ASME 2016 Summer Heat Transfer Conference
SAND2016-3057C
HT2016
July 10-14, 2016, Washington, DC, USA

HT2016-7010

REDUCED ORDER MODELING APPLIED TO THE DISCRETE ORDINATES METHOD FOR RADIATION HEAT TRANSFER IN PARTICIPATING MEDIA

John Tencer
Sandia National Laboratories
Albuquerque, NM, USA

Kevin Carlberg
Sandia National Laboratories
Livermore, CA, USA

Roy Hogan
Sandia National Laboratories
Albuquerque, NM, USA

Marvin Larsen
Sandia National Laboratories
Albuquerque, NM, USA

ABSTRACT

Radiation heat transfer is an important phenomenon in many physical systems of practical interest. When participating media is important, the radiative transfer equation (RTE) must be solved for the radiative intensity as a function of location, time, direction, and wavelength. In many heat transfer applications, a quasi-steady assumption is valid. The dependence on wavelength is often treated through a weighted sum of gray gases type approach. The discrete ordinates method is the most common method for approximating the angular dependence. In the discrete ordinates method, the intensity is solved exactly for a finite number of discrete directions, and integrals over the angular space are accomplished through a quadrature rule. In this work, a projection-based model reduction approach is applied to the discrete ordinates method. A small number or ordinate directions are used to construct the reduced basis. The reduced model is then queried at the quadrature points for a high order quadrature in order to inexpensively approximate this highly accurate solution. This results in a much more accurate solution than can be achieved by the low-order quadrature alone. One-, two-, and three-dimensional test problems are presented.

INTRODUCTION

The steady-state gray radiation transport equation (RTE) with isotropic scattering is given by

$$\vec{\Omega}\cdot\vec{\nabla}I(\vec{\Omega})+\sigma_T I(\vec{\Omega})=\kappa\frac{\sigma T^4}{\pi}+\frac{\sigma_s}{4\pi}\int I\left(\vec{\tilde{\Omega}}\right)d\vec{\tilde{\Omega}}. \qquad (1)$$

The boundary conditions for Eq. 1 are Dirichlet; they specify the outgoing intensity at a surface to be equal to the sum of the surface emission and the reflected intensity. This equation defines the radiative intensity as a function of 5 independent variables: 3 in space and 2 in angle. The inputs that parameterize this equation are the absorption and scattering coefficients and the emission source (or alternatively the temperature). The intensities in different directions are only coupled through the right-most (in-scattering) term in Eq. 1 and reflective boundary conditions.

Model-reduction techniques are useful for decreasing the computational cost of many-query problems and are increasingly popular in the areas of optimal design, optimal control, uncertainty quantification, and inverse problems. At first glance, the solution of the RTE does not appear to belong to this class of problems. However, if the angular discretization is performed according to the discrete ordinates method, the angular coordinates may be viewed as independent parameters that must be sampled in much the same way as the parameters in the previously mentioned applications.

NOMENCLATURE

I	is the radiative intensity
I_n	is the angular intensity at quadrature point n
σ_T	is the macroscopic total cross section or extinction coefficient
w_n	is the nth quadrature weight
σ_s	is the macroscopic scattering cross section or scattering coefficient
κ	is the macroscopic absorption cross section or absorption coefficient

σ is the Stefan-Boltzmann constant
T is the material temperature
I_b is the black-body intensity $I_b = \frac{\sigma T^4}{\pi}$
$\vec{\Omega}_i$ is a unit vector pointing in the ordinate direction corresponding to quadrature point i
ε is the surface emissivity
$\vec{n}$ is the surface normal unit vector
$\bar{\bar{\phi}}$ is the reduced basis

The discrete ordinates method is a common approach for representing the angular dependence of the radiative intensity. In the discrete ordinates method, the RTE is satisfied along a set of discrete directions and a quadrature rule is used to evaluate integrals over angle. The choice of quadrature rule is somewhat problem dependent and there are many options available [1-12]. The choice of quadrature rule defines a set of N directions, $\vec{\Omega}_i$ and weights, w_i $i = 1,2,\cdots,N$. Eq. 1 is then approximated as a set of N first-order PDEs in only the three spatial dimensions.

$$\vec{\Omega}_i \cdot \vec{\nabla} I_i + \sigma_T I_i = S \qquad i = 1,2,\cdots,N \tag{2}$$

The source, S is generally the sum of the emission and in-scattering terms. This couples the solution in each ordinate direction with that in every other ordinate direction. This added complexity may be avoided by the use of scattering source iteration [7, 13] in which Eq. 2 is solved repeatedly with different source distributions via fixed point iteration. The uncollided intensity is given by:

$$\begin{aligned} \vec{\Omega}_i \cdot \vec{\nabla} I_i^0 + \sigma_T I_i^0 &= \kappa \frac{\sigma T^4}{\pi} && i = 1,2,\cdots,N \\ I_i^0 &= \varepsilon \frac{\sigma T_w^4}{\pi} && \vec{n} \cdot \vec{\Omega}_i > 0, \end{aligned} \tag{3}$$

while, the nth-time collided intensity (with n>0) is given by:

$$\begin{aligned} &\vec{\Omega}_i \cdot \vec{\nabla} I_i^n + \sigma_T I_i^n = \frac{\sigma_S}{4\pi} \sum_{j=1}^{N} w_j I_j^{n-1} \qquad i = 1,2,\cdots,N \\ &I_i^n = (1-\varepsilon) \sum_{\vec{n}\cdot\vec{\Omega}_j<0} \left|\vec{n} \cdot \vec{\Omega}_j\right| w_j I_j^{n-1} \qquad \vec{n} \cdot \vec{\Omega}_i > 0. \end{aligned} \tag{4}$$

Thus, each source iteration can be viewed as one set of collisions. The solution to Eq. 2 is then approached as n approaches infinity.

$$I_i = \sum_{j=0}^{\infty} I_i^j \qquad i = 1,2,\cdots,N \tag{5}$$

For a given source distribution S, Eq. 2 may be solved independently for each ordinate direction. These solutions are all similar, varying only in the direction parameter, $\vec{\Omega}_i$. For 1D slab geometry problems, this parameter space is one-dimensional as only the component of the direction of travel aligned with the spatial dimension enters into the governing equation. For 2D and 3D problems, this parameter space is two-dimensional. Any point on the unit sphere (or hemisphere in the 2D case) may be described by two variables. The goal of the parameterized model reduction strategy used is to generate a reduced basis for possible solution vectors (i.e. intensity distributions) across the parameter (angular) space. The first step is to discretize Eq. 2. In this paper, the spatial discretization is accomplished through linear finite elements. The discrete problem is given by

$$\bar{\bar{K}}(\vec{\Omega})\vec{I}(\vec{\Omega}) = \vec{S}, \tag{6}$$

where $\bar{\bar{K}}$ is an m x m matrix where m is the number of degrees of freedom in the discretized problem. For a linear finite element discretization m is the number of nodes. Eq. 6 is linear with respect to both the unknown intensity as well as the directional parameter. The existence of this reduced basis may be confirmed by sampling the parameter space and computing the dimensionality of the span of the solution vectors generated. For a fixed source distribution, for several values of the directional parameter, $\overrightarrow{\Omega_1}, \overrightarrow{\Omega_2}, \cdots, \overrightarrow{\Omega_K}$ the solution vector $\overrightarrow{I_1}, \overrightarrow{I_2}, \cdots, \overrightarrow{I_K}$ is computed. These solution vectors are then concatenated to construct the global m x K basis matrix.

$$\bar{\bar{M}} = \left[\overrightarrow{I_1}, \overrightarrow{I_2}, \cdots, \overrightarrow{I_K}\right] \tag{7}$$

If a large number of samples K are taken, it is likely that the column rank of $\bar{\bar{M}}$ is less than K. To address this situation, the singular value decomposition (SVD) of $\bar{\bar{M}}$ is generated [15] yielding 3 matrices $\bar{\bar{U}}$, $\bar{\bar{S}}$, and $\bar{\bar{V}}$ where $\bar{\bar{M}} = \bar{\bar{U}}\bar{\bar{S}}\bar{\bar{V}}^T$. $\bar{\bar{S}}$ is a diagonal matrix containing the singular values and $\bar{\bar{U}}$ is a full matrix containing the modes as its columns. The presence of a reduced basis may be confirmed by examining the decay of these singular values. In the 1D example cases, 99.9999% of the energy is captured by the first 5 or 6 modes depending on the source distribution. In the 2D example cases, capturing the same fraction of the total energy requires between 10 and 35 modes depending on the source distribution.

Fig. 1 shows the decay for a number of different cases. The precise behavior of the singular values is problem dependent, however, the qualitative behavior appears to be similar for problems of the same spatial dimension. The rapid decay in singular values suggests that a relatively small number of modes is likely to be sufficient to capture the vast majority of the solution behavior.

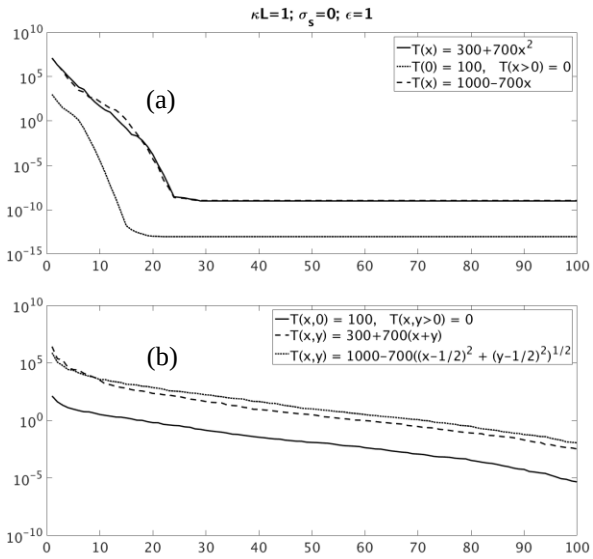

Fig. 1 Decay of singular values demonstrates the presence of a reduced order basis for the intensity solution for several source distributions. (a) 20 modes are sufficient to capture all behavior in the 1D geometry. (b) 50 modes would capture a very large percentage of the solution behavior in the 2D geometry.

As a result, only the primary modes of $\bar{\bar{M}}$ corresponding to the first k<K columns of $\bar{\bar{U}}$ are used. This defines the reduced basis, $\bar{\bar{\phi}}$ through proper orthogonal decomposition (POD). The discretized intensity in any direction is then approximated as

$$\vec{I}(\vec{\Omega}) \approx \bar{\bar{\phi}}\vec{x} \tag{8}$$

where $\vec{x}$ is an n-dimensional vector and $\bar{\bar{\phi}}$ is a tall skinny m x k matrix. The reduced order model (ROM) is then given by substituting Eq. 8 into the full-order model (Eq. 6).

$$\left(\bar{\bar{K}}(\vec{\Omega})\bar{\bar{\phi}}\right)\vec{x} = \vec{S} \tag{9}$$

Eq. 9 is over-determined, i.e., a solution may not exist. Instead, a unique solution can be computed by enforcing the Galerkin orthogonality condition or by applying least-squares Petrov–Galerkin projection [29], which minimizes the residual and solves

$$\left(\bar{\bar{K}}(\vec{\Omega})\bar{\bar{\phi}}\right)^T\left(\bar{\bar{K}}(\vec{\Omega})\bar{\bar{\phi}}\right)\vec{x} = \left(\bar{\bar{K}}(\vec{\Omega})\bar{\bar{\phi}}\right)^T\vec{S} \tag{10}$$

1-D EXAMPLE

Consider the 1D case of a purely absorbing slab surrounded by black walls. The temperature profile is assumed quadratic, $T(x) = 300 + 700x^2$. For 1D geometries, the Gauss-Legendre quadrature rules are typically employed. Consider the angle-integrated intensity generated by the Gauss-Legendre quadrature with 6 directions (S_6). The error in this case is approximately 2%. Contrast this with the distribution generated by the Legendre-Gauss quadrature with 200 directions (S_{200}) which has an expected error of about 0.001%. The S_6 model will be referred to as the low-order model (LOM). The S_{200} model will be referred to as the high-order model (HOM). The intensities from the S_6, LOM solution may be used to generate a ROM capable of predicting intensities in any arbitrary direction. If this ROM is then used to generate intensities for the S_{200} quadrature, the resulting error in the angle-integrated intensity is less than 0.3%. The convergence of the LOM and ROM solutions as the LOM quadrature order is increased is plotted in Fig 2.

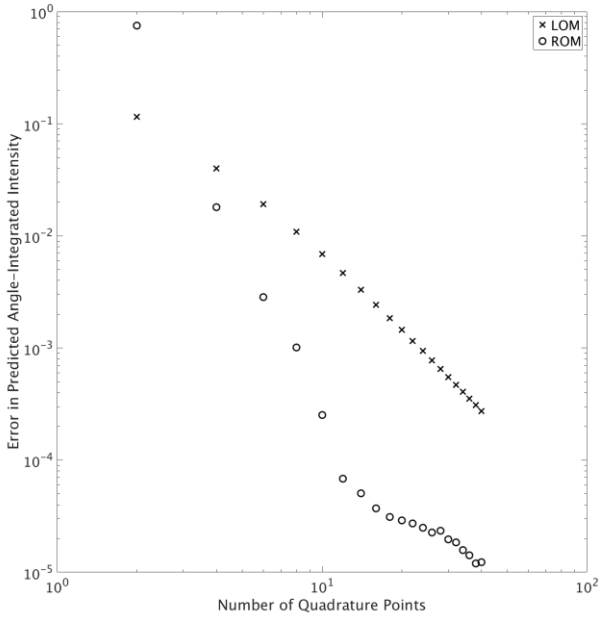

Fig. 2 Convergence of angle-integrated intensity distributions for quadratic temperature profile with increasing LOM quadrature order.

The results in Fig 2 do not tell the whole story since the ROM is both more accurate and necessarily more expensive than the LOM with the same number of quadrature points. Fig 3 shows the comparison of execution time and accuracy for each method. The trends in Fig 3 only hold up for sufficiently large problems (m>>k) otherwise model reduction provides little advantage. Often it is possible in 1D problems to get away with an extremely coarse mesh such that this condition is not satisfied. However, this condition is almost always true for any practical 2D or 3D problem. These results show a great deal of promise for ROMs to reduce the computational expense and increase the accuracy of discrete ordinates simulations in 1D.

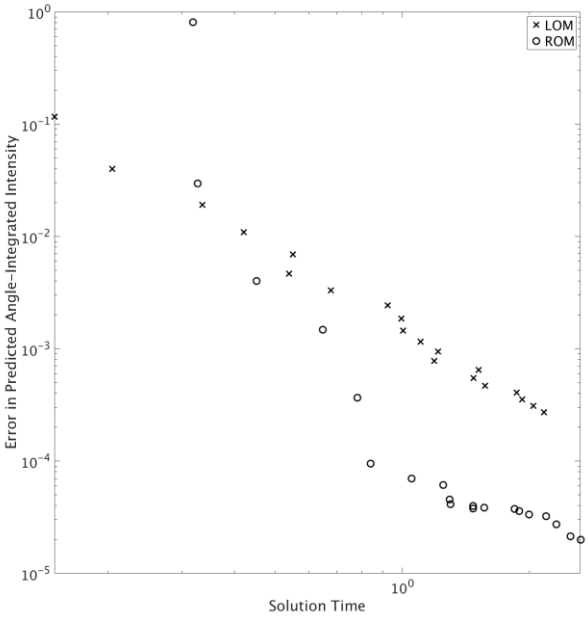

Fig. 3 Convergence of angle-integrated intensity distributions for quadratic temperature profile with increasing LOM quadrature order.

2-D EXAMPLES

In order to demonstrate this solution methodology in higher dimensional geometries, consider the case of a square surrounded by black walls and filled with a purely absorbing medium with opacity equal to the inverse of the side length of the square. Consideration of a purely absorbing medium is sufficient as this is mathematically equivalent to considering any individual source iteration step for a scattering problem. The LOM considered here is the 14th order PNTN quadrature which still contains 112 ordinate directions. The HOM is the 32nd order quadrature of the same type which contains 544 ordinate directions. The PNTN quadrature rule [9] is defined by using the Gauss-Legendre quadrature set to define the levels along the z-axis as well as the total weight for each level. The azimuthal angles for each level are set equal to the roots of the Chebyshev polynomials.

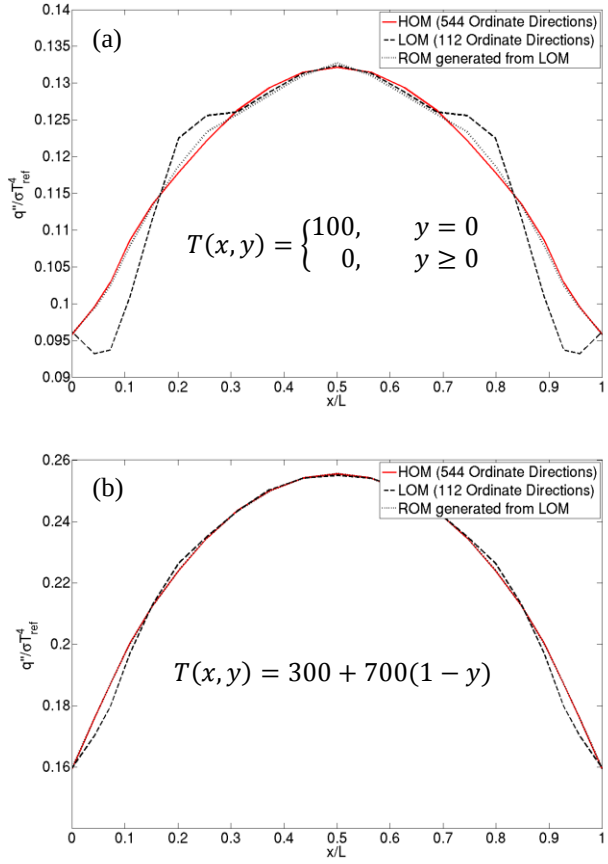

Fig. 4 Normalized heat flux predictions q”(x,1) for (a) a discontinuous temperature profile and (b) a linear temperature profile for the 14th order PNTN quadrature, 32nd order PNTN quadrature, and the ROM derived from the 14th order quadrature solutions but evaluated at the 32nd order quadrature points

Fig 4 shows the heat flux evaluated at the top (y=1) surface for a pair of temperature distributions. Both the accuracy of the solution and the improvement provided by the ROM relative to the LOM are seen to depend strongly on the source distribution, the order of the LOM, and the mesh resolution. However, in all cases a rapid increase in accuracy was observed once the LOM achieved a sufficient order.

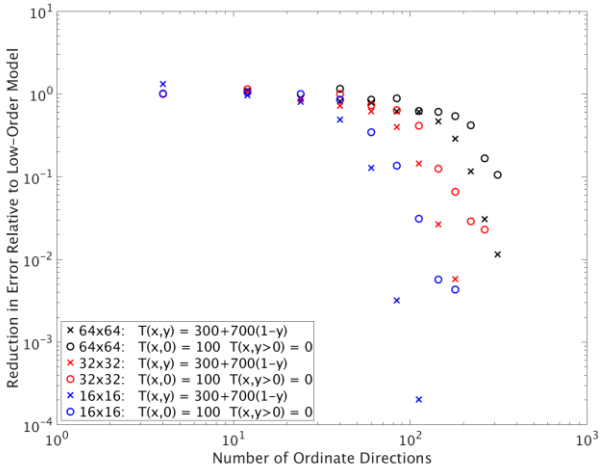

Fig. 5 Reduction in error provided by the ROM relative to the low-order model for various spatial mesh resolutions. This is the L2 error of the ROM solution divided by the L2 error of the low-order model solution

Fig 5 shows the reduction in error in the angularly integrated intensity distribution relative to the LOM for a variety of source distributions and mesh resolutions. In all of these cases a regular rectangular grid is used. The relative improvement is seen to be more pronounced for the linear temperature profile than for the discontinuous temperature profile. This is related to the fact that the LOM error tends to be smaller for the linear temperature profile. The number of ordinate directions required in the LOM to sufficiently inform the ROM and achieve the greatly enhanced accuracy is observed to increase with mesh refinement. This is likely due to the reduction in false scattering with increasing mesh resolution.

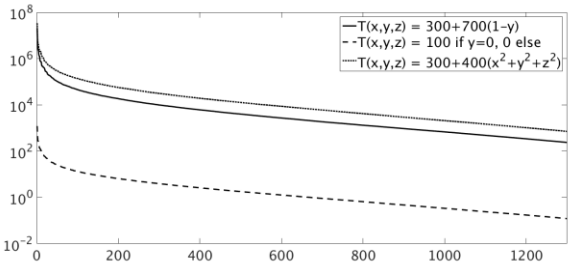

Fig. 6 Decay of singular values demonstrates that many additional modes are required to capture the behavior of the fully 3D problem

3-D EXAMPLES

We now move on to 3D applications. Consider a purely absorbing cube with an optical side-length of one surrounded by black walls. Fig 6 shows the relatively slow decay of the singular values for this geometry for several temperature distributions. Fig 6 shows that the methodology is significantly less efficient for the 3D problem than was demonstrated for the 2D problem. Between 711 and 1284 modes are required to capture 99.9999% of the energy depending on the source distribution while between 42 and 296 are required to capture even 99.9% of the energy. This is partially due to reduced symmetry and partially due to the mesh size. The same quadrature includes twice as many angles in 3D as it does in the 2D case. Also, the number of degrees of freedom in the 3D simulation is significantly greater than the number in the 2D simulation. The end result is that the LOM may be required to be impractically large in order to adequately inform the ROM and achieve significant accuracy gains.

However, this obstacle is not insurmountable. The failure of the previously described technique is largely due to inefficient sampling of the angular parameter space. This may be overcome by using a greedy search algorithm [17 - 20] to sample the parameter space. In the greedy search algorithm, samples are chosen adaptively by placing the new sample point at the location where the estimated error in the ROM prediction is maximum.

This process bears some similarity to previously developed adaptive quadrature methods [21-27]. However, it is fundamentally different from those techniques, which seek to enrich a low order quadrature through local refinement. In the proposed method, a ROM is constructed to represent a very high order quadrature and the ROM training points (ordinate directions) are chosen adaptively. The proposed method is demonstrated to be effective on 3D unstructured meshes.

As has been exploited in ROM error modeling methods [30], the ROM residual for a given direction is closely related to the ROM error for that direction and the residual is selected as an effective error indicator. The ROM residual tends to be a highly oscillatory function of the angle with many local minima and maxima making the continuous optimization problem inherent to each step of the greedy algorithm difficult. To avoid this difficulty, a discrete optimization problem is solved instead where the sample points are constrained to belong to a high order quadrature.

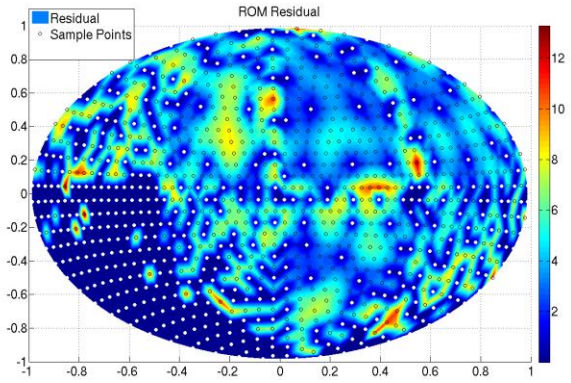

Fig. 7 Distribution of sample points generated through greedy search algorithm for a radial temperature distribution $T(r)=300+400r^2$

Fig 7 shows the result of one such greedy search after 440 sample points have been selected. The greedy search algorithm preferentially places sample points in the directions with large directional intensity values. In this case, those are directions pointing back towards the origin. In practice, acceptable levels of accuracy may be reached with far fewer sample points than are shown in Fig 7.

The greedy search algorithm is inherently suboptimal and so there is little value added by finding the exact location of the maximum error at each step. Significant time savings may be achieved by choosing the maximum of as few as 10 randomly selected points and only evaluating the ROM residual at these points and then sampling at the point with the largest residual norm. This greatly reduces the number of ROM evaluations required. Although a ROM evaluation is significantly less expensive than a FOM evaluation, evaluating the ROM for every quadrature point of the HOM prior to generating each sample can become expensive. Additional computational savings may be had by seeding the ROM with a number of predetermined sample points prior to initiating the greedy search algorithm. In the examples to follow, the ROM is seeded with 48 sample points analogous to the P6-T6 quadrature.

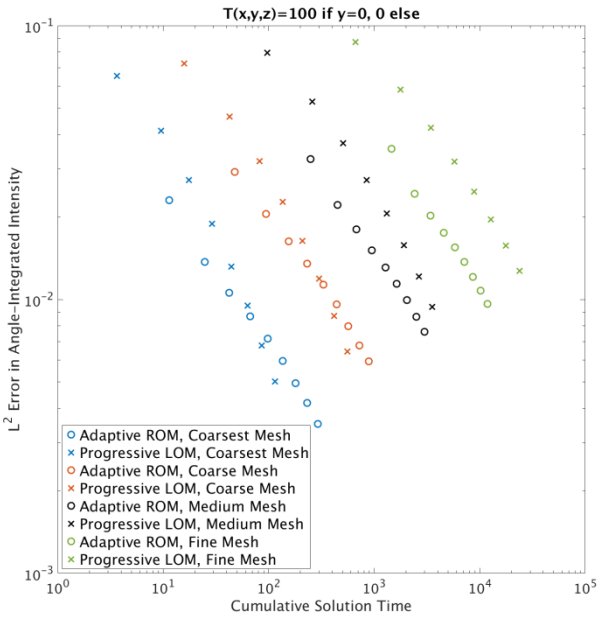

Fig. 8 Accuracy as a function of cumulative solution time for a discontinuous temperature distribution T(x,y,z)=100 if y=0, 0 else

The HOM is chosen to be the 40th order PN-TN quadrature (P40-T40) which includes 1680 ordinate directions. This number of ordinate directions has been shown to be sufficient to eliminate ray effects from the solution for the spatial mesh resolutions considered here [28]. The adaptive ROM performs better as the mesh is refined. This is due to the decreasing relative cost of ROM evaluations relative to FOM evaluations as mesh size increases.

Fig 8 shows the results of a timing study for a 3D example problem. Because the adaptive ROM yields a much more accurate solution than the LOM when given the same number of FOM evaluations but is slower because the adaptive ROM evaluates the ROM a large number of times in addition to the set number of FOM evaluations, it makes sense to compare the two approaches on the basis of error and solution time.

The error in the adaptive ROM may be (inexpensively) estimated at any time and used as a stopping criterion for the greedy search algorithm. The timings shown in Fig 8 include this error estimation although a maximum number of FOM evaluations was used as the stopping criteria rather than a predetermined error level. There is presently no analogous technique for estimating the error incurred by the discrete ordinates method too small a quadrature set is used as the LOM. Instead, an understanding of the magnitude of the error is typically obtained by performing multiple LOM evaluations with increasing quadrature orders. For example, one might evaluate the LOM using the P6-T6, P8-T8, and P10-T10 quadratures which include 48, 80, and 120 ordinate directions respectively and use the differences between the solutions to infer an approximate level of accuracy for the P10-T10 solution. If the inferred level of accuracy in inadequate, successively higher order quadratures may be invoked until the accuracy is acceptable.

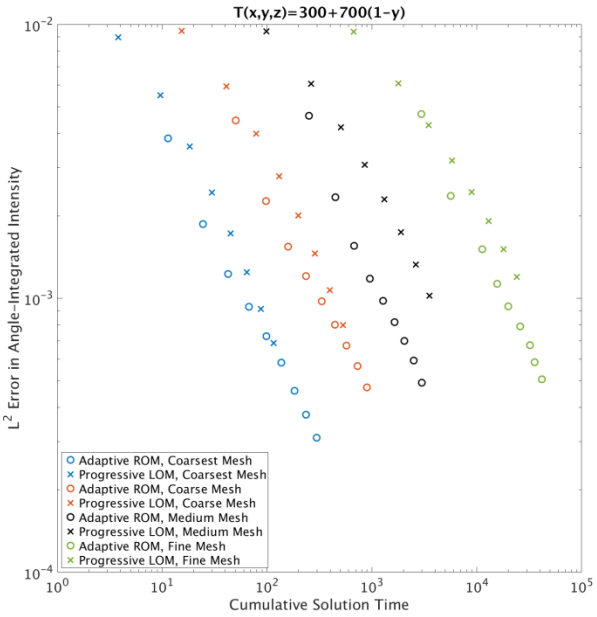

Fig. 9 Accuracy as a function of cumulative solution time for a linear temperature distribution, T(x,y,z)=300+700(1-y)

For this reason, the cumulative solution time is used in Figs 8, 9, and 10. It is shown in all three figures that seeding the adaptive ROM with 48 ordinate directions and applying no refinement (the first data point) is both more accurate than the

P10-T10 solution and significantly faster than evaluating the 3 LOMs listed above. This initial speed advantage increases with mesh resolution.

These conclusions hold true for a wide variety of temperature distributions. The series of meshes used in this set of example problems included approximately 3.4k, 12.2k, 36.5k, and 111.8k nodes. In all cases, the medium is purely absorbing and the walls are black. The wall temperature is equivalent to the temperature of the adjacent medium.

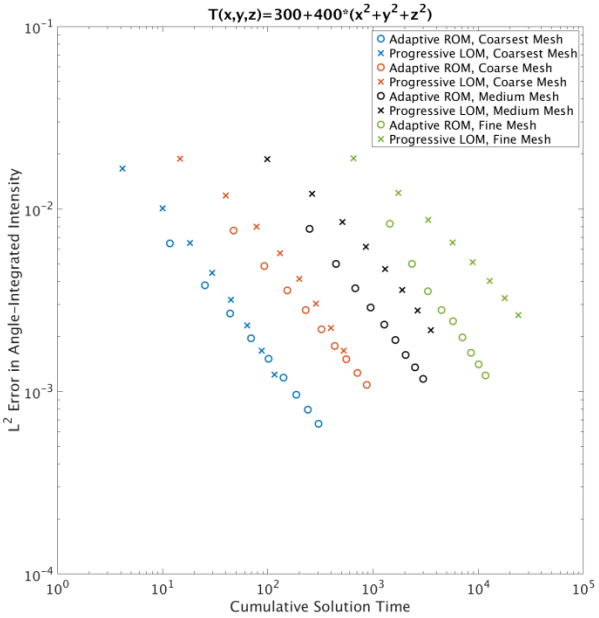

Fig. 10 Accuracy as a function of cumulative solution time for a radially quadratic temperature distribution T(x,y,z) = 300+400(x^2+y^2+z^2)

ERROR ESTIMATION

The adaptive ROM error may be estimated at any time by constructing a relationship between the ROM residual and the ROM error. At each step in the greedy search algorithm, both of these quantities are known for what will become the next sample point. The form of the relationship between the residual norm and the error norm is roughly a power law as shown in Fig 11. The training data in Fig 11 represents the residual norm – error norm pairs that are known from previous iterations of the greedy search algorithm. These errors are the L2 errors in the directional intensity. The prediction data are the errors calculated for all quadrature points in the HOM not included in the adaptive quadrature. The training data are used to generate a curve fit that allows for the approximation of the error at any point for which the ROM residual norm is known; this can be interpreted as a reduced-order model error surrogate (ROMES) model [30].

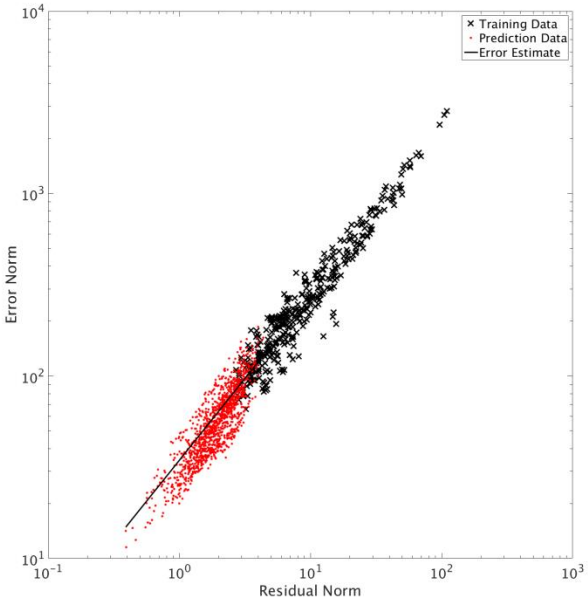

Fig. 11 Example of construction of ROM error estimate from training data acquired during the greedy search algorithm for a linear temperature distribution, T(x,y,z)=300+700(1-y)

This process provides an estimated distribution of the ROM error over the angular space. The resulting error in angularly integrated quantities may then be estimated as well. Fig 12 shows the convergence of one such quantity.

$$\text{Error} = \sum w_i^2 \int (I_i^{FOM} - I_i^{ROM})^2 dV \qquad (12)$$

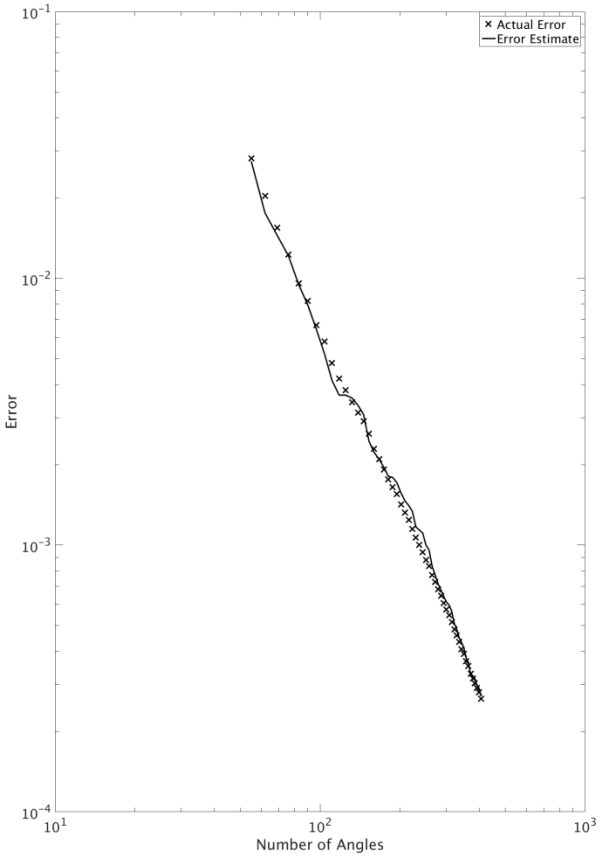

Fig. 12 Convergence of adaptive ROM and associated error estimate of an angularly integrated quantity for a linear temperature distribution, T(x,y,z)=300+700(1-y)

CONCLUSIONS

The discrete ordinates method is shown to be amenable for reduced order modeling. For 1D and 2D problems, any sufficiently large sampling of ordinate directions results in adequate sampling of the angular domain to construct a highly accurate reduced order model. For 3D problems more care must be taken to sample the angular domain efficiently. A greedy sampling approach is proposed using the ROM residual as an error indicator. Timing studies show this to be highly effective relative to successively increasing the angular quadrature order. The advantages of this approach increase with increasing mesh resolution since the cost of additional ROM evaluations is only weakly related to the mesh size (unlike the cost of additional FOM evaluations). Additionally, an error estimate is proposed that closely matches the actual error for the problems considered. This error estimate is inexpensive to generate and evaluate and may be easily used as a stopping criteria for the greedy search algorithm.

Although the deterministic error model presented performs well in predicting the mean behavior, additional utility is likely possible through the use of stochastic error models which would provide a distribution of error estimates. The source iteration scheme described by Eqns. 3 and 4 converges rapidly for problems with small scattering albedos. However, it is highly inefficient for highly scattering problems. There is potential for incorporating the ROM into a KP acceleration method which uses earlier ROM estimates to generate the low-order operator. This is currently under development.

ACKNOWLEDGMENTS

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. This document has been reviewed and approved for unclassified, unlimited release under 2015-XXXX

REFERENCES

[1] B. G. Carlson, "Mechanical Quadrature and the Transport Equation," Los Alamos Scientific Laboratory, Los Alamos, NM, 1961.

[2] B. G. Carlson, "Transport Theory: Discrete Ordinates Quadrature over the Unit Sphere," Los Alamos Scientific Laboratory, Los Alamos, NM, 1970.

[3] S. H. Kim and K. Y. Huh, "A New Angular Discretization Scheme of the Finite Volume Method for 3-D Radiative Heat Transfer in Absorbing, Emitting and Anisotropically Scattering Media," *International Journal of Heat and Mass Transfer,* vol. 43, pp. 1233-1242, 2000.

[4] R. Koch and R. Becker, "Evaluation of Quadrature Schemes for the Discrete Ordinates Method," *J. Quant. Spectrosc. & Radiat. Transfer,* vol. 84, pp. 423-435, 2004.

[5] R. Koch, W. Krebs, S. Wittig and R. Viskanta, "Discrete Ordinates Quadrature Schemes for Multidimensional Radiative Transfer," *J. Quant. Spectrosc. & Radiat. Transfer,* vol. 53, no. 4, pp. 252-372, 1995.

[6] K. D. Lathrop and B. G. Carlson, "Discrete Ordinates Angular Quadrature of the Neutron Transport Equation," Los Alamos Scientific Laboratory, Los Alamos, NM, 1965.

[7] E. E. Lewis and W. F. Miller, Jr., Computational Methods of Neutron Transport, La Grange Park, Illinois: American Nuclear Society, Inc., 1993.

[8] B. W. Li, J. H. Zhou, X. Y. Cao, H. G. Chen and K. F. Cen, "The Spherical Surface Symmetrical Equal Dividing Angular Quadrature Scheme for Discrete Ordinates Method," *ASME Journal of Heat Transfer,* vol. 124, no. 3, pp. 482-490, 2002.

[9] G. Longoni, Advanced Quadrature Sets, Acceleration and Preconditioning Techniques for the Discrete Ordinates Method in Parallel Computing Environments, Gainsville, FL: University of Florida, 2004.

[10] G. Longoni and A. Haghighhat, "Development and Application of the Regional Angular Refinement Technique and its Applications to Non-Conventional Problems," in *Proc. PHYSOR*, Seoul, Korea, 2002.

[11] S. A. Rukolaine and V. S. Yuferev, "Discrete Ordinates Quadrature Schemes Based on the Angular Interpolation of Radiation Intensity," *J. Quant. Spectrosc. & Radiat. Transfer*, vol. 69, pp. 257-275, 2001.

[12] C. Thurgood, A. Pollard and P. Rubini, "Development of TN Quadrature Sets and HEART Solution Method for Calculating Radiative Heat Transfer," in *Int. Symp. on Steel Reheat Furnace Technology*, Hamilton, Ontario, Canada, 1990.

[13] M. L. Adams and E. W. Larsen, "Fast Iterative Methods for Discrete-Ordinates Particle Transport Calculations," Progress in Nuclear Energy, vol. 40, no. 1, pp. 3-159, 2002.

[14] P. Benner, S. Gugercin and K. Willcox, "A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems," *SIAM Review*, 2015.

[15] D. Ryckelynck, "Hyper-Reduction of Mechanical Models Involving Internal Variables," *Int. J. Numer. Meth. Engng*, vol. 77, pp. 75-89, 2009.

[16] D. Amsallem, M. Zahr, Y. Choi and C. Farhat, "Design Optimization Using Hyper-Reduced-Order Models," *Structural and Multidisciplinary Optimization*, vol. 51, pp. 919-940, 2015.

[17] M. Grepl, "Reduced-Basis Approximation and A Posteriori Error Estimation for Parabolic Partial Differential Equations," Ph. D. thesis, MIT, Cambridge, MA, 2005.

[18] M. Grepl and A Patera, "A Posteriori Error Bounds for Reduced-Basis Approximations of Parametrized Parabolic Partial Differential Equations," *ESAIM Math. Modelling Numer. Anal.*, vol. 39, pp. 157-181, 2005.

[19] K. Veroy, C. Prud'homme, D.V. Rovas, and A.T. Patera, "A Posteriori Error Bounds for Reduced-Basis Approximation of Parametrized Noncoercive and Nonlinear Elliptic Partial Differential Equations," in *Proceedings of the 16th AIAA Computational Fluid Dynamics Conference*, Orlando, FL, 2003.

[20] T. Bui-Thanh, K. Willcox, and O. Ghattas, "Model Reduction for Large-Scale Systems with High-Dimensional Parametric Input Space," *SAIM J. Sci. Comput.*, vol. 30, no. 6, pp. 3270-3288, 2008.

[21] J. Stone and M. Adams, "Adaptive Discrete-Ordinates Algorithms and Strategies," in *Mathematics and Computation, Supercomputing, Reactor Physics and Nuclear and Biological Applications*, Palais des Papes, Avignon, France, 2005.

[22] D. Koeze, "Goal-Oriented Angular Adaptive Neutron Transport using a Spherical Discontinuous Galerkin Method," M.S. thesis, Delf University of Technology, Netherlands, 2011.

[23] J. Stone, "Adaptive Discrete-Ordinates Algorithms and Strategies," Ph.D. thesis, Texas A&M University, 2007.

[24] J. Stone and M. Adams, "Progress on Adaptive Discrete-Ordinates Algorithms and Strategies," in *Nuclear Mathematical and Computational Sciences: A Century in Review, A Century Anew*, Gatlinburg, Tennessee, 2003.

[25] G. Longoni and A. Haghighat, "Development and Application of the Regional Angular Refinement Technique and its Applications to Non-Conventional Problems," in *Proc. PHYSOR 2002*, Seoul, Korea, 2002.

[26] P. N. Brown and B. Change, "Locally Refined Quadrature Rules for SN Transport," Report UCRL-JRNL-220755, Lawrence Livermore National Laboratory, 2006.

[27] G. Longoni and A. Haghighat, "Development of New Quadrature Sets with the Ordinate Splitting Technique," in *Proceedings of the 2001 ANS International Meeting on Mathematical Methods for Nuclear Applications*, Salt Lake City, UT, 2001.

[28] J. Tencer, "A Comparison of Angular Discretization Techniques for the Radiative Transport Equation," in *Proceedings of the ASME 2015 International Mechanical Engineering Congress & Exposition*, Houston, TX, 2015.

[29] K. Carlberg, C. Bou- Mosleh, and C. Farhat. "Efficient non-linear model reduction via a least- squares Petrov–Galerkin projection and compressive tensor approximations." *International Journal for Numerical Methods in Engineering* 86, no. 2 (2011): 155-181.

[30] M. Drohmann and K. Carlberg. "The ROMES method for statistical modeling of reduced-order-model error." *SIAM/ASA Journal on Uncertainty Quantification* 3, no. 1 (2015): 116-145.