECKER and R. RANNACHER, "A Feed-Back Approach to Error Control in Finite Element Methods: Basic Analysis and Examples," *East-West J. Numer. Math.*, **4**, 237 (1996).

18. R. BECKER and R. RANNACHER, "An Optimal Control Approach to Error Estimation and Mesh Adaptation in Finite Element Methods," *Acta Numerica*, **10**, 1 (2001).

19. A. HAGHIGHAT and J. C. WAGNER, "Monte Carlo Variance Reduction with Deterministic Importance Functions," *Prog. Nucl. Energy*, **42**, 25 (2003).

20. S. A. TURNER and E. W. LARSEN, "Automatic Variance Reduction for Three-Dimensional Monte Carlo Simulations by the Local Importance Function Transform—II: Numerical Results," *Nucl. Sci. Eng.*, **127**, 36 (2007).

21. E. W. LARSEN, G. THÖMMES, A. KLAR, M. SEAID, and T. GÖTZ, "Simplified P_N approximations to the Equations of Radiative Heat Transfer and Applications," *J. Comput. Phys.*, **183**, 652 (2002).

22. Y. WANG, J. RAGUSA, and M. DEHART, "Xuthos: A Discontinuous Galerkin Transport Solver for Newt Based on Unstructured Triangular Meshes," *Proc. Int. Conf. Physics of Reactors: Nuclear Power: A Sustainable Resource (PHYSOR 2008)*, Interlaken, Switzerland, September 14–19, 2008.

23. H. PARK and C. R. E. de OLIVEIRA, "Coupled Space-Angle Adaptivity for Radiation Transport Calculations," *Nucl. Sci. Eng.*, **161**, 216 (2009).

24. V. HEUVELINE and R. RANNACHER, "Adaptive FEM for Eigenvalue Problems with Application in Hydrodynamic Stability Analysis," *J. Numer. Math.*, 1 (2006).

25. K. A. CLIFFE, E. J. C. HALL, and P. HOUSTON, "Adaptive Discontinuous Galerkin Methods for Eigenvalue Problems Arising in Incompressible Fluid Flows," *SIAM J. Sci. Comput.*, **31**, 4607 (2010).

26. G. C. POMRANING, "Asymptotic and Variational Derivations of the Simplified P_N Equations," *Ann. Nucl. Energy*, **20**, 623 (1993).

27. E. W. LARSEN, J. E. MOREL, and J. M. McGHEE, "Asymptotic Derivation of the Multigroup P_1 and Simplified P_N Equations with Anisotropic Scattering," *Nucl. Sci. Eng.*, **123**, 328 (1996).

28. P. S. BRANTLEY and E. W. LARSEN, "The Simplified P_3 Approximation," *Nucl. Sci. Eng.*, **134**, 1 (2000).

29. E. E. LEWIS and G. PALMIOTTI, "Simplified Spherical Harmonics in the Variational Nodal Method," *Nucl. Sci. Eng.*, **126**, 48 (1997).

30. J. J. LAUTARD, D. SCHNEIDER, and A. BAUDRON, "Mixed-Dual Methods for Neutronic Reactor Core Calculations in the Cronos System," In *Proc. Int. Conf. Mathematics and Computation, Reactor Physics and Environmental Analysis of Nuclear Systems*, Madrid, Spain, 1999 (CD-ROM).

31. G. F. CAREY, *Computational Grids: Generation, Adaptation and Solution Strategies*, Taylor & Francis, London (1997).

32. D. W. KELLY, J. P. de S. R. GAGO, O. C. ZIENKIEWICZ, and I. BABUŠKA, "A Posteriori Error Analysis and Adaptive Processes in the Finite Element Method: Part I–Error Analysis," *Int. J. Numer. Method. Eng.*, **19**, 1593 (1983).

33. W. BANGERTH, R. HARTMANN, and G. KANSCHAT, "deal.II: A Finite Element Differential Equations Analysis Library," Technical Reference (2008); http://www.dealii.org/.

34. W. BANGERTH, R. HARTMANN, and G. KANSCHAT, "deal.II—A General-Purpose Object-Oriented Finite Element Library," *ACM Trans. Math. Software*, **33**, *4*, Article 24 (2007).