International Conference on Mathematics, Computational Methods & Reactor Physics (M&C 2009)
Saratoga Springs, New York, May 3-7, 2009, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2009)

Relaxation scheme for Monte Carlo explicitly restarted Arnoldi's method for criticality calculations

Jeremy Lloyd Conlin and James Paul Holloway
University of Michigan
Ann Arbor, MI 48109
jlconlin@umich.edu; hagar@umich.edu

Abstract

In a recent paper we have shown that an explicitly restarted Arnoldi's method can be used to calculate eigenvalues and eigenvectors for Monte Carlo criticality applications. In this paper we present further research in a Monte Carlo application of Arnoldi's method of minimized iterations. We have implemented a relaxed matrix-vector product to reduce the uncertainty of the mean of eigenvalue estimates while maintaining the convergence properties of Arnoldi's method.

Key Words: Arnoldi, Monte Carlo, Relaxed, Matrix-vector product, Eigenvalue

1 Introduction

Monte Carlo criticality codes use the power method to calculate the eigenvalue of a neutron multiplying material. As an alternative we have implemented Arnoldi's method of minimized iterations [1] for Monte Carlo particle transport to calculate eigenvalues and eigenvectors of a multiplying medium. In our previous report [4] we have shown that our implementation is able to accurately calculate the fundamental and first harmonic eigenpairs of the transport operator.

In 2000, Bouras and Frayssé [2] reported that a relaxation strategy for matrix-vector products can be used for Arnoldi's method, thereby reducing the computational expense while maintaining a desired level of precision. In this paper, we show that a relaxed Arnoldi's method can be used in Monte Carlo particle transport calculations to reduce the uncertainty of mean eigenvalue estimates for the same number of particles tracked as compared to a non-relaxed method.

2 Transport Equation

The Boltzmann equation describing neutron transport for a k-effective eigenvalue problem,

$$\mathbf{\Omega}\cdot\nabla\psi(\mathrm{r},\mathbf{\Omega})+\Sigma_\mathrm{t}\psi(\mathrm{r},\mathbf{\Omega})=\frac{\Sigma_\mathrm{s}}{4\pi}\int\psi(\mathrm{r},\mathbf{\Omega})\,\mathrm{d}\Omega+\frac{1}{\mathrm{k}}\frac{\nu\Sigma_\mathrm{f}}{4\pi}\int\psi(\mathrm{r},\mathbf{\Omega})\,\mathrm{d}\Omega, \tag{1}$$

can be written in operator form

$$(\mathrm{L}+\mathrm{C}-\mathrm{S})\psi=\frac{1}{\mathrm{k}}\mathrm{F}\psi \tag{2a}$$

$$\mathrm{T}\psi=\frac{1}{\mathrm{k}}\mathrm{F}\psi \tag{2b}$$

where L, C, and S are the leakage, collision, and scattering operators respectively and F is the fission operator; $T = L + C - S$ is the transport-collision operator. The left-hand side of Eq. 2 represents the neutron loss and scattering mechanisms and the right-hand side represents the neutron gain mechanism.

If we define

$$q \equiv F\psi, \tag{3a}$$

$$A \equiv F\,T^{-1} \tag{3b}$$

and manipulate Eq. 2, we find

$$Aq = kq. \tag{4}$$

This is a standard eigenvalue problem with eigenvalue k and eigenvector q for the transport-fission operator, A. The linear operator A is used to generate a Krylov subspace as discussed in the next section.

3 Krylov Methods

Krylov subspace methods [6] approximate solutions (e.g. eigenvectors) from a subspace of the form

$$\mathscr{K}_m = \text{span}\left\{v, Av, A^2v, \ldots, A^{m-1}v\right\}, \tag{5}$$

where v is some function/vector and A is a linear operator. Krylov subspace methods are iterative methods which increase the size of the subspace by one dimension for each iteration by application of the operator A.

Krylov subspace methods don't require the linear operator A to be known explicitly, rather the application of that operator on a vector must be defined. This is known as a matrix-free method. This makes Krylov subspace methods attractive to Monte Carlo criticality calculations, as in Monte Carlo we don't have an explicit form of our linear operator, but its application on a source vector is straightforward.

In Monte Carlo methods, the fission source q is a collection of fission sites. To apply A to a fission source, one samples a neutron position from the source, transports the particle, and stores the coordinates of the neutron when it causes fission in a new fission source.

3.1 Power Method

The power method is a simple implementation of a Krylov method and is used in most Monte Carlo criticality applications. At each iteration the linear operator, A, is applied to a vector and a new estimate of the fundamental eigenpair is calculated;

$$q_{m+1} = \frac{1}{k_m} A q_m, \tag{6a}$$

$$k_{m+1} = k_m \frac{|Aq_m|}{|q_m|}. \tag{6b}$$

In Eqs. 6, q_m and k_m are the estimates of the eigenvector and eigenvalue, respectively, at the m-th iteration. Only one estimate of the eigenpair is stored at any time.

Because an eigenvalue estimate is calculated at each iteration, a mean value of these estimates can be calculated over many iterations and the variance of the mean of the estimates gives us the statistical uncertainty of our calculation. As the method proceeds and the dimension of the Krylov subspace increases, the estimate of the eigenvalue and eigenvector converge to their true values—assuming no bias in the methodology[3].

3.2 Arnoldi's Method

Arnoldi's method is another Krylov method, but not so simple as the power method. Arnoldi's method applies the linear operator A to a vector in the same manner as the power method. It differs from the power method by storing the results of the previously calculated applications of the linear operator. These results are the basis vectors for the Krylov subspace. This basis is denoted $\{q_i\}, i = 1, \ldots, m$, and the basis vectors are called Arnoldi vectors.

At each iteration of Arnoldi's method the linear operator is applied to the previously calculated Arnoldi vector. The result is then orthogonalized against all the Arnoldi vectors,

$$\tilde{q}_{m+1} = Aq_m - \sum_{j=1}^{m} q_j h_{jm}, \tag{7}$$

and normalized,

$$q_{m+1} = \frac{\tilde{q}_{m+1}}{h_{m+1,m}}. \tag{8}$$

In Eqs. 7 and 8, h_{jm} is the inner product of the new Arnoldi vector against a previous Arnoldi vector

$$h_{jm} = \langle Aq_m, q_j \rangle = \langle q_{m+1}, q_j \rangle. \tag{9}$$

The process of orthogonalizing and normalizing the Arnoldi vectors creates an upper Hessenberg matrix, H with elements h_{jk}. After n iterations of the Arnoldi process we have

$$AQ_n = Q_n H_n + q_{n+1} h_{n+1,n} e_n^T \tag{10}$$

where the columns of Q_n are the Krylov subspace basis vectors (Arnoldi vectors) and e_n is the n-th standard basis vector.

The Arnoldi factorization shown in Eq. 10 is a similarity transformation of the linear operator, A, to the upper Hessenberg matrix H_n plus the term $q_{n+1} h_{n+1,n} e_n^T$. H_n is small so its eigenvalues and eigenvectors

can easily be calculated with whatever algorithm is preferred. Let (λ, x_n) be an eigenpair of H_n. We can show that λ is an estimate of the eigenvalue of A: beginning with the Eq. 10 multiplied by x_n on the right,

$$AQ_n x_n = Q_n H_n x_n + q_{n+1} h_{n+1,n} e_n^T x_n \tag{11a}$$

$$AQ_n x_n = \lambda Q_n x_n + q_{n+1} h_{n+1,n} e_n^T x_n \tag{11b}$$

$$Ay_n = \lambda y_n + q_{n+1} h_{n+1,n} e_n^T x_n \tag{11c}$$

where $y_n = Q_n x_n$. We see, therefore, that (λ, y_n) is nearly an eigenpair of A, excepting only a term $q_{n+1} h_{n+1,n} e_n^T x_n$ that we desire to be small. The pair (λ, y_n) is called a Ritz pair of A; it is an estimate of an eigenpair of A. Note also that y_n is a linear combination of the Arnoldi vectors with coefficients that are the elements of the eigenvector x_n.

The extra term in Eq. 10, $q_{n+1} h_{n+1,n} e_n^T$ is the difference between the eigenvalues of H_n and A is called the residual and is

$$r_n = \left| q_{n+1} h_{n+1,n} e_n^T x_n \right| = \left| h_{n+1,n} \right| \left| e_n^T x_n \right| \tag{12}$$

In simpler terms, the residual is equal to the absolute value of the last component of the eigenvector x_n multiplied by $\left| h_{m+1,m} \right|$.

The number of iterations of Arnoldi's method required for convergence to the correct solution depends upon the initial vector supplied to the method. If the initial vector is the desired eigenvector then Arnoldi's method will have converged to the true solution in one iteration. If we are interested in n eigenpairs and the initial vector is a linear combination of the n desired eigenvectors then Arnoldi's method will converge to the solution in n iterations.

We are never so fortunate to know the answer a priori. It is to our advantage to find an initial vector as close as possible to the desired eigenvectors. After a few iterations of Arnoldi's method we have a better estimate of our desired eigenvectors than our initial vector. Arnoldi's method can be restarted using this improved estimate as the initial vector for a new set of iterations. Restarting Arnoldi's method reduces computational expense and improves the convergence of calculation.

3.2.1 Monte Carlo Implementation

In our implementation [4], we treat an Arnoldi restart similarly to a power iteration. At the end of each Arnoldi restart, an estimate of the eigenpairs are calculated and stored. The mean and variance of these estimates can be compared to the mean and variance calculated with the power method.

The orthogonalization of the basis vectors in Arnoldi's method requires taking the inner product of two vectors as shown in Eq. 7; we must define what it means to take the inner product of two fission sources. In our studies, we have discretized the fission source into spatial bins and so have represented the infinite-

dimensional fission source as a B-dimensional vector

$$q(x) = \sum_{i=1}^{B} a_i f_i(x), \tag{13}$$

where B is the number of spatial bins and

$$f_i(x) = \begin{cases} 1, & x_i \leq x < x_{i+1} \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

Representing the fission source vector in this manner, taking the inner product of two fission sources is simply the dot product of two vectors of coefficients $\{a_i\}_{i=1}^{B}$.

Applying A to a fission source in Arnoldi's method is done the same way as for the power method; particles are sampled from a fission source, transported, and tallied when they cause fission to create a new fission source. In Arnoldi's method particles are sampled from a piecewise constant fission source function represented by the vector of coefficients $\{a_i\}_{i=1}^{B}$, while in the power method the fission source is a collection of fission sites.

Orthogonalizing the Arnoldi vectors will inevitably lead to elements of these vectors being negative. We have developed a method of sampling from such a source and transporting negative weight particles. For details on this sampling strategy see our earlier paper [4].

4 Relaxed Krylov Methods

Recently there has been some interest in Krylov methods where the application of the linear operator or the matrix-vector product is performed inexactly. An inexact application of the linear operator is one where the application of the linear operator is terminated early; for example when the application of the operator is an iterative process the iterations are terminated before the calculation converges. By terminating early, computation time is saved at the expense of the precision of the calculation.

Bouras and Frayssé [2] have discovered that under certain circumstances the inexact application of the linear operator has little or no effect upon the convergence of the eigenpair calculation. They discovered that as Arnoldi's method proceeds and the size of the Krylov subspace increases, the precision to which the matrix-vector product is calculated can be decreased, or relaxed. The strategy by which they relaxed their matrix-vector product is shown below.

Let η be the final desired tolerance for the solution of the eigenvalue and α_k be a scalar defined as

$$\alpha_k = \frac{1}{\min\left(\left\|r_{k-1}\right\|, 1\right)}, \tag{15}$$

where r_{k-1} is the residual associated with the previously computed eigenpair. When relaxing the matrix-vector product a perturbation term is added

$$\hat{q}_{m+1} = (A + \Delta A)\, q_m \tag{16}$$

where ΔA represents the error in the matrix-vector product. In order to compute the Ritz pair with a residual less than η, the matrix-vector product must be computed so that

$$||\Delta A_k|| = \varepsilon_k ||A||, \tag{17}$$

where $\varepsilon_k = \min(\alpha_k \eta, 1)$.

This methodology allows the precision of the matrix-vector product to be relaxed when the residual is small while still remaining within the desired tolerance of for the solution of the eigenvalue. The smaller the residual the closer the estimate is to the true eigenvalue and the more aggressively the precision of the matrix-vector product is relaxed.

4.1 Monte Carlo Implementation

The perturbation term introduced in Eq. 16 should be familiar to a Monte Carlo researcher. The statistical uncertainty from a Monte Carlo application of a linear operator can be treated mathematically by adding a random matrix to our linear operator. We know the uncertainty of a Monte Carlo particle transport calculation is inversely proportional to the square root of the number of particles tracked (N) to apply the operator,

$$||\Delta A|| \propto \frac{1}{\sqrt{N}}. \tag{18}$$

In the Monte Carlo sense, relaxing the application of the operator is done simply by tracking fewer particles in one cycle than in the previous cycle.

The relaxation scheme we have used for our Monte Carlo implementation is similar to that of Bouras and Frayssé; the smaller the residual, the more we can relax the precision of the application of the transport-fission operator. In our Monte Carlo implementation of Arnoldi's method the number of particles tracked in iteration k is defined as

$$N_k = \begin{cases} N_0 & , \quad \eta < r_{k-1} \\ N_0 (r_{k-1}/\eta) & , \quad \eta \geq r_{k-1}, \end{cases} \tag{19}$$

where r_{k-1} is the residual from iteration $k-1$, N_0 is the number of histories to track in each cycle if there was no relaxing the matrix-vector product, and η is a given parameter used to control how aggressively we relax. This scheme prevents the number of histories tracked in a cycle from getting too large, but reduces the number of histories tracked when the residual becomes small.

The first iteration of each Arnoldi restart always tracks N_0 particles. The number of particles tracked in subsequent iterations may be relaxed by the fraction r_{k-1}/η as shown in Eq. 19. The total number of particles tracked in a restart is reduced as well as the computational expense of tracking particles in each iteration. The calculation of the residual requires calculating the eigenvectors of H_n at each iteration. In non-relaxed Arnoldi, the calculation of the eigenvectors is performed at the end of each restart. In relaxed Arnoldi, we must calculate the residual—and also the eigenvectors—at every iteration. This contributes to the time required for Arnoldi's method. The magnitude of the extra expense will be shown later.

5 Numerical Results

We have implemented a relaxed Arnoldi algorithm in a slab-geometry Monte Carlo criticality code. Here we present the results of a few simulations. The geometry of the problem consists of a single semi-infinite, homogeneous slab of width 20 mfp. We have used this geometry with two different materials, a highly absorbing material with cross sections $\Sigma_s = 0.5$, $\nu\Sigma_f = 0.5$, $\Sigma_t = 1.0$ and a highly scattering material with cross sections $\Sigma_s = 0.9$, $\nu\Sigma_f = 0.1$, $\Sigma_t = 1.0$ The geometry and homogeneous absorbing material were chosen to compare with results published in Rathkopf and Martin [5]. The fundamental eigenvalue for this material is published as 0.985928 and will be referred to as λ_0 in this paper. The fundamental eigenvalue for the scattering material was calculated with an S_N code and is 0.933387.

In our simulation with the absorbing material, we track 1E6 particles per iteration. In the Arnoldi methods, we use 10 iterations per restart with 5 inactive restarts. In the relaxed Arnoldi's method there are 150 active restarts while in the non-relaxed Arnoldi there are 28 active restarts. The power method has the same number of iterations as the non-relaxed Arnoldi with 50 inactive iterations and 280 active iterations. The relaxation parameter $\eta = 0.1$ for relaxed Arnoldi. All methods have the same number of inactive iterations and track roughly the same number of particles

In the simulation with the scattering material, we also track 1E6 particles per iteration. In the Arnoldi methods, we use 10 iterations per restart with 15 inactive restarts. In relaxed Arnoldi's method there are 150 active restarts while in the non-relaxed Arnoldi there are 25 active restarts. The power method has the same number of iterations as the non-relaxed Arnoldi with 150 inactive iterations and 250 active iterations. The relaxation parameter, $\eta = 0.1$ for relaxed Arnoldi. All methods have the same number of inactive iterations and track roughly the same number of particles

Figure 1 shows the convergence of the eigenvalue estimates for the absorbing material and Figure 2 shows the eigenvalue convergence for the scattering material. A dark line shows the reference solution for each material. Each method agrees with the reference solution within statistical uncertainty. Plots of the estimated eigenvector are not included in this paper because they were not interesting. Each method calculated the fundamental eigenvector accurately and the eigenvectors calculated with each method are indistinguishable from each other.

Table I shows the results from the power, Arnoldi, and relaxed Arnoldi simulations for both the absorbing and scattering materials. For the absorbing material the power method uncertainty is the smallest followed by the relaxed Arnoldi's method, and the uncertainty of the non-relaxed Arnoldi's method is the largest. For the scattering material we notice again that the non-relaxed Arnoldi's method has the largest uncertainty, but that the relaxed Arnoldi has an uncertainty equal to that of the power method.

We know the uncertainty of the mean goes as $1/\sqrt{N_e}$ where N_e is the number of estimates of the mean. Arnoldi's method suffers a little bit in this regard because it only calculates an estimate at the end of every

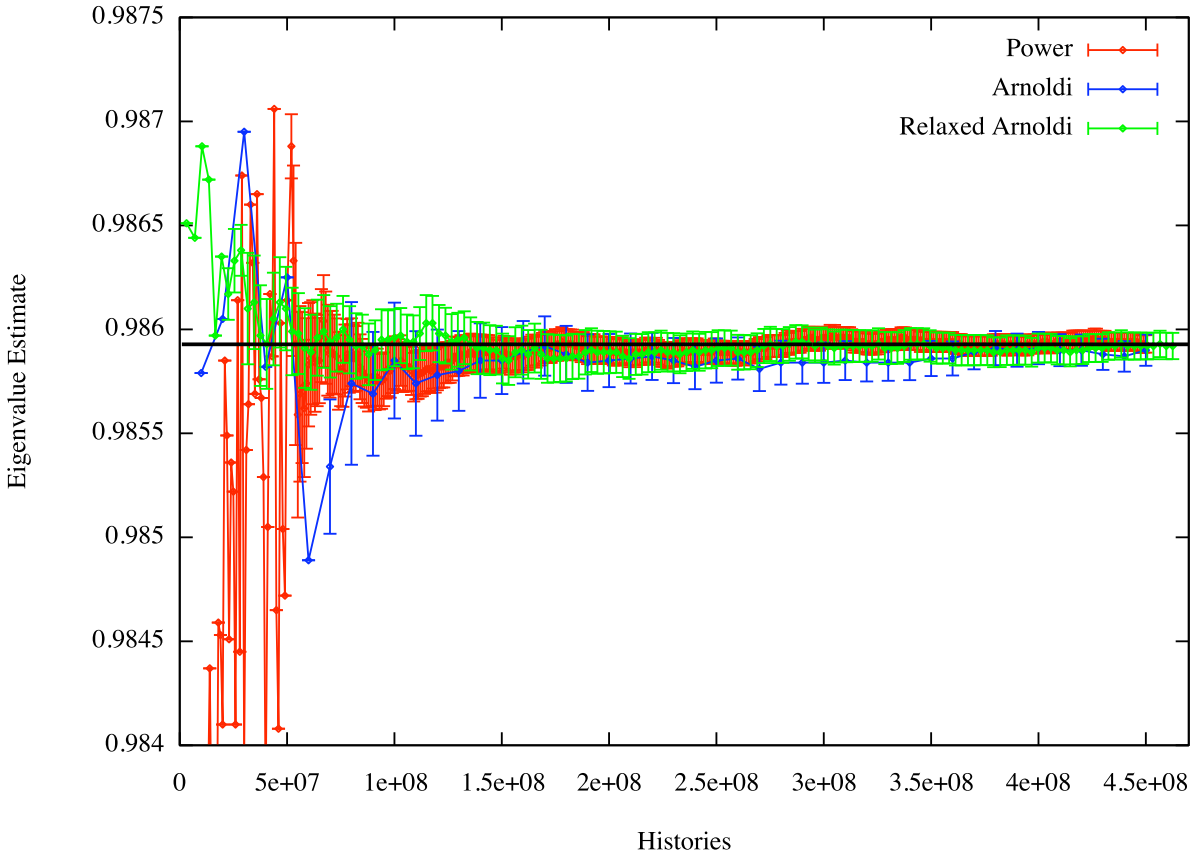

Figure 1. Eigenvalue convergence from power, Arnoldi, and relaxed Arnoldi methods for the absorbing material. The power method is shown in red, Arnoldi's method in blue, and relaxed Arnoldi's method in green. In the Arnoldi methods we use 10 iterations per restart with 5 inactive restarts. In relaxed Arnoldi there are 150 active restarts, while non-relaxed Arnoldi has 28 active restarts. The power method has 50 inactive iterations and 280 active. The relaxation parameter for relaxed Arnoldi is 0.1. All methods have the same number of inactive iterations and roughly the same total number of particles tracked. A heavy black line is shown representing the benchmark taken from the published report by Rathkopf and Martin [5].

Table I. Mean eigenvalue estimates for power, Arnoldi, and relaxed Arnoldi methods. The reference eigenvalue, λ_0, is from Rathkopf and Martin [5] for the Absorbing material.

Material	λ_0	Method	Eigenvalue
Absorbing	0.98593	Power	0.98593 ± 0.00004
		Arnoldi	0.98590 ± 0.00008
		Relaxed Arnoldi	0.98592 ± 0.00006
Scattering	0.93339	Power	0.93335 ± 0.00008
		Arnoldi	0.93331 ± 0.00023
		Relaxed Arnoldi	0.93338 ± 0.00008

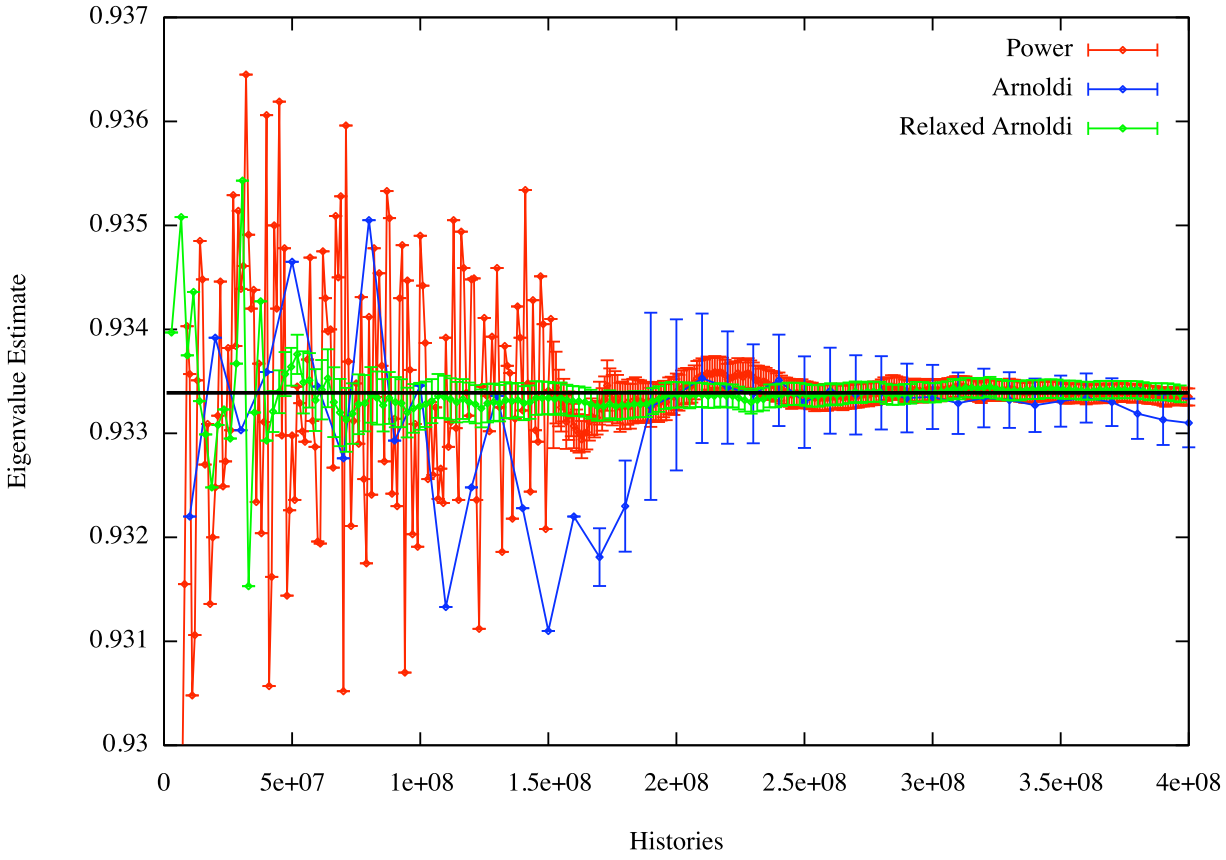

Figure 2. Eigenvalue convergence from power, Arnoldi, and relaxed Arnoldi methods for a scattering material. The power method is shown in red, Arnoldi's method in blue, and relaxed Arnoldi's method in green. In the Arnoldi methods we use 10 iterations per restart with 5 inactive restarts. In relaxed Arnoldi there are 150 active restarts, while non-relaxed Arnoldi has 25 active restarts. The power method has 50 inactive iterations and 250 active. The relaxation parameter for relaxed Arnoldi is 0.1. All methods have the same number of inactive iterations and roughly the same total number of particles tracked. A heavy black line is drawn indicating the reference solution.

restart while the power method calculates an estimate at every iteration. For this problem, Arnoldi's method has fifteen times fewer estimates than the power method, thus reducing the uncertainty for a given number of particles tracked. The relaxed Arnoldi improves somewhat upon this deficiency. By reducing or relaxing the number of particles tracked in a restart (and an iteration) relaxed Arnoldi has more restarts for the same number of particles tracked in a simulation. Increasing the number of restarts and the number of eigenvalue estimates reduces the uncertainty for relaxed Arnoldi's method.

Figure 3 shows a plot of the uncertainty of the mean of the eigenvalue estimates as a function of the number of particles tracked in the simulation for the absorbing material. At first glance it appears from Figure 1 that Arnoldi's method is less noisy than the power method, but—as was previously mentioned—the actual number of estimates is much less, thus the uncertainty for Arnoldi's method is greater. For a

small number of particles relaxed Arnoldi's method has the smallest uncertainty. Relaxed Arnoldi tracks fewer particles on average in each iteration and so for the same number of particles tracked will have more eigenvalue estimates than non-relaxed Arnoldi and thus a smaller uncertainty. We see a similar trend for the scattering material in Figure 4.

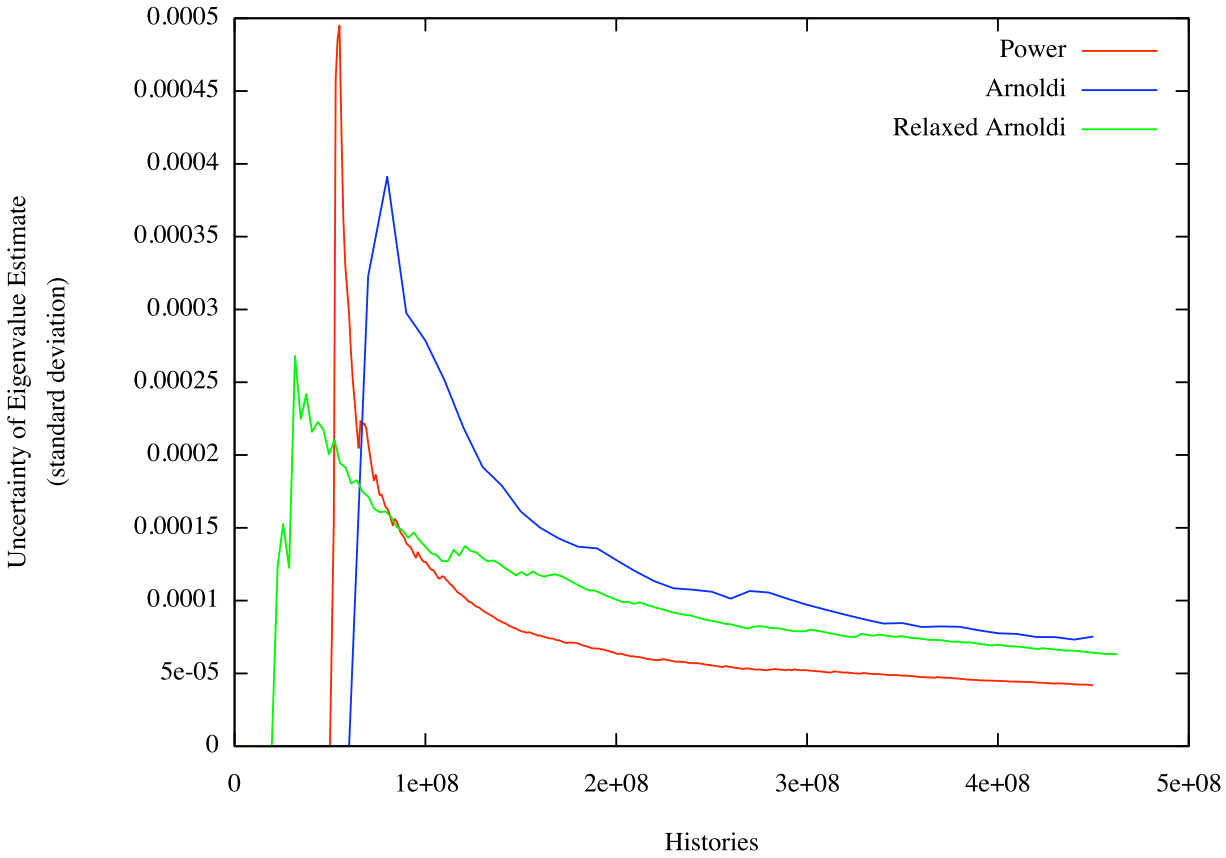

Figure 3. Uncertainty of mean of eigenvalue estimates from power, Arnoldi, and relaxed Arnoldi methods for the absorbing material. The power method is shown in red, Arnoldi's method in blue, and relaxed Arnoldi's method in green. In the Arnoldi methods we use 10 iterations per restart with 5 inactive restarts. In relaxed Arnoldi there are 150 active restarts, while non-relaxed Arnoldi has 25 active restarts. The power method has 50 inactive iterations and 250 active. The relaxation parameter for relaxed Arnoldi is 0.1. All methods have the same number of inactive iterations and roughly the same total number of particles tracked.

Table II shows the timing results for our simulation. All the simulations were run on an Intel Core Duo running at 2.0GHz with only one CPU. The cost values shown were calculated by dividing the number of particles tracked per second of the power method by the method in the table. This gives us a relative efficiency compared to the power method.

We can see that for the absorbing material, non-relaxed Arnoldi's method is the fastest and the relaxed Arnoldi's method is the slowest. Arnoldi's method does require the calculation of eigenvalues of a matrix

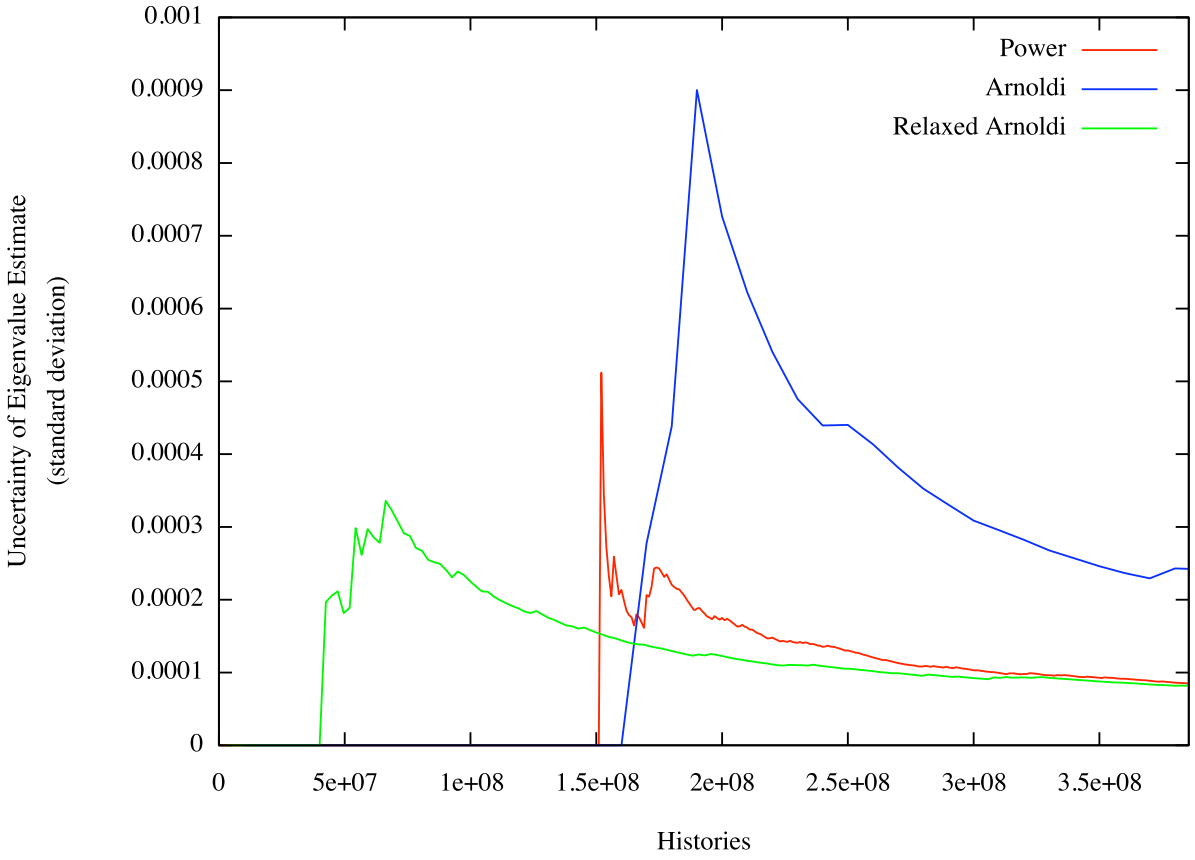

Figure 4. Uncertainty of mean of eigenvalue estimates from power, Arnoldi, and relaxed Arnoldi methods for the highly scattering material. The power method is shown in red, Arnoldi's method in blue, and relaxed Arnoldi's method in green. In the Arnoldi methods we use 10 iterations per restart with 15 inactive restarts. In relaxed Arnoldi there are 150 active restarts, while non-relaxed Arnoldi has 25 active restarts. The power method has 150 inactive iterations and 250 active. The relaxation parameter for relaxed Arnoldi is 0.1. All methods have the same number of inactive iterations and roughly the same total number of particles tracked.

as well as the orthogonalization and normalization of the Arnoldi vectors in addition to the tracking of particles. However Arnoldi's method can take advantage of the fact that it doesn't have to store the particles from the previous iteration as the power method must do. This storage is computationally expensive and consumes memory resources. The extra computational expense required in Arnoldi's method is less than the expense of storing more particles as in the power method.

The relaxed Arnoldi's method, however, is slower than the power method. It has the same advantages as Arnoldi's method in that it doesn't have to store the particles, but must also calculate the residual in each iteration in order to determine how much to relax the application of the linear operator. Calculating the residual requires calculating the eigenvalues and eigenvectors of the upper Hessenberg matrix H_n as shown in Eqs. 12. In Arnoldi's method, the calculation of the eigenpairs of H_n is performed only once

every restart. In relaxed Arnoldi's method, this calculation must be performed at every iteration. For the particular simulation described in this paper the eigenpairs are calculated ten times more frequently in relaxed Arnoldi's method than in non-relaxed Arnoldi's method. Relaxed Arnoldi is therefore more computationally expensive in this example by 10% over the power method.

The timing results from the highly scattering material differ from that of the absorbing material. Here we see that the computational expense for the power method and the Arnoldi methods are nearly the same. It should be noted that in all cases, the Arnoldi methods can computer higher eigenmodes with no increase in computational cost.

Table II. Timing results for power, Arnoldi, and relaxed Arnoldi methods. These simulations were run on an Intel Core Duo running at 2.0GHz, using only one CPU. The cost of each method is the average number of particles per second compared with the power method.

Material	Method	Runtime (s)	Total # Particles	Particles/sec	Cost
Absorbing	Power	1742	4.5×10^8	2.58×10^5	1
	Arnoldi	1669	4.5×10^8	2.70×10^5	0.956
	Relaxed Arnoldi	1972	4.63×10^8	2.34×10^5	1.100
Scattering	Power	2310	4.00×10^8	1.73×10^5	1
	Arnoldi	2340	4.0×10^8	1.71×10^5	1.01
	Relaxed Arnoldi	2272	3.97×10^8	1.75×10^5	0.989

6 Conclusion

We have shown in our paper that we can employ a relaxation scheme for the application of the transport-fission operator in Arnoldi's method for Monte Carlo applications. The scheme reduces the number of particles used in each iteration. While this reduces the computational expense of tracking particles, additional computational effort is required to determine the amount of relaxation. In these simulations the extra effort for the absorbing material was greater than the savings for relaxing and the overall computational expense was greater. For the highly scattering material, the extra computational effort was small and the relaxed Arnoldi's method was only marginally faster than the power method and non-relaxed Arnoldi's method. More work remains to be done to fully understand why the computational expense is so much less for the highly scattering material and not for the absorbing material.

Relaxing the precision of the application of the linear operator in Arnoldi's method can reduce the uncertainty of the mean of the fundamental eigenvalue estimate for the same number of particles tracked between the two Arnoldi methods. Relaxed Arnoldi's method has a smaller uncertainty because the number of eigenvalue estimates for a given number of tracked particles is greater. In the authors' opinion, Arnoldi's method should always be relaxed in a Monte Carlo criticality application. Much can be gained through this approach.

REFERENCES

[1] W. E. Arnoldi. The principle of minimized iterations in the solution of the matrix eigenvalue problem. Quarterly of Applied Mathematics, 9(1):17–29, 1951.

[2] A Bouras and V Frayssé. A relaxation strategy for the arnoldi method in eigenproblems. Technical report, 2000.

[3] R. J. Brissenden and A. R. Garlick. Biases in the estimation of keff and its error by monte carlo methods. ANNALS OF NUCLEAR ENERGY, 13(2):63–83, 1986.

[4] Jeremy Lloyd Conlin and James Paul Holloway. Arnoldi's method of minimized iterations for monte carlo criticality calculations. volume International Conference on the Physics of Nuclear Power: A Sustainable Resource, Sep 2008.

[5] J. A. Rathkopf and W. R. Martin. The finite element response matrix method for the solution of the neuron transport equation. Progress in Nuclear Energy, 18(1-2):237–250, 1986.

[6] Yousef Saad. Numerical Methods for Large Eigenvalue Problems. 1992.