PHYSOR 2012 Advances in Reactor Physics Linking Research, Industry, and Education
Knoxville, Tennessee, USA, April 15-20, 2012, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2012)

RAYLEIGH QUOTIENT ITERATION IN 3D, DETERMINISTIC NEUTRON TRANSPORT

R. N. Slaybaugh*
*Department of Nuclear Engineering and Engineering Physics
University of Wisconsin - Madison
1500 Engineering Drive, Madison, WI 53703
r.slaybaugh@gmail.com

T. M. Evans and G. G. Davidson
Radiation Transport Group
Oak Ridge National Laboratory
PO Box 2008 MS6182, Oak Ridge, TN 37831
evanstm@ornl.gov; davidsongg@ornl.gov

P. P. H. Wilson*
wilsonp@engr.wisc.edu

ABSTRACT

Today's "grand challenge" neutron transport problems require 3-D meshes with billions of cells, hundreds of energy groups, and accurate quadratures and scattering expansions. Leadership-class computers provide platforms on which high-fidelity fluxes can be calculated. However, appropriate methods are needed that can use these machines effectively. Such methods must be able to to use hundreds of thousands of cores and have good convergence properties. Rayleigh quotient iteration (RQI) is an eigenvalue solver that has been added to the S_N code Denovo to address convergence.

Rayleigh quotient iteration is an optimal shifted inverse iteration method that should converge in fewer iterations than the more common power method and other shifted inverse iteration methods for many problems of interest. Denovo's RQI uses a new multigroup Krylov solver for the fixed source solutions inside every iteration that allows parallelization in energy in addition to space and angle. This Krylov solver has been shown to scale successfully to 200,000 cores: for example one test problem scaled from 69,120 cores to 190,080 cores with 98% efficiency. This paper shows that RQI works for some small problems. However, the Krylov method upon which it relies does not always converge because RQI creates ill-conditioned systems. This result leads to the conclusion that preconditioning is needed to allow this method to be applicable to a wider variety of problems.

Key Words: Rayleigh quotient, shifted inverse iteration, Krylov methods, energy decomposition, massively parallel

1. INTRODUCTION

The steady-state Boltzmann equation for neutron transport covers six dimensions of phase space. Typical transport problems today are three-dimensional, have up to thousands $\times$ thousands $\times$ thousands of mesh points, use up to $\sim$150 energy groups, include accurate expansions of scattering terms, and are solved over many angular directions. The next generation of challenging problems are even more highly refined. High-fidelity, coupled, multi-physics calculations are the

new "grand challenge" problems for reactor analysis, requiring that the finely-resolved neutron flux be calculated quickly and accurately.

Very large computers, such as the Jaguar machine [1], are now available to perform such high-fidelity calculations. Most existing solution methods are not able to take full advantage of new computer architectures, or they have convergence properties that limit their usefulness for difficult problems. This work aims to accelerate the eigenvalue transport solution with a method that can use leadership-class machines fully and alleviate some of the convergence issues.

To accomplish the goal of this work, an eigenvalue method was added to the code Denovo [2]. This paper describes the addition of Rayleigh Quotient Iteration (RQI) and discusses how it could be useful in the future. The basis of RQI indicates that it should converge in fewer iterations than the traditionally-used power method and other shifted inverse power methods. A recently-added multigroup (MG) Krylov solver enables the use of RQI for large problems in massively parallel environments [3]. However, because a Krylov method is used to converge the eigenvector within each eigenvalue iteration, RQI may not always be successful in practice without preconditioning.

This paper confirms that RQI uses fewer iterations than power iteration (PI) for some problems and previews that it needs preconditioning to be successful for others. To address this need, we have developed a parallel, multigrid preconditioner. Full details of this method will be addressed in a future paper, but we show some preliminary results on the C5G7 benchmark [4] to illustrate its potential effectiveness. This paper is organized as follows, Section 2 gives pertinent background information. Section 3 provides an overview of Denovo's new multigroup solver and the implementation of RQI. Section 4 discusses some potential limitations on RQI's performance. Some test results and the conclusions are presented in Sections 5 and 6, respectively.

2. BACKGROUND

The multigroup S_N equations can be written in operator form as

$$\mathbf{L}\psi = \mathbf{MS}\phi + \frac{1}{k}\mathbf{M}\chi f^T\phi \,. \tag{1}$$

Here the angular flux is ψ, and the angular flux moments are ϕ. $\mathbf{L}$ is the first-order linear differential transport operator, and $\mathbf{M}$ is the moment-to-discrete operator that projects the angular flux moments onto discrete angles. $\mathbf{S}$ is the scattering matrix, χ is the fission spectrum vector, and f is the vector of fission cross sections. k is the asymptotic ratio of the number of neutrons in one generation to the number in the next. Eq. (1) can be converted to a fixed source problem by replacing the fission term with an external source, q_e.

The angular flux and angular flux moments can be related through the discrete-to-moment operator: $\phi = \mathbf{D}\psi$. Combining this relationship with Eq. (1) and rearranging gives an equation only in terms of ϕ. The formulation is aided by defining $\mathbf{T} = \mathbf{DL}^{-1}$, and $\mathbf{F} = \chi f^T$ [3]:

$$(\mathbf{I} - \mathbf{TMS})\phi = \frac{1}{k}\mathbf{TMF}\phi \,. \tag{2}$$

A common way to solve eigenvalue problems is with PI. This method is attractive because it only requires matrix-vector products and two vectors of storage space. PI uses the form of the problem

seen in Eq. (3) and then iterates as shown in Eq. (4), where i is the iteration index. This converts the generalized form of the eigenvalue problem seen in Eq. (2) to the ordinary form. In the generalized form the eigenvector-value pair is $(\phi, \frac{1}{k})$ and in the ordinary form it is (ϕ, k). In legacy applications the eigenvector is often the fission source rather than the flux moments [3], [5].

$$\mathbf{A}\phi = k\phi \,, \tag{3}$$

$$\text{where} \qquad \mathbf{A} = (\mathbf{I} - \mathbf{TMS})^{-1}\mathbf{TMF} \,,$$

$$\phi^{i+1} = \frac{1}{k^i}\mathbf{A}\phi^i \,; \qquad k^{i+1} = k^i \frac{f^T\phi^{i+1}}{f^T\phi^i} \,. \tag{4}$$

Inside of PI, the application of $\mathbf{A}$ to ϕ requires the solution of a multigroup problem that looks like a fixed source problem,

$$(\mathbf{I} - \mathbf{TMS})y^i = \mathbf{TMF}\phi^i \,. \tag{5}$$

PI's convergence can be very slow for problems of interest. For an $n \times n$ matrix $\mathbf{A}$, an eigenvalue-vector pair satisfies $\mathbf{A}x_i = \lambda_i x_i$ for $i = 1, ..., n$. Let $\sigma(\mathbf{A}) \equiv \{\lambda \in \mathbb{C} : rank(\mathbf{A} - \lambda\mathbf{I}) \leq n\}$ be the spectrum of $\mathbf{A}$ and the eigenvalues be ordered as $|\lambda_1| > |\lambda_2| \geq \ldots \geq |\lambda_n| \geq 0$. The error from PI is reduced in each iteration by a factor of $\mathbf{A}$'s dominance ratio, $\frac{\lambda_2}{\lambda_1}$. For loosely coupled systems PI will converge slowly because λ_2 is close to λ_1 [6].

Shifted inverse iteration (SII) typically converges more quickly than PI. SII capitalizes on the fact that for some shift μ, $(\mathbf{A} - \mu\mathbf{I})$ will have the same eigenvectors as $\mathbf{A}$. If $\mu \notin \sigma(\mathbf{A})$, then $(\mathbf{A} - \mu\mathbf{I})$ is invertible and $\sigma([\mathbf{A} - \mu\mathbf{I}]^{-1}) = \{1/(\lambda - \mu) : \lambda \in \sigma(\mathbf{A})\}$. Eigenvalues of $\mathbf{A}$ that are near the shift will be transformed to extremal eigenvalues that are well separated from the others. The shifted and inverted matrix is used in a power iteration-type scheme. Given a good shift, $\mu \approx \lambda_1$, SII usually converges more quickly than PI, especially for loosely coupled systems [6], [7].

Wielandt's method is a flavor of SII that has been used widely in the neutron transport community [8], [9], [10]. In the transport equation formulation, Wielandt's method changes the eigenvalue problem to $(\mathbf{I} - \mathbf{TM}(\mathbf{S} + \gamma_e\mathbf{F}))\phi = \delta\gamma\mathbf{TMF}\phi$, where γ_e is the current estimate for the dominant eigenvalue, $\gamma_1 = \frac{1}{k}$, and $\delta\gamma = \gamma_1 - \gamma_e$. The power method is applied to this, giving [11]

$$\phi^{i+1} = \delta\gamma^i(\mathbf{I} - \mathbf{TM}(\mathbf{S} + \gamma_e\mathbf{F}))^{-1}\mathbf{TMF}\phi^i \,. \tag{6}$$

SII is often preferable to PI. Allen and Berry [7] compared the power method and the shifted inverse power method for the one-group, 1-D transport equation and found the shifted inverse method converges much more quickly. Itagaki [9], [10] used Wielandt's spectral shift technique to solve the one-group, 3-D diffusion equation for various problem types. Itagaki's numerical results show that when the estimate γ_e is not good, the correct eigenvalue may never be found.

An analysis of the Ringhals reactor was done using two-group, finite differenced diffusion with Wielandt's method. When doing actual core analysis it was found that the local flux shape around the control rods made the flux converge slowly. The only way to improve convergence was to use a good initial guess for the flux [12]. More recently Zinzani et al. [8] modified the Wielandt method for the two-group diffusion equation to be able to find multiple eigenmodes. The method removes the requirement of a good initial guess. Numerical results showed their method to be accurate and robust, but the algorithm is complicated and quite slow.

These studies demonstrate, as is expected given the theoretical considerations discussed above, that shifted inverse iteration can perform better than power iteration and the degree to which it is better often depends upon the quality of the shift and the initial guess for the eigenvector. All of these examples, however, only used one or two groups.

3. RAYLEIGH QUOTIENT ITERATION IN DENOVO

The code being used in this work is Denovo, a massively parallel discrete ordinates code being developed at Oak Ridge National Laboratory. Rayleigh Quotient Iteration is a shifted inverse iteration method that uses an optimal shift. RQI applies a shift to the equations via the scattering matrix in a manner similar to Wielandt's method. This can cause some complications for common fixed source solvers.

The scattering matrix is lower triangular for groups that have downscattering. There are entries above the diagonal only when there is upscattering. The fission matrix has an energy-block filled column for every group with fission. Applying the shift as $(\mathbf{S} + \gamma_e \mathbf{F})$ makes the scattering matrix energy-block dense even when there is very little upscattering. Traditional solution methods for the fixed source part of the equation do not handle dense scattering matrices well. This has hampered the implementation of SII in multi-group, 3-D codes because solving many groups that have upper-triangular scattering entries will take a long time. To overcome this barrier, a new solver is needed.

3.1. Multigroup Solver

The new MG Krylov solver removes the traditional inner-outer iteration structure from fixed source, or multigroup, solution. This change lets the solver handle upscattering efficiently and enables parallelization in the energy dimension. The solver has been shown to successfully scale to hundreds of thousands of cores. For example, a test scaled from 69,102 cores to 190,080 cores with 98% efficiency [13]. The fixed source problem has traditionally been solved with Gauss Seidel (GS). GS is an iterative method in energy, where a space-angle solve is done for each energy group individually using a within-group solver. The energy groups are solved in series, where $g = 1$ is the highest energy group and $g = G$ is the lowest. For a group g and an energy iteration index k this is [2]

$$\big(\mathbf{I} - \mathbf{TMS}_{gg}\big)\phi_g^{k+1} = \mathbf{TM}\Big(\sum_{g'=1}^{g-1} \mathbf{S}_{gg'}\phi_{g'}^{k+1} + \sum_{g'=g+1}^{G} \mathbf{S}_{gg'}\phi_{g'}^{k} + q_{e,g}\Big) \,. \tag{7}$$

The first term on the right includes downscattering contributions from higher energy groups and the second term represents upscattering contributions from lower energy groups that have not yet been converged for this energy iteration. All downscattering groups only need to be solved once since the second term on the right is zero. Groups with upscattering, however, must be repeatedly solved until they converge. Convergence of GS is governed by its spectral radius, so the method can be very slow when upscattering has a large influence on the solution [14]. GS is fundamentally serial in energy because of how the group-to-group coupling in treated.

The MG Krylov solver combines the space-angle and energy iterations to make one iteration level. This allows the energy groups to be decomposed such that they can be solved in parallel. The space-angle-energy iterations are much like the space-angle within-group iterations, but the iteration is over a block of groups instead of just one group. The block can include all groups or just upscattering groups, in which case the downscattering groups are treated as before. An added benefit is that Krylov methods generally converge more quickly than GS [6].

The multigroup Krylov method applied to the upscattering block is shown here, where $\mathbf{S}_{\text{up_block}}$ contains the upscattering groups and $\mathbf{S}_{\text{up_source}}$ has the downscattering only groups

$$\underbrace{(\mathbf{I} - \mathbf{TMS}_{\text{up_block}})}_{\tilde{\mathbf{A}}} \phi^{k+1}_{\text{up_block}} = \mathbf{TM}(\mathbf{S}_{\text{up_source}} \phi^{k+1}_{\text{up_source}} + q_e) \ . \tag{8}$$

Trilinos [15] provides Denovo's Krylov solver, with a choice of either GMRES or BiCGSTAB [2]. The Krylov solver is given an operator that implements the action of $\tilde{\mathbf{A}}$, or the matrix-vector multiply and sweep. In the new solver $\tilde{\mathbf{A}}$ is applied to an iteration vector, v, containing the entire upscattering block instead of just one group:

1. matrix-vector multiply: $y = \mathbf{MS}_{\text{up_block}} v$,
2. sweep: $z = \mathbf{T} y$,
3. return: $v \leftarrow v - z$.

To implement the energy parallelization, the problem is divided into energy sets, with groups distributed evenly among sets. After each set performs its part of the matrix-vector multiply, a global reduce-plus-scatter is the only required inter-set communication. Since each set uses the entire spatial mesh with the same spatial decomposition, the established performance of spatial scaling does not change. Energy decomposition offers the ability to further decompose a problem, even once the practical limit of spatial performance has been reached. The total number of cores is equal to the number of computational domains, that is, the product of the number of energy sets and the number of spatial blocks. For 20,000 spatial blocks and 10 energy sets, which is a reasonable decomposition, 200,000 cores will be used. See Ref. [3] for more details.

3.2. Rayleigh Quotient Iteration

The RQI solver uses the MG Krylov solver and as a result can be decomposed in energy as well. In addition to being able to scale to hundreds of thousands of cores, it is hypothesized that the superior convergence properties of RQI will prove beneficial in practice. The reason the shift in RQI is optimal comes from the Rayleigh quotient (RQ) itself. For a generalized eigenvalue problem $\beta \mathbf{A} x = \alpha \mathbf{B} x$, where $\gamma = \frac{\alpha}{\beta}$, the RQ is

$$\rho = \frac{y^T \mathbf{A} x}{y^T \mathbf{B} x} \ . \tag{9}$$

If x and y are right and left eigenvectors corresponding to α and β, respectively, then $\alpha = y^T \mathbf{A} x$ and $\beta = y^T \mathbf{B} x$. In this case, $\rho = \gamma$ [16].

The RQ provides the minimum residual for the eigenvalue problem in the least squares sense, where equivalence holds only when $\tau = \rho$ and $x \neq 0$:

$$||(\mathbf{A} - \tau\mathbf{B})x||^2 \geq ||\mathbf{A}x||^2 - ||\rho\mathbf{B}x||^2 \,. \tag{10}$$

The Rayleigh quotient is thus an optimal guess for an eigenvalue given a vector close to the corresponding eigenvector. Rayleigh quotient iteration is like SII but rather than using a manually chosen shift, the shift for iteration i is

$$\mu_i = \gamma_{i-1} - \rho_i \,. \tag{11}$$

RQI has better convergence properties than SII since the RQ is an optimal guess for the eigenvalue of interest. For more details on the Rayleigh quotient and RQI in general refer to Ref. [17].

RQI has been implemented in Denovo by subtracting $\rho\mathbf{TMF}$ from both sides of Eq. (2). This gives the following shifted system where $\gamma \equiv \frac{1}{k}$:

$$(\mathbf{I} - \mathbf{TM}\tilde{\mathbf{S}})\phi = (\gamma - \rho)\mathbf{TMF}\phi \,, \tag{12}$$
$$\text{where } \tilde{\mathbf{S}} \equiv \mathbf{S} + \rho\mathbf{F} \,.$$

The new matrix, $\tilde{\mathbf{S}}$, is energy-block dense since the fission matrix is dense. $\tilde{\mathbf{S}}$ looks like one big upscattering block and can be treated similarly. By using the new MG Krylov solver over all groups the entire problem can be decomposed in energy and solved efficiently. Note that using GS to solve a problem with a dense $\tilde{\mathbf{S}}$ would likely take a very long time to converge.

Algorithm 1 Rayleigh Quotient Iteration in Denovo

Get initial guesses, ϕ_0 and k_0
Calculate the Rayleigh quotient, ρ_1
For $i = 1, 2, ...$ until convergence:

1. Calculate the shift using Eq. (11)
2. Apply the shift for each energy set as seen in Eq. (12)
3. Solve $[\mathbf{I} - \mathbf{TM}\tilde{\mathbf{S}}]^{\tilde{g}}\phi^{\tilde{g}} = \mu[\mathbf{TMF}]^{\tilde{g}}\phi^{\tilde{g}}$ for an updated $\phi^{\tilde{g}}_{i+1}$ with the multigroup Krylov solver
4. Global barrier so all processes finish calculating $\phi^{\tilde{g}}_{i+1}$; all processes get the new ϕ_{i+1}
5. Set $k_i = \frac{1}{\rho_i}$
6. Calculate ρ_{i+1} via Eq. (13); $k_{i+1} = \frac{1}{\rho_{i+1}}$
7. Check for convergence: $\frac{k_{i+1} - k_i}{k_i} < tol$

The energy-parallelized RQI method is shown in Algorithm 1. Here $\mathbf{A} = (\mathbf{I} - \mathbf{TMS})$, and a block of energy groups corresponding to an energy set is denoted by superscript $\tilde{g}$. The moments being computed by a given energy set are indicated by $\phi^{\tilde{g}}$. The Rayleigh quotient is

$$\rho = \frac{\phi^T\mathbf{A}\phi}{\phi^T\mathbf{TMF}\phi} \,. \tag{13}$$

4. Convergence Concerns

Using RQI in combination with a Krylov method raises some concerns about whether or not it can converge. A method's behavior can be affected by characteristics of the mathematical problem as well as by characterisitics of the algorithm. Conditioning is one way to express the perturbation behavior of the mathematical problem. An ill-conditioned problem is one in which some small perturbation of x leads to a large change in $f(x)$. A large condition number corresponds to an ill-conditioned problem, and vice versa.

The condition number of a matrix $\mathbf{A}$ for a norm p is defined as

$$\kappa_p(\mathbf{A}) = ||\mathbf{A}||_p \, ||\mathbf{A}^{-1}||_p \, . \tag{14}$$

For an $n \times m$ matrix $\mathbf{A}$, the singular values are ordered such that $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n > 0$. If the two-norm is used, then $||\mathbf{A}||_2 = \sigma_1$, $||\mathbf{A}^{-1}||_2 = \frac{1}{\sigma_m}$, and $\kappa_2(\mathbf{A}) = \frac{\sigma_1}{\sigma_m}$, where σ_m is the mth singular value of $\mathbf{A}$. If $\mathbf{A}$ is singular, its condition number is infinity [6].

On initial consideration it would seem shifted inverse methods might not work well when the shift is very good because the matrix becomes so ill-conditioned. If the shift is exact, i.e. when $\mu = \gamma_1$, the matrix is singular. It turns out this concern is unfounded. Peters and Wilkinson [18] proved that ill-conditioning is not a problem for inverse iteration methods. Trefethen and Bau [6] also assert that this is the case as long as the fixed source portion is solved with a backwards stable algorithm.

A backwards-stable algorithm will "give exactly the right answer to nearly the right question"[6]. An algorithm $\tilde{f} : X \to Y$ for a problem $f : X \to Y$ is backwards stable if, for each $x \in X$ [6],

$$\tilde{f}(x) = f(\tilde{x}) \qquad \text{for some } \tilde{x} \text{ with } \frac{||\tilde{x} - x||}{||x||} = O(\epsilon_{machine}) \, . \tag{15}$$

4.1. Krylov Methods

With this in mind it seems that for RQI to succeed, the Krylov method used in it must (1) be backwards stable and (2) converge the eigenvector so that a good approximation to the eigenvalue can be obtained. Convergence and stability properties differ from Krylov method to Krylov method. Because GMRES is the most widely used Krylov method in this work, convergence and stability discussions will focus on it.

Paige et al. [19] demonstrate that GMRES is backwards stable when finding x in $\mathbf{A}x = b$ for a "sufficiently nonsingular $\mathbf{A}$." They derive some specific criteria that show $\mathbf{A}$ is non-singular enough when its minimum singular value is below a certain size and when the maximum and minimum singular values of the orthonormal matrix created in the Arnoldi process are not too far apart. Thus it seems that GMRES is only guaranteed to be backward stable for sufficiently well-conditioned systems [19].

Nachtigal et al. [20] perform convergence analysis for GMRES using an arbitrary $\mathbf{A}$. They show the bound on the convergence rate is related to the pseudospectrum of $\mathbf{A} + \mathbf{E}$ where $||\mathbf{E}|| \leq \epsilon$ and $\epsilon > 0$. This indicates that GMRES might converge very slowly for systems with poor spectral

properties. However, GMRES is monotonically convergent: $||r_{n+1}||_2 \leq ||r_n||_2$, so given a sufficient amount of time and resources it will eventually converge [20], [6].

While these assertions are for GMRES, the experience of the computational community indicates the conclusions extend to most Krylov methods. For ill-conditioned systems Krylov methods may not be backwards stable and tend to converge very slowly. As a result, many researchers have found that Krylov methods must be preconditioned to be able to get good results in practice [21], [22], [6] , [19].

5. RESULTS

A few calculations were done to investigate the improvement gained from RQI. The quality of improvement is measured by the total number of Krylov iterations rather than timing. This is a more meaningful measure because some calculations were so small and ran so fast that timing was inaccurate, this is a metric that can be compared across computers, and there may be room for some optimization in the future that would change the timing.

The first Denovo problem chosen was a small, few-group system intended to demonstrate the new solver's correctness. This problem used a $3 \times 3 \times 3$ grid with 0.1 cm spacing, vacuum boundary conditions, two materials, four groups with only downscattering, S_2, P_0, and had a dominance ratio of 0.1396 and a reference k of 0.11752. The small dominance ratio of this and many of the test problems used implies that PI should perform reasonably well on them. The debug version of Denovo was used and the calculation was not parallelized in any way.

This problem was solved with RQI and PI, both using the MG Krylov solver for the multigroup iterations. Note that when PI is used with no upscattering, the groups are solved one at a time using a Krylov method within each group. The RQI method uses MG Krylov over the whole block; recall that this system has effective upscattering added by the shift. GMRES was the Krylov method used in all cases.

Table I. Small, Two Material, Four Group, Vacuum Boundary Problem RQI Results

Solvers	k	group iters	eigen iters	total Krylov
PI	0.11752	2 or 3 / group	7	73
RQI	0.11752	4 to 8	6	39

The results are given in Table I. The term "group iters" either refers to the number of within group iterations (in the case of PI) or multigroup iterations where the Krylov vector is over all groups (in the case of RQI). Both RQI and PI got the correct answer. PI used a total of 73 Krylov iterations while RQI used a total of 39 Krylov iterations. The subspace sizes were not the same so the work per Krylov iteration was different, but RQI still used fewer Krylov iterations than PI.

The second problem was nearly the same as the first, but was intended to test the reflecting boundary case. For this problem all boundaries were reflecting and only one material was used.

The changes made the dominance ratio 1.630×10^{-15} and k became 2. These results can be seen in Table II. Again RQI needed fewer Krylov iterations than PI and both methods converged to the correct answer.

Table II. Small, One Material, Four Group, Reflecting Boundary Problem RQI Results

Solvers	k	group iters	eigen iters	total Krylov
PI	2	9 / group	2	72
RQI	2	17 and 18	2	35

These first two small tests show that RQI can get the right answer and that it can converge in fewer iterations than power iteration, even for problems with a small dominance ratio. To continue investigation, an intermediately-sized problem was run next. This system had a $9 \times 9 \times 9$ spatial grid with 0.5 cm spacing, reflecting boundary conditions, one material, 27 groups, 13 upscattering groups, S_2, and P_1. The dominance ratio was 6.32×10^{-13} and the reference k was 0.4.

When PI with MG Krylov was used the correct answer was obtained in 90 Krylov iterations per eigenvalue iteration and two outer iterations for a total of 180 Krylov iterations. When RQI was used, k did not converge. The maximum number of Krylov multigroup iterations per eigenvalue iteration was set to 1,000 and this was taken every time except the first, which took 25. This means the eigenvector did not converge after the first iteration. The value of k oscillated between 0.3966 and 0.3967 until the calculation was manually terminated. While the problem did not converge, the estimated eigenvalue was close to the reference.

With an even more difficult problem RQI could not get even nearly the right eigenvalue. The two-dimensional C5G7 mox benchmark problem [23] was selected for investigation with an optimized version of Denovo. This problem used three enrichment levels, an external mesh file, two reflecting boundaries to make it 2-D, 10 materials, 7 groups, 4 upscattering groups, S_2, P_0, k_0 set to 1.14, and had a dominance ratio of 0.7709. The reference k is 1.18655 ± 0.008.

Table III. 2-D MOX C5G7 Benchmark RQI Results

Solvers	k	mg iters	eigen iters	total Krylov
PI + TTG GS	1.18538	157 GS	32	21,365
PI + MG Krylov	1.18702	$\sim$100	32	3,129
RQI	$\sim$0.95	1,000	120*	119,006

*Terminated manually

The results from this calculation are presented in Table III. Here power iteration was used with both transport two grid accelerated Gauss Seidel (TTG GS) [2] and with MG Krylov. PI + TTG GS needed 21,365 single-group-sized Krylov iterations and PI + MG Krylov needed 3,129

all-group-sized Krylov iterations to converge. With RQI the eigenvector did not converge and the eigenvalue was wrong. Except for the first set, the multigroup iterations did not converge. A 3-D version of the C5G7 mox benchmark problem was also tried with RQI. This, too, did not converge.

When the Krylov iterations do not converge, the flux estimate is not good. Without a good approximation to the eigenvector, the RQ is no longer a valid approximation to the eigenvalue and thus the eigenvalue problem does not converge. It is likely that the eigenvector does not converge because RQI creates poorly conditioned systems. Krylov methods do not handle ill-conditioned systems very well and may converge extremely slowly when trying to solve them. These observations suggest that preconditioning the Krylov solve will resolve these issues.

It has in fact been found that preconditioning the multigroup solves can allow the Krylov solver to converge the eigenvector and RQI to subsequently converge [13]. A new multigrid in energy preconditioner [13] was used with the MG Krylov solver inside both RQI and PI. The preconditioner is a geometric multigrid scheme over energy and works by applying one or more V-cycles per multigroup GMRES iteration. Detailed discussion of these results is beyond the scope of this paper, but one illustrative example, the 3-D C5G7 benchmark problem, is given here. This problem has a dominance ratio of 0.773 and was solved on 720 cores.

Preconditioned RQI converged in 19 eigenvalue iterations and required a total of 187 Krylov iterations. Preconditioned PI needed 32 eigenvalue and 192 total Krylov iterations. Unpreconditioned PI also took 32 eigenvalue iteration, but needed 1,224 Krylov iterations. Note that the Krylov subspace sizes were the same in all cases. This result indicates that with preconditioning not only can RQI converge, but it can do so in fewer iterations than power iteration for problems that are more interesting than the very simple tests discussed above. Additional results for more challenging problems using up to 198,528 cores will be published shortly, and some results can be found now in reference [13].

6. CONCLUSIONS

Rayleigh quotient iteration is an old method that is an adaptation of shifted inverse iteration. Because the RQ provides an optimal shift, it is expected that using RQI in Denovo will converge in fewer iterations than power iteration for loosely coupled problems, provided that the eigenvector is converged.

Further, this solver is wrapped around the new MG Krylov solver. It has been shown that the energy decomposition improves solution time and scales well [13]. It is therefore reasonable to expect that when RQI converges, decomposing the entire matrix in energy for eigenvalue calculations will work and provide the same kind of scaling improvement.

RQI has not been applied to the transport equation before. This is likely because it takes $O(n^3)$ operations for full dense matrices, and without parallelization in energy it could be prohibitively expensive [16]. In the past, adding a shift to make the scattering matrix energy-block dense was difficult to handle. The system would have been solved with Gauss Seidel, which would have been restrictively slow. The multigroup Krylov algorithm has enabled energy parallelization and made the calculation of the eigenvector tractable.

Computers that facilitate enough parallelization to decompose in energy and have enough memory to store Krylov subspaces for the full transport equation make combining RQI and MG Krylov possible. The resulting convergence of the eigenvalue should be faster than PI for at least some problems. The energy decomposition will allow eigenvalue calculations to overcome the limitations of scaling in space and angle alone such that the code can be scaled to many more cores. These ideas have never before been used together in this way.

The intermediate and large problems that were tested demonstrate that RQI does not converge for even slightly challenging problems. This is probably caused by the poor condition number created by RQI. When the multigroup Krylov solver is not preconditioned it cannot converge the eigenvector in many cases. The incorrect eigenvector subsequently creates an incorrect eigenvalue.

The small problems showed that RQI can converge in fewer iteration than power iteration and therefore has the potential to be quite beneficial if the multigroup iterations can be converged. If the MG Krylov solver is preconditioned, then it may be able to converge the eigenvector for cases of interest. Results in reference [13] in addition to the 3-D C5G7 results mentioned here show that RQI works when preconditioned with a multigrid in energy preconditioner. With preconditioning RQI has the potential to converge in many fewer iterations that the traditional power iteration method.

ACKNOWLEDGEMENTS

This research was performed under appointment to the Rickover Fellowship Program in Nuclear Engineering sponsored by Naval Reactors Division of the U.S. Department of Energy. The authors would also like to thank Oak Ridge National Laboratory and their computing facilities (OLCF) for computer time and project support.

REFERENCES

[1] S. N. C. F. Computational. "World's Most Powerful Computer For Science." (2010). URL http://www.nccs.gov/jaguar/.

[2] T. M. Evans, A. S. Stafford, R. N. Slaybaugh, and K. T. Clarno. "Denovo – New Three-Dimensional Parallel Discrete Ordinates Code in SCALE." *Nuclear Technology*, **volume 171**(2), pp. 171–200 (2010).

[3] T. M. Evans, G. G. Davidson, and R. N. Slaybaugh. "Three-Dimensional Full Core Power Calculations for Pressurized Water Reactors." *Journal of Physics: Conference Series*, **volume SciDAC 2010** (expected 2012).

[4] OECD-NEA. "Benchmark on Deterministic Transport Calculations Without Spatial Homogenisation, MOX Fuel Assembly 3-D Extension Case." Technical Report 5420, Paris, France (2005).

[5] E. E. Lewis and W. F. Miller Jr. *Computational Methods of Neutron Transport*. American Nuclear Society, La Grange Park, IL (1993).

[6] L. N. Trefethen and D. Bau, III. *Numerical Linear Algebra*. SIAM, Philadelphia, PA (1997).

[7] E. J. Allen and R. M. Berry. "The Inverse Power Method for Calculation of Multiplication Factors." *Annals of Nuclear Energy*, **volume 29**(8), pp. 929–935 (2002).

[8] F. Zinzani, C. Demaziere, and C. Sunde. "Calculation of the Eigenfunctions of the Two-Group Neutron Diffusion Equation and Application to Modal Decomposition of BWR Instabilities." *Annals of Nuclear Energy*, **volume 35**(11), pp. 2109–2125 (2008).

[9] M. Itagaki and N. Sahashi. "Three-Dimensional Multiple Reciprocity Boundary Element Method for One-Group Neutron Diffusion Eigenvale Computations." *Nuclear Science and Technology*, **volume 33**(2), pp. 101–109 (1996).

[10] M. Itagaki. "Matrix-Type Multiple Reciprocity Method Applied to the Modified Helmholtz Neutron Flux Mode in Nuclear Criticality System." *Engineering Analysis with Boundary Elements*, **volume 26**, pp. 807–812 (2002).

[11] S. Nakamura. *Computational Methods in Engineering and Science, With Applications to Fluid Dynamics and Nuclear Systems*. John Wiley & Sons, Inc., New York (1977).

[12] A. Hotta, Y. Suzawa, and H. Takeuchi. "Development of BWR Regional Instability Model and Verification Based on Ringhals 1 Test." *Annals of Nuclear Energy*, **volume 24**(17), pp. 1403–1427 (1997).

[13] R. N. Slaybaugh. *Acceleration Methods for Massively Parallel Deterministic Transport, a PhD Dissertation*. University of Wisconsin, Madison, WI (2011).

[14] M. L. Adams and E. W. Larsen. "Fast Iterative Methods for Discrete-Ordinates Particle Transport Calculations." *Progress in Nuclear Energy*, **volume 40**(1), pp. 3–159 (2002).

[15] M. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu, T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, H. Thornquist, R. Tuminaro, J. Willenbring, and A. Williams. "An Overview of Trilinos." Technical Report SAND2003-2927, Sandia National Laboratories, Albuquerque, NM (2003).

[16] G. W. Stewart. *Matrix Algorithms Volume II: Eigensystems*. SIAM, Philadelphia, PA (2001).

[17] B. N. Parlett. "Rayleigh Quotient Iteration and Some Generalizations for Nonnormal Matrices." *Mathematics of Computation*, **volume 28**(127), pp. 679–693 (1974).

[18] G. Peters and J. H. Wilkinson. "Inverse Iteration, Ill-Conditioned Equations and Newton's Method." *Society for Industrial and Applied Mathematics Review*, **volume 21**(3), pp. 339–360 (1979). URL http://www.jstor.org/stable/2029572.

[19] C. C. Paige, M. Rozloznik, and Z. Strakos. "Modified Gram-Schmidt (MGS), Least Squares, and Backward Stability of MGS-GMRES." *SIAM Journal on Matrix Analysis and Applications*, **volume 28**(1), pp. 264–284 (2006).

[20] N. M. Nachtigal, S. C. Reddy, and L. N. Trefethen. "How Fast are Nonsymmetric Matrix Iterations?" *SIAM Journal on Matrix Analysis and Applications*, **volume 13**(3), pp. 778–795 (1992).

[21] M. Benzi. "Preconditioning Techniques for Large Linear Systems: A Survey." *Journal of Computational Physics*, **volume 182**, pp. 418–477 (2002).

[22] Y. Saad and M. H. Schultz. "Conjugate Gradient-Like Algorithms for Solving Nonsymmetric Linear Systems." *Mathematics of Computation*, **volume 44**(170), pp. 417–424 (1985).

[23] OECD-NEA. "Benchmark on Deterministic Transport Calculations Without Spatial Homogenisation, A 2-D/3-D MOX Fuel Assembly Benchmark." Technical report, Paris, France (2003).